



Stigler Center

for the Study of the Economy and the State

Stigler Committee on Digital Platforms

Final Report

2019

Stigler Committee on Digital Platforms

Organization:

- **Luigi Zingales**, Robert C. McCormack Distinguished Service Professor of Entrepreneurship and Finance, University of Chicago Booth School of Business. Director, the George J. Stigler Center for the Study of the Economy and the State
- **Guy Rolnik**, Clinical Associate Professor of Strategic Management, University of Chicago Booth School of Business
- **Filippo Maria Lancieri**, Fellow, George J. Stigler Center for the Study of the Economy and the State. JSD Candidate, University of Chicago Law School

Subcommittee on Market Structure and Antitrust

- **Chair: Fiona Scott Morton**, Theodore Nierenberg Professor of Economics, Yale University School of Management
- **Pascal Bouvier**, Managing Partner and co-founder, MiddleGame Ventures
- **Ariel Ezrachi**, Slaughter and May Professor of Competition Law and Fellow of Pembroke College, University of Oxford
- **Bruno Jullien**, Research Faculty, Toulouse School of Economics
- **Roberta Katz**, Senior Research Scholar, Stanford University
- **Gene Kimmelman**, President and CEO, Public Knowledge
- **Douglas Melamed**, Professor of the Practice of Law, Stanford Law School
- **Jamie Morgenstern**, Assistant Professor, School of Computer Science, Georgia Tech

Subcommittee on the News Media Industry

- **Chair: Guy Rolnik**, Clinical Associate Professor of Strategic Management, University of Chicago Booth School of Business
- **Julia Cage**, Assistant Professor of Economics, Sciences Po Paris
- **Joshua Gans**, Professor of Strategic Management and Jeffrey S. Skoll Chair of Technical Innovation and Entrepreneurship, Rotman School of Management, University of Toronto
- **Ellen Goodman**, Professor of Law, Rutgers University
- **Brian Knight**, Professor of Economics, Brown University
- **Andrea Prat**, Richard Paul Richman Professor of Business and Professor of Economics, Columbia University
- **Anya Schiffrin**, Director of the Technology, Media, and Communications specialization, School of International and Public Affairs, Columbia University

Subcommittee on Privacy and Data Protection

- **Chair: Lior Strahilevitz**, Sidley Austin Professor of Law, University of Chicago
- **Lorrie Cranor**, Director and Bosch Distinguished Professor in Security and Privacy Technologies, CyLab Privacy and Security Institute; FORE Systems Professor, Computer Science and Engineering & Public Policy, Carnegie Mellon University
- **Florencia Marotta-Wurgler**, Professor of Law, New York University School of Law
- **Jonathan Mayer**, Assistant Professor of Computer Science and Public Affairs, Princeton University
- **Paul Ohm**, Professor of Law and Associate Dean for Academic Affairs, Georgetown University Law Center
- **Katherine Strandburg**, Alfred B. Engelberg Professor of Law, New York University School of Law
- **Blase Ur**, Neubauer Family Assistant Professor of Computer Sciences, University of Chicago

Subcommittee on Politics

- **Chair: Nolan McCarty**, Susan Dod Brown Professor of Politics and Public Affairs, Princeton University
- **Rana Foroohar**, Global Business Columnist and Associate Editor, *Financial Times*
- **Andrew Guess**, Assistant Professor of Politics and Public Affairs, Princeton University
- **David Lazer**, Professor of Political Science, Northeastern University
- **Alexandra Siegel**, Postdoctoral Fellow, Stanford University
- **Nick Stephanopoulos**, Professor of Law and Herbert and Marjorie Fried Research Scholar, University of Chicago
- **Joshua Tucker**, Professor of Politics, New York University

ACKNOWLEDGMENTS:

First and foremost, we thank the chairs and all the Committee members listed above. These incredible academics and policymakers dedicated, without compensation, a significant amount of their spare time to think about Digital Platforms and how they impact our modern society. We also would like to thank Simone Cavallaro, Sebastian Burca, and Rachel Piontek for their invaluable support – without them this project would not have gotten off the ground; Asher Schechter for media support; the George J. Stigler Center for the Study of the Economy and the State at the University of Chicago Booth for financial and logistics support and the Alfred P. Sloan Foundation for financial support that helped us organize the in-person meetings critical to the success of the Committee.

Suggested Citation:

Stigler Committee on Digital Platforms, Final Report, September 2019, available at <https://research.chicagobooth.edu/stigler/media/news/committee-on-digital-platforms-final-report>

Table of Contents

Policy Brief for Regulators p. 6

Market Structure and Antitrust Subcommittee Report p. 23

News Industry Subcommittee Report p. 139

Privacy and Data Protection Subcommittee Report p. 206

Politics Subcommittee Report p. 268

Stigler Committee on Digital Platforms: *Policy Brief*Luigi Zingales¹Filippo Maria Lancieri²

September 2019

I. INTRODUCTION

One of the key defining factors of the past decade is the rise of Digital Platforms (DPs), such as Google, Facebook, Amazon, Apple. As more and more of our economy and society moved online, these companies ascended from non-existent or nearly bankrupt in the early 2000s to join Microsoft as global behemoths, exceeding (as of August 2019) more than 4 trillion dollars in market capitalization.

This meteoric rise is not surprising. These companies invented new products and services that revolutionized the way we work, study, travel, communicate, shop, and even date. In the process, they created trillions of dollars in consumer surplus. Nonetheless, recognizing the enormous gains brought about by these companies to date does not equate to saying that: (i) these gains will endure, especially if markets are no longer competitive; and (ii) there is no room for welfare gains by reducing some of the downsides brought about by them. Cars dramatically improved our way of life. Nonetheless, they also introduced new risks which demanded new laws and regulations. Traffic lights and roundabouts have not destroyed the benefits of cars, but they have dramatically reduced their negative impact on society.

Whether it is the novelty of their product or the consumer surplus they created (or both), so far these companies have largely avoided any regulation. In the past few years, however, as the number of scandals involving DPs increased, concerns about their unchecked power started to emerge. These concerns were not limited to economic aspects (are these companies moving to prevent any competition?) or privacy (are we in an age of surveillance capitalism?). They include the impact DPs have on our political arena and democratic values: Are they helping promote hate and/or are they a threat to the working of our democratic system?

As these important discussions multiplied, so did the proposals to intervene. Abroad, these proposals were the result of government-appointed committees—from the EU to the UK or Australia. In the United States—where no government committee was formed—the proposals were reactions to the perceived threat posed by DPs, with little to no analysis of the underlying root problems, let alone a link between market failures and remedies.

To fill this void, the George J. Stigler Center at the University of Chicago Booth School of Business decided to organize an independent Committee on Digital Platforms. The Committee brought together a group of more than 30 highly qualified, independent academics and

¹ Robert C. McCormack Distinguished Service Professor of Entrepreneurship and Finance, and Charles M. Harper Faculty Fellow, University of Chicago Booth School of Business. Director, George J. Stigler Center for the Study of the Economy and the State.

² Fellow, George J. Stigler Center for the Study of the Economy and the State. JSD Candidate, University of Chicago Law School.

policymakers³ from different disciplines to think holistically about how DPs impact: (i) the economy, (ii) privacy and data security, (iii) the news media industry, (iv) the functioning of our democracy.

For over a year, the members of each subcommittee dedicated a significant amount of time to develop a set of cohesive, independent studies on how DPs impact modern society. Draft versions of each subcommittee’s white paper were featured at the Stigler Center’s 2019 Antitrust and Competition Conference, which brought together more than 130 highly regarded academics and policy experts to discuss these topics.⁴ At the conference, each white paper received detailed feedback by two independent commentators representing different points of view, along with more general feedback from the audience. Overall, the studies presented herein represent the most comprehensive independent analysis of Digital Platforms to date.

This *Policy Brief*, aimed at a non-specialized audience, summarizes the main concerns identified by these studies and provides a viable path forward to address the identified concerns.⁵ It tries to do so in the least intrusive way possible. Section II presents the novel concerns raised by DPs. Section III describes the various policy solutions. Section IV concludes.

II. SOURCES OF CONCERNS

The term “Digital Platform” lacks a consistent definition—different companies may be characterized as a platform in different environments. For example, Google, Facebook, Amazon, Apple, and Microsoft raise different concerns regarding how their “bottleneck power” impacts the markets in which they operate.⁶ Considerations on market power involve all five companies mentioned above. By contrast, considerations about the news media or democracy are more specific to companies such as Google and Facebook and—to a lesser extent—Twitter. For this reason, the focus of our analysis in this Brief will be primarily Google and Facebook.

II.1 MARKET STRUCTURE/ANTITRUST

Digital Platforms tend to monopolies: The markets where DPs operate exhibit several economic features that, while not novel per se, appear together for the first time and push these markets towards monopolization by a single company. These features are: i) strong network effects (the more people use a product, the more appealing this product becomes for other users); ii) strong economies of scale and scope (the cost of producing more or of expanding in other sectors decreases with company’s size); iii) marginal costs close to zero (the cost of servicing another consumer is close to zero); (iv) high and increasing returns to the use of data (the more

³ See <https://research.chicagobooth.edu/stigler/events/single-events/antitrust-competition-conference/digital-platforms-committee>.

⁴ See the agenda for the 2019 ANTITRUST AND COMPETITION CONFERENCE - DIGITAL PLATFORMS, MARKETS, AND DEMOCRACY: A PATH FORWARD, available at <https://research.chicagobooth.edu/stigler/events/single-events/antitrust-competition-conference>

⁵ For example, most of our reference footnotes are to accessible articles in the main press. The reports all have multiple technical references for more specialized audiences.

⁶ Or their power to funnel user attention. Bottleneck power arises when “consumers primarily single-home and rely upon a single service provider.” For example, most sites depend on Google to receive traffic—hence saying that Google is a bottleneck in internet traffic.

data you control, the better your product); and v) low distribution costs that allow for a global reach. This confluence of features means that these markets are prone to tipping; that is, they reach a point where the market will naturally tend towards a single, very dominant player (also known as “winner takes all markets”). An entrant will most likely be unable to overcome the barriers to entry represented by scale economies and data control, as they are difficult to achieve in a quick, cost-effective manner.

When free is not necessarily good for consumers: DPs defend themselves by saying that, since most consumers do not pay for their services, how can they be hurt? This statement is incorrect in many ways. First, there is nothing special about a zero price—if the business is so successful, consumers could be charged a negative price to use Google (think of miles awarded for credit card use).⁷ Second, house buyers do not pay for their real estate broker out of their pocket, but that does not mean they do not pay for the service nor that they cannot be hurt by high real estate brokers’ commissions. Two-sided platforms, like real estate brokers, often charge more on one side to subsidize the other. In equilibrium, a higher real estate fee will be reflected in higher house prices, which will hurt buyers. The same is true for DPs like Facebook and Google. Second, only the monetary price consumers pay is zero. Consumers pay in kind, by transferring their data. Finally, market power may manifest itself through lower quality, lower privacy protection, less creation of new business/entry, less variety of political viewpoints, and, importantly, less investments in innovation. For example, a recent paper demonstrates how Facebook became much more aggressive in data collection after it faced less competition from MySpace.⁸

Market power in ads can lead to monopolization in other markets: DPs can increase the prices paid by advertisers, many of them small businesses, diverting more and more income to platforms. Have you ever noticed how Amazon buys the ads for the search “Amazon” on Google despite it being the first organic result? This shows how much power Google has even over gigantic corporations. Through their power in the ads market, DPs can also block entry of potential competitors. For example, Facebook banned cryptocurrency ads on its platform just a year before announcing its entry in the crypto space with Libra.

Consumer harm is greatest when market power is combined with behavioral biases: Consumers tend to stick with default options. If forced to choose, they opt for the most salient alternative. Highlighting an option in red or putting it in the first position nudges consumers in that direction. Google recognizes the power of defaults and pays Apple an estimated 12 billion dollars per year to be the default search engine on the iPhone.⁹ Manipulations are common even in brick-and-mortar shops, yet they are especially harmful when i) the manipulator knows a lot about the potential customers; and ii) there are limited (or no) alternatives, as is the case for most DPs. Framing, nudges, and default options can direct consumers to choices they regret. In addition, there is increasing evidence that many online products are designed to be as addictive

⁷ Indeed, Microsoft Rewards pays for searches using Bing, and a very small search engine plants trees the more someone searches. The fact that both companies cannot obtain market share from Google even in this context shows how high entry barriers are in search markets. See <https://www.microsoft.com/en-us/rewards> and <https://info.ecosia.org/what>.

⁸ See Dina Srinivasan, “The Antitrust Case Against Facebook: A Monopolist’s Journey Towards Pervasive Surveillance in Spite of Consumers’ Preference for Privacy,” *Berkeley Business Law Journal* 16, no. 1 (2019): 39.

⁹ See <https://fortune.com/2018/09/29/google-apple-safari-search-engine/>.

as possible, or to keep consumers “hooked” on the platform to increase sales without consideration to well-being.¹⁰ The combination of addiction and monopoly is probably the worst possible.

If there is a lot of smoke, there is probably fire: Many online markets present extremely high profit margins and no new relevant entry, a sign of significant barriers to entry. Furthermore, DPs bought hundreds of companies over the past years, most without any form of scrutiny by regulators. Finally, there is evidence that venture capitalists are reluctant to fund businesses in sectors that compete directly with DPs. Venture capitalists label this phenomenon “kill zones.” When combined with the structural characteristics presented above, there is sufficient evidence to justify an in-depth investigation on these companies, as those started by the Federal Trade Commission, the Department of Justice, and the European Commission. This is particularly true because, as discussed more below, DPs refuse to provide independent researchers with the data necessary to understand whether their behavior is indeed harmful. It is paradoxical that companies refuse access to the data necessary for in-depth, independent studies and then use the lack of in-depth, independent studies as evidence of lack of harm.¹¹

II.2 NEWS MEDIA

Concentration in the news media market is a problem for democracy: The news market is unlike most others. News has some aspects of a public good. In addition, a vibrant, free, and plural media industry is necessary for a true democracy. Thus, in studying the impact of DPs on the media industry, we cannot restrict ourselves to standard economic measures like consumer welfare, but rather we need to think in terms of *citizen welfare*: how democracy functioning is impacted.

Digital platforms disintermediate newspapers and monopolize news markets: DPs are quickly controlling news distribution. Facebook is now the second largest news provider in terms of attention share in the United States. In the UK, Facebook is third, Google is fifth, and Twitter is tenth. By curating the news viewers receive, DPs have effectively appropriated the role that newspaper editors used to have in influencing readers’ attention. This poses a concentration issue, as thousands of different viewpoints have now been replaced by a duopoly. These concerns are exacerbated by three additional problematic features. First, the editing is aimed at maximizing a viewer’s time on the platform, with little attention to quality of content. Second, this news editing is personalized, potentially promoting a fragmentation of citizens into separate news bubbles, jeopardizing the ability of different political groups to talk to each other. Last but

¹⁰ See Adam Alter, *Irresistible: The Rise of Addictive Technology and the Business of Keeping Us Hooked* (Penguin, 2017).

¹¹ For example, there is prima facie evidence of an increase of fatal car accidents after the introduction of car sharing services in a city. Uber and Lyft hold the data to disprove they are responsible for this increase, but so far they have not allowed any independent inquiry on the topic, nor did they produce evidence to the contrary. Given their interest in doing so, at some point we need to start thinking about inverting the burden of proof: Prima facie evidence of responsibility that cannot be further scrutinized because the companies refuse to share the data that would prove or disprove the claims should be considered strong evidence they are responsible. See <https://promarket.org/uber-kill-real-cost-ride-sharing/> citing John Manuel Barrios, Yael V. Hochberg, and Hanyi Yi, “The Cost of Convenience: Ridesharing and Traffic Fatalities,” *Available at SSRN 3361227*, 2019.

not least, Google voting stock is controlled by two individuals, Sergey Brin and Larry Page; Facebook by one, Mark Zuckerberg. Thus, three people have total control over the personalized, obscure news feeds of billions of human beings.

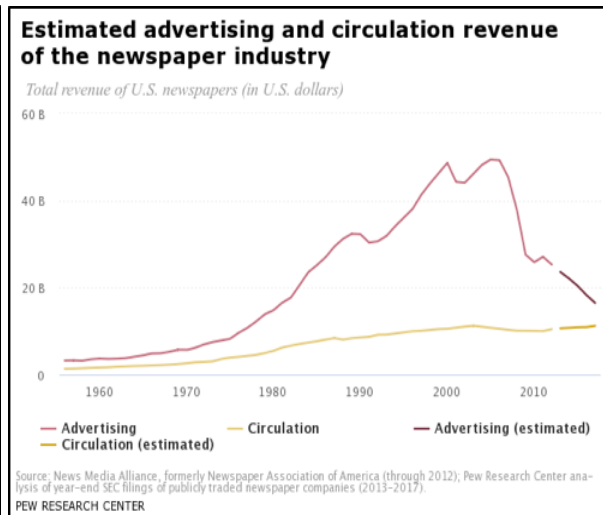

Digital Platforms have weak incentives to prioritize quality content and limit false information: These companies have weak economic and legal incentives to promote quality journalism.¹² First, journalism itself is a small part of the total content distributed by these platforms. While the data is not made public—again a problem in itself—estimates indicate that, despite its name, news is only a small fraction of Facebook News Feed.¹³ Second, the goal of all these DPs is to maximize engagement, often through extreme or divisive content, as recognized by Facebook itself.¹⁴ Unlike other media, however, DPs do not have any legal liability for promoting this content, thanks to Section 230 of the Communications Decency Act. This immunity, combined with the limited competition these platforms face, means that DPs have very weak incentives to promote quality content or to limit the spread of false or divisive information.

Digital Platforms are devastating the newspaper industry: Newspapers are a collateral damage of the digital platform revolution. Craigslist destroyed the lucrative newspaper classified ads, and Google and Facebook dramatically reduced the revenues newspapers could get from traditional advertising. Local newspapers have been hit particularly hard: At least 1800 newspapers closed in the United States since 2004, leaving more than 50% of US counties without a daily local paper. Every technological revolution destroys pre-existing business models. Creative destruction is the essence of a vibrant economy. In this respect, there is nothing new and nothing worrisome about this process. Yet, a vibrant, free, and plural media industry is necessary for a true democracy. The newspapers of yesteryear played an essential function in a democratic system. How can we make sure this function is still performed by somebody (not necessarily yesteryear newspapers)?

¹² For a more detailed account of how the ad-tech creates a race to the bottom in terms of disinformation, see <https://go.shr.lc/2YyDr8U>.

¹³ See, for example, <http://bit.ly/2zluSE9>.

¹⁴ See <https://www.facebook.com/notes/mark-zuckerberg/a-blueprint-for-content-governance-and-enforcement/10156443129621634/>.



The growing gap in investigative journalism: Ever since the muckraking magazines arose at the beginning of the 20th century, printed media has supported investigative journalism. Investigative journalism requires significant long-term investments, whose return cannot be easily appropriated by the investor, since the news report can be repeated by other sources. In oligopolistic markets, newspapers were able to finance some investigative reporting with their profits and were able to capture some of its benefits via enhanced reputation. The reduction in the number and the profitability of newspapers has severely curtailed the funds dedicated to this activity. Thus far, not-for-profit investigative reporting outlets, like Pro-Publica, have not fully covered the shortage.

The dearth of local news: The problem of lack of investigative reporting is particularly severe at the local level. *The New York Times* and the *Washington Post* have the resources to pursue national stories, but local corruption in Topeka, Kansas, or Montgomery, Alabama, is hardly of national interest. As a result, accountability at the local level may suffer. Consistent with this fear, a recent paper shows that the closure of a local newspaper increases the long-term borrowing cost of a municipality, interpreted as a sign of the inefficiencies produced by lack of accountability. Closures of local newspapers also tend to diminish both the amount of information voters have in local elections and voter turnout. Thus, there is a concern that local democracy might die in the darkness.

II.3 PRIVACY AND DATA PROTECTION

Market incentives alone will not solve privacy and data security problems: Firms that collect and process private information do not internalize the harms associated with consumer privacy and security breaches. Nor do they internalize negative externalities, or potential misuses of data that impact people who are not their own consumers.¹⁵ Notice and choice, or asking consumers to click “I accept” in extremely long terms-of-service, places solely on consumers the burden to

¹⁵ See, for example, a report describing how DNA tests done by a family historian unveiled that an uncle had an extra-marital daughter without the uncle submitting any information. <https://www.wsj.com/articles/dna-testing-creates-wrenching-dilemmas-for-the-family-historian-11563595261>.

anticipate all the downsides of their online activities. Consumers are ill-equipped to do so—they lack the time, knowledge, and capacity. This problem is only getting worse as firms become ever-more skilled in developing interfaces that manipulate choice.

This is not just a theoretical concern. Companies from small to sophisticated lack basic data protection features: Facebook stored hundreds of millions of passwords in plain text files in 2019.¹⁶ Over-reliance on industry self-regulation and “notice-and-choice” mechanisms, sparse state and federal laws with different obligations, and the lack of a regulator with a clear mandate and enforcement powers is not the correct path for a country embracing digitization like the United States. The best examples are the recent FTC enforcement actions. The settlement with Facebook—the strongest enforcement in data protection ever—does little to prevent Facebook from collecting as much data as possible and freely using these data in any way it deems appropriate.¹⁷ Facebook’s stocks even went up at the announcement of the settlement, as most commentators saw the FTC punishment as just a slap on Facebook’s wrist.¹⁸ The FTC fine against Google—charged for baiting children with targeted ads that violated their privacy—was even milder. It was less than the profits Google earned with the sanctioned practice. Furthermore, the only significant behavioral change Google committed to was to abide by a Federal Law it should have been abiding by in the first place.

Dark patterns are a particularly powerful, largely ignored, problem and they mostly impact poor and uneducated consumers: Dark patterns are “user interfaces that make it difficult for users to express their actual preferences or that manipulate users into taking actions that do not comport with their preferences or expectations.” Companies can pre-select choices, highlight or hide buttons, or constantly nag consumers in order to push them to make decisions against their preferences or expectations.¹⁹ While using interfaces or promotions that help sell a product is not illegal, doing so in an extremely manipulative way can be, as these companies are strongly manipulating vulnerable consumers into buying products and services (or watching another cat video) they ultimately do not want. Simple manipulation of user interfaces can increase acceptance rates of a data protection plan by 228% without companies facing significant consumer backlash.²⁰ While dark patterns work across the board, the effects are particularly pronounced with less-educated, vulnerable users.

Dark patterns are particularly pervasive when combined with market power: Extreme dark patterns—the ones that truly annoy consumers but can increase acceptance rates by 371%—lead to a consumer backlash against the companies employing them. Thus, consumers punish the most abusive companies. The problem is that, as seen above, many markets where DPs operate are prone to monopolization. The lack of meaningful competitors enables these companies to use very aggressive persuasion strategies. For example, studies have shown a strong link between

¹⁶ See <https://www.theverge.com/2019/3/21/18275837/facebook-plain-text-password-storage-hundreds-millions-users>.

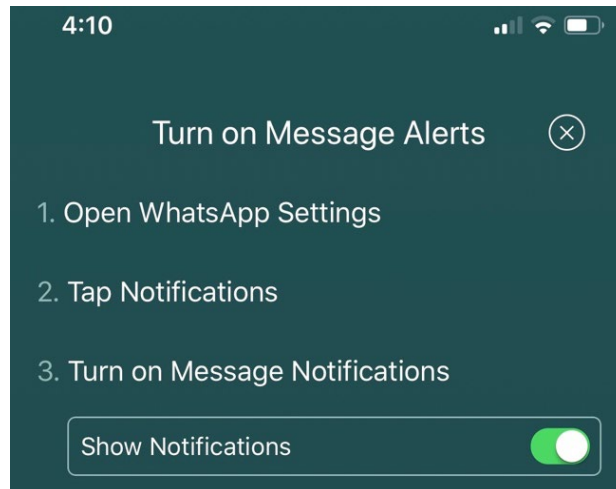
¹⁷ See <https://www.wired.com/story/off-facebook-activity-privacy/>

¹⁸ See <https://www.theverge.com/2019/7/12/20692524/facebook-five-billion-ftc-fine-embarrassing-joke>.

¹⁹ See <https://www.wsj.com/articles/how-tech-giants-get-you-to-click-this-and-not-that-11559315900>.

²⁰ See Jamie Luguri and Lior Strahilevitz, “Shining a Light on Dark Patterns,” *University of Chicago Coase-Sandor Institute for Law & Economics Research Paper No. 879*, August 2019, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3431205 for a more detailed analysis.

constant notifications and extreme anxiety, in particular on teenagers.²¹ Nonetheless, if iOS users try to turn off all WhatsApp notifications, they will be constantly bombarded by a screen commanding them to turn the notifications back on—there is no option to simply state: “*I do not wish to receive WhatsApp notifications, thank you.*”²² As WhatsApp is now the primary means of communication in many countries,²³ users cannot simply abandon WhatsApp either.



Dark patterns will probably only get worse: Even companies with comparatively little consumer information can design user interfaces that benefit them.²⁴ Sophisticated companies such as DPs collect an enormous amount of personal data. They have outsized powers to design interfaces in very manipulative ways with little to no oversight by regulators. For example, internet companies, focused only on increasing engagement to sell advertising, constantly promote interfaces that make users addicted to their products.²⁵

We need to talk about internet addiction and the combination with market power: Internet addiction is an extremely important topic that deserves much more attention than it is currently receiving. As one report puts it, “Strategies such as offering addictive content at moments when consumers lack self-control increase time spent on the platform and profitable ad sales even as the platform lowers the quality of content. These tactics increase the welfare costs of market power.” DPs are in a zero-sum race for our attention and using the most pervasive tactics to ensure they win. Since most societies regulate addictive products—drugs, alcohol, tobacco, and gambling—to protect the consumers, it is time we discuss how to regulate DPs with the same goal in mind.

²¹ See <https://www.latimes.com/nation/la-na-smartphones-causing-student-anxiety-20190607-story.html> and <https://www.nytimes.com/2019/05/20/well/family/is-digital-addiction-a-real-threat-to-kids.html>.

²² As of August 2019.

²³ Even for businesses, see <https://www.zdnet.com/article/whatsapp-is-the-main-digital-channel-for-brazilian-smes/>.

²⁴ The Wired Magazine describes how they increase subscriptions by 9% by asking people to “place order” instead of “start subscription.” See <https://www.wired.com/story/wired-paywall-one-year-later/>.

²⁵ See, <https://www.youtube.com/watch?v=NUMa0QkPzns> for an explanation of design choices aimed to be addictive; and <https://ledger.humanetech.com> for a list of peer-reviewed studies presenting negative impacts of consumer-facing mobile tech.

II.4 POLITICS

Social media companies are responsible for mass democratization of speech: Social media are rightly heralded as democratizing platforms that have greatly increased the voice of excluded members of different societies. Social media companies such as Facebook and Twitter played key roles in helping organize the Arab Spring revolutions, the Black Lives Matter movement in the United States, and, currently, the democracy protesters in Hong Kong.

Social media is different from other information technologies: The advent of many information technologies caused revolutions in political governance: The printing press, radio, and TV have profoundly transformed our democratic governance and accountability. Yet, social media's unprecedented scale, ease of anonymity, and capacity to by-pass traditional gatekeepers may be unique. When combined with DPs' tendency to monopolization, it can become problematic because it removes the accountability afforded by competition. As the report states, "in the end, the technology with the most potential to reshape modern political institutions and outcomes falls under the control of just a few firms, who themselves are enormously powerful political actors."

Digital platforms are uniquely powerful political actors: Google and Facebook may be the most powerful political agents of our time. They congregate five key characteristics that normally enable the capture of politicians and that hinder effective democratic oversight:

- i. **Money:** They have immense economic power, allowing them to effectively lobby politicians and regulators. As the five most valuable publicly listed corporations in the world and with combined cash reserves of hundreds of billions of dollars, DPs are widely using this economic power to influence politics. According to [opensecrets.org](https://www.opensecrets.org), Alphabet (Google), Amazon, and Facebook were the second, sixth, and ninth largest spenders on direct lobbying among US corporations in 2018.²⁶ In this aspect, DPs are like oil or tobacco companies in the resources they can mobilize in their defense.
- ii. **Media:** Their increasingly powerful role as a media outlet not only allows DPs to shape public discourse and to define how politicians can reach their constituents; more importantly, it allows platforms to claim both Section 230 immunity and first amendment exemptions to ward off any regulations that try to control their actions. In this aspect, DPs are similar to very powerful newspapers.
- iii. **Complexity:** Their sheer size, complexity, and absolute opacity complicates the development of effective regulatory tools, as platforms can always use information asymmetries to by-pass regulations without much awareness. Complexity also diminishes the potential talent pool for governments and incentivizes revolving doors, complicating oversight even further. In this aspect, DPs are similar to large banks in their ability to potentially dodge the most powerful regulators.

²⁶ See <https://www.opensecrets.org/lobby/top.php?indexType=s&showYear=2018>.

- iv. **Connectivity:** Their connectivity and membership allows DPs to engage their user base to challenge any political initiative that disadvantages them (think of Uber drivers' protests). In this aspect, DPs have "membership powers" similar to the National Rifle Association or to the American Association of Retired Persons in their ability to directly mobilize voters in their defense.
- v. **National Champions:** Finally, DPs constantly play the "national interest" card whenever their own interests are threatened.²⁷ In this way, DPs are similar to the steel and airplane industry, which demand preferential treatment for their strategic role.

In sum, Google and Facebook have the power of ExxonMobil, the *New York Times*, JPMorgan Chase, the NRA, and Boeing combined. Furthermore, all this combined power rests in the hands of just three people.

Digital platforms are incredibly opaque—this is a problem in itself: Finally, the lack of transparency is something that has to be stressed again, as it also impacts our political arena. If we do not know whether social media has overall positive or negative effects on our polity, we have to blame the DPs themselves. All the data they generate is proprietary and they deny outside, independent researchers access to almost all of it. In doing so, they also prevent our societies from comprehending their true impacts. For example, conservatives are constantly accusing Google and Facebook of bias, something the companies deny. Only Google and Facebook have the data that would allow an independent researcher to prove or disprove the conservatives' claim, *and* they block access to this data. Thus, we have to rely on their word. Are we sure that Elizabeth Warren's posts on Facebook receive equal distribution when compared to other candidates? This is a major problem, pervasive in all areas analyzed by the reports, and one that must be immediately addressed.

SUMMARY: This concentration of economic, media, data, and political power is potentially dangerous for our democracies. Our summary of the main concerns around DPs demonstrates why all the attention they are receiving is justified. Indeed, the conversation has barely started on some of the most worrisome topics, such as dark patterns and manipulation, addiction, or the platforms' incredible political power. To make matters worse, as more of our lives move online, the more commanding these companies will become. We are currently placing the ability to shape our democracies into the hands of a couple of unaccountable individuals. It is clear that something has to be done.

²⁷ See President Trump threatening to tax French wine in response to a tax targeting US DPs, or President Trump also trying to protect Apple from the tariff spat with China because "Apple is a great American company," at <https://reut.rs/2ZgqIMg> and <https://cnb.cx/2TUY7pR>.

III. POSSIBLE SOLUTIONS

The four subcommittees were charged with proposing an array of possible solutions in each of their specific interest areas. Some of these solutions were mutually exclusive, and sometimes different committees arrived at different conclusions. Thus, we are left with the difficult task to prioritize the various solutions and integrate them into a coherent whole.²⁸ For sake of brevity, we make no claim of being comprehensive. For a complete picture, we refer the interested reader to the four reports.

1. Forcing Interoperability

The cause of most of the problems described above is the lack of meaningful competition in many key digital markets. A major cause of this lack of competition is the presence of very sizable network externalities: that is, I want to be on the social media where my friends are. Network externalities as a potential barrier to entry are not a new phenomenon: It plagued the early phone industry. To eliminate this problem, the United States forced interoperability among the various phone companies—AT&T is obliged to connect calls started by T-Mobile consumers. The same should be done with social media. Mandating not only an open but also a common Application Program Interface (API) would allow different messaging systems to connect to one another. In so doing, a common API guarantees interoperability and eliminates the network externalities that drive the winner-take-all nature of the social media market. Facebook, for example, used all its power to kill potential interoperability solutions in order to gain market power. In 2008, it even used Federal Criminal Law to successfully attack a young startup called Power Ventures that was trying to connect different social media platforms.²⁹ The Open Banking Directive in the UK and the Brazilian Good Payer's Credit Act are good examples of cases where tailored interventions on data disclosure and open standards are increasing competition.³⁰

2. A More Aggressive Antitrust

Changing Merger Guidelines for Digital Platforms: DPs acquired hundreds of companies over the past years, most without facing any scrutiny from antitrust regulators. In traditional markets, the cost of delaying an intervention might be limited. In a market with strong tendencies toward monopolization, a mistake in the approval of a merger can condemn an industry to a monopoly. If you add the political power of these monopolies, the mistake could become irreversible. Therefore, we need to change the threshold for merger review in markets where DPs operate, basing it on transaction value or some other criteria that allows regulators to scrutinize transactions between DPs and startups—simply focusing on turnover is not enough. In addition,

²⁸ In doing so, we inevitably introduce our own views and biases. Thus, neither the participants of each committee nor the chairs necessarily agree with our conclusions.

²⁹ See <https://www.eff.org/cases/facebook-v-power-ventures> and <https://www.eff.org/deeplinks/2018/07/facing-facebook-data-portability-and-interoperability-are-anti-monopoly-medicine>.

³⁰ See <https://www.openbanking.org.uk/customers/what-is-open-banking/> and <https://iapp.org/news/a/new-changes-to-brazilian-good-payers-act-includes-automatic-registration-into-database/>.

when an acquisition involves a dominant platform, authorities should shift the burden of proof, requiring the company to prove that the acquisition will not harm competition.³¹

Strengthening antitrust enforcement: Traditional antitrust tools are not applicable in multi-sided markets, where one side pays zero prices. If Uber were to prohibit its drivers from working for Lyft, it would make it much harder for Lyft to compete. Similarly, when Uber offers a loyalty discount, it makes it more difficult for its customers to switch to Lyft. Exclusive dealings and loyalty discounts, which are common in most markets, deserve much closer antitrust scrutiny in DP markets because these markets have a natural tendency to monopolization: Many practices that are benign in other markets could easily become the straw that breaks the camel's back in DP markets.

3. Reducing the Power of Data

DPs use their control over specific types of data to increase their market power and, more importantly, their political power. There is a desperate need to better understand how DPs are impacting our political environment. This understanding is greatly impeded by the fact that independent academics and regulators do not have access to all DPs' data, nor do they have the possibility of doing tests on the various ways that different interfaces may impact behavior. The effects of Dark Patterns discussed above are startling. Platforms have been doing similar A/B testing for years and know much more about the effectiveness of these practices. Even more worrisome, DPs hire hand-picked academics to undertake the studies they want, selectively releasing them to the public. This double selection severely distorts the evidence available in the public domain. Either access is opened up broadly, or the burden of proof should be inverted: Qualitative or imperfect evidence of harm, when combined with deliberate and severe data restrictions, should be considered *prima facie* evidence of harm.³²

The FTC should be empowered to implement a data access mandate: Congress should empower the FTC to: (i) have access to DPs' internal databases and studies, (ii) perform their own independent research on how platforms impact different areas of our society, and (iii) moderate independent researchers' access to these databases. The FTC is a well-established agency that is accustomed to conducting in-depth investigations and whose Bureau of Economics and Office of Technology Research is amongst the better staffed in the country.

Different types of openness for different types of data: Data openness does not mean a total disregard for privacy protection. While re-identification is a serious problem,³³ the openness of

³¹ Importantly, Saikrishna Kamepalli, Raghuram Rajan, and Luigi Zingales, "Kill Zone" (n.d.), shows that acquisition by incumbents may be problematic when markets present network externalities and switching costs. Expecting that new social media would be acquired by the incumbent, which will incorporate all the new desirable features, customers will be reluctant to switch to new entrants, even when these are technologically superior. This reluctance will severely reduce the market price at which these new entrants will be acquired, discouraging new entry, as the paper seems to find empirically.

³² See Luigi Zingales, "Preventing Economists Capture," in *Preventing Regulatory Capture: Special Interest Influence and How to Limit It* (Cambridge University Press, 2013).

³³ See Luc Rocher, Julien M. Hendrickx, and Yves-Alexandre de Montjoye, "Estimating the Success of Re-Identifications in Incomplete Datasets Using Generative Models," *Nature Communications* 10, no. 1 (July 23, 2019): 3069, <https://doi.org/10.1038/s41467-019-10933-3>.

larger databases enables anonymization techniques that, while not perfect, are important steps in the right direction. There are promising computational technologies—like differential privacy and secure multiparty computation—that can help in this dimension. More importantly, one must acknowledge that different types of data may allow different openness rules for different purposes. For example, regulators should have ample access to DPs’ internal databases. They can also intermediate the access for third parties who respect privacy considerations, such as outside, independent researchers. Regulators can then ensure that a small but significant anonymized sample is made available for a larger use, depending on the trade-off between re-identification risks and the gains from openness. In some cases, it will be up to elected officials to settle the trade-off between privacy protection and competition, as done in Brazil in the Good Payer’s Credit Act, where the data sharing was mandated by law.

Longer-term—the creation of a Digital Authority: The strongest indication emerging from the four reports is the importance of having a single powerful regulator capable of overseeing all aspects of DPs. DPs generate several concerns across different fields, all linked to the power of data. To address these concerns in a holistic way, there needs to be a single regulator able to impose open standards, to mandate portability of and accessibility to data, to monitor the use of dark patterns and the risks of addiction, and to complement the FTC and the DoJ in merger reviews. Countries like the UK are considering the set-up of a Digital Markets Unit. The United States and other nations should follow their example.³⁴

The Dangers of a Digital Authority: As George Stigler would readily point out, a new Digital Authority runs the risk of being captured by industry, becoming a new barrier to entry rather than a promoter of competition. This risk can be minimized, albeit not eliminated, by a careful institutional design. This is one reason why we envision—at least initially—to have the Digital Authority as a subdivision of the FTC, an across-industry authority with a better-than-average record of avoiding capture. Most importantly, the Digital Authority will have to be very transparent in all its activities. The Reports discuss a range of different institutional design mechanisms that can be explored to protect the Digital Authority from capture.

4. Reducing the Political Power of Digital Platforms

Disclosure obligations: As mentioned above, DPs are formidable political machines. The Honest Ads Act, which extends traditional campaign disclosure obligations to the internet, seems the first obvious step in limiting this excessive political power. The second, and probably most important, is to understand how these companies are acting as political agents. To comprehend the political role DPs play, some new disclosure obligations should be in place:

- i. **Non-neutrality:** Platforms should disclose when they voluntarily adopt non-neutral policies for content. For example, if platforms are deliberately demoting content related to specific topics, they should make clear what types of content they are demoting and why. Failure to disclose this information should result in fines or sanctions.

³⁴ See <https://www.gov.uk/government/speeches/pm-speech-opening-london-tech-week-10-june-2019>.

- ii. **Relationship with politicians:** Platforms should disclose when they provide specific support or technical assistance to political parties, candidates, or interest advocacy groups, outlining what type of support has been provided and the outcome of this support. For example, during the 2016 elections, Facebook and Twitter embedded some of their own employees in the campaign teams of the various candidates.³⁵ It is not clear to what extent this is still a practice. There should be full transparency on those efforts, including whether candidates are charged equal prices, receive equal treatment, are granted equal distribution of their contents, etc. In particular, electoral regulators must ensure that candidates who are criticizing companies like Facebook and Google receive equal treatment to candidates who support them.
- iii. **Academic funding:** Platforms should disclose their direct funding to academia and their relationship with academics. This disclosure should include DPs charitable arms and the donations of key individuals, so as to prevent easy gaming of the requirements.³⁶

These are initial recommendations that must be reviewed in 3 years: Overall, the combined disclosure recommendations will go a long way in allowing us as a society to better understand the real positive and negative impacts of DPs. Once we know more, these obligations should be reviewed to both allow for flexibility if they are unnecessary or for more stringent requirements if additional concerns appear.

5. Pro-Consumer Default Rules

The current “notice and consent” system on the type of information that can be collected or shared does not effectively protect individual privacy. At the same time, top-down regulation, which allows bureaucrats to impose their will on how markets should behave, can be overly rigid. An alternative can be the adoption of “consumertarian default rules”; that is, default rules on data protection that follow the preferences of a majority of US consumers. Such defaults should be based on “the results of well-designed, scientifically rigorous studies that elicit consumer preferences, opt-out costs, and knowledge of the rules and alternatives, as well as ignorance and biases of such rules’ potential costs and benefits.” These default rules should also be revisited periodically to account for updates in consumers’ preferences due to technological changes or better education.

Default rules only work if dark patterns are addressed: The growing use of dark patterns implies that the default rules should be “sticky,” that is, there should be stringent constraints on waiving the default in favor of a less data protective setting. Click-through or simple pop-up boxes do not satisfy the waiver. When it comes to data externalities,³⁷ however, sticky defaults

³⁵ See <https://www.reuters.com/article/us-facebook-election-usa/facebook-to-drop-on-site-support-for-political-campaigns-idUSKCN1M101Q>.

³⁶ A good example is the controversy around the New America Foundation, when Eric Schmidt forced the think-tank to stop criticizing Google. See <https://www.nytimes.com/2017/08/30/us/politics/eric-schmidt-google-new-america.html>.

³⁷ See Omri Ben-Shahar, “Data Pollution,” *University of Chicago Coase-Sandor Institute for Law & Economics Research Paper No. 854*, August 2018, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3191231.

might not be enough. There is a need to design top-down regulations to limit the overall amount of data collected and shared.

Who Monitors the Defaults? The consumertarian approach is very appealing, but it has one potential weakness: Who would be responsible for defining/monitoring these default rules? Competition and experimentation in this area are important, as they can allow the testing of different approaches. For example, some states may wish to delegate this responsibility to regulators. Others may privilege courts. An interesting alternative may be for some authorities to establish a safe harbor for companies that pre-commit in advance to the result of product-specific studies, which must be periodically rerun. If a company fails to qualify for the safe harbor, it is exposed to additional legal liability in litigation if a plaintiff can prove that the default fails the consumertarian standard. Finally, federal regulations should be a floor—states should be free to establish different, more protective requirements as they deem appropriate.

6. Changing Section 230:

Section 230 of the Communications Decency Act was conceived for early internet providers. These companies did not edit content and thus were not liable for it, as phone companies were not liable for the content of the phone calls they transmitted.³⁸ However, many social media platforms, like Facebook and YouTube, actively recommend content and monetize it. If YouTube's auto-play keeps leading people to extreme views or deeply deceptive content, as independent research and YouTube itself seem to indicate,³⁹ and YouTube is monetizing this specific content with ads, YouTube should be considered responsible for biases of its own algorithms, just as banks are responsible for discriminatory bias in algorithmic lending. In this respect, Section 230 is a major subsidy to DPs, favoring them in their competition with traditional media companies.

Connecting promotion and liability: When social media actively promote content and make money out of the promoted content, we think they should be subjected to the same editorial responsibility as newspapers. By contrast, if they limit themselves to solely transmitting information, with no editing or promotion, like WhatsApp, then they should be treated like phone companies and be exempted from content liability. Such a system would go a long way in levelling the playing field between DPs and traditional media.

7. Preserving Diversity in News Provision

Prevent further media concentration: Media markets are concentrating fast. To preserve diversity in news sources, the FTC and the DoJ should incorporate media plurality as a key metric in merger reviews, as done in the UK. Media plurality would be measured as the share of the attention devoted by consumers to different media sources.

³⁸ Even for phone companies this is starting to change, as shown by the new regulations preventing robocalls. See <https://www.wsj.com/articles/large-telecoms-state-enforcers-make-pact-to-combat-robocalls-11566489602>

³⁹ See <https://www.nytimes.com/2019/08/11/world/americas/youtube-brazil.html> and <https://www.theverge.com/2019/1/25/18197301/youtube-algorithm-conspiracy-theories-misinformation>.

Need for experimentation: The potential reduction in political accountability due to the demise of local newspapers is a serious issue. Yet, some of the potential solutions (like government subsidies to existing newspapers) might cause bigger distortions than the ones they are meant to fix. For this reason, we advocate a two-pronged approach. On the one hand, there is the need for more empirical research to assess the actual welfare costs produced by lack of political accountability. On the other hand, various alternative solutions to generate political accountability should be experimented on a small scale: from prize money for the best investigative reporting pieces to the voucher system for newspapers proposed by the News Media Report, and from prizes for local whistleblowers to a system of random federal auditing of local governments, as done in Brazil.⁴⁰ This is a major topic—concentration in local news is real and growing.⁴¹ We need pilot projects and experimentation now so we can scale-up the most successful interventions.

8. If all else fails ...

The winner-take-all characteristics of many digital markets suggest that even if all the proposed policies are implemented, in some markets we would still find ourselves in a world of few companies (sometimes just one) with outsized market and political power. This is particularly true in the search provider market, where there are increasing returns to scale and thus it is efficient to have a single search provider.

The imposition of fiduciary duties on these companies is an interesting alternative: Monopoly is worrisome in general, but it is particularly problematic in the case of companies that also enjoy great political power. Even Milton Friedman (1962), a father of the idea that a board's sole responsibility is to maximize profits, recognizes that this duty should apply only in competitive markets. In monopolies, the maximization of profits can lead to severe distortions. For this reason, Hart and Zingales (2019) propose the imposition of an additional fiduciary duty on the boards of monopolies: a fiduciary duty towards society.⁴² Policymakers should seriously consider imposing such obligations on DPs like Google, which operate in markets that are or tend to become natural monopolies.

⁴⁰ See Eric Avis, Claudio Ferraz, and Frederico Finan, "Do Government Audits Reduce Corruption? Estimating the Impacts of Exposing Corrupt Politicians," *Journal of Political Economy* 126, no. 5 (2018): 1912–64. , as representative of a growing literature on the topic.

⁴¹ See, <https://www.newyorker.com/magazine/2018/10/22/the-growth-of-sinclairs-conservative-media-empire>. Of course, the problem is not that Sinclair is conservative. The problem is that monopolies in media suppress dissent and lead to uniformity. They should be prevented regardless of whether they are liberal or conservative.

⁴² One might argue that a company like Google or Facebook would never tweak its algorithms to adversely impact its political adversaries out of concern for its reputation. If reputation is indeed sufficient to make monopolies behave in the interest of society as a whole, then this proposal is, at worst, redundant. Yet, it could be used as an insurance policy to prevent bad behavior when concern for reputation falls short.

IV. CONCLUSION

DPs have produced trillions of dollars of consumer surplus, but they have also raised novel policy challenges. These challenges are not confined to the economic sphere, but touch multiple disciplines. To try and address these challenges, the Stigler Center has gathered more than 30 leading academics across multiple disciplines to conduct a year-long analysis on the nature of the problems posed by DPs and their potential solutions. The full product of this effort can be found [here](#). This Policy Brief summarizes the key findings for a non-technical audience and assembles a coherent set of policy proposals.

Some will regard our proposals as too timid, while others as too radical: We regard them as the minimum response to address the new challenges raised by DPs. As the Digital Revolution is advancing, the political system is called to manage the effects of this revolution on society. Without a public debate, the policy response risks being dominated by the interests of the DPs themselves. The independent nature of this report makes it the ideal starting point for such debate, which we hope will be intense and fruitful.

**George J. Stigler Center for the Study of the Economy and the State
The University of Chicago Booth School of Business**

Stigler Committee on Digital Platforms

Market Structure and Antitrust Subcommittee

Report

July 2019

Subcommittee members

Fiona Scott Morton (chair), Theodore Nierenberg Professor of Economics, Yale School of Management

Pascal Bouvier, Managing Partner, MiddleGame Ventures

Ariel Ezrachi, Slaughter and May Professor of Competition Law, The University of Oxford

Bruno Jullien, Senior Researcher, CNRS, Toulouse School of Economics

Roberta Katz, Senior Research Scholar, Center for Advanced Study in the Behavioral Sciences, Stanford University

Gene Kimmelman, President and CEO of Public Knowledge

A. Douglas Melamed, Professor of the Practice of Law, Stanford Law School

Jamie Morgenstern, Assistant Professor of Computer Science, Georgia Institute of Technology

Acknowledgements

We thank Luigi Zingales and Guy Rolnik for inviting us to spend our time on this interesting and complex project, and all the wonderful staff at the Stigler Center. We particularly thank Rachel Piontek for her logistic and document skills, and Filippo Lancieri for guiding the intellectual work along from the very start. Many thanks for outstanding help with content go to students Doni Bloomfield, Rachel Cheong, and Steffi Ostrowski of the Yale Law School. We are grateful to the reviewers at other institutions who took time to give wise and helpful comments. We additionally want to thank our colleagues in other nations whose reports came out before ours and provided stimulating ideas and frameworks. We particularly build on the “*Unlocking digital competition: Report of the Digital Competition Expert Panel*” chaired by Jason Furman for the Government of the United Kingdom and the European Commission’s Special Advisors’ report “*Competition Policy for the Digital Era*.”

DISCLAIMER:

The purpose of these preliminary reports is to identify what are the new challenges digital platforms pose to the economic and political structure of our countries. These reports also try to identify the set of possible tools that might address these challenges. Yet, there is potential disagreement among the members of the committees on which of these problems is most troubling, which tools might work best, whether some tools will work at all or even whether the damage they might produce is larger than the problem they are trying to fix. Not all committee members agree with all the findings or proposals contained in this report. The purpose of these preliminary reports, thus, is not to unanimously provide a perfect list of policy fixes but to identify conceptual problems and solutions and start an academic discussion from which robust policy recommendations can eventually be drafted.

*** The Committee in-person meetings were partially supported by a grant from the Alfred P. Sloan Foundation, whom we also thank for supporting this project**

Table of Contents

Introduction..... 27

Executive Summary 29

A. Attributes 29

B. Harms..... 30

C. Solutions..... 31

Important Features of the Digital Business Environment 34

A. The Structure of Digital Markets 34

 1. An Economic Perspective on the Digital Market Structure 34

 2. Barriers to Entry and Expansion 40

D. How Big Data Affects Digital Products and the Sale of Advertisements 44

 1. What Makes Data Big? And Does Big Data Harm Consumers? 44

 2. Data Control and Ownership..... 51

 3. Digital Platforms are Characterized by Free Services 54

Problems Arising in Digital Markets..... 57

A. Quality Harms..... 58

 1. Personal Data and Behavioral Economics 58

 2. The Harms from an Advertising-Supported Business Model 61

 3. Online Exploitation and Addiction 64

B. The Assessment Problem..... 66

 1. Measuring Consumer Welfare in a Behavioral World..... 66

 2. Assessing the Social Welfare of Advertising..... 67

C. Harms to Investment and Innovation 68

 1. Rents 68

D. Harm to Entry, Including Disintermediation..... 71

 1. The Practical Consequence of Entry Barriers 71

 2. Incumbent Incentive to Leverage Entry Barriers 71

 3. Disintermediation and Foreclosure of Potential Platform Entrants 72

 4. Foreclosure of Complements to Capture Rents 73

E. Harm to Innovation 74

 1. Competition Promotes Innovation 74

 2. Entry Barriers and Innovation..... 75

 3. The Reward for Innovation 76

4. Current Trends in Innovation..... 77

F. Old Wine in a New Bottle, or a New Reality?..... 79

Solutions..... 80

A. Will Markets Self-Correct When Competition Problems Arise? 80

B. US Antitrust..... 81

1. Basic Principles..... 81

C. Special Challenges Presented by Technology Platforms 87

D. Could a Reformed Antitrust Law be Effective in Dealing with Digital Platforms?..... 92

1. Reform by What Means? 92

2. Reform of Antitrust Law through a General Tightening 94

3. Specific Areas of Possible Antitrust Reform [designed for the specialist reader] 95

4. A Competition Court..... 99

5. A Regulatory Partner Could Enhance Effective Antitrust Enforcement..... 99

E. Regulation 100

1. The Digital Authority..... 104

2. Menu of Regulations..... 107

3. DA-Enforced Remedies for Antitrust Violations..... 117

4. Aligning other policies with competition..... 119

5. Adjudication Process 119

Conclusion..... 119

References 122

Introduction

When the global, interactive, many-to-many communications network called the internet became available for broad public use in 1995, people were overjoyed with the new power at their fingertips. It is easy to forget that in the early years of the internet, that power was limited to visiting shops and getting news from the relatively small number of enterprises that had built websites on the World Wide Web. Over the past 25 years, that power has exploded with head-spinning velocity: Today, there is no area of human life that has not been affected by the technological innovations made possible by the internet. We now buy goods and services, do banking, pay bills, find information, and talk with multiple groups of friends and acquaintances on the web. The speed, scale, and scope of the internet, and of the ever-more powerful technologies it has spawned, have been of unprecedented value to human society.

History teaches us that social institutions must adapt after major technological advances. In just the past century, we saw this after the introduction of the automobile, the airplane, radio and television, and cell phones, to name but a few examples. New occupations—for example, bus drivers—arise to replace others—for example, buggy whip producers. At the same time, new behavioral norms and expectations replace those associated with past technologies—consider the social and economic changes wrought by the automobile, or the changes in both the rules and norms of the workplace in the age of mobile devices. As unintended and unforeseen social problems and harms arose, society responded with governance mechanisms aimed at addressing the problems or harms without impeding the clear benefits associated with the advances. These have ranged from the adoption of new social norms to the creation of new laws and regulations. Examples abound, encompassing everything from crosswalks and traffic lights to legal remedies addressing unfair competition in the marketplace. Ideally, the goal is to steer technological advances to ensure widespread benefit without attendant widespread harms—to protect and preserve innovation and advancement while minimizing harms so that all of society reaps net benefits. Such a goal often involves government intervention and always involves tradeoffs as society wrestles with the prevalence and cost of harms and how to balance them against the prevalence and size of benefits. While often messy, this is a healthy and desirable debate.

This working group came together to address specific problems arising from the digital platforms' reach, scale, scope, and use of data. We were asked to examine concerns stemming from the market structure contemporary platforms have created, and to investigate their competitive behavior, including the consequences of network effects that can create barriers to entry for new innovators and entrench incumbents. The global nature of many of today's platforms, a result of their scale, scope, and business models, creates novel complexities and considerations, particularly a concern that the digital platform may be a unique combination of economic forces that require both new analysis and new public policy. Regulatory authorities throughout the world are now turning their attention to these same questions. This report

contributes to this international analytical project by providing some of the necessary frameworks and inputs. We intend it to be a complement to other recent work, as experts across the world wrestle with how to ensure that markets remain open and healthy, allowing beneficial technological and social advancements to continue. Many of our conclusions and suggestions echo the findings of reports that have come out in the past year, and we hope they will be helpful to those reports not yet released. The list of antitrust experts and agencies working on this problem includes Australia, the United Kingdom, Germany, the European Commission, France, Israel, and Japan.⁴³

The issues are global in their scope, and these various jurisdictions are all engaged in analyzing how best to ensure that societies in general, and competition in particular, continue to thrive in the Digital Age. Our charge was restricted to market structure and competition, while other committees considered the equally important topics of the impact of digital platforms on politics, the media, and the nature of privacy. We note that monopolies can concentrate political power, reduce media plurality, and provide insufficient competition on dimensions such as privacy. In this way the findings of this report and the others are linked and quite consistent. Digital markets and platforms have already delivered great benefits to consumers, and the global concerns that have surfaced relating to actual or potential consumer harms may require action to ensure that the benefits are not undermined. Our report concludes that with deliberate government action to protect competition and consumers, the benefits from innovative firms could be even greater and more equitably spread, ensuring that the public is not short-changed in firms' pursuit of profit. Accordingly, this report is offered in the spirit of ensuring a future of continued technological and economic progress and social well-being as we move further forward into the Digital Age.

⁴³ Crémer (2019); Furman (2019); Australian Competition (2018); Japanese Ministry (2018); Schweitzer (2018); French Competition (2018); Reishut HaTachrut (2018).

Executive Summary

A. Attributes

The market structure and antitrust report begins by discussing the characteristics of digital markets. These markets often have extremely strong economies of scale and scope due to low marginal costs and the returns to data. Moreover, they often are two-sided and have strong network externalities and are therefore prone to tipping. If so, the competitive process shifts from competition *in* the market to competition *for* the market. This combination of features means many digital markets feature large barriers to entry. The winner in these settings often has a large cost advantage from its scale of operations and a large benefit advantage from the scale of its data. An entrant cannot generally overcome these without either a similar installed base (network effects) or a similar scale (scale economies), both of which are difficult to obtain quickly and cost-effectively.

Additional barriers to entry are, ironically, generated by the very consumers who are harmed by them. Consumers do not scroll down to see more search results, they agree to settings chosen by the service, they single-home on one platform, and they generally take actions that favor the status quo and make it difficult for an entrant to attract consumers. In general, the findings from the behavioral economics literature demonstrate an under-recognized market power held by incumbent digital platforms.

The theme that runs throughout the report is the difficulty of entry into digital platform businesses once an incumbent is established. Whether the entrant is vertical or horizontal, has succeeded to some degree, is nascent, is a potential entrant, or is a large platform in an adjacent space, its existence improves consumer welfare. Either the entrant provides more choice, different features, and a chance of higher quality, or the threat of those outcomes spurs the incumbent to provide lower prices, higher quality and innovation, and to do so more quickly.

The role of data in digital sectors is critical. Personal data of all types allows for targeted advertising to consumers, a common revenue model for platforms. The report shows that the returns to more dimensions and types of data may be increasing, which again advantages incumbents. Consumer data in the United States is not regulated in any way that gives useful control or privacy to consumers; and additionally, most consumers have little idea what is being collected about them and re-sold. One way in which digital platforms often exploit their market power – and increase their profits – is by requiring consumers to agree to terms and conditions that are unclear, difficult to understand, and constantly changing, but which give the platform freedom to monetize consumers' personal data.

Digital platforms are characterized by free services. “Free” is not a special zone where economics or antitrust do not apply. Rather, a free good is one where the seller has chosen to set a monetary price of zero and may set other, non-monetary, conditions or duties. It is possible that a digital market has an equilibrium price that is negative; in other words, because of the value of target advertising, the consumer’s data is so valuable that the platform would pay for it. But the difficulty of making micropayments might lead a platform to mark up this negative competitive price to zero. As a result, barter is a common way in which consumers pay for digital services. They barter their privacy and information about what restaurants they would like to eat in and what goods they would like to buy in exchange for digital services. However, in principle, that information has a market price that can be analyzed.

B. Harms

Market power, consumer biases and an ad-supported platform model can generate significant consumer harms. First, market power in advertising markets will result in markups paid by advertisers. Secondly, while behavioral economists have studied consumer biases and firm responses in offline markets, these are swamped by what digital businesses can learn by using high-dimensional, large datasets to explore every nook and cranny of consumers’ many behavioral shortcomings and biases in real time. Framing, nudges, and defaults can direct a consumer to the choice that is most profitable for the platform. A platform can analyze a user’s data in real time to determine when she is in an emotional “hot state” and then offer targeted sales. These tactics reduce the quality of the zero-price content the user experiences on the platform.

In addition to *de novo* entry, platforms fear disintermediation by a partner or complement. If a platform’s partner is able to directly access and serve the platform’s customers, it might take them off the platform entirely, reducing the platform’s profit. A platform that has total control of demand due to control over framing of consumer choices, policies for complements, and technical standards can steer customers to content and complements of most benefit to it. The most privately beneficial content might be owned by the platform itself rather than provided by independent firms that could extract rent or even challenge the platform’s market power in the future. To the extent that consumers single-home, they may not be aware of such steering, or may not have competitive alternatives to which they can turn if they are aware.

Today’s platforms understand that in some settings they can obtain higher margins if they either, make all of the necessary complements themselves, or, position themselves as a mandatory bottleneck between partners and customers. In particular, digital platforms are often very careful to maintain complete control over the user relationship so that they do not face any threat of disintermediation from a complement. These technological and policy choices can be used to reduce the possibility of successful entry by direct competitor. Other strategies such as

exclusive contracts, bundling, or technical incompatibilities can also be used by platforms to restrict entry of competitors. Some of these strategies could be violations of existing antitrust law, as discussed below.

Insufficient competition and entry result in harms to investment and innovation. There is significant theoretical and empirical research that concludes that anticompetitive creation or maintenance of market power will cause a reduction in the pace of innovation. The lessening or blocking of innovative entry is of particular concern given its value to consumers. A VC will usually be wary of outright investing in an innovative startup that will implicitly or explicitly compete head-on with a tech giant. In that case, the best hope might be to be the preferred innovator of a complement and sell its business to the platform at an early stage. However, if entry barriers were reduced, the entrepreneur would not have to settle for a small fraction of the platform's profits, but could compete for all of them and try to replace the platform. This possibility would generate a much larger incentive to innovate.

C. Solutions

While some markets may self-correct, the findings of this report suggest that rapid self-correction in markets dominated by large digital platforms is unlikely.

While US antitrust law has long been flexible in combatting anticompetitive conduct, there is increasing concern that it has been underenforced in recent years. Antitrust law and its application by the courts over the past several decades have reflected the now outdated learning of an earlier era of economic thought, and they appear in some respects inhospitable to new learning. Antitrust enforcement better suited to the challenges of the Digital Age may therefore require new legislation.

Technology platforms present particular challenges for antitrust enforcement. Markets tip and the resulting market power is durable, so even effective antitrust enforcement is unlikely to generate fragmented markets. Nonetheless, enforcement that protects competition on the merits in the first stage and prevents exclusionary conduct in the second stage will help ensure that market-participants make unfettered choices among competing platforms and that entry and innovation are not inhibited by private rent-seeking.

Economists and lawyers will have to develop tools to explain to courts the role of behavioral biases in the creation of market power and in their effect on the quality of content. The existence of zero money prices means that measurement of quality will be critical. The law needs better analytical tools to take into account the impact of potential and nascent competitors and competition. Market definition will vary according to what consumers are substituting between, whether there is competition on the platform between complements, or competition between platforms, or competition between a platform and potential or nascent competitors

regarding possible future markets. The need to identify the specific anticompetitive exclusionary conduct and analyze it may raise enforcement costs given all the possible variants of exclusionary conduct possible in digital markets.

This report details the particular areas where antitrust law will need reform in order to adequately deliver competition to consumers. The report also suggests the establishment of a specialist competition court to hear all private and public antitrust cases which would allow judges to develop some expertise. The committee believes that vigorously enforcing the antitrust laws under these conditions would be likely to increase entry in digital platform industries, competition, and consumer welfare. Moreover, such enforcement would likely result in remedies to restore competition that has already been lost, as well as serve as a deterrent to future anticompetitive conduct.

However, because technology platforms present the enforcement challenges detailed above, even effective enforcement may not be enough to generate competitive digital markets in a timely fashion. Therefore, the report suggests that Congress should consider creating a specialist regulator, the Digital Authority. The regulator could be tasked with creating general conditions conducive to competition. The committee also suggests separating out some types of regulation that will apply to virtually all market participants while other regulation will apply only to companies with bottleneck power. “Bottleneck power” describes a situation where consumers primarily single-home and rely upon a single service provider, which makes obtaining access to those consumers for the relevant activity by other service providers prohibitively costly.

The Digital Authority could routinely collect data on digital transactions and interactions, with an emphasis on data from businesses with bottleneck power. These data – made public to the extent possible – would allow policy makers and researchers to assess the performance of the sector. The DA could have a mandate to create “light touch” behavioral nudges when they will make markets more competitive. An example of a regulation that would enhance competition is data portability. The DA could set up rules that allow users to easily port their data from one service provider to another and monitor compliance. The DA may also promote open standards in such areas as micro-payments and digital identities. Should Congress request it, the DA could oversee a mandate for interoperability in any market where market power has become entrenched and threatens long term harm to competition. The Report also suggests that the DA could carry out a parallel merger review that would be set up to incorporate necessary antitrust reforms and modern standards.

Some regulations could apply only to firms that meet the DA’s definition for bottleneck power. Because the cost of false negatives is high and there is uncertainty, the public interest requires the DA to take a more interventionist approach in these settings. The DA could have

merger review authority over even the smallest transactions involving digital businesses with bottleneck power because nascent competition against these entities is very valuable for consumers. Non-discrimination rules could protect against a complement that is a potential competitor of the platform itself, or one that operates only on the platform as a rival provider of content.

When a company has been found liable for violating the antitrust laws, part of the current process is that antitrust authority devises a remedy to restore the lost competition. Data sharing, full protocol interoperability, non-discrimination requirements, and the unbundling of content from a platform are all tools that the regulator, in conjunction with the antitrust authority, could apply and monitor over time in order to restore competitive markets.

Important Features of the Digital Business Environment

Digital technologies are a central driver of future prosperity, hopefully delivering waves of innovation, efficiencies, and consumer welfare. These technologies have revolutionized the way consumers shop and communicate, the way businesses deliver value, the way people work together in collaboration, and—the subject of this report—the overall dynamics of competition.

These transformative changes, while immensely beneficial, have also triggered growing concerns about the power of a small number of firms to control and influence billions of lives. As an increasing volume and range of commercial activities have been digitalized, society has witnessed the emergence of certain key platforms and gatekeepers and a shift in market dynamics.

This section outlines the key features of the digital environment to set the foundation for the discussion of problems and solutions.

A. The Structure of Digital Markets

1. An Economic Perspective on the Digital Market Structure

Despite the predictions of some early observers that the internet would create competitive—even perfect—markets,⁴⁴ certain digital firms have been able to acquire significant market positions and preserve them over time. Many of the most innovative internet-derived digital markets, such as search engines, social networks, network operating systems, ecommerce, and ride-sharing, are highly concentrated and have been dominated by one or a few firms for a number of years. The lack of entry of competitors in these important markets—despite high profits—suggests either barriers to entry or exclusionary conduct, or both. We first discuss the nature of entry barriers in digital platforms and why they are difficult for an entrant to surmount.

1) A Unique Setting Subject to Tipping

From an economic perspective, there is no single new characteristic that would make competition in digital platforms different from more traditional markets. Rather, it is the coincidence of several factors at a scale that has not been encountered before that makes the problem unique and requires new analysis of market structure and market power. In particular, the platforms with which this report is most concerned demonstrate extremely strong network effects, very strong economies of scale, remarkable economies of scope due to the role of data,

⁴⁴ “The conventional wisdom regarding Internet competition . . . is that the unique characteristics of the Internet will bring about a nearly perfect market.” Brynjolfsson (2000): 563; “The traditional economic view suggests that . . . the Internet should reduce search costs for consumers and thereby reduce prices and make markets more competitive.” Brown and Goolsbee (2002): 482.

marginal costs close to zero, drastically lower distribution costs than brick and mortar firms, and a global reach.

Markets with these combined features are prone to tipping—a cycle leading to a dominant firm and high concentration. Digital markets are prone to tipping for two primary reasons. First, because fixed costs play such an important role in digital markets, these markets feature especially large returns to scale. Second, many digital markets are driven by network effects that strengthen large incumbents and weaken new entrants.

When markets are prone to tipping, the competitive process shifts from competition *in* the market to competition *for* the market. In that case, consumers may only benefit from competition among several firms for the relatively short time period in which the firms compete to be the ultimate winner of very large economic profits. The winner’s monopoly profits serve as the inducement for entry and investment. After a market has tipped, a potential rival for the market can only overcome the incumbency advantage of established networks through significant innovation. However, even an innovative entrant may not be able to create competition that benefits consumers in the presence of the factors that led to the tipping, leaving open a role for public policy to allow for competition for the market. Moreover, a competitor is even less likely to enter the market if the incumbent platform is able to leverage its powerful position to disadvantage or exclude potential entrants.

There are many well-known problems that follow from lack of competition, including higher prices, less innovation, and lower quality in all its forms. Policy may be needed to address the cause of such symptomatic problems. In the view of this committee, protecting entry for existing and potential competitors is the most important way to protect or improve consumer welfare in digital platforms. Initial competition for a market should be conducted on the merits without any anticompetitive practices, and later entrants should face a level playing field and no exclusionary conduct as they contest the market. Regulation may be required to prevent incumbents from erecting improper barriers to entry.

As it is traditionally recognized, the reward for the winner’s innovative activity is the ability to extract rent from the platform through the exercise of market power.⁴⁵ However, we highlight three deviations from this principle. First, the winner must have “won” on the merits of its product, without the use of any anticompetitive conduct. Correcting illegal practices by dominant firms that have won in a tipped digital market can be difficult, but is necessary or the firm will earn profits from its illegal behavior—and harm consumers along the way. Second, the reward of the entire market that propelled the winner to innovate must be available to the next entrant. If the incumbent is able to withhold those rents by excluding the entrant, or reduce those

⁴⁵ See *United States v. Alcoa*, 148 F.2d 416, 430 (2d Cir. 1945) (“The successful competitor, having been urged to compete, must not be turned upon when he wins.”) (Hand, J.).

rents by limiting the entrant's share, then the pace of innovation will slow. The economic literature demonstrates that vibrant innovation and entry is the most important source of consumer welfare over time—hence the focus of the report on ensuring that entry functions well both when a market is created and thereafter.⁴⁶ Third, there are sources of platform rents that society may determine through appropriate regulation should not be part of the winner's reward. For example, investments in complements offered on the platform may be best incentivized and encouraged to compete if the platform owner is not entitled to exclude them. (As used in this report, a complement is a good or service offered on or through the platform that increases the value of the platform to a consumer.) Digital markets at issue today may not have featured free entry and lack of anticompetitive conduct in the past. This report covers that case by describing how antitrust can be used to look backward, as well as forward, for antitrust violations. We also suggest regulation that can look forward to protect entrants, and thereby consumers, in these other cases.

2) Increasing Returns to Scale

Digital markets are used to exchange information goods and services. Typically, information goods involve increasing returns to scale because their production requires a fixed cost and no or little variable cost.⁴⁷ In other words, when an additional user is served, costs do not go up proportionately. For example, an eBook, once produced, can be distributed at almost no cost to all users with access to the internet. The same holds for information services that are subject to fixed design and development costs and fixed maintenance and updating costs; Google can update Google Calendar for 100 million users with similar fixed expenses as would be needed for only a fraction of such users.

In contrast to traditional media or cultural markets that have had to incur physical distribution costs, digital markets are largely able to avoid such costs. The fact that information services can be delivered to any geographic location with no or minimal cost is one reason for the abundance of such services. The capacity to reach a large scale at small cost changes the nature of business growth. While a traditional business often starts with local implementation followed by gradual expansion through investment as reputation and financial resources increase, many online businesses aim at rapid large-scale expansion. This rapid growth may reduce the length of the competition-for-the-market phase, as market winners can establish dominance and begin exercising their market power quickly. It took only five years for Facebook, the “move fast

⁴⁶ See, e.g., Farrell and Klemperer (2007); Wen and Zhu (2018).

⁴⁷ Increasing returns to scale occur when the average unit cost decreases when sales expand.

and break things” company, to go from a million users in 2004, the year of its founding, to more than 350 million users in 2009, when it overtook MySpace for good.⁴⁸

The implications for market structure are well known: increasing returns to scale lead companies to invest in fixed costs in order to have the best product to attract customers.⁴⁹ Then, with a larger customer base, the firm can enjoy lower average costs per consumer, allowing it to make an offer to consumers that is attractive in both quality and price. The increasing returns to scale create barriers to entry: New firms cannot offer the quality of the incumbent without the same large-scale operation to pay for the fixed costs. But the firm can only achieve a large scale if quality is high. Thus, a potential entrant, foreseeing that it will not be profitable at the smaller scale, will not enter the market to challenge the incumbent.

3) Economies of Scope

Platforms also have powerful economies of scope in the form of their relationships with users and brands. In digital markets specifically, scale offers an additional advantage. Firms can apply machine learning to extensive data sets to improve their products and expand their activities into new areas. Because machine learning yields better insights when it is trained on larger datasets, firms with access to large amounts of data can raise the quality of their services in ways that smaller firms cannot. This creates a form of dynamic economies of scale, allowing large firms with large amounts of data to raise product quality at lower costs than small firms. Firms may also be able to leverage the data, or the insights due to machine learning, that they receive from an existing service to enter into an adjacent market with a higher quality product, demonstrating a novel form of economies of scope. Combining mapping software in a platform that already offers email, for example, allows that platform to offer a higher quality restaurant recommendation product. Moreover, firms serving a larger customer base with a greater variety of products are able to generate more advertising revenue per consumer through more effective targeting.⁵⁰ The development of machine learning technologies and data analysis is a source of increasing returns to scale and scope that can contribute to digital market concentration.

⁴⁸ Albanesius (2009); Sedghi (2014). *See also* Levy (2014), reporting Mark Zuckerberg noting that “[w]e’ve changed our internal motto from ‘Move fast and break things’ to ‘Move fast with stable infrastructure’.”

⁴⁹ *See generally* Sutton (1991).

⁵⁰ Such economies of scope are leading to a rapid expansion of the capture and use of personal information by firms. For example, many large firms are starting to expand beyond internet platforms in order to collect more data from consumers. The rise of voice-assistant products (including Microsoft’s Cortana, Apple’s Siri, Google’s Google Voice, and Amazon’s Alexa) can enable firms to gather information from offline consumer behavior and then feed that information into online advertisement-targeting algorithms. One study estimates that major internet platforms, data brokers, credit card companies, and healthcare data companies derived nearly \$76 billion in 2018 from selling personal consumer information directly or indirectly via targeted ads. Shapiro and Aneja (2019).

4) Network Effects

Most information services involve some form of positive network effects, in which consumer benefit grows as the number of users increases. The simplest network effects flow directly from interacting with other users, as in social networks or peer-to-peer services, such as eBay or Venmo, where the more users there are on the network, the richer the users' experience is likely to be. (In these examples the platforms choose not to be interoperable, so the network effects apply to a single firm, rather than an industry.) Multi-homing lessens network effects because a consumer can enjoy the size of both networks, rather than having to choose one. For example, a consumer that carries both Visa and American Express credit cards can shop at stores that accept either card, or both. Many other network effects are indirect, in that they are mediated by a “complement” to the network. A complement is a good or service that increases the value of another good or service to a consumer. For example, the ability to not only make calls, but also to play music on a handset increases the value of the handset to users. Complements today often come in the form of applications (“apps”) or a specific type of content. The more complements a platform has, the more popular it is with users. Just as customers of a popular app store receive an indirect network effect when more and better developers are attracted to app stores with big customer bases, customers of a widely-used social media site benefit from the many games designed for that social media site, which in turn are driven by the large number of consumers. These network effects can also be seen in recommendation systems or driving directions that exploit larger datasets of users' purchasing behavior or travel paths to offer higher quality advice.

Markets with network effects are prone to concentration because consumers benefit from being on the same network as other users. No one wants to be on their own social media site. However, when network effects are exhausted relatively quickly (as can occur, for example, in messaging apps⁵¹) or when there is heterogeneity in preferences (e.g. teenagers prefer not to be on the same social network as their parents⁵²), the market structure may be oligopolistic.

Some indirect network effects are multi-sided: A user of one type (e.g., a buyer) benefits from the participation of other types of users (e.g., sellers). For instance, in ecommerce platforms, which intermediate trade between sellers and buyers, a buyer does not directly benefit from the presence of other buyers but does benefit from the presence of more sellers—who are in turn attracted by the presence of the buyers. Multi-sided network externalities are prominent on the internet for two reasons. First, business models like ecommerce are plentiful. Second, and

⁵¹ To send a message to someone, a user only needs that single person to be on a particular messaging app, rather than everyone they would want to send a message to. This explains why many messaging apps can live alongside one another—WhatsApp, Snapchat, SMS, and Facebook Messenger all have significant customer bases. *See* “Most Popular Mobile Messaging Apps” (2019).

⁵² Sweney (2018).

more importantly, a great deal of activity on the internet is financed by paid advertising on the relevant site. Advertising-financed services are platforms exploiting two-sided network effects between advertisers and consumers. Advertising can be more valuable when there are more consumers viewing a site, while consumers are attracted by content that may be of higher quality when financed by many advertisers competing for space.

On multi-sided platforms, one or more sides may be subsidized if their participation attracts paying customers on the other side. Thus, we may see a monetary price of zero in one part of a very lucrative business model. Shoppers on eBay can buy without paying fees because eBay understands that fees would drive them away, whereas their presence draws in sellers who can be charged fees. Users of Gmail pay a monetary price of zero but allow Google to read their email so that advertisers can market to them based on personal information. Google can then charge a high price for the ads. Payment platforms such as credit cards or PayPal similarly charge end consumers low or negative fees because a large group of consumers bring in retailers who pay the payment platform for access.

Network effects can lead to consumer-friendly competition at early stages. Economic theory and market observation indicate that during the phase when competitors are all trying to tip the market towards themselves, they compete intensely. This phase is characterized by vigorous competition between firms trying to build market shares and generate bandwagon effects. For example, Uber and Lyft have hotly contested the market for ride-sharing—and spent billions of dollars subsidizing riders’ fares along the way.⁵³ One 2016 estimate suggested that payments from Uber customers covered only about 40% of the cost of their rides.⁵⁴ If network effects are strong, however, the market will tip in favor of one competitor, who then becomes the monopolist.

5) Low Marginal and Distribution Costs

Digital goods often have low to zero marginal costs of expansion to another user, as mentioned above. Distribution costs, one of the major expenses of expanding in a brick and mortar world, are largely also zero on the internet. Indeed, poor internet infrastructure in the United States and elsewhere may be the main distribution cost for many digital platforms.⁵⁵ Some digital platforms do have positive distribution or marginal costs such as a piece of hardware, ecommerce warehouses, or maintenance of scooters for example.

⁵³ Uber “has lost \$12bn from its operations in total since 2014. ... ‘We will not shy away from making short-term financial sacrifices where we see clear long-term benefits,’ wrote Dara Khosrowshahi, [Uber’s] chief executive.” Bond and Bullock (2019). In 2018, Lyft “lost more than \$900 million after expenses.” McArdle (2019).

⁵⁴ Kaminska (2016). Many of the paradigmatic firms of the dot.com bubble also followed this strategy of subsidizing for growth—often sending the firm to oblivion. Bensinger (2012).

⁵⁵ See, e.g., FitzGerald (2019); Sheetz (2019); Google (2010).

6) *Global Reach*

Lastly, many of the digital platforms we discuss in this report have a global reach. Those that are focused on licensed digital content, rather than user-generated content, must negotiate licensing and other legal issues in new jurisdictions. Reaching consumers in many languages is another cost of expanding globally, as is dealing with local transportation regulations. However, the total cost of expansion is generally lower than in traditional brick and mortar businesses.

2. **Barriers to Entry and Expansion**

In this section, we explore the implications of the characteristics discussed above and the possible effects of the increased market power created by barriers to entry. Outlined below are some of the drivers that make market power more sustainable and disruption less likely in digital markets.

1) *Cost and Benefit Barriers to Entry*

As mentioned above, network effects and scale economies are two significant barriers to entry. The winner in these settings has a massive cost advantage from its scale of operations, and a massive benefit advantage from the scale of its data. An entrant cannot generally overcome these without either a similar installed base (network effects) or a similar scale (scale economies), both of which are difficult to obtain quickly and cost-effectively. It is possible for an entrant to arrive with scale and an installed base because it is a competing digital platform. Such entrants may create effective competition for an incumbent platform that benefits consumers.

Barriers to equivalent data resources, a side effect of not having the history, scale, or scope of the incumbent, can inhibit entry, expansion, and innovation.⁵⁶ The same effects that drive the quality of digital services higher as more users join—a positive feedback loop—makes the strong stronger and the weak weaker.⁵⁷ Data feeds the development of algorithmic and AI training processes that enables more profitable exploitation of consumer attention through advertising. A data advantage over rivals can enable a company to achieve a virtuous circle of critical economies of scale leading to network effects, and a competitive balance in its favor, leading to the gathering of yet more data. A new entrant is likely to experience this in reverse—a vicious cycle—as it fails to surmount the entrance barrier.

Choosing a business model that is scalable, and has strong economies of scale, is of paramount importance to creating entry barriers. A social media platform that chose to monitor harmful content might need to hire many workers to keep ahead of users that game the

⁵⁶ The key themes below are discussed in greater detail in Ezrachi and Stucke (2016). *See also* Ezrachi and Stucke (2018) (on file with authors).

⁵⁷ OECD (2014): 29 (citing Shapiro and Varian 1999).

algorithms. Hiring workers is costly, slows down the ability to grow, and makes the platform less profitable both in the short run and in the long run by limiting economies of scale. A platform aiming to tip the market in its favor quickly will not choose this business model if it can gain consumers at lower cost.

Because cost and benefit barriers are so powerful, digital platforms have powerful incentives to pull slightly ahead of any rival. Due to the positive feedback loop described above, a small advantage can turn into a large one. This powerful incentive to disadvantage a competitor raises the possibility that some competitors will choose to violate the antitrust laws in order to stay ahead. In general, to maintain or improve their competitive advantage, incumbents have strong incentives to limit openness or interoperability and to be averse to data-portability policies. For example, in a recent case, the upstart ethernet switch firm Arista alleged that the incumbent Cisco encouraged openness and interoperability while it was dominant, but once a rival mounted a serious challenge, claimed copyright protection on the previously open user interface, thus cutting off the entrant from locked-in users.⁵⁸

2) Barriers to Entry Created by Consumer Behavior

Market power is, ironically, generated by the very consumers who are harmed by it. Consumers do not replace the default apps on their phones, do not scroll down to see more results, agree to settings chosen by the service, and take other actions that may look like poor decisions if those consumers like to choose among options and experience competition. Often the actions needed to generate choice for the consumer seem trivial, such as a download and installation, opening another app, or a few clicks. Consumers make these “mistakes” because of inherent behavioral biases such as discounting the future too much and being too optimistic. The situation is worse when the information needed to counteract bias is hard to obtain. For example, consumers tend not to run the same search on a different search engine to compare the results, so they may never find out the relative quality of the default search engine they use.

Research into behavioral biases and their consequences is several decades old and has already produced ample and convincing evidence of the nature and existence of these biases on the part of consumers. In 2002, Daniel Kahneman received the Nobel Prize in Economic Sciences for his pioneering work on behavioral economics.⁵⁹ Fifteen years later, the Nobel Prize was awarded to Richard Thaler in recognition of the progress of research in this domain and of the huge amount of knowledge that has been produced.⁶⁰

⁵⁸ Wolfe (2018); Clark (2016).

⁵⁹ “Daniel Kahneman—Facts” (2019).

⁶⁰ “Richard Thaler—Facts” (2019).

In his Nobel lecture, Thaler pointed to the key insight from Kahneman and Amos Tversky: people exhibit “bounded rationality,” meaning that in a complex world they use consistent rules of thumb to make predictions and decisions. Thaler called on economists to use “psychological realism” to improve their understanding of human decision making.⁶¹ That psychological realism combined with the economics toolkit produced the field of behavioral economics.

Behavioral economics has had a profound influence in the conduct of economic policy that will become even more prevalent as more knowledge is digested and applied.⁶² It is of great relevance for our understanding of internet economics because, as information flows improve and some physical barriers are removed, human factors are more likely to provide the frictions that have increasing effects on market outcomes. These frictions in decision-making, and the fact that consumers can be manipulated to take advantage of their biases, render consumers sticky—that is, people are slower to move to a superior product than they would be absent the manipulation. This in turn makes demand less contestable and less favorable for an entrant.

One lesson from behavioral economics is that small differences (nudges) in how choices are presented can have large effects on what choices people make.⁶³ A nudge to use a particular browser as a default, for example, can entrench a platform’s browser. Another lesson is that consumers overweight their immediate benefit relative to their welfare in the future. A consumer searching for a solution to a particular problem will be inclined to click or use the first result or recommendation, rather than searching on another page or scrolling down to examine many listings. The tendency to choose in this manner entrenches the market power of the platform that can control the display of content. Similarly, consumers’ preference for instant gratification may lead them to sign away privacy rights they otherwise say they value.⁶⁴ This allows incumbent platforms to gather data from these consumers that further entrenches their market position. In general, the findings from behavioral economics demonstrate an under-recognized market power held by incumbent digital platforms.

A second way consumers create entrenched market power is by single-homing. A multi-homing user, for example, checks the price of a ride on both Uber and Lyft each time she needs a car. A user that single-homes bestows market power on the platform she uses exclusively because advertisers and other content providers *can only* get the user’s attention by going through that platform. While users sometimes have the ability to employ multiple services, there

⁶¹ Thaler (2018).

⁶² See, e.g., Sunstein (2013); Thaler and Sunstein (2003).

⁶³ Thaler (2018): 1283.

⁶⁴ Acquisiti, Taylor, and Wagman (2016); Acquisti, Brandimarte, and Loewenstein (2015); Acquisiti (2004); see also “Americans’ Attitudes About Privacy” (2015): 17 (reporting that 74% of survey respondents believe it is “very important” to be “in control of who can get info about you”).

is usually a convenience cost to doing so. Making multi-homing easier will be a key element in encouraging competition.

3) Barriers to Entry Created by an Incumbent Rival

Because platform market power is higher when users single-home, platforms try to get users to do so. Sometimes this results in higher quality services: A mapping service can alert a user that it will take 45 minutes to get to the airport when it sees a flight entry on the user's calendar. If the user buys those services from different platforms, she does not get the alert and experiences lower quality. However, it is often possible for the data from one service to be read by a rival, so platforms may also encourage single-homing by preventing interoperability.

By limiting a rival's access to data, dominant firms can make exclusive reliance on their service either inevitable or the clearly best decision. Sharing data or allowing access to certain pieces of information is often feasible at a technical level, but it is not normally in the interest of the platform that could lose its users. For example, Google has been able to limit its users from visiting (multi-homing on) competitors such as Yelp by displaying their information in the search window directly.⁶⁵ To some extent, the limited switchability to the rival is driven by asymmetric information in two places. First, consumers are generally not aware of, or attentive to, the costs of failing to switch; for example, the quality of the content and the extent to which the platform steers consumers to inferior content or product choices may be hard to see. Second, there is often no way to compensate the consumer directly with lower prices—as an entrant such as Warby Parker can do by selling eyeglasses directly to consumers at a lower price—because the money price of services from a platform is often fixed at zero for both the incumbent and the rival.⁶⁶

For all of these reasons, digital platform market power can become entrenched. The United Kingdom, the European Commission, Australia, and Germany have all published reports concluding that digital platforms' market power has indeed become entrenched.⁶⁷ Surmounting the existing barriers to entry created by consumer behavior, cost structure, public policy, and any past anticompetitive conduct is extremely difficult. This fact has direct effects on consumers: *without entry or the credible threat of entry, digital platforms need not work hard to serve consumers because they do not risk losing their consumers to a rival.*

⁶⁵ Duhigg (2018).

⁶⁶ See, e.g., European Commission (2017). Indeed, as the U.K. Competition and Markets Authority has found, even in industries where prices exist but there are substantial search frictions, consumers do not make choices that would give them substantial savings. See note 74, *infra*, and associated text.

⁶⁷ See Furman (2019): 75; Crémer (2019): 112; Australian Competition (2018): 35; Schweitzer (2018): 2.

B. How Big Data Affects Digital Products and the Sale of Advertisements**1. What Makes Data Big? And Does Big Data Harm Consumers?***1) Targeted Advertising*

Technology firms claim that they perform machine learning on big data, and that doing so gives them both a competitive edge over rivals and allows them to better tailor their services to their consumers. How might this work in practice? Consider a search engine with information about a given user's search history, including the fact that this user recently searched for a specific pair of Nike running shoes. The ad service can more effectively spend advertisers' budgets by showing the user ads focused on running-related products (or similar shoes to those the user searched for). Advertisers will get a greater return on investment because more of their ads will be shown to users who have demonstrated interest in running products—possibly very similar running products to the ones being advertised. Moreover, users may actually prefer seeing ads for other running products compared to more generic ads. The more personalized matching of advertisement to potential customer would appear to help both the producer and consumer of the advertisement.

This simplified example of personalized advertisement uses very little information about a particular user (only that she had a particular search query) and little to nothing about other users or another augmenting dataset, with the exception of a database of running-related items. Indeed, this level of personalization could be similarly achieved in many low-tech spaces (e.g., running products advertised in running magazines, whose readers have expressed interest in running simply by reading the magazine). What is different about personalized advertisement if the ad server has much more data at its disposal?

Box I**Platform vs. Brick-and-Mortar Advertising and Targeting**

Traditional brick-and-mortar stores and online platforms differ greatly in their advertising and personalization capabilities.

At the highest level, local grocers tend not to force shoppers to identify themselves when they shop, rarely verify identification if used, and rarely have the ability to merge purchase history with other detailed information from other aspects of their customers' lives to design targeted advertising. Online retailers, on the other hand, almost always require account creation for purchasing, verify this information for each transaction, and have direct or easy access to detailed non-shopping information about their customers.

Local grocery stores are capable of some data collection and personalization. For example, they normally know that a majority of their customers live relatively close to the grocery store, so they can rely on their knowledge about the general demographics of the neighborhood population. The stores may also ask that consumers use a loyalty card to receive discounts, which then allows them to track per-customer purchasing patterns, offer particular customers with certain purchasing patterns particular coupons or free products, and see the result of such offers on a user's purchase history. Nonetheless, the effectiveness of these methods is limited in comparison to digital shops. Consumers normally retain the option to buy full-priced groceries without tracking (pay in cash), or to intermittently switch their cards with other people they know (for example). More importantly, even if grocers used fingerprints or other unique identifiers to remove the possibility of anonymity, the data they would track would still be entirely comprised of grocery purchase history for its customers.

This limitation does not apply to digital platforms. Online groceries require a customer to make an account prior to allowing any purchase. Thus, all purchases belong to a particular account, which also contains identifying information about a customer including their search history, billing, shipping, and email addresses. This ensures that (potentially) all search history and all purchases sent to a particular address (or paid for with a certain mailing address) can be collated and analyzed by the seller and used to design offers or other services for a customer. If this seller has a broader set of services than just grocery retail (for example, the seller is Google or Amazon), the information they can collate together with this purchase history might include email transcripts, calendar information, or search and purchase history of non-grocery products, amongst others.

All of this information can allow for much more aggressive marketing to a particular consumer, based on many aspects of their lives beyond their historical grocery shopping habits. A grocery store can take advantage of hungry shoppers by placing a representative selection of junk food in the checkout line. A digital platform, by contrast, can design an individualized tempting "checkout offer" based on that particular consumer's purchase history, current behavior, time of day, and the emotional content of recent communications.

2) Dimensions Along Which Data Can Be Big

In order to discuss the use-cases of large-scale datasets in online personalized advertising, it helps to define in what sense a dataset can be big. Big data commonly refers to two very different properties of a dataset: either that the dataset has many people's data in it, or that the dataset has a great deal of information about each person in the dataset. For simplicity we refer to the former as "large population datasets" and the latter as "high dimensional data." These distinct ways in which a dataset can be big enable very different uses.

Large population datasets allow the possessor to infer both unknown attributes of current users, and statistical facts about individuals not currently in the dataset. If the dataset contains many instances of users' queries along with their locations at the time they made those queries, an advertiser could use future users' locations to help predict what those users might search for or be interested in seeing advertisements about. Numerous statistical techniques can be used to show that, for a large population dataset, simple statistics that hold true on the dataset should also hold true for fresh users, assuming they come from a similar pool as those in the dataset. For example, if 30% of queries in the database originating from Cape Cod searched for lobster rolls, there is a 30% chance that a query made by a future user based in Cape Cod will be for lobster rolls, assuming the user is visiting in a similar time frame that the dataset was gathered in, arrived at the search engine in a manner similar to other dataset users, and the number of queries in the dataset originating in Cape Code is sufficiently large. If the dataset has many queries but fairly few from Cape Cod, statistical techniques will provide lower confidence in their ability to predict future queries emanating from Cape Cod.

High-dimensional datasets allow for different uses than large-population datasets do. Suppose a dataset contains only a few users, but each user's entry contains their entire email history. A quick read of a user's recent emails could give an advertiser a very clear window into what *that* user might be looking to purchase. The richness of a given user's data entry can describe in great detail many facts about her, including her future travel itineraries, plans for large purchases, information about her career and social networks, and so forth. Looking only at that user's information, a deft advertiser could likely select any number of products this user would be much more likely than an average person to find interesting. Hotels in Cape Cod would much prefer to advertise to a user if her recent emails describe travel plans to Cape Cod. However, if only three or four users belong to the dataset, an advertiser can learn very little about new users from studying the rich but poorly-populated dataset; again, statistical techniques have much lower confidence when the datasets they operate on are small.

The most useful datasets, from the perspective of an advertiser or other service provider, are large in both senses: they contain rich information about a huge number of people. This is even more true when the rich information contains different types of data, for example, email,

location, and search queries. A dataset of this type allows the provider to both learn high-level population statistics (for which it needs a large population in its dataset) and to carefully tailor its ads to each individual in its dataset (because it has very rich information about users in the dataset). Even more interestingly, the complexity of population-level statements one can make from such a dataset increases. This occurs both because each user's data has more dimensions, so there are more relevant hypotheses to explore, and because as the number of users in the dataset grows, so does the statistical significance of any particular statement that holds for the dataset. It is possible that we will see in the future a trend towards better algorithms allowing platforms to use less data and still target well.⁶⁸ However, it seems that data with enough volume and specificity will always be needed to develop these algorithms.

3) Accuracy (and Utility) May Display Increasing Returns

Entrants could have an opportunity if marginal returns were to fall as providers' datasets grow, because the incumbent's marginal cost of acquiring new data would eventually exceed the marginal value. But a simple model shows this may not be the case. Given a fixed dataset of a particular number n of individuals with a certain number k of features (also referred to as attributes), the dataset owner does not face decreasing marginal returns with respect to either n or k .⁶⁹

Why might datasets show increasing marginal returns in either the number of features or the number of rows already present in the dataset? For a formal treatment of this question, please see the discussion below which works out in detail a particular example in the advertising domain. Informally, imagine you are trying to sell a service to people who live in Manhattan and are planning a wedding. Knowing either that a person lives in Manhattan, or that they are planning a wedding, may make almost no difference in determining their willingness to use the service because both are very low-probability events, while having both pieces of information together allows the Manhattan wedding planner to easily pinpoint interested consumers. More generally, as a firm accumulates more information about more people, its marginal returns to new data need not fall.

This last point helps us understand why data has increasing marginal returns. Even if the dataset is large enough that it allows a company to make accurate inferences about a given population, the company will always benefit from having specific information about a given

⁶⁸ For example, during the 2019 ANTITRUST AND COMPETITION CONFERENCE - DIGITAL PLATFORMS, MARKETS, AND DEMOCRACY: A PATH FORWARD, Hal Varian, Google's chief economist, pointed out to the Imagenet Large Scale Visual Recognition Challenge 2017 as an example of this trend – see <http://image-net.org/challenges/LSVRC/2017/index> ; See also Varian's presentation at <https://www.youtube.com/watch?v=mlBWWHouH8>.

⁶⁹ This holds even in "natural" settings, where each member of the dataset is drawn identically from a normal distribution.

individual, allowing it to become more and more confident about what the consumer wants, and to better tailor its services and ads.

Companies therefore have no incentives to stop looking for and accumulating new pieces of data, entrenching incumbents with large datasets vis-à-vis entrants with smaller databases. Consumers on the internet leave numerous traces of their activities across a range of applications (for example, their location, what they buy, who they talk to, and what they say), and technology allows platforms to identify and analyze these traces. The amount of data on individual behavior that can be collected, merged, analyzed, and stored is rising, and the combination of different dimensions of data generates valuable information about individuals' tastes and behaviors. As individuals rely more and more on a platform to organize their lives through their online social, cultural, or economic activity, their data become more informative about their future choices and firms are willing to pay to influence those choices. Furthermore, the emergence of the Internet of Things means that platforms will have access to yet more data generated by home appliances, cars, and other devices. Indeed, consumers' devices can now track eye movement, mouse movement, body movement, and body position. In parallel with the evolution of the internet that made tracking of billions of individuals possible, advances in data mining and artificial intelligence have enabled firms to learn more from data than was conceivable a few decades ago.⁷⁰

Many digital markets have tipped and therefore there are only a few entrenched platforms able to gather this breadth of data. As Alessandro Acquisti and colleagues concluded in a recent study, "a few 'gatekeeper' firms [will be left] in a position to control the tracking and linking of . . . behaviors across platforms, online services, and sites."⁷¹ When data exhibit increasing returns to dimensionality and size, platforms that can track many users across those dimensions will have economies of scale and scope; they will be able to sell more valuable advertising.

Nonetheless, this targeting can also raise the quality of services provided by platforms. When they can identify individual tastes at fine levels and personalize their services to this taste, they often improve people's lives. Search engines can better answer queries or find a nearby destination, cultural and news websites are able to suggest well-suited content, and ecommerce websites can improve matching between buyers and sellers. These are all part of the consumer benefit described previously.

⁷⁰ Jaron Lanier, a creator of virtual reality, has warned of the problems associated with a "surveillance economy" in which users of digital tools and platforms would be enticed to give up personal data in exchange for "free" products and access. Lanier (2013). The data, Lanier argues, could then be monetized by the owners of the platforms and applications, largely through the sale of the data to advertisers and others finding value in the users' personal characteristics and proclivities.

⁷¹ Acquisti, Taylor, and Wagman (2016); Acquisti, Brandimarte, and Loewenstein (2015): 444.

4) *Types of Data*

The definition of what constitutes a consumer's data can become complicated. Suppose the consumer's data and those of similar consumers allowed the service to infer characteristics about the consumer (for example, that she has a tattoo) without any action or communication on the topic by the consumer. Is this piece of information part of the "consumer's" data, or is it the intellectual property of the algorithm owner?

The Vestager Report and the Furman Report categorize data as volunteered, observed, or inferred. Volunteered data is intentionally provided by the user to the service -- for example, when a user provides their favorite TV shows to a service in order to receive recommendations. Observed data, such as the history of shows that the consumer actually watches, are automatically gathered by a service. Some observed data may not be intentionally provided by users if they do not understand the privacy protections in the service -- for example, location tracking of a person using a video app. Finally, the service can process volunteered and observed data to *infer* additional information about the user or a group of users.⁷² We include (but do not always distinguish among) all these types of data in this report's discussions.

⁷² See Crémer (2019): 25. Famously, Target claims to have been able to predict a customer's pregnancy before she knew about it. Duhigg (2012).

Box II

Why Data Has Increasing Marginal Returns

For a particular dataset, suppose the first feature is the one the owner wants to predict for future users. For example, this feature might refer to whether or not a user will book a particular hotel if shown an advertisement for it. If $k=2$, and the second feature represents the home zip code of a user, there might be some limited ability to predict interest in a hotel based on this zip code being sufficiently far from the hotel. If $k=3$, consider the case where the third feature refers to the annual household income of each person.

With both zip code *and* income available, we can express much more complex prediction rules for interest in hotel X . For example, we could predict that a user will book the hotel if

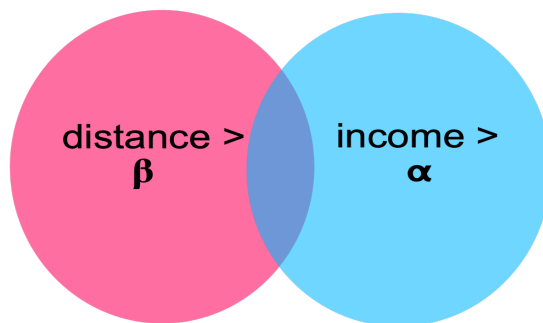
$$\alpha \cdot distance + \beta \cdot income \geq \gamma,$$

for some $\alpha, \beta, \gamma \geq 0$. If the *true* relationship between these three features is in fact linear, then the utility of gathering the income variable will be marginally decreasing in the set of other (linearly relevant) features.

However, it might be that the best prediction of whether a hotel will be booked is a nonlinear function of distance and income; perhaps a person will book exactly when they live far enough away *and* make at least a certain amount of money:

$$distance \geq \alpha \text{ and } income \geq \beta$$

where users with $distance > \alpha$ and $income > \beta$ make up 10% of the distribution. Suppose 40% of the users have $distance > \alpha$, and 40% of the users have $income > \beta$.



Why Data Has Increasing Marginal Returns (Cont.)

Knowing a user's zip code (and corresponding distance) gives some limited ability to predict when both the income and distance are sufficiently large, but knowing both the zip code and income will allow them to perfectly predict the 10% of the population that might have interest in staying at the hotel. For this reason, the utility of data owners need not be marginally decreasing in the set of features they have.

To be explicit, fix an advertising firm for the hotel. If the advertising firm purchases the ability to see each potential advertisee's zip code, then showing the advertisement to customers as a function of their distance is now possible. For example, the advertiser could only show the advertisement to customers whose distance is more than β , increasing the probability that the ad is shown to an interested party from 10% to 25%. If the firm then acquired the ability to estimate each advertisee's income, it could use both features together and only show the advertisement to customers with both income $> \alpha$ and distance $> \beta$.

This would raise the probability that a targeted advertisee would be interested in the ad from 25% to 100%. The two features are otherwise identical, and so it would also be the case that first acquiring income and then distance would raise the probability of an advertisement reaching an interested customer from 10% to 25% to 100% if one started with no features, then added income, then distance. Therefore, the utility increase which comes from adding distance first is $.15 * (\text{value of showing the ad to an interested customer})$, while the utility increase from adding distance after first acquiring income is $.75 * (\text{value of showing the ad to an interested customer})$. That is, the distance feature is more valuable *after* first learning about income (and vice versa). This shows that the value of features may not be marginally decreasing in the set of features accrued thus far.

One can also use the same sort of relationship to argue that the dataset owner's utility need not be marginally decreasing in n , the number of people in the dataset. Again, suppose one can only predict interest in the hotel with the conjunction of zip code and income information, but that this relationship is not a priori known to the dataset owner, and they are instead trying to *learn* how to predict hotel interest based on a dataset they have available to them. Using statistical techniques, it won't be possible to learn rules of this form (if x and y then z) to a high degree of accuracy unless the dataset has sufficiently large n . The utility of adding additional people to the dataset can *increase* as the dataset grows, depending on the owner's utility for accuracy of the learned model.

2. Data Control and Ownership

1) Status Quo

In the United States at present, a consumer has no property or control rights over their data. This stands in contrast to Europe, where the GDPR sets a standard for what a digital business can do with a consumer's data.

In the U.K., regulators have sought to give consumers *useful* control over their financial data. After a lengthy investigation showed that consumers rarely switch banks despite large gains to doing so, regulators sought to use Open Banking to give people the ability to seamlessly move between banks.⁷³ The Open Banking Initiative allows consumers to obtain, see, and transmit their banking activity in a standardized and secure fashion to regulated and approved third-party firms.⁷⁴ This is an example of data portability. Theory suggests that consumers will use that power to move their business to banks that lower prices and improve services. If Open Banking causes more competitive outcomes, it may provide a strong model for regulated portability and interoperability in other markets.⁷⁵ The Vestager Report defines protocol interoperability and data interoperability as stronger than data portability because they allow continuing communication between two services. It defines full protocol interoperability as a complete linking of the two services in a way that reduces network effects. We will return to these concepts below.

Data intermediaries collect consumers' information that they then sell to third parties. These intermediaries may be large websites obtaining the information through their service to consumers, or data brokers. Data have specific features that make this market unusual. For example, data are “non-rivalrous” meaning a broker can sell the same dataset to many buyers and still retain it, unlike, for example, the sale of apples. Once you've sold an apple, you're out an apple; once you've sold a given piece of information to one buyer, you can sell it again to another buyer. A key feature of data in this context is that data may be either directly shared with the buyer, or withheld (to prevent the buyer from achieving its own economies of scale and scope) but embedded into a service the buyer wants, such as targeting advertising.⁷⁶

2) Externalities

What can a person do with ownership over her data? The previous section describes two ways in which data might be used: First, it might be used to tailor services to a given person; second, it might be used to learn patterns that hold on average for the population from which the dataset was drawn. The former does not necessarily transfer any knowledge about one customer to the treatment of another, while the latter aims to learn about some fraction of a dataset and use that information to affect interactions with future customers. Note that this use of data creates externalities between consumers. Purchase or travel patterns by one person are used to create recommendations or suggestions for similar people, where “similar” is determined by machine learning. For this reason, the value of an individual person's data is more than the value of using

⁷³ Smith et al. (2016).

⁷⁴ U.K. Competition (2017); *see also* Open Banking Ltd. (2018); Manthorpe (2018). Third party firms that customers can elect to give data access are regulated and approved by the Financial Conduct Authority. U.K. Financial (2017).

⁷⁵ There are some signs that Open Banking has seen early success. Furman (2019): 70.

⁷⁶ Bergemann and Bonatti (2018b); Bergemann and Bonatti (2018a).

it to market to them. A customer whose data predicts the behavior of many other consumers is very valuable to a platform.

3) *Deletions*

One might simplistically think that “owning” one’s own data gives a person the ability or right to delete it. The ease with which a customer’s data can truly be deleted depends on the way the firm has created and is using that data. If a user’s own data is used only to personalize her experience, deleting her data will mean that the dataset—and the service’s subsequent behavior—has no trace of the customer ever belonging to the dataset. If, however, the service used the customer’s data to inform population-wide behavior (by analyzing the dataset including this customer and storing the results of those analyses), merely deleting the customer’s data from the dataset does not truly erase her presence—her data will still have a lasting effect on the future behavior of the service.

For example, suppose the service scans the dataset for one representative user from each zip code and remembers the purchasing history of each representative user. Even if the service removes the initial entry corresponding to a particular user in the original dataset, if the service still uses this set of representative users’ purchasing histories, some of the user’s data can still remain in the system. While this example may seem contrived, many machine learning methods do some amount of memorization of some subset of their training data.⁷⁷

Beyond memorization, there are other, more subtle ways in which a user’s data can affect the long-term behavior of a system and what information the system holds on to. The natural tool to restrict how much a system’s information is affected by one user’s data is to impose *differential privacy*. Loosely speaking, differential privacy restricts the statistics created from the dataset to not reveal if any particular person or observation is in the dataset.

4) *Transparency, Choice, and Fraud*

According to the Federal Trade Commission, the market for data suffers from a lack of transparency.⁷⁸ Most consumers have no idea how much information is being collected about them, sold, and used to make a profit.⁷⁹ One way in which digital platforms exploit their market power is by requiring consumers to agree to terms and conditions that are unclear, difficult to understand, and constantly changing. The terminology in these contracts is legal and the documents are often lengthy; the consequence of the different clauses is difficult to understand and foresee. Moreover, the user of a device or platform makes a choice to sink investments

⁷⁷ For example, support vector machines are explicitly recorded as a small number of datapoints on which the model was learned.

⁷⁸ Federal Trade Commission (2014).

⁷⁹ Federal Trade Commission (2014): 42.

(posts, calendars, media, and so on) in a particular platform at a moment in time when a particular user agreement is in force. After the device or platform updates its terms of service, a user may be locked in. Having bought a phone, they won't immediately want to switch; having built a network of friends, they often won't want to leave. Thus, the user does not have the same set of choices as she did the first time she hit "agree," and some agencies have argued this is no longer a free choice.⁸⁰ Lastly, a digital platform may describe its data-use policy in its terms of service and then deviate from that, rendering the initial statement fraudulent.⁸¹

Box III Digital Identities

Perhaps the next major shift in digital competition will be the quest to control the identification market. Once we create an account with any digital platform, we create a digital identity which incorporates select data on age, sex, address, email address, preferences, and, frequently, much more. These digital identities help companies identify and tag users to the data they generate, be it transactional, social, simple web navigation or even meta data. Digital identities work like access and tracking mechanisms, allowing a user to surpass a wall in exchange for enabling the company to link the data to a given dataset (e.g., using a Facebook login to access a website).

Identity data and control is highly valuable and platforms can monetize it in many ways. They can use it to personalize services and charge subscription fees, provide advertising or market intelligence, or as a way to increase bottleneck power and charge companies for user access. The more privacy protection technologies grow and limit widespread data collection online, the more being the single identification point for users will grow in importance and value.

3. Digital Platforms are Characterized by Free Services

"Free" is not a special zone where economics or antitrust do not apply.⁸² A free good is one where the seller has chosen to set a monetary price of zero and may set other, non-monetary, conditions or duties. Zero is a number, just like 10 is a number. If a competitive price is \$10 while the realized price is \$15, there is a \$5 markup above the competitive price. This is the same harm the consumer bears when the competitive price is -\$5 (at this price the digital platform is actually paying the consumer for her data and information in addition to providing her with services) and the realized price is \$0.

⁸⁰ Bundeskartellamt (2019) (determining that Facebook's take-it-or-leave-it under agreements constitute "abuse of market power").

⁸¹ See, e.g., Romm and Dwoskin (2019).

⁸² See, for example, the current head of the U.S. Department of Justice's Antitrust Division describing the challenge faced by his division of defining markets when goods are free. Delrahim (2019) (arguing, without regard for quality effects, that traditional antitrust market definition cannot work because "[w]e cannot look at the effects of a five percent increase in price because five percent of zero is still zero"). See also *id.* (arguing that "[c]hoosing variables for measuring market shares also can be more complicated where shares of revenue is not an option").

Barter is a common way in which consumers pay for digital services. They barter their privacy and information about what restaurants they would like to eat in and what goods they would like to buy in exchange for digital services. The platform then sells targeted advertising, which is made valuable by the bartered information. But, in principle, that information has a market price. It is not easy to see if the value of any one consumer's information is exactly equal to the value of the services she receives from the platform. However, many digital platforms are enormously profitable, and have been for many years, which suggests that in aggregate we do know the answer: the information is more valuable than the cost of the services. The economics literature has modeled this setting and is able to define a data markup.⁸³

The current inability to use both positive and negative prices for digital goods means that the policy discussion cannot focus on dollars alone as the unit of cost. Rather, digital platforms should be analyzed using both price and quality. "Quality-adjusted price" is a metric often used by economists in this situation. If a platform's price is fixed at zero and the quality of the service improves, then its quality-adjusted price has fallen. Conversely, if a platform's price remains zero but its quality falls, its quality-adjusted price has risen. When the price is fixed at zero, it is possible to track quality-adjusted price over time: the movement in quality accurately reflects quality-adjusted price.⁸⁴

Online platforms offer many services for zero monetary price while they try to raise participation in order to generate advertising revenue. Free services are prevalent on the internet in part because internet firms can harness multi-sided network externalities. While the low price can be a blessing for consumers, it has drawbacks for competition and market structure in a world where institutions have not arisen to manage negative prices. Because there is currently no convenient way to pay consumers with money, platforms are able to mark up the competitive price all the way to zero. This constraint can effectively eliminate price competition, shifting the competitive process to quality and the ability of each competitor to generate network externalities. Depending on the context this may favor or impede entry of new products. For example, entry will be encouraged when a price of zero leads to supra-competitive profits, and impeded when a zero price prevents entrants from building a customer base through low price. Moreover, unlike traditional markets where several quality layers may coexist at different price levels (provided that some consumers favor lower quality at low price), markets where goods are free will be dominated by the best quality firm and others may compete only in so far as they can differentiate their offers and target different customers. This strengthens the firm's incentive to

⁸³ Bergemann, Bonatti, and Smolin (2018).

⁸⁴ The European Commission has noted the importance of assessing product quality in zero-price settings. Esayas (2018).

increase quality through increasing fixed costs in order to attract customers (known as the Sutton sunk cost effect) and further pushes the market toward a concentrated market structure.⁸⁵

It is a puzzle that, to date, no entrepreneur or business has found a way to pay consumers for their data in money. For example, a consumer's wireless carrier could aggregate micropayments across all manner of digital destinations and apply the credit to her bill each month. It may be that adverse selection, transaction costs, and coordination difficulties create too large a barrier for today's entrepreneurs, though technical solutions like a verifiable digital identity, mentioned above, would combat adverse selection.⁸⁶ Furthermore, a carrier that could bargain effectively with platforms on behalf of its subscribers for high payments would likely gain subscribers. Notice that an easy method to pay consumers, combined with price competition for those consumers, might significantly erode the high profits of many incumbent platforms. Platforms likely have no economic incentive to work diligently to operationalize negative prices.

⁸⁵ Sutton (1991).

⁸⁶ Posner and Weyl (2018).

Problems Arising in Digital Markets

The changing market dynamics, outlined above, raise doubts about the market's ability to ensure ongoing innovation and competition. Increased concentration levels, market power, network effects, and control over data and analytics have in many digital markets tipped the market in favor of the incumbents.

The theme we return to throughout this report is the difficulty of entry into digital platform businesses once an incumbent is established. Whether the entrant is vertical or horizontal, has succeeded to some degree, is nascent, or is a potential entrant, its existence improves consumer welfare. Either the entrant provides more choice, different features, and a chance of higher quality, or the threat of those outcomes spurs the incumbent to provide lower prices, higher quality and innovation, and to do so more quickly.

By focusing on this feature of digital platforms, we are highlighting the *cause* of the market power, not its result. Absent entry barriers of the type discussed above, the tremendous amount of profit available in these markets would stimulate entry. Protecting competition in these markets requires protecting competitors. Entry and potential entry create more competitors (in expectation), and that increase depends on competition working effectively so that a meritorious entrant can successfully dethrone the incumbent.

The categories of economic harms to consumer welfare from digital platforms are the standard ones: price, quality, and innovation. The report will primarily emphasize quality and innovation harms due to their greater complexity and generality. For example, by excluding competitors, dominant firms do not need to innovate as hard as they otherwise would be required to keep their customers. Likewise, when platforms do not face competition, they will be able to reduce quality, for example, by decreasing privacy protections, without losing customers or revenue.⁸⁷ When a service reduces quality without lowering price, it is *raising* quality-adjusted prices, which harms consumers. For example, if a phone service were to lower the quality of service but keep monthly fees the same, it would have raised quality-adjusted price. Because many digital services are purchased with barter, the monetary price paid by consumers is zero, and quality-adjusted prices cannot be directly seen the way a nominal price can be. On the advertiser side of the platform, where monetary prices *are* charged, harms to competition exist in the more ordinary form of higher markups for ads and other services. These markups are eventually paid by consumers because they are built in to the prices of the goods and services that are advertised online.

⁸⁷ Reyna (2018).

A. *Quality Harms*

1. Personal Data and Behavioral Economics

1) *Behavioral Economics*

As discussed above, behavioral economics helps improve our understanding of real consumer choices and suggests that consumer exploitation is common. There are a number of systematic consumer biases that, when incorporated into economic analysis, affect outcomes and welfare. For instance, individuals can be subject to salience effects, putting excessive weight on the most salient information. Confirmation bias can lead them to change their preferences to conform with past choices.⁸⁸ Consumers are often biased toward the status quo even when it is no longer optimal.⁸⁹ Perhaps the most important consumer biases are impatience and lack of self-control. The former refers to the discounting of any payoff that occurs further in the future than the present. The latter is closely related and refers to the extent to which individuals fail to resist short-term impulses in order to achieve long-term goals. The literature in behavioral antitrust argues that status quo, salience, and impatience are the most relevant for antitrust analysis.⁹⁰ Platforms that analyze their consumers' behavior can exploit these biases by framing choices to make certain information salient, designing a status quo that is profitable, inducing addictive behaviors, generating sales through impulsive consumption, and exploiting consumers' disinclination to search. These strategies are common in the brick and mortar world. For example, the candy aisle in supermarket check-out lines, or the rug store that has a special 50%-off sale every day.

2) *Using Machine Learning to Take Advantage of Consumers*

There are now decades of economic research demonstrating consumer bias and firm responses in offline markets.⁹¹ For example, gyms offer subscription memberships rather than pay-per-visit knowing that members will not come as often as they anticipate,⁹² and credit cards offer teaser rates knowing that consumers plan to have no debt in 6 months' time.⁹³ However, the strategies firms have used to date are swamped by what digital businesses can learn by using

⁸⁸ "Confirmation Bias" (2019) ("Confirmation bias is the tendency to search for, interpret, favor, and recall information in a way that confirms one's preexisting beliefs or hypotheses."); *see, e.g.*, Thaler (2018): 1266 ("People guess that in the United States today gun deaths by homicide are more frequent than gun deaths by suicide, although the latter are about twice as common. The bias comes because homicides are more publicized than suicides, and thus more 'available' in memory.").

⁸⁹ Samuelson and Zeckhauser (1988) ("[D]ecision makers exhibit a significant status quo *bias*. Subjects in our experiments adhered to status quo choices more frequently than would be predicted by the canonical model.").

⁹⁰ Fletcher (2019).

⁹¹ *See, e.g.*, Kahneman and Gilovich (2002); Barberis and Thaler (2003); *see also, e.g.*, Kahneman, Knetsch, and Thaler (1991).

⁹² DellaVigna and Malmendier (2006).

⁹³ Bar-Gill and Bubb (2012).

high-dimensional, large datasets to explore every nook and cranny of consumers' many behavioral shortcomings and biases in real time. This strategy is of serious concern when combined with an understanding of the digital platform business model discussed above. The platform's goal is to use its knowledge of consumers, combined with its market power (and the resulting lack of consumer choice), to extract profit. Some of that profit may come from value creation, or "expanding the pie," but some will come from transfers to the platform from both the consumer side and the advertiser side. Because individuals are subject to behavioral biases, consumers are vulnerable to a platform's exploitative behavior. Additionally, when individuals are workers, as they often are in the "gig" economy, these tools can be used to advantage the platform against the worker.

Digital businesses not only have more information than traditional firms, but they have more variations of products or services and the ability to control the environment and the timing of choices and offers. For example, a firm can hide a component of a good's price to let consumers discover prices only once they have invested significant time and effort into buying the good. This strategy can be differentially employed depending on the consumer's past willingness to pay. Framing, nudges, and defaults can direct a consumer to the choice that is most profitable for the platform. A platform can analyze a user's data in real time to determine when she is in an emotional "hot state" and offer a good that the user would not purchase when her self-control was higher. Consider a supercomputer tracking a consumer, via her cell phone, around the town until she is tired and frustrated in some way, and at that moment presenting her ads and information about junk food. This type of exploitation could depend on input from devices such as eye-tracking sensors, the ability of AI to understand the emotion expressed in texts and email, and all the other data the platform has about the consumer combined at a very large scale. This same tactic can be used to gain advantage against an independent contractor, e.g. a driver, whose behavior and location can be tracked for long periods.⁹⁴ In addition, machine learning applied to big data may help differentiate well-informed and sophisticated consumers or workers from poorly informed or more naïve consumers, raising the possibility of further exploitation of those least prepared to resist it.

Internet firms make frequent use of digital defaults, framings, and nudges. When a user is signing up for a new service, the sign-up form may automatically check the box that permits the service to send the user emails.⁹⁵ A user can opt-out of receiving emails by unchecking the box, but doing so is harder than sticking with the default. The results of a search that fits onto one page may all be sponsored, whereas finding the organic links requires paging down. *Homo*

⁹⁴ <https://www.nytimes.com/interactive/2017/04/02/technology/uber-drivers-psychological-tricks.html>

⁹⁵ This is no longer permitted in the European Union. See Regulation (EU) 2016/679 of the European Parliament and the Council of 27 April 2016 on the Protection of Natural Persons with Regard to the Processing of Personal Data and on the Free Movement of Such Data, and Repealing Directive 95/46/EC (General Data Protection Regulation), 2016 O.J. (L 119/1), § 32 ("Silence, pre-ticked boxes or inactivity should not . . . constitute consent.").

economicus is hardly influenced by defaults—to a rational agent, scrolling down or unchecking a box is trivial—but real people are influenced. Nudges are not unique to digital products; for example, an employer offering a default health care plan to its employees nudges employees towards choosing that plan—employees can select a different plan, but doing so is harder than sticking with the default.⁹⁶ What is noteworthy, however, is the platform’s detailed, personalized, minute-by-minute control over their interface. This control enables platforms to create a façade of competition, choice, and autonomy when in fact users are being directed with behavioral techniques.⁹⁷

With big data and machine learning, firms are able to understand and manipulate individual preferences at a scale that goes far beyond what is possible in traditional markets. This capability is qualitatively new. The environment is characterized by extreme asymmetries of information and analytical capacity between the platform and the user. This enables firms to charge higher prices (for goods purchased and for advertising) and engage in behavioral discrimination, allowing platforms to extract more value from users where they are weak.⁹⁸ The problem is only growing; platforms continue to make investments to extract data, encourage stickiness and addiction, and promote ever-greater use, in order to run data analytics and enable more precise targeting.⁹⁹

The economic literature suggests that competition by itself cannot resolve the issue raised by the exploitation of behavioral biases or poor consumer information. This is because staying profitable in a competitive environment may force firms to exploit behavioral bias to achieve maximal profitability. Firms abstaining from doing so may be driven out of the market. Rather, competition causes a shift of surplus to wiser consumers; profit from exploitation of biased consumers is used to compete for well-informed consumers. For consumers who can guard against exploitation, there are therefore significant gains to be had from competition. But this will not be true for all people; some will be taken advantage of under perfect competition if consumer protection regulations are insufficient. Hence, while this report focuses on the competition problems created by powerful platforms, related issues raise broader consumer protection concerns that cannot be solved through greater competition.

⁹⁶ Thaler (2018): 1283. Default settings have very large effects, for example, in voluntary organ donation decisions. Countries with an op-in default generally have the vast majority of the nation’s adults enrolled to donate, while countries with default opt-out decisions see the opposite result. *See* Johnson and Goldstein (2003).

⁹⁷ Ezrachi and Stucke (2016) (comparing the façade of choice and autonomy on digital platforms to the Truman Show).

⁹⁸ For a discussion of the use of online price and behavioral discrimination, see Stucke and Ezrachi (2018b); Stucke and Ezrachi (2018a).

⁹⁹ Edwards (2018).

2. The Harms from an Advertising-Supported Business Model

1) Market Power Leads to Markups

Access to high-quality data, scale, and scope has allowed a few large platforms to offer efficient targeting services for advertising and to dominate the advertising market. Facebook, Google, and, increasingly, Amazon act as gatekeepers to the online advertising market. The three platforms provide access to billions of users, as well as a data-rich environment, essential for modern online advertising. Being in control of the data and the assets—users—creates market power. That market power generates a profit margin which, for ad-supported platforms, comes from the sale of advertising. One of the characteristics of the digital advertising environment is its opacity: major platforms are able to leave bidders and publishers in the dark with respect to the true success, costs, and profits from placement of advertising. This can be exacerbated when the platform also supplies buyer or publisher tools and analytics. Opacity is partially a consequence of market power. Market power can be used to discourage, or even prevent, multi-homing by buyers. One additional concern is that this opacity may give firms not only the normal ability to exercise market power in intermediation but also to engage in fraud.¹⁰⁰ Lack of transparency also undermines buyers' ability to measure the effectiveness of digital advertising and therefore to understand its true value relative to price.¹⁰¹

A report by the French Competition Authority estimated that publishers received 40% of advertiser sales, with intermediaries collecting the rest.¹⁰² What justifies such a large markup for intermediaries? Without detailed study we cannot know for sure. However, reasons likely include the fact that very few platforms can target customers (who may be single-homing) using detailed and accurate data, so advertisers cannot take advantage of competition. The platforms do not sell the data to advertisers, but promise to place the ad in front of the requested demographic. The advertiser lacks transparency or any ability to learn about its customers and potential customers. Moreover, a platform that operates in the advertising placement business holds the keys to a second black box through its control of the pricing process.¹⁰³ The integration of the business of running the price-discovery mechanisms as well as tools for each side of the platform enables opacity, which helps maintain market power. When a platform bundles services such as advertising placement and return-on-investment analysis with advertiser data such as completed purchases, it can further enhance its market power.¹⁰⁴ The platform generally never shares the

¹⁰⁰ See, e.g., Vranica and Marshall (2016); Shields (2016); Farivar (2018) (reporting allegations that Facebook's "average viewership metrics were not inflated by only 60-80 percent; they were inflated by some 150-900 percent").

¹⁰¹ This in turn obscures the size of the platforms' markup. See Select Committee on Communications (2018); Pidgeon (2016).

¹⁰² French Competition (2018): 40, § 82.

¹⁰³ For example, Google's ownership of the ad platforms formerly known as DoubleClick and AdWords. See Spangler (2018).

¹⁰⁴ See generally Elhauge (2009).

data with advertisers or publishers, but keeps it to itself, preventing disintermediation by a brand or publisher, further sustaining its large profit margin.

2) *Incentives Created by Markups*

Ad-supported platforms' high markups provide a powerful reason to try and keep users online for another minute in order to show more ads. These profits push platforms to design their firms around "engagement"—an obsession with keeping users on their system for as much time, and with as much attention, as possible.¹⁰⁵ As much of the behavioral literature cited above shows, advertising-supported digital businesses can use consumer biases to hold people's attention in ways that ultimately harm them. Early empirical work has found such effects in social media. In a recent working paper, Hunt Allcott and colleagues found that Facebook users who were paid to leave the site for four weeks wound up with higher subjective well-being than similarly situated people randomly assigned to a group not offered the payment.¹⁰⁶ Users who took the break from Facebook had a "large and persistent reduction in Facebook use after the experiment," along with reduced political polarization and news knowledge.¹⁰⁷ For the platform, engagement serves two reinforcing purposes. First, the more time a user spends on a platform, the more the platform knows about her. Second, the longer a user is on the platform, the greater its income from ads and services. Thus, the more time a user is on the platform, the more ads the platform can sell, and the more it can charge per ad.

This financial incentive may explain the growing use of tracking and the creation of ecosystems that are based on the ability to manipulate the user into staying longer on the platform.¹⁰⁸ This opens the door to exploitative tactics or content as a method of increasing engagement. Furthermore, a second element to the large margin earned by a platform may be the low cost of the purchased content that keeps users on the platform. An important question (addressed more fully by the Media committee) is whether the prices for that content are set competitively. In a bargaining environment, the market power of the platform may allow it to negotiate a price for content that is below competitive levels, potentially creating dynamic harm to input providers.¹⁰⁹

The financial incentive created by a large markup can lead to anticompetitive behaviors as well as exploitative ones. Platforms may seek to reduce interoperability and awareness of outside options. For example, platforms may exclude certain services or increase friction in accessing third parties' services. High search and switching costs are used to "lock in" users and

¹⁰⁵ See, e.g., Bergen (2019) ("The company spent years chasing one business goal above others: 'Engagement,' a measure of the views, time spent and interactions with online videos.").

¹⁰⁶ Allcott et al. (2019).

¹⁰⁷ *Id.* at 1.

¹⁰⁸ Stucke and Ezrachi (2016).

¹⁰⁹ Stucke and Ezrachi (2017).

reduce the ability of competitors to access those users. Platforms may adopt strategies to reduce multi-homing to obtain more market power over their users. We will return to this theme below.

3) Resulting Quality of Content

As discussed in detail below, the cheapest way to keep users on a platform is to present content as being more attractive than users thought at the moment when they were ready to leave the platform, or to make departure more costly in other ways, so that users stay longer. (See a similar discussion in the Privacy Report.) The platform is essentially degrading the quality of the content offered in a way that present-biased human beings find engaging. As discussed further below, content that instantly engages most effectively is content that generates outrage, not necessarily content that is truthful or thoughtful. Simple strategies such as more advertising minutes per minute of content also lower quality. Low-quality content represents an increase in the quality-adjusted price of platform services experienced by consumers.

This business model, and the opacity that accompanies it, makes it difficult for either advertisers or consumers to realize they are being charged a markup and makes entry into these advertising markets extremely difficult. Because the advertiser side is where a platform earns its revenue, this is obviously a key entry barrier. The question of how an enforcer or regulator might enable entrants to overcome this barrier to entry is an important part of the discussion in Section III.B of the report.

4) Welfare and Efficiency

When analyzing the impact of online advertising and the behaviors it creates, one should keep in mind that antitrust aims to promote effective competition for the benefit of consumers. Competition authorities tend to treat advertisers as any other customers of a service. The premise is that market power in advertising channels impedes the ability of brands to reach consumers and inform them about their products. Such a treatment of the welfare from advertising relies on the implicit view that i) the role of advertising is solely to inform consumers; and ii) competition in the advertising market leads to efficiency. The economic literature tells us that both assumptions are questionable, as does observation of digital markets.¹¹⁰ In competition policy, when assessing the impact of a decision or a regulation involving advertising, it is important to focus on the final impact on consumers and market efficiency rather than on the advertisers' surplus.

On the one hand, targeted advertising to wise and well-informed consumers is welfare-improving insofar as it allows advertisers to send the right information to the right people, improving their choices and fostering competition among suppliers. On the other hand, in the

¹¹⁰ See, e.g., Anderson, Waldfoegel, and Stromberg (2016). Goldfarb and Tucker (2019).

modern economy this simple model becomes more complex because of the cost to the consumer, namely loss of privacy. There is an open empirical question as to whether the tradeoff is worthwhile to consumers.¹¹¹ We discuss the problem of measuring welfare in a setting where consumers are manipulated below.

3. Online Exploitation and Addiction

1) Human Reward Systems

Digital platforms' manipulation of their users—in part designed to get users addicted—is getting a considerable amount of attention from the public. These actions lower the quality of a platform and harm users.

As discussed above, some platforms have deliberately incorporated features that feed human “reward” centers into their products to induce users to give more and more of their time—and data—to the platform. These tools are designed for scale—they become even more valuable the more traffic they carry and the more users they garner—and hence the competition among producers has been described as a competition for eyeballs.¹¹² Because the digital tools and networks have been designed for use at scale, and because there is relatively little cost and considerable benefit associated with adding more users, producers want as much user engagement as they can get. The creators of digital products have benefited from social science and neuroscience findings that concern, for example, how certain colors or mechanics can feed a user's dopamine, much as nicotine does.¹¹³ The Center for Humane Technology, which calls the problem the “hijacking of our society,” describes the issues as follows:

Facebook, Twitter, Instagram, and Google have produced amazing products that have benefited the world enormously. But these companies are also caught in a *zero-sum race for our finite attention*, which they need to make money. Constantly forced to outperform their competitors, they must use increasingly persuasive techniques to keep us glued. They point AI-driven news feeds, content, and notifications at our minds, continually learning how to hook us more deeply—from our own behavior.¹¹⁴

This business model is based on acquiring a large volume of data to generate income, and it has led to unprecedented investment in addiction—ensuring continuing use of the interface. From diaper apps,¹¹⁵ to rewarded ads (in which viewers are given in-app rewards, such as another “life” in a video game, in return for watching an ad)¹¹⁶ internet firms harness their

¹¹¹ See, e.g., Aguirre et al. (2015); Johnson (2013).

¹¹² See, e.g., Bloomfield (2014).

¹¹³ Edwards (2018).

¹¹⁴ Center for Humane Technology (2018).

¹¹⁵ Edwards (2018).

¹¹⁶ Google AdMob (2018).

knowledge of second-to-second individual responses alongside social science learnings about human biases for maximum attention.¹¹⁷ Researchers do not yet know the full extent of the harm (or potential harm) that may come from this sort of addiction and manipulation, but they are now actively engaged in relevant research.¹¹⁸ What is known is that a business that depends on users staying online to watch ads and have their preferences harvested will focus its resources on keeping users online—for example, with intelligent and flexible algorithms.¹¹⁹ The algorithm will learn from, and respond to, these basic human preferences, thereby delivering—perhaps along with good content—a large quantity of low-quality content.¹²⁰

2) Examples

There are many examples of exploitation of consumers who have sunk costs in a platform and may not be fully informed. For example, app designers have enticed children into playing free games that are built around in-app purchases, leading children to make large purchases without parental knowledge or permission. The U.K.’s consumer protection agency, the Office of Fair Trading, found that such purchases could be pricey: “A My Little Pony game, for example, offer[ed] users a virtual ‘mountain of gems’ for a real-life £69.99.”¹²¹ Regulators around the world have cracked down to force companies to refund large payments that children make on their parents’ phones.¹²² More disturbing examples of low-quality content are YouTube recommended videos that lead the viewer towards false or dangerous content.¹²³ Prior to having these patterns made public and criticized, a Google search about the earth’s geology would lead to a chain of recommendations that resulted in “flat earth” content;¹²⁴ YouTube would offer

¹¹⁷ The video app TikTok may have taken this approach to its logical conclusion. See Herrman John (2019) (“TikTok assertively answers anyone’s *what should I watch* with a flood . . . [the app] has stepped over the midpoint between the familiar self-directed feed and an experience based first on algorithmic observation and inference. . . . It’s an algorithmic feed based on videos you’ve interacted with, or even just watched. It never runs out of material. It is not, unless you train it to be, full of people you know, or things you’ve explicitly told it you want to see. It’s full of things that you seem to have demonstrated you want to watch, no matter what you actually say you want to watch.”).

¹¹⁸ See, e.g., Kuss (2018); De-Sola Gutiérrez, de Fonseca, and Rubio (2016) (noting “a consensus about the existence of cell-phone addiction, but the delimitation and criteria used by various researchers vary”). The field is still nascent and there is no consensus that internet addiction exists or is a distinct psychological condition. See Ryding and Kaye (2018) (criticizing the cavalier use of the term “internet addiction” without a firmer grasp of the problem); Zajac et al. (2017).

¹¹⁹ Beginning in 2014, Twitter followed Facebook in centering around algorithm-selected content rather than simply displaying tweets from accounts a user chooses to follow in reverse chronological order. See Oremus (2017).

¹²⁰ Lewis (2018); Lewis and McCormick (2018); see also “Aim of the Project” (n.d.).

¹²¹ Osborne (2013).

¹²² U.K. Competition (2015); Federal Trade Commission (2015); European Commission. 2014.

¹²³ Bergen (2019).

¹²⁴ Newton (2019); Google (2019) (noting that the company was planning to “begin reducing recommendations of borderline content and content that could misinform users in harmful ways—such as videos promoting a phony miracle cure for a serious illness, claiming the earth is flat, or making blatantly false claims about historic events like 9/11”); see also Roose (2019).

teenage girls interested in diets videos about how to get anorexia, and so forth.¹²⁵ It is important to realize that this content is not chosen by human curators at the platform. Rather, the algorithm learns what content people will click on, and what content will cause them to stay on the platform longer, through many millions of small experiments; that is the content that is suggested and viewed. Exploitation and addiction caused by the optimization of the platform is a harm to consumers because they are likely watching lower-quality content than they would choose if they were fully informed in advance about how the content is chosen, or perhaps if they had alternative platforms to choose among.

Privacy

Another worry is that the privacy of consumers on digital platforms is violated; this too is a decline in the quality of the product. For example, Facebook recently announced that it will merge the infrastructures of Facebook Messenger, WhatsApp, and Instagram. As the New York Times noted, “[t]he integration plan raises privacy questions because of how users’ data may be shared between services. WhatsApp historically required only a phone number when new users signed up. By contrast, Facebook and Facebook Messenger ask users to provide their true identities. Matching Facebook and Instagram users to their WhatsApp handles could harm those who prefer to keep their use of each app separate.”¹²⁶ Germany’s competition regulator responded to this announcement by prohibiting Facebook from combining data from different sources (such as WhatsApp or Instagram) with data from Facebook.com without a user’s explicit and voluntary consent.¹²⁷ Facebook’s eagerness to get third-party apps connected to its network has led to mass data leaks, exposing sensitive information from hundreds of millions of people.¹²⁸ And Facebook is hardly alone.¹²⁹

B. The Assessment Problem

1. Measuring Consumer Welfare in a Behavioral World

Each of these issues results in fundamental difficulties in applying standard antitrust analysis— which is related to the assessment of the welfare effects of various practices or of a merger—to digital markets. Adequate measures of volume, quality, and consumer surplus may be difficult to obtain. The number of users choosing a certain option may not reflect their true preferences if the platform can make that choice a default that is difficult to see or to change. The

¹²⁵ Tompson (2018) (reporting Tristan Harris, head of the Center for Humane Technology, arguing that “[t]he problem is [YouTube] doesn’t actually care about what you want, it just cares about what will keep you next on the screen. The thing that works best at keeping a teenage girl watching a dieting video on YouTube the longest is to say here’s an anorexia video”).

¹²⁶ Isaac (2019).

¹²⁷ Bundeskartellamt (2019).

¹²⁸ “Data on 540 Million Facebook Users Exposed” (2019).

¹²⁹ See, e.g., MacMillan and McMillan (2018); Gartenberg (2018).

number of clicks on ads may not correlate with greater welfare if higher volume of clicks is obtained by exploiting lack of self-control and addictive behaviors.

Given the prevalence of behavioral effects in the digital economy, the measurement of consumer welfare must be carried out very carefully. As we have mentioned, behavioral economics is now a well-established discipline that can help sort different online behaviors and business practices. Incorporating this knowledge into the legal practice's toolbox may help develop better measures of output and quality.¹³⁰ We caution, however, that the legal structure of US antitrust law is not well set up to accommodate this complexity as it opens the door for judges to weigh all manner of social concerns as well as traditional economic effects. We see two approaches that might be more fruitful. First, the fastest route to more accurate measures of welfare might well be for a digital regulator to limit the business models that serve harmful content, the way regulators limit the harm from mortgages by restricting debt to income ratios, the harm from prescription drugs by requiring access through a physician, and the harm from automobiles by requiring airbags and crash tests, to name a few examples. If platforms had little or no incentive to deliver harmful content to consumers, the standard consumer welfare toolkit would be more accurate in this market. Second, harmful content is also, from the viewpoint of the consumer's long run or ex ante self, low-quality content. Analytical paradigms in antitrust analysis commonly deal with low quality and quality-adjusted prices, and could incorporate the role of exploitative content in this way.

Another reason to be pessimistic about measuring traditional surplus concepts is related to the barter nature of the exchange: Users barter attention and personal data for services. With a "free" service, consumers are paying for any expansion of activity with their attention to content. When facing a zero-money price, and when quality is difficult to observe, consumers are not receiving salient signals about the social value of their consumption because the price they believe they face does not reflect the economics of the transaction, and they are ignorant of those numbers.

2. Assessing the Social Welfare of Advertising

Behavioral economics also calls for careful use of standard economic surplus measures when applied to the supply of advertising. Unlike supply of goods, higher value of advertising need not imply higher social value. As discussed above, when advertising efficacy relies on psychological nudges that bring people into consumption they would consciously avoid otherwise, more advertising may well mean less consumer welfare.

¹³⁰ For a view of the role behavioral economics can play in legal practice, see Zamir and Teichman (2014).

Moreover, while price competition directly benefits consumers, competition through persuasive advertising may waste resources.¹³¹ To see this point, consider the following example: Two firms share the market equally without advertising. If one then uses advertising and shifts the demand toward that firm (perhaps by some psychological effect), it will then have a larger share of the market. However, if both firms decide to advertise, they will continue to share the market equally. If the advertising cost is not too large, both firms will advertise in order to preserve their market share. Therefore, despite advertising spending, the equilibrium sales and consumption are the same as without advertising. Here, allowing for advertising only induces wasteful spending, which in our setting is captured by the platform as profit. This is an extreme example, but it highlights that a higher volume of advertising—and greater platform profits—may not indicate an increase in social welfare in the same way we normally think a higher output of, for example, shoes, would indicate.

C. Harms to Investment and Innovation

1. Rents

Successful platform-style strategies pre-date the internet. For example, the key to Microsoft's business was the orchestration of ecosystem-wide innovation—the personal computer industry—to benefit its own core offering as well as complements provided by third-party partners.¹³² The strategy we highlight in this section of the report is a successful platform's choice of how much rent to expropriate from these complementors. For example, Microsoft turned PC hardware—an essential complement to an operating system—into a commodity business with the main exception of the microprocessor made by Intel. However, PC applications software—again an essential complement to the OS—remained (in part) an area where other firms could enter, compete, and earn profit.

The level of market power attached to a successful platform is so high that it often gives the platform owner the ability to expropriate almost the entire surplus available on its platform. And the ability to add that surplus to its existing core profit is the incentive to do so. For example, Microsoft's actions to favor Word over WordPerfect expropriated surplus from an existing software complement to its operating system. Importantly, in this example, WordPerfect was not a potential entrant into, or substitute for, a PC operating system, so there was not an exclusionary theme. The distinction between complementors that could disintermediate the platform and are therefore potential rivals, and those that cannot is critical for the antitrust

¹³¹ Although subtler, the same point has been shown to hold in the economic literature about informative advertising that conveys useful information about products such as existence, location, and other product characteristics. *See* Bagwell (2007); Anderson and Renault (2006). In this case advertising brings real value to consumers. But competition to steer consumers from other suppliers and/or to preserve market share may still lead to excessive spending on advertising.

¹³² Gawer and Cusumano (2002) (detailing the characteristics of these types of platforms).

analysis below. When the platform owner takes steps to disadvantage complementors who have no market power, the complementors may resist by using non-market (political and legal) strategies. Such complainants were part of the antitrust cases against Microsoft. We see this response to current platforms by firms in Europe today.

Chamath Palihapitiya, a venture capitalist, has quoted Bill Gates as arguing that a platform exists whenever “the economic value of everybody that uses it exceeds the value of the company that creates it.”¹³³ In that world, complementors earn rents. By contrast, Ben Thompson, a tech journalist, argues that Facebook and Google are what he calls aggregators—firms that completely control the relationship between suppliers and users.¹³⁴ His point is that this control allows the aggregator to exercise market power over one side of the platform, control access by the other side, and extract all the rents. This emphasis on the creation of complementor rents and their distribution or expropriation is a theme that runs through this Report.

When and how a platform appropriates the rents of its complementors varies and may engender different responses from complementors. At root, the reason the complementor is there in the first place is because the platform originally needed content to attract consumers, so it invited the complementors onto the platform. Businesses selling widgets on ecommerce sites, games on social media, and mapping apps on handsets, are all examples of complements that were critical to successfully launching a platform. The complementors make these investments thinking they will obtain a return, and that expected return leads to efficient levels of investment. If investments were made knowing that the returns would be zero, there might well be dynamic harm because the platform would not be able to get off the ground. (Or alternatively, the owner of the platform would have to vertically integrate into many applications and provide them itself.) A significant source of discontent today seems to come from complementors who invested believing that they could capture surplus with a good product and who are later expropriated by the platform. A second discontented group are the complementors that existed in the brick and mortar world and now have no choice but to use a dominant platform, so they have no bargaining power despite providing valuable products or services.

If the expropriation is accomplished using a tool that is anticompetitive, it may violate competition laws. In the United States, this antitrust violation can be established when the complementor may be able to disintermediate the platform and is therefore a potential rival, for example, Netscape and Windows.¹³⁵ If such a move is not a violation of competition law but

¹³³ Shah (2015).

¹³⁴ Thompson (2018) (contrasting platforms, such as Microsoft and Apple, which “need 3rd parties to make them useful and build their moat through the creation of ecosystems,” with aggregators, such as Facebook and Google, which “attract end users by virtue of their inherent usefulness and, over time, leave suppliers no choice but to follow the aggregators’ dictates if they wish to reach end users”).

¹³⁵ The antitrust theories of harm that cover this case are discussed below.

violates the expectations of market participants, it may be viewed as unfair. On the other hand, if the platform has significantly improved quality or engaged in innovation, market participants may think the platform fairly earned those rents.

Digital platforms have a variety of rent-extraction strategies that seem to be loosely related to the social scrutiny they are receiving.¹³⁶ For example, the news industry had no choice but to use Facebook. But Facebook has been reluctant to share any of its profits with news companies—and its market power has meant it hasn't had to.¹³⁷ Interestingly, when content providers have market power, the platform cannot always extract rents, even when it has a high market share. Take the hypothetical case of a fragmented travel platform industry in a geography where the airline market is concentrated. In that setting airlines could disintermediate a travel site by withholding their participation in the site, or by favoring their own sales channels. The equilibrium division of platform rents in this case will favor the powerful content providers, the airlines, rather than the platforms.¹³⁸ Market participants observe that Facebook has slowly commoditized most companies which supply attention to its users—from news and content markets to games and apps, companies' profits deriving from Facebook have slowly diminished as Facebook appropriated most of the gains.¹³⁹ These trends may be part of the reason there is growing attention to Facebook's business model today.

Today's prominent tech companies most likely learned from older tech businesses that those older businesses had “lost,” or shared, a too-high percentage of profits with their ecosystems. While older businesses were continually confronted with questions about what to own versus what to cede to partners on their platform, without today's technological advancements (including advanced data analytics and greater computational power), older platforms were not as swift and flexible in capturing value as current digital platforms can be. The increased scale and scope of control has provided modern digital platform owners with increased power over their ecosystems. Today's platforms understand that they can obtain higher margins if they either make all of the necessary complements themselves or position themselves as a mandatory bottleneck between partners and customers—leading to many platforms taking a significant commission on sales on their platform or extracting value through barter of information. In particular, today's digital platforms are very careful to maintain complete control over the user relationship so that they do not face any threat of disintermediation. These choices

¹³⁶ See, e.g., Dzieza (2018) (discussing the fear Amazon sellers have of the firm and its ability to shut them down with little or no explanation); Kaminska (2019).

¹³⁷ Thompson (2019) (“[C]ontent suppliers are absolutely commoditized: Facebook doesn't need to do anything to keep them on the platform, because where else will they go? Might as well keep the money for itself.”).

¹³⁸ For a related setting, see Morton et al. (2015) (on file with authors).

¹³⁹ See, e.g., Thompson (2019); Brown (2018); Weber (2016).

can be used to reduce the possibility of successful entry by an innovator in the platform's space. The next section turns to this problem of entry.

D. Harm to Entry, Including Disintermediation

1. The Practical Consequence of Entry Barriers

When evaluating entry as the main source of competition against a platform, as discussed above, the entrant has a significant disadvantage relative to the platform. Venture capitalist investors will often evaluate a startup based on its ability to either access or build enough data swiftly enough, all with the aim of reaching enough insights to take advantage of all the forces discussed above. In other words, a new entrant starved of data, quantitatively and qualitatively speaking, relative to a tech giant, is at a significant competitive disadvantage, and investors will be unlikely to invest if they view that data deficit as insurmountable. Although it can be attractive for a VC to invest in a firm that may be acquired by Facebook, Google, or Amazon, the road to a successful acquisition is fraught with danger, most notably the ability of any of these three giants to replicate a specific feature, functionality, or business model should other considerations such as intellectual property rights, team quality, defensibility, or time to market not weigh against such copycat strategy. Additionally, investors do not always fare well in these acquisitions.¹⁴⁰

2. Incumbent Incentive to Leverage Entry Barriers

There is growing evidence that conglomerate digital platforms are in an advantaged position to stop or block entry by more focused rivals when compared to traditional businesses. Companies like Alphabet, Amazon, and Facebook operate in multiple business verticals (for example, mail, maps, and search), collecting different dimensions of data on a consumer (for example, identity, location, and purchase intent) which give faster intelligence on competitive threats and new chinks in the platform's competitive armor. These companies can then derive superior insights into what firms they should block, which they should buy, and how they should grow strategically. This gives the platform an advantage over a rival entrant considering the same set of opportunities, and increases their abilities to exclude such rivals. A rival platform with similar economies of scope, data insights, and installed base may be a more formidable entrant.

If large digital platforms have both the incentive and ability to purchase and block entrants that compete with them, or might compete in the future, the question is whether they have done so. The evidence that platforms have bought a series of potential competitors in recent

¹⁴⁰ See, e.g., Kraus (2019) ("Ultimately, thanks to a 'last in, first out' philosophy, Eero's Series D investors, led by Qualcomm, will recoup 84 percent of their investments. The seed round and Series A-C investors will all get back 31 cents on the dollar.").

years is anecdotal but fairly robust.¹⁴¹ For example, many observers believe that Instagram and WhatsApp might well have been serious competitors for Facebook.¹⁴² The evidence that platforms have blocked potential entrants is likewise anecdotal¹⁴³ and was used in the Android and AdSense cases the European Commission has brought against Google.¹⁴⁴ More formal research in this area is essential.

3. Disintermediation and Foreclosure of Potential Platform Entrants

In addition to *de novo* entry, platforms fear disintermediation by a partner. If a platform's partner is able to directly access and serve the platform's customers, it can overtake the platform. Disintermediation can also occur through commoditizing services so that one side (normally the end consumer) is willing to substitute away from the platform—leading to a loss of profits.

The threat of disintermediation has important implications for the analysis of market entry and foreclosure. Modern platforms have an incentive to regularly thwart companies that compete with them for user demand. A platform that has total control of demand can steer customers to content and complements it owns rather than to those provided by independent firms that might challenge its market power. And because of potential harm to competition and complements, US regulators have often been tasked with preventing discrimination, foreclosure, and similar strategies in specific industries.¹⁴⁵ Without a great deal of insight into the technology and the strategy of the platform, foreclosure will be difficult to observe by outsiders. Technological tools such as mobile hardware technologies, advances in computing power, communications technologies, application programming interfaces (API), cloud computing technologies, and data analytics enable a strategy of keeping the attractive content from establishing a relationship with the user and thereby possibly entering as a competitor. Platforms have bluntly moved to prevent disintermediation and have engaged in foreclosure to block potential rivals. For example, Facebook acted to suppress the growth video-capture-and-sharing app Vine when Vine attempted to link its users to their Facebook friends.¹⁴⁶ Facebook CEO Mark Zuckerberg personally approved the decision to prevent Vine users from finding friends on the app via Facebook.¹⁴⁷

Exclusive contracts and loyalty contracts can also be used to achieve exclusion. For example, a long-term contract that requires an advertiser not to use an entrant can foreclose demand from that entrant, leading to exit. An exclusive contract with a global reach can prevent an able niche competitor from growing larger and obtaining economies of scale. Bundling of

¹⁴¹ See Furman (2019): 49, tbl.1A.

¹⁴² Wu (2018).

¹⁴³ “American Tech Giants” (2018).

¹⁴⁴ European Commission (2019).

¹⁴⁵ Khan (2019).

¹⁴⁶ Robertson (2018).

¹⁴⁷ See *id.* Several years after buying Vine, Twitter shut it down. See Newton (2016).

services by the incumbent platform can be designed to exclude entry or foreclose existing rivals. Contracts between platforms and advertisers that allow for individual negotiation can protect an incumbent from losing individual targeted sales to an entrant without requiring the incumbent to lower its prices across the board. An incumbent platform with market power will often have the incentive and ability to undertake these strategies and thereby preserve its profit.

4. Foreclosure of Complements to Capture Rents

One critical place to control the relationship is platform access. Amazon and Facebook regularly make decisions over which app or vendor is able to sell or is denied access to their stores and customers. Platforms often have a financial incentive to steer customers to particularly profitable products and can use the power of defaults and ordering to accomplish that effectively. Vendors operate in a risky environment where the platform's whims can determine its future as much or more than consumer satisfaction. If it chooses, a platform can steer demand elsewhere and the vendor loses access to its customers—because these customers are another company's users who single-home and buy from the default choice at the top of the page. Provided the consumer continues to find the totality of the platform experience positive, this bargaining power allows the platform to dictate business terms. The Vestager Report describes the setting clearly:

Other platforms impose rules and institutions that reach beyond the pure matching services and shape the functioning of the marketplace and, potentially, the relationship between the various platform sides, e.g. by regulating access to and exclusion from the platform, by regulating the way in which sellers can present their offers, the data and APIs they can access, setting up grading systems, regulating access to information that is generated on the platform, imposing minimum standards for delivery and return policies, providing for model contracts, imposing price controls and MFN clauses, etc. Such rule-setting and “market design” determine the way in which competition takes place.¹⁴⁸

The way competition takes place determines the level of profit achieved by the platform and each complement. In and of itself, a platform setting the terms of trade, quality levels, services, and so on may not be problematic if the purpose of the change is to “grow the pie” in a way that complementors view as fair (i.e., not involving expropriation). For example, Amazon, Facebook, or Google know in real time which products are sold to whom, at what price, and which packaging or incentives work, which may drive their rules. However, if these rules become opaque and uncertain or the insights gleaned from an app or vendor are biased or used against it in an asymmetrical manner, then the rule changes may not be about increasing everyone's revenue, but about moving a larger share to the platform.

¹⁴⁸ Crémer (2019): 60.

Merchants or vendors can find themselves banned, demoted in search results, or required to bear higher costs without the ability to move to a competing platform because either there is none or because the customers single-home, will not depart the platform because of the loss of one vendor, and cannot be reached elsewhere. The EC's Android case describes the disadvantage faced by independent apps that compete with the Google apps that are included in the mandatory bundle.¹⁴⁹ Twitter has vertically integrated into video streaming by foreclosing the rival service Meerkat.¹⁵⁰ Likewise, by selling logistics services to many of its sellers, Amazon gains an advantage when it wishes to launch a store brand. It can analyze the data from its rivals to develop an entry plan against those sellers. It is important to measure whether, and how much, quality increases with these strategies.

It is not clear what profits such complementors expected or achieved from the platform relationship, nor if there was significant relationship-specific investment required. Vendors may be less likely to enter the market at all, or to innovate, if they know they must distribute through a particular platform and their most successful products will be quickly copied. Businesses that could grow on a platform and increase the platform's attractiveness to consumers will be unwilling to invest if their profits are not secure, and this may be a source of dynamic inefficiency. However, platforms have an incentive to attract good complements in order to attract users. This incentive limits the platform's desire to expropriate complementor rents under some circumstances.

E. Harm to Innovation

1. Competition Promotes Innovation

There is significant theoretical and empirical research that concludes that anticompetitive creation or maintenance of market power will cause a reduction in the pace of innovation.¹⁵¹ This result is intuitive in the sense that firms "run faster" when they face competitors; competing firms will try to offer a better product on any dimension consumers care about, including innovation. Engaging in successful innovation is certainly both feasible and common for a large platform with its enormous collection of data and other assets. These companies routinely spend large sums on R&D, launch new products and services, and are more able than other competitors to derive superior insights into how they should innovate based on the data collected from

¹⁴⁹ The recent Spotify complaint against Apple in Europe is another example. *See* Vincent (2019).

¹⁵⁰ Zhu (2018) ("Meerkat, a mobile app that enabled Twitter users to broadcast live video streaming to their followers, vanished after Twitter acquired its competitor Periscope and cut off Meerkat's access to Twitter's social graph."); *see also* Welch (2012). Twitter acquired Periscope, a competing startup, and then decided to shut out Meerkat. On October 2016, Meerkat was shut down. Currently, Periscope is fully integrated with the Twitter platform and claims 1.9 million daily users (last updated and released info as of 2017).

¹⁵¹ For a comprehensive survey of the innovation literature, see the literature summarized in Federico, Morton, and Shapiro (2019) (on file with authors).

aggregating demand and advances in machine learning and advanced data analytics. However, the relevant counterfactual is whether the pace of innovation would be faster if platforms faced more robust competition.

2. Entry Barriers and Innovation

The lessening or blocking of innovative entry is of particular concern given its value. A VC will usually be wary of outright investing in an innovative startup that will implicitly or explicitly compete head-on with a tech giant. Given the tech incumbents' ability to block or foreclose a threatening entrant, the chance of successful entry is tiny. VCs would rather invest in businesses that are creating new categories or solving common technical issues. Take Google's search engine as an example. To our knowledge there is only one search engine that has reached a successful market size as a standalone business: Duck Duck Go. Google search is the dominant firm in this category with high entry barriers; despite the enormous size of this market, VC investors apparently do not want to fund an entrant. By contrast, VCs are attracted to startup teams that solve a specific issue for Google (and have a chance of buyout) as opposed to funding a team that wants to compete head to head. This dynamic leads to a self-fulfilling prophecy. Reduced VC investment due to the inability to enter successfully in fact causes less entry; and less investment also causes less differentiated innovation in the sector. This can be seen increasingly by evidence of platform acquisitions.¹⁵² Despite very high and stable profit margins, markets like social media and search have faced little entry.

As discussed above, large tech platforms can combine enormous financial resources with data resources. Digital platforms have an ability to produce free cash flows at a speed and level that is entirely new because of the combination of almost zero marginal cost, instant distribution, and global reach. This advantage creates both the incentive and the ability for the digital platform to outspend, to out-invest, or to acquire incumbents or new competitors. Incumbents have the incentive and ability to stand in the way of possibly disruptive innovation. With deep pockets, they can purchase possible future disruptors in order to align the path of innovation with their strategies or otherwise control it. This story is widely believed to be the reason that Facebook purchased Instagram and WhatsApp.¹⁵³

¹⁵² The number of potential competitors purchased by the tech giants is large. For example, Amazon has purchased Zappos, Fabric, CDNow, Quorus, Audible, Goodreads, and Quisdi; Facebook has acquired WhatsApp, Instagram, FriendFeed, and tbh; Google has bought Nest Labs, DoubleClick, YouTube, Waze, AdMob, Teracent, BeatThatQuote.com, Admeld, and Tenor. See "List of Mergers and Acquisitions by Amazon" (2019); "List of Mergers and Acquisitions by Facebook" (2019); "List of Mergers and Acquisitions by Alphabet" (2019). At times, these firms have barely bothered to disguise their anticompetitive purpose. See Kastrenakes (2018).

¹⁵³ Wu (2018).

Alternatively, platforms may create “kill-zones” around themselves.¹⁵⁴ For example, Facebook and Twitter’s aggressive API foreclosure, acquisition of competitors, and copying of new services have boosted their market power.¹⁵⁵ Facebook even acquired a mobile phone monitoring app, Onavo, which allows it to spot up-and-coming rivals and buy them or snuff them out.¹⁵⁶ While investment in innovation will continue, the type of innovation that will be funded will be broadly determined by the incumbent and its strategies. Disruptive innovation in markets that are characterized by high concentration levels and network effects is likely to be reduced compared to a competitive market. One of the few sources of entry in digital platforms comes from rival platforms that enter each other’s markets, as these large firms are more able to overcome entry barriers of all kinds.

3. The Reward for Innovation

Entrepreneurs may expect a low payoff to developing a free-standing product because of entry barriers and exclusionary conduct by the incumbent platform. In that case, its best hope is to be the preferred innovator of a complement and sell its business to the platform at an early stage. This source of financial reward is often cited as a reason why large platform acquisitions are good for society. Certainly, a reward of this type (a share in the platform’s business model and flow of rents) is better than no payoff and will stimulate some level of innovation. However, it is important to see that this incentive is distorted relative to that of the competitive context. Suppose public policy could reduce entry barriers and prevent anticompetitive exclusionary conduct towards entrants. In that case, the entrepreneur would not have to settle for a small fraction of the platform’s profits, but could compete for all of them and try to replace the platform. When a young firm has a chance to compete for the entire market on the merits of its innovation because there is no exclusionary conduct holding it back, success could lead to up to 100% of market profits, rather than a smaller acquisition payoff.¹⁵⁷ The availability of this option would allow entrepreneurs to develop improvements, complements, or replacements for a platform according to what has the highest returns. The inability to innovate in pursuit of the whole market leads to lower entry in tech sectors that are already dominated by a single large company.

¹⁵⁴ Google’s, Facebook’s, and Amazon’s annual conferences “held to announce new tools, features, and acquisitions, always ‘send shock waves of fear through entrepreneurs,’” according to one investment firm, and “[v]enture capitalists attend to see which of their companies are going to get killed next.” “The Future of Tech” (2018).

¹⁵⁵ Khan (2019); Robertson (2018).; *see, also, e.g.*, Griffith (2017).

¹⁵⁶ Seetharaman and Morris (2017).

¹⁵⁷ As previously noted, many of these markets will tend toward concentration naturally, so the size of profits is not affected by effective competition enforcement, but the number of years they can be earned will be affected when an incumbent cannot block entrants.

4. Current Trends in Innovation

The incipient but growing technical research supports a concern for the impact of big tech on innovation. Much of this research uses startup formation and its fuel, venture capital, as a proxy for innovation. In 2018, Facebook commissioned the consulting firm Oliver Wyman to write a report on this topic.¹⁵⁸ Using Facebook, Google, and Amazon as a proxy for big tech platforms, the authors drew four conclusions: i) Facebook, Google, and Amazon contribute a very small portion of the total venture capital in tech; ii) Facebook, Google, and Amazon M&A activity has no impact on aggregate investment levels; iii) Facebook, Google, and Amazon R&D has had no impact on venture capital investment levels; and iv) the presence of Facebook, Google, and Amazon does not dampen venture capital investment in technology relative to other mature sectors.

The first three conclusions are factual and unsurprising. But the fourth conclusion, that the presence of Facebook, Google, and Amazon does not dampen VC activity related to other sectors, is surprising. Ian Hathaway has summarized a straightforward rebuttal.¹⁵⁹ The Facebook-commissioned report looked at industry-wide effects. However, as Hathaway points out, viewing the industry in the aggregate masks the effects of dominating firms. By looking at the sub-industries associated with each firm—social platforms (Facebook), internet software (Google), and internet retail (Amazon)—a different trend emerges. Since 2009, change in startup investing in these sub-industries has fared poorly compared to the rest of software for Google and Facebook, the rest of retail for Amazon, and the rest of all VC for each of Google, Facebook, and Amazon. This suggests the existence of so-called “kill-zones,” that is, areas where venture capitalists are reluctant to enter due to small prospects of future profits.¹⁶⁰

In a study of the mobile app market, Wen Wen and Feng Zhu come to a similar conclusion: Big tech platforms do dampen innovation at the margin.¹⁶¹ Their study analyzed how Android app developers adjust their innovation strategies in response to entry (or threat of entry) by Google:

[A]fter Google’s entry threat increases, affected developers reduce innovation and raise the prices for the affected apps. Once Google enters, the developers reduce innovation and increase prices further. However, app developers’ incentives to innovate are not completely suppressed; rather, they shift innovation to unaffected and new apps. Given many apps already offering similar features, Google’s entry may reduce social inefficiency.¹⁶²

¹⁵⁸ “Assessing the Impact” (2018).

¹⁵⁹ Hathaway (2018).

¹⁶⁰ Hathaway (2018).

¹⁶¹ Wen and Zhu (2018).

¹⁶² Wen and Zhu (2018).

Ultimately, these studies are suggestive but not determinative. Much more research needs to be done to properly identify the existence and extent of “kill zones” for market entry and innovation. Nonetheless, the evidence thus far does suggest that current digital platforms face very little threat of entry and are negatively impacting investment in key digital areas. This is reinforced by the fact that the key players in this industry remained the same over the last two technology waves, staying dominant through the shift to mobile and the rise of AI.¹⁶³ In the past, dominant business found it difficult to navigate innovation or disruption waves. By contrast, Facebook, Google, Amazon, Apple, and even Microsoft were able to ride these waves without significant impact on market share or profit margins. This indirect evidence corroborates the argument that these companies are facing few competitive threats.

There is an informal sense from practitioners that in mobile, cloud computing, and AI the result is the same: Innovation at the core has slowed and is centered around the tech giants.¹⁶⁴ As a result, new innovation waves have emerged, centering around blockchain, distributed ledgers, and decentralized models (the new Silicon Valley growth paradigm); and with specialized AI, niche retail applications, or, in the enterprise sector, b2b models. It should be noted that, apart from blockchain, whose decentralization eliminates network effects (in theory), the other areas of innovation do not necessarily herald a change in future market structure. Some have also argued that, in markets that are characterized by significant market power, innovation will likely shift from focusing on consumer benefit (to entice a rise in market share), to consumer exploitation.¹⁶⁵

In summary, unlike last-century’s businesses, digital businesses will often lack competition *in* the market. This absence may be redressed by intensive competition *for* the market, as firms compete to be the winner that takes all. However, if dominant firms are able to maintain their dominance, even as the external environment and tastes change over time, consumers may be denied the benefits of competition for the market as well.

¹⁶³ This dynamic affects where we see innovation, and by whom. While distributed technology seems to be a relatively open playing field with the potential for many entrants, the bulk of innovation in other spaces is dominated by tech giants. This is the case, for example, in the AI and voice-assistant space. In the AI field, innovation is driven by startups leveraging and customizing the existing technology and libraries developed by Microsoft, Amazon, Facebook, and Google. These companies have been quick to acquire promising ventures. Since 2011, Facebook has acquired seven AI startups: Face, Protogeo, Wit.ai, Bloomsbury AI, Dreambit, Grokstyle, and Ozlo. Since 2010, Amazon has also acquired seven startups in the AI space: Yap, ORbeu, GoButler, harvest.ai, Graphiq, Body Labs, and Dispatch. Finally, Google has acquired 20 AI startups since 2011, the most notable being Kraggle and DeepMind. See “List of Mergers and Acquisitions by Facebook” (2019); “List of Mergers and Acquisitions by Alphabet” (2019). The AI space may not see a successful IPO because of these dynamics. Similarly, innovation in voice-assistant technology has been dominated by tech giants. Microsoft’s Cortana, Apple’s Siri, Google’s Google Voice, and Amazon’s Alexa are examples of dominant firms leveraging their previous offerings to retain a first-mover advantage in the market.

¹⁶⁴ See, e.g., Griffith (2017); “Tech Monopolies” (2018); “Andreeson Horowitz” (2018).

¹⁶⁵ Ezrachi and Stucke (2016).

F. Old Wine in a New Bottle, or a New Reality?

The market characteristics and problems discussed above in Part I and Part II raise an important question: Is there something new or different about the digital economy, which justifies a revised approach to intervention? Is this simply old wine in a new bottle, the same old problems of network effects and concentration, or do we face a new reality? As we outlined above, we believe the digital platform may be a unique combination of economic forces that requires both new analysis and new public policy.

We expect some to argue that there is no need for action. It is in the interest of the powerful incumbent to deny there is any problem to solve in order to delay while the market remains in its hands. The waiting game allows incumbents to collect profit from the status quo and use those profits to raise doubts or marginalize concerns in order to delay or prevent effective intervention.¹⁶⁶ However, the body of research that indicates increasing problems of underenforcement is growing rapidly.¹⁶⁷

Of course, the rapid developments of our digital world are not yet fully understood and merit ongoing study. However, the material above makes clear that when market power becomes entrenched, it is difficult to return to a competitive market. Inaction could create decades of market power leading to weakened innovation, increased rent extraction, and social domination by gatekeeper tech firms. Making no policy change and standing back to let platforms carry out the strategies of their choice carries unusual risk. The new reality of the digital landscape presents new challenges that will not easily be resolved through organic competition or current US antitrust enforcement norms. The United States now risks missing the moment to confront a serious challenge to competition. That risk spurs our search for possible solutions, which we outline in Part IV.

¹⁶⁶ See, e.g., Mullins Jack Nicas (2017) (relating Google's funding of professors for, among other things, research opposing antitrust intervention in the search engine market).

¹⁶⁷ Morton (2019).

Solutions

In this part, we outline a range of solutions to address the issues raised above. Needless to say, there is no magic pill. The legislature, judiciary, and regulators will each need to play a role. Each of the proposed solutions brings with it benefits and costs and should be considered as part of a balanced policy. Still, with these limitations in mind, these solutions have the power to address changing market dynamics and enable society to move toward a more efficient, open, and dynamic market environment.

Competition policy is only one facet of society's wider goals. Alongside the promotion of competition policy, policymakers may wish to implement policies to promote privacy, democratic accountability, and fairness. While these goals may be advanced indirectly through the measures proposed below, they can also be addressed more directly through laws and regulations that we do not consider here. Other committees provide valuable insights on possible advancements on that front.

Having established the impact of high entry barriers, we open with a fundamental question about the ability of a market to self-correct.

A. Will Markets Self-Correct When Competition Problems Arise?

The question whether the market will self-correct is central to antitrust policy. It affects the scope of illegality, as well as the zeal with which competition agencies approach intervention. It also affects the gain from implementing both ex ante and ex post solutions to insufficient competition. When there is a reasonable chance that technological progress and innovation will speedily prevent or remove bottlenecks and maintain an active competitive process, intervention should be limited. For example, the stock trading app Robinhood Markets, which slashed per-trade fees to zero, has forced competitors to scramble—even JP Morgan is now offering a free stock trading app for consumers.¹⁶⁸ Likewise, Amazon and Wal-Mart are currently vigorously competing for fast and cheap delivery services, forcing others to follow.¹⁶⁹

Acolytes of the Chicago School have persuaded many federal officials and judges that markets will generally quickly self-correct, and that antitrust scrutiny should therefore be limited. This conviction has extended into digital markets. Opponents of government intervention point to the dynamic nature of technology, the role of disruptive innovation, and firms' strong investment in research and development. They defend the adequacy of a free-market,

¹⁶⁸ Son (2018).

¹⁶⁹ Bhattacharyya (2019); Cheng (2018); Statt (2018).

noninterventionist approach, and argue that the cost of over-intervention outweighs the benefits of intervention.

While some markets may self-correct, the findings of this report suggest that rapid self-correction in markets dominated by large digital platforms is unlikely, and that harms to economic welfare from the exercise of market power in such markets are substantial. As discussed above, entrants find it difficult to overcome the high barriers to take on digital platform incumbents. Economies of scale, economies of scope, network effects, and negligible marginal cost all work together to make entry difficult in existing markets. Moreover, while monopoly profits are a lure to competitors, incumbents can use those very profits to entrench themselves and protect their position. No matter how dynamic the technology, an entrant will not unseat a monopolist if the monopolist is permitted to buy the dynamic entrant for a share of monopoly profits. Both parties gain from such a transaction—and the public loses.

The result is less entry than a more competitive environment would create. Less entry into digital markets means fewer choices for consumers, stunted development of alternative paths of innovation, higher prices, and lower quality. Self-correction is not a realistic expectation in this environment—indeed, the available evidence suggests it has not happened—and public policy should not rely exclusively on it. Effective antitrust enforcement and regulation must take account of this reality. If there is a force toward self-correction, it may require active promotion to succeed, and in this way public intervention can be complementary rather than antagonistic to market forces. Indeed, the other reports that have addressed this problem around the world have accepted that policy changes are necessary in order to avoid stagnant and harmful digital markets. We now turn to what policy options exist, and which of these the government might adopt.

B. US Antitrust

1. Basic Principles

Antitrust law is intended to prohibit private conduct by firms that reduces economic welfare. Although application of antitrust principles can be complex and the specific doctrinal embodiments of those principles can seem arcane and arbitrary, at its core US antitrust law is simple.¹⁷⁰ It can be summarized in a single sentence: *Private conduct that creates or increases market power, other than by efficiency-based competition on the merits, is illegal.*

There are two fundamental components of any antitrust violation. The first is bad conduct, commonly called “anticompetitive conduct.” The second is a resulting harm to economic welfare from the creation of more market power than would otherwise exist. A firm

¹⁷⁰ Melamed (2019).

gains market power when the competitive discipline imposed upon that firm by either actual or potential future rivals is reduced. One way to achieve market power is through price-fixing and other forms of collusion, but we will not address such behaviors here. Rather, this report will focus on antitrust law as it applies to mergers and non-merger exclusionary conduct.

Importantly, a firm does not violate the antitrust laws if it gains market power by competing on the merits rather than as a result of anticompetitive conduct. For example, if a firm enters the market with a fantastic new product that has strong network effects, competes solely on the merits of that product with no anticompetitive conduct, and finds itself with 95% market share, it would not have violated the antitrust laws. Similarly, a firm does not violate the antitrust laws if it engages in anticompetitive conduct but that conduct does not harm the market as a whole and fails to result in additional market power.¹⁷¹

Box IV

More About US Antitrust Law

There are four substantive elements to any US antitrust violation: Private, as opposed to government, conduct; anticompetitive conduct; creation or increase in market power; and a causal connection between the conduct and the market power. There are three basic types of conduct that can be anticompetitive. They are (i) mergers that lessen competition, (ii) conduct that excludes or weakens actual or potential rivals, and (iii) conduct that constitutes or facilitates collusion (e.g. cartels) among firms that would otherwise compete more vigorously. Possible antitrust problems involving digital platforms are most likely to arise from the first two types of conduct.

Anticompetitive conduct is conduct that is likely to lead to the creation or maintenance of market power for reasons other than an increase in allocative efficiency (i.e., by increasing output or decreasing price where price is not less than marginal cost) or productive efficiency (i.e., by reducing cost or increasing product quality). The range of potential anticompetitive conduct is nearly boundless and can include burning down a rival's factory, designing products to be incompatible with rivals' products, acquiring smaller rivals to shut them down or deny them to other suitors, exclusive dealing and Most Favored Nation clauses (MFNs) in some circumstances, and some forms of tying or bundling multiple products or services. Digital markets are prone to new and innovative violations of the antitrust laws because marginal costs are often close to zero, the business models themselves are often new and innovative, and firms are able to be creative with new product bundles, contracts, and transactions.

¹⁷¹ Certain types of conduct, notably agreements among competitors fixing prices or allocating customers or service areas, are regarded as unlawful per se, which means that they are unlawful without proof of harm to competition in the market as a whole. Per se illegality in these circumstances is not a rejection of the principle described in this text. Rather, per se illegality reflects a pragmatic judgment that those types of conduct are so likely to harm competition and so unlikely to provide any welfare benefits that it would be needlessly costly and burdensome to require proof of harm to the market as a whole on a case-by-case basis.

More About US Antitrust Law (Cont.)

Current popular debate often seems to imply that high market share in and of itself is a violation of the antitrust laws. It is important to understand that a firm can violate the antitrust laws *only if it engages in anticompetitive conduct*, even if its conduct causes the firm to gain monopoly power. Thus, for example, a firm does not violate the antitrust laws by gaining market power solely because it has a better product that consumers choose to buy for that reason or because it develops a better distribution channels that consumers find convenient.

Since at least the middle of the twentieth century, US antitrust law has embraced the view that those who do not engage in anticompetitive conduct are entitled to the fruits of their labor, including any market power or dominance that they might have gained. This view rests on the concern that breaking up or restricting firms that obtain their success by competition on the merits would deter the very kind of aggressive but procompetitive conduct that the antitrust laws are intended to encourage. Such no-fault intervention would deter such conduct, both because firms would fear that too much success would end up hurting them and because they would be uncertain about the antitrust implications of their conduct. While firms that achieve durable market power, especially in industries that are regarded as “natural monopolies,” are sometimes subject to sectoral regulation, durable market power is not itself a sufficient basis for antitrust intervention. Society has other legal mechanisms to regulate an industry that is not delivering on social goals.

US antitrust law has also long required that antitrust plaintiffs prove, not just anticompetitive conduct, but also that the conduct harmed competition in the market as a whole.¹ This requirement limits antitrust enforcement to those matters that are important enough to warrant costly and burdensome antitrust proceedings and ensures that antitrust law remains focused on making markets work.

Both the anticompetitive requirement and the market power requirement protect the competitive process, another concern that arises in the popular debate. The conduct requirement does this by focusing antitrust enforcement on conduct that is not efficiency-based and can thus create market power only by distorting competition on the merits. The market power requirement does this by focusing antitrust enforcement on conduct that impairs the competitive discipline of rivals and, in that way, harms the competitive process. Many of the current proposals for more aggressive antitrust enforcement, although couched in antitrust language, would dispense with either the bad conduct requirement or the market power requirement.

Conduct that harms competition without creating any efficiencies can readily be characterized as anticompetitive. Examples of such conduct include mergers among competitors that do not create efficiencies by combining complementary assets, conduct whose only purpose is to harm or exclude competitors, and agreements among competitors about their prices or other terms of trade.

¹ See *Brunswick Corp. v. Pueblo Bowl-O-Mat, Inc.*, 429 U.S. 477, 488 (1977) (“The antitrust laws . . . were enacted for the protection of competition, not competitors.”) (internal citation and quote marks omitted).

More About US Antitrust Law (Cont.)

When conduct both generates efficiencies and impairs the competitive efficacy of rivals, the question whether it is anticompetitive is less clear and will depend on specific facts. The conduct is more likely to be unlawful if substantially all of the efficiencies could be achieved by an alternative that is less harmful to rivals or if the efficiencies are insubstantial and the harm to rivals is great. Courts rarely deem conduct anticompetitive if it generates substantial efficiencies that cannot be realized by an available alternative, regardless of the impact of that conduct on rivals.

As explained in greater detail below, antitrust enforcement decisions are often made with great uncertainty about the efficiencies and market impact of the conduct at issue. For the past 30 years or so, antitrust law has reflected the view that mistakes in the direction of over-enforcement are more problematic than those in the direction of under-enforcement. One premise of this view is that under-enforcement results in market problems that are likely to be corrected by innovation or other market developments. As noted above, this report assesses the state of the evidence on entry and innovation and concludes that such self-correction is unlikely to be prompt and effective, in part because digital platforms that have achieved market power possess formidable barriers that inhibit entry and innovation. The assumption that over-enforcement is costly is reflected in aspects of antitrust doctrine applicable to exclusionary conduct, vertical contracts, and merger review as well as in courts holding antitrust plaintiffs to demanding standards of proof.

The relative concerns about the risks and cost of over-enforcement and under-enforcement embodied in current antitrust law has been subject to increasing criticism. This report concludes that the tolerance of antitrust law for under-enforcement should be reassessed and recalibrated. Careful recalibration of the benefits and costs of different mistakes is likely essential to the development of more aggressive antitrust law that remains consistent with the principles described above.

While US antitrust law has long been flexible in combatting anti-competitive conduct, there is increasing concern that it has been underenforced in recent years. First, there is increasing evidence that the enforcement agencies and courts have permitted too many mergers between competing firms that have led to post-merger price increases and other indications of increased market power.¹⁷² Vertical mergers are rarely challenged by the enforcement

¹⁷² See Kwoka (2014); Baker (2018): 39 (listing six reasons to believe that industry concentration cannot be explained by scale economics, including the fact that various forms of anticompetitive behavior are “insufficiently deterred” and that “market power is durable”); Kwoka (2017); Kwoka (2015); Dafny, Duggan, and Ramanarayanan (2012) (“[T]he mean increase in local market [firm concentration] between 1998 and 2006 (inclusive) raised premiums by roughly 7 percent from their 1998 baseline”); Ashenfelter and Hosken (2010). *But see* Vita and Osinski (2018). *See also* Kulick (2017); Blonigen and Pierce (2016): 24.

agencies,¹⁷³ and claimed or expected merger-related efficiencies are often not realized. Second, monopsony—market power of buyers or employers—appears to be a growing problem.¹⁷⁴ Recent studies suggest that labor markets are less competitive than previously thought and employers have exercised market power against workers in those markets.¹⁷⁵ Third, there have been few antitrust challenges to exclusionary conduct since the government’s 1998 case against Microsoft, and courts have in several instances been hostile to such cases and have imposed daunting proof requirements on plaintiffs. Apparent under-enforcement is in part due to courts’ reliance on so-called Chicago School assumptions that do not have a sound theoretical or empirical basis.

Regardless of whether or not antitrust enforcement has failed to keep up with conduct in the economy more generally, the challenge of enforcing in the area of digital platforms presents new issues. The platforms create new competitive environments; they provide opportunities for new types of anticompetitive conduct; and they create new economic and conceptual challenges for antitrust enforcement. This section is focused on analysis and recommendations designed to help such future enforcement, though that enforcement will often be addressed to conduct that occurred in the past.

The challenges facing future antitrust enforcement are more than just analytical and intellectual. Antitrust law and its application by the courts over the past several decades have reflected the now outdated learning of an earlier era of economic thought, and they appear in some respects inhospitable to new learning. Antitrust enforcement better suited to the challenges of the Digital Age may therefore require new legislation.

¹⁷³ Hoffman (2018) (admitting that “vertical merger enforcement is . . . a small part of [the FTC’s] merger workload”).

¹⁷⁴ *See, e.g.*, Kades (2018).

¹⁷⁵ Preger and Schmitt (2019).

Box V

EU Competition Law

EU competition law, like US antitrust law and the competition laws of most nations, is intended in large part to promote economic welfare by prohibiting private conduct that injures or is likely to injure competition. Not surprisingly in light of this shared objective, US and EU law are very similar.

With respect to market power, similar to US antitrust law, a firm does not violate the EU Competition laws if it gains power by competing on the merits rather than as a result of anticompetitive conduct. EU Competition law only condemns the abuse of a dominant position. An abuse, under EU law, may however include a wider range of prohibited exclusionary and exploitative practices than under US antitrust law. As a result, some practices which may not trigger enforcement action under US antitrust law, may nonetheless be regarded as infringing EU competition law.

In this context, it is also worth noting the (wider) goals of EU competition law. According to the European Commission, competition on the market is protected ‘as a means of enhancing consumer welfare and of ensuring an efficient allocation of resources.’¹ This notwithstanding, EU competition law has also consistently been held to protect ‘not only the interests of competitors or of consumers, but also the structure of the market and, in so doing, competition as such.’² Moreover, a genuinely indigenous objective is worthy of note, namely that of promoting European market integration.³

The multitude of competition goals, and their position within the wider normative EU values may sometimes contribute to possible inconsistencies between the EU and US analysis.⁴

EU Competition law is enforced both at public and private levels; public enforcement is carried out by the European Commission and by the Competition Authorities of the member states. Private enforcement takes place in the courts of the member states, where private parties may bring follow-on or stand-alone damage claims.

1 European Commission, ‘Guidelines on the Application of Article 81(3) of the Treaty’ [2004] OJ C101/97, para.13 (hereinafter ‘the General Guidelines’).

2 Case C-501/06 P *GlaxoSmithKline Services Unlimited v Commission and Others* [2009] ECR I-9291, para 63. See also Case C-8/08 *T-Mobile Netherlands and Others* [2009] ECR I-4529, paras 31, 36, 38-39; Council Regulation (EC) 1/2003 on the Implementation of the Rules on Competition Laid Down in Articles 81 and 82 of the Treaty [2003] OJ L1/1, Recital 9; European Commission, ‘Green Paper on Vertical Restraints in EC Competition Policy’ COM(96) 721 final, para 180.

3 Information Service High Authority of the European Community for Coal and Steel Luxembourg, ‘The Brussels Report on the General Common Market’ (June 1956) (Spaak Report); David J Gerber, ‘The Transformation of European Community Competition Law?’ [1994] Harvard Intl LJ 97, 102.

4 Ezrachi, Ariel, EU Competition Law Goals and the Digital Economy (June 6, 2018). Oxford Legal Studies Research Paper No. 17/2018. Available at SSRN: <https://ssrn.com/abstract=3191766>

C. Special Challenges Presented by Technology Platforms

Technology platforms present particular challenges for antitrust enforcement. Markets tip and resulting market power is durable, so even effective antitrust enforcement is unlikely to generate fragmented markets. Nonetheless, enforcement that protects competition on the merits in the first stage and prevents exclusionary conduct in the second stage will help ensure that market-participants make unfettered choices among competing platforms and that entry and innovation are not inhibited by private rent-seeking. Additionally, these markets move very quickly in areas such as new product introduction, foreclosure, and tipping. Antitrust litigation does not move quickly. Effective antitrust enforcement should move as rapidly as is practicable. However, enforcers will be most effective when they choose enforcement priorities and remedies to generate optimal deterrence of anticompetitive conduct.

As explained above, digital platforms are able to exploit behavioral biases to their advantage. Economists describe these strategies as falling into two conceptually distinct categories of conduct: exclusive (strategies that foreclose competition) and collusive (strategies that cause higher prices).¹⁷⁶ The former occurs when platforms exploit behavioral biases to keep consumers attached to their platforms and make switching to alternatives more difficult. These tactics generally make consumers less receptive to competitive alternatives—they lower contestability—and thus raise entry barriers. Platforms also exploit behavioral biases—such as hyperbolic discounting and limited self-control—to extract surplus from both consumers and content providers. Strategies such as offering addictive content at moments when consumers lack self-control increase time spent on the platform and profitable ad sales even as the platform lowers the quality of content. These tactics increase the welfare costs of market power.

Many technology platforms are distinctive because they provide valued services to consumers without charging a monetary price. Instead, consumers barter their attention and data to the platforms in exchange for these services. The platforms use that attention and data to generate monetary payments from advertisers. While a barter transaction is, in principle, subject to antitrust scrutiny just like any other transaction, antitrust enforcement has had vastly more experience with transactions based on monetary prices, and that experience has prompted the development of sophisticated tools to analyze money prices.¹⁷⁷ Where monetary prices are fixed at zero while quality changes over time—in response to changes to the nature of the services, privacy protections, content offerings and the like—the quality-adjusted prices change. Because economists, antitrust agencies, and courts have less experience with quality, they lack equally sophisticated tools for analyzing changes in quality-based prices. This is partially because quality

¹⁷⁶ Salop (1986).

¹⁷⁷ For example, the US agencies' traditional definition of an antitrust market is one in which a monopolist could profitably raise price by a non-insignificant amount for a significant period of time. See "Horizontal Merger Guidelines" (2010): 9, § 4.1 (defining the hypothetical monopolist test).

naturally presents itself in a less quantified form than price, but also because in regular markets price often adjusts to quality, not the other way around, so economists have been able to rely on price as one measure of quality. When enforcers are able to quantitatively link quality to price, they will be better equipped to fit digital cases into older, price-centered jurisprudence.

Technology platforms also pose unusual challenges for antitrust merger enforcement. To the extent that platforms are in winner-take-all or winner-take-most markets, mergers among significant, existing competitors, which are the mergers most subject to antitrust challenge, are likely to be rare. Instead, competition in such markets is largely for the future, often in evolving and very different markets. This competition is sometimes called “competition for the market” or “leapfrog” competition. In this context, acquisition by a dominant platform of a much smaller and possibly nascent firm could be very damaging to competition if, absent the acquisition, the smaller firm would develop into a major competitive threat or would lead to significant change in the nature of the market. In a concentrated market structure, this potential competition from very small entrants may be the most important source of competition faced by the incumbent firm.¹⁷⁸

The problem is that it is very difficult to know at the time of an acquisition whether the acquired firm is likely to develop into a competitor or whether, to the contrary, acquisition by the platform offers the most promising path to the commercial development and use of the acquired firm’s new technology or an essential exit strategy for investors in the acquired firm. Antitrust agencies and antitrust law need to develop a better understanding of the circumstances under which acquisitions of nascent competitors might be anticompetitive.¹⁷⁹ Enforcers may have to learn to think more as venture capitalists do and understand their analytics in order to better identify harms. A pattern of repeated purchases of small potential competitors that could have

¹⁷⁸ The D.C. Circuit in *Microsoft* endorsed this principle. See *United States v. Microsoft Corp.*, 253 F.3d 34, 79 (D.C. Cir. 2001) (en banc) (per curiam) (“[I]t would be inimical to the purpose of the Sherman Act to allow monopolists free reign to squash nascent, albeit unproven, competitors at will.”).

¹⁷⁹ In addition, because small acquisitions are not subject to pre-merger review under the Hart-Scott-Rodino Act, agencies are often unaware of the acquisitions until after they are consummated. The threshold for pre-merger notification was raised in December 2000 from \$10 million to \$50 million. While transactions under the threshold are still subject to Clayton Act enforcement, agencies do not see them in advance, may not find them at all, and can only sue to unwind them after the fact. After the filing threshold increase, there was a sharp uptick of newly non-notified mergers (between \$10 and \$50 million) between direct competitors—the type of mergers that likely would have been blocked during HSR review, had it occurred. Small technology mergers fall in this category also and are rarely reviewed or challenged. See Wollman (2019). An example of a problematic merger of this type that was challenged by DOJ after consummation is *Bazaarvoice-PowerReviews*. *United States v. Bazaarvoice*, 3:13-cv-00133-WHO, 2014 WL 203966 (N.D. Cal., Jan. 8, 2014). While the government prevailed in that litigation, because the two software firms had gone some way down the integration path, the result was arguably not as favorable for consumers as if the transaction had been notified and blocked ex ante which would have preserved two robust competitors. In general, the evidence demonstrates that firms act quickly to obtain market power in the absence of government enforcement, and that the requirement to notify the government serves as an important deterrent to anticompetitive mergers.

developed into substitutes, or an acquisition price that reflects a sharing of monopoly rents, might be useful indicators of possible competitive risks.

Antitrust law might also have difficulty policing conduct by established technology firms that tends to exclude or marginalize smaller rivals. Recall that at launch a platform invites and encourages complementors to provide content and functionality on its platform. It does this when it is competing for consumers' attention against other platforms because it wants its platform to be attractive. A successful platform creates an ecosystem that is valuable to consumers. However, one or both of the following issues may arise. In the first, a complement to the platform seeks to become its horizontal competitor. In the second, the platform seeks to become a horizontal competitor of one of its complements. Let us take these one at a time.

First, a complement can develop the ability to form a relationship with the end user that is sufficiently free-standing and valuable to take the user off the platform and into a separate relationship with the complement. The platform has an incentive to foreclose the complement to prevent this loss of market power and profit. Because the complement is transitioning into direct horizontal rivalry with the platform, US law does not have any trouble recognizing antitrust violations of this form. For example, content providers like Yelp are Google's complements—people want to be able to find Yelp reviews via Google—but Yelp is also a search engine that could grow to rival Google search. Google has entered the restaurant review market with its map product, and Yelp now alleges that Google is engaging in foreclosure.¹⁸⁰

Second, if the platform observes that a complement (say, complement Z) is earning strong profits, the platform may seek to enter that complementary market. Because the platform and rival complement Z providers are now horizontal competitors, the platform has the incentive—and, often, the ability—to foreclose those competitors. This could take the form of banning rival complements from the platform, reducing their ability to interoperate, raising their costs, steering customers elsewhere, and so forth. If the platform's new product Z is a wonderful innovation, there will be an efficiency (in the form of higher quality) to weigh against any harms from less competition.¹⁸¹ At root, there is a fundamental question about whether all the rents of the platform are part of the competitive return to the creator of the platform, particularly if the complementors' added-value can be driven to zero once the network effects are operational. In a setting with single-homing users, those complements may have no bargaining power *ex post* and the platform may have the incentive and ability to give them a share of zero. This is a

¹⁸⁰ Duhigg (2018).

¹⁸¹ Khan (2017). *See also* Zhu and Qihong (2018) (finding that “[w]hile Amazon's entry discourages affected third-party sellers from subsequently pursuing growth on the platform, it increases product demand and reduces shipping costs for consumers”).

particularly interesting problem when those complementors added value to the platform at a critical competitive moment before network effects and entry barriers protected it.

Addressing this issue, the U.K. Furman Report and the EC Vestager Report both devote significant discussion to platforms being “fair” to complementors. First, a sectoral regulator is likely to be better than the antitrust laws at enforcing fairness norms. Antitrust law is focused, not on fairness itself, but on anticompetitive conduct that creates market power and might therefore permit various types of conduct that some might regard as “unfair” because, for example, a platform extracts the rents of partners that offer a good product but have no bargaining power. Because the complementor is not threatening to replace the platform, the simple exclusion theory of harm explained above does not apply directly. The situation is informally akin to the familiar “open early, closed late” strategy.¹⁸² In the “open early, closed late,” strategy a firm at first encourages others to become dependent on connecting to it and relying on it, and later uses their dependence to shut out competitors or extract monopoly rents.¹⁸³ This type of behavior might be regarded as exploitation or excessive pricing and thus deemed to be an unlawful abuse of dominance under EU law. It would violate US law only if, among other things, it enabled the platform to gain or preserve market power it otherwise would not have in either the platform market or the market in which the excluded firm did business. In the latter case under US antitrust law, this conduct might be captured under the “duty to deal” framework.

Because large technology platforms have huge scale and benefit from network effects, they are often able to engage in aggressive conduct targeted at rivals without violating existing antitrust standards. The platforms might be able to copy rivals’ innovations or otherwise increase the value of their services to consumers without pricing below cost, and they might be able to insist that rivals using their platforms enter into agreements, such as agreements regarding access to consumer data, that enhance the platform’s ability to compete. A platform might be able to require exclusive contracts or loyalty-based contracts that cause single-homing by one side (e.g. drivers or consumers). Such single-homing might cause the market to tip in the platform’s favor and exclude a competitor. This kind of conduct often has efficiency benefits that make it difficult to challenge the conduct under the antitrust laws. Certain aspects of antitrust law might be adjusted, for example antitrust rules intended to prohibit dominant firms from engaging in conduct that would exclude an equally efficient competitor could be revised to better protect smaller competitors. Nonetheless, it is unlikely that such adjustments would entirely eliminate the competitive advantages inherent in large firms with substantial scale and scope economies.

¹⁸² Sher and Tennis (2016); Shapiro (2005): 15 (“[I]n a network industry, a firm might obtain a dominant position based in part on certain ‘open’ policies that induce reliance by complementary firms, and then later exploit that position by offering less favorable interconnection terms or by refusing to interconnect with them altogether.”).

¹⁸³ For an example, see the discussion of Cisco-Arista at note 58, *supra*, and accompanying text.

Pinpointing the locus of competition and therefore the relevant market in which technology platforms compete can also be challenging because the markets are multisided and are often ones with which economists and lawyers have little experience. This complexity can make market definition another hurdle to effective enforcement. For example, two platforms might compete in general search, while also each offering social media and mapping functionalities, among other services. Advertisers that buy ads on searches may be a common set of customers. While courts and agencies have substantial experience analyzing advertising markets, for example, they are less knowledgeable about markets for attention or barter transactions involving data made available to providers as an unintended byproduct of using a digital platform. The problems are compounded by the facts that technologies surrounding the products' functions in digital markets are continually changing and changes in quality-adjusted prices are difficult to observe.

Only one litigated US case has explicitly addressed these issues. That case culminated in a 2018 decision by the US Supreme Court.¹⁸⁴ Although the decision has been praised by some conservative commentators,¹⁸⁵ it has been widely criticized by others.¹⁸⁶ The case itself involved the credit card business and what the Court called a “transaction platform,” in which the platform (American Express) facilitated simultaneous transactions between consumers and merchants.¹⁸⁷ It should not, therefore, be legal precedent applicable to other kinds of platforms like Google and Facebook. But the case does suggest that the five-Justice majority on the Court is hostile to antitrust enforcement (at least in vertical and exclusion cases), does not understand multi-sided markets very well, and might be more influenced by ideological preconceptions than by evidence in the case or fact-finding by district court judges.¹⁸⁸

The harm from lack of competition in digital markets will manifest itself in quality and innovation, as well as from higher prices to advertisers. As detailed by Giulio Federico and colleagues, the impact on consumer welfare of a decline in innovation due to lack of competition is likely to be large, especially in the case of fast-moving technologies that affect many consumers and related businesses.¹⁸⁹ Very often the uncertainty involved in evaluating harms to innovation will be high, especially in contrast to the analysis of price forecasts. It is possible to measure pipeline projects and current R&D to obtain a sense of competitive overlap or

¹⁸⁴ Ohio v. American Express Co., 138 S. Ct. 2274 (2018).

¹⁸⁵ See, e.g., Federal Trade Commission (2018): 228-235 (Vinson & Elkins partner Darren Tucker praising the *American Express* decision).

¹⁸⁶ See, e.g., Carlton (Forthcoming); Hovenkamp (2019); Melamed and Petit (2019); Katz and Sallet (2018). See also Federal Trade Commission (2018): 226 (Goldstein & Russell partner Eric Citron describing the *American Express* decision as “economically illiterate”).

¹⁸⁷ *American Express*, 138 S. Ct. at 2277.

¹⁸⁸ Anthony Kennedy, one of the members of the five-justice majority in *American Express*, has since retired, but his replacement, Brett Kavanaugh, is likely even more ideological on antitrust matters. See Calkins (2018).

¹⁸⁹ Federico, Morton, and Shapiro (2019) (on file with authors).

trajectory, but the tools do not yet exist to accurately forecast the speed and direction of innovation in the longer run. Likewise, obtaining quantitative evidence about the innovations or products that *would* have been offered to consumers in the absence of the conduct is often not possible.

Perhaps in part as a result of these challenges, US antitrust has not been active in policing allegedly anticompetitive conduct by technology platforms. The government's last monopolization case involving issues raised by platforms was the Microsoft case in 1998. Today, the European Commission and the European National Competition Authorities effectively act as the global enforcers for allegedly exclusionary conduct that operates at a global level in industries such as software, chips, and digital platforms.¹⁹⁰

D. Could a Reformed Antitrust Law be Effective in Dealing with Digital Platforms?

There are many ways that US antitrust law could be revised to make it more aggressive in addressing competition problems while maintaining the objective of prohibiting private conduct that reduces economic welfare. Such changes could improve competition enforcement in digital markets as well as others.

1. Reform by What Means?

With few exceptions, antitrust law has in the past evolved in a common-law-like process by which it has reflected new learning and judicial and market experience. This process is continuing, at least to some extent, as antitrust law and enforcement have recognized, for example, previously unnoticed competition problems in labor markets and doctrine has evolved to incorporate new learning about competitive problems that can be created by most favored nation (MFN) and other vertical agreements. The challenges posed by the big technology platforms and the current populist political climate have, however, put the issue of antitrust reform before Congress in various legislative proposals. There are advantages and disadvantages to both common law evolution and new legislation.

Evolution by a common law-like process takes time. It took the Chicago School roughly 20 years to refocus antitrust law, from the early expressions of its perspective in 1950s and 1960s¹⁹¹ until the Supreme Court's seminal decision in *GTE Sylvania* in 1977¹⁹²—and the obstacles are probably greater now. While there had been ebbs and flows in antitrust enforcement in the mid-twentieth century, there was not a well-formed conceptual framework calling for aggressive enforcement that the Chicago School proponents had to overcome. The

¹⁹⁰ See, e.g., European Commission (2018a); European Commission (2009); European Commission (2017b).

¹⁹¹ See, e.g., Bowman (1957); Bork (1966); Bork and Bowman (1963): 138.

¹⁹² *Continental T.V., Inc. v. GTE Sylvania Inc.*, 433 U.S. 36 (1977).

structure-conduct-performance paradigm was widely accepted as an economic proposition, but it was not embedded in a rigorous normative and conceptual antitrust framework. By contrast, oversimplified Chicago School thinking has provided a widely accepted framework for antitrust analysis for more than thirty years. Perhaps more importantly, many federal judges, appointed by an increasingly ideological vetting process, are trained in and adherents of that framework. Many seem unaware of new economic research that calls into question many of the tenets of that framework and continue to cite outdated Chicago School publications of the 1970s and 1980s. And, while there has been a great deal of economic research and literature on which a new antitrust paradigm could be constructed, there is not a widely accepted, alternative paradigm that is comprehensible to and administrable by lawyers and judges. Even if such a paradigm were written tomorrow and rapidly became widely accepted, it would likely take years for that paradigm to be manifest in doctrinal changes and market outcomes.

New legislation could in principle be adopted and take effect much more rapidly. New legislation would not need to depart from the dual requirements of bad conduct and harm to competition in the market as a whole. Such legislation might, instead, implement a recalibration of the relative tolerance of antitrust law for the risk of over-enforcement and under-enforcement by prescribing rebuttable presumptions that would ease the high proof requirements currently imposed on antitrust plaintiffs and place on defendants a more rigorous burden of proving efficiencies. Some possible new presumptions and similar reforms are outlined below.

The risk, of course, is that new legislation will not be enacted by experts committed to sound, economically-focused antitrust. It will be designed by Congress in a politically charged environment subject to pressure from the very companies who stand to lose their market power if subject to increased antitrust oversight, or who benefit if their trading partners are subjected to excessive oversight.

There is more at stake than the risk of flawed legislation. Antitrust law has maintained legitimacy and widespread support for nearly 130 years in part because it applies to all forms of commercial activity and is not perceived as special interest legislation. In our view it is very important that antitrust law not have different rules aimed at different sectors—such as technology¹⁹³ or agriculture¹⁹⁴—that would differentiate industries and undermine political support for antitrust law in general. For this reason, the report outlines a number of useful digital platform interventions that can be undertaken by a sectoral regulator rather than falling to the task of antitrust enforcement.

Equally important, antitrust law has benefited immensely from the brevity and vagueness of the key statutory provisions because they have enabled antitrust law to evolve in response to

¹⁹³ As suggested, by, e.g., Warren (2019a).

¹⁹⁴ As suggested by, e.g., Warren (2019b).

new learning. The challenge with new legislation is to embrace enduring normative principles without codifying current economic learning in a way that will prevent the law from evolving to take account of newer economic findings.

2. Reform of Antitrust Law through a General Tightening

As noted, most antitrust cases require uncertain decisions about unknowable future events like innovation or entry or about unobservable economic elements like demand curves, marginal cost, and product quality. Much US antitrust law reflects judgments about how to deal with such uncertainty. Those judgments are embodied in rules regarding burdens of proof, evidentiary presumptions, and decision-theoretic approaches to fact finding. They are also embodied in substantive legal doctrine, such as rules regarding predatory pricing and unilateral refusals to deal that are intended to reduce the likelihood of false positives (erroneously finding a business has violated the law) even at the risk of false negatives (erroneously deciding that a business has not violated the law). Digital markets typically have high levels of uncertainty and move quickly.

Given uncertainty, courts must determine how much weight to put on the risk of enforcement mistakes: both the likelihood of a mistake and its cost. Much US antitrust law is driven by a judgment, embraced by the Chicago School, that avoiding false positives (good conduct judged to be bad) is more beneficial to society than avoiding false negatives (anticompetitive conduct judged to be good).¹⁹⁵ This judgment rests on the beliefs that false positives are difficult to correct but that false negatives will be quickly corrected by market forces. These beliefs seemed plausible in 1975 in a Chicago School framework, but they have never been empirically demonstrated and have fallen into disrepute. Moreover, their logic is incomplete because the cost to society of a false negative, for example, depends not only on its likelihood, but also on both the magnitude and the duration of the resulting harm.

It is time for antitrust law to recalibrate the balance it strikes between the risks of false positives and false negatives. Underenforcement is likely to be costlier than previously thought because, among other things, market power of large technology platforms is more enduring. False negatives are almost certainly more common than previously thought because certain types of conduct that were previously thought to be benign are now understood to be anticompetitive. Especially in technology markets, the most important competitive threats to incumbent firms are likely to come from new entrants that might be vulnerable to exclusionary conduct or anticompetitive acquisitions when their competitive prospects are uncertain. In addition, false positives might be less common than previously thought because of the development in the past few decades of more sophisticated and reliable econometric and simulation tools for assessing

¹⁹⁵ For a classic statement of this proposition, see Easterbrook (1984) (“There are limits on the ability of courts to sort the beneficial from the deleterious manifestations of [restrictive dealing] practices, and most of the time it is better not to try than to try and fail.”). For a criticism of this view, see Baker (2015).

conduct and market power effects. New work by economists studying multi-sided markets, network effects, the economics of nominally “free” goods and services, and restrictive vertical agreements suggests that antitrust agencies and courts should be able to continue to improve in their ability to make sound enforcement decisions. And false positives might be less costly than previously thought because firms are finding new and different ways to realize efficiencies.

A recalibration of this type should influence antitrust law in two basic ways. First, it should provide a basis for revising certain aspects of antitrust doctrine that were adopted explicitly in order to minimize the risk of over-enforcement. Second, it could more broadly provide a basis for courts to impose less demanding proof requirements on antitrust plaintiffs, especially where facts are difficult to observe or prove directly and indirect or circumstantial proof is available. Again, we pursue our theme of harm to entry through the next section. Exclusion of existing or potential entrants is well established in both the economics literature and the antitrust jurisprudence as a harm to competition, but the law is not well calibrated to recognize this familiar tactic in its new setting.

Some of the specific ways in which the law might be revised are set forth below. They could be achieved by common law-like evolution of antitrust law or by new legislation.

3. Specific Areas of Possible Antitrust Reform [designed for the specialist reader]

To address the issues raised by technology platforms, antitrust enforcement agencies and courts will need to understand the unusual factual context that those platforms often present and be receptive to recent and future economic learning about the implications of that context and how to apply antitrust principles to it. Economists and other experts will need to develop new understanding and new tools to aid agencies and courts in addressing these matters, among others:

- How to assess the quality-adjusted price paid for a good or service sold in a barter transaction with zero or close to zero monetary price, and how to define and analyze markets in which a substantial portion of the sale take the form of barter transactions.
- How learning from behavioral economics and related disciplines about addictive or exploitative content should be considered in addressing issues regarding economic welfare.
- How technology platforms are able to take advantage of consumer biases (such as salience, status quo bias, or impatience) to bind consumers to their platforms and make

switching to alternatives more difficult than imagined by lay intuition (“competition is one click away”).¹⁹⁶

- How market circumstances affect the likelihood and nature of innovation and how to evaluate innovation, whether using qualitative or quantitative tools, in both its magnitude and direction.
- How to assess potential competition from new or small firms or not-yet-identified future innovators and entrants. This is especially important in markets that depend on technological change and in which competition in the market is less important than leap-frog competition for the market.
- How to assess consumer welfare in a two-sided market. Learnings from platform economics suggests that users on different sides of a platform generally have divergent interests and thus that defining a single two-sided market—rather than two, closely interrelated ones—obscures the analysis. Platform economics also demonstrates that neither the change in the two sided price nor the change in the transaction volume is a sufficient statistic for a how a firm’s conduct affects consumer welfare.¹⁹⁷

Other antitrust reforms, including those below, would require changes to doctrine:

- Antitrust law prohibits unilateral refusals to deal only under very unusual circumstances. Current law reflects, among other things, concerns about the difficulty of determining the required terms of trade and incentive effects of required dealing on both the dominant firm and its rivals.¹⁹⁸ The law thus gives platforms substantial freedom both to refuse to deal with actual or potential rivals, including complements, and to deal with them only on onerous terms. Those terms might include access to customer data, interoperability, and other terms that raise the costs of rivals and enable the platforms to reinforce their

¹⁹⁶ Wismer (2012).

¹⁹⁷ One of the features of such platforms is that, because of the feedback effects between the two sides of a platform and the importance of the relative prices charged on the two sides, increased output by the platform does not necessarily imply increased economic welfare. Whether increased output by the platform will increase economic welfare depends on how the benefits and costs of that output are allocated across the two sides. *See* Katz, Michael L. (2019); Rochet and Tirole (2003). There is another sense in which increased output of digital platforms does not necessarily increase economic welfare. Because of some of the unique attributes of digital platforms discussed above, consumer demand for digital services does not necessarily reflect consumer welfare from those services. *See* Section II.1, *supra*. This latter concern applies to other products as well, such as mortgages and prescription drugs; and it is very relevant to possible regulation of digital platforms, as discussed below. By contrast, while antitrust enforcers and courts need to understand those attributes of digital platforms in order to understand how the relevant markets work, antitrust law is for several reasons based on the assumption that consumer preferences reflect consumer welfare.

¹⁹⁸ *See, e.g.,* Verizon Commc’ns Inc. v. Law Offices of Curtis V. Trinko, 540 U.S. 398 (2004); *see also* Melamed (2006).

dominant positions. This doctrine should be reconsidered in light of the substantial importance of large technology platforms and, in particular, their central role as distribution channels, both of which suggest that the benefits of antitrust intervention might be greater than previously appreciated.

- Predatory pricing law has been shaped in large part to avoid over-enforcement and with explicit acknowledgement that the law permits some forms of anticompetitive pricing conduct.¹⁹⁹ Courts have adopted a narrow and rigid notion of recoupment²⁰⁰ and have made it almost impossible to prove that prices are below cost, even where it seems likely.²⁰¹ Digital goods often have a marginal cost close to zero, which makes tests that require prices to be below incremental or variable cost almost impossible for a plaintiff to meet. The law has also been construed to protect only rivals that are equally efficient at the time of the conduct at issue and thus to disadvantage smaller rivals that have not yet reached efficient scale. Predatory pricing law should be modified so that it will be better able to combat anticompetitive pricing by digital platforms and other firms.
- The paradigm of predatory pricing law has also been relied upon to assess more complex pricing strategies, such as loyalty discounts. Loyalty discounts and similar contracts can be used to drive one side of a platform to single-home, which can cause a market to tip and enhance market power. The efficiency benefits of loyalty discounts are very different from those of low prices in general; for example, unlike low prices themselves, even above-cost loyalty discounts do not necessarily increase static welfare. Antitrust law should not rely exclusively upon predatory pricing standards to assess loyalty discounts.²⁰²
- The Supreme Court held in *American Express* that a plaintiff in a case involving a vertical restraint must define and prove a relevant market and may not rely on direct proof of harm to competition.²⁰³ This holding was based on the notion that vertical restraints almost always enhance efficiency and almost never harm competition. Scholars over the past 30 years have demonstrated that that notion is false and therefore, that vertical restraints must be evaluated individually on the specific facts.²⁰⁴ Where there is direct evidence of harm to competition, antitrust law should not require circumstantial evidence via a defined relevant market.

¹⁹⁹ See *Barry Wright Corp. v. ITT Grinnell Corp.*, 724 F.2d 227, 231, 234 (1st Cir. 1983) (Breyer, J.).

²⁰⁰ Hemphill and Weiser (2018).

²⁰¹ See, e.g., *United States v. AMR Corp.*, 335 F.3d 1109 (10th Cir. 2003).

²⁰² Hemphill and Weiser (2018).

²⁰³ *Ohio v. American Express Co.*, 138 S. Ct. 2274 (2018).

²⁰⁴ Salop (2018b); Salop (2018a).

- Courts are generally very reluctant to second guess a defendant's product design decisions, even where the design harms competitors.²⁰⁵ Yet product design decisions involving, among other things, app stores, mobile device screen layouts, data storage and analysis, and interface design are often key elements in digital market competition. Antitrust courts should be more willing to assess product design decisions where appropriate.
- There should be no safe harbor based on the short-term of exclusive dealing agreements and other restrictive vertical agreements when they are used by dominant platforms or firms. Such firms can use their market power to induce desired behavior in trading partners without relying on long-term contracts.²⁰⁶

Perhaps most importantly, antitrust law might be revised to relax the proof requirements imposed upon antitrust plaintiffs in appropriate cases or to reverse burdens of proof. Burdens of proof might be switched by adopting rules that will presume anticompetitive harm on the basis of preliminary showings by antitrust plaintiffs and shift a burden of exculpation to the defendant or by ensuring that plaintiffs are not required to prove matters to which the defendants have greater knowledge and better access to relevant information. These proof requirements include the following, which are likely to be important in the application of antitrust standards to technology platforms:

- Mergers between dominant firms and substantial competitors or uniquely likely future competitors should be presumed to be unlawful, subject to rebuttal by defendants. This presumption would be valuable, not because it would identify anticompetitive mergers with precision, but because it would shift the burden to the party with the best access to relevant information on issues of competitive effects and efficiencies from the merger.
- Courts should not presume efficiencies from vertical transactions. Crediting of efficiencies should require strong supporting evidence showing merger-specificity and verifiability.²⁰⁷

²⁰⁵ See *United States v. Microsoft Corp.*, 253 F.3d 34 (D.C. Cir. 2001) (per curiam).

²⁰⁶ Compare *United States v. Dentsply Int'l, Inc.*, 399 F.3d 181 (3d Cir. 2005), with *Omega Env'tl., Inc. v. Gilbarco, Inc.*, 127 F.3d 1157 (9th Cir. 1997).

²⁰⁷ For an example of the Court presuming efficiencies from vertical transactions with minimal evidence, see *American Express*, 138 S. Ct. at 2289.

- Courts should be more willing to permit plaintiffs to prove harm to competition by circumstantial evidence, especially where the propositions in question are not observable and there thus cannot be direct evidence.²⁰⁸

4. A Competition Court

Revisions to the law may have little effect to the extent that judges see antitrust cases only rarely and have difficulty understanding the economic underpinning of antitrust law. One way to ameliorate this problem would be to establish a specialized antitrust court on which a certain number of Article III judges would sit for a specified term of several years. These judges could be expected to hear multiple antitrust cases and to develop substantial antitrust expertise. And, because they would be chosen from the general federal bench, they would bring the broader perspective of generalist judges. This model could be used at the trial court level, the appellate level, or both. If judges are selected from the larger pool of Article III judges and rotated on and off the specialized court, judges on the specialized court will be less likely to develop an overly narrow intellectual interpretation of antitrust law.

5. A Regulatory Partner Could Enhance Effective Antitrust Enforcement

Digital markets move quickly. Using the internet as a distribution channel to reach an installed base of billions around the globe means that new products can be rolled out quickly—as can anticompetitive conduct. The pace of antitrust enforcement is far slower—judging by the evidence, too slow—to protect small entrants except by creating deterrence. A sectoral regulator, by contrast, could be endowed with the authority to move quickly. By taking steps to preserve competition before markets have tipped or entrants have been purchased, a regulator could prove a valuable complement to antitrust enforcement.

Antitrust enforcement proceeds on a case-by-case basis, and, apart from merger enforcement, it is largely backward looking—it looks at conduct already undertaken that is alleged to harm competition. A regulator, by contrast, can look forward and establish rules to constrain future conduct before there has been harm to competition, or before investments by the defendant and third parties have been shaped by the conduct. This is particularly useful in cases where ex post conduct remedies would be costly or ineffective.

Effective antitrust enforcement requires effective remedies. Treble damages and financial penalties can compensate for past harms and deter future bad conduct, but they do not restore competition to markets in which competition has been harmed. Even an injunction to forbear from the same or similar anticompetitive conduct going forward will not restore the lost competition if entry barriers are high. For example, if the market has tipped and network

²⁰⁸ See, e.g., *id.* (injury to competition in a two-sided market); *United States v. AMR Corp.*, 335 F.3d 1109 (10th Cir. 2003) (predatory pricing).

externalities are very strong, the firm that became a monopolist through violations of the antitrust laws could stop the conduct at issue and yet retain its monopoly position and the associated stream of profits. An antitrust authority that wants to restore lost competition must induce entry and/or impose structural remedies. Such remedies are likely to require detailed and often technical monitoring and years of effort. Antitrust enforcers are not suited to that type of oversight whereas a sectoral regulator is. One possible solution is to permit antitrust authorities and courts to design antitrust remedies and rely on an expert sectoral regulator to oversee their implementation. A similar approach was approved by the Supreme Court in the *Otter Tail* case.²⁰⁹

E. Regulation

For the reasons above, we believe the establishment of a sectoral regulator should be seriously considered. Given the tasks detailed by the other committees in this project as well as this one, there would be much for a regulator to do that would improve the impact of digital platforms on society. A digital regulator, should Congress choose to create one, could consider regulations and actions along the following lines. We call this potential regulator the Digital Authority (DA).

The general harm identified above is insufficient entry, and therefore insufficient competition, in digital platforms. Many of the regulations below are designed to lower barriers to entry directly. Others are responsive to the difficulty of effectively prosecuting antitrust cases today in digital markets. Regulations that mimic the antitrust laws but lower the burden of proof for the regulator and allow it to move faster are a way to gain effective enforcement in this sector, if not others.

Regulation offers a valuable addition to antitrust enforcement. It can help design the digital landscape and align the interests and incentives of platforms and key providers with those of consumers and society. When carefully designed, a regulatory regime can limit or even preempt the harmful effects detailed above, while minimizing its impact on the dynamic nature of digital markets. It can offer clarity and legal certainty as to the boundaries of acceptable competition. Moreover, some of the problems discussed above may have only one structural solution: breakup of the platform. An enforcer might not want to choose that option because it is very disruptive. But less disruptive ex post remedies require ongoing monitoring, which antitrust enforcers are not well-positioned to do. Handing that job off to a regulator might better serve consumers.

Therefore, the committee suggests considering regulatory steps in conjunction with improved antitrust enforcement. It is important that regulation support, and indeed enhance,

²⁰⁹ *Otter Tail Power Co. v. United States*, 410 U.S. 366 (1973).

competition. The report focuses here on regulations targeted at improving competition, not other policy challenges related to digital platforms. And it is critical that any new regulatory framework avoid the well-documented pitfalls that regulation often invites (e.g., agency capture, revolving door syndrome, or incumbent protection). By having a pro-competition mandate, our hope is that the DA will be able to use the tools at its disposal in ways that resist the natural impediments to entry identified in Part I, rather than entrench powerful incumbents.

This combination of enforcement tools applied to one industry is nothing new in the US economy. Virtually every sector of our economy has required both antitrust and specific regulatory oversight, from banking to agriculture to communications, in order to promote competition and other public interest goals. The regulations we propose below are focused on, and limited to, those that will enhance competition.

The communications sector may offer the best guidance for how to approach public accountability for digital platforms. Telephone, cable, and wireless communications networks have many of the same attributes as digital platforms. Large capital expenditures to build networks, declining costs as consumers use the networks and buy more services on the networks, and the difficulties for new players to enter these markets and compete have led to a world of one or few players in each market. And just like social networking, online search, and web-delivered content, communications networks have been the lifeblood of how we communicate and practice our democracy by delivering TV, radio, text, and conversation, making competitive outcomes in the sector critical. The FCC has served as the sector-specific regulator for telecommunications. In complex industries, a sector-specific regulator can have a wider remit than an antitrust authority as is detailed in the box below.

Box VI The FCC Model

In the Telecommunications Act of 1996, Congress sought to affirmatively promote, not just protect, competition.¹ Congress' vision included service availability that was universal and affordable to all and a commitment to local and diverse ownership of news sources to support a robust marketplace of ideas.² It wanted to prevent undue discrimination, limit ownership to prevent excess market power (both horizontally and vertically in some instances), and police against abusive contract provisions that distort fair market practices.³ Some of these goals were achieved with more success than others. And in general, the regulatory process inherently restricts behavior and therefore efficiencies. The net benefits of regulation should factor in these lost efficiencies.

The phone number portability rule was a clearly pro-competitive regulation. The FCC published the Wireless Local Number Portability rule in 2003.⁴ The rule allows cell phone users to keep their phone number when they switch between wireless carriers. Prior to the rule, a consumer who wanted to take advantage of a low price or better quality with a competing provider would have had to change her phone number. This significant switching cost dampened competition between carriers. The number portability rule made switching easier and thereby strengthened competition between carriers.⁵ Similarly, when portability of 1-800 numbers (e.g., 1-800-flowers) was introduced, prices of those services fell.⁶

After Congress found the cable companies to be local monopolies in 1992, Congress identified a dearth in cable competition partially caused by vertically-integrated cable operators' refusal to sell their programming to potential competitors. It therefore temporarily prohibited exclusive programming contracts, banned a variety of abusive contracting practices, and required that vertically-integrated firms sell their content to competitors under reasonable prices, terms, and conditions.⁷

1 H.R. Rep. No. 104-458 (1996) (Conf. Rep.). Even before the 1996 Act, the FCC had already begun the work of promoting competition through the broad authority it had been granted in earlier statutes.

2 47 U.S.C. § 151 (2018).

3 *Id.*; 47 U.S.C. § 160 (2018).

4 See generally *Wireless Local Number Portability (WLNP)*, FED. TRADE COMM'N, <https://www.fcc.gov/general/wireless-local-number-portability-wlnp>.

5 Press Release, Federal Communications Commission, FCC Observes First Anniversary of Wireless Local Number Portability (Nov. 24, 2004), <https://www.fcc.gov/document/fcc-observes-first-anniversary-wireless-local-number-portability> (“Wireless local number portability (LNP) eliminated a barrier to full competition in mobile telephone services and between landline and wireless services”).

6 V. Brian Viard, *Do Switching Costs Make Markets More or Less Competitive? The Case of 800-Number Portability* (Cheung Kong Graduate School of Business Research Paper No. 1773R2, 2004), <https://ssrn.com/abstract=371921>.

7 47 U.S.C. § 251 (2018).

The FCC Model (Cont.)

The resulting increased competition in the cable industry enabled the Direct Broadcasting Satellite industry to grow,¹ opened the door for telephone companies and smaller cable providers to compete in the provision of distribution,² and ultimately drove the digital video market to develop today's cable-modem-powered broadband internet services.

The 1996 Telecommunications Act included interconnection requirements between competing carriers to expand competition. The Act outlined a regulatory regime of duties to connect, of parity in quality between connections offered to the incumbent's own affiliates and competitors, and of rates and contract terms that were just, reasonable, and nondiscriminatory.³ The regulation was designed to protect all of the businesses that needed to connect to the long-distance wire for their business but competed with a vertically-integrated local wire. Not only was raising rivals' costs prohibited, but complete foreclosure was also prohibited. This duty to deal in a non-discriminatory way is an example of a policy that arguably promoted entry of cable, local telephone, and long-distance competitors.

Similarly, the FCC developed Customer Proprietary Network Information (CPNI) rules, designed to protect sensitive business data transmitted through public telecommunications networks. The FCC recognized that a dominant platform, such as the telephone company, on which businesses rely in order to reach their customers, could easily gather and take advantage of sensitive business user data to promote the phone company's business and harm competition.⁴ Phone service competitors, as well as data-driven businesses (such as home-security monitoring firms, hotels, and airlines), rely upon these rules to grow their businesses without interference from telecom network owners. The rule was so effective at promoting competition that Congress codified it in the 1996 Telecommunications Act. The harvesting of data flowing through a platform—generated by and belonging to others—is a standard business practice among today's digital platforms.

Of course, these past regulatory efforts have also helped us learn what not to do. Regulations could be and have been used to entrench incumbent firms' market power, erecting regulatory barriers to entry for new or innovative competitors.

1 47 U.S.C. § 548(a) (2018).

2 47 U.S.C. § 521(6) (2018); 47 U.S.C. § 548 (2018).

3 47 U.S.C. § 251 (2018).

4 47 U.S.C. § 222 (2018).

The FCC Model (Cont.)

Regulatory capture is a common problem.¹ For example, AT&T has historically had a symbiotic relationship with the US government.² Perhaps the height of this relationship was the explicit Kingsbury Commitment in 1913, which allowed AT&T, rather than the government, to set the “solution” for managing competition.³ The current Chairman of the FCC, Ajit Pai, has called this “a cautionary tale about the dangers of regulatory capture.”⁴ It was likely this close relationship that allowed the AT&T monopoly to persist for so long, eventually requiring an antitrust case to address the problem.

¹ See George J. Stigler, *The Theory of Economic Regulation*, 2 BELL J. ECON. REG. 3 (1971) (arguing that, “as a rule, regulation is acquired by the industry and is designed and operated primarily for its benefit.”); see also Sheldon Whitehouse, *Legal Responses to Regulatory Capture*, ACS BLOG (Nov. 9, 2018), <https://www.acslaw.org/acsblog/blog-post-for-the-american-constitution-society-legal-responses-to-regulatory-capture> (“We unfortunately live in a time of widespread regulatory capture. We should not pretend otherwise.”); Elizabeth Warren, *Corporate Capture of the Rulemaking Process*, REG. REV. (June 14, 2016), <https://www.theregreview.org/2016/06/14/warren-corporate-capture-of-the-rulemaking-process> (“Under the law, it is easy for business groups to challenge a rule for being too strong or too restrictive. But it is much harder for public interest groups or ordinary citizens to challenge a rule for being too weak or riddled with loopholes.”); Daniel Carpenter, *Challenges in Measuring Regulatory Capture*, REG. REV. (June 22, 2016), <https://www.theregreview.org/2016/06/22/carpenter-challenges-measuring-regulatory-capture> (“[T]here are limitations when it comes to preventing capture. Perhaps the most important is that capture is difficult to measure”).

² Chris Matthews, *AT&T and the Government Have Been ‘Friends’ for a Really Long Time*, FORTUNE (Aug. 18, 2015), <http://fortune.com/2015/08/18/att-nsa>.

³ Ajit Pai, Chairman, Fed. Comm’n Comm’n, Remarks on the 100th Anniversary of the Kingsbury Commitment (Dec. 19, 2013), <https://www.fcc.gov/document/pai-remarks-100th-anniversary-kingsbury-commitment>.

⁴ *Id.*

1. The Digital Authority

To be effective, a proposed regulatory regime requires an enforcement body capable of carefully designing and enforcing the relevant regulations. We start therefore with a proposal for Congress to pass legislation creating a Digital Authority with the mandate to develop targeted regulation to achieve the goals described above and subsequently engage in monitoring and enforcement.

We anticipate that this regulator will also be tasked with non-competition digital goals, such as those in the areas of privacy, media, data-use restrictions, and consumer protection. While the antitrust agencies will employ structural interventions to protect competitive markets wherever possible, the focus of this regulator will be on both carrying out remedies for the

antitrust authority that require ongoing oversight, and on developing regulations going forward that are a combination of structural safeguards, such as unbundling or separation, with limited behavioral interventions in areas where traditional antitrust tools are insufficient. Other jurisdictions that are assessing competition in digital platforms all propose some form of regulation.²¹⁰ Having forward-looking regulations in place will increase business certainty about what conduct is permitted and how enforcement actions are likely to proceed. Ideally, this predictability and clarity will encourage companies to comply with the law, thus requiring fewer government resources for enforcement.

The DA legislation will require Congress to define the scope of regulatory power. The definition must include digital businesses that facilitate transactions of any kind (including the sale of advertising). It should have clear and broad authority over digital business models in order to prevent firms subject to regulation from evading its oversight.

We offer a menu of potential regulation, starting with the least intrusive form and building to the more severe interventions. Each type of regulation comes with costs and benefits which we briefly detail. The size of the costs and benefits will be different across settings, meaning that the choice of the best tool for any particular case may vary. The committee also suggests separating out some types of regulation that will apply to virtually all market participants while others are only appropriate tools to apply to companies with bottleneck power.

1) Bottleneck power

“Bottleneck power” describes a situation where consumers primarily single-home and rely upon a single service provider (a “bottleneck”), which makes obtaining access to those consumers for the relevant activity by other service providers prohibitively costly. As the U.K. Furman Report put it,

[O]ne, or in some cases two firms in certain digital markets have a high degree of control and influence over the relationship between buyers and sellers, or over access by advertisers to potential buyers. As these markets are frequently important routes to market, or gateways for other firms, such bottlenecks are then able to act as a gatekeeper between businesses and their prospective customers.²¹¹

The finding of bottleneck power will employ consideration of the forces that, as discussed above in this Report, tend to impede entry and lead to foreclosure. The Furman Report similarly explains that this single-homing foreclosure tends to happen when users experience high switching costs, such as loss of valued personal data or reputational indicators at the point

²¹⁰ See, e.g., Australian Competition (2018): 13-14; Crémer (2019): 8-10; Furman (2019): 60-61.

²¹¹ Furman (2019): 41.

of switching; contract terms that deter switching; technical barriers to switching, such as complex switching processes or a lack of interoperability between the old service and the new or second service; tying services, which can be by contract or technical; and the inertia of defaults.²¹² Digital businesses that have this incentive and ability to develop and preserve a single-homing environment should be considered entities with bottleneck power. The DA, pursuant to congressional guidance, should have the sole authority to define bottleneck power and should update the definition either regularly or on an “as needed” basis. This is not a determination that should be left for definition by generalist judges. If Congress establishes a specialized panel of judges who review antitrust matters (as described above), it would be wise to grant the same panel sole judicial review authority over DA regulatory actions.

2) *Data*

Should Congress pass any laws concerning consumer data, the authority could set forth pro-competitive rules concerning the consequences of consumers’ control over their data and about user choice in the sharing of data. Where users are simply and clearly informed and given the opportunity to make viable choices about which companies get their data, this valuable competitive information can be shared as the consumer prefers. The DA could design data sharing rules with the general goal of reducing single-homing and promoting entry.

3) *Partnership with the Antitrust Agencies*

Antitrust enforcement agencies or courts could designate the DA as the administrator and/or architect of remedies in antitrust cases. In cases where structural remedies are not appropriate to restore the lost competition, the DA could be directed to either carry out a remedy designed by the competition authority, or provide an effective behavioral monitoring remedy consistent with the DA’s mandate. Since the difficulty of designing an appropriate and administrable antitrust remedy that can be enforced by a court often limits antitrust enforcement, having the option of an expert regulator to craft the remedy, or to simply be responsible for ongoing monitoring or other execution of the remedy, could better ensure cost-effective enforcement of pro-competition policies.

The DA can partner with the antitrust agency when it observes anticompetitive conduct. Due to their frequent interaction with firms in the industry and their real-time data feed, regulators may observe exclusionary conduct before the antitrust authority, and be able to give it relevant information and data.

²¹² Furman (2019): 36.

2. Menu of Regulations

We lay out below a menu of regulations that could be used to solve the problems identified above from least interventionist to most interventionist. For each, we describe the likely benefits and the costs.

1) Broadly Applicable Regulations

Some of these regulatory tools should be applied broadly to all firms in the industry, not only to bottleneck firms. It would be appropriate, however, to include a small business exception and perhaps even a new business exception, to allow very small entrants, who may benefit competition, time to ramp up against larger established companies.

a. Data collection

The Digital Authority could regularly collect data on market transactions, with an emphasis on data from businesses with bottleneck power. The information would allow policy makers and researchers to assess the performance of the sector and improve rule enforcement. For example, the data may include, inter alia, a sample of searches at a set of websites, a sample of queries followed by purchases at other sites, a sample of downloads of applications at an app store, or a sample of activity and ads shown on a social media site. The DA may further regularly collect information, including what types of data the business collects, how it uses that data, and who else is bartered or sold access to which elements of the data. It may include experiments to evaluate the true portability of data, the ads generated by certain user information, and other information relevant to assessing how and if competition is taking hold in the market. This program will require ongoing costs to run, but these should be reasonable compared to the benefits.

Box VII**Real Time Regulation in Financial Services**

Technological innovation and the data revolution are also disrupting the financial services market, spurring regulators into action. Banks and other companies are undergoing an innovation and digitization spree, partially in response to competition by non-financial institutions and startups. This rapid change increases the challenges for regulators to follow and analyze what regulated companies are doing.

The response by some financial regulators has been to embrace new technologies focused on machine readable regulation (with the help of machine learning and AI) and digital reporting (with the help of data sharing frameworks, APIs, cloud computing, and advanced analytics). This is leading to an important shift from manual and template driven compliance and regulatory reporting to a near real time or real time regulatory reporting and supervision. Examples of authorities adopting this system range from the European Central Bank, various national central banks in continental Europe, the Monetary Authority of Singapore, or more notably, the Bank of England and the Financial Conduct Authority (FCA) in the UK.

The Bank of England and the FCA are arguably the most advanced when it comes to digital regulatory reporting. Both were the first financial regulators to set up innovation labs and programs to interact with startups and new technologies. Their aim is to assess which regulation is suitable for machine reading, what data could be included in a standardized fashion within a digital regulatory reporting framework and how these changes will impact data models currently used for regulatory stress testing, leading to upgrades. This standardization then allows for near real time data sharing, lowering compliance costs to small and big firms alike.

Closer to home, in the United States, the SEC is implementing rules requiring trading firms to report intra-day trading data in real time. The SEC, the CFTC and the FDIC have also each set up fintech/regtech innovation labs to assess fintech solutions in general and regulatory tech solutions in particular. A Digital Authority could learn from, and expand on, these experiences of near real time regulation when overseeing an industry that has data at its core.

These data may be requested by an antitrust agency to aid in an investigation or prosecution. This ability to transfer existing data will speed up the enforcement of antitrust laws. After an investigation is opened, rather than subpoenaing the relevant data from the targeted firms and spending months arguing over definitions and formats, the antitrust authority could immediately begin analysis with the DA dataset.

In addition, the DA should make as much of these data public as possible, subject to keeping personally identifiable information and business secrets confidential. This will allow academics and nonprofits to study particular markets and consumer behavior as they do in

industries such as airlines (where “Data Bank 1A” is publicly available)²¹³ to great public benefit. Leveraging the nonprofit and educational sector to help the DA and elected officials understand these markets is likely to be helpful and cost-effective.

b. Restrictions on Practices That Enhance Behavioral “Mistakes”

As described above, it is well known that behavioral “nudges” can lead consumers to make better choices. Better choices promote competition because they generate increased market share for firms that make better offers to consumers. The DA should have a mandate to create such “light touch” rules when they will make markets more competitive. For example, automatic renewals can discourage consumers from comparison shopping when a contract ends. A restriction on automatic renewals in that setting could lower prices market-wide. Regulations that require firms to make salient the most important terms of an offer can improve competition (e.g., a credit card offer must show the APR in large font). The U.K. likely has the most advanced regulatory regime in this regard, and the DA may be able to learn from it and other jurisdictions that have already taken these steps.²¹⁴

c. Data Portability and Mobility

Congress may, at some point, pass a data law of some type that gives consumers control of their data. For this control to translate into more competitive markets, it must be used to lower switching costs and facilitate entry. The DA has a role to ensure that users can easily transfer their data from one service to another in industries where there is a common business model (e.g., social media, banking, or online grocery shopping). The DA will identify industries where porting is likely to aid the competitive process. Being able to port one’s data directly lowers the cost of moving from one service to another, which in turn causes businesses to compete harder to keep those customers.²¹⁵ Consumers who control their data and have the right to receive it in a standardized format from the business will be able to take advantage of a new entrant by porting themselves to it, along with their own data. With information about a user’s past purchases, likes, friends, and so forth, the entrant can provide a higher-quality service and grow more quickly.

The DA could propose a standard for exchanging the data, but remain open to options that industry favors, provided the format is not itself an entry barrier. The data porting standard should be updated frequently to accommodate new innovations in the industry. New innovators can think broadly about what services users might like that rely on this data, or are compensated

²¹³ Bureau of Transportation Statistics (2017).

²¹⁴ Fletcher (2019).

²¹⁵ This is why incumbents may create or maintain systems that make data portability difficult. *See* Section I.1.B.1, *supra*.

through access to the data.²¹⁶ The Vestager Report divides data into personal and not personal; and content into volunteered, observed, and inferred.²¹⁷ The DA could determine which of these types of data must be included in the portability standard.

While a porting regulation lowers consumer switching costs greatly, they may still be high enough that demand is not sufficiently contestable to induce entry.²¹⁸ The DA could also set up a process by which a customer can choose to send her data to an entrant by authorizing it to be transferred directly from her former service provider. The DA would need to authorize the entrant to offer this facility to its consumers and establish regulations to require the incumbent to transfer the consumer's data upon the authorized request from the entrant. This may be particularly useful as the Internet of Things becomes more important; a consumer may wish to port the food supply service that was bundled with her new refrigerator from Amazon to an entrant. Being able to authorize the entrant to obtain all her data from Amazon will lower the cost of switching. Conceptually, automatic porting is no different from manual porting, but it is mechanized in a way that is likely to raise contestability and therefore can make entry more profitable. The Vestager Report notes that the GDPR Article 20 provides these data portability rights to Europeans. If consumers have the right to quickly and easily patronize an entrant without data lock-in, there will be more incentive to enter into these markets. However, other entry barriers remain, which we discuss below.

d. Open Standards to Promote Competition

The DA should move preemptively to prevent the consolidation of control over users' identities, as this would create a large new source of market power. The DA could create an open standard so that new entrants can easily offer their own digital identity product that allows users to access goods and services online. One example of this type of product is Solid, by Tim Berners-Lee, often named as the creator of the World Wide Web. Solid offers users a "POD" that safeguards their digital identity that they can use to connect with different services.²¹⁹ Several government-backed efforts at identity portability are underway across the world. These include Estonia's e-Estonia initiative to give citizens a unique digital identifier; India's Aadhaar, a verifiable 12-digit identity number issued for each citizen which serves as an identifier and authenticator for a variety of offline and online services; Sweden's and Norway's BankId, which allows companies, banks, and governmental agencies to identify and conclude agreements with individuals over the internet; and even self-sovereign identity solutions studied by start-ups that

²¹⁶ See, e.g., the RadicalxChange group. "Mission and Values" (2019). The group's chair, Glen Weyl, has argued in a paper written with colleagues that it's a mistake to view internet services priced at \$0 as simply free; instead, data can be viewed as a form of labor or barter. See Arrieta-Ibarra et al. (2018) (exploring "whether and how treating the market for data like a labor market could serve as a radical market that is practical in the near term").

²¹⁷ See Crémer (2019): 8; *id.* at 24.

²¹⁸ See Furman (2019): 129, § 5.11.

²¹⁹ "How It Works" (2019).

would use blockchain to allow individuals to own their identity credentials and control who can access their data in online services. If an individual could then port their identity to the platforms and providers they wish to use, this would again promote entry of new services and erode the switching costs of established platforms.

The DA could consider creating an open standard that would facilitate micro-payments among consumers and digital entities. The coordination needed among stakeholders to create a successful micro-payment system is substantial, and it likely will require assistance and oversight from a regulator.

e. Merger Review

The behavior that may be of greatest concern to the many policymakers studying powerful digital businesses is their acquisition of potential competitors. These acquisitions often fall below the value threshold under which the buyer would need to notify competition authorities in advance of the deal. As a consequence, authorities have limited or no ability to assess whether a given deal is procompetitive or harmful to competition before it closes.²²⁰ Markets move quickly and a competitor's window of opportunity to gain traction against the incumbent is short. For these reasons, Congress could give the DA merger review authority. Similar to the FCC's merger review role, this would be conducted concurrently with the antitrust review done by the FTC or DOJ, but with different standards and tools. It would not be prudent to alter the nation's antitrust laws to accommodate one difficult and fast-moving sector where false negatives are particularly costly. Therefore, giving additional power over merger review to the sectoral regulator is a good solution.

These specific merger regulations should require merging firms to demonstrate that the combination will affirmatively promote competition. This shifting of the burden of proof from the government (to prove harm) to the parties (to prove benefit) will assist the DA by placing the job of demonstrating efficiencies on the parties, who have a greater ability to know what they are. In some cases, the DA's review may be the only merger review conducted, as it should not be subject to the minimum size limitations on HSR filings. In particular, notification and pre-clearance could be required for any acquisition by a business designated as having bottleneck power.²²¹

In its merger review process, the DA could be explicitly tasked with evaluating a given merger's likely harm to existing as well as potential competition. Another example discussed

²²⁰ See Section III.2.B, *supra*, for a further discussion of these difficulties.

²²¹ When network effects are strong, a digital business with bottleneck power will likely only have very small competitors. Therefore, even small transactions can neutralize an important potential competitor that is poised to grow. See Section III.3.B, *supra*.

above is the case of content or complements that could expand from that position to compete with the digital bottleneck business itself. As already noted, entry from elsewhere in the vertical (or conglomerate) chain may be the most effective and promising entry point to challenge an established bottleneck business. Mergers with either of these types of entrants have the effect of neutralizing companies that might one day have posed a competitive challenge to the bottleneck business. This view of potential competition should drive DA merger review.

The decision in *Credit Suisse* and the dicta, or language, in *Trinko* greatly expanded the industries and conduct that have become, for practical purposes, exempt from antitrust scrutiny.²²² The agencies have understandably been skittish to expend limited resources bringing cases that risk being thrown out on *Trinko* grounds.²²³ It is important that *Trinko* not be used to create a no-man's land where neither regulation nor antitrust are applied to harmful behavior. A traditional antitrust savings clause can no longer be relied upon, as the clause in the 1996 Telecom Act was found insufficient to protect antitrust enforcement in *Trinko*. Legislation creating regulation and antitrust enforcement for digital businesses should address this concern head on. The statute must be extremely specific, explaining for each tool and goal whether it is intended to supersede antitrust or not. Antitrust enforcers and other agencies can share dual authority with different review standards and goals. They can account for each other's determinations in a manner that will minimize inconsistencies without having one always take priority over the other. Antitrust must remain in full force except where Congress explicitly says otherwise.

The merger review process must move rapidly. The agency will need a simple and efficient merger review process so that businesses can move forward without undue delay, and the agency does not expend more resources than necessary. These concerns indicate that the burden of proof must primarily be placed on the merging parties who have the incentive, data, and resources to quickly deliver the right information to the authority. Decisions should be subject to judicial review, like a rulemaking process under the Administrative Procedure Act,

²²² See *Credit Suisse Securities (USA) LLC v. Billing*, 551 U.S. 264 (2007); *Verizon Communications v. Law Offices of Curtis V. Trinko, LLP*, 540 U.S. 398 (2004).

²²³ Howard Shelanski, at the time the FTC's top antitrust economist, testified in Congress on behalf of the Commission a few years after the *Credit Suisse* decision. He argued that a narrow interpretation of *Trinko* was possible. See Federal Trade Commission (2010) (statement of Howard Shelanski, Deputy Director for Antitrust in the Bureau of Economics at the FTC). The key facts in *Trinko* were that the legislation at issue, the 1996 Telecom Act, went farther than antitrust law; an agency, the FCC, had issued rules directly regulating the conduct at issue; and the FCC actively administered those rules. See *Trinko*, 540 U.S. 398. Shelanski argued that "[w]here a competent agency actively administers a rule whose standard for the competitive conduct at issue in litigation is more demanding on the defendant than antitrust law, the Court was right to find it relevant whether the marginal gains outweigh the potential costs of antitrust enforcement against the same conduct." Federal Trade Commission (2010): 9. Yet he expressed concern that courts may use much broader interpretations of the line of cases. The Court in *Trinko* expressed concern about misuse of antitrust law by impudent plaintiffs, so some preemption could be limited to private plaintiffs, with expert agencies being given greater leeway.

with reliance on the specialist Competition Court discussed above. However, in light of the difficulty of winning antitrust cases in current US courts, judicial review will be effective at supporting the mission of the DA when combined with a more explicit and tougher enforcement statute or a Competition Court, or both.

In order for the agency to be most efficient, it should use a combination of adjudication and rulemaking for these merger transaction reviews. Rules help businesses know what behavior is proscribed, and help the agency clarify its intentions without having to wait for a good test case—or multiple test cases—to fully explain the issue. Clear rules will help companies avoid proposing anticompetitive transactions to begin with.

The cost of such regulation is duplicative merger reviews. However, given the importance of markets that tip and the uncertainty about whether an acquisition is a substitute or complement, having two reviews is perhaps helpful in getting to the right answer. If the regulator can block an incumbent platform's acquisition of potential competitors, and those competitors have the ability to quickly get their customers' data, entrants will plan to enter and compete with the platform, rather than enter and be bought. If the competitive environment is improved by data portability and the elimination of exclusionary conduct, the chance of success will rise and consumers will benefit from increased competition.

f. Interoperability

There may be settings where the DA is charged by Congress to prevent the creation of market power in the first place due to the importance of the market and the potential harms from the market power. Such a market would be a candidate for open interoperability standards that could be used by all competitors and would promote entry. The DA could oversee the creation of an open standard and its adoption in the market of interest. For example, devices in the home might be required to adhere to an open standard so that any platform could connect with any device. Without such a standard, an Apple thermostat would not function in conjunction with a Google home assistant. Under an open standard, if a consumer had purchased a number of Amazon appliances and then switched her home assistant from Amazon to Apple, she would not need to purchase a new refrigerator, thermostat, and security system because all of those devices could connect to the new platform and stay functional. It is possible that such open standards can slow down innovation that depends on the interface, but open standards will drastically reduce lock-in and market power, leading to greater incentive to innovate on the service itself.

2) *Regulations that apply to firms with bottleneck power*

Some regulations should apply only to firms that meet the DA's definition for bottleneck power. In these cases, the bottleneck firm has the incentive and ability to harm competition and is an important source of consumer welfare. These firms require extra monitoring to be sure they

are not violating antitrust, or other laws, because of the uncertainties in technology and demand, the speed at which platforms tip, the irreversibility of tipping, and the need for expert evaluation of the design of algorithms. For all of these reasons, the cost of false negatives is high and therefore, under conditions of uncertainty, the public interest requires the DA to take a more interventionist approach.²²⁴

a. Mergers

The DA could be given merger review authority over all transactions involving digital businesses with bottleneck power because new competition against these entities is the most valuable for consumers. Businesses with bottleneck power would notify the DA and obtain pre-clearance for an acquisition of any size. While this would potentially result in many reviews for the DA, their number will be limited by the number of platforms with bottleneck power.

The DA may want to use its merger review authority over bottleneck firms to assess consummated mergers, just as the antitrust authority can under section 7 of the Clayton Act. Past mergers of potential competitors that created monopoly positions could be assessed in this way. Similarly, past vertical mergers (e.g. a business that develops tools for brands to place ads and a business that runs mechanisms to set ad prices) may be found to cause higher prices or otherwise lessen competition. Any such anticompetitive mergers could be unwound by the DA.

b. Non-discrimination and foreclosure

Discrimination against current or future rivals is an important tool in a foreclosure strategy by a digital business with bottleneck market power.²²⁵ As explained above, there are broadly two types of foreclosure of a complement: one that operates against a complement that is a potential competitor of the platform itself, and one that operates only on the platform between rival providers of content. Because large digital platforms today are conglomerates with large eco-systems of complements, this is an important area in which to enforce competition. Non-discrimination can be a helpful tool in creating a competitive environment in which entrants are protected and can thrive, while allowing a platform to vertically integrate to some degree. Non-discrimination requirements should be used only after careful study because they can also prevent efficient forms of service that enhance competition.

Anticompetitive foreclosure of a firm that is a current or potential competitor can be addressed by the antitrust authority if the antitrust law is strong enough. If not, there will be a role for the DA to develop a stronger rule in order to successfully enforce against such

²²⁴ If these regulations fail to create or maintain competitive digital markets, the DA should alert Congress that neither antitrust nor existing regulatory tools are effective so that Congress may consider stronger steps.

²²⁵ For a deeper discussion of the power of discrimination in preventing entry by disintermediation, see Section II.2.B.3, *supra*.

foreclosure by firms with bottleneck power. Furthermore, the DA could use similar rules in a forward-looking regulatory context. Effective non-discrimination rules can foster entry and diversity, create potential sources of disruptive innovation and protect start-ups and other entrants. The goal of the forward-looking regulation is to prevent a digital business with bottleneck power from exercising it in order to protect *entrants* on the platform so that they have the chance to become *competitors* of the platform.

Platform strategies to prevent multi-homing are an important category for the DA to include in its analysis of foreclosure. A platform contract to induce single-homing on one side of a multi-sided market can be used to reduce competition (e.g., a loyalty payment to drivers of a ride-sharing service). Likewise, a platform can make it costlier for users to multi-home among applications on the platform by, for example, limiting data sharing. Such strategies can foreclose entrants and harm competition; appropriate regulation could limit their use by firms with bottleneck power.

The second reason for forward-looking regulation is to prevent digital businesses with bottleneck power from inefficiently expropriating rents created by complements on their platform. As described above, this harm is less well-protected by antitrust laws so the need for enhanced regulation will be greater when considering foreclosure *on* the platform. The DA may be able to build on the concept of “business to platform” regulation that is developing in Europe to create effective non-discrimination rules.²²⁶ The motivation of this EC regulatory effort is to create a “fair, predictable, sustainable, and trusted legal environment” in which complementors and content providers can invest safely and contribute to social welfare.²²⁷ The Vestager Report states that platforms “[i]f dominant . . . have a responsibility to ensure that they regulate in a pro-competitive way. Dominant platforms should be subject to a duty to ensure interoperability with suppliers of complementary services.”²²⁸ Likewise, the Furman Report recommends developing a platform code of conduct to ensure fairness.²²⁹

The DA could promulgate regulations prohibiting the foreclosure of a competing content provider on a platform that is vertically integrated. The authority would need to develop rules to identify foreclosure that might depend on finding certain anticompetitive conduct, market share, or market power to make such determinations. The DA must also account for potential pro-competitive innovations that a company with bottleneck power seeks to provide. The data

²²⁶ For regulatory developments, see European Commission (2018b).

²²⁷ European Commission, Directorate-General for CNECT (2018) (“The present proposal aims at ensuring a fair, predictable, sustainable and trusted legal environment for business users, corporate website users, providers of online intermediation services and online search engines alike, which will limit the occurrence and the impact of harmful platform-to-business trading practices occurring in certain online activities, thereby safeguarding trust in the online platform economy and preventing further legal fragmentation of the Digital Single Market.”).

²²⁸ Crémer (2019): 71.

²²⁹ See Furman (2019): 5.

collected by the DA will allow it to examine such outcomes and weigh all factors to make pro-competitive determinations. This balancing will be costly and, despite its best efforts, the regulator may err. However, as the report emphasizes, non-intervention is also costly. Non-discrimination rules in the past suffered from slow and expensive adjudication, which limited their usefulness.²³⁰ Strategies for speedy adjudication are addressed in detail below.

c. Bundling

A digital platform with bottleneck power may have a contract with complementors (e.g., retailers on an ecommerce platform) that bundles together access to their transaction data along with logistics services. This could have harmful anticompetitive effects.²³¹ As described above, the business may also compete against those retailers on its ecommerce site. The business could use the retailers' data to learn which products are selling well and expropriate the ideas and strategies of the retailer. That data advantage over rivals can enable a company to achieve and/or maintain critical economies of scale, better predict consumer behavior, and form a powerful barrier to entry for potential competitors.²³² Bundling may also discourage multi-homing.

Another example of possibly harmful bundling could occur when a platform owner requires installation of a bundle of applications. Those apps might be chosen to block the growth of rival apps that were extracting rents from the platform or threatening to be a future competitor of the platform. The Internet of Things will create more settings in which bundling policy will be critical. Will a consumer's new fridge arrive with a supply contract from Amazon's WholeFoods? Will the consumer be able to change that contract in some period of time, or is the consumer permitted to purchase the fridge without any contract?

An antitrust case in these settings may be ineffective in protecting entrants and competition in digital bottleneck businesses due to the complexity of the problem and the slow pace of litigation. However, the DA could establish regulations that prohibit anticompetitive bundling by firms with bottleneck power. Such a firm would be required to demonstrate that its bundle was on balance procompetitive if foreclosure was alleged. The DA could require unbundling and an offer to business customers of a choice of contracts in the case of anticompetitive bundling. The DA would need to enforce such contracts.

²³⁰ See Federal Communications Commission (2018): 1, n.2 and accompanying text (noting that "the [Federal Communication] Commission's 2008 Leased Access Order . . . has [been] stayed for a decade in conjunction with several judicial appeals").

²³¹ Retailers may accept what otherwise seems like a bad bargain due to the importance of being available to customers through the bottleneck. See Section II.2.B.4., *supra* for a discussion of bottleneck firms' ability to dictate business terms.

²³² See Section I.1.B.1, and Section I.2.A.3, *supra*.

3. DA-Enforced Remedies for Antitrust Violations

When a company has been found liable for violating the antitrust laws, the antitrust authority is tasked with devising a remedy to restore the lost competition. A fine does not restore lost competition. No longer engaging in the illegal conduct may help the next entrant or complement that wishes to interoperate, but it will typically not restore the competition that has already been lost, particularly in the face of durable barriers to entry that protect incumbent digital platforms.

Antitrust authorities are good at enforcing structural remedies that require no ongoing monitoring, such as requiring a divestiture between a platform and its content, or the sharing of a dataset or intellectual property with the entrant, royalty-free. However, reducing entry barriers often requires a remedy that involves ongoing monitoring, as do behavioral remedies, such as firewalls between platforms and content. As mentioned above, the DA could also enforce remedies for antitrust violations identified and addressed by existing antitrust agencies.²³³ Below are some types of remedies for which the DA would be a more appropriate body to enforce.²³⁴

1) Data Sharing

Anticompetitive conduct may result in a market that has tipped in favor of a single provider which then benefits from unparalleled access to data. In those cases, a new entrant may find it impossible to service users with new products as it lacks the scale needed for effective operation. Data sharing could restore the lost competition. The relevant data to share may not be just historical data, but present and future data also. Because data are non-rivalrous, an incumbent can both share its data with a competitor and also keep it. Thus, access to data forms a very important remedy in the toolkit of both the antitrust authority and the DA. The Furman Report recommends that agencies mandate “data openness” (which leads to data sharing) to enhance competition.²³⁵ The Vestager Report likewise recognized that data sharing can help level the playing field.²³⁶

2) Full Protocol Interoperability

Another useful tool that could restore lost competition is an open protocol and interoperability standard that would be available for entrants to use on a continuing basis and allow them to overcome network effects. A bottleneck business whose anticompetitive conduct created a monopoly position could be required to interoperate with its competitors. Entrants, previously rendered uncompetitive by network effects, could use the APIs to bring information

²³³ See Section III.3.A., *supra*.

²³⁴ For example, ongoing monitoring. See Section III.3.A, *supra*.

²³⁵ See Furman (2019): 10, 74, § 2.79.

²³⁶ See Crémer (2019): 98-107.

from the incumbent bottleneck firm to its own users. In a social media context this would allow the users of the new service to see not only all the content on their own service, but also content from friends on an incumbent site that was subject to an interoperability requirement. The network barrier to entry would no longer protect the incumbent firm, which would then encourage entry into the industry. Interoperability would facilitate ongoing competition on the merits of the user experience, rather than on the size of the installed base, and potentially stimulate robust competition.

Interoperability managed by the DA would be necessary due to the ongoing monitoring needed, the likelihood of technical change, and the incentive for non-cooperation by the incumbent firm. The DA could mandate the standard protocols or APIs to be applied and tightly control the process to avoid having competition undermined by actions of the dominant firm. The DA would need a process to update protocols at the time of the launch of new functionality or innovation. It would need rules to protect the privacy and choice of users on one service as some form of access to them is granted to users of another service. With easy interoperability, users will be free to make a real choice about which service they prefer. This will encourage new market entry and vigorous competition between providers.

3) Non-discrimination

The clear and simple remedy for a case when a bottleneck digital business favors its own content or complement is divestiture of one of the businesses, either the bottleneck business or the content/applications. This removes both the incentive and the ability for the conduct. However, this structural remedy could be costly to consumers in various ways, leading to the conclusion that a behavioral non-discrimination remedy might be more appropriate. Requiring a dominant bottleneck to abide by a non-discrimination rule could induce competitive entry by allowing complementary businesses to thrive and eventually become horizontal competitors to the bottleneck.²³⁷ The ongoing monitoring necessary to enforce this type of remedy in a specific antitrust case is not an ideal role for an antitrust agency. However, if the antitrust agency determines that such a remedy run by the DA would restore and protect competition, the law would allow it the option of requesting the DA to carry out the remedy. A speedy mechanism to adjudicate complaints would be of key.

4) Un-Bundling

As described above, the requirement to unbundle contracts could be an antitrust remedy that is less onerous than divestiture. Such a remedy would require ongoing monitoring that would best be performed by the DA.

²³⁷ See Section III.2.B.3.E.

4. Aligning other policies with competition

In addition to the structural competition tools, the authority should be empowered to align privacy protection, as well as AI and algorithmic oversight, with competition goals. Privacy protections that shield consumers from misuse or over-collection of their data can be set up in a way that raises or lowers entry barriers. This point is often forgotten inside a specialist agency and therefore we strongly suggest that the DA's mandate include evaluating and then directing regulatory solutions in a pro-competitive direction.²³⁸ Mandating that the DA have vigorous competition as one of its goals will help to make sure that potentially complex analysis takes place and that the needs of entrants are taken seriously when making policy decisions. The authority may also need to examine the development and use of algorithms to capture consumer attention, maximize advertising revenue, and drive consumer purchases or information selection, and examine how structural or other tools can promote competition in that space as well.

5. Adjudication Process

Adjudication of disputes by this new authority must be quick. Due to the fast pace of change in these industries, the short amount of time it takes to destabilize or eliminate an entrant, the substantial discrepancy in bargaining power between digital bottlenecks and their business customers, and the necessity to use government resources efficiently, a speedy process is crucial. This could be achieved through mandatory deadlines for dispute resolution or other procedural rules requiring the authority to produce a decision in a fixed number of days, as well as by crafting clear standards that are simple to enforce where possible.

The regulatory framework outlined above would coexist with and complement antitrust enforcement, only blocking transactions and preventing behavior that harm the public interest in ways clearly defined by Congress and that strengthen the reach of antitrust. The authority could help elucidate the common line-drawing problem of what behaviors are in or out of the reach of antitrust through communication with the antitrust enforcement agencies and with the public. An effective regulator will devise rules that promote competition and new avenues for innovation.

Conclusion

This report has discussed the unique combination of attributes of large digital platforms and their tendency toward entrenched market power. The entry barriers that result are in part due to certain characteristics of digital technology, but in part also due to behaviors of market participants. Consumers create entry barriers with their behavioral biases, and incumbents create entry barriers through strategic use of contracts and technologies, as well as by engaging in

²³⁸ It is worth considering whether the competition function and staff of the DA will be more effective when incorporated into other agency tasks or separated into a distinct division.

various other activities. The resulting monopoly or concentrated market structures do not serve consumers as well as would a market in which entry is a credible, or actual, reality.

Public policy has been slow to respond to economic harms resulting from these conditions. Such harms include advertising prices that are higher than would be expected in a more competitive environment. The markups distort decision-making and are subsequently passed through to consumers in the goods and services they buy. There are also quality harms. Platforms have incentives to provide low quality in order to keep users “engaged” and sell more ads. More generally, a lack of competition lessens the pressure on any platform to deliver high quality to its customers for fear that they will move to a rival platform. Perhaps most importantly, insufficient competition among and for digital platform position distorts and reduces innovation in a sector that has been—and, under the right conditions, will continue to be—the source of huge benefits for consumers and society.

It is unlikely that these problems will self-correct, meaning new and revised rules and incentives will be needed to prevent market power from entrenching a few dominant tech firms as economic and social gatekeepers. The United States is very far behind the frontier in antitrust enforcement, both because courts have taken a conservative view of what constitutes anticompetitive conduct and because agencies have not yet developed expertise in digital competition cases. Considerable work can be done by academics to help provide new relevant knowledge and tools to both agencies and courts. It also may be necessary for Congress to pass new legislation that revises the antitrust laws, establishes a specialist Competition Court, or both. The committee believes that vigorously enforcing the antitrust laws under these conditions would be likely to increase entry in digital platform industries, competition, and consumer welfare. Moreover, such enforcement would result in remedies to restore competition that has already been lost as well as serve as a deterrent to future anticompetitive conduct. Finally, because the problems we identify may require action beyond antitrust, we also propose the establishment of a new digital regulatory agency, or Digital Authority.

Such an agency could increase social welfare by establishing baseline market rules and conditions that both promote competition and limit the dimensions of competition to those that benefit consumers, rather than exploit them. In addition, the agency could assist in carrying out modern competition enforcement, using burdens of proof that reflect new understandings of market behaviors.

Whether the platforms that are dominant today have achieved that position entirely on the merits bears further examination. If illegalities are determined, appropriate remedies will be needed to restore the lost competition. However strong antitrust remedies likely will require coupling with market-opening regulations to overcome market power that prevented competitive market entry. Going forward, two important goals for public policy are ensuring entry and

competition *for* platforms, as well as creating competition *on* platforms. Forward-looking laws and regulations should be geared to lowering entry barriers as much as possible and vigorously guarding against exclusion, lock-in, and foreclosure in the platform context. Without a swift policy pivot to apply these combined tools, digital markets may tip toward levels of entrenched power that undermine the benefits of innovative digital businesses for American consumers.

References

- Acquisiti, Alessandro, Curtis R. Taylor, and Liad Wagman. 2016. "The Economics of Privacy." *Journal of Economic Literature* 54: 442.
- Acquisiti, Alessandro. 2004. "Privacy in Electronic Commerce and the Economics of Immediate Gratification." Pre-proceedings paper for the Fifth ACM Conference on Electronic Commerce, New York, May 17–24. <https://www.cs.drexel.edu/~greenie/privacy/privacy-gratification.pdf>.
- Acquisti, Alessandro, Laura Brandimarte, and George Loewenstein. 2015. "Privacy and Human Behavior in the Age of Information." *Science Magazine* 347: 509.
- Aguirre, Elizabeth M., Dominik Mahr, Dhruv Grewal, Ko de Ruyter, and Martin Wetzels. 2015. "Unraveling the Personalization Paradox: The Effect of Information Collection and Trust-Building Strategies on Online Advertisement Effectiveness." *Journal of Retailing* 91: 34.
- "Aim of the Project." n.d. AlgoTransparency.org. <https://algotransparency.org/methodology.html?candidat=Francois%20Fillon&file=ytreceos-presidentielle-2017-06-10>.
- Albanesius, Chloe. 2009. "More Americans Go to Facebook Than MySpace." *PC Magazine*, June 16. <https://www.pcmag.com/news/241432/more-americans-go-to-facebook-than-myspace>.
- Alcott, Hunt, Luca Braghieri, Sarah Eichmeyer, and Matthew Gentzkow. 2019. "The Welfare Effects of Social Media." NBER Working Paper No. 25514, April. <https://web.stanford.edu/~gentzkow/research/facebook.pdf>.
- "American Tech Giants Are Making Life Tough for Startups." 2018. *Economist*, June 2. <https://www.economist.com/business/2018/06/02/american-tech-giants-are-making-life-tough-for-startups>.
- "Americans' Attitudes About Privacy, Security, and Surveillance." 2015. Pew Research Center, May 20. https://www.pewresearch.org/wp-content/uploads/sites/9/2015/05/Privacy-and-Security-Attitudes-5.19.15_FINAL.pdf.
- Anderson, S. P., J. Waldfogel, and D. Stromberg, eds. 2016. *Handbook of Media Economics*, vol. 1A. North Holland.
- Anderson, Simon P., and Régis Renault. 2006. "Advertising Content." *American Economic Review* 96(1).
- "Andreessen Horowitz Makes Big Bet on Blockchain." 2018. *RealDaily*, February 13. <https://realdaily.com/andreessen-horowitz-makes-big-bet-on-blockchain>.

- Arrieta-Ibarra, Imanol, Leonard Goff, Diego Jiménez-Hernández, Jaron Lanier, and E. Glen Weyl. 2018. "Should We Treat Data as Labor? Moving Beyond 'Free'." *American Economic Association Papers and Proceedings* 108: 38. <https://pubs.aeaweb.org/doi/pdfplus/10.1257/pandp.20181003>.
- Ashenfelter, Orley, and Daniel Hosken. 2010. "The Effect of Mergers on Consumer Prices: Evidence from Five Mergers on the Enforcement Margin." *Journal of Law and Economics* 53: 417.
- "Assessing the Impact of Big Tech on Venture Investment." 2018. *Oliver Wyman*, July 11. <https://www.oliverwyman.com/content/dam/oliver-wyman/v2/publications/2018/july/assessing-impact.pdf>.
- Australian Competition and Consumer Commission. 2018. "Digital Platforms Inquiry: Preliminary Report." December. <https://www.accc.gov.au/system/files/ACCC%20Digital%20Platforms%20Inquiry%20-%20Preliminary%20Report.pdf>.
- Bagwell, Kyle. 2007. "The Economic Analysis of Advertising." In *3 Handbook of Industrial Organization*, ed. Mark Armstrong and Robert Porter. [https://doi.org/10.1016/S1573-448X\(06\)03028-7](https://doi.org/10.1016/S1573-448X(06)03028-7).
- Baker, Jonathan B. 2015. "Taking the Error Out of 'Error Cost' Analysis: What's Wrong with Antitrust's Right." *Antitrust Law Journal* 80: 1.
- Baker, Jonathan B. 2018. "Has the US Economy Become More Concentrated and Less Competitive: A Review of the Data." Paper presented at Georgetown University Law Center, September 13. https://www.ftc.gov/system/files/documents/public_events/1398386/georgetown-deck-913.pdf.
- Barberis, Nicholas, and Richard Thaler. 2003. "A Survey of Behavioral Finance." In *Handbook of the Economics of Finance*, ed. René M. Stulz and George Constantinides.
- Bar-Gill, Oren, and Ryan Bubb. 2012. "Credit Card Pricing: The CARD Act and Beyond." *Cornell Law Review* 97: 967.
- Bensinger, Greg. 2012. "In Kozmo.com's Failure, Lessons for Same-Day Delivery." *Wall Street Journal*, December 2. <https://blogs.wsj.com/digits/2012/12/03/in-kozmo-coms-failure-lessons-for-same-day-delivery>.
- Bergemann, Dirk, Alessandro Bonatti, and Alex Smolin. 2018. "The Design and Price of Information." *American Economic Review* 108: 1.

- Bergemann, Dirk, and Alessandro Bonatti. 2018a. "Markets for information: An Introduction." Cowles Foundation Discussion Paper No. 2142, August 29. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3240310.
- Bergemann, Dirk, and Alessandro Bonatti. 2018b. "The Wild West of Information Markets: What We Need to Know Before Law and Order Can Rule." *Vox*, October 11. <https://voxeu.org/article/wild-west-information-markets>;
- Bergen, Mark. 2019. "YouTube Executives Ignored Warnings, Letting Toxic Videos Run Rampant." *Bloomberg*, April 2. <https://www.bloomberg.com/news/features/2019-04-02/youtube-executives-ignored-warnings-letting-toxic-videos-run-rampant>.
- Bhattacharyya, Suman. 2019. "To Compete with Amazon, Walmart and eBay are Spending on Logistics." *Digiday*, January 15. <https://digiday.com/retail/compete-amazon-walmart-ebay-investing-logistics>.
- Blonigen, Bruce A., and Justin R. Pierce. 2016. "Evidence for the Effects of Mergers on Market Power and Efficiency." NBER Working Paper No. 22750. <https://www.nber.org/papers/w22750>.
- Bloomfield, Doni. 2014. "High-Speed Ad Traders Profit by Arbitraging Your Eyeballs." *Bloomberg*, November 7. <https://www.bloomberg.com/news/articles/2014-11-07/high-speed-ad-traders-profit-by-arbitraging-your-eyeballs>.
- Bond, Shannon, and Nicole Bullock. 2019. "Uber IPO Prospectus Shows Ride-Hailing Revenues Stalled." *Financial Times*, April 11. <https://www.ft.com/content/c68d3662-5c76-11e9-939a-341f5ada9d40>.
- Bork, Robert H. 1966. "Legislative Intent and the Policy of the Sherman Act." *Journal of Law and Economics* 9(1): article 5.
- Bork, Robert H., and Ward S. Bowman, Jr. 1963. "The Crisis in Antitrust." *Fortune*, December. Reprinted in *Columbia Law Review* 65: 363.
- Bowman, Ward S. Jr. 1957. "Tying Arrangements and the Leverage Problem." *Yale Law Journal* 67.
- Brown, Jeffrey R., and Austan Goolsbee. 2002. "Does the Internet Make Markets More Competitive? Evidence from the Life Insurance Industry." *Journal of Political Economy* 110.
- Brown, Pete. 2018. "More Than Half of Facebook Instant Articles Partners Have Abandoned It." *Columbia Journalism Review*, February 2. https://www.cjr.org/tow_center/are-facebook-instant-articles-worth-it.php.

Brynjolfsson, Erik, and Michael D. Smith. 2000. "Frictionless Commerce? A Comparison of Internet and Conventional Retailers." *Management Science* 46: 563.

Bundeskartellamt. 2019. "Bundeskartellamt Prohibits Facebook from Combining User Data from Different Sources." Press release, February 7.

https://www.bundeskartellamt.de/SharedDocs/Meldung/EN/Pressemitteilungen/2019/07_02_2019_Facebook.html.

Bureau of Transportation Statistics. 2017. "Sources of Aviation Data." U.S. Department of Transportation. <https://www.bts.gov/explore-topics-and-geography/topics/sources-aviation-data>.

Calkins, Stephen. 2018. "How Might a Justice Kavanaugh Impact Antitrust Jurisprudence?" Promarket Blog, July 20. <https://promarket.org/might-justice-kavanaugh-impact-antitrust-jurisprudence>.

Carlton, Dennis W. Forthcoming. "The Anticompetitive Effects of Vertical Most-Favored-Nation Restraints and the Error of *Amex*." *Columbia Business Law Review*.

https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3328628.

Center for Humane Technology. 2018. "Our Society is Being Hijacked by Technology." Humane Tech. <http://humanetech.com/problem>.

Cheng, Andrea. 2018. "Amazon, Facing Growing Competition, Offers Free Shipping on All Orders for the First Time." *Forbes*, November 5.

<https://www.forbes.com/sites/andriacheng/2018/11/05/amazon-free-shipping-holiday-minimum-order-walmart-target/#6c459dc555e1>.

Clark, Don. 2016. "Arista Networks Files Antitrust Claims Against Cisco." *Wall Street Journal*, January 25. <https://www.wsj.com/articles/arista-files-antitrust-charges-against-cisco-1453754642>.

"Confirmation Bias." 2019. *Wikipedia*, March 22. https://en.wikipedia.org/wiki/Confirmation_bias.

Crémer, Jacques, Yves-Alexandre de Montjoye, and Heike Schweitzer. 2019. "Competition Policy for the Digital Era." European Commission, Directorate-General for Competition, April 4. <http://ec.europa.eu/competition/publications/reports/kd0419345enn.pdf>.

Dafny, Leemore, Mark Duggan, and Subramaniam Ramanarayanan. 2012. "Paying a Premium on Your Premium? Consolidation in the US Health Insurance Industry." *American Economic Review* 102(2): 1161.

"Daniel Kahneman—Facts." 2019. Nobel Media AB. <https://www.nobelprize.org/prizes/economic-sciences/2002/kahneman/facts>.

“Data on 540 Million Facebook Users Exposed.” 2019. *BBC*, April 4.
<https://www.bbc.com/news/technology-47812470>.

DellaVigna, Stefano, and Ulrike Malmendier. 2006. “Paying Not to Go to the Gym.” *American Economic Review* 96: 694.

Delrahim, Makan, U.S. Assistant Attorney General. 2019. “‘I’m Free’: Platforms and Antitrust Enforcement in the Zero-Price Economy.” Address at Silicon Flatirons Annual Technology Policy Conference at The University of Colorado Law School, February 11.
<https://www.justice.gov/opa/speech/assistant-attorney-general-makan-delrahim-delivers-keynote-address-silicon-flatirons>.

De-Sola Gutiérrez, José, Fernando Rodríguez de Fonseca, and Gabriel Rubio. 2016. “Cell-Phone Addiction: A Review” *Frontiers Psychiatry* 7: 175.

Duhigg, Charles. 2012. “How Companies Learn Your Secrets.” *New York Times*, February 16.
<https://www.nytimes.com/2012/02/19/magazine/shopping-habits.html>.

Duhigg, Charles. 2018. “The Case Against Google.” *New York Times Magazine*, February 20.
<https://www.nytimes.com/2018/02/20/magazine/the-case-against-google.html>.

Dzieza, Josh. 2018. “Prime and Punishment: Dirty Dealing in the \$175 Billion Amazon Marketplace.” *Verge*, December 19.
<https://www.theverge.com/2018/12/19/18140799/amazon-marketplace-scams-seller-court-appeal-reinstatement>.

Easterbrook, Frank H. 1984. “Vertical Arrangements and the Rule of Reason.” *Antitrust Law Journal* 53: 134.

Edwards, Haley Sweetland. 2018. “You’re Addicted to Your Smartphone. This Company Thinks It Can Change That.” *Time*, April 13. <http://time.com/5237434/youre-addicted-to-your-smartphone-this-company-thinks-it-can-change-that>.

Elhauge, Einer. 2009. “Tying, Bundled Discounts, and the Death of the Single Monopoly Profit Theory.” *Harvard Law Review* 123: 397.
<https://www.bloomberg.com/news/features/2019-04-02/youtube-executives-ignored-warnings-letting-toxic-videos-run-rampant>.

Esayas, Samson Y. 2018. “Competition in (Data) Privacy: ‘Zero’-Price Markets, Market Power, and the Role of Competition Law.” *International Data Privacy Law* 8:181.

European Commission, Directorate-General for CNECT. 2018. “Proposal for a Regulation of the European Parliament and of the Council on Promoting Fairness and Transparency for Business Users of Online Intermediation Services.” 238 final, April 26. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A52018PC0238>.

- European Commission. 2009. “Antitrust: Commission Closes Formal Proceedings Against Qualcomm.” Press release, November 24. http://europa.eu/rapid/press-release_MEMO-09-516_en.htm.
- European Commission. 2014. “In-app Purchases: Joint Action by the European Commission and Member States is Leading to Better Protection for Consumers in Online Games.” Press release, July 18. http://europa.eu/rapid/press-release_IP-14-847_en.htm.
- European Commission. 2017a. “Competition Decision, Case COMP/AT.39740—Google Search(Shopping).” Antitrust Procedure Council Regulation, June 27. http://ec.europa.eu/competition/antitrust/cases/dec_docs/39740/39740_14996_3.pdf.
- European Commission. 2017b. “Report on the Monitoring Exercise Carried Out in the Online Hotel Booking Sector by EU Competition Authorities in 2016.” http://ec.europa.eu/competition/ecn/hotel_monitoring_report_en.pdf.
- European Commission. 2018a. “Antitrust: Commission Fines Google €4.34 Billion for Illegal Practices Regarding Android Mobile Devices to Strengthen Dominance of Google’s Search Engine.” Press release, July 18. http://europa.eu/rapid/press-release_IP-18-4581_en.htm.
- European Commission. 2018b. “Fairness in platform-to-business relations.” Digital economy and society initiative, provisional data, April 26. https://ec.europa.eu/info/law/better-regulation/initiatives/ares-2017-5222469_en.
- European Commission. 2019. “Antitrust: Commission Fines Google €1.49 Billion for Abusive Practices in Online Advertising.” Press release, March 20. http://europa.eu/rapid/press-release_IP-19-1770_en.htm.
- Ezrachi, Ariel, and Maurice E. Stucke. 2016. *Virtual Competition: The Promise and Perils of the Algorithm Driven Economy*. Harvard University Press.
- Ezrachi, Ariel, and Maurice E. Stucke. 2018. “Digitalisation and Its Impact on Innovation: Report Prepared for the European Commission.” Working paper, DG Research and Innovation, July.
- “Fairness in Platform-to-Business Relations.” 2018. European Commission. https://ec.europa.eu/info/law/better-regulation/initiatives/ares-2017-5222469_en.
- Farivar, Cyrus. 2018. “Advertisers Allege Facebook Hid the Fact That No One Watches Video Ads.” *Arstechnica*, October 17. <https://arstechnica.com/tech-policy/2018/10/advertisers-allege-facebook-hid-the-fact-that-no-one-watches-video-ads>.
- Farrell, Joseph, and Paul Klemperer. 2007. “Coordination and Lock-In: Competition with Switching Costs and Network Effects.” In *3 Handbook of Industrial Organization*, ed.

Mark Armstrong and Robert Porter.

<https://www.sciencedirect.com/science/article/pii/S1573448X06030317>.

Federal Communications Commission. 2018. “Leased Commercial Access: Modernization of Media Regulation Initiative.” Further Notice of Proposed Rulemaking, June 8.

<https://docs.fcc.gov/public/attachments/FCC-18-80A1.pdf>.

Federal Trade Commission. 2010. “Is There Life After *Trinko* and *Credit Suisse*? The Role of Antitrust in Regulated Industries.” Hearing Before the Subcommittee on Courts and Competition Policy of the House of Representatives Committee of the Judiciary, 111th United States Congress, July 15.

https://www.ftc.gov/sites/default/files/documents/public_statements/prepared-statement-federal-trade-commission-courts-and-competition-policy-committee-judiciary-united/100615antitrusttestimony.pdf.

Federal Trade Commission. 2014. “Data Brokers: A Call For Transparency and Accountability: A Report of the Federal Trade Commission.”

<https://www.ftc.gov/system/files/documents/reports/data-brokers-call-transparency-accountability-report-federal-trade-commission-may-2014/140527databrokerreport.pdf>.

Federal Trade Commission. 2015. “Apple Inc. Will Provide Full Consumer Refunds of At Least \$32.5 Million to Settle FTC Complaint It Charged for Kids’ In-App Purchases Without Parental Consent.” Press release, January 15. <https://www.ftc.gov/news-events/press-releases/2014/01/apple-inc-will-provide-full-consumer-refunds-least-325-million>.

Federal Trade Commission. 2018. “Transcript: FTC Hearing #3: Competition and Consumer Protection in the 21st Century.” Session 4, October 15.

https://www.ftc.gov/system/files/documents/public_events/1413712/ftc_hearings_session_3_transcript_day_1_10-15-18.pdf.

Federico, Giulio, Fiona Scott Morton, and Carl Shapiro. 2019. “Antitrust and Innovation: Welcoming and Protecting Disruption.” Working paper, NBER Innovation Policy and the Economy.

FitzGerald, Drew. 2019. “Facebook Looks to Build Underwater Ring Around Africa.” *Wall Street Journal*, April 7.

Fletcher, Amelia. 2019. “The EU Google Decisions: Extreme Enforcement or the Tip of the Behavioral Iceberg?” *Competition Policy International Antitrust Chronicle*, January.

<https://www.competitionpolicyinternational.com/wp-content/uploads/2019/01/CPI-Fletcher.pdf>.

French Competition Authority [L’Autorité de la Concurrence]. 2018. “Portant sur l’exploitation des données dans le secteur de la publicité sur internet [On the Exploitation of Data in the Internet Advertising Sector].” March 6.

<http://www.autoritedelaconcurrence.fr/pdf/avis/18a03.pdf>.

- Hoffman, D. Bruce. 2018. "Vertical Merger Enforcement at the FTC." Address at the Credit Suisse 2018 Washington Perspectives Conference, January 10.
https://www.ftc.gov/system/files/documents/public_statements/1304213/hoffman_vertical_merger_speech_final.pdf.
- "Horizontal Merger Guidelines." 2010. U.S. Department of Justice and the Federal Trade Commission. <https://www.ftc.gov/sites/default/files/attachments/merger-review/100819hmg.pdf>.
- Hovenkamp, Herbert. 2019. "Platforms and the Rule of Reason: The *American Express* Case." *Columbia Business Law Review* 2019(2).
https://scholarship.law.upenn.edu/cgi/viewcontent.cgi?article=3060&context=faculty_scholarship.
- "How It Works." 2019. *Solid*. <https://solid.inrupt.com/how-it-works>.
- Isaac, Mike. 2019. "Zuckerberg Plans to Integrate WhatsApp, Instagram and Facebook Messenger." *New York Times*, January 25.
<https://www.nytimes.com/2019/01/25/technology/facebook-instagram-whatsapp-messenger.html>.
- Japanese Ministry of Economy, Trade, and Industry. 2018. "Fundamental Principles for Rule Making to Address the Rise of Platform Businesses Formulated." December 18.
https://www.meti.go.jp/english/press/2018/1218_002.html.
- Johnson, Eric J., and Daniel Goldstein. 2003. "Do Defaults Save Lives?" *Science* 302: 1338.
- Johnson, Justin P. 2013. "Targeted Advertising and Advertising Avoidance." *Rand Journal of Economics* 44: 128.
- Kades, Michael. 2018. "In Conversation with Herbert Hovenkamp and Ioana Marinescu." *Equitable Growth*, October 11. <https://equitablegrowth.org/in-conversation-with-herbert-hovenkamp-and-ioana-marinescu>.
- Kahneman, Daniel, and Thomas Gilovich. 2002. *Heuristics and Biases: The Psychology of Intuitive Judgment*. Cambridge University Press.
- Kahneman, Daniel, Jack L. Knetsch, and Richard H. Thaler. 1991. "Anomalies: The Endowment Effect, Loss Aversion, and Status Quo Bias." *Journal of Economic Perspectives* 5: 193.
- Kaminska, Izabella. 2016. "The Taxi Unicorn's New Clothes." *Financial Times*, December 1.
<https://ftalphaville.ft.com/2016/12/01/2180647/the-taxi-unicorns-new-clothes>.
- Kaminska, Izabella. 2019. "Amazon (sub)Prime?" *Financial Times*, March 20.
<https://ftalphaville.ft.com/2019/03/20/1553085361000/Amazon--sub-Prime->

- Kastrenakes, Jacob. 2018. "Facebook is Shutting Down a Teen App it Bought Eight Months Ago, Along with Two Others That Languished." *Verge*, July 2. <https://www.theverge.com/2018/7/2/17528896/facebook-tbh-moves-hello-shut-down-low-usage>.
- Katz, Michael and Jonathan Sallet. 2018. "Multisided Platforms and Antitrust Enforcement." *Yale Law Journal* 127: 1724.
- Katz, Michael L. 2019. "Platform Economics and Antitrust Enforcement: A Little Knowledge is a Dangerous Thing." *Journal of Economics and Management Strategy* 28: 138.
- Khan, Lina. 2017. "Amazon's Antitrust Paradox." *Yale Law Journal* 126: 710.
- Khan, Lina. 2019. "The Separation of Platforms and Commerce." *Columbia Law Review* 119: 973. <https://papers.ssrn.com/id=3180174>.
- Kraus, Rachel. 2019. "Amazon Bought Eero for \$97 Million and Employees Still Got Screwed." *Mashable*, April 5. <https://mashable.com/article/amazon-eero-wifi-router-sale/#HN0veV33jPqc>.
- Kulick, Robert B. 2017. "Ready-to-Mix: Horizontal Mergers, Prices, and Productivity." U.S. Census Bureau, Center for Economic Studies, Working Paper No. 17-38. <https://papers.ssrn.com/abstract=2637961>.
- Kuss, Daria J., Eiman Kanjo, Mark Crook-Rumsey, Fraenze Kibowski, Grace Y. Wang, and Alex Sumich. 2018. "Problematic Mobile Phone Use and Addiction Across Generations: The Roles of Psychopathological Symptoms and Smartphone Use." *Journal of Technology in Behavioral Science* 3: 141.
- Kwoka, John E. 2017. "The Structural Presumption and the Safe Harbor in Merger Review: False Positives or Unwarranted Concerns?" *Antitrust Law Journal* 81: 837.
- Kwoka, John. 2014. *Mergers, Merger Control, and Remedies: A Retrospective Analysis of U.S. Policy*. MIT Press.
- Kwoka, John. 2015. "The Changing Nature of Efficiencies in Mergers and in Merger Analysis." *Antitrust Bull* 60(3): 231.
- Lanier, Jaron. 2013. *Who Owns the Future?* Simon and Schuster.
- Levy, Steven. 2014. "Mark Zuckerberg on Facebook's Future, From Virtual Reality to Anonymity." *Wired*, April 30. <https://www.wired.com/2014/04/zuckerberg-f8-interview>.
- Lewis, Paul, and Erin McCormick. 2018. "How an Ex-YouTube Insider Investigated its Secret Algorithm." *Guardian*. February 2. <https://www.theguardian.com/technology/2018/feb/02/youtube-algorithm-election-clinton-trump-guillaume-chaslot>.

- Lewis, Paul. 2018. “‘Fiction is Outperforming Reality’: How YouTube’s Algorithm Distorts Truth.” *Guardian*, February 2.
<https://www.theguardian.com/technology/2018/feb/02/how-youtubes-algorithm-distorts-truth>.
- “List of Mergers and Acquisitions by Alphabet.” 2019. *Wikipedia*, April 12.
https://en.wikipedia.org/wiki/List_of_mergers_and_acquisitions_by_Alphabet.
- “List of Mergers and Acquisitions by Amazon.” 2019. *Wikipedia*, April 12.
https://en.wikipedia.org/wiki/List_of_mergers_and_acquisitions_by_Amazon.
- “List of Mergers and Acquisitions by Facebook.” 2019. *Wikipedia*, March 25.
https://en.wikipedia.org/wiki/List_of_mergers_and_acquisitions_by_Facebook.
- MacMillan, Douglas, and Robert McMillan. 2018. “Google Exposed User Data, Feared Repercussions of Disclosing to Public.” *Wall Street Journal*, October 8.
<https://www.wsj.com/articles/google-exposed-user-data-feared-repercussions-of-disclosing-to-public-1539017194>.
- Manthorpe, Rowland. 2018. “What is Open Banking and PSD2? WIRED Explains.” *Wired*, April 17. <https://www.wired.co.uk/article/open-banking-cma-psd2-explained>.
- McArdle, Megan. 2019. “Uber and Lyft Are Losing Money. At Some Point, We’ll Pay for It.” *Washington Post*, March 5. https://www.washingtonpost.com/opinions/uber-and-lyft-are-losing-money-at-some-point-well-pay-for-it/2019/03/05/addd607c-3f95-11e9-a0d3-1210e58a94cf_story.html.
- Melamed, A. Douglas, and Nicolas Petit. 2019. “The Misguided Assault on the Consumer Welfare Standard in the Age of Platform Markets.” *Review of Industrial Organization* 54: 1.
- Melamed, A. Douglas. 2006. “Exclusionary Conduct Under the Antitrust Laws: Balancing, Sacrifice and Refusals to Deal.” *Berkeley Technology Law Journal* 20: 1247.
- Melamed, A. Douglas. 2019. “Response: Antitrust Law is Not That Complicated.” *Harvard Law Review Forum* 130: 163.
- “Millions of Facebook Passwords Exposed Internally.” 2019. *BBC*, March 21.
<https://www.bbc.com/news/technology-47653656>.
- “Mission and Values.” 2019. *RadicalxChange*. <https://radicalxchange.org/about/>.
- “More Americans Go to Facebook Than MySpace.” 2009. *PC Magazine*, June 16.
<https://www.pcmag.com/news/241432/more-americans-go-to-facebook-than-myspace>.

- Morton, Fiona Scott, R. Craig Romaine, and Spencer Graf. 2015. "Benefits of Preserving Consumers' Ability to Compare Airline Fares." Unpublished manuscript, prepared for Travel Technology Association, May 15.
- Morton, Fiona Scott. 2019. "Modern U.S. antitrust theory and evidence amid rising concerns of market power and its effects." Washington Center for Equitable Growth. <https://equitablegrowth.org/research-paper/modern-u-s-antitrust-theory-and-evidence-amid-rising-concerns-of-market-power-and-its-effects/>.
- "Most Popular Mobile Messaging Apps Worldwide as of January 2019." 2019. *Statista*, January. <https://www.statista.com/statistics/258749/most-popular-global-mobile-messenger-apps>.
- Mullins, Brody, and Jack Nicas. 2017. "Pay Professors: Inside Google's Academic Influence Campaign." *Wall Street Journal*, July 14. <https://www.wsj.com/articles/paying-professors-inside-googles-academic-influence-campaign-1499785286>.
- Newton, Casey. 2016. "Why Vine Died." *Verge*, October 28. <https://www.theverge.com/2016/10/28/13456208/why-vine-died-twitter-shutdown>.
- Newton, Casey. 2019. "How YouTube Helps Flat-Earthers Organize." *Verge*, February 20. <https://www.theverge.com/interface/2019/2/20/18232524/youtube-flat-earth-recommendation-algorithm-conspiracy>.
- OECD. 2014. "Data-Driven Innovation for Growth and Well-Being: Interim Synthesis Report." *New Sources of Growth: Knowledge-Based Capital*, October. <http://www.oecd.org/sti/inno/data-driven-innovation-interim-synthesis.pdf>.
- Open Banking Ltd. 2018. "Open Banking: Guidelines for Participants." <https://www.openbanking.org.uk/wp-content/uploads/Guidelines-for-Open-Data-Participants.pdf>.
- Oremus, Will. 2017. "Twitter's New Order." *Slate*, March 5. http://www.slate.com/articles/technology/cover_story/2017/03/twitter_s_timeline_algorithm_and_its_effect_on_us_explained.html.
- Osborne, Hilary. 2013. "OFT Warns Free Online Games Pressure Children Into In-app Purchases." *Guardian*, September 26. <https://www.theguardian.com/money/2013/sep/26/free-online-games-children-apps>.
- Pidgeon, David. 2016. "Where Did the Money Go? Guardian Buys its Own Ad Inventory." *Media Newsline*, October 4. <https://mediatel.co.uk/newsline/2016/10/04/where-did-the-money-go-guardian-buys-its-own-ad-inventory>.
- Posner, Eric A., and E. Glen Weyl. 2018. *Radical Markets: Uprooting Capitalism and Democracy for a Just Society*. Princeton University Press.

- Preger, Elena, and Matt Schmitt. 2019. “Employer Consolidation and Wages: Evidence from Hospitals.” Working paper, Washington Center for Equitable Growth.
<https://equitablegrowth.org/working-papers/employer-consolidation-and-wages-evidence-from-hospitals>.
- Reishut HaTachrut [Israel Antitrust Authority]. 2018. “The Israel Antitrust Authority invites interested members of the public to submit their comments, towards an inquiry into competition issues in the digital economy.” Press release, April 9.
<http://www.antitrust.gov.il/eng/subject/177/item/35246.aspx>.
- Reyna, Agustín. 2018. “The Psychology of Privacy—What Can Behavioural Economics Contribute to Competition in Digital Markets?” *International Data Privacy Law* 8: 240.
- “Richard Thaler—Facts.” 2019. Nobel Media AB. <https://www.nobelprize.org/prizes/economic-sciences/2017/thaler/facts>.
- Robertson, Adi. 2018. “Mark Zuckerberg Personally Approved Cutting Off Vine’s Friend-finding Feature.” *Verge*, December 18.
<https://www.theverge.com/2018/12/5/18127202/mark-zuckerberg-facebook-vine-friends-api-block-parliament-documents>.
- Rochet, Jean-Charles, and Jean Tirole. 2003. “Platform Competition in Two-Sided Markets.” *Journal of the European Economic Association* 1: 990.
- Romm, Tony and Elizabeth Dwoskin. 2019. “U.S. Regulators Have Met to Discuss Imposing a Record-Setting Fine Against Facebook for Privacy Violations.” *Washington Post*, January 18. <https://www.washingtonpost.com/technology/2019/01/18/us-regulators-have-met-discuss-imposing-record-setting-fine-against-facebook-some-its-privacy-violations>.
- Roose, Kevin. 2019. “YouTube Unleashed a Conspiracy Theory Boom. Can it Be Contained?” *New York Times*, February 19.
<https://www.nytimes.com/2019/02/19/technology/youtube-conspiracy-stars.html>.
- Ryding, Francesca C., and Linda K. Kaye. 2018. “‘Internet Addiction’: A Conceptual Minefield.” *International Journal of Mental Health and Addiction* 16: 225.
- Salop, Steven C. 1986. “Practices that (Credibly) Facilitate Oligopoly Co-ordination.” In *New Developments in the Analysis of Market Structure: Proceedings of a conference held by the International Economic Association in Ottawa, Canada*, ed. Joseph E. Stiglitz and G. Frank Mathewson.
- Salop, Steven C. 2018a. “Invigorating Vertical Merger Enforcement.” *Yale Law Journal* 127: 1742.
- Salop, Steven C. 2018b. “Revising the Vertical Merger Guidelines: Presentation at the FTC Hearings on Competition and Consumer Protection in the 21st Century,” November 1.

https://www.ftc.gov/system/files/documents/public_events/1415284/ftc_hearings_5_georgetown_slides.pdf.

- Samuelson, William and Richard Zeckhauser. 1988. "Status Quo Bias in Decision Making." *Journal of Risk and Uncertainty* 1.
- Schweitzer, Heike, Justus Haucap, Wolfgang Kerber, and Robert Welker. 2018. "Modernising the Law on Abuse of Market Power: Report for the Federal Ministry for Economic Affairs and Energy." German Bundesministerium Wirtschaft und Energie, April 9. <https://www.bmwi.de/Redaktion/DE/Downloads/Studien/modernisierung-der-missbrauchsaufsicht-fuer-marktmaechtige-unternehmen-zusammenfassung-englisch.pdf>.
- Sedghi, Ami. 2014. "Facebook: 10 Years of Social Networking, in Numbers." *Guardian*, February 4. <https://www.theguardian.com/news/datablog/2014/feb/04/facebook-in-numbers-statistics>.
- Seetharaman, Deepa, and Betsy Morris. 2017. "Facebook's Onavo Gives Social-Media Firm Inside Peek at Rivals' Users." *Wall Street Journal*, August 13. <https://www.wsj.com/articles/facebook-onavo-gives-social-media-firm-inside-peek-at-rivals-users-15026220031>.
- Select Committee on Communications, House of Lords. 2018. "U.K. Advertising in a Digital Age." 1st Report of Session 2017–19, HL Paper 116. <https://publications.parliament.uk/pa/ld201719/ldselect/ldcomuni/116/116.pdf>;
- Shah, Semile. 2015. "Interview with Chamath Palihapitiya, CEO." *Social Capital*, September 16. <http://haystack.vc/2015/09/17/transcript-chamath-at-strictlyvcs-insider-series>.
- Shapiro, Carl and Hal R. Varian. 1999. *Information Rules: A Strategic Guide to the Network Economy*. Harvard Business School Press.
- Shapiro, Carl. 2005. "Exclusionary Conduct: Testimony Before the Antitrust Modernization Commission." September 29. <http://faculty.haas.berkeley.edu/Shapiro/amcexclusion.pdf>.
- Shapiro, Robert, and Siddhartha Aneja. 2019. "Who Owns Americans' Personal Information and What Is It Worth?" *Future Majority*, March 8. <https://assets.futuremajority.org/uploads/report-for-future-majority-on-the-value-of-people-s-personal-data-shapiro-aneja-march-8-2019.pdf>.
- Sheetz, Michael. 2019. "Here's Why Amazon Is Trying to Reach Every Inch of the World With Satellites Providing Internet." *CNBC*, April 7. <https://www.cnbc.com/2019/04/05/jeff-bezos-amazon-internet-satellites-4-billion-new-customers.html>.
- Sher, Scott A., and Bradley T. Tennis. 2016. "Exploiting Others' Investments in Open Standards." *Competition Policy International*, September 15.

<https://www.competitionpolicyinternational.com/exploiting-others-investments-in-open-standards/>.

- Shields, Mike. 2016. "Facebook Says it Found More Miscalculated Metrics." *Wall Street Journal*, November 16. <https://www.wsj.com/articles/facebook-says-it-found-more-miscalculated-metrics-1479303984>;
- Smith, Alasdair, Tom Hoehn, Philip Marsden, Jill May, and Ed Smith. 2016. "Retail Banking Market Investigation: Final Report." U.K. Competition and Markets Authority, August. <https://assets.publishing.service.gov.uk/media/57ac9667e5274a0f6c00007a/retail-banking-market-investigation-full-final-report.pdf>.
- Son, Hugh. 2018. "JP Morgan to unveil new investing app with an eye-catching, disruptive price: Free." *CNBC*, August 21. <https://www.cnbc.com/2018/08/21/jp-morgan-to-unveil-new-investing-app-with-an-eye-catching-disruptive-price-free.html>.
- Spangler, Todd. 2018. "Google Killing Off DoubleClick, AdWords Names in Rebranding of Ad Products," *Variety*, July 27. <https://variety.com/2018/digital/news/google-doubleclick-adwords-rebranding-1202859088>.
- Statt, Nick. 2018. "Amazon and Walmart's Rivalry is Reshaping How We'll Buy Everything in the Future." *Verge*, August 3. <https://www.theverge.com/2018/8/3/17630604/amazon-walmart-competition-tech-partnerships-grocery-delivery>.
- Stucke, Maurice E., and Ariel Ezrachi. 2016. "The Subtle Ways Your Digital Assistant Might Manipulate You." *Wired*, November 29. <https://www.wired.com/2016/11/subtle-ways-digital-assistant-might-manipulate>.
- Stucke, Maurice E., and Ariel Ezrachi. 2017. "Looking Up in the Data-Driven Economy." *Competition Policy International*, May 15. <https://www.competitionpolicyinternational.com/looking-up-in-the-data-driven-economy>.
- Stucke, Maurice E., and Ariel Ezrachi. 2018a. "Alexa et. al., What Are You Doing with My Data?" *Critical Analysis of Law* 5: 148.
- Stucke, Maurice E., and Ariel Ezrachi. 2018b. "How Digital Assistants Can Harm Our Economy, Privacy, and Democracy." *Berkeley Technology Law Journal* 32: 1239.
- Sunstein, Cass R. 2013. "The Storrs Lectures: Behavioral Economics and Paternalism." *Yale Law Journal* 122: 1826.
- Sutton, John. 1991. *Sunk Costs and Market Structure*. MIT Press.
- Sweney, Mark. 2018. "Is Facebook for Old People? Over-55s Flock in as the Young Leave." *Guardian*, February 12. <https://www.theguardian.com/technology/2018/feb/12/is-facebook-for-old-people-over-55s-flock-in-as-the-young-leave>.

- “Tech Monopolies Are Stifling U.S. Innovation. Antitrust Enforcement May Help.” 2018. *Brink News*, December 11. <http://www.brinknews.com/is-antitrust-enforcement-the-wests-best-weapon-against-chinese-tech>.
- Thaler, Richard H. 2018. “From Cashews to Nudges: The Evolution of Behavioral Economics.” *American Economic Review* 108:1265.
- Thaler, Richard H., and Cass Sunstein. 2003. “Libertarian Paternalism.” *American Economic Review Papers and Proceedings* 93:175.
- Thompson, Ben. 2018. “Tech’s Two Philosophies.” *Stratechery*, May 9. <https://stratechery.com/2018/techs-two-philosophies>.
- Thompson, Ben. 2019. “The BuzzFeed Lesson.” *Stratechery*, January 28. <https://stratechery.com/2019/the-buzzfeed-lesson>.
- Tompson, Nicholas. 2018. “When Tech Knows You Better Than You Know Yourself.” *Wired*, October 4. <https://www.wired.com/story/artificial-intelligence-yuval-noah-harari-tristan-harris/>.
- U.K. Competition and Markets Authority. 2015. “CMA Refers 3 Games to the ASA.” Press release, June 4. <https://www.gov.uk/government/news/cma-refers-three-childrens-online-games-to-the-asa>.
- U.K. Competition and Markets Authority. 2017. “The Retail Banking Market Investigation Order 2017.” https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/600842/retail-banking-market-investigation-order-2017.pdf.
- U.K. Financial Conduct Authority. 2017. “FCA Finalises Revised Payment Services Directive (PSD2) Requirements.” Press release, September 9. <https://www.fca.org.uk/news/press-releases/fca-finalises-revised-psd2-requirements>.
- Vincent, James. 2019. “Spotify Files Antitrust Complaint Over ‘Apple Tax,’” *Verge*, March 13. <https://www.theverge.com/2019/3/13/18263453/spotify-apple-app-store-antitrust-complaint-ec-30-percent-cut-unfair>.
- Vita, Michael and F. David Osinski. 2018. “John Kwoka's Mergers, Merger Control, and Remedies: A Critical Review.” *Antitrust Law Journal* 82: 361.
- Vranica, Suzanne and Jack Marshall. 2016. “Facebook Overestimated Key Video Metric for Two Years.” *Wall Street Journal*, September 22. <https://www.wsj.com/articles/facebook-overestimated-key-video-metric-for-two-years-1474586951>.
- Warren, Elizabeth. 2019a. “Here’s How We Can Break Up Big Tech.” *Medium*, March 8. <https://medium.com/@teamwarren/heres-how-we-can-break-up-big-tech-9ad9e0da324c>.

- Warren, Elizabeth. 2019b. "Leveling the Playing Field for America's Family Farmers." *Medium*, March 27. <https://medium.com/@teamwarren/leveling-the-playing-field-for-americas-family-farmers-823d1994f067>.
- Weber, Harrison. 2016. "Facebook Kicked Zynga to the Curb, Publishers Are Next." *Venturebeat*. June 30. <https://venturebeat.com/2016/06/30/facebook-kicked-zynga-to-the-curb-publishers-are-next>.
- Welch, Chris. 2012. "Tumblr 'Truly Disappointed by Twitter's Decision' to Block Friend-Finding On the Blogging Service." *Verge*, August 22. <https://www.theverge.com/2012/8/22/3261270/tumblr-removes-twitter-find-people-you-know>.
- Wen, Wen and Feng Zhu. 2018. "Threat of Platform-Owner Entry and Complementor Responses: Evidence from the Mobile App Market." Harvard Business School Working Paper 18-036. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=284853.
- Wisner, David. 2012. "Google's Larry Page: 'Competition is One Click Away,' (and Other Quotes of the Week)" *Forbes*, October 14. <https://www.forbes.com/sites/davidwisner/2012/10/14/googles-larry-page-competition-is-one-click-away-and-other-quotes-of-the-week>.
- Wolfe, Jan. 2018. "Arista to Pay \$400 Million to Cisco to Resolve Court Fight." *Reuters*, August 6. <https://www.reuters.com/article/us-cisco-arista-settlement/arista-to-pay-400-million-to-cisco-to-resolve-court-fight-idUSKBN1KR1PI>.
- Wollman, Thomas. 2019. "Stealth Consolidation: Evidence from an Amendment to the Hart-Scott-Rodino Act," *American Economic Review Insights* 1(1). <https://www.aeaweb.org/articles?id=10.1257/aeri.20180137>.
- Wu, Tim. 2018. "The Case for Breaking Up Facebook and Instagram." *Washington Post*, September 28. <https://www.washingtonpost.com/outlook/2018/09/28/case-breaking-up-facebook-instagram>.
- Zajac, Kristyn, M. K. Ginley, R. Chang, and N. M. Petry. 2017. "Treatments for Internet Gaming Disorder and Internet Addiction: A Systematic Review." *Psychology of Addictive Behaviors* 31: 979.
- Zamir, Eyal, and Doron Teichman, eds. 2014. *The Oxford Handbook of Behavioral Economics and the Law*. Oxford University Press.
- Zhu, Feng, and Qihong Liu. 2018. "Competing with Complementors: An Empirical Look at Amazon.com." *Strategic Management Journal* 39: 2618.
- Zhu, Feng. 2018. "Friends or Foes? Examining Platform Owners' Entry Into Complementors' Spaces." *Journal of Economics and Management Strategy* 28: 23.

Protecting Journalism in the Age of Digital Platforms²³⁹

**Stigler Committee on Digital Platforms
News Industry Subcommittee**

Media subcommittee members²⁴⁰

Guy Rolnik (Chair)

University of Chicago

Julia Cagé

Sciences Po Paris

Joshua Gans

University of Toronto

Ellen Goodman

Rutgers University

Brian Knight

Brown University

Andrea Prat

Columbia University

Anya Schiffrin

Columbia University

Additional contributors

Prateek Raj

IIM Bangalore

²³⁹ We would like to thank Luigi Zingales, Fiona Scott Morton, James Hamilton, Robert McChesney, Sally Hubbard, Jesse Eisinger, Rodney Benson, Alberto Ibarguen, Betsy Reed and Asher Schechter. We will also like to thank Filippo Lancieri for help in organizing the project, Samantha Eyster-Driscoll for copyediting, and Sebastian Burca, Simone Cavallaro and Rachel Piontek for their invaluable support at the Stigler Center. The Committee in-person meetings were partially supported by a grant from the Alfred P. Sloan Foundation, whom we also thank for supporting this project. Corresponding author: guy.rolnik@chicagobooth.edu

²⁴⁰ The main proposals of the report represent the opinion of all the Media subcommittee members but some may have disagreements on specific parts of the report.

Table of Contents

Protecting Journalism in the Age of Digital Platforms 139

Executive Summary..... 142

Report from the Media Subcommittee..... 146

The Precarious Economics of News..... 153

Digital Platforms and Disruption of News 156

 Production of News..... 156

 Reducing the Incentive to Produce Original Content 156

 Platform Duopoly and the Business Model of the News Media..... 157

 Distribution of News..... 158

 Disaggregation of the Customer 158

 Advertising Networks and Attention Aggregation 159

 Subscriptions and Attention Aggregation..... 159

 Social Media and Attention Aggregation 160

 Market Power in Advertising..... 160

 Consumption of News..... 161

 Atomization of News..... 162

 Moderation of News Content..... 162

The Market Response..... 167

 The Multifaceted Donation Model..... 167

 The Newsrooms Collaboration: An Alternative Path for Non-profit Journalism 170

 The Subscription Model..... 171

 Rating the Sources: The Solution to the Spread of Disinformation? 172

 Summary 172

Recommendations 173

Public Funding of News..... 174

 Existing Subsidies Schemes..... 175

 Our Proposal: Media Vouchers..... 176

 Which media outlets could benefit from the media vouchers?..... 176

 A focus on local media outlets..... 178

What happens if a citizen decides not to allocate her voucher?..... 178

How can concentration be avoided? 179

Will the vouchers make a difference?..... 179

Why do we advocate media vouchers? 179

Media vouchers and the status quo 180

Media vouchers and alternative funding models 181

Lessons from the Seattle voucher experience 183

The limits of media vouchers..... 183

Subsidizing inputs versus outputs..... 184

Regulation of Media Concentration..... 184

 Citizen Welfare Standard 184

 Concentration and Media Capture 185

 Inadequacy of Existing Regulation to Prevent Concentration in News Provision..... 185

 Our Proposal: A Quantifiable Citizen Welfare Criterion for Reviewing Mergers between News Providers
 186

Increasing the Transparency of Digital Platforms 187

 Our Proposal: Transparency Requirements and Voluntary Labeling 188

 Transparency Requirements 188

 Labeling 189

Increasing Accountability of Digital Platforms 190

 Section 230 Background 191

 Proposals to Amend Section 230 192

 Our Proposal: Section 230 as a Quid Pro Quo Benefit 194

Conclusion 196

References 198

Executive Summary

While the Internet has contributed immensely to access to and diffusion of information and has opened numerous opportunities to improve life across the globe, it has also brought challenges, risks and harms that may endanger the very democratic and liberal order that many believed it would advance.

Perhaps nowhere are these contradictory trends more present than in the fourth estate—the news media and journalism. The technology that made news creation and diffusion cheaper and faster and gave billions of people a voice has also become a tool used by state and private powers to manipulate and propagate disinformation and hate. It has also disrupted the business model of original news creators, disintermediated them from their consumers and created a new news ecosystem.

Growing concentration in the business sector in the US and the accumulation of market and political power by large corporations across the developed world has attracted more scrutiny in the last decade. It is now clear that the natural tendency toward concentration in modern capitalism is magnified in digital markets, where a handful of corporations enjoying network effects today exercise more power and influence globally than any other private entities have in the last century. Two of those companies—Google and Facebook—are not only giant economic players that have changed most industries, but are also the largest media companies in history. While they maintain that they are technology companies, they not only have unprecedented influence on news production, distribution and consumption, but also are rapidly changing the

incentives, behavior and norms of all players in the news media ecosystem.

Headlines evoking a “crisis in the news” and “crisis in journalism” have appeared for more than a decade. To be sure, media scholars and practitioners have described journalism and the news industry in crisis terms time and again in history, mostly after technological shocks. This report does not adopt a crisis narrative. Rather, it soberly reckons with an era of profound change. We believe that changing technologies always warranted updates to the laws and regulations that shape news media. The digital revolution and ascent of dominant digital platforms call for a significant renewal of the rules in this important sphere again.

We demarcate two periods in the digital revolution with regard to its impact on the news ecosystem: the first two decades of the spread of the Internet, and the last decade, characterized by the rise of a handful of digital platforms. Technology had a profound influence on journalism and the news in both periods, but there are important distinctions between the two. The first twenty years after the invention of the World Wide Web saw a dramatic decline in the cost of information distribution and an increase in information accessibility. The news industry had to adjust to advertising and readers shifting to the digital world, which caused a decline in revenues and profitability and the loss of the important business of classified ads. The last decade has seen still more advances in technology, but with a growing share of digital activity and news consumption moving to digital platforms. This decade has been characterized by a

profound influence of the platforms on the relationship between news producers and the public and on the very nature of the public sphere.

The introduction of new technologies to mass media has always had an influence on the character of news, politics and society. From the printing press and telegraph to radio, television and cable television, each technology brought opportunities and challenges and in turn public demand for new laws and regulations. The challenges brought by the platforms run deep: unbundling of news products; personalization and targeting tools unprecedented in their sophistication and precision; and atomization of the news media. Together these trends have created a new ecosystem of news consumption, more complicated and fragmented than ever—and most importantly—split into billions of individual “feeds” and “editions” for each user.

The news media and journalism are broad subjects with many categories and definitions. This report is focused mainly on what we think is the most important for the functioning of democracies: accountability and investigative journalism. This type of news gathering, investigation and analysis reveals information that is crucial for readers as citizens, and information that powerful actors like to be concealed. Hence, it produces not only private benefits for the consumer but also positive externalities benefitting society at large.

While the authors of this report do not believe that there was ever a “golden age” of quality, independent journalism that can be revived, we believe that digital platforms present formidable new threats to the news media that market forces, left to their own devices, will not be sufficient. In the report we review some of the main market responses that try to improve the sustainability of independent

journalism, their contributions and their shortcomings.

Our report is based on the assumptions that independent journalism is a crucial pillar of democracy, but that the production of investigative and accountability journalism was always underfunded and underproduced by the market—as original producers of this kind of journalism can at best capture only a small fraction of the benefits to society.

The report identifies four areas of immediate concern to the news media:

1. The gradual decimation of the business model that enabled many news outlets to produce accountability and investigative journalism for decades. Especially acute is the collapse in revenues of local news outlets and the closure of such news outlets across the developed world.
2. The shift in news distribution from the traditional news organization to algorithms controlled by digital platforms and the growing concentration, power and control that a handful of these platforms have as gatekeepers of the news across the globe.
3. The opacity of the algorithms that control news distribution and the lack of publicly available information on news consumption in the platforms’ ecosystem.
4. The weak economic and legal incentives that these powerful gatekeepers of the news have to prioritize quality content and limit false information.

While the threats to quality news ecosystems are significant, this report does not recommend direct intervention in the management of the platforms and their relationship with users and news producers. Because a handful of platforms exercise gatekeeping power over information, regulatory intervention must be very careful not to put even more power in the hands of those platforms. Any state or regulatory intervention should be measured and limited.

Our policy recommendations are limited to topics that directly relate to the news. Yet they should be read together with the policy recommendations of the subcommittee on market structures listing proposals meant to increase competition in the digital world, give users more power and control over their data and limit the market power of the platforms and their ability to entrench their dominant market position. The subcommittee members think that opening platforms for competition through interoperability, giving users ownership of their data, and the potential breakup of platforms may contribute to reducing the gatekeeping power of these platforms and positively impact the type of information that users consume.

The dramatic shift of advertising dollars from traditional news outlets to a handful of digital platforms has many stakeholders in the news industry calling for regulatory intervention to reverse or halt this trend. This report takes a different approach: we do not focus on finding ways to return to a “glorious” past when a larger share of advertising was allocated to traditional news outlets.

The report’s starting point is that the marriage between quality accountability journalism and advertising revenues was always fraught with conflicts of interests, biases, battles for attention and challenges to the autonomy and integrity of news organizations. A large body of research, evidence and surveys documents the influence of advertisers on the agenda, content and framing of reporting, and direct and indirect bias, censorship and self-censorship caused by dependence on advertising. There is also evidence of biases and distortions in news reporting in the pre-platform era caused by ownership and control of news outlets by tycoons, oligarchs and politically connected business groups. The shift of readers to the

Internet and the rise of digital platforms have exacerbated these biases as the business model of many news outlets collapsed; publishers became more dependent on a few large advertisers, and newsrooms were presented for the first time with granular real-time data on the virality of single stories—which enabled them to adopt editorial strategies that market single stories instead of full editions.

Reversing the shift of advertising dollars from the digital platforms back to traditional media may not only prove to be like swimming against the stream—it may further incentivize news outlets to chase clicks and virality. In the race to get more clicks and exposure through the sophisticated, targeted, personalized, advertising-maximizing digital platforms, publishers may give the platforms more power and editorial sway in the curation of the news. Nevertheless, they will always trail behind platforms in the competition to monetize those clicks, as they will find it difficult to compete against the vast data troves and artificial intelligence capabilities held by giant tech companies.

This report sees the seismic shift in the advertising dollars to the online world as an opportunity to create a news ecosystem supported more by paid subscriptions and public funding, and less by advertising. The report does not seek to protect, subsidize or prioritize existing news outlets, but it does assume that journalists will continue to play a central role in production of accountability journalism.

Our main policy recommendations are as follows:

1. Introducing some public funding for news organizations, relying on citizen choice, to support journalism. The allocation mechanism of the funds should be designed to promote competition and entry and limit the entrenchment of incumbent large news media outlets. The funds should be

allocated directly by the citizens, independently of any government intervention. Special consideration should be given to the funding of local journalism, where we see most of the aforementioned problems concentrated today. This funding mechanism is highly cost effective: \$50 per US adult is likely to be sufficient.

2. All mergers and acquisitions involving news companies should be subject, in addition to the standard antitrust review, to a news plurality review. Standard competition policy protects direct consumer welfare, and therefore does not take into account the indirect effect that excessive media concentration can cause on citizen welfare. We propose an approach to quantifying news plurality that is neutral to the identities of the owners of the merging entities and to the platform on which news content is delivered. The proposed approach, based on attention shares, has been used in a recent merger decision in the UK.
3. Developing a new regulatory system that will ensure necessary transparency regarding information flows and algorithms. This can be done through a new regulatory framework and oversight body that sets standards for the disclosure of information and news sources, develop source-based reputational mechanisms and bring to light biases and choices in editorial decisions and algorithms for the presentation of the news. These regulators should produce periodical reports on news consumption and the influence of algorithm

design on the distribution of news and the behavior of users.

4. Digital platforms enjoy a hidden subsidy worth billions of dollars by being exempted from any liability for most of the speech on their platforms (Section 230). We do not propose to repeal Section 230 but rather propose that platforms that would like to enjoy this protection should have to agree to take clear measures to prioritize content according to criteria other than the maximization of ad revenue.

The pace of change brought by the Internet is unlike any previous technological shock. The proposals in this report aim to address the main threats we see today to the news media ecosystem, but are far from offering complete solutions to an ecosystem that is changing every year. We believe that after rolling out the main policy recommendations above—additional public funding of journalism, disclosing the vast data that platforms have on news consumption, and taking steps to limit excess concentration of political power by tech and media players—experts, regulators and legislators will be equipped with much more information that will enable us to consider further updates to the regulations governing the news media.

Amid growing threats to democratic values and institutions across many liberal democracies around the globe, a bold plan for strengthening independent, strong and rigorous accountability journalism is needed more than at any time since the dawn of the modern liberal democracies.

Report from the Media Subcommittee

Were it left to me to decide whether we should have a government without newspapers, or newspapers without a government, I should not hesitate a moment to prefer the latter.
—Thomas Jefferson

Nothing but a newspaper can drop the same thought into a thousand minds at the same moment. . . . To suppose that they only serve to protect freedom would be to diminish their importance: they maintain civilization.
—Alexis de Tocqueville

There is not a crime, there is not a dodge, there is not a trick, there is not a swindle, there is not a vice which does not live by secrecy.
—Joseph Pulitzer

Journalism is printing what someone else doesn't want printed: everything else is public relations
—George Orwell (attributed)

What is a democracy? The fundamental principle of our modern political system is “one person, one vote.” We believe it should be “one *informed* person, one vote.” Hence free, unbiased, high-quality information is indispensable to democratic debate, institutions and processes. It matters for the quality of elections and the accountability of elected representatives. Journalism, by revealing previously undisclosed information, plays a crucial role in combating and reducing corruption and holding the powerful to account, and is also central to the proper functioning of markets and governance of firms.

With the rise of the Internet, information has become more accessible to the public around the world. The Internet gave voice to hundreds of millions of people and enabled them to connect, come together and form digital communities and networks to express their shared interests. However, as accessing information has become easier for the public, there has been an explosion in

information, and organizing and filtering it has emerged as a major challenge. Up until a decade ago, it was mostly agreed that the benefits of the digital revolution were significantly higher than the negative impact. But with the rise of the dominance of the digital platforms, we are gradually shifting to a network architecture that consolidates much of the power, activity and resources on the Internet in a handful of platforms—a situation that calls for rethinking the rules of the game in the news media and on the Internet.

Production of high quality news with journalistic rigor has always been costly. While there are large public benefits from news production, the private benefit for news producers has been limited. Traditionally, journalists at news media organizations did the job of producing original reporting, and editors bundled this information into news editions for the public. Classified and display advertising and subscriptions were the primary ways of sustaining the traditional business model. After their rise, digital platforms such as Google and Facebook emerged as organizers and bundlers

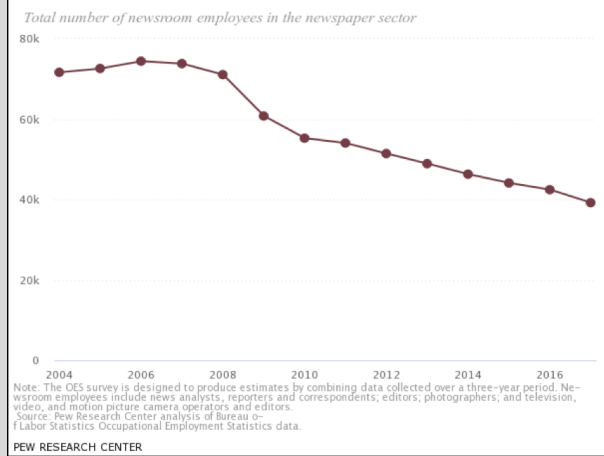
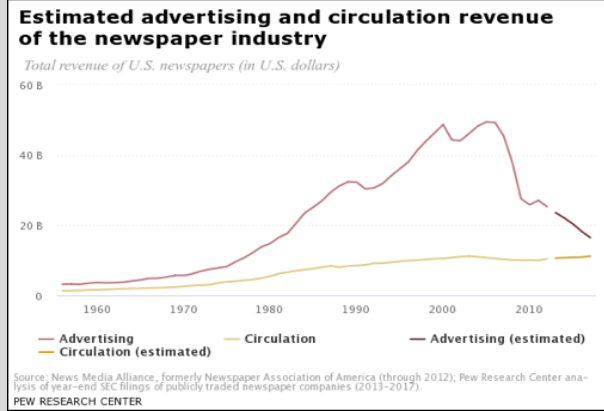
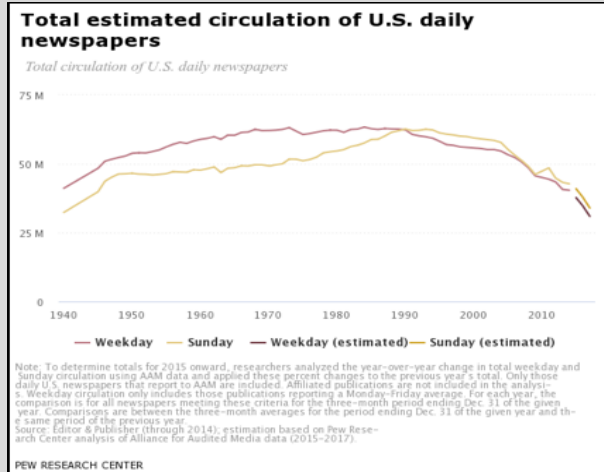
of information. They aggregate content from original information producers such as news media companies and bundle this information as curated “feeds” and search results to users. A rising number of users, especially those who are young, get most of their news directly from social media feeds. The disintermediation between readers and original news producers has disrupted the way in which news is produced, organized and consumed in the digital age.

Digital platforms, and the Internet more generally, have disrupted the advertising market for media outlets. Traditionally, classified and display advertising was the major source of revenue and profitability for newspapers. In the 1990s, the entry of online marketplaces like Craigslist in the United States was disruptive to the classified advertising market and led to an increase in subscription prices of newspapers and a decline in their readership.²⁴¹ Likewise, with the loss of classified advertising and the associated increased reliance on print advertising, there was an increased bias in news reporting toward these advertisers.²⁴² Separately, there has been a reduction in demand for print newspapers among readers due to the availability of digital news and the consumption of news via social media platforms. That this is especially true among young individuals suggests that this trend will only accelerate in future years.

Box 1: Decline of Newspapers in Numbers

The overall circulation of newspapers has declined since the 1990s with the rise of the Internet in the US. Advertising revenue especially plummeted after 2008, when the fallout from the financial crisis and the rise of digital platforms coincided. The industry did not recover from this decline in advertising revenue, and growth in

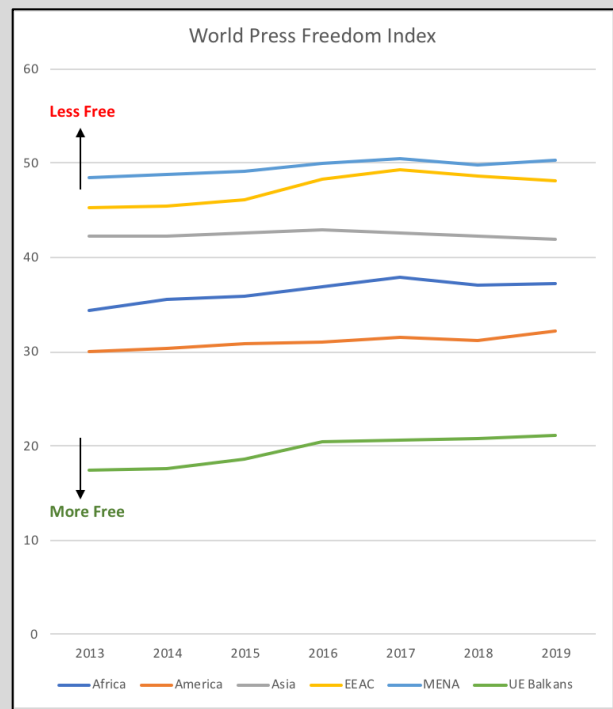
circulation revenue was too slow to arrest overall revenue decline. A decline in revenues and circulation has led to a loss of newsroom employment, a critical measure of journalistic depth that has declined in the US over time.



²⁴¹ Seamans & Zhu (2013).

²⁴² Beattie et al. (2018).

Not only have newspapers declined; press freedom is also under threat around the world, as indexed by the World Press Freedom Index.



Taken together, the reduction in demand on both sides of this market—from paying readers and from advertisers—has put severe financial pressures on traditional media outlets. Since 2007, while digital platforms like Facebook and Google grew exponentially, the advertising revenue of newspapers has dwindled, leading to severe financial strain and a sharp decline in newsroom employment. This disruption of the financing model might also have changed the mix of local versus national news. Given that digital platforms (Facebook, Twitter, Google, etc.) have national distribution, there is a concern that local news has been crowded out in favor of national news. This is consistent with the conjecture that the Internet, and communications technology more generally, may lead to a “death of distance,” that is, the

overcoming of physical distance and a new ability to connect more isolated areas to less isolated ones. This is consistent with evidence the other settings: The introduction of television in the US led to a reduction in local newspaper circulation, and the entry of the *New York Times* into metro areas led to reductions in local newspaper readership.²⁴³

These issues are particularly acute for newspapers that have traditionally served local markets. Since 2004, 1,800 papers have closed in the US. Six percent of US counties currently have no newspapers, and an additional 46 percent have only one newspaper, usually a weekly. Over one-half of counties are not served by a daily newspaper. A similar trend is seen in democracies like Australia, where the number of journalists in traditional print industries fell by 20 percent from 2014 to 2017, and among regional publishers and broadcasters cost-reduction measures range from the closure of newspapers to the consolidation of broadcasting operations. Similarly, in the UK, 321 local press have seen closure in the last ten years.²⁴⁴

This decline in the number of newspapers has reduced the degree of competition in local news markets, in terms of both readership and advertising. This is in contrast to the promise of digital platforms, which had the potential to reduce barriers to entry and facilitate the sharing of information. Instead, by disrupting newspaper advertising markets and shifting demand from print to digital sources, platforms have reduced pluralism and increased concentration in local newspaper markets.

The loss of local newspapers and the emergence of “news deserts” has important consequences for local governance. For example, local newspaper closures between

²⁴³ Gentzkow (2006); George & Waldfogel (2006).

²⁴⁴ <http://newspaperownership.com>; Wilding et al. (2018); Cairncross (2019).

1996 and 2015 in the US led to higher borrowing costs for municipalities in the long-run, even in localities with high Internet usage, as local governments were held less accountable for their public financing decisions.²⁴⁵ Similarly, a study of newspapers in California found that when there are fewer reporters who cover an area, fewer people run for mayor, and fewer people vote.²⁴⁶ In other words, a decline in local journalism due to the emergence of digital platforms can have far-reaching consequences for politics and the economy. Again, these findings are in line with evidence from other settings. The introduction of television in the US, for example, led to a reduction in political knowledge and voter turnout.²⁴⁷ Similarly, increased newspaper coverage of local Congressional representatives is associated with better informed constituents and enhanced representation.²⁴⁸

Taken together, the demise of local newspapers, along with evidence on their social benefits, suggests significant challenges for local governance in coming years. While the aforementioned evidence is not directly linked to digital platforms, there is also some direct evidence that the entry of the Internet and digital platforms has displaced traditional media outlets and the associated news coverage, including investigative journalism, and changed political outcomes.²⁴⁹ For example, the rollout of the Internet in Germany led to a reduction in voter turnout, and researchers²⁵⁰ attribute this effect to a reduction in television viewership following broadband Internet availability. Studying the rollout of broadband Internet in Italy, researchers²⁵¹ documented an initial reduction in voter turnout followed by a later increase as parties

harnessed this new technology. In a recent study, researchers²⁵² have found that broadband development in the UK has displaced other traditional media with greater news content such as radio and newspapers, and has also decreased voter turnout. This effect, which is most pronounced among the less educated and the young, also leads to lower local government expenditures and taxes, particularly expenditures targeted at less-educated voters. Taken together, emerging evidence suggests that the entry of digital media has displaced, rather than enhanced, traditional news reporting in these areas, leading to reductions in voter turnout and changes in policy outcomes.

While digital platforms' dominance is a relatively new phenomenon of the last decade, it is important to remember that news media have long been ridden with market and non-market forces that subverted and biased their reporting. Political parties, governments, advertisers, large corporations, funders, and audiences are a few of the forces that influenced news media. At the cusp of the twentieth century, advances in technology gave rise to greater independence of news media, as they were able to produce and transmit news at a much lower cost, and led to a shift from partisan to professional journalism.

With the rise of digital platforms, the cost of distributing information went down, which increased the entry of new information producers and increased diversity of voices. Digital platforms became powerful intermediaries between original information producers and readers and unbundled and "atomized" news. They use algorithms to curate personalized content for users based on

²⁴⁵ Gao, Lee & Murphy (2018).

²⁴⁶ Rubado & Jennings (2019).

²⁴⁷ Gentzkow (2006).

²⁴⁸ Snyder & Stromberg (2010).

²⁴⁹ Falck, Gold & Heblich (2014); Gavazza, Nardotto & Valletti (2018).

²⁵⁰ Falck et al. (2014).

²⁵¹ Campante, Durante, & Sobbrío (2017).

²⁵² Gavazza et al. (2018).

the unprecedented level of data these platforms have over each individual’s private history and preferences. With billions of users, they can thus create billions of “bundles” or editions of news stories with the objective of maximizing advertising revenues. These algorithms are opaque, and while platforms know exactly which individuals are exposed to which stories and why, publishers and the public have very little knowledge about such information. This gives rise to a huge asymmetry between the data and knowledge that the platforms and the public have on news consumption.

Along with atomization of news, platforms have—through their sheer size and market share as curators and aggregators—acquired unprecedented gatekeeping power over news media outlets all over the world, wielding a huge influence on the version of reality that readers see.


As noted above, traditional media were always plagued with biases, and chased attention to get advertising dollars. But those incentives were disciplined partially by reputation concerns, professional norms, and legal liabilities. Digital platforms, in contrast, are not disciplined by such forces. They are protected from most legal liabilities, and their reputation is not tied directly to the content they present as feeds to their users.

Recent research and multiple investigations by news organizations²⁵³ support the assertion that platforms do not have incentives to prioritize quality content. A recent study found that disinformation can spread faster than true news on social media such as Twitter.²⁵⁴ Not only are users not good at distinguishing reliable and unreliable news; digital platforms at the same time have access to private information about users, enabling them to selectively target visceral, addictive

and at times extremist content. This, coupled with the fact that digital platforms are not held accountable for the published content, has made digital platforms powerful tools of influence, having insight into people’s private behavior, but enjoying immunity from any consequences.

Box 2: Rise of Digital Platforms in Numbers

The revenue growth of digital media is nothing short of spectacular, as evidenced by the market valuations of companies like Alphabet. The share of advertising attributable to digital advertising has roughly doubled since 2010. In 2018, the share of digital advertising (38 percent) was higher than the advertising shares of television (34 percent) and newspapers and magazines (12.2 percent), and it is projected to keep growing in the coming years. Within the digital platforms, advertising revenue is highly concentrated, with two companies controlling over half of it. In particular, in 2018, Google had a 37.1 percent share and Facebook had a 20.6 percent share. Assuming that these trends continue, as predicted, the degree of concentration in advertising markets will dramatically increase in the coming years.



253

<https://www.theverge.com/interface/2019/4/3/18293293/youtube-extremism-criticism-bloomberg>

254 Cagé & Mazoyer (2019).

News Consumption by Media

How many people get their political news from a digital source? The table below shows that the share of Americans who report regular use of the digital platform is large and increasing. Similar patterns are observed in virtually every democratic country in the world. However, the table also shows that television is still the dominant platform for news. These aggregate shares also hide enormous heterogeneity. Television is four times more popular among older adults than among younger adults, and social media consumption is much more popular among younger adults.

Share of US Adults Who Get News Often on Each Platform

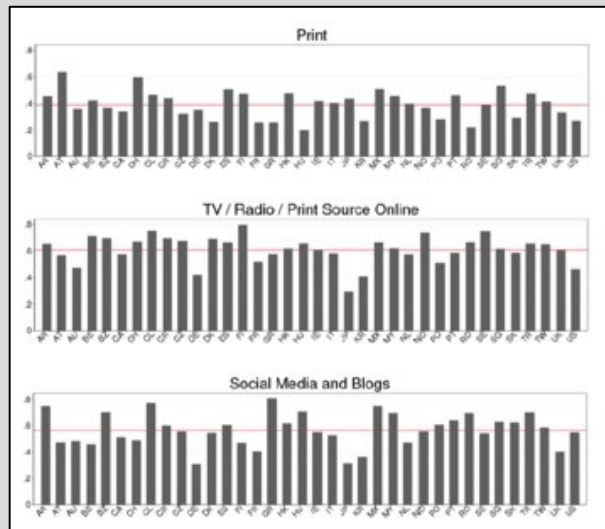

	TV	News website	Radio	Social Media	Print Newspapers
2016	57%	28%	25%	18%	20%
2018	49%	33%	26%	20%	16%
2018 age 18-29	16%	27%	13%	36%	2%
2018 age 50-64	65%	28%	28%	14%	18%

Source: Pew Research Center

News Consumption by Source

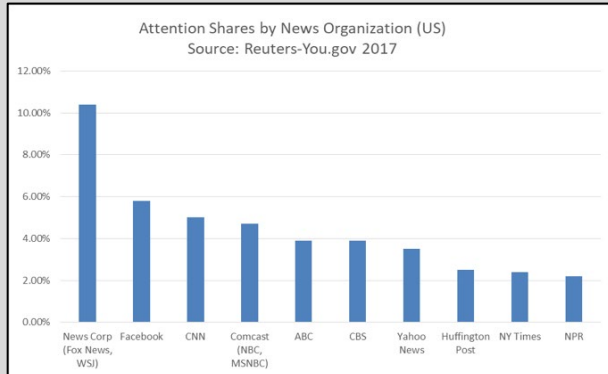
A robust pattern observed in data is that a large share of digital news is actually the online version of traditional media, highlighting that news producers still continue to work in traditional media. The table below covers 36 countries. It shows penetration shares for traditional channels (TV, radio, print) as well as digital media disaggregated by whether the user is viewing a pure Internet source (e.g.,

Huffington Post), a social media or blog platform (e.g., Facebook) or the online version of a TV, radio or print source. As we can see the last modality is more prevalent in most countries.



Three pure digital platforms appear among the top ten US news providers in terms of attention share: Facebook at #2, Yahoo News at #7, and *Huffington Post* at #8. Only the last one produces original content. In the UK, three digital platforms are found among the top ten news organizations: Facebook at #3, Google at #5, and Twitter at #10. Indeed, similar patterns emerge in the 36 countries for

which data are available.²⁵⁵ Facebook is by far the dominant pure-digital news source, although it is critical to highlight that Facebook is an aggregator and not a producer of original content. It is among the top three in 14 of those countries.



Summary of Findings

The media landscape is fast-changing. Available evidence indicates the following patterns:

- The share of advertising revenues going to digital platforms is large and increasing. Facebook and Google receive over half of it.
- Television is still the dominant news medium, though the role of digital news is increasing and already dominant in younger generations.
- Although news may be delivered digitally, the content is most likely to come from traditional print and television providers.
- Among the pure-digital news providers, Facebook is by far the dominant player.

The concern with harmful externalities of concentration of power and biases in the news media related with the emergence of new technology is by no means a new phenomenon. Throughout history, such concerns have been

answered with new regulations. For example, with the development of radio, the Federal Radio Commission was founded in 1927 in the US, which evolved into the Federal Communications Commission (FCC) in 1934. As media technology evolved, FCC regulations evolved, too, including the 1941 National TV Ownership Rule, the 1970 Radio/TV Cross-Ownership Restriction, and the 1975 Newspaper/Broadcast Cross-Ownership Prohibition. These regulations attempted to prevent the concentration of ownership of news media in order to support diversity in the market for ideas. As Internet and digital platforms have disrupted the media industry, we have reason to believe that government should again look into the ways in which the negative externalities of the media can be constrained.

The influence of the digital platforms on the news media has been under increased scrutiny and focus since the last US presidential elections, yet much of the focus has been on fake news and the interference of foreign governments in elections through such platforms. But the influence of digital platforms on the news and journalism ecosystem goes much deeper than just the spread of fake news. The business model of news has been severely disrupted by the rise of digital platforms, and news production and consumption have been disintermediated. The business model disruption has reduced the incentive to produce original reporting, and the platform algorithms have rewarded the production of visceral and emotive content.

While there is a great deal of data on the decline of revenues, profits and number of journalists employed by news organizations, the potential public harms from the new news and journalism ecosystems are inflicted on very large and dispersed groups, and they are much more difficult to analyze and measure. This

²⁵⁵ Kennedy & Prat (forthcoming).

report will focus on these potential harms and recommend ways to develop a more sustainable and competitive economic model for a news media ecosystem that produces quality journalism. These recommendations include a new way to increase funding and competition in the news market as well as the transparency and accountability of digital platforms.

The Precarious Economics of News

News is an information good—perhaps the purest form of it. Whether it be day-to-day decisions regarding what products to buy, how to manage health, how to prepare for the weather, or political decisions regarding whom to vote for or whether to attend a protest, the news provides information that allows people to make those decisions facing less uncertainty. However, information goods, especially news content such as investigative journalism, have some unique characteristics that give rise to underproduction of news and lower demand. In this section, we explain the unique economics of newsworthy information.

Information, once produced, can be consumed widely without constraint, making it non-rivalrous. Unlike a physical good, when one agent consumes information this does not prevent another agent from doing the same. Also, information, once disseminated, can be distributed by agents other than the agent responsible for its production, making it non-excludable. Both these characteristics of information give rise to underproduction of information, as the producer of information, who incurs the fixed cost of producing news, is unable to accrue the full benefit from producing it, as it is easy to copy and share.

Another important characteristic of information that arises from the demand side is

uncertainty. Almost by definition, information is a good whose value is not necessarily known to the consumer at the time of purchase. If information is revealed to the purchaser, which is what efficient purchases of a good would entail, the purchaser has no incentive to actually pay for the information once disclosed. In such a case the information producer has to resort to trying to sell information without disclosure. What this implies is that the demand for information will be lower than it would otherwise be and will not reflect the value consumers actually place on that information, which will also drive diminished returns for information producers. However, this also means that the production of many information goods will depend on finding means of payment—such as advertising in case of news—that do not involve direct payments from consumers themselves.

Box 3: Kenneth Arrow on Economics of Information

The Nobel prizewinning economist Kenneth Arrow identified indivisibility and inappropriability as characteristics of information that would lead to insufficient production. These generate the challenge of ensuring that the returns to those responsible for producing information (that is, those bearing the costs of production) are closer to the social return for information production. In modern parlance, this is often broken down into two dimensions of the public nature of goods under the terms *non-rivalry* and *non-excludability*.

For information with value that cannot be so easily inferred from past experience with the information provider, there is a special challenge.

Arrow (1962) put it this way:

[T]here is a fundamental paradox in the determination of demand for information;

its value for the purchaser is not known until he has the information, but then he has in effect acquired it without cost.²⁵⁶ Of course, if the seller can retain property rights in the use of the information, this would be no problem, but given incomplete appropriability, the potential buyer will base his decision to purchase information on less than optimal criteria. He may act, for example, on the average value of information in that class as revealed by past experience. If any particular item of information has differing value for different economic agents, this procedure will lead both to a nonoptimal purchase of information at any given price and also to a nonoptimal allocation of the information purchased.

Undersupply and underdemand of information may be more acute for certain type of news. Some news—such as that related to weather, traffic conditions, impending or actual disasters, product reviews, or scientific breakthroughs—is of primary use for individual decision-making. By contrast, other news is of primary use for social decision-making, including how to vote, how to protest and whether to avoid or support particular businesses or organizations. In this situation, the decisions one person is making are part of a collective decision-making process, and, hence, each person will be interested and place value on others who are participating in the decision having access to that news. Apart from the private versus social dimension of the news, news also differs in terms of its timeliness or, more critically, its longevity, that is, how close in time it is to the moment a relevant decision has to be made. This is most obvious with respect to weather and traffic decisions, but may also be of importance for protests or disasters. By contrast, news about corruption or poor policy-making might be of use for the next

election, and, therefore, its value does not necessarily depend on its timely provision.

Consider a case of investigative journalism that exposes the corruption of a government official in a particular county. Such news has limited private benefit for readers, and very few users are likely to buy this news story. Yet, the public benefit from exposing this story is large, as it not only exposes a corrupt official but at the same time creates a deterrent for corruption in the future as officials fear similar exposés. Hence, investigative journalism is a public good with limited private benefit. Such investigative journalism is costly to produce and delivers limited immediate private benefits for news media outlets.²⁵⁷

There are good reasons to believe that the economic issues associated with the supply and demand for information become stronger as news moves from the private to the social and, perhaps paradoxically, from being of immediate to longer-term use.

A look at the case of investigative journalism helps reveal how the production of such journalism is acutely ridden with issues of underdemand and undersupply. Investigative journalism is the provision of news as the result of a long, complex and often very costly investigation. As a result, it is unlikely to be the kind of news that requires quick action but is released well ahead of a decision point. In addition, it is more likely to involve a social element, whereby the consumption of that news by others raises its value to individuals. For example, an accusation of sexual misconduct against a powerful figure is likely to have a greater impact if it is widely read. Additionally, the longevity of news represents a problem for private provision, as there is time for that news to be provided by others. To see this, imagine that the output of investigative

²⁵⁶ See Roth (2002), and Gans & Stern (2010).

²⁵⁷ Hamilton (2016).

journalism is a ‘scoop’—a news output that others do not have. A news outlet might publish a scoop and for a period of time may be the only outlet with that news. If the news has a short half-life, then that period would be valuable in that consumers would have an incentive to consume news in the originating outlet. However, if the news is long-lived, there is no similar time pressure. People will be able to consume the news more easily as it becomes widely reported. Even if this is done with attribution to the original outlet, it is not clear there is any mechanism by which that outlet will benefit disproportionately in terms of consumers willing to consume the news on that originating outlet. Thus, there is a clear production externality for such news caused by the knowledge spillover of that news to other outlets. Consequently, there is limited incentive to become an outlet that is able to invest in generating scoops of this kind.

Box 4: Biases in News

While this report and many other reports highlight the various potential harms of digital platforms to the business model, distribution and consumption of news, and stress the important contribution of journalism to democracy, it is important to stress that there was no agreed “golden age” of journalism, and it was always subject to multiple forms of capture and biases. Some sources of bias in media are:

- Bias toward political or corporate owners
- Bias toward funders
- Bias toward advertisers
- Bias toward newsmakers who provide access

- Bias toward slant of the audience
- Bias toward a particular ideology
- Bias toward sensational news that boosts ratings

An example of media capture could be due to banks. Recent research²⁵⁸ showed that Italian newspapers that were more indebted were more more likely to agree with banks, whatever the interest of banks may be. In other words, as newspapers become less financially healthy, the more likely they are to be riddled with biases. This trend is also confirmed in historical research²⁵⁹, where researchers find that, in the absence of a thriving subscription and advertising base, US newspapers were more likely to be politically captured in the 19th century. Similarly, in Argentina, newspapers that carried more government advertising covered government corruption scandals less between 1998 and 2007, showing evidence of advertiser bias.²⁶⁰

Another challenge that news presents that makes it distinctive in terms of the economics of information is that there is always potentially an interested party who will have some control over the news. Thus, news can be pursued independently up to a point, but there is always some area where conflict is to be expected. The only way around this is to have a diversity of news outlets with a consequent diversity of operating interests.

The above challenges of social news and special interest are especially severe in the case of investigative journalism. Coupled with the fact that such journalism—due to its non-rivalrous, non-excludable and uncertain nature—is subject to underdemand and undersupply, we conclude that production and

²⁵⁸ Zingales (2016).

²⁵⁹ Gentzkow, Glaeser & Goldin (2006); Petrova (2011).

²⁶⁰ Di Tella & Franceschelli (2011).

consumption of investigative journalism-type content face many challenges.

Box 5: Why is investigative journalism difficult to produce?

Investigative journalism has large public benefits, all of which cannot be captured by producers, but its production incurs *high fixed costs* (e.g., careful collection of evidence and analysis) with a long gestation period, which may lead to dead ends and failure. Such journalism is also *hazardous*, and prone to lawsuits, and may also antagonize stakeholders (advertisers, owners, newsmakers etc.).²⁶¹ Finally, due to its public nature, investigative journalism has uncertain benefits. It has been suggested that a key benefit from successful investigative journalism stories tends to be of reputation, and that “investigative journalism is like haute couture It isn’t highly profitable per se, but it helps create brand awareness and it excites the most talented designers... [and is done] when margins are high.”²⁶²

Overall, the production of investigative journalism is like taking a risky bet with few upsides and many downsides. News media firms may take such risky bets to produce investigative coverage only if they have the capacity to produce such “haute couture” content, the power to bundle and distribute such news, to best monetize it, and the incentive to differentiate from the competition.²⁶³

Digital Platforms and Disruption of News

In this section, we look at the business model of news and discuss how the Internet and the digital platforms have influenced the business model of news at its various stages of production, distribution and consumption of news.

A. Production of News

1. Reducing the Incentive to Produce Original Content

As consumers move principally to online consumption, it is becoming easier to “steal content” from competitors. This affects media outlets’ incentives to produce high-quality (and costly) news content in the Internet era. Recent studies of audience news consumption behavior have indicated that news users increasingly rely on multiple news media and seem to shop for the best news across outlets online.²⁶⁴ As a consequence, they follow the news on multiple media platforms.²⁶⁵ It has been well-documented that the Internet has reduced loyalty to any single outlet, in particular for technological reasons.²⁶⁶ Revealing is the fact that online when coming to a news website through search or social media, most users cannot recall the name of the website’s news brand after their visit.²⁶⁷ According to Reuters data, in France in 2018, consumers of at least one offline media outlet consume on average 2.83 outlets online.²⁶⁸

News in online media is not only copied by many, but it is also copied fast. An analysis of French media showed that on average news was delivered to readers of different media

²⁶¹ Hamilton (2016).

²⁶² Zingales & Rolnik (n.d.).

²⁶³ Raj & Rolnik (2018)

²⁶⁴ Athey, Calvano, & Gans (2013).

²⁶⁵ Picone, Courtois, & Paulussen (2015); Yuan (2011).

²⁶⁶ Athey, Calvano, & Gans (2013).

²⁶⁷ Reuters Institute (2017).

²⁶⁸ Reuters Institute (2018).

outlets in less than 4 minutes in 25 percent of the cases. Also, the analysis found that such high reactivity came with high verbatim copying, as only 32.6 percent of the online content was original.²⁶⁹ Such a scale of copying online might potentially negatively affect media outlets' newsgathering incentives, as original news producers would capture only a fraction of the audience and of the economic returns to original news production.²⁷⁰

While a reactive online media reduces the incentives of news producers to invest in original content, in the long run, producers can gather reputation effects, whereby users do indeed share content from the original content producer more frequently. However, media outlets with a larger fraction of original content are still losing part of the audience they would receive absent copying and consumers' switching across outlets. Furthermore, this negative effect is accentuated by the impact of the platforms on the advertising markets for news media: when attention is spread across publishers, switching consumers actually see fewer ads than their loyal counterparts on a given publisher.

2. Platform Duopoly and the Business Model of the News Media

Consumers increasingly consume information on news aggregators such as Google News or Yahoo News, and there is a debate whether these aggregators act on the consumption side as a complement, bringing additional traffic to traditional media outlets online, or as a substitute, stealing the audience for these outlets, and then negatively affecting their news production incentives. Empirical evidence seems to indicate that news aggregators act as a complement on the consumption side. For example, analysis using a shutdown of Google News in Spain in

December 2014–January 2015 as a natural experiment found that the removal of Google News reduced overall news consumption by about 20 percent for users affected by the shutdown, and visits to news publishers declined by about 10 percent, a negative shock that particularly affected small publishers.²⁷¹ In other words, Google News seems to act as a complement rather than as a substitute, at least for small publishers.

However, even if the “pure aggregators” act as a complement to the traditional media outlets on the consumption side, they may hurt media outlets' incentives to produce original content through their negative impact on the advertising market. Here one may think of the digital platforms all together (i.e., not only Google and Yahoo but also social media platforms such as Facebook and Twitter) and more broadly of all the tech giants rather than just the aggregators. Indeed, with the rise of the digital platforms, the supply of available ad space online has increased far more rapidly than the demand for it, owing mainly to advertising on digital platforms, so that the price has dropped precipitously. As a consequence, traditional media are devoting more and more space for online ads, but are winning fewer and fewer of them. In 2018, Google and Facebook were the dominant digital advertising companies, with a combined 58 percent of the US market, followed by Amazon, whose advertising business is expanding quickly. According to the latest estimates from eMarketer, by 2020, Amazon will have captured a 7 percent share of US digital ad spending, compared with Facebook's 20.8 percent and Google's 35.1 percent.²⁷²

The growth in digital ad space is far from being shared equally among players in the online advertising ecosystem. Google and Facebook act as a digital duopoly that

²⁶⁹ Cagé, Hervé, & Viaud (2017).

²⁷⁰ Anderson (2012).

²⁷¹ Joan & Gil (2016).

²⁷² Anderson (2012).

represented up to 85 percent of all digital advertising growth in 2016. According to the European Audiovisual Observatory (2017) and a number of other studies, this digital duopoly even represented all digital advertising growth in the United States by capturing 99 percent of digital ad growth in 2016, and up to 92 percent in France. That is, the share of the digital advertising growth left for traditional news media is nearly zero (and even negative for some news media).

This may even become worse in the future as a consequence of the European Union's General Data Protection Regulation (GDPR). The GDPR will reduce traditional media outlets' ability to collect data online—to protect consumers' privacy, the GDPR requires marketers to secure explicit permission for data-use activities—and thus their capacity to create targeted online advertising. But platforms such as Facebook will continue to collect tons of personal information on their users (with a monopoly on these data) and so become even more competitive on the targeted online advertising market.

B. Distribution of News

Once it has been created, news content needs to be distributed to create value. The fundamental issue that shapes this activity is the fact that, over any time interval, consumer attention is limited. Thus, even if news content is freely available, only a fraction of it will be consumed.

The ascent of digital platforms has negatively impacted the news distribution model. The fact that nowadays the vast majority of consumers prefer to get to news through social platforms and search, rather than going directly to a news website, has been well documented. But given that these access points are limited—again Facebook and Google are here in a nearly duopolistic situation—this gives them market power vis-à-vis the online

news media and more generally vis-à-vis all the newsrooms. De facto, the digital platforms aim at dictating the terms of distribution and all dealings with the news media. News publishers have lost control over distribution; the news is increasingly filtered through algorithms and platforms that are opaque and unpredictable. In particular, each change in Facebook's algorithm has a huge effect on the size of the news websites' audience (implying changes in their revenues). This has been particularly striking in recent years when Facebook has decided to reduce exposure to news, instead prioritizing interactions with family and friends, and leading to a huge drop in the traffic from Facebook to news publishers' websites.

1. Disaggregation of the Customer

Advertising played an important role in funding traditional news media. Prior to digitization, advertisements would be placed physically in the newspaper or intermittently on television and radio. In equilibrium, advertisers and outlets would come to understand the make-up of consumers and be able to adjust advertising content accordingly. Moreover, there was a sense in which a newspaper or a program could result in the bundling of attention on a regular basis—that is, news consumers might read the paper or watch the nightly news every day. Thus, an advertiser looking to place ads in front of those consumers would know precisely where to find them. This assisted in making each ad more valuable, and the advertiser and outlet would benefit and divide value from such matches.

In principle, digitization would not change anything with regard to this type of product. Indeed, as it became possible to know even more about individual consumers (through data collected about them through, say, their browsing and click behavior), the ability of outlets to match consumers and

advertisers should have been enhanced. In non-news related advertising such as search, this promise of more efficient matching was brought about. However, in news related advertising, digitization brought more choice for consumers. In other words, consumers split their attention across outlets by a substantially increased degree and, moreover, would not necessarily follow the editorialized priority for content on the same outlet—picking and choosing what they wanted to pay attention to rather than passively accepting the “flow” of content chosen by editors.

While such fragmentation of attention was a natural and efficient response for consumers, it also meant that the advertising product that outlets were selling became far less straightforward and, instead, the issue of how to put an ad in front of particular types of consumers potentially became harder rather than easier to address. In effect, while before a sales department of a news outlet could tell advertisers about the consumers that it, almost exclusively, could bring to them, with the fragmentation of attention, that sales pitch involved consumers that may also appear in the sales pitch of other outlets. For advertisers, it became harder to identify when consumer attention might be sold and, moreover, who might be selling it. This combination of lower match quality along with greater competition between outlets at the margin is a potential explanation for the dramatic loss in advertising revenue (even independent of classified ads) that occurred from 2000 to the present day.

The breakdown in the coherence of advertising products around news did, however, invite changes to reconstitute it. In each case, these changes were designed to re-aggregate consumers into bundles of attention that could be described, understood and sold to advertisers. We describe each in turn.

2. Advertising Networks and Attention Aggregation

One of the issues that created difficulties for the advertising product of news outlets was that it was difficult to track users and which ads they had seen (both within but mostly between outlets). Ad networks (such as DoubleClick, which was acquired by Google) were developed that allowed for tracking—at least when consumers used a single browser on a single device—and the promise of such tracking was to ensure that consumers received the “right” number of ads from a given advertiser and were not “missed” or served up too many ads, leading to “wasted” impressions.

Such advertising networks allowed advertising markets to become reorganized in a way that was not outlet-centric. The challenge, however, is that this took away another piece of information useful for matching consumers to ads—that is, the self-selection that comes from consumers deciding which content to devote their attention to. Ad networks are very efficient at matching relatively generic ads with consumers or targeting consumers with ads based on their browsing behavior. However, this happens at a higher degree of abstraction than what might attract them in terms of news. News outlets—especially local ones—may have been better able to match local consumers with local businesses. While that is possible for advertising networks, it is possible that something was lost in the transition.

3. Subscriptions and Attention Aggregation

As the business model of the media is in crisis, with falling advertising revenues and print subscribers, in recent years, news media firms have been transitioning to charging subscription fees for their digital content. While some news outlets are better able to generate subscription revenue than others—in particular, national or global outlets—this has

flow-on effects to the organization of the advertising market. This is because, despite those subscription fees, advertisements continue to be placed in front of subscribers.

Although the free Internet fragmented consumer attention across outlets, when a consumer subscribes to an outlet, it signals that that outlet will grab a higher share of his or her attention. As a result, this makes subscribers' attention a more straightforward product to sell to advertisers. In other words, it can counter that disaggregation that might otherwise occur, while at the same time making subscription and advertising revenue (to a degree at least) positively associated.

4. Social Media and Attention Aggregation

Another way in which attention has been aggregated in a way that makes the advertising product more coherent is social media. Social media has the quality—like the newspapers or nightly news of older times—of managing to regularly and reliably grab a share of consumer attention each day. That means that social media networks can sell advertising products that more consistently match ads and consumers without missed opportunities or waste. This ability of social media to grab the attention of customers gives them an editorial function that curates that news. In other words, the aggregate attention that comes from being able to manage consumer information overload—something that used to be performed by news outlets exclusively—can now be undertaken by these networks and their related aggregators like Google News or Apple News.

5. Market Power in Advertising

The analysis of market power in advertising markets related to news (and potentially other) content has always been made more complicated by the two-sided nature of media

markets. On the one side, outlets attract consumer attention and compete for it. On the other side, they sell that attention to advertisers. The question is: Having obtained some share of consumer attention, if an outlet chose to decrease the price of ad space, would that put pressure on other outlets to do the same?

The traditional answer is no. Having obtained consumer attention, an outlet is essentially a monopolist over reselling that attention to advertisers. In that sense, regardless of the prices they set, it will have no impact and not be impacted by the ad prices set for other outlets. In that sense, outlets have market power in the advertising market; to the extent that generates rents, those rents may be wholly or partially dissipated as those outlets compete for consumer attention.

The traditional answer, however, relies on an assumption that each consumer, over a relevant time period, gives all of her attention to a single outlet (which is called *single-homing*). This, in turn, motivates advertisers to advertise wherever consumers happen to be. However, when consumers fragment their attention (what is called *multi-homing*), this assumption no longer holds, and outlets do not have a monopoly over access to that consumer. In this situation, alongside the matching difficulties mentioned above, each outlet is no longer a monopolist in dealing with each advertiser and thus, outlets compete with one another. In this case, as one outlet lowers its ad price, it will put pressure on others to do the same.

C. Consumption of News

Box 6: Bundling and Architecture of Serendipity

As news is public information, bundling of that information is an important manner in which traditional news outlets have attracted and retained customers. Traditional news outlets sold different news content types as a bundle. To reach a large audience, outlets had an incentive to bundle diverse news. So, a typical newspaper covered content including national, international and local politics, business, sports and page 3 culture, along with classified ads.

The bundling and curating by experts (editors) created additional value, as this curation built an “architecture of serendipity.”²⁷³ Sunstein (2008) notes that, “For good lives, good universities, and good societies, the power of self-sorting is at best a mixed blessing. However unpleasant and jarring they can be, unchosen, unanticipated encounters play a crucial role; they are indispensable not only to education but also to citizenship itself. Far from wishing them away, we should welcome them.”

Curated bundles of news promoted the discovery of news. While customers chose the quality, slant and niche of their newspapers, they received news as bundles, and so were exposed to news that may be “unpleasant,” “jarring,” “unchosen,” or “unanticipated.” Such a system of serendipity limited the degree of self-sorting.

A lab study²⁷⁴ found that news that catches public attention may be biased towards negative or “horserace” related political content. Given such preferences, news bundling helped the spread (and production) of new content such as investigative journalism, which is of public value but not designed to catch attention and go “viral.” Thus, bundling reduced underinvestment in public goods like investigative journalism.²⁷⁵

Platforms have for most of their existence insisted that they are not media companies.²⁷⁶ They have described what they do as offering neutral platforms for connectivity, allowing users to find information of relevance to them. It has now become clear that platforms’ moderation of content creates salience. How they do this—what content platforms promote and what they hide, who is speaking and with what credibility—is not transparent. This opacity works hand in hand with moderation to put people in the flows of content that they cannot assess and cannot escape. The principal method by which platforms create media salience is through their algorithmic design and recommendation engines. Their algorithms are a form of editorial judgment that privileges particular forms and sources of media content.²⁷⁷ In this way, algorithms shape consumption on an individualized basis.²⁷⁸ Platforms also exercise editorial judgment by blocking content. Platform moderation, whether by algorithmic design or by human intervention, whether by prioritizing content or blocking it, is an “essential, constant, and definitional part of what platforms do.”²⁷⁹

²⁷³ Sunstein (2008).

²⁷⁴ Trussler & Soroka (2014).

²⁷⁵ Hamilton (2016).

²⁷⁶ See generally Napoli & Caplan (2017) and Pasquale (2016).

²⁷⁷ Carlson (2018).

²⁷⁸ Caplan & Boyd (2018).

²⁷⁹ Gillespie (2018).

1. Atomization of News

Traditional news media would produce editions of news that would bundle news of multiple types. To attract a large base, such a bundle would offer a variety of content and viewpoints, and also provided editors the ability to bundle stories of public relevance such as local investigative journalism, such as on local corruption, which would not have grabbed readers' attention otherwise if left to compete for attention on its own. In the digital age, news has been atomized, as users often consume content curated by the algorithmic editing by digital platforms.

The editorial power of digital platforms also influences the editorial decisions of news producers. An analysis of an online news dataset obtained from an Indian English daily newspaper showed that editors give more coverage to news stories whose articles receive more clicks and that this effect is quantitatively important.²⁸⁰ Digital platforms and “virality” have become so important to editorial decisions that researchers have found that a number of news stories first originate on social media, and absent their propagation on social media, these stories would never make it to the website of the traditional news publishers.²⁸¹

2. Moderation of News Content

Content filtering

Most Americans encounter a substantial portion of their news media through information platforms. According to the Pew survey, more than 68 percent of American adults get some news from social platforms, with 20 percent doing it often.²⁸² As of the end

of 2016, 45 percent of all traffic to publisher sites came from Facebook. Google was responsible for 31 percent.²⁸³ It may be the case that these numbers may be beginning to decline across the globe, according to Reuters Institute. In some countries, especially authoritarian ones, messaging apps like WhatsApp are becoming more important for news circulation; WhatsApp is used for news by about half of surveyed online users in Malaysia (54 percent) and Brazil (48 percent), and by about one-third in Spain (36 percent) and Turkey (30 percent).²⁸⁴

As discussed above, the dominance of information platforms as a distribution mechanism for news impacts the production side of journalism, in terms of reducing the advertising base to fund journalism and incentivizing news media to produce content that will survive algorithmic sorting.²⁸⁵ These algorithmic filters also influence what news content is consumed.

Like traditional news publishers, platforms are in the business of selling audience attention to advertisers. They are able to do this with unprecedented efficiency by using personal data to promote content predicted to engage users and thereby provide more value to advertisers. Platforms offer advertisers access to the “data exhaust” of individuals as they move in real space and across devices so they can target the most receptive audience segments.²⁸⁶ Online advertising “has evolved rapidly from a digital version of conventional ad placement involving agencies and publishers, to what is now a data-driven market focused on audience segmentation and targeted messaging.”²⁸⁷

²⁸⁰ Sen and Yildirim (2015).

²⁸¹ Cagé & Mazoyer (2019).

²⁸² <http://www.journalism.org/2018/09/10/news-use-across-social-media-platforms-2018/>.

²⁸³ Bell et al. (2017).

²⁸⁴ Reuters Institute (2018).

²⁸⁵ See Bell et al. (2018): 28 (news media companies must devote resources to accommodating the platform algorithmic changes); Marwick & Lewis (2017).

²⁸⁶ Wu (2016), Ghosh & Scott (2018): 13.

²⁸⁷ Ghosh & Scott (2018): 5.

Platform companies feed user data into models that produce an advertising technology platform. Using this platform, advertisers can find narrowly segmented audiences and target them through social media feeds and websites with ads ever more precisely tailored to their perceived personal preferences. Platforms develop their predictive models based on inferences from user data including preferences revealed through past consumption or likes.²⁸⁸ Where advertisers have data, in the form of customer lists or other personal data, the platforms can find audiences that share characteristics and thereby deliver to advertisers what Facebook calls a “lookalike” audience.²⁸⁹ Advertising includes not just product promotions but also paid content. Information producers can use data profiling to target audiences using the same approach as product advertisers. There has been considerable controversy, for example, with disinformation providers using these tools to “deepen engagement with known audience segments and broaden engagement to new ones.”²⁹⁰

Platforms will only be successful in generating user engagement with advertising if they can generate engagement with content. Serving up editorial content, whether user-generated or professionally generated, follows the same logic as serving up advertising. Platforms target individuals with content that will be most engaging based on predictive inferences. The platform may apply content filters based on direct signals from the user, collaborative filters based on the preferences of similar users, or some hybrid of the two.²⁹¹ The

platforms do not disclose how they filter content. Twitter’s “trending topics” are already popular. Facebook News Feed and YouTube’s Suggested Videos seek to predict what will become viral in a user’s network and amplify it with that use.²⁹² Facebook explains that its News Feed algorithm attaches a “relevance score” to content based on predictions about a user’s likelihood to click, likelihood to spend time with the content, likelihood to like, comment and share, likelihood that the user will find the content informative, likelihood that the content is “clickbait”, and likelihood that the content links to a low-quality web page.²⁹³

Users can customize their news feeds on social media platforms to a degree, within the constraints of the algorithmic filters that are applied. Facebook friends and Twitter follows shape content exposure.²⁹⁴ Friends and other influencers people choose are important funnels for what news reaches them.²⁹⁵ Google allows users to customize their Google News settings and subscribe to channels on YouTube. Individual choice, however, can push back only so far against the forces of algorithmic filtering. YouTube autoplay queues up the next video to carry viewers from one video to the next to keep them on the platform. The addictive qualities of social media platforms keep people attached to the flow of content long after they have left the confines of their “selected” content. Social bots are one way that content providers can hack people’s attention to push content on them that they might not have chosen and cannot choose.²⁹⁶

²⁸⁸ Wilding et al. (2018).

²⁸⁹ Ghosh & Scott (2018): 16.

²⁹⁰ Ghosh & Scott (2018): 17).

²⁹¹ Ricci, Rokach, & Shapira (2011).

²⁹² Grimmelmann (2018).

²⁹³

https://www.facebook.com/help/publisher/718033381901819?helpref=faq_content

²⁹⁴ See DeVito (2017) (finding friends on Facebook to be the most important determinant of News Feed choices).

²⁹⁵ Bergström & Jervelycke Belfrage (2018).

²⁹⁶ Shao et al. (2018): (“[B]ots are particularly active in amplifying fake news in the very early spreading moments, ... target influential users ... [and] may disguise their geographic locations.”).

Algorithmic filtering does not care in principle about the quality and type of content it promotes. Relevance and engagement are what it cares about. The theory is that if consumers do not like and will not engage with low-quality information, then presumably they will see less of it and vice versa. There are at least two caveats to this revealed preference theory. One is that the preferences accounted for algorithmically are only revealed preferences, not the higher-order considered preferences of public service media theories. The second is that algorithmic filtering stops offering consumers content that they are not predicted to want. Eli Pariser calls this the *filter bubble*: algorithms drive people into narrower homologous information spaces where the content confirms biases and does not expose them to differences.²⁹⁷ This theory is challenged by other research that shows algorithmic exposure to multiple viewpoints.²⁹⁸

Whether or not algorithmic filtering reduces exposure to alternative viewpoints, it privileges a certain kind of content. This is content that provokes outrage and emotion and tends to extremity. Studies show that filtering algorithms funnel people into more extreme expressions of their particular preferences, including political and cultural ones. Viewpoint amplification encourages engagement.²⁹⁹ With respect to political viewpoints, this tendency seems to be more pronounced on the right than on the left, with the consumption of highly partisan information

asymmetrically concentrated among those with more conservative views.³⁰⁰

There is not always a line between human and algorithmic filtering on the platform. The algorithm is created by humans and changed by humans. For example, in January 2018, Facebook announced changes to its News Feed algorithm to prioritize “meaningful content posted by friends and family over the news, videos and posts from brands.”³⁰¹ As a result, the amount of news shrank from 5 percent to 4 percent of the content on feeds. Facebook also changed its algorithm to prioritize local news.³⁰² We do not have good data on what these tweaks do to news consumption, nor are these changes to be relied upon as long-term strategies. Facebook has altered its strategies before, most notably with respect to its “pivot to video.” News producers put resources into accommodating the new algorithmic strategy, only to see the strategy change again. Not only can platforms like Facebook or Twitter alter media consumption through algorithmic tweaks, but they can also alter behavior by favoring certain messages.³⁰³ The opacity in the system means that we only know about these tweaks when they are disclosed or, rarely, discovered.

Content blocking

While algorithmic sorting prioritizes information, another mechanism blocks it and ensures that it will not be consumed on the platform. Blocking often, but not always, involves human intervention. Platforms moderate content by two means and at two

²⁹⁷ Pariser (2011). See also Sunstein & Vermeule (2009).

²⁹⁸ Bakshy, Messing, & Adamic (2015); Fletcher & Kleis Nielsen (2017) (contesting evidentiary basis for the proposition that online audiences are more polarized than offline audiences).

²⁹⁹ Tufekci (2018).

³⁰⁰ See Guess, Nyhan, & Reifler (2018) (“pattern of selective exposure was heavily concentrated among a

small subset of people—almost six in ten visits to fake news websites came from the 10 percent of Americans with the most conservative information diets”).

³⁰¹ Beckett (2018).

³⁰² <https://newsroom.fb.com/news/2018/01/news-feed-fyi-local-news/>

³⁰³ See Zittrain (2014) (describing how Facebook and Google can alter voter turnout by tweaking news feed and search results).

stages. The means are human or machine. The moments are before and after publication. Before publication, software will block content that can reliably be identified as illegal or otherwise prohibited.³⁰⁴ This kind of automatic blocking is used to prevent the circulation of content that allegedly violates copyright, has been identified as violating local laws (e.g., child pornography), or violates the platform’s terms of service. This form of ex ante content removal is more relevant to user-generated content than to news producers.

Most content moderation takes place after it is posted, and is conducted through a combination of human and machine algorithms. The platform, for the most part, reacts to content users have flagged for review. There is little transparency into how they make these decisions or what the results are.³⁰⁵ “Each social media platform has cobbled together a content moderation labor force consisting of company employees, temporary crowd workers, outsourced review teams, legal and expert consultants, community managers, flaggers, administrators, moderators, super flaggers, nonprofits, activist organizations, and the entire user population.”³⁰⁶ Increasingly, in the wake of public outrage over the use of platforms to incite violence, spread disinformation, recruit terrorists, and otherwise propagate “bad” content, platforms moderate proactively. They remove content and accounts that violate their terms of service without relying on users to tell them to. This moderation is also opaque. Kate Klonick’s research suggests that content moderators adopt traditional analogical reasoning, apply multifactor tests, and conduct balancing. Casey Newton’s investigative reporting has revealed that Facebook content moderators work under

high pressure, often exploitative, conditions that are harmful to their mental health.³⁰⁷ Their decisions about content are guided by Facebook’s public community guidelines, internal supplemental guidance, and episodically updated interpretations in real time that may override that guidance.

Box 7: Trust in the Age of the Internet

The advent of the Internet and consumption of information online changed how people view, understand and trust the information they receive. Old relationships were upended, and traditional journalism’s authority was undermined.

In the absence of the traditional signals of authority, how do audiences gauge trustworthiness? A 2003 study by JD Greer found that they consider whether a site belongs to a person or a well-known outlet, but not whether the advertisements represent reputable organizations.³⁰⁸ Similarly, a 2007 study of user behavior defined two key elements of credibility as being “site” credibility and “sponsor” credibility, and found that respondents trust news sites more than personal sites.³⁰⁹ Lack of transparency³¹⁰ and use of native advertising are said by consumers to make them less trusting of the media.³¹¹

Researchers³¹² have further found that because traditional clues of credibility (bylines, trusted brands) no longer prevail and it is often not clear on aggregation sites where information originated, online articles with direct quotes from named sources were viewed as more credible than those without.

³⁰⁴ Klonick (2018).

³⁰⁵ Klonick (2018); Chen (2017).

³⁰⁶ Gillespie (2018).

³⁰⁷ Newton (2019).

³⁰⁸ Greer (2003).

³⁰⁹ Flanagin & Metzger (2007).

³¹⁰ Milhorange & Singer (2018).

³¹¹ Amazeen & Muddiman (2018).

³¹² Sundar et al. (1998).

Just as they did in the Middle Ages, audiences trust information that is familiar and/or comes from friends. Coverage of something that people have experienced may also make them more likely to trust media reports.³¹³ One study³¹⁴ found that fake news headlines that were familiar were perceived as substantially more accurate even when they were clearly implausible or contradicted the respondents' beliefs. Warning labels about the headlines being incorrect had no effect on perceptions of credibility or even caused people to share the information more as readers assumed that a Facebook warning meant the story was true.³¹⁵

Based on a survey using Facebook data on graduate students about non-partisan news issues, researchers found that recommendations from Facebook friends/opinion leaders caused respondents to trust both the news article and the outlet it came from more, and caused respondents to say they would be more likely to read an outlet in the future.³¹⁶

Through content filtering and blocking, digital platforms have become increasingly influential in determining what information people

consume. Hence, while the Internet emerged with the promise of democratizing information access, with the rise of digital platforms as information gatekeepers, information flow and curation has become more concentrated.

The figure³¹⁷ below summarizes the differences between the functioning of traditional news media, and the news media after the rise of digital platforms.

A look at the economics of news media shows us that the news media industry has been facing three distinct disruptions with the rise of digital platforms.

- **Advertising disruption:** The production of investigative journalism has become difficult because ad revenues have dwindled with the rise of digital giants, and cash-strapped news media firms cannot afford to produce original journalism-type content.
- **Atomization disruption:** Newspapers have lost the power to bundle news, and news no longer remains picked by professional news editors. Instead, news gets bundled by opaque algorithms designed by a few digital giants, who have become the new gatekeepers of news and



³¹³ Livio & Cohen (2016).

³¹⁴ Pennycook & Rand (2017).

³¹⁵ Levin (2017).

³¹⁶ Turcotte et al. (2015).

³¹⁷ Raj and Rolnik (2018).

information, and whose only goal is to maximize engagement.

- **Accountability disruption:** In the traditional model, editors were responsible for the news they published for public consumption. In the digital model, algorithms designed by digital giants to filter and curate content have little incentive to be public-spirited, as algorithms neither produce the content, nor do they pay the negative externalities. Instead, they are designed to maximize “engagement,” which pushes them to prioritize visceral and viral content over news of public interest.

The Market Response

The market has responded in the last decade to the sharp decline in revenues of news outlets in various ways. Most prominent was the surge in the number of news outlets financed by philanthropists and foundations. Another market response was a gradual shift of outlets to a revenues model based solely or mostly on subscribers. While these market responses try to tackle the revenues or the financing sources, other market responses try to reduce the costs side. Among these are efforts to use technology and computation to perform some of the journalistic work and collaboration between large groups of newspapers on global investigative projects. In the next section we review these market responses and discuss their various shortcomings.

A. The Multifaceted Donation Model

Philanthropy is booming in our democracies, in particular in the US, where we see a growing role of private funders in the provision of public goods as government retrenches. Such a phenomenon is not specific to the media industry. As highlighted in Reich, Cordelli and

Bernholz (2016), “in the United States and most other countries, we see philanthropy in all areas of modern life,” but philanthropy also increasingly supports the provision of information.

The growing role of philanthropy in media funding has been well documented. The *Growth in Foundation Support for Media in the United States* report published in 2013 by Media Impact Funders reports that \$1.86 billion was awarded in media-related grants from 2009 to 2011. The investigative website *ProPublica*, created in 2008 by the billionaires Herbert and Marion Sandler, is funded entirely through philanthropy; its French counterpart, *Disclose*, launched in November 2018, is similarly raising money through crowdfunding and larger donations, including from US foundations (e.g., Open Society). Other examples include First Look Media and *The Intercept*, created by Pierre Omidyar; and recently *The Markup*, a news site to investigate big tech, subsidized by Craig Newmark (the Craigslist founder).

As of today, there are more than 150 nonprofit centers doing investigative journalism in the US, and for-profit newspapers like the *New York Times*—just like foundation-owned newspapers like *The Guardian* in the United Kingdom—have recently set up nonprofit ventures to support their journalism. Interestingly, *The Guardian* has decided to implement a unique business model, where there is no paywall (news is available online for free for all consumers), but where consumers are invited to nonetheless subscribe or donate to the newspaper so as to preserve independent journalism. As of today, *The Guardian* gets more revenue from consumers than from advertising thanks to the success of its membership and contribution model. More than a million people worldwide

contributed to *The Guardian* between 2015 and 2018.³¹⁸

Out of the 160 member organizations of the Institute for Nonprofit News (an association founded in 2009 with just 27 members), more than 100 were created between 2007 and 2017.³¹⁹ In France, the *Le Monde Afrique* website, launched in 2015 by the daily newspaper *Le Monde*, has received financial support from the Bill & Melinda Gates Foundation. Overall, philanthropy is becoming a very large part of the revenue streams of a growing number of news companies. In a series of articles published in the *Columbia Journalism Review*, David Westphal defines philanthropy as “journalism’s new patrons.”³²⁰

Concurrently, during the last decade, we have also observed an increasing tendency of out-of-market billionaires to acquire media outlets, often at a very low cost but with even lower profit expectations. Jeff Bezos (Amazon) and *The Washington Post*, Patrick Soon-Shiong (a biotech billionaire entrepreneur) and *The Los Angeles Times*, Marc Benioff (Salesforce) and the *Time* magazine are but a few examples of this new “taste” of tech entrepreneurs with deep pockets for the media industry. While these new media moguls publicly claim that they are acting as philanthropists, it is more accurate to call them “new media patrons.” The development of this patronage model is far from specific to the US, as is apparent from the recent entry of telecommunications billionaires on the French media market (e.g., Xavier Niel and *Le Monde*; Patrick Drahi and *Libération*, BFMTV, RMC, etc.), and most recently of the Czech billionaire Kretinski (who made a fortune in the energy sector, and is now buying

shares in *Le Monde* and other media outlets). Furthermore, this model has a historical precedent. In the 19th century, before the appearance of the penny papers and the development of mass media, “out of their own funds, wealthy political leaders sometimes provided start-up capital for newspapers.”³²¹ The main difference with today’s situation is that while historically the patronage was *political* inasmuch as these newspapers were endorsing political parties, nowadays, a large share of the new media moguls seem to care much less about politics (ensuring that a Republican or a Democrat candidate is elected) but much more about regulation. Or, more precisely, about the assurance of the absence of regulation (in particular in the case of the e-commerce and of the telecommunication sector).

It is important to distinguish between philanthropic funding (via charitable donations) of the media on the one hand, and the patronage model on the other. The philanthropic model consists of creating nonprofit news organizations that are then funded via charitable donations. The patronage model, while it also claims to be philanthropic in spirit (in particular in view of the low profitability of the sector), consists in buying and controlling news media organizations, keeping them as for-profit entities. But in the end, these two models pose similar problems regarding journalists’ independence and the disproportionate weight given to the preference of the wealthy. While philanthropy may offer one resource with the potential to fund the production of high-quality journalism, media outlets must resist potential hidden agendas. This is not specific to the media, and the risks

³¹⁸

<https://www.theguardian.com/media/2018/nov/05/guardian-passes-1m-mark-in-reader-donations-katharine-viner>

³¹⁹ Birnbauer (2018).

³²⁰ Westphal (2018).

³²¹ Hamilton (2004) similarly highlights that, before the emergence of nonpartisan reporting as a commercial product in the American newspaper markets in the 1870s, the type and amount of information provided depended on the value of the readers as that derived from political patronage.

of philanthropist funding have already been highlighted in the context of the funding of education, with questions about the power of donors to set research agendas. As highlighted by Reich (2018), we must consider philanthropy “as an act with political dimensions, in the sense that philanthropy can be an expression of political power. ... Wealthy elites can pose problems for democratic politics, even—and perhaps especially—when elites direct their wealth toward the public sphere” (p.64).

Both the philanthropic and the patronage model raise the same issue: Power resides where the money is. The media have all too often served as toys for billionaires in search of influence. From this point of view, there is no difference between private ownership of for-profit entities and the funding of foundations. Most often, donors indeed retain control over the governance and the purpose of the foundation, and in particular over how the funds are spent. If—to take only one example that shows how complex the situation is, since absent this external funding the newspaper would have cut off its newsroom—being owned by the founder of Amazon raised an independence issue and auto-censorship risk for journalists working at the *Washington Post*, who may for example less easily cover issues linked to e-commerce, what is the difference between being owned directly by Jeff Bezos (as is the case today) or being funded on a daily basis by a hypothetical “Bezos Foundation for the Media” created, funded, and governed by the same Bezos? This answer is simple: There is no difference.

The limits of the foundation model for the media have been well described. Benson (2016) has documented that “foundation donations are not ‘free’ but rather constitute a redirection of public resources ... to nontransparent and unaccountable foundations that have assumed media policy

responsibilities.” Moreover, foundations prefer funding specific projects rather than general operations, which creates the possibility of conflicts of interest. Obviously, founders will always claim that they never impose changes to the content of the investigations they have funded—and it may well be the case—but do we really expect media outlets to apply for funding to investigate the funders?

Underlining the limits of the foundation model does not mean, however, that we do not need nonprofit journalism. On the contrary. The central question is not one of the corporate form of the news organization (for-profit or not-for-profit) but the one of its governance. A number of interesting initiatives have emerged in recent years, such as the Civil Media Company in 2018, a startup that aims to use blockchain technology and crypto-economics (more precisely a cryptocurrency based on the Ethereum blockchain) to start hundreds of publications in the United States.

Regarding the donation model, note in conclusion that a growing number of donations are made today by the digital platforms themselves. Google’s “Digital News Initiative,” for example—initially launched for three years in 2015 with a \$150 million fund and relaunched in 2018 with \$300 million to be spent over the next three years—can be considered as a foundation-like initiative to support the media. Similarly, the Facebook “Journalism Project” aims at helping local news outlets make use of social media. Another example is the Google News Lab, whose catchphrase is as follows: “We collaborate with journalists and entrepreneurs to help build the future of media.” While it is now clear that digital platforms should contribute to the funding of journalism (given they are weakening the economic fundamentals of high-quality news production, and they are making money out of it)—an issue we will come back to in the solution part of this report—it is

unclear they should do so as if they were benevolent donors. Because they are not. Furthermore, they should not be free to choose which media outlets to help or not.

B. The Newsrooms Collaboration: An Alternative Path for Non-profit Journalism

While we have just highlighted the pros and cons of the donation model for the future of the news media, it is interesting to focus on a new form of nonprofit news organizations, the consortiums of journalisms. The most famous is the International Consortium for Investigative Journalists, a global network of more than 190 investigative journalists in more than 65 countries around the world, which recently exposed the Panama Papers and the Paradise Papers.

More generally, collaborative journalism is growing all around the world. Collaborative journalism is defined by the Center for the Cooperative Media as “the practice of executing journalistic endeavors using a cross-entity approach.” Already in 2014, the Pew Research Center noted these collaborations defined “a new era of interest.”³²² The website Medium recently listed the best collaborative journalism projects of 2018, among them the BBC Local News Partnership, which gathers together 843 newsrooms within 90 news organizations in the UK sharing local content.³²³ In a recent report, the Center for Cooperative Media of Montclair State University identifies six models of collaborative journalism: (i) temporary and separate, (ii) temporary and co-creating, (iii) temporary and integrated, (iv) ongoing and

separate, (v) ongoing and co-creating, and (vi) and ongoing and integrated.³²⁴

What are the advantages of collaborative journalism? The very first one is economically driven. As highlighted above, the media incentives to produce original news are negatively affected nowadays by extensive copying. Collaborative journalism and consortiums of journalists can be an interesting solution for media outlets to reduce the fixed costs associated with costly investigative journalism by sharing them. (In a sense, this is the exact same logic as the one behind the Associated Press at the time of the creation of this nonprofit cooperative.)

Furthermore, investigative journalism increasingly relies on the use of big data, which necessitates costly—and sometimes complicated to use—data-driven technology. The Panama Papers investigation, based on a 2.6 terabyte trove of data, would not have been possible without these new technologies such as automation, algorithms, OCR, etc. Here again, it is much easier for journalists to collaborate across newsrooms. Note, however, that the Panama Papers leaks also would not have been possible without journalists. They involved more than 100 media partners and several hundred journalists. Similarly, for the Paradise Papers investigation, with its files including far more information about US citizens, the ICIJ collaborated with more than 380 journalists working on six continents in 30 languages.

Finally, at a time when there is a growing threat to journalistic independence and press freedom, in particular due to recent changes in media ownership, collaborative journalism can be seen as a way to avoid

³²² <http://www.journalism.org/2014/12/04/journalism-partnerships/>

³²³ <https://medium.com/centerforcooperativemedia/a-look-at-nine-of-the-best-collaborative-journalism-projects-of-2018-cfd49b3c4865>

³²⁴ <https://centerforcooperativemedia.org/center-cooperative-media-identifies-6-models-collaborative-journalism-revolution-media/>

ensorship. The example of the Panama Papers is particularly relevant from this point of view. Given that tens of newsrooms in many different countries were involved, it was impossible for each country to censor some of the findings.

Obviously, all these advantages do not imply that collaborative journalism is the sole solution to the media crisis. There are some downsides to the consortiums of journalists. As highlighted by the Pew Research Center (2014), “things can easily go wrong.” The report gives the example of a Knight-funded grant series to pilot eight collaborations between news outlets that had only one active participant when the seed money ran out. But overall, it seems necessary in the future to provide more funding to these initiatives that allow the production of investigative journalism.

C. The Subscription Model

Despite the observed drop in advertising revenues in recent years, the core business model for effective financing of publishers’ websites is still through advertising. Advertising is indeed the largest contributor to publishers’ online revenues. Even if pay models are becoming an important part of the business of digital news nowadays, in most countries there is still only a minority of news lovers who pay for online news.³²⁵

However, it is interesting to highlight that a number of recent and successful media outlets have made the choice to rely only on subscriptions. One of the best illustrations of such a successful strategy is the French pure online publication *Mediapart*. This publication, specialized in investigative journalism, was created in 2008 with a hard paywall model. It has in recent years played a key role in uncovering several corruption scandals involving politicians of both the left and right.

At the end of 2012, for example, *Mediapart* revealed that the French budget minister evaded paying tax in France on sums deposited in undeclared Swiss bank accounts. Following *Mediapart*’s allegations, a legal investigation was opened into the tax fraud accusations, and Jérôme Cahuzac resigned before being charged with tax fraud. As of today, *Mediapart* has more than 140,000 subscribers providing revenue of €13.7 million (in 2017). With its 4,700,000 unique visitors per month and 85 staff members, the publication is highly profitable (and has been making a profit for seven years now).

As of today, *Mediapart* can be considered a model for the whole news industry. Why does the subscription model seem to be an interesting path to follow for the future of the news? First, because the collapse of advertising revenues for newspapers is not new. Even in the United States, where advertising is king, newspaper ad revenues have been declining as a percentage of GDP since 1956—and will continue to do so in the future.³²⁶ Hence publishers need to find alternative sources of revenues. As we have seen above, in today’s online world, publishers are competing with a duopoly online (Facebook and Google) and they are no longer competitive, in particular regarding targeted online advertising.

Second, media outlets such as *Mediapart* that are behind a paywall are much less dependent on the digital platforms regarding their traffic. Given the subscription model, their traffic is indeed mostly direct (while we saw before that the vast majority of consumers now prefer to get to news through a side door). This lower reliance on platforms implies that when Facebook decides to modify its algorithm—as it did for example in the summer 2018—a medium such as *Mediapart*,

³²⁵ Cornia et al. (2017).

³²⁶ Cagé (2015).

contrary to the majority of the French media, was barely affected.

Other innovative business models in recent years include the one of *The Correspondent*, which is entirely member-funded. Originally launched in 2013 as a Dutch news website funded through a successful crowdfunding campaign, *The Correspondent* just terminated a successful US\$2.5 million campaign to launch an English-language “unbreaking news” platform in the summer of 2019.

D. Rating the Sources: The Solution to the Spread of Disinformation?

In the vast majority of Western democracies nowadays, a challenge is to fight against the spread of disinformation. The market has developed a number of solutions.

One of them is NewsGuard, a browser extension that labels news sources with either a green (for trustworthy) or red (for unreliable) icon. NewsGuard was founded in 2018 and financed by the Knight Foundation and Publicis (among others). In 2019, it has expanded its partnership with Microsoft and is now accessible to users of Microsoft Edge mobile apps on iOS and Android. The company—with a team of roughly 50 journalists—rates more than 2,000 websites.³²⁷ (Similar initiatives in the US include Snopes and PolitiFact.)

In the same spirit, the French daily newspaper *Le Monde* has recently developed Décodex, a database of around 1,000 websites compiled by *Le Monde*’s Décodeurs project in the course of their fact-checking. The Décodex divides websites into four categories: (i) satirical websites, (ii) websites that have published a significant amount of false

information, (iii) websites whose approach to verification is questionable, and (iv) news websites.

The main downside of a database like Décodex is its reliability. Obviously, *Le Monde*’s Décodeurs project can be considered trustworthy, and *Le Monde*’s journalists are recognized worldwide as high-quality independent journalists. But other initiatives in other countries can be less reliable. And more importantly, a number of these initiatives may have a hard time convincing citizens that they are actually reliable. Breitbart, for instance, calls NewsGuard “media blacklists” that “[promote] fake news.”³²⁸

An open question remains: Who is going to investigate the reliability of the newsroom in charge of rating the other newsrooms? From this point of view, an initiative that may be more attractive than NewsGuard or Décodex is CrossCheck: French media outlets—from the Agence France Press to BuzzFeed through *Libération* and *Le Monde*—decided to team up on a fact-checking initiative. In other words, this initiative combines both collaborative journalism and the rating of sources.

E. Summary

News organizations around the world have shown different degrees of success in adjusting to the new realities of the digital world. While some news outlets or media groups were able to gradually transform their businesses into the digital era by investing in online businesses, most news outlets in the developed world had to aggressively cut the number of journalists. We surveyed a number of market responses of news organizations and entrepreneurial journalists to the new reality of the digital age.

³²⁷

<https://www.nytimes.com/2019/01/16/business/media/media-steve-brill-fake-news.html>

³²⁸ Nolte (2019).

Chief among these responses is the surge in philanthropy-funded journalism. As noted in the section, this form of funding has created many successful initiatives, famous among them ProPublica in the US. Yet philanthropic-funded journalism can present the same problems that corporate control of news media presented in the past. Philanthropists may be benevolent players in the news ecosystem, but they may have their own agenda and limited interests. In most countries very wealthy individuals or billionaires would be wary or reluctant to finance or be involved in news outlets that pick fights with powerful politicians or business groups. A philanthropy-funded news ecosystem can result in a public discourse and media agenda that is in line with the point of view of a handful of billionaires.

The subscriber model that is proving itself for large established news outlets like the *New York Times*, *Wall Street Journal*, or *Financial Times* and new digital native initiatives like *MediaPart* in France offers hope for quality news outlets that have a very distinct brand reputation in their market. But this model will be a partial solution for most countries. News outlets that were able to get significant subscribers revenues to support large newsrooms are usually the top national or global outlets that target very large audiences. The smaller the audience and the more local the reporting, the lower is the likelihood of a subscriber model achieving much success.

Recommendations

As this report acknowledges in different sections, democratic journalism has always been in “crisis.” Nonetheless, the rise of digital platforms and the subsequent disintermediation of the relationship between publishers and

users raises some important issues for the future of newsrooms. The question remains of what, if anything, can be done to directly address the platforms’ control of the relationship between publishers and readers (a form of bottleneck power).

Digital platforms’ bottleneck power manifests in their ability to use monopsony power to pay news outlets less than the competitive price for the news.³²⁹ The previously described changes in news consumption patterns means that news outlets became increasingly dependent on online platforms to access readers, also becoming more dependent on the platforms to tap ever scarcer advertising revenue. This increased platforms’ economic power: Not only do they carry news for free, with the only compensation for sharing snippets and other reports being the increased traffic and attention diverted to news outlets’ websites. They also squeeze online advertising margins through their control over this associated ecosystem. To make matters worse, the dependency is one-sided: while platforms like Google and Facebook control ad exchanges and account for the lion’s share of traffic to most newspapers, the opposite is not true: Facebook reports that only 4 percent of its News Feed posts are news. This further weakens newspapers’ bargaining power, as well exemplified by the examples of Germany and Spain, where platforms simply stopped carrying news after changes in copyright laws required them to pay for news reports. The significant drop in traffic to most newsrooms was such that they were soon acquiescing to the older free-carrying terms in order to stay afloat. Another version of this power is the platforms’ ability to maintain as proprietary most user data associated with news consumption—in particular when it is done through the platforms’ API instead of on publishers’ websites (a process that should rise as

³²⁹ See Cairncross (2019).

aggregators such as Apple News increases in importance).

Countries are struggling to develop tools that can effectively address this bottleneck power, which mostly reflects platforms' tight controls over their ecosystem—a process that may benefit users. Europe's revamped copyright laws hope to force platforms to the bargaining table by granting publishers more control over how their products are shared. While these efforts are laudable, this report is based on the view that the golden era of advertising-funded journalism never existed, such that attempts to return this glorious past seem misguided. That is why we defend a series of mechanisms to help newsrooms survive in a new technological environment: from alternative vouchers to fund the news media to a series of news monitoring obligations that should ameliorate the platforms' bottleneck power by providing more transparency about their actions. Nonetheless, if digital platforms continue to control the interaction between newsrooms and users, a regulator should be empowered to take stronger action to reign in platform power and ensure that citizens can continue to access relevant news.

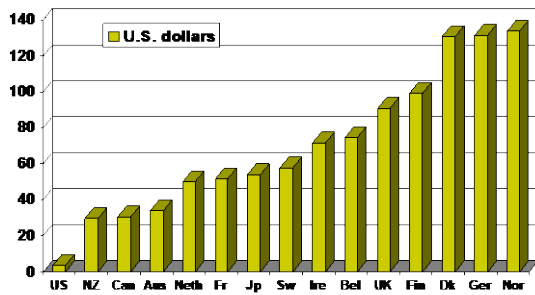
This system would be structured as part of the quid pro quo bargain that shields platforms from liability in exchange for more responsive and public-oriented companies. We endorse some of the suggestions of the Cairncross report—an independent report published in February 2019 that offers an overview of the challenges facing high-quality journalism in the UK—to require platforms to voluntarily adopt a code of conduct in which they clearly state the basis of their relationship with publishers. We would go further, recommending that such a code of conduct include not only economic terms (whether news outlets are rewarded for snippet sharing,

etc.), but also a description of what types of user data the platforms will share with publishers and what efforts platforms are taking to distribute meaningful and relevant journalism within their ecosystems. This transparency should level the playing field and allow for competition between publishers to reach the audience. Access to data is particularly important, as it will help publishers establish more meaningful independent subscriber relationships. More data allows for better personalization in general, which will be reflected not only in better ads but, most importantly, in a publisher that is more responsive to the demands of its subscribers. Publishers may, for example, personalize digital editions to reflect the preferences of their readers—a process similar to what services like Apple News promise to do. It may also boost newsroom collaboration efforts, equally benefiting publishers.

Public Funding of News

Public support of journalism is not a novel concept. In the US, public funding of the press was common in the nineteenth century in the form of postal and printing subsidies given to publishers to print and distribute newspapers, which were estimated to be around 0.21 percent of the US GDP between 1840 and 1844 (equivalent to \$43 billion as 0.2 percent of 2018 US GDP).³³⁰ As the figure below makes clear, in most developed countries, the government financially supports the media one way or another, while the US is an outlier.

³³⁰ McChesney & Nichols (2010): 310.



Per capita public funding of public media in US dollars. Source: Benson and Powers (2011), estimates from 2007 to 2009.

In the US and elsewhere, a growing number of researchers are advocating increased public funding.³³¹ Most recently, the Cairncross Review recommends first direct funding for local public-interest news outlets, and second the launch of a new innovation fund. A salient argument in opposition to such public funding is the threat to editorial independence. In this section, we propose a “private media voucher” system—funded with public money—to remedy some of the biases that increased public funding of media may induce. We first present our proposal, describe its advantages and shortcomings, and then discuss why we do think that the media voucher system is better not only than the status quo, but also than other propositions that have been made recently to guarantee a sustainable future for journalism.

A. Existing Subsidies Schemes

There are plenty of examples around the world of government support of the media. While American subsidies to both the press and audiovisual media (e.g., the postal subsidy) are low and have fallen sharply over the past few decades, the US was historically the very first country to introduce subsidies to the press.

The most commonly used subsidy scheme today is the reduced value-added tax (VAT). In most European countries, newspapers indeed benefit from a reduced VAT (although a challenge in many countries remains the extension of the reduced VAT to digital publications). Although it is widely believed that UK newspapers are not subsidized by the government, they in fact pay zero VAT, which amounts to an effective subsidy of several million pounds each year. Furthermore, a number of European countries have also introduced both indirect and direct subsidies to newspapers. Subsidies can be either neutral or discriminatory. Sweden, for example, offers both an operating subsidy and a distribution subsidy; Norway subsidizes the newspaper that is second in circulation in each local market, the smallest paper in certain isolated regions, as well as a national paper that offers dissident and controversial political views; while France offers delivery subsidies and pluralism subsidies, as well as modernization subsidies.³³²

Even in countries where press subsidies are relatively low—as in the UK—the state intervenes in the media through subsidies to the audiovisual sector. The funding of public broadcasting can take different forms, e.g., license fees, income tax charges, parliamentary grants, etc., depending on the countries.

In late 2018, Canada’s federal government introduced a tax package worth CAN\$595 million, rolled out over five years, in support of journalism.³³³ It includes (i) a temporary, non-refundable tax credit that will allow subscribers to claim 15 percent of the cost of subscriptions of eligible digital news media; (ii) a new category of “qualified” donor for nonprofit journalism; and (iii) a refundable tax credit for qualifying news organizations

³³¹ McChesney & Nichols (2010); McChesney & Pickard (2011); Bollinger (2010).

³³² Cagé (2015).

³³³ <https://www.niemanlab.org/2018/11/canada-introduces-a-595-million-package-in-support-of-journalism/>

that “produce a wide variety of news and information of interest to Canadians.”

At the local level in the US, in July 2018, the state of New Jersey decided to give \$5 million in funding for innovative projects to improve local news in the state.³³⁴ The public funding will go through the New Jersey Civic Information Consortium, a nonprofit news incubator.

Notwithstanding this growing interest for government subsidies to the media, we think that there are good reasons why there is little public aid for the press in the US as of today. Public funding of news can indeed threaten media editorial independence. We accept this criticism. Given that we have documented in this report that media are in need of new financial resources, we propose an innovative model: publicly funded media vouchers that are privately allocated.

B. Our Proposal: Media Vouchers

Our proposal to provide public funding for the news media can be summarized as follows: We propose to give each adult a media voucher worth a certain value—\$50 in our favorite proposal—per year from the US Treasury, to donate to her favored media outlet(s).³³⁵ This proposal is in the spirit of the “vouchers for equal democracy” proposed by Cagé (2018) (a €7 voucher given each year to each citizen to fund the political party of her choice), as well as the “democracy vouchers” advocated by Lessig (2015) (with a specific focus on elections) and implemented in Seattle for local elections (we discuss below the lessons from the Seattle experience).

This system would work as follows. Every year, when filling her tax returns, each

citizen indicates to the tax administration the media outlet(s) to which she wants to allocate her media voucher. The vouchers can be split in up to 10 different \$5 vouchers. Technically, to preserve anonymity, this should work like electronic voting: each citizen is provided with a token and the allocation choices should not be linked to the addresses of the token holders (there exist many protocols of anonymous voting on blockchain-based networks that could be used). Resale of the media vouchers will be forbidden by law.

1. Which media outlets could benefit from the media vouchers?

We want to guarantee that the list of media outlets eligible to receive voucher funding is as extensive as possible, so as to be sure there is no threat regarding media independence, and that the vouchers are used to fund the production of information (and not of entertainment, for example). We also want to guarantee a high degree of pluralism.

To guarantee that these public subsidies actually subsidize the production of information—and in particular the production of high-quality information—we impose a small number of conditions. To benefit from the media vouchers, the media outlets have to:

- Hire at least one journalist. This will in practice exclude from the benefit of the media vouchers all aggregators that do not produce any original content. While this threshold may seem too low, we have chosen it on purpose, following the criticisms faced by the Canadian subsidy model described above. To qualify for the subsidies in the Canadian model, media have indeed to “employ more than two

³³⁴ <https://www.niemanlab.org/2018/07/water-in-a-news-desert-new-jersey-is-spending-5-million-to-fund-innovation-in-local-news/>

³³⁵ We propose here a \$50 voucher to keep the total cost reasonable. As we have highlighted in this report, high-quality information is a public good, and this public good suffers from critical underfunding.

journalists as employees.” This excludes small news startups.³³⁶ We think that small startups should not be excluded; this is all the more important given that a key priority should be to fund local journalism; at the city or even the county level in the US, media outlets have very small newsrooms.

- Mostly produce general-interest news. In the majority of the countries where public subsidies exist, there are some conditions associated with who can benefit from them. In France, for instance, to qualify for the majority of the subsidies, media outlets should obtain the “*Information politique et générale*” label. Our approach to “general-interest news” is rather large here; the media vouchers could, for example, be used to fund media outlets that specifically focus on a given topic, e.g., the economy or the environment. Moreover, importantly, we do not condition the format of this news (it can be print, but also audiovisual, etc.). What is important is that general-interest news is actually produced; we think that it may be a good thing to have various production formats. We simply want to guarantee that the vouchers are not used to mostly finance the production of non-news content.

- Be transparent. The modern media industry suffers from a lack of transparency. This lack of transparency—in particular regarding ownership—partly explains the very low trust in the media (which we observe not only in the US but also all over the world). Hence for the media outlets to be able to benefit from the media vouchers, they will have to annually publish online (as well as in their print version, for the newspapers) the following information: (i) the detailed list of their owners, for each shareholder with more than 1 percent of the capital shares; (ii) for each of these owners, their main source of revenues (to avoid

conflicts of interest). Moreover, they will also have to publish their annual balance sheet. Finally, they will have to make public their articles of association, with detailed information regarding the governance structure. We think that introducing such transparency is important; this condition is along the lines of the public funding of political parties. In Italy, for example, to benefit from the “*due per mille*” system, political parties have to publish their accounts, their status and the list of their donors.

- Be ethical. Finally, while the media outlets meeting the previous criteria will be eligible for media voucher funding independently of their political bias or the tone of their coverage, we think it is nonetheless important to introduce an ethical code created by news media stakeholders. The idea behind such a code is simply to avoid the public funding of disinformation and other harmful content. Compliance with the code will be assessed on an annual basis and the media outlet will have due process and a chance to appeal any adjudication of violation. If a media outlet loses its media vouchers status in a given year, it will nonetheless be considered again the following year (so as to avoid the creation of an opposition from a number of media outlets, which could use their “victimization” status).

Finally, an independent body—which we could call, for example, the Independent News Monitor—will be in charge of validating on an annual basis the exact list of the media outlets that meet these criteria and administering the ethical code. Key is the independence of this body; we believe that it should include representatives of journalists and of media owners, as well as scholars. These members should be named for a four-year nonrenewable term. (This will be, for example,

³³⁶ <https://www.niemanlab.org/2019/03/instead-of-helping-canadian-news-startups-a-new-government-subsidy-will-only-prop-up-failed-models/>

in the spirit of the aforementioned Cairncross Review’s proposal to create a new Institute for Public Interest News.)

Obviously, other conditions could be in principle introduced, for example, that to qualify for the media vouchers, a media outlet may not rely for more than a certain share of its revenues on advertising revenues. We leave this dimension for discussion, but we believe nonetheless that the lower the number of conditions, the better.

2. A focus on local media outlets

In the baseline version of our proposal, all media outlets, as long as they satisfy the previously listed conditions, could be eligible for media voucher funding. We are well aware, however, that local media are currently struggling much more than national media. While a number of national outlets have recovered their audience in recent years—in particular newspapers through the rise in digital subscriptions—and while they still rely on advertising revenues (even if declining), the business model for local news is gone. In print media, the traditional economic model of a local newspaper was to bundle diverse types of content, such as local news, national news, classified ads, etc., into a single product to sell to consumers.³³⁷ But digital platforms have weakened this approach. Classified ads have moved to specialized online outlets (e.g., craigslist.com or monster.com), and soft news about local communities is now provided free of charge on social networks such as Facebook. Similarly, national and international news is now provided almost exclusively by a few of the largest news outlets.

According to a study released in 2018 by the University of North Carolina’s School of Media and Journalism,³³⁸ more than one in five

local papers has closed in the US over the past decade and a half. Almost 200 counties in the US have no newspaper at all. Importantly, as highlighted in the study, “the people with the least access to local news are often the most vulnerable—the poorest, least educated and most isolated.” In a recent conversation (May 2019), *New York Times* executive editor Dean Baquet said, “The greatest crisis in American journalism is the death of local news.”³³⁹

Hence, we think that financial support is of particular urgency for local media. We thus propose the alternative allocation mechanism: to favor the funding of local journalism, citizens should allocate at least half of their vouchers—and can allocate up to 100 percent—to fund local media outlets.

3. What happens if a citizen decides not to allocate her voucher?

Obviously, it may happen—even if allocating the media vouchers costs citizens nothing—that a significant fraction of taxpayers decide not to allocate their vouchers (similarly to what happens, for example, with the “presidential fund” in the US: most people do not check the box). Given that we believe it is important that there is a high enough amount of public funding devoted to the production of high-quality news each year, we think that the vouchers should nonetheless be allocated in this case.

Hence, we propose the following rule: in case a citizen does not choose a media outlet to which to allocate her voucher, then her voucher will be allocated as a function of the allocation of the other vouchers. We think this is the best allocation rule: it relies on the preferences expressed by the citizens and avoids any government intervention.

³³⁷ See Angelucci, Cagé & Sinkinson (2017).

³³⁸ <https://www.usnewsdeserts.com/reports/expanding-news-desert/>

³³⁹ <https://www.inma.org/blogs/world-congress/post.cfm/trump-vs-new-york-times-the-executive-editor-s-perspective>

An alternative allocation rule, to favor the production of local news, would be to allocate these vouchers as a function of the allocation of the other vouchers but only among the local media outlets (in the spirit of the scheme described above).

4. How can concentration be avoided?

A possible caveat of our scheme is that it could potentially lead to the allocation of the large majority of the media vouchers to a small number of media outlets, in particular to the best known outlets at the time of the scheme's implementation, while these outlets may not be the ones most in need of public funding, and while we want to favor competition. To avoid such hyperconcentration and guarantee pluralism, we introduce the following threshold: a given media outlet cannot receive more than 1 percent of the total number of media vouchers. In the US, which includes around 260 million adults, no media outlet would be allowed to receive more than 2.6 million vouchers (i.e., \$130 million).

What will happen in the event of an “over-allocation”, i.e., when more than 1 percent of the adult population decides to (simultaneously) allocate its media vouchers to the same outlet? Then, in this case, we follow the same rule as the one described above in the case of non-allocation: the media vouchers are allocated as a function of the allocation of the other vouchers, but only among the media outlets that are below the 1 percent threshold. (Obviously, we can discuss the “optimal” threshold; 1 percent may be too high, and a 0.5 percent threshold may be seen as preferable. We may need a higher threshold especially for smaller countries.)

Note also that, as we highlighted above, we allow the citizens to split their vouchers; while the face value of one voucher is \$50, if they wish to, they can split it into up to ten

different \$5 vouchers. This should stimulate pluralism, as well as buttress the scheme where at least half of the voucher has to be allocated to fund local media (see above).

5. Will the vouchers make a difference?

Note that while we want to make sure that the vouchers won't be concentrated among a small number of successful outlets, we nonetheless also want to assure that the vouchers will be an efficient tool that will allow existing or new news outlets to produce high-quality information that citizens can consume.

The \$50 media vouchers will guarantee that this is indeed the case. Take the “extreme” example of a large and successful newspaper, the *New York Times*. Given that a single media outlet cannot receive more than 1 percent of all the vouchers, the maximum amount that could be received by the newspaper is equal to \$130 million. The typical *New York Times* reporter's annual salary is around \$110K (which involves a full cost of around \$130K for the newspaper). Hence \$130 million corresponds to around 1,000 reporters, a number that has to be compared to the 1,450 journalists who work as of today for the newspaper. The \$130 million figure also corresponds to less than a third of the total spending on wages and benefits of the company. Hence, even for a large, profitable media outlet such as the *New York Times* (\$112.4 million in operating profits in 2018), the media vouchers could make a difference.

A. Why do we advocate media vouchers?

We believe that media vouchers are the best scheme that could be implemented in the future to provide funding for the media. They present a number of important advantages compared to existing schemes. In this section, we will argue that they are better not only than the status quo,

but also than other proposals that have been made to sustain journalism in the future.

1. Media vouchers and the status quo

The status quo obviously varies from one country to the other. We have highlighted above the existing public subsidy schemes. These schemes are relatively important for existing print media in European countries such as the Nordic countries and France, while they are nearly nonexistent in the UK. The UK government nonetheless spends a great deal of public money to fund the BBC each year.

From a direct public funding point of view, the status quo in the US is almost the complete absence of funding. However, public money is invested indirectly in the media through tax deductions associated with philanthropy (see below). In the first part of this report, we have highlighted the gradual decimation of the business model that enabled many news outlets to produce investigative journalism for decades. If one believes—and we believe—that the journalism crisis raises growing threats to democratic values and institutions, then we need to confront the status quo. In a word, we need to find alternative resources for journalism.

One could argue that in this case, it will be enough in European countries that already subsidize their media to increase the existing amount of public funding (while preserving the existing schemes). We think that we can do better. The main advantage of our public funding (media vouchers) scheme over existing ones is that it offers a solution to the threat usually—and sometimes quite rightly—

associated with **public funding of news**: With media vouchers, dependence on public funding won't compromise the independence of the media. The state intervention will indeed be "neutral": all media outlets will be treated equally, without distinction as to content or opinion. In other words, there will be increased *public funding without application of government discretion*. This advantage results from the fact that in our model the media vouchers are allocated by citizens themselves. Hence we believe that in countries that today provide direct subsidies to their media outlets, it will be preferable in the future to use the same amount of money but allocate it through media vouchers.³⁴⁰

Some maintain that from this point of view **tax relief** does not compromise media independence. That is strictly right. But we think that vouchers are nonetheless preferable to tax relief for at least two reasons. First, tax relief disproportionately benefits large media and, depending on the tax relief scheme, large media that are profitable. Media vouchers, on the other hand, will help financially support not only large but also small media, from their emergence onward, and help ensure pluralism. Second, we believe that an additional advantage of vouchers is that they may help reconcile citizens with media. Media suffer from a lack of trust; part of the distrust comes from the fact that citizens doubt that media are independent, in particular because of their funding (either advertising revenues or ownership). Conversely, with media vouchers, citizens would be directly involved in media funding. They may obtain information on the ownership of each media outlet (given that outlets will have to publish the detailed list of

³⁴⁰ Our media voucher scheme could potentially be extended to the aid flows donors direct to the media sector abroad (around \$454 million per year of official development assistance), and in particular in sub-Saharan Africa. While there are concerns regarding the way this support is allocated, an innovative scheme—

which would reinforce media accountability with respect to citizens—would be to allocate these funds directly to the citizens in the form of vouchers, so that they could subscribe to or sustain the outlet(s) of their choice.

their owners—see above). They may decide to support the more independent outlets. This may help at least to partly reduce the **confidence crisis** and reconcile citizens with their media.

In the US, while there is a shortage of *direct* public funding of the media, we have observed in recent years a rapid increase in philanthropic funding. Hence some argue in favor of the status quo and defend a reform that would facilitate and give more incentives to **philanthropic funding of the media**. One way to do so could be to increase the magnitude of already existing tax deductions attached to philanthropic donations (which include donations to nonprofit media). We have already highlighted above the shortcomings of philanthropic media funding: philanthropists may be benevolent players in the news ecosystem, but they may also have their own agenda and limited interests.

Of course, we are not arguing against philanthropic funding. But rather than using even more public money to financially support wealthy individuals willing to finance the news ecosystem, we think it is preferable to spend additional public money to give to *all citizens*—rather than only the wealthiest—the resources to sustain financially the media outlets of their choice. The issue here is again the one of optimal allocation of resources. A handful of billionaires are no more legitimate than the state to allocate public resources to the media of their choice. To ensure pluralism and preserve media independence, we think it better to decentralize this allocation and leave it with all citizens through media vouchers.

2. Media vouchers and alternative funding models

Another alternative to reconcile citizens with media and to sustain the journalism industry—

in particular in countries where information is not only underproduced but also underconsumed—would be for the government to **offer to each citizen a free subscription** to a media outlet.³⁴¹ Compared to that alternative, our media vouchers proposal has at least two advantages. On the one hand, privacy is preserved (given anonymity, there will be no public information on who allocates vouchers to which media). Second, many media outlets offer free content and do not rely on a subscription model. These outlets would be de facto excluded from the benefit of the subscription alternative, while they could benefit from the vouchers (and we think they should).

Obviously, one possibility could be to combine the voucher model with the free subscription one; hence, we could decide that when a citizen allocates her vouchers to a media outlet, if the outlet's content is behind a paywall online or is a print medium to which one needs to subscribe, then she will automatically obtain a free subscription to this media. We think it would be interesting to also consider this dual model.

Note also that we believe that the proposals made recently (e.g., in the Cairncross Review), as well the current philanthropic support of news, have been focused too much on “**innovation**.” We believe that even traditional journalism needs to be funded; the focus should be more on quality news, and less on the technical means used to deliver it.

We also believe that there is no reason to focus exclusively on large investigations that require a great deal of resources. Obviously, we need such investigations but, in particular at the local level, there is also a shortage of day-to-day journalism with deeply locally involved

³⁴¹ For example, in 2009, the French government gave teenagers a year's free newspaper subscription on their

18th birthday. This experience was considered a success; however, it has not been renewed.

journalists, who sit through most often “boring” government meetings.


We want to emphasize this need for day-to-day journalism. Local journalism is of particular importance in fighting corruption at the local level. Absent local journalism, important events such as council meetings and court hearings are not covered. At the local level, journalists are the only democratic watchdogs—while at the national level there are many more whistleblowers and social media are used as a way to propagate information. Furthermore, local journalism matters not only with respect to politics, but also daily life activities and quality of life more broadly. Historically, local reporters have been the one in charge to publish on the front page a plan to close a local swimming pool, a school or a library. Finally, more local news is also associated with higher voter turnout at local elections.

Large prizes for investigative journalism can be an efficient mechanism to incentivize the production of the specific kind of journalism that democracy needs. They can be complementary to our proposal of increased public funding—but they are not enough. Media outlets are not suffering as of today from a lack of incentives to get a Pulitzer prize; they are suffering from a lack of resources to do journalism, and we observe journalists who at the time of receiving their Pulitzer are working as PR consultants, either because they have been laid off or because journalism’s level of pay has driven them out of their job.

Box 8: Investigative Journalism and Prizes

Relying solely on large prizes to promote investigative journalism creates the risk of creating a superstar media industry, where a few media firms with ample resources hog the limelight. Investigative journalism today has already become like a superstar

industry: Until the start of 21st century, local newspapers were the regular winners of the Pulitzer Prize for Investigative Reporting. Since the 2000s (with the dawn of the the Internet age), the Pulitzers have been dominated by superstar newspapers like the *New York Times*. In other words, large prize based incentives cannot be a cure for one of the key aspects of the media crisis: death of local investigative journalism.



Note finally that the \$50 voucher level advocated here is not meant to be the only possible level. One could think of higher or lower levels. This is an issue that requires extensive deliberation and experimentation and that we do not plan to settle here.

For the sake of comparison, the licence fee that is used to finance the BBC in the UK is currently equal to £155 per citizen/year. In France, the corresponding *redevance audiovisuelle* amount is €139. In Germany, the public TV licence fee is €210 per year (€17.5 per month).

However, as highlighted above, the main issue with European-style public media funding is that public subsidies may open the door for manipulating journalists and inducing pro-government media bias. This threat to editorial independence may explain why the US has been reluctant in recent decades to publicly fund the media. Our proposal opens a

new path to supporting the media from the US Treasury without government discretion.

3. Lessons from the Seattle voucher experience

As noted above, the inspiration for the media vouchers proposal comes from both the “vouchers for equal democracy” proposed by Cagé (2018) and the “democracy vouchers” implemented in Seattle for local elections. Hence, it is of interest to draw the lessons from the Seattle experience.

In January 2017, the city of Seattle implemented a “democracy vouchers” system to fund the local elections. In January 2018, all registered voters were mailed an envelope containing four \$25 democracy vouchers that they could donate to local political candidates of their choice. These vouchers were taxpayer-funded. To benefit from these vouchers, candidates had to agree to certain limits (with the maximum contribution to candidates set to \$500; the initial goal of the vouchers was to get big money out of politics).

Even if hard to evaluate scientifically, this experience is considered overall as a success. First, six candidates benefited from the vouchers and 46,000 vouchers (corresponding to around \$1.1 million) were allocated. Second, the candidates who benefited from the vouchers were generally successful at ballot box. Finally, the vouchers shifted the donor pool in an egalitarian direction relative to the pool of cash donors.³⁴² In particular, Brian J. McCabe and co-authors documented an increase in involvement by underrepresented groups. This increase is of particular importance for us here. Indeed, the objective of our media vouchers is not only to bring

necessary resources to the media, but also to rebuild confidence between citizens and media, and to incentivize people to consume more information. This is particularly important for those who tend not to consume much information.

4. The limits of media vouchers

Of course, media vouchers alone are not going to solve the journalism crisis entirely. In particular, while they bring a solution on the production side, they would have little *direct* impact on the consumption side. In other words, while media vouchers will lead to the production of more high-quality information, we cannot ensure that it will be consumed. Nonetheless, the voucher allocation may lead to an increase in news consumption. Having to choose on an annual basis which media outlets they want to allocate their media vouchers to may indeed raise citizens’ interest in news. The vouchers may also help reconcile citizens with media.

Should we be worried that citizens will “misallocate” their vouchers, in the sense that they may decide to allocate them to media outlets that produce relatively more entertainment and relatively less information, or to media outlets that are biased toward a political party? In other words, can we be sure that the media vouchers system will only support the quality media a social planner would like to support? We cannot know *ex ante* what will be the result of the allocation—and this is why it could make sense to pilot the voucher system in a small number of states to begin with. But we think it is essential to have such a *decentralized* allocation, with no government discretion, for two reasons: First, we do not think that the government will do

³⁴² See in particular McCabe (2017) and McCabe & Heerwig (2018).

better than citizens at allocating media vouchers, and will likely do worse. Second, we think that this decentralized allocation will be at least as good as what the market alone would do. Note that if the citizens most interested in news are also the ones who most likely to allocate their vouchers, then the decentralized media voucher allocation would most probably lead to more funding of high-quality information than today's market outcome.

5. Subsidizing inputs versus outputs

Production of journalistic content involves resources—talent and capital—representing fixed costs that must be incurred to generate news of value. Like many informational goods, the benefits and costs of this endeavor can be aligned in two broad ways. First, additional incentives can be created to compensate journalists upon the production of news content. Second, subsidies can be provided to compensate for the costs of producing that content.

In terms of approaches each has tradeoffs. Rewarding journalists based on some measure of the value of news output has the advantage of creating relatively high-powered incentives to actually produce news output with the ascribed value. By contrast, subsidizing journalists' costs risks insufficient effort being directed at providing news content.

The flip side of this tradeoff is a practical matter—one of measurement. How do you measure the value of news content produced in order for it to represent a clear reward to journalists? At present, that is done by attracting and monetizing attention. However, as noted already, the rewards from this are likely to be insufficient for certain types of content with high public value but low mass market appeal. By contrast, while costs pose their own measurement challenges, costs vary across journalistic endeavors in relatively

known ways—salaries of journalists, travel, etc.—and certainly by less than the variation in the value of news output. Thus, a relatively simple subsidy can stimulate activity.

For these reasons, we have posited a mechanism that is primarily based on subsidizing inputs rather than directly rewarding outputs. That said, there is an indirect mechanism in our approach. Citizens will be loath to provide subsidies to news outlets that are not producing news content. In other words, while there are static concerns that subsidies may be ill spent, the dynamic consequences should mitigate those concerns. This is the very same balance that we strike when subsidizing other information public goods such as basic scientific knowledge.

Regulation of Media Concentration

A. Citizen Welfare Standard

Concentration in the media sector creates standard problems associated with a lack of competition in economic markets. These problems include, among other potential harms, lower quality products and higher prices for consumers. Given the two-sided nature of media markets, high levels of concentration in the media sector might also create inefficiencies in the advertising market, with advertising prices set too low. Increasing competition makes prices for consumers fall and product quality increase. These changes lead to increases in standard measures of consumer welfare.

The subcommittee argues that these standard measures of consumer welfare miss key additional benefits of competition in the media sector. These benefits are incorporated into citizen welfare, which recognizes the centrality of the media, and increasingly digital

media, to the democratic process. Moreover, these benefits are not typically incorporated into existing regulatory practices.

B. Concentration and Media Capture

A distinguishing feature of media is that they should provide citizens with the information needed to keep governments and other political and economic powers accountable. Of course, powerful individuals usually prefer not to be held accountable and may use threats, promises, and other tactics to induce private media outlets to distort and suppress information. This phenomenon of media capture has been documented in a number of contexts and countries.³⁴³ It often involves an informal understanding between owners of private media outlets and public officials: media outlets will produce news coverage that is favorable to the government, and the government will reciprocate with a policy that is favorable to media outlets. Media capture is not limited to political news: Commercial interests, too, can try to distort financial reporting to their advantage via advertising and other means. Once in a while, the mechanisms behind media capture are revealed to the public, as in the case of the Leveson Inquiry in the UK or the secret recording of the conversation between Benjamin Netanyahu and Arnon Mozes in Israel.

How can media capture be prevented? The kind of informal understanding between media owners and public officials described above is more difficult to achieve if it requires securing the assent of a large number of owners, especially because a news source that maintains appropriate distance from the state and powerful business actors might gain an audience by being trusted to speak truth to power.³⁴⁴ Keeping a low level of concentration in the media industry is thus likely to be an

effective defense against media capture. These benefits of competition are captured in citizen welfare.

C. Inadequacy of Existing Regulation to Prevent Concentration in News Provision

Media concentration is usually achieved through serial acquisitions. Rupert Murdoch's media empire in the US, UK, and other countries is a case in point. The natural remedy is regulation against excessive concentration. However, there is a growing consensus that current regulation is inadequate to achieve this goal, especially in the context of digital platforms.³⁴⁵ This point is recognized by economists, regulatory authorities, and law scholars.

The current debate on the risk posed by digital platforms sometimes seems to imply that before the Internet era the world was in a golden age of free, pluralistic and independent media. However, almost every democracy in the world had high levels of concentration for television, newspapers and radio during the 1990s. Ownership was typically in the hands of either a state-owned entity under the direct control of the government or a small number of private companies, often owned by government-friendly oligarchs.³⁴⁶ The most notable exception to this situation was arguably the United States in the second part of the 20th century, when the news industry displayed low levels of concentration and was mainly perceived to be non-captured. As we shall see next, that was the product of a technological and regulatory environment that no longer exists.

The existing media merger regulatory regime contains two sets of rules, both of which

³⁴³ Schiffrin (2017).

³⁴⁴ Besley and Prat (2003).

³⁴⁵ Polo (2005), Ofcom (2009), and Prat (2015).

³⁴⁶ Djankov et al. (2003).

are inadequate for different reasons. The first set of rules involves platform-specific restrictions on media ownership. Many countries, including the US, have prohibitions against excessive ownership within one particular medium, like terrestrial television. These rules were effective when the set of media platform was small and stable. But difficulties with this approach emerged when cable television was introduced, and then further intensified with the advent of the Internet. In a world where the number of media platforms is expanding and the distinctions between them are blurring, those rules look arbitrary and obsolete. Even worse, there is no obvious way of extending them to include all news-providing platforms.

The other source of regulation is standard competition policy. The problem is that this powerful principle-based set of rules is chiefly meant to protect consumer welfare, as described above. In particular, a merger between media companies is blocked if it is likely to lead to higher prices, lower quality, or some other direct harm to consumers. While this criterion is important and should continue to exist and be applied, it does not cover the indirect harm that a media merger can impose on citizens through an increase in the risk of media capture, as described above.

D. Our Proposal: A Quantifiable Citizen Welfare Criterion for Reviewing Mergers between News Providers

This subcommittee believes that transactions that affect the level of media concentration should be evaluated according to a citizen welfare criterion, alongside the standard consumer welfare criterion. Namely the relevant authority (to be discussed below)

should block mergers that significantly increase concentration in news provision, not because they would increase prices or improve product quality but because they impose an increased risk of media capture.

This assessment can in part be made on the basis of objective measures. Recently, both economic theory and regulatory practice have proposed a similar index of media concentration based on attention shares.

On the theory front, researchers like Prat (2018) have proposed a media power index, based on the idea of attention share. The attention share of a news source is first defined at the individual level as the percentage of time that the individual devoted to a media source divided by the total time the individual devotes to all sources. The overall attention share of a source is then defined as the average attention share that the source commands across all voters in the country.

On the regulatory front, in 2018 the UK's Competition and Markets Authority made a landmark decision.³⁴⁷ The CMA blocked the proposed acquisition of Sky by 21st Century Fox because of a "media plurality consideration." Crucially, the hypothetical effect on plurality was assessed quantitatively. A key element of the assessment was the share of reference metric which is virtually identical to Prat's attention share definition. It showed that the proposed merger would have created a new entity with a larger attention share than of the existing commercial news providers.

The attention share approach has three advantages. First, it is platform-neutral. The unit over which concentration is measured is not a specific platform, but rather citizens' news processing bandwidth. This makes the index robust to the addition of new media or the blurring of borders between two existing media. Second, citizen welfare provides micro-

³⁴⁷ Competition and Markets Authority (2018).

foundations for the measure. In particular, Prat demonstrated that the attention share of a news source determines the upper bound on the ability of that source to influence the voting process through media capture: by putting a cap on attention shares, the regulatory authority can control the worst-case scenario for voters. Third, attention shares can be computed—or at least approximated—with existing data, as the CMA did in the UK and as Kennedy and Prat did for 36 countries.

The attention share approach can be applied to local media, too. Suppose two local news sources wish to merge, or perhaps two media conglomerates that both own a large number of media sources wish to merge, as in the proposed acquisition of Tronc by Sinclair. National-level concentration is obviously not the right measure. Attention shares must be computed in each of the local media markets that would be affected by the merger, with specific reference to local news.

Obviously, the attention share approach is no panacea. At least two important caveats apply. First, just like other concentration indices, such as the Herfindahl index, it should be the beginning, not the end. For instance, it could define thresholds above which the regulatory authority must initiate a plurality investigation.

Second, ownership fragmentation is not the only form of plurality. It is also important that news providers represent the diversity of views and interests in the population. For that goal, however, the right tool is increased public funding, which is discussed in the previous section.

Box 9: Which Regulator Should Enforce Media Plurality?

Which agency should be given the authority to block media mergers on the basis of plurality considerations? The two natural candidates are the default antitrust

authority and the media regulator. The advantages of the latter option are that a media regulator could acquire specialized knowledge in this area and that it would not impose an additional task on the antitrust authority. The potential drawback is that, because of revolving-door practices, industry regulators are often more prone to capture than the antitrust regulator. Joint responsibility between the two agencies is also a reasonable option. In the UK case, the analysis was carried out by Ofcom and the final decision was made by the CMA.

Increasing the Transparency of Digital Platforms

Today’s Internet has brought with it a fragmentation of where people get their news. There is no longer a single or small set of outlets that command consumer attention. Instead, consumers divide their attention across outlets. This has been facilitated by social media, which aggregate news from a variety of sources. The benefit of that is that there is potential for a greater diversity of outlets and also, potentially, a diversity of where news can come from. The cost is a lack of transparency. To see this, consider individual news items aggregated by social media. While some broad source for those news items is given, because consumers may be one-off or, at best, highly intermittent consumers of the media outlet, the source of those news items is not readily apparent or top of mind for consumers. It is not easy to evaluate bias in the news or the reputation of the journalists who generated it. As a consequence, the reputation mechanisms that might ensure that high-accuracy news is screened from low-accuracy news break down, with the adverse consequences outlined above.

Moreover, there is even less transparency with respect to editorial content. As mentioned earlier, the choice of editors in

terms of what to prioritize was very transparent when news outlets had a single product. Today, those choices are hidden within algorithmically optimized personalization used by social media for that purpose. Thus, if there are biases in the presentation of news or omissions of certain types of news content, it is not readily amenable to investigation. To be sure, personalization has many other benefits, but one of the costs is a loss in transparency whereby the choices made by editors or algorithms can be laid bare for scrutiny. This further exacerbates the impact of information asymmetries in the news.

Box 10: Economics of News Quality

News dissemination invariably involves someone other than the news consumer evaluating evidence and facts and filtering those to provide the essence of the news to consumers who have limited time and attention to delve deeper into the subject. Hence, news providers know more about how the news was filtered and how well the facts support the summarized contentions than do consumers of the news.

An efficient news market would allow the consumer to take the story at face value, but, even in inefficient markets, if there are other voices in the market, those voices can conduct their own investigations and check for distortions. So long as those voices are active, the distortions can be brought to light. This has the initial effect of reducing the ‘harm’ that the distorted news story itself might generate. However, it also has a longer-term effect of reducing the reputation of the news outlet that produced distorted news. This causes fewer consumers to rely on it as a source of news. In other words, outlets have an incentive to develop a reputation for quality, but this incentive only operates if

multiple voices and interests can check for deviations from that quality and bring them to light.

In performing its functions as a check on power and authority—whether public or private—the reliability and trustworthiness of the news are critical. However, since the digitization of news content, it has become harder to determine the sources of particular news items or why they might be given priority by certain outlets. Prior to digitization, most news came from outlets that offered a single product to large numbers of consumers. While particular news sources might be hard to identify, the responsibility for the publication lay with the outlet, and, hence, the outlet could be the bearer of a reputation for accuracy. The issue in the reputation mechanism working was whether exit (that is, consumers leaving outlets with a reputation for lower accuracy) served to discipline all outlets for higher accuracy. In other words, the challenge was having enough outlets to provide competitive discipline for each of them. Moreover, ownership and other aspects of bias could be identified and examined. Finally, the editorial decisions as to what news received priority were transparent because the same product was received by all consumers.

A. Our Proposal: Transparency Requirements and Voluntary Labeling

1. Transparency Requirements

We propose two broad mechanisms that can be used to reduce the crisis in transparency we see today—transparency requirements that provide longer-term information on editorial credibility and other factors, particularly in terms of algorithmic editorial choices, and voluntary

labeling that provides underlying information on news sources and carries information on trustworthiness.

There should be transparency around platform editorial decisions. Recall that because consumer attention is scarce, editorial decisions about prioritizing information can direct attention in ways that effectively set the news agenda for the individual. To deal with this issue we propose mandatory transparency about aggregate news targeting, audience reach, ownership, and sponsorship so that individuals, researchers, and civil society groups can understand the forces shaping the news agenda. Covered entities would include existing dedicated news outlets of a certain size, as well as the largest social media sites such as Facebook and Twitter and the largest aggregators such as Google and Apple. The goal here is not to regulate those decisions but to provide visibility into them so that consumers and others can understand editorial decisions and policy—potentially impacting the choice of where they obtain their news information.

The outputs of this transparency would be threefold:

1. The production of periodical reports on aggregate news distribution and consumption on the Internet and how it is influenced by the changing design of the algorithms that control the news feed. For example: anonymized data on the reach of media content and data explaining those information flows.
2. Alerting the public to patterns of news and information consumption (typically of a viral nature) that amplify disinformation, hatred, incitement, or seek to harm democratic processes and institutions.
3. Making certain kinds of internal reports and analyses that platforms have on news consumption available to regulatory agencies and legislators.

To minimize the natural tendency of any media regulator to be captured by platforms, we propose that in addition to the periodic real-time reports, all data obtained by the regulators, its internal decisions and its interactions with the platforms should also be available to the public after a three-year delay.

Box 11: Possible Transparency Tools

1. A public feed, with minimal algorithmic optimization.
2. A private dashboard to users so they can summarily see the nature of their feed (including the content they see, and the content they don't see).
3. An API to let users customize their own news feed based on their preferences and knowledge about algorithmic priorities.
4. A public dashboard visible to the public so they can summarily see the nature of a feed in a given location or globally (like Twitter Trends, but with less editorialization).
5. A more granular dashboard visible to a regulator so it can see how the algorithm customizes a news feed for different clusters/groups of users.

2. Labeling

The information asymmetries that pervade news dissemination resemble others consumers face in evaluating product quality. News is a “post-experience” good whose quality often cannot be assessed until after it has been experienced as true, prescient, enlightening, or not. Labeling of a news source—that is, who has produced it and what entity owns or funds it—is a strategy to reduce post-experience costs. It allows consumers to evaluate the news product before consuming it. Newspaper bylines and mastheads, as well as broadcast ownership reports, all served the function of

signaling to consumers the quality of the news source they were getting. Digital intermediation has made these signals fainter, and new digitally native sources often have no labels at all. For instance, it is often difficult for consumers to rate the energy consumption of household appliances. For this reason, governments and nonprofits around the world have developed certification processes whereby manufacturers can have their products tested and a certified label affixed. This gives consumers the confidence to purchase products based on more information.

For the news, it would not be possible to certify the accuracy of individual news items. However, as the goal is to ensure that consumers have the information necessary to evaluate the news, what could be certified is the source of the news and whether the providers meet certain standards. To be sure, what we are proposing here is not government licensing of news outlets. Indeed, the quality certification could be conducted entirely privately. To give consumers confidence to purchase or consume news products with more understanding of the source of the news, we propose a voluntary labeling scheme administered by an independent news monitor. This will allow consumers to quickly obtain information about the interests of the provider (including ownership and conflicts) as well as provide a way to affix market-based reputations to different news sources.

News outlets that choose to can submit ownership and sponsorship information to the independent news monitor to run the labeling mark. When examining any individual news item, what consumers will see is whether the source of that item has submitted the information. Not having a label does not affect publication of news or even representation of content as news on the Internet. So, it is not a free speech impediment, and it is fully voluntary. For news outlets that wish to exhibit

independence so as to provide confidence in the accuracy of their reporting, labeling will be valuable. Some outlets will not want such certification. Consumers will decide how much weight they place on this factor in news consumption. The broad goal is to give consumers, who need to rely on news for debates and decisions, a fighting chance at understanding the interests of those providing the information. This label, of course, could cover other forms of information, including declared interests in advertising and other revenue. However, we do not take a particular position on the complexity of the labeling scheme here.

The government or a standard-setting body might have to help provide a technical standard by which labeling could be associated with news items published on the Internet. This is a technical challenge, but it is one that has been solved previously on the Internet with respect to e-commerce. In summary, we propose establishment of an independent news monitor that can provide trusted information on news sources at the point of publication for any news item.

Increasing Accountability of Digital Platforms

Section 230 of the Communications Decency Act provides digital platforms with immunity from certain legal claims arising from the content on their networks. This immunity has proven to be an enormous effective subsidy for these intermediaries not enjoyed by legacy media entities like broadcasters, newspapers, or digital native journalistic enterprises that have to contend with more legal exposure. As a libertarian enactment, Section 230 has made the Internet a place of unbridled speech. This should not be confused with freedom from censorship or gatekeeping. As described above, the platforms make many design and content

choices to amplify and circulate certain kinds of content and not others. Moreover, the free-for-all that Section 230 nominally creates has the effect of silencing some voices because of harassment, bullying, and other discourse harms.

Jurisdictions outside the US have adopted versions of Section 230, but none provides as much protection.³⁴⁸ In Europe, platforms have borne more liability and responsibility for removing illegal content. Under the European E-Commerce Directive, for example, intermediaries are exempt from liability for content they host so long as they “play a neutral, merely technical and passive role towards the hosted content.” Once they become aware that any hosted content is illegal, the intermediaries “need to remove it or disable access to it expeditiously.”³⁴⁹ Germany enacted the NetzDG law in 2018, enabling courts to fine social media companies with more than 2 million users up to €50 million if they do not delete posts contravening German hate speech law within 24 hours of receiving a complaint or seven days in more ambiguous cases. There are a number of EU and member state proposals to hold platforms responsible not only for illegal content but also for harmful content and to impose a “duty of care” for managing content in the public’s interest.³⁵⁰

In the US, public upset over platform amplification of harmful content, the failure to filter out harmful and illegal (child pornography, terrorist) content, political disagreements over content salience, and platform usurpation of legacy media advertising revenue has led to calls to reduce the protections of Section 230 and cognate intermediary liability immunities. There seems

to be confusion about what exactly Section 230 does and, correspondingly, mistaken views about what limiting the immunity might achieve in terms of platform design. Section 230 has indeed been a windfall for platforms, allowing them to grow at the expense of legacy media. However, the provision is **not** directly responsible for most of the content that today reduces the salience of fact-based journalism and increases various kinds of noise. It is possible that reducing the scope of protection of Section 230 could incentivize platforms to boost exposure to credible news sources while depressing circulation of, or deplatforming entirely, disinformation that might trigger liability. There would be free speech costs (overmoderation) associated with such a move. For this reason, we propose below to preserve Section 230 protections, but to convert it to a safe harbor subsidy available on condition of compliance with various public interest requirements drawn from media and telecommunications policy traditions.

A. Section 230 Background

Section 230 was enacted as part of the Telecommunications Act of 1996 to govern “Internet service providers.” The ISP to the ordinary publisher in 1996 was something like the scooter to automobiles today: a useful invention, but one hardly on the verge of dominance. Tarleton Gillespie writes: “At the time Section 230 was crafted, few social media platforms existed. US lawmakers were regulating a web largely populated by ISPs and web ‘publishers’—amateurs posting personal pages, companies designing stand-alone websites, and online communities having

³⁴⁸ <https://wilmap.law.stanford.edu/map>

³⁴⁹ Directive 2000/31/EC of the European Parliament and of the Council of 8 June 2000 on certain legal aspects of information society services, in particular

electronic commerce, in the Internal Market (“Directive on electronic commerce”).

³⁵⁰ See, e.g., UK Secretary of State for Digital, Culture, Media & Sport, Online Harms White Paper (April 2019).

discussions.”³⁵¹ Although sometimes viewed as a sweeping libertarian intervention, Section 230 actually began life as a smut-busting provision: an amendment for the “Protection for Private Blocking and Screening of Offensive Material.”³⁵² Its purpose was to allow and encourage Internet service providers to create safe spaces, free of pornography, for children.^{353,354}

The goals at the time of adoption were (1) to give new “interactive computer services” breathing room to develop without lawsuits “to promote the continued development of the Internet,”³⁵⁵ while (2) also encouraging them to filter out harmful content without fear of getting into trouble for under- or overfiltering. Thus Section 230 is both a shield to protect speech and innovation and sword to attack speech abuses on platforms.

The shield part is embodied in Section 230(c)(1): “No provider or user of an interactive computer service shall be treated as the publisher or speaker of any information provided by another information content provider.” This is not blanket immunity for the distribution of content, and indeed platforms are still liable for their *own* content, and for federal crimes and copyright violations related to third-party content. The immunity is really limited to the speech-related harms that publishers ordinarily face such as defamation and intentional infliction of emotional distress. In other words, a platform like Facebook remains liable for distributing child pornography, which is federal criminal content. It also remains liable for Facebook-authored defamatory content. Facebook cannot,

however, be held secondarily liable for defamatory content posted by its users.³⁵⁶

The sword part of Section 230 is contained in Section 230(c)(2)(A): “No provider or user of an interactive computer service shall be held liable on account of—any action voluntarily taken in good faith to restrict access to or availability of material that the provider or user considers being obscene, lewd, lascivious, filthy, excessively violent, harassing, or otherwise objectionable.” This was designed to avoid the paradoxical situation in which an intermediary tries to moderate its platform to reduce harmful content, but then is subject to liability because it has exercised editorial control.

Section 230 is sometimes characterized as a “get out of jail free card” for platforms. According to Chesney and Citron, “Section 230 has evolved into a kind of super-immunity that, among other things, prevents the civil liability system from incentivizing the best-positioned entities to take action against the most harmful content.”³⁵⁷ To be sure, courts have extended immunity in situations that almost certainly go beyond what Congress originally intended. For example, platforms have been excused from transmitting otherwise illegal content even when they have solicited and have clear knowledge of that content. Moreover, platforms that do not function as publishers or distributors (e.g., Airbnb) have also invoked Section 230 to relieve them of liability that has nothing to do with free speech.

B. Proposals to Amend Section 230

There are currently proposals to revise Section 230, weakening the protections it affords to

³⁵¹ Gillespie (2018).

³⁵² H.R. REP. No. 104-223, Amendment No. 2-3 (1995) (proposed to be codified at 47 USC. § 230).

³⁵³ S. REP. No. 104-23, at 59.

³⁵⁴ Pub. L. No. 104-104; see H. Conf. Rep. No. 104-458 (1996).

³⁵⁵ 47 USC. § 230(b)(1).

³⁵⁶ *Zeran v. Am. Online, Inc.*, 129 F.3d 327, 331 (4th Cir. 1997).

³⁵⁷ Chesney & Citron (forthcoming 2019).

platforms. One possibility is to insist on more sword—more care to block or deemphasize harmful speech.³⁵⁸ There are also proposals to weaken the shield and expose platform intermediaries to more liability.³⁵⁹ Senator Mark Warner has floated a relatively narrow proposal to make platforms liable for state-law torts (defamation, false light, public disclosure of private facts), for failure to take down a deep fake or other manipulated audio/video content after the victim had already secured a judgment against the creator of the offending content.³⁶⁰ There have been attempts in the past, including from US state Attorneys General, to carve out other exceptions for the enforcement of state law.³⁶¹ These have foundered on concerns about the incentives this would create for platforms to block too much speech, including speech of importance to journalists like the distribution of gun permits, which some states prohibit publishing.³⁶²

Proposals to weaken Section 230 certainly entail risks to free speech. If platforms

are liable for the content they transmit, they will likely behave in risk-averse ways to remove content that creates or has the potential to create legal exposure.³⁶³ This could silence speakers and have a disproportionate impact on marginal voices or, in some countries, political dissidents or minority groups. In the copyright context, where the Digital Millennium Copyright Act makes platforms potentially liable for copyright violations, they respond overzealously to take-down requests and implement aggressive filtering technologies that block more content than necessary. Early efforts to increase liability for harmful speech may be headed in the same direction. In Germany, the platforms' transparency reports show that they are blocking content in response to the NetzDG law, although there is not yet evidence of overblocking.³⁶⁴

In weighing the costs of any Section 230 contraction against the benefits, it is necessary to identify what those benefits are. Section 230 sceptics may exaggerate these

³⁵⁸ One possibility would be to adopt in the US what has already been adopted in Europe, which is a notice-and-takedown regime that requires platforms to remove content upon notice.

³⁵⁹ See, e.g., Keats Citron & Wittes (2017) (arguing for the conditioning of the benefits of Section 230(c)(1) on reasonable precautions: “No provider or user of an interactive computer service that takes reasonable steps to address unlawful uses of its services shall be treated as the publisher or speaker of any information provided by another information content provider in any action arising out of the publication of content provided by that information content provider.”) The first weakening of Section 230 has already occurred with the Allow States and Victims to Fight Online Sex Trafficking Act (FOSTA), H.R. 1865, 115th Cong. (2018).

³⁶⁰

<https://www.techdirt.com/articles/20180731/10183940336/senator-mark-warner-lays-out-ideas-regulating-internet-platforms.shtml>

³⁶¹ <https://www.eff.org/files/cda-ag-letter.pdf>.

³⁶² <https://www.eff.org/deeplinks/2013/07/state-ags-threaten-gut-cda-230-speech-protections>.

³⁶³ See, e.g., Kumar Katyal (2001) (“Because an ISP derives little utility from providing access to a risky subscriber, a legal regime that places liability on an ISP for the acts of its subscribers will quickly lead the ISP to purge risky ones from its system”); and *Zeran v. America Online*, 129 F.3d 327, 333 (4th Cir. 1997) (“Because service providers would be subject to liability only for the publication of information, and not for its removal, they would have a natural incentive simply to remove messages upon [accusation], whether the contents were defamatory or not”).

³⁶⁴

<https://publications.parliament.uk/pa/cm201719/cmselect/cmsctech/822/82208.htm#footnote-077> (Facebook in the first half of 2018 reported “‘886 NetzDG reports identifying a total of 1,704 pieces of content’, with ‘218 NetzDG reports’ resulting in the deletion or blocking of content. This, Facebook noted, ‘amounted to a total of 362 deleted or blocked pieces of content’... Twitter’s transparency report, covering the same period, indicated that they received a total of 264,818 complaints of which ‘action’ was taken on 28,645. ... Google, meanwhile, received reports relating to 214,827 ‘items’ on YouTube ... of which 56,297 resulted in action, either the item being removed or blocked.”).

benefits in connection with the speech that, while harmful, is not already circumscribed by law (unlike, say, defamation or incitement). This is because Section 230 is not directly related to some of the most problematic discourse harms such as disinformation, outrage, hate speech, radicalization, bullying, and intimidation. To take just one example, Facebook's erstwhile program that allowed advertisers to target users based on phrases like "how to burn Jews" and "Jew hater" did not benefit from Section 230 immunity, and so would not likely be affected by its absence. The same is true for most of the algorithmic amplification discussed above. The UK Online Harms White Paper has recognized that removal or substantial weakening of platform immunity (in that case, the EU E-Commerce Directive immunity, which is less generous than that which Section 230 provides) is a disproportionate response, unlikely to create the right incentives, and very likely to harm free speech.

The focus of this report is the platforms' impact on the production, distribution, and consumption of responsible journalism. Does Section 230 increase the amplification of "noise" (e.g., disinformation and outrage) that crowds out and starves good journalism? It is not clear. As a doctrinal matter, this noise is not illegal. Therefore, because Section 230 does not relieve platforms of otherwise applicable liability, it should not, in theory, increase platform tolerance for noise. On the other hand, the availability of Section 230 to excuse other kinds of speech traffic for

which there would be a liability, such as incitement or defamation, may create a more permissive and careless platform design than would otherwise exist. It is possible that if Section 230 benefits were less generous, platforms would implement safer design features that would also—as a byproduct—amplify truthful or otherwise "beneficial" information. The plan by the UK government to impose a risk-based "duty of care" on platforms seeks to improve the speech environment overall by increasing accountability for certain classes of speech.³⁶⁵

C. Our Proposal: Section 230 as a Quid Pro Quo Benefit

For reasons having to do with risk reduction and harm prevention, legislators may well amend Section 230 to raise the standard of care that platforms take and increase exposure to tort and criminal liability. Since our focus is on improving the conditions for the production, distribution and consumption of responsible journalism, we have a different kind of proposal. We look at Section 230 as a speech subsidy that ought to be conditioned on public interest requirements, at least for the largest intermediaries who benefit most and need it least. It is a speech subsidy not altogether different from the provision of spectrum licenses to broadcasters or rights of way to cable providers or orbital slots to satellite operators.³⁶⁶

The public and news producers pay for this subsidy. The public foregoes legal recourse

³⁶⁵ The new regime's covered "online harms" include terrorist content, child sexual exploitation and abuse, incitement of violence, harassment and cyberstalking, hate crime, encouraging or assisting suicide, the sale of illegal goods or services, revenge pornography, cyberbullying, and children accessing inappropriate material.

³⁶⁶ Cleland (2018) estimates that the net benefit to Google, Facebook and Amazon from Section 230 was \$510 billion in "riskless disruptive innovation per

immunity from civil liability." To this, he adds \$755 billion in "socialized costs of platforms' uneconomic riskless disruptions," although this category reaches far beyond direct effects of Section 230 (e.g., addiction, polarization, election manipulation, devaluation of intellectual property, privacy harms). Of note, he also estimates the benefit of "exemption from all FCC economic and public interest regulation" to be \$31 billion (e.g., public safety, privacy, children and consumer protection, content requirements).

against platforms and otherwise sustains the costs of harmful speech. News producers bear the risk of actionable speech, while at the same time losing advertising revenue to the platforms freed of that risk. Media entities have to spend significant resources to avoid legal exposure, including by instituting fact-checking and editing procedures and by defending against lawsuits. These lawsuits can be fatal, as in the case of Gawker Media.³⁶⁷ More commonly, they face the threat of “death by ten thousand duck-bites” of lawsuits even if those suits are ultimately meritless.³⁶⁸ The monetary value of Section 230 to platforms is substantial, if unquantifiable.³⁶⁹

Section 230 subsidies for the largest intermediaries should be conditioned on the fulfilment of public interest obligations.³⁷⁰ We address transparency and data sharing requirements in this report that should apply generally to the companies within scope. However, there may be additional requirements that a regulator would not want to, or could not, impose across the board. For example, the UK has proposed to require platform companies to ensure that their algorithms do not skew towards extreme and unreliable material to boost user engagement. We would not recommend such a regulation, but it might be appropriate to condition Section 230 immunity on such a commitment for the largest intermediaries.

We believe that Section 230 immunity for the largest intermediaries should be premised on requirements that are well-

developed in media and telecoms law:

- Transparency obligations: Whichever of these requirements did not become part of a mandatory regulatory regime could be made a condition of Section 230 immunity. Platforms should give a regulator and/or the public data on what content is being promoted to whom, data on the process and policies of content moderation, and data on advertising practices. These obligations would go some way towards replicating what already exists in the off-platform media environment by virtue of custom and law. Newspaper mastheads, voluntary codes of standards and practices, use of ombudsmen, standardized circulation metrics, and publicly traceable versioning all provide some level of transparency that platforms lack. For broadcasters, there are reporting and public file requirements especially with respect to political advertising and children’s programming.

- Subsidy obligations: Platforms should be required to pay a certain percentage of gross revenue to support the voucher system discussed in this chapter. In the US, telecoms providers have been required to pay into a Universal Service fund in order to advance public interest goals of connectivity. Public media has not been funded by commercial broadcasters. However, there was once a serious proposal at the advent of digital broadcast television for commercial broadcasters to “pay or play”—that is, to allow them to serve the public interest by subsidizing others to create programming in lieu of doing it themselves.³⁷¹ The nexus between the Section

³⁶⁷ *Bollea v. Gawker Media*, 913 F.Supp.2d 1325 (2012).

³⁶⁸ *Fair Hous. Council of San Fernando Valley v. Roommates.Com, LLC*, 521 F.3d 1157, 1174 (9th Cir. 2008).

³⁶⁹ The difference between the early dismissal of a lawsuit under Section 230 versus later dismissal in the absence of Section 230 protection has been estimated to be worth \$485, 000 each.
<https://static1.squarespace.com/static/571681753c44d8>

35a440c8b5/t/5c6c5649e2c483b67d518293/1550603849958/Section+230+cost+study.pdf

³⁷⁰ The size of the intermediary covered should be internationally harmonized to correspond with definitions in the new UK duty of care and European Copyright Directive.

³⁷¹ Advisory Committee on the Public Interest Obligations of Digital Television Broadcasters (1996). See also Sunstein (1999): 40 (“What if a broadcaster was willing to give \$10 million to PBS in return for

230 benefit and the journalism subsidy is essentially this: Section 230 unavoidably allows many discourse harms and “noise” in the information environment. A subsidy for more quality “signal” goes some distance towards compensating for these harms. It is important to stress that any tax that would be levied on the platforms has to go to the general budget and not be tied in any way to the voucher scheme proposed here—to prevent a situation whereby news outlets have an incentive to lobby directly or indirectly through their editorial agenda to protect the power and revenues of the platforms.

Conclusion

Media is central to the democratic process because it provides the information citizens need to make voting decisions and keep government and powerful players in the public and private sector accountable. This report summarizes the evidence on the effect that the growth of digital platforms is having on news provision and reviews possible policy interventions.

The available evidence highlights a number of points:

- The role of digital media in the democratic process cannot be understood in isolation, as voters, citizens and readers obtain their political information from both old and new platforms.
- There is ample evidence from a variety of historical and geographical contexts that news quality affects voting outcomes and accountability of decisions makers in both the private and public sphere. In turn, news quality—both coverage and impartiality—

worsens when the media industry is concentrated and not independent.

- News is a public good. In the 20th century this good was mostly paid for by advertisers. The most direct effect of the rise of the digital economy has been a crisis of that revenue model. Advertising dollars have inexorably shifted from news producers—mostly newspapers—to digital platforms that do not produce news. This effect is particularly strong for local news. Some subscription-based news providers are thriving, but most citizens are not willing to pay for them.

- There is also more nuanced evidence for a reduction in citizens’ information and engagement and the echo chamber effect. Evidence on the effect of “fake news” is instead very limited. When we analyze policy interventions, we focus on countries where freedom of expression enjoys strong constitutional protection, and we do not consider generalized content regulation policies. We group possible policy interventions into four categories:

- Revenue model. We consider the pros and cons of a number of possible revenue sources: advertising, subscription, philanthropy, and public funding. The most cost-effective and robust solution is a publicly funded voucher system. The funds are allocated directly by the citizens— independently of any government intervention. A cap is placed on the amount that a particular news organization can receive to promote competition.

- Media plurality. All mergers and acquisitions involving news companies should be subject to a news plurality review in addition to the standard antitrust review. Standard competition policy protects direct consumer welfare, and therefore does not take into

every minute, or every thirty seconds, of relief from a public interest responsibility?”).

account the indirect effect that excessive media concentration can cause on citizen welfare. We propose an approach to quantifying news plurality that is neutral to the identity of the owner of the merging entities and to the platform on which news content is delivered. The proposed approach, based on attention shares, has been used in a recent merger decision in the UK.

- **Algorithm regulation.** Developing a new regulatory system that will ensure transparency regarding information flows from news sources, practices and algorithms to the public. This will be done by a new regulatory framework and oversight body to set standards for the disclosure of information and news sources, to develop a means for source-based reputational mechanisms and to bring light to biases and choices in editorial decisions and algorithms for the presentation of the news. These regulators will produce periodical reports on news consumption and the algorithms' design influence on the distribution of news and the behavior of users. The digital regulator will set terms and conditions that will allow publishers and original content creators to negotiate with the digital platforms.

- **Liability exemption.** Digital platforms enjoy a hidden subsidy worth billions of dollars by being exempted from any liability to most of the speech on their platforms (Section 230). We do not propose to repeal Section 230 but rather propose that platforms that would like to enjoy this protection would have to agree to take measures to prioritize content by criteria other than maximization of revenues.

The existence of vibrant independent journalism is an essential part of liberal democracy. Powerful actors in government, politics and the private sector have always tried to capture, weaken or manipulate the press. As the business model of news media outlets has disintegrated and the production and distribution of news has become more

concentrated, there is a growing risk that in many democracies and localities politicians and powerful private actors will not be held accountable for their actions.

Technology will reduce the costs of collecting and analyzing information, but journalists will continue to play the central role in informing the public and holding the powerful to account.

The aim of this report is to offer policy recommendations to preserve independent journalism—not to find ways to aid or subsidize existing news outlets. While other reports and initiatives try to find ways to preserve the news media's revenues from advertising, we believe that this approach is mistaken. Any policy to preserve the free press should try to reduce or eliminate the news media's reliance on politicians, governments, advertisers, large business groups or billionaires.

Our recommendations represent a significant shift from the status quo. We believe that increased public funding of news media through a competitive process that caps the amount which can be allocated to a single news outlet will create a more competitive and independent news ecosystem.

The success of a handful of large national publications or small niche publications in reaching adequate profitability through the subscription model is encouraging. Nonetheless, for the most part citizens are not willing to pay for this public good—meaning we need to rethink the economic model of the news media.

Recent events across the Western world have demonstrated the fragility of the liberal democratic order, and we believe that waiting longer to see if market forces alone can maintain the free press in the 21st century may be a risky choice

References

- Amazeen, Michelle A., and Ashley R. Muddiman. 2018. "Saving media or trading on trust? The effects of native advertising on audience perceptions of legacy and online news publishers." *Digital Journalism* 6(2): 176–195.
- Anderson, Simon P. 2012. "Advertising on the Internet." In *The Oxford Handbook of the Digital Economy*, ed. Martin Peitz and Joel Waldfogel. Oxford University Press.
- Angelucci, Charles, Julia Cagé, and Michael Sinkinson. 2017. "Media Competition and News Diets." CEPR Discussion Paper #12198.
- Arrow, K. J. 1962. "Economic Welfare and the Allocation of Resources for Invention." In *The Rate and Direction of Inventive Activity: Economic and Social Factors*. Princeton University Press.
- Athey, Susan, and Mark Mobius. 2012. "The Impact of News Aggregators on Internet News Consumption: The Case of Localization." Working Paper No. 3353. Stanford Graduate School of Business, January 11.
- Athey, Susan, Emilio Calvano, and Joshua Gans. 2018. "The Impact of the Internet on Advertising Markets for News Media." NBER Working Paper 19419. National Bureau of Economic Research. <https://doi.org/10.3386/w19419>.
- Bakshy, Eytan, Solomon Messing, and Lada A. Adamic. 2015. "Exposure to Ideologically Diverse News and Opinion on Facebook." *Science* 348:1130.
- Beattie, Graham, Ruben Durante, Brian Knight, and Ananya Sen. 2017. "Advertising Spending and Media Bias: Evidence from News Coverage of Car Safety Recalls." NBER Working Paper No. 23940. National Bureau of Economic Research.
- Beckett, C. 2018. "Facebook's newsfeed changes: a disaster or an opportunity for news publishers?" *London School of Economics Blog*, January 20.
- Bell, Emily. 2016. "Facebook Is Eating the World." *Columbia Journalism Review*, March 7.
- Bell, E. J., T. Owen, P. D. Brown, C. Hauka, and N. Rashidian, N. 2017. "The platform press: How Silicon Valley reengineered journalism." Tow Center for Digital Journalism, Columbia University. https://www.cjr.org/tow_center_reports/platform-press-how-silicon-valley-reengineered-journalism.php/.
- Benson, Rodney. 2016. "Are Foundations the Solution to the American Journalistic Crisis?" Media Ownership Project Working Paper 2016-001, New York University.
- Benson, R., and M. Powers. 2011. *Public media and political independence. Lessons for the Future of Journalism from Around the World*. New York: New York University.
- Bergström, Annika, and Maria Jervelycke Belfrage. 2018. "News in Social Media." *Digital Journalism* 6(5): 583–598. <https://doi.org/10.1080/21670811.2018.1423625>.
- Birnbaumer, Bill. 2018. *The Rise of Nonprofit Investigative Journalism in the United States*. Routledge Research in Journalism. Routledge.

- Besley, T. J., and Prat, A. 2003. "Pension fund governance and the choice between defined benefit and defined contribution plans." Working paper, Institute for Fiscal Studies WP03/09.
- Bollinger, L. C. 2010. *Uninhibited, Robust, and Wide-Open. A Free Press for a New Century*. Oxford University Press.
- Cagé, Julia. 2015. *Sauver Les Médias: Capitalisme, Financement Participatif et Démocratie*. La République Des Idées. Seuil (English version: *Saving the Media. Capitalism, Crowdfunding and Democracy*, Harvard University Press, 2016).
- Cagé, Julia. 2017. "Media Competition, Information Provision and Political Participation: Evidence from French Local Newspapers and Elections, 1944-2014." CEPR Discussion Papers 12198.
- Cagé, Julia. 2018. *Le Prix de La Démocratie*. Fayard (English version: *The Price of Democracy*, Harvard University Press, forthcoming).
- Cagé, Julia, and Béatrice Mazoyer. 2019. "News Breaker or Audience Raider? The Ambivalent Relationship between Twitter and News Media."
- Cagé, Julia, Nicolas Hervé, and Marie-Luce Viaud. 2019. "The Production of Information in an Online World: Is Copy Right?"
https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2672050.
- Cairncross, F. 2019. "The Cairncross Review: A sustainable future for journalism."
https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/779882/021919_DCMS_Cairncross_Review_.pdf.
- Calzada, Joan, and Ricard Gil. 2016. "What Do News Aggregators Do? Evidence from Google News in Spain and Germany." *SSRN*, September 13.
https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2837553.
- Campante, Filipe, Ruben Durante, and Francesco Sobbrío. 2017. "Politics 2.0: The multifaceted effect of broadband Internet on political participation." *Journal of the European Economic Association* 16(4): 1094–1136.
- Caplan, R., and Boyd. 2018. "Isomorphism through algorithms: Institutional dependencies in the case of Facebook." *Big Data & Society*, February 14.
- Carlson, M. 2018. "Automating judgment? Algorithmic judgment, news knowledge, and journalistic professionalism." *New Media and Society* 20(5): 1755–1772.
<https://doi.org/10.1177/1461444817706684>.
- Chen, Adrian. 2017. "The Human Toll of Protecting the Internet from the Worst of Humanity." *New Yorker*, Jan. 28. <http://perma.cc/AAK8-ZHEV>.
- Chen, Yuyu, and David Yang. 2019. "The Impact of Media Censorship: 1984 or Brave New World?" *American Economic Review*.
- Chesney, Bobby, and Danielle Citron. Forthcoming. "Deep Fakes: A Looming Challenge for Privacy, Democracy, and National Security."

- Cleland, Scott. 2018. "The Unfair and Deceptive Online-Offline Playing Field of Divergent US Competition and Consumer Protection Policy." Submission for US FTC Fall 2018 Hearings on "Competition and Consumer Protection in the 21st Century" Topic #2: "Competition and Consumer Protection in Communication, Information, and Media Technology Networks," FTC Project # P181201, August 7.
- Competitions and Markets Authority. 2018. "21st Century Fox / Sky Merger Inquiry: Final Report." June 5. <https://www.gov.uk/cma-cases/twenty-first-century-fox-sky-merger-european-intervention-notice>.
- Cornia, Alessio, Annika Sehl, Felix Simon, and Rasmus Kleis Nielsen. 2017. "Pay Models in European News." Reuters Institute for the Study of Journalism, University of Oxford. <http://www.digitalnewsreport.org/publications/2017/pay-models-european-news>.
- DeVito, Michael A. 2017. "From Editors to Algorithms." *Digital Journalism* 5(6): 753–773
- Di Tella, R., and I. Franceschelli. 2011. "Government advertising and media coverage of corruption scandals." *American Economic Journal: Applied Economics* 3(4): 119–51.
- Djankov, S., C. McLiesh, T. Nenova, and A. Shleifer. 2003. "Who owns the media?" *Journal of Law and Economics* 46. <https://scholar.harvard.edu/files/shleifer/files/media.pdf>.
- Drago, Francesco, Tommaso Nannicini, and Francesco Sobbrino. 2014. "Meet the Press: How Voters and Politicians Respond to Newspaper Entry and Exit." *American Economic Journal: Applied Economics* 6(3): 159–88. <https://ideas.repec.org/a/aea/aejapp/v6y2014i3p159-88.html>.
- European Audiovisual Observatory. 2017. "The EU Online Advertising Market – Update 2017." <https://rm.coe.int/the-eu-online-advertising-market-update-2017/168078f2b3>.
- Falck, Oliver, Robert Gold, and Stephan Heblich. 2014. "E-Lectons: Voting Behavior and the Internet." *American Economic Review* 104(7): 2238–65. <https://doi.org/10.1257/aer.104.7.2238>.
- Flanagin, A. J., and M. J. Metzger. 2007. "The role of site features, user attributes, and information verification behaviors on the perceived credibility of web-based information." *New Media and Society* 9(2): 319–342.
- Fletcher, Richard, and Rasmus Kleis Nielsen. 2017. "Are News Audiences Increasingly Fragmented? A Cross-National Comparative Analysis of Cross-Platform News Audience Fragmentation and Duplication." *Journal of Communication* 67(4): 476.
- Gans, J. S., and S. Stern. 2010. "Is there a market for ideas?" *Industrial and Corporate Change* 19(3): 805–837.
- Gao, Pengjie, Chang Lee, and Dermot Murphy. 2018. "Financing Dies in Darkness? The Impact of Newspaper Closures on Public Finance." <https://www.brookings.edu/wp-content/uploads/2018/04/Murphy-et-al..pdf>.
- Gavazza, Alessandro, Mattia Nardotto, and Tommaso Valletti. 2018. "Internet and Politics: Evidence from U.K. Local Elections and Local Government Policies." *Review of Economic Studies*. <https://doi.org/10.1093/restud/rdy028>.

- Gentzkow, Matthew. 2006. "Television and Voter Turnout." *Quarterly Journal of Economics* 121(3): 931–72.
- Gentzkow, Matthew, Edward Glaeser, and Claudia Goldin. 2006. "The Rise of the Fourth Estate: How Newspapers Became Informative and Why It Mattered." In *Corruption and Reform: Lessons from America's Economic History*, ed. Edward L. Glaeser and Claudia Goldin. National Bureau of Economic Research.
- Gentzkow, Matthew, and Jesse M. Shapiro. 2008. "Competition and Truth in the Market for News." *Journal of Economic Perspectives* 22(2): 133–54.
<http://ideas.repec.org/a/aea/jecper/v22y2008i2p133-154.html>.
- Gentzkow, Matthew, Jesse M. Shapiro, and Michael Sinkinson. 2011. "The Effect of Newspaper Entry and Exit on Electoral Politics." *American Economic Review* 101(7): 2980–3018.
- George, L. M., and J. Waldfogel. 2006. "The New York Times and the market for local newspapers." *American Economic Review* 96(1): 435–447.
- Ghosh, Dipayan, and Ben Scott. 2018. "#DIGITALDECEIT: The Technologies Behind Precision Propaganda on the Internet." *New America*, January. <https://www.newamerica.org/public-interest-technology/policy-papers/digitaldeceit/>.
- Gillespie, Tarleton. 2018. "Platforms Are Not Intermediaries." *Georgetown Law Technology Review* 2(198): 201.
- Greer, J. D. 2003. "Evaluating the credibility of online information: A test of source and advertising influence." *Mass Communication and Society* 6(1): 11–28.
- Grimmelmann, James. 2018. "The Platform Is the Message." *Georgetown Law Technology Review* 2(217): 227.
- Guess, Andrew, Brendan Nyhan, and Jason Reifler. 2018. "Selective Exposure to Misinformation: Evidence from the consumption of fake news during the 2016 US presidential campaign." European Research Council, January 9.
<http://www.dartmouth.edu/~nyhan/fake-news-2016.pdf>.
- Hamilton, James T. 2004. *All the News That's Fit to Shell: How the Market Transforms Information into News*. Princeton University Press.
- Hamilton, James T. 2016. *Democracy's Detectives: The Economics of Investigative Journalism*. Harvard University Press.
- Keats Citron, Danielle, and Benjamin Wittes. 2017. "The Internet Will Not Break: Fixing Section 230 Immunity for Bad Samaritans." *Fordham Law Review* 86(401).
- Kennedy, Patrick, and Andrea Prat. Forthcoming. "Where Do People Get Their News?" *Economic Policy*.
- Klonick, Kate. 2018. "The New Governors: The People, Rules, and Processes Governing Online Speech." *Harvard Law Review* 131(1598): 1637.
- Kumar Katyal, Neal. 2001. "Criminal Law in Cyberspace." *University of Pennsylvania Law Review* 149(1003): 1007–08.

- Lessig, L. 2015. *Republic, Lost: The Corruption of Equality and the Steps to End It*. Grand Central Publishing.
- Levin, Sam. 2017. "Facebook told advertisers it can identify teens feeling 'insecure' and 'worthless'." *Guardian*, May 1.
- Livio, Oren, and J. Cohen. 2018. "'Fool me once, shame on you': Direct personal experience and media trust." *Journalism* 19(5): 684–694.
- Marwick, Alice, and Rebecca Lewis. 2017. "Media Manipulation and Disinformation Online." *Data & Society*, May 15.
- McCabe, Brian. 2017. "High-Dollar Donors and Donor-Rich Neighborhoods: Representational Distortion in Financing a Municipal Election in Seattle." *Urban Affairs Review* 55(4).
- McCabe, Brian J., and Jennifer A. Heerwig. 2018. "Diversifying the Donor Pool: Did Seattle's Democracy Vouchers Program Reshape Participation in Municipal Campaign Finance?" <https://osf.io/zu8wh/#>.
- McChesney, R., and J. Nichols. 2010. *The Death and Life of American Journalism: The Media Revolution that Will Begin the World Again*. Nation Book.
- McChesney, R. W., and V. W. Pickard. 2011. *Will the Last Reporter Please Turn Out the Lights: The Collapse of Journalism and what Can be Done to Fix It*. New Press.
- Milhorance, Flávia, and J. Singer. 2018. "Media Trust and Use among Urban News Consumers in Brazil." *Ethical Space: The International Journal of Communication Ethics* 15(3/4).
- Napoli, Philip, and Robyn Caplan. 2017. "Why Media Companies Insist They're Not Media Companies, Why They're Wrong, and Why It Matters." *First Monday* 22. <https://perma.cc/LC7N-8JX6>.
- Newton, Casey. 2019. "The Trauma Floor: The secret lives of Facebook moderators in America." *Verge* (Feb. 25), <https://www.theverge.com/2019/2/25/18229714/cognizant-facebook-content-moderator-interviews-trauma-working-conditions-arizona>.
- Nolte. 2019. "Fact Check: Media Blacklisters NewsGuard Promotes Fake News, Repeatedly Violates Own Standards." *Breitbart*, Jan. 24. <https://www.breitbart.com/the-media/2019/01/24/fact-check-media-blacklisters-news-guard-promotes-fake-news-violates-standards/>.
- Ofcom. 2009. "Second public service broadcasting review: Putting viewers first." Technical report. Ofcom's Second Public Service Broadcasting Review, January. https://www.ofcom.org.uk/_data/assets/pdf_file/0016/51325/psb2statement.pdf.
- Pariser, Eli. 2011. *The Filter Bubble: How the New Personalized Web Is Changing What We Read and How We Think*. Penguin Publishing Group.
- Pasquale, Frank. 2016. "Platform Neutrality: Enhancing Freedom of Expression in Spheres of Private Power." *Theoretical Inquiries in Law* 17:487. <https://perma.cc/9X5W-793A>.
- Pennycook, Gordon, and David G. Rand. 2017. "Who falls for fake news? The roles of analytic thinking, motivated reasoning, political ideology and bullshit receptivity." *SSRN*, January. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3023545.

- Petrova, M. 2011. “Newspapers and parties: How advertising revenues created an independent press.” *American Political Science Review* 105(4): 790–808.
- Pew Research Center. 2014. “Journalism Partnerships. A New Era of Interest.” December 4. <https://www.journalism.org/2014/12/04/journalism-partnerships/>.
- Pew Research Center. 2016. “State of the News Media Report 2016.” June 15.
- Picone, Ike, Cédric Courtois, and Steve Paulussen. 2015. “When News Is Everywhere.” *Journalism Practice* 9(1): 35–49. <https://doi.org/10.1080/17512786.2014.928464>.
- Polo, Michele. 2005. “Regulation for pluralism in the media markets.” In *The Economic Regulation of Broadcasting Markets: Evolving Technology and the Challenges for Policy*, ed. Paul Seabright and Jurgen von Hagen. Cambridge University Press.
- Prat, Andrea. 2015. “Media capture and media power.” In *Handbook of Media Economics*, ed. S. P. [Anderson](#), J. [Waldfogel](#), and D. [Strömberg](#). Elsevier, North Holland.
- Prat, Andrea. 2018. “Media Power.” *Journal of Political Economy* 4(126).
- Prat, Andrea, and David Strömberg. 2005. “Commercial television and voter information.” CEPR Discussion Paper No. 4989. <https://ssrn.com/abstract=772002>.
- Prior, Markus. 2007. *Post-Broadcast Democracy: How Media Choice Increases Inequality in Political Involvement and Polarizes Elections*. Cambridge Studies in Public Opinion and Political Psychology. Cambridge University Press.
- Raj, P., and G. Rolnik. 2018. “Internet and the Business Model of Journalism.” Working paper.
- Reich, R., C. Cordelli, and L. Bernholz. 2016. *Philanthropy in Democratic Societies: History, Institutions, Values*. University of Chicago Press.
- Reich, R. 2018. *Just Giving: Why Philanthropy Is Failing Democracy and How It Can Do Better*. Princeton University Press.
- Reuters Institute. 2017. “Digital News Report 2017.”
- Reuters Institute. 2018. “Digital News Report 2018.”
- Ricci, Francesco, Lior Rokach, and Bracha Shapira. 2011. “Recommender Systems: Introduction and Challenges.” In *Recommender Systems Handbook*, ed. Francesco Ricci, Lior Rokach, Bracha Shapira, and Paul Kantor.
- Roth, A. E. 2002. “The economist as engineer: Game theory, experimentation, and computation as tools for design economics.” *Econometrica* 70(4): 1341–1378.
- Rubado, M. E., and J. T. Jennings. 2019. “Political Consequences of the Endangered Local Watchdog: Newspaper Decline and Mayoral Elections in the United States.” *Urban Affairs Review*, April 3. <https://doi.org/10.1177/1078087419838058>.
- Schiffrin, Anya, ed. 2017. “In the Service of Power: Media Capture and the Threat to Democracy.” *CIMA*, August 16. <https://www.cima.ned.org/publication/media-capture-in-the-service-of-power/>.
- Seamans, Robert, and Feng Zhu. 2013. “Responses to entry in multi-sided markets: The impact of Craigslist on local newspapers.” *Management Science* 60(2): 476–493.

- Sen, Ananya, and Pinar Yildirim. 2015. "Clicks and Editorial Decisions: How Does Popularity Shape Online News Coverage?" Digital seminar, Toulouse School of Economics, November 4. https://editorialexpress.com/cgi-bin/conference/download.cgi?db_name=EARIE42&paper_id=37.
- Shao, Chengcheng, Giovanni Luca Ciampaglia, Onur Varol, Kaicheng Yang, Alessandro Flammini, and Filippo Menczer. 2018. "The Spread of Low-Credibility Content by Social Bots." *Nature Communications* 9(1): 4787. <https://doi.org/10.1038/s41467-018-06930-7>.
- Snyder, James M., and David Stromberg. 2010. "Press Coverage and Political Accountability." *Journal of Political Economy* 118(2): 355–408.
- Stromberg, David. 2004. "Radio's Impact on Public Spending." *Quarterly Journal of Economics* 119(1).
- Sundar, S. Shyam, Sunetra Narayan, Rafael Obregon, and Charu Uppal. 1998. "Does web advertising work? Memory for print vs. online media." *Journalism & Mass Communication Quarterly* 75(4): 822–835.
- Sunstein, Cass. 1999. "Private Broadcasters and the Public Interest: Notes Toward a 'Third Way'." University of Chicago Law School, Olin Law and Economics Program, Discussion Paper No. 65.
- Sunstein, Cass. 2008. "The Architecture of Serendipity." *Harvard Crimson*, June 4. <https://www.thecrimson.com/article/2008/6/4/the-architecture-of-serendipity-this-years/>.
- Sunstein, Cass R., and Adrian Vermeule. 2009. "Conspiracy Theories: Causes and Cures." *Journal of Political Philosophy* 17:202.
- Trussler, M., and S. Soroka. 2014. "Consumer demand for cynical and negative news frames." *International Journal of Press/Politics* 19(3): 360–379.
- Tufekci, Zeynep. 2018. "It's the (Democracy-Poisoning) Golden Age of Free Speech." *WIRED*, January 16. <https://www.wired.com/story/free-speech-issue-tech-turmoil-new-censorship/>.
- Turcotte, Jason, C. York, J. Irving, R. M. Scholl, and R. J. Pingree. 2015. "News recommendations from social media opinion leaders: Effects on media trust and information seeking." *Journal of Computer-Mediated Communication*, 20(5): 520–535.
- Westphal, David. 2018. "Journalism's New Patrons: Newspapers deepen embrace of philanthropy." *Columbia Journalism Review*, January 8.
- Wilding, D., P. Fray, S. Molitorisz, and E. McKewon. 2018. "The Impact of Digital Platforms on News and Journalistic Content." *Centre for Media Transition* 52. <https://www.accc.gov.au/system/files/ACCC%20commissioned%20report%20-%20The%20impact%20of%20digital%20platforms%20on%20news%20and%20journalistic%20content%2C%20Centre%20for%20Media%20Transition%20%282%29.pdf>
- Wu, Tim. 2016. *The Attention Merchants*. New York: Vintage Books.

- Yuan, Elaine. 2011. "News Consumption Across Multiple Media Platforms." *Information, Communication & Society* 14(7): 998–1016.
<https://doi.org/10.1080/1369118X.2010.549235>.
- Zingales, Luigi. 2016. "Are Newspapers Captured by Banks? Evidence From Italy." *ProMarket*, Stigler Center at the University of Chicago Booth School of Business, May 12.
<https://promarket.org/are-newspapers-captured-by-banks/>.
- Zingales, Luigi, and Guy Rolnik. n.d. "The Role of Media in Information Age." Working paper.
- Zittrain, Jonathan. 2014. "Response, Engineering an Election: Digital Gerrymandering Poses a Threat to Democracy." *Harvard Law Review Forum*, June 20.

George J. Stigler Center for the Study of the Economy and the State
The University of Chicago Booth School of Business

Stigler Committee on Digital Platforms

Privacy and Data Protection Subcommittee

Report

July 2019

Subcommittee members

Lior Strahilevitz, University of Chicago (Chair)
Lorrie Faith Cranor, Carnegie Mellon University
Florencia Marotta-Wurgler, New York University
Jonathan Mayer, Princeton University
Paul Ohm, Georgetown University
Katherine Strandburg, New York University
Blase Ur, University of Chicago

Contributing Co-Authors

Sebastian Benthall, New York University
Filippo Maria Lancieri, University of Chicago
Jamie Luguri, University of Chicago
Mark Verstraete, New York University

DISCLAIMER: The purpose of this report is to foster discussion and illuminate key problems in privacy and security. The report is a document reflecting the subcommittee consensus, but not every part of the report reflects the views of each individual author.

*** The Committee in-person meetings were partially supported by a grant from the Alfred P. Sloan Foundation, whom we also thank for supporting this project**

Table of Contents

Executive Summary 209

Introduction..... 214

I. An Overview of Some Key Problems in Privacy and Security..... 215

 A. Why Market Solutions Often Fail 216

 B. Shortcomings of Notice and Choice 217

 C. Harmonization Challenges 218

 D. Limited Resources for Deterring Privacy Violations..... 219

 E. Data Security Threats..... 220

II. Existing Approaches to Protecting Privacy and Security..... 223

 A. FTC Enforcement..... 223

 B. California law..... 224

 C. Proposed Domestic Legislation 224

 D. GDPR..... 225

 E. ALI Privacy Principles Project 226

 F. Industry self-regulation..... 227

III. Paths Forward: Default Rules, Constraints on Dark Patterns, and Improved Authentication 228

 A. Default Rules for Privacy Policies 228

 1. Determining the Contents of Privacy Defaults 231

 2. Data about Consumer Privacy Preferences and Expectations..... 231

 3. Toward “Consumertarian Default Rules” 234

 B. Dark Patterns and Manipulation 237

 1. Defining Actionable Manipulation 238

 2. A Data-Driven Approach to Spotting Problematic Dark Patterns 242

 3. Legal Framework 248

 C. Institutional Implementation Considerations for Default Rules and Dark Pattern Regulations ... 255

 D. Mitigation for Security Threats Caused by Data Breaches 257

 E. Privacy-Preserving Measurement and Coordination 262

Conclusion..... 264

Executive Summary

Any serious effort to think systematically about the emerging policy issues surrounding digital platforms must contend with a host of privacy and security challenges. It is becoming increasingly clear that core aspects of the American approach to privacy, such as faith in industry self-regulation, the embrace of a sectoral approach involving overlays of state and federal laws, heavy reliance on notice and choice, and the lack of remedial powers for the nation's primary privacy regulator, leave the world's largest economy ill-equipped to vindicate important interests in privacy and security. The weaknesses and idiosyncrasies of the American approach impede efforts to harmonize global privacy law, and that lack of harmonization threatens the free flow of data in international e-commerce.

There are several key reasons why neither industry self-regulation nor notice-and-choice approaches to protecting privacy and security succeeded. Firms that collect personal information do not internalize all the harms associated with consumer privacy and security breaches, nor do they have adequate incentives to take account of how decisions they make may affect the interests of consumers who are not their customers. These problems are compounded by the highly technical nature of privacy and security decisions, the difficulty of establishing clear causal links between particular practices and subsequent consumer harms, and the challenges of monitoring and evaluating firms' key investments in data security. Notice and choice, by contrast, largely fails as a strategy because it does not live up to the promises its name suggests. More precisely, the standard approach is to provide consumers with a lengthy wall of legalese text that they will not read and then to pretend that consumers are making meaningful choices that vindicate their interests in response. Moreover, firms have become skilled at developing choice architectures that nudge consumers toward options that benefit company profitability but may not reflect consumers' actual preferences or expectations.

There are a number of existing reform proposals in the privacy and security domain, some quite laudable, others more dubious in terms of their benefits. Our subcommittee's approach is not to provide a roadmap for comprehensive national privacy legislation. Rather, we have focused our attention on developing three complementary approaches to protecting privacy and security interests. These approaches are modular, so they can be integrated into any existing proposals, though each would also be effective as a stand-alone reform.

The first of our three proposals advocates for the use of data-driven contractual default rules in privacy and security. Default rules are starting points for contractual relations between parties. For example, they govern the scope of data collection and retention of a consumer's personal information by a platform. By definition, any default rule can be modified by the mutual agreement of the contracting parties. We can contrast these default rules with mandatory rules, which provide contractual rights and responsibilities that cannot be altered by the parties. While there is an important role for mandatory rules to play in instances where consumer

preferences are very homogenous, where collective action problems arise, or where significant information asymmetries exist, as with highly technical and complex matters, default rules are a vital tool in the regulatory arsenal. Default rules are particularly useful in the many instances where consumers have heterogeneous preferences, where consumers or firms possess relevant information to which regulators do not have ready access, and where substantial externalities are not present in a transaction.

We propose that the content of contractual default provisions will depend on the articulated preferences of ordinary consumers as measured by scientifically rigorous survey instruments. Based on our own pilot testing of such instruments involving the privacy and security practices of Facebook, Google, Amazon, and other platforms, consumers will often – but not always – prefer and expect default provisions that enhance their privacy and security. In privacy and security settings there will be many instances in which it is appropriate for the law to use “consumertarian” default rules – i.e., the legal defaults preferred or expected by a majority of consumers. While these default rules will not always reflect the preferences of platforms and other firms that are contracting with consumers, these firms will have the ability to convince consumers to waive the protections required by default if they can convince consumers that waiving such protections will make consumers better off. In that sense, consumertarian default rules and subsequent efforts to convince consumers to waive protections that firms find problematic for their business models will function as information-forcing devices that reduce unexpected surprises, encourage dialogue, and prevent firms from seeking to engage in privacy and security practices that might prompt customer backlash. Because many default protections will be sticky under these circumstances, and firms will have incentives to be selective about which rights they ask consumers to waive, the result of our proposal on net will be to heighten privacy and security protections for consumers. This proposal for consumertarian default rules is somewhat similar to the Privacy-by-Default regime that was enacted as part of the European General Data Protection Regulation (GDPR), though we believe our approach does a better job of creating predictable, transparent benchmarks.

Under our approach, the protections granted by default rules can be waived by consumers, but only via a process that takes the notion of informed consumer consent more seriously than the law does currently. Our discussion of dark patterns, which we will now explain, provides a yardstick for measuring the sufficiency of such waivers.

Our second set of proposals concerns dark patterns. Dark patterns are user interfaces that make it difficult for users to express their actual preferences or that manipulate users into taking actions that do not comport with their preferences or expectations. Examples of dark patterns abound in privacy and security. For example, Google Maps repeatedly asks users whether a site that they regularly return to should be labeled “home” or “work.” If the user agrees to label the geolocation, then the pop-up queries will cease. If the user clicks on “Not Now” then there will be more queries a few days later. The result is that the application may be so persistent in asking

for users to confirm personal information that they will eventually relent to prevent further nagging, and not because they want to share this information. Alternatively, many firms make it easy to alter one's defaults in a way that benefits the firm financially but quite cumbersome to alter default settings in a way that will make the customer less profitable for the firm. Some firms attempt to hide what they are doing, for example when firms add items to a consumer's online shopping cart without asking the consumer, or when they add a hidden service to an online transaction by preselecting the addition so as to ensnare consumers who they know will quickly click through a screen to complete the transaction. Firms may employ intentionally ambiguous terminology in an effort to confuse consumers into opting for a service they do not want, or they may manipulate consumers by targeting acute emotional vulnerabilities. One attribute all these dark patterns share is a tendency to exploit "System 1" (quick, instinctive) decision-making and suppress more deliberative "System 2" thought processes.

We surmise that firms have done rigorous beta testing of dark pattern techniques because they have become prevalent in recent years. Academics and non-profit organizations have documented the prevalence of dark patterns, but until now no researchers have published research that examines the efficacy of these dark patterns in prompting consumers to make choices that are inconsistent with what consumers would choose in a neutral decision architecture. Our report fills a significant gap in the literature, thanks to an extensive data collection effort by Jamie Luguri and Lior Strahilevitz. Those authors, after obtaining IRB approval, exposed a census-weighted, nationally representative sample of 1,762 Americans to a decision-making framework in which the control group was offered an easy yes / no choice over whether to sign up for an expensive identity theft protection plan, an experimental group was subjected to rather mild dark pattern interventions, and a second experimental group was subjected to the aggressive use of dark patterns. Both dark pattern conditions were designed to prompt consumers to agree to pay for an identity theft protection plan that few members of the control group wanted.

The bottom-line results from this dark pattern experiment were striking. Employing mild dark patterns increased the percentage of consumers who ultimately agreed to accept the data protection plan by 228% (from roughly 11% to 26%). Employing aggressive dark patterns increased the percentage of consumers who agreed to accept the data protection plan by 371% (from roughly 11% to 42%). In other words, in both the mild and aggressive dark pattern environments, it was more likely than not that consumers were agreeing to sign up for the plan our researchers were selling them because of the dark pattern, and not because of an underlying desire to purchase the plan itself.

Notably, the experiment's use of aggressive dark patterns generated the equivalent of a customer backlash. Consumers in the aggressive dark pattern condition had their moods adversely affected, they were less likely to agree to participate in follow-up research by the same researchers when given the opportunity to do so, and they were more likely to withdraw from the

experiment, forfeiting their entitlement to be compensated for taking the survey. This data provide evidence that the market itself constrains the use of aggressive dark patterns somewhat. On the other hand, the use of mild dark patterns generated either no such effects or far smaller effects – with respect to their mood at the conclusion of the experiment, consumers exposed to the control group decision architecture and the mild dark pattern decision architecture were statistically indistinguishable. To summarize, then, firms face significant incentives to avoid using the most blatant and annoying dark pattern strategies, but the use of more subtle dark patterns seems to be all upside from the perspective of a firm’s bottom line. These mild strategies seem to substantially increase the percentage of consumers who will sign up for a good or service without alienating those customers, at least in the short run.

The more highly educated consumers were, the less vulnerable they were to having their choices manipulated via dark patterns. These effects were statistically significant, and highly troubling. Dark patterns work on many people, but lower socio-economic status individuals are especially vulnerable to them. Less educated individuals were particularly susceptible to mild dark patterns.

In light of these findings, our report advocates a per se legal rule that will apply to many situations involving the use of dark patterns to prompt consumers to share personal information. Where a firm’s choice architecture more than doubles the percentage of users who agree to share information, when compared with a neutral choice architecture, consumers’ consent to share such information is not valid. Moreover, dark pattern tactics that satisfy this “more likely than not” test should be treated as unfair and deceptive practices in trade, which are proscribed by federal and state consumer protection laws. There may be other domains in which a dark pattern is highly problematic but does not satisfy this per se test. To deal with those situations, we offer a multi-factor balancing test that can be used to identify the dark patterns that are most likely to diminish consumer well-being. That test looks to the extent to which a dark pattern unnecessarily raises the transaction costs of opting out of a protection that large numbers of consumers expect or prefer, the extent to which the dark pattern targets problematic consumer vulnerabilities, or the extent to which a dark pattern is hidden rather than transparent.

The final proposal in our report focuses on mitigating the security threats that are caused by data breaches. A single data breach at one platform or digital service can present major problems for other platforms and services. The reason stems from password reuse. Consumers frequently use identical or very similar login credentials at multiple sites. As a result, hackers may obtain credentials from one site and then quickly try to use those same pilfered credentials to gain access to various other sites and platforms. One main way that firms currently try to protect themselves is by purchasing stolen credentials, which creates perverse incentives.

Among the various reforms to be considered, private data breach clearinghouses are preferable, given existing technological constraints. Ideally, such clearinghouses could use techniques that employ advanced mathematics to test whether user passwords are repeated across

multiple sites without disclosing login credentials. Firms would be required or strongly incentivized to make their own data available for queries in order to ping a centralized database. The clearinghouse proposal would encounter some challenges, ranging from the paramount need to protect the clearinghouse as a single source of failure to the technical challenges associated with identifying similar but not identical passwords that are being used across multiple sites. Still, the subcommittee has concluded that, on balance, such an approach is superior to the viable alternatives that the report discusses.

The subcommittee report concludes with an overview of the need to balance privacy interests against transparency considerations. The other subcommittee reports on digital platforms laud the benefits of greater data transparency to evaluate the role of platforms in markets, media, and politics. These are worthwhile goals. At the same time, well-meaning efforts to make data transparent can expose individuals' sensitive private information. Differential privacy and secure multiparty computation are two of the most promising technologies for protecting privacy in these domains, and policymakers should ensure that these strategies are implemented alongside transparency initiatives.

Introduction

It is rare for a week to pass in the United States without some new privacy scandal. Sometimes the scandal involves a massive breach of sensitive personal information that was inadequately protected. Other times it involves a deviation from promises to keep information confidential or secure. And other instances involve unexpected or unwanted data collection, retention, or use, or the repurposing of data to further a different, unanticipated objective. These privacy scandals are costly to society. They result in embarrassment, psychological harm, a loss of control, and financial damage. As a result of these harms, consumers may disengage from using digital platforms, they may engage in less online commerce, they may become more reluctant to say what they think, they may curtail their communications with intimates, and they may take steps to disguise their identities through the use of deception and other self-help strategies. Failing to protect privacy generates significant social costs; that much is clear. What's less clear is the answer to the question, "What should be done?"

In addressing this question it is important to recognize that there are issues about which consumers' preferences and needs are heterogeneous. Some consumers are quite concerned about the privacy of certain personal information and others are not. Some would gladly trade cost or convenience or appealing product features for less privacy, and others would strongly object to such swaps. Some care a great deal about keeping information about their political beliefs confidential but are less reluctant to share information about their intimate associations. For others this hierarchy is reversed. Sometimes consumers want to exercise choices or control regarding their private information but find the options confusing or difficult to exercise.

A response to this question must also recognize that differences in individual privacy expectations, preferences, and needs may be attributable to societal or cultural factors. Individuals have differing experiences with technology, socioeconomic backgrounds, and nationalities, any of which may shape their distinct privacy perspectives. A response must be sensitive to these differences and must avoid magnifying inequities and other societal concerns. While the sheer complexity of the privacy and security decisions people are confronted with on a daily basis may justify a significant role for technocratic expertise, it is important to remain cognizant of the dangers of regulatory capture, the slowness of bureaucratic decision-making, and the dangers of paternalistic decision-making that is out of step with evolving consumer preferences and practices.

I. An Overview of Some Key Problems in Privacy and Security

The American approach to privacy makes the country something of an outlier, at least in comparison to other liberal democracies. Notable features of the present American approach include: (1) its failure to enact comprehensive consumer privacy legislation at the federal level, relying instead on sector-specific federal privacy law and state privacy law; (2) the relatively weak regulatory powers of the primary federal enforcement agency in the United States; (3) a reliance on class action suits to deter and punish various privacy violations and security breaches; and (4) a hierarchy of values that generally emphasizes free speech interests. As a result, significant gaps exist in American privacy law, and whether a particular act generates legal consequences depends on which part of the United States a party finds itself in and what the incentives are for regulatory agencies or private litigants to file suit. Moreover, in some instances, significant privacy harms may not lend themselves to constitutionally permissible remedies in the federal courts.

Foreign observers sometimes believe that American privacy protections are nonexistent because there is not a single authoritative and universal law that protects a particular set of privacy interests. That is an error. Some key components of global privacy law—like data breach notification laws or the Fair Information Practices—have emerged in the United States and been widely adopted elsewhere. There are some American states that have embraced or are moving towards comprehensive consumer privacy legislation, such as the recently adopted California Consumer Privacy Act. And there are some areas of the economy—such as the telecommunications, health care, criminal justice, and education sectors—where federal privacy protections are relatively robust by global standards. American privacy laws are relatively complex, and in most of the country they are relatively lax or incomplete, but legal protections for privacy are far from nonexistent.

Another distinct feature of the American approach to privacy is a relatively high level of stated government confidence that market incentives will safeguard consumer privacy and security interests without the need for regulatory interventions. From that perspective, firms will compete over privacy and security, and the companies that fail to deliver what consumers want will be driven out of business. Whatever the merits of a hands-off approach to privacy in theory, in practice it has largely failed. To be sure, we do see evidence that firms in industries where privacy and security are particularly important (such as financial services or cloud computing) invest more resources in privacy and security than firms in industries where those concerns are less salient.³⁷² In some extreme and salient cases, executives have lost their jobs because of privacy snafus. But prominent market failures persist, justifying smart regulatory interventions.

³⁷² Marotta-Wurgler (2016).

A. Why Market Solutions Often Fail

There are several important potential sources of market failure. First, many of the harms associated with privacy and security breaches are not internalized by firms. A breach might not be discovered until months or years after it occurs—if ever. Another reason for this dynamic relates to the difficulties of tying many secondary harms to a particular failure by a company. For example, a consumer might be victimized by identity theft but unable to trace the identity theft to any particular breach. Perhaps the same consumer information was breached several times in recent years, making it difficult to establish a causal connection between a corporate practice and a negative consequence. The more breaches occur, the more vexing this problem becomes.

A second reason why harms aren't fully internalized involves the costs to consumers of monitoring the consequences of privacy snafus. As corporate errors become more commonplace, there is less sustained attention devoted to each one in the media and other sources. With the exception of catastrophes like the Equifax breach or Cambridge Analytica scandal, it is likely that a new privacy goof will come along quickly to push the most recent breach out of a consumer's news feed. Even individualized data breach notifications, whose prompt dissemination is mandated by law, can get lost or ignored in a sea of other mandated disclosures.

Third, decision-making over privacy is so complicated that consumers often do not understand the nature of the bargains they are making when they agree to share personal information in order to use a non-priced good or service. A major reason is the problem of unanticipated uses and failures of imagination. Thanks to advances in computing power and data mining techniques, new uses of old data are regularly discovered. Many consumers do not even understand how data they surrender is used contemporaneously. It is therefore unreasonable to expect them to anticipate the sorts of uses that data they share today will be put to in three, five, or ten years. If and when such data is used in a way that disadvantages consumers, they will be unable to tie that harm to the company they shared it with, even if the company still survives as a going concern. When consumers come into direct contact with dozens of companies and indirect contact with hundreds, it is unrealistic to expect them to perform due diligence by reading privacy policies and devoting resources to understanding all their implications. As a result, markets function poorly to constrain companies that engage in controversial privacy practices.

Fourth, a great deal of personal information is held by firms with which consumers have no direct contact. These may be data brokers, credit agencies, and vendors who have purchased data; or prospective employers, landlords, and insurers with whom a contractual relationship never materialized. Any consumer efforts to penalize firms that have behaved badly will be dampened by the indirect nature of these relationships, and intermediary firms therefore will not internalize fully the externalities associated with data breaches.

Fifth, in many instances markets may fail to satisfy consumer preferences because of the “collective privacy” problem.³⁷³ That is, information about each consumer is distributed, leaving the person whose information is at issue with incomplete control over the dissemination of that information, even if she takes all reasonable precautions to preserve its confidentiality. For example, suppose one sibling decides to share her information with a genetic testing firm. The other sibling, whose genetic attributes can now be inferred, may object strenuously to the placement of this information in the hands of a proprietary firm. But the law provides the objecting sibling with essentially no recourse. Along similar lines, even if an individual elects not to join a social media platform, her friends and relatives may reveal a good deal of information about her that can then be used in ways that affect her life. Photos of her may appear, and stories about what she said or did may be widely shared. The problem is particularly acute for children, whose parents and peers often share sensitive information about them online without fully understanding the future ramifications.

Sixth, firms often face dampened incentives to differentiate themselves with respect to privacy, especially via consumer-facing advertising. Except in rare instances, privacy is not the most salient attribute for consumers when considering a good or service. A firm that seeks to raise awareness about the relative strengths of its privacy practices compared to its competitors has to worry about spooking consumers into avoiding the product or service line altogether. As a result, while there have occasionally been splashy marketing campaigns emphasizing a particular company’s commitment to privacy, these sorts of campaigns have been either short-lived or implemented on behalf of niche products with small market shares.

Finally, consumers are often left to bear the burden of data breaches themselves. While consumers may hear about data breaches after they occur, they rarely know what steps to take in reaction. Data breaches can have cascading consequences as hackers leverage the information stolen to cause further damage. Consumers face difficulties evaluating these cascading consequences and mitigating the harms. Furthermore, it is difficult for consumers to evaluate a firm’s investments in data security. Indeed, firms have incentives to obfuscate their efforts on the data security front because information they share with consumers will inevitably make its way to hackers, who can use this information to discover vulnerabilities.

In summary, while firms plainly have some incentives to deliver privacy and security to consumers, there are lots of reasons to explore regulatory interventions.

B. Shortcomings of Notice and Choice

One variant of a market-oriented approach to protecting privacy is notice and choice. The notice and choice approach dictates that consumers be informed of what firms are doing with personal information so they can make their own choices about how and whether that data

³⁷³ Strahilevitz (2010); Barocas & Levy (2018).

should be used. There is nothing wrong with this approach in the abstract. Indeed, the approach has the potential to accommodate the heterogeneity of consumer preferences and values. In practice, however, notice and choice has largely failed as a regulatory strategy for some of the aforementioned reasons. For example, the kinds of choices that consumers are being asked to make are too complex and too speculative, especially given the open-ended nature of the privacy policy language that is placed in front of them.

Notice and choice, as presently practiced, fails for a broader reason as well. The information that gets disclosed to consumers is too voluminous and intricate for consumers to read, understand, and then make a well-informed choice that reflects their values, preferences, and interests.³⁷⁴ Were consumers to have relationships with just a few entities, such investment of time and cognitive resources might be realistic. In a world where consumers have dozens of apps on their smartphones, visit numerous web sites, and bring an increasing number of connected devices into their homes and workplaces, notice and choice fails. Consumers click “I agree” without reading and hope for the best.

Various proposals have emerged to enhance notice and choice through techniques like visceral notice or highlighting unexpected terms.³⁷⁵ While these approaches represent an improvement over the status quo, they too run up against constraints. For example, visceral notices that grab consumers’ attention may work well initially, but once the technique is employed in even a modest number of interactions it appears to lose its force. When everything is visceral, nothing is. The limited empirical testing that has been done on these sorts of techniques also suggests they are of limited utility.

A further problem with notice and choice is that a consumer’s choice might not be freely exercised. Companies have become increasingly sophisticated at nudging consumers into acts and omissions that result in privacy practices that are good for corporate bottom lines but may be at odds with what consumers would prefer were they to understand fully the nature of their interactions. The legal system in the United States has been quick to deem consumers as having consented to terms and conditions despite genuine reason to believe that their consent is the product of manipulation rather than free choice.³⁷⁶ Below, we will examine this issue in much more depth and identify appropriate interventions.

C. Harmonization Challenges

Digital Platforms operate in global marketplaces. This presents real challenges for firms. First, because national and multinational regimes for regulating privacy and security vary, practices that are lawful in one country may be unlawful in another. Second, for efficiency reasons, it may make sense for personal information to be gathered in one country and then

³⁷⁴ McDonald & Cranor (2008); Bakos et al. (2014).

³⁷⁵ Calo (2012).

³⁷⁶ Lemley (2006).

analyzed in another. Third, even firms that primarily operate domestically in the United States must deal with increasingly diverse state law regimes. This means that cumulative obligations regularly arise under state law, compliance burdens increase, and legal conflicts can occur—especially where extraterritorial law enforcement or national security interests are involved.

International tensions involving data privacy have already manifested. Most prominently, the Court of Justice of the European Union’s decision in *Schrems* struck down the US-EU Safe Harbor Agreement that permitted personal data of EU citizens to flow across the Atlantic. The quick negotiation and implementation of the EU-US Privacy Shield Framework as a replacement prevented an immediate crisis, but there are reasons to worry that these kinds of issues will recur, particularly if the GDPR and American law continue to diverge. Indeed, new legal challenges to Privacy Shield are presently pending in *Schrems II* before the Court of Justice. The results could be quite negative for American technology companies that have substantial business activities overseas.

The costs of complying with different sets of state laws also can be high. States may legislate without giving adequate consideration to the compliance burdens that are imposed on firms that operate nationally. These kinds of costs may be borne more easily by well-capitalized digital platforms that enjoy economies of scale. For new start-ups the costs may be much more significant, and the burdens associated with complying with every state law could function as barriers to entry. In other instances, the existence of an unusual law in one state may create surprises for firms. For example, Illinois has a unique law limiting the collection and use of biometric information, one that enables people whose rights are violated to file a civil suit and recover minimum statutory damages. The law has given rise to a number of class action suits involving the use of facial recognition and fingerprint applications. In the first half of 2019, class actions invoking the Illinois law were being filed at a rate of approximately one per day. While some companies were no doubt aware of the Illinois law when they were designing their products to integrate biometrics, other companies were caught unaware and now face the prospect of significant legal liability. Ignorance of the law is no excuse, but imposing obligations of this sort at the national level provides greater notice to firms that wish to operate in a national marketplace.

D. Limited Resources for Deterring Privacy Violations

To the extent that the American legal system has a primary cop on the beat of privacy and security, that entity is the Federal Trade Commission (FTC).³⁷⁷ But the Commission is hamstrung by its limited authority to police privacy. The FTC primarily addresses privacy issues pursuant to its power to regulate unfair and deceptive practices in trade. Congress has imposed significant procedural burdens on administrative rulemaking under that authority, such that the FTC must resort to individual enforcement actions and case-by-case articulation of privacy law.

³⁷⁷ Hoofnagle (2016).

Congress has also constrained the FTC's monetary penalty authority in unfairness and deception cases to violations of consent decrees, effectively creating a "two-strikes" system—a firm's first privacy violation results in a consent decree, and only subsequent violations can result in penalties. Thus, for example, the FTC's ability to fine Facebook \$5 billion in 2019 for its actions with respect to Cambridge Analytica depends entirely on a 2012 consent decree between Facebook and the FTC, which arose from a prior privacy snafu. Were it not for that earlier consent decree, Facebook would not face monetary penalties for its actions, at least not insofar as the FTC is concerned.

It is also widely believed that the FTC lacks the resources to regulate privacy effectively. Only 46, or four percent, of the agency's 1,141 employees work in the Division of Privacy and Identity Protection. And only about five of the FTC's employees are technology experts. The FTC's professional consumer protection staff is commendable but as the role of digital platforms in our economy has grown, regulatory resources have not kept up. As a result, many major problems involving privacy and security are discovered and publicized by other entities, such as investigative journalists, state attorneys general, academic researchers, and experts hired by class action attorneys. Resource constraints render the FTC a reactive regulator rather than a proactive one.

Class action litigation has achieved some success, mainly from a deterrence perspective. However, class actions rarely result in the victims of privacy and security snafus being made whole for the associated psychological costs. Such litigation is lucrative for the attorneys who bring suits successfully and has incidental benefits for the privacy-related non-profits that sometimes receive large payouts pursuant to settlements. Plaintiffs' lawyers will gravitate to bringing cases under statutes that contain minimum statutory damages provisions. As a result, privacy interests vindicated by those statutes are enforced rather aggressively, and comparable interests that are embedded in legal frameworks without such provisions tend to get inadequate attention. Two results are skewed priorities and a risk of over- and under-deterrence in areas of privacy law.

E. Data Security Threats

Repositories of consumer data have come under increasing attack in recent years due to a unique confluence of technological and market factors. Rather than providing a framework for omnibus interventions that would tackle the myriad of problems associated with data security, our subcommittee has chosen to focus on a very important but relatively soluble problem, one revolving around the use of passwords as access credentials for digital platforms. The reason for this focus is that problems around access credentials have similar roots to many of the problems pinpointed above in relation to data privacy: the over-reliance on the user to address all the shortcomings of their personal choices.

Despite repeated efforts to advance alternative schemes for authentication, passwords remain ubiquitous on the web because no single alternative scheme is superior to passwords on every relevant metric, and passwords have advantages in familiarity and deployment. Large-scale data breaches have unfortunately become a common occurrence, ranging from the breach of credit reporting agency Equifax to the breaches of hundreds of different websites, including major online services like Yahoo. Collectively, billions of credentials (pairs of usernames and passwords) have been stolen. The most obvious consequences of these breaches are to the service that was breached. Once a breach is discovered, which unfortunately can take years, the breached company typically forces a password reset for all affected accounts. However, in many cases other companies are at risk. While each website or service typically runs its own authentication system (i.e., most users have an account unique to each service, rather than a federated identity controlled by a single service-agnostic entity), information stolen from data breaches can inextricably bind the fates of otherwise independent accounts at competing companies.

Most websites limit the number of log-in attempts for a particular account, which might normally prevent guessing attacks for all accounts other than those with the most predictable passwords. However, when an account holder has had their account credentials on another service breached and the accounts can be associated across services (e.g., because they share the same username or registration email address), guessing attacks can become a major threat. This is because users frequently reuse passwords across accounts. Even when the passwords are not exactly the same, attackers can often guess similar passwords in fewer guesses than would set off alarms on most services. As a result, attackers use credentials stolen in data breaches to make highly targeted guesses, frequently compromising accounts successfully.

A core issue is that the whole authentication system overly relies on users policing their own password use. Crucially, companies and organizations do not know which of their users have chosen a similar password for their account with a competitor—something that prevents “security coordination” among companies. Compounding this problem, users themselves are many times not fully aware of the risks involved in password re-use. The makers of password manager software have tended to design the software not to castigate users for reusing passwords, worried that being too heavy-handed would encourage the user to stop using that password manager and potentially turn to a competitor. The abundance of websites that require logins and passwords for trivial matters, and the feeling that these are low-stakes decisions, encourage re-use and diminish incentives for monitoring data breaches. In addition, companies may not disclose data breaches to users in sufficient detail or in a timely manner. Even when notified of a password-related data breach, users may not fully understand the causes of, or mitigations for, their vulnerability to account compromises related to password reuse.

In short, the system places a significant burden on users to constantly police the security of their accounts in a context where many might not even recall that they reused the same password elsewhere, or they might not understand why doing so is problematic. As a harm-mitigation

strategy, companies like Facebook seek out (and in some instances purchase) lists of stolen credentials from hackers to proactively lock their own users' accounts. This state of affairs represents a market failure. Authentication systems are siloed, each controlled by a single company. Following data breaches, other companies and organizations are at risk because of password reuse, yet they do not know for which accounts this is the case, nor do the users themselves. And when companies buy information from hackers to try to fill these gaps, they provide a (marginal) incentive for hackers to obtain sensitive personal information in the future.

This dynamic takes place in a context where data breaches are becoming increasingly costly. Data about an individual stolen in a data breach can also be weaponized against data subjects or even their employers, friends, and relatives in other ways, such as by enabling an attacker to engender misplaced trust by demonstrating knowledge of this information. For example, phishing attacks (fraudulent emails designed to cause an individual to disclose personal information or credentials, or to provide access to a resource) have become increasingly targeted through the inclusion of user-specific personal information (e.g., knowledge of friends' names, the inclusion of the user's phone number or past purchases) in recent years. These targeted attacks are known as spear phishing. We expect that spear phishing attacks will become even more targeted in the near future as attackers amass data stolen in data breaches, including an individual's personal information lost directly in the data breach, information gleaned from accessing accounts (e.g., an email inbox) using credentials stolen in a data breach, and the ability to send messages from breached accounts in the first place. These types of information hold the potential to fool a victim into mistakenly thinking an email is legitimate. Spear phishing can have disastrous consequences for the victim, as exemplified by the exfiltration of John Podesta's emails during the 2016 presidential campaign. Unfortunately, individuals will often know, at best, the broad types of information stolen in a data breach, not the specific information nor comprehensive steps for remediating the breach. As a result, targeted phishing attacks using stolen information will become increasingly difficult for individuals to recognize.

II. Existing Approaches to Protecting Privacy and Security

A. FTC Enforcement

The Federal Trade Commission seeks to deter unfair and deceptive consumer privacy and data security practices through enforcement of section 5 of the FTC Act against entities engaged in practices that are likely to cause substantial injury to consumers. In addition, the FTC is charged with enforcing the Children’s Online Privacy Protection Act of 1998 (COPPA), the US-EU Privacy Shield, and, joint with another agency, the Fair Credit Reporting Act (FCRA) and Gramm-Leach-Bliley Act (GLBA). These last two statutes, among other things, seek to protect financial information. Almost all FTC actions result in a settlement or default judgment. As a result, there are barely any judicial decisions involving the FTC in this area.

Since 1995 the FTC has enforced the self-imposed duties that firms laid out in privacy policies. Though the FTC did not require any specific terms in such policies, it did require that firms live up to the promises made in those agreements as failure to do so would be considered a deceptive trade practice. Most FTC activity in this space involved actions of this type. These include actions against Facebook for failing to keep its promise to keep consumer information private, against AshleyMadison.com for suffering a massive data breach and thus failing to keep its promise to keep their customer information secure, and against Uber after it breached its promise to monitor employee access to consumer data, among others. As parts of the consent decree with the FTC, the violating companies agree to correct their violating behavior by, for example, implementing comprehensive data-security assessment in the case of AshleyMadison.com, or undergoing a comprehensive privacy program and obtaining regular independent audits in the case of Uber. The consent decree also will usually provide “fencing-in relief” that sweeps more broadly than the initial violation, providing the FTC a foundation for subsequent enforcement actions involving the same firm.

A limitation to this approach is the fact that firms making few or no promises in their privacy policies can avoid some FTC actions. In a handful of cases, the FTC has brought actions against firms for deceptive omissions, as it did in an action against Nomi Technologies, a firm providing in-store tracking technology that promised to give consumers an opt-out choice but later failed to do so. In a minority of cases, the FTC has brought actions against firms that it alleged were engaging in unfair practices. While “unfairness” is an intentionally broad concept, section 5 requires additional elements for an unfairness claim (including harm to consumers) that can be difficult to prove in some privacy cases.

The FTC has brought around 200 privacy cases to date, representing an average of about ten cases a year. While the FTC is strategic in the actions that it brings, often focusing on large companies, there is reason to doubt the overall deterrent effect given the small number of cases and the small penalties.

B. California law

Perhaps the most notable state regulation is the California Consumer Privacy Act (CCPA), which becomes effective on January 20, 2020. The CCPA applies to for-profit businesses that have annual gross revenue in excess of \$25 million, receive or disclose information of more than 50,000 California residents, or derive more than 50% of their annual revenue from selling California residents' data. Personal data is defined broadly, and includes information that relates to, describes, or could be used to identify California residents or households. The CCPA provides exceptions for de-identified or aggregate data, but the scope of those terms remains ambiguous.

The CCPA requires businesses to inform consumers, at or before the point of collection, of the categories of personal data being collected and the purposes the data will be used for. In addition, businesses must provide notice and the opportunity to opt out of third-party data sharing. Online privacy policies must describe a consumer's CCPA rights, and businesses are prohibited from including contract provisions that limit a consumer's enforcement remedies. Consumers also have the right to request information concerning the source of the data, the specific types of data being collected, and the categories of data being shared with third parties. In general, the CCPA does not create specific data security rules, though it does require businesses to maintain reasonable security procedures appropriate to the nature of personal information being collected.

The California Attorney General is given primary enforcement authority, though consumers also have a (very) limited individual and class action right of action regarding the law's data breach provisions. The AG may seek injunctive remedies, or impose fines of up to \$7,500 for each intentional violation and \$2,500 for unintentional violations.

Even before the CCPA was enacted, California had heightened privacy requirements on the books. The California Online Privacy Protection Act (CalOPPA), for example, requires online services to post privacy policies, which sets up subsequent FTC or attorney general enforcement. The state's data safeguard law also requires reasonable data security precautions and is enforceable as a violation of California's unfair competition law.

C. Proposed Domestic Legislation

Privacy and security concerns have helped spark creative legislative proposals. Numerous pieces of proposed legislation have been introduced within the last year at both the state and federal level. The list of pending bills is long and grows by the month. Here we wish to highlight a couple of significant recent proposals that are particularly germane to this report's emphasis and recommendations.

In January 2019, Washington State Senator Carlyle introduced the Washington Privacy Act, which was approved by a majority of the state Senate but stalled in the state's House of

Representatives. The law was similar to the CCPA in many respects, though the version approved by the Senate differed from the California law in that it lacked a private cause of action, ceding enforcement to the state attorney general's office. The proposed legislation would also restrict the use of facial recognition technologies in both the private and public sectors. It is likely that such legislation will be reintroduced in the 2020 legislative session. Other states such as New York, Texas, and Massachusetts are also considering similar legislation, and Maine and Nevada have enacted more limited versions.

In April 2019 Senators Warner and Fischer introduced the federal Deceptive Experiences to Online Users Reduction (DETOUR) Act. The bill is, as far as we know, the first federal legislation that takes aim at the problem of dark patterns online and their use to convince consumers to part with personal information. The legislation would only apply to online services that have more than 100 million monthly users. Among other provisions, the bill would make it unlawful "to design, modify, or manipulate a user interface with the purpose or substantial effect of obscuring, subverting, or impairing user autonomy, decision-making, or choice to obtain consent or user data." The proposed legislation also includes a safe harbor for firms engaging in conduct that establishes "default settings that provide enhanced privacy protections to users or otherwise enhance[s] their autonomy and decision-making ability." The FTC would have authority to enforce the law, which treats violations as unfair or deceptive practices in trade. A key dimension of the bill is how to identify user interfaces that have the "substantial effect of impairing user autonomy, decision-making, or choice." That is a problem to which our report devotes significant attention. We provide a proof of concept for a relatively clear rule that can differentiate permissible persuasion from impermissible dark patterns.

D. GDPR

The European Union's General Data Protection Regulation ("GDPR") is perhaps the most ambitious privacy and security law to date. The GDPR imposes strict rules for data processing and creates significant rights for data subjects. The GDPR defines personal data as any information relating to an identifiable or identified person, and provides heightened protections for sensitive personal data, including race, health, sexual orientation, and criminal records.

The GDPR applies to all EU entities and to non-EU entities that process personal data of EU residents either in connection with an offering of goods or services or for the purposes of behavioral tracking. EU member state data protection authorities will be the primary enforcers and can assess fines as high as 4% of an entity's global annual revenue or €20 million, whichever is greater. The GDPR also includes a private right of action, a class action right, and a third-party right of action that allows non-profits and trade associations to sue on behalf of consumers.

The GDPR distinguishes between data "controllers" and "processors." Data controllers determine the purpose of data collection and the means of data processing. Data processors

receive and handle data from controllers. Controllers generally must comply with heightened obligations. Controllers, for example, must demonstrate compliance not only for themselves, but also for the processor. In addition, controllers are also generally responsible for monitoring their processor's behavior. The GDPR requires that data controllers inform consumers about how their data will be used, how long the data will be kept, and how they can exercise their data-related rights. Controllers must also disclose any third party that will handle consumer data. It also requires that all privacy policies be written in clear, plain language. Furthermore, the GDPR incorporates principles of privacy by design and privacy by default. The former requires data controllers to implement privacy protective measures like pseudonymization and data minimization into the design of products and services. The latter obligates controllers to establish default rules that reduce the risks associated with the unnecessary collection, processing, retention, or dissemination of personal information.

The GDPR also provides consumers with a number of rights. Consumers can request that data controllers provide an explanation of what data they are collecting and how they are using it. In addition, consumers can request that data be corrected or erased. Though controllers may share personal data with third parties to fulfill the original purpose of processing, they may not do so for a different purpose unless the consumer consents or the controller uses a new legal basis. Personal data can also be transferred outside the EU, but generally only under circumstances where substantially equivalent privacy protections remain in effect. One important such ground is an "adequacy decision" by the European Commission, which certifies that the non-EU country has adequate personal data protection.

Finally, the GDPR requires firms to maintain certain data security practices. Data controllers, for example, must engage in data minimization by collecting only the data necessary to carry out a particular task. Firms must also implement appropriate technical measures to ensure security appropriate to the risk of accidental or unlawful destruction, loss, or unauthorized disclosure or access of personal data. Although the GDPR does not impose any specific security practices, it recommends pseudonymization, encryption, and regular security testing, among other things. In addition, personal data can only be retained for as long as necessary to fulfill the original basis for collection and processing, though there are a few exceptions for when data is used for the public interest.

E. ALI Privacy Principles Project

The American Law Institute's (ALI) Data Privacy Principles project began in 2013 under the leadership of Paul Schwartz and Daniel Solove. The ALI project (which the subcommittee chair has worked on as Adviser) is an ambitious effort to identify best practices in data privacy and security and to use these principles to help lawmakers and regulated entities develop sound principles for managing personal information. The project has made considerable progress in the six years since it began and was approved by the ALI Membership in 2019, marking the

project's completion. Since the process for drafting ALI Principles is highly collaborative, and the norms of the organization nudge its work product towards consensus ideas, its drafting has offered opportunities for many stakeholders to shape its content.

The core provisions of the current draft of the Privacy Principles deal with the transparency of use of personal information, the appropriate role and process for providing notice and obtaining individual consent, the scope of confidentiality duties, the importance of use limitations, and the placement of limits on data retention and downstream transfer. The Principles also advocate for user access, error correction, and data portability rights that are grounded in existing frameworks like the Fair Information Practices and the GDPR. Other provisions in the principles would impose obligations on firms that collect or process personal data to employ reasonable data security protections. The Privacy Principles project embraces the imposition of some new duties on the operators of digital platforms, and it attempts to do so in a relatively comprehensive way. Given its recent adoption, we do not endeavor here to replicate the comprehensiveness of that effort. Nor are we as constrained as the ALI is to identify existing approaches that are already working in some parts of the world or nation. We hope the approach we lay out below will supplement the ALI's work by being both narrower in its scope and also more ambitious, nimble, and imaginative about how thoughtful legal regulation can proceed.

F. Industry self-regulation

There have been a number of attempts by firms, sectors, and trade groups to self-regulate, including attempt to adopt "trustmarks" (such as TRUSTe and BBB Online) and codes of conduct. These efforts have mostly failed for lack of adoption, limited consumer protections, and lax enforcement and monitoring.

A related approach has been the development of privacy protections through multi-stakeholder processes. The National Telecommunications and Information Administration (NTIA), for example, worked with industry groups to develop privacy guidelines for the Organization for Economic Cooperation and Development (OECD). More recently, through cooperation between government, industry leaders, and consumer advocates, the NTIA has developed reports outlining best practices for data security and the commercial use of facial recognition, among other technology topics. Multi-stakeholder cooperation is attractive in theory because solutions would reflect industry needs and knowhow as well as consumer privacy interests. Yet multi-stakeholder regulation can only work if there is scope for agreement among key stakeholder groups, incentive to reach agreement, and self-monitoring and policing. These conditions have yet to manifest themselves. Our candid impression is that multi-stakeholder processes have often been employed more as a pretext for delay than as a substantive alternative to privacy regulation.

III. Paths Forward: Default Rules, Constraints on Dark Patterns, and Improved Authentication

The enactment of a comprehensive privacy and data security law at the federal level is a major item on the congressional agenda. Various thoughtful proposals have emerged from legislators and from organizations with substantial expertise. Our hope in this section is to offer a possible central organizing principle for such legislation, one that differs in emphasis from both the American status quo and a European approach that combines prescriptive rules with enormous discretion to bureaucratic decision-makers. One aim of our proposal is to accommodate the significant heterogeneity that exists among consumers with respect to personal information while condemning corporate practices that use that heterogeneity as a pretext for manipulating consumers into making choices that undermine their welfare.

Our second goal is to create legal safe harbors for competing firms to collaborate with respect to a common threat stemming from state- and non-state actors seeking to penetrate existing data security protections. Existing efforts to facilitate the sharing of information between competitor firms have not solved the legal problems associated with this sharing, nor have they developed adequate incentives to facilitate such sharing, like well-designed clearinghouses for personal information that may be mirrored to control access to the interfaces and databases of many companies. Our proposal attempts to make progress on that key challenge.

A. Default Rules for Privacy Policies

In the United States, contractual relations between individuals and firms are typically governed by bodies of law that include both mandatory and default rules. Mandatory rules govern parts of a contractual relationship that are determined by law and that provide rights and obligations that cannot be waived. An example of such rules is the duty of good faith under Article 2 of the Uniform Commercial Code (UCC), which governs the sales of goods. Default rules are also determined by law, but they can be modified contractually by the parties. Default rules can be created by various organizations, be they public or private, and then adopted contractually by the parties. Yet it is the state that typically selects the content of both default rules and mandatory rules.

Most contract rules are default rather than mandatory, since default rules allow transacting parties to opt out and customize aspects of their contractual relationships to suit their preferences. Their main function is to provide off-the-rack rules that allocate rights and risks between parties. Oftentimes default rules mimic the terms that most contracting parties would have agreed to mutually, leaving the cost of opting out to be borne by an idiosyncratic minority. Other times default rules are information-forcing—they penalize contracting parties with private information by encouraging them to reveal such information by contracting out of the unfavorable default. Another vital function of default rules is to fill in gaps in incomplete

agreements. Default rules govern when contracts are silent as to a particular obligation. Parties that fail to specify an alternative obligation are thus governed by the default, even when their agreement says nothing about it.

In their traditional setting, default rules are designed to minimize contracting costs by supplying gap fillers and minimizing the costs of opt outs. These purposes fit somewhat awkwardly into the consumer setting because consumers are typically ignorant of both the content of default rules and many terms of the agreements they enter. Consumers might even be mistaken about the relative value of different rules. In some settings, consumers also will lack the bargaining power necessary to prompt firms to change the terms of a contract. Most consumer contracts are offered in long, standardized forms, usually on a take-it-or-leave-it basis. Privacy policies are perfect examples of this.

In other contexts, as a result of insights from behavioral economics, default rules have been designed to be “sticky” and to encourage socially desirable behaviors. Sticky defaults take advantage of the costs of opting out. A prominent example of this can be found in the context of retirement savings, where the employees are defaulted into saving plans and only a few employees will be willing to incur the cost of opting out or cancelling.³⁷⁸

As noted, default rules are designed to minimize opt-out costs and, in the context of nudges, can be sticky. A recent study has identified an additional cost that should factor into the optimal design of default rules: information costs.³⁷⁹ When the knowledge consumers need to make informed choices is high, consumers might opt-out, even from a sticky default, when it is not in their best interest to do so. Thus, when thinking about default rules, regulators should consider both the cost of mechanical opt outs (i.e., the cost of clicking, or itemized signing, or reading long forms) as well as the costs of becoming informed. As we explain below, firms that have a self-regarding interest in having consumers opt out of consumer-friendly defaults can manipulate consumers into waiving such protections through various types of behavioral nudges.

While default rules have been a preferred approach, reliance on mandatory rules rather than default rules may be appropriate when significant externalities are present, or when there is reason to believe that the entities formulating the mandatory rules can determine what collective choices are socially optimal.

In addition to the aforementioned reasons cautioning against a wholesale use of default rules, recent events have suggested that while relying on default rules could accommodate the heterogeneity of consumer preferences with respect to what should be done with their personal information and what tradeoffs they are comfortable with, constructing a choice architecture that facilitates well-informed and autonomous decision-making by consumers is no easy task. There is a continuum of privacy and security choices that consumers make every day, and they do so

³⁷⁸ Thaler & Benartzi (2004).

³⁷⁹ Bar-Gill & Ben-Shahar (2019).

subject to constraints on their time, energy, information processing capacity, and cognitive ability. Some decisions that consumers are charged with making are straightforward enough that boundedly rational consumers can express preferences that are meaningful and likely to reflect their own values. Under these circumstances the legal system ought to respect their autonomous choices. For example, even if consumers do not know all the consequences of turning on location services while using a navigation app like Google Maps or Waze, many of them likely understand enough of the tradeoffs to render their expressed authentic preference a meaningful and relevant data point in predicting what they want. Similarly, a consumer faced with a choice of whether to store their credit card information at a shopping site they use frequently can make a reasonably well-informed choice as to the convenience versus data-security tradeoff involved. As long as consumers are adequately informed about the relative value of the default and opt-out, empowering consumers to make these kinds of decisions, rather than enabling technocratic decisionmakers to decide for consumers, has real virtues. That being said, regulators should consider the aggregate effect of forced-choice regimes on consumers' time and cognitive bandwidth.

At the same time there are other contexts, like those involving privacy and security choices, that are very technical or complex (and thus costly for consumers to learn about). In these contexts, consumers may be especially prone to exploitation, significant collective action problems may arise, and bargaining and preference asymmetries can emerge. In other environments, the overwhelming majority of consumers may have a known, homogenous preference with respect to aspects of a transaction that are not especially salient. Relying on consumers to vindicate their interests through some combination of notice and "choice" is unrealistic, given the other demands on consumers' time and attention. These kinds of problems are sufficiently prevalent to render mandatory rules a key toolkit for privacy regulators, precisely because in those settings we can be less certain that whatever preferences individual consumers express via assent to boilerplate language actually further their interests.

Instances involving externalities and collective action problems present perhaps the clearest case for shifting from default rules to mandatory rules. An individual consumer's decision to share her DNA or location may also reveal a great deal about her relatives' genetic information or her spouse's whereabouts. Or an individual may feel compelled to reveal personal information he would otherwise prefer to keep confidential because of an unraveling dynamic, such as when failure to disclose will induce others to assume the concealer has something negative to hide.³⁸⁰ In these instances, deferring to even authentically selected private choices may not produce socially optimal outcomes, and there is a strong case to be made for employing mandatory rules, especially where society has confidence in the ability of the state to identify welfare-enhancing privacy and security choices.

³⁸⁰ Peppet (2011).

Our emphasis here is towards developing privacy default rules that reflect the tradeoffs that consumers actually want to make, and that they are capable of making, given the costs of educating themselves, evaluating different options, and making choices. Default rules should also aid in the important task of filling gaps in the face of contractual silence. It is not uncommon for firms to hide their information practices by disclosing next to nothing, or by making it difficult for regulators to police unfair and deceptive practices. In such cases, market forces sometimes (but not always and not always adequately) penalize firms for engaging in practices that do not reflect consumer welfare interests. A complete set of rules, default and mandatory, would fill that significant void.

1. Determining the Contents of Privacy Defaults

Existing scholarship has identified several basic approaches to determining the content of default rules generally. The most common approach is the “majoritarian default rule.”³⁸¹ A second approach to determining the content of default rules is alternatively referred to as an “information-forcing default rule” or “penalty default rule.”³⁸² Such an approach to setting default rules arises in settings where it is important to elicit private information from one of the parties. The information-forcing default rule imposes a choice by default that the parties are very unlikely to prefer. Because the consequences for a consumer of sticking with the default rule are likely to be quite negative, the consumer will have a strong incentive to reveal what he prefers to the company with which he is doing business, and the company can then satisfy the consumer’s preference.

In the absence of mandatory rules or other benchmarks, firms are presenting consumers with terms that maximize profits for the company rather than offering what well-informed consumers would prefer if given enough time to make a decision and understand the consequences. This strategy of selecting producer-friendly terms is quite plausibly not social welfare maximizing, given the potential for disconnect between what consumers say they want and what they appear to be receiving from digital platforms. Other times, firms might choose to say nothing about particular data practices, leaving consumers (and regulators) uninformed about what such practices are.

2. Data about Consumer Privacy Preferences and Expectations

The foregoing discussion suggests that it is worth conducting empirical research to assess whether it is possible to identify both a consensus among consumers over firms’ practices with regard to personal information and sufficient heterogeneity in consumer preferences to warrant the use of default rules rather than mandatory rules. In the domain of privacy, consumers have well-formed preferences and expectations, particularly in their interactions with familiar digital

³⁸¹ Listokin (2010) Porat & Strahilevitz (2014).

³⁸² Ayres & Gertner (1989).

platforms. Frequently those preferences and expectations diverge. Other times, consumers hold incorrect beliefs regarding firms' practices or their own attributes. To fill in what would otherwise be knowledge gaps in this white paper, Jamie Luguri and Lior Strahilevitz launched an empirical study of American consumers' preferred default rules in the consumer privacy domain. The study relied on a survey of a census-weighted representative sample of nearly 2000 American consumers recruited by the survey research firm Dynata. These survey respondents were randomly assigned to answer a series of questions about either the level of privacy they expect or the level of privacy they desire, and they were asked to confront some of the tradeoffs associated with life in the modern world, where consumers often trade personal information in exchange for a zero-cost service. The Luguri and Strahilevitz study will be published separately this year.

The study, like previous research, reveals that there was a great deal of heterogeneity in terms of how consumers expect digital platforms to collect, use, and safeguard personal information.³⁸³ For example, a supermajority of respondents understood that Amazon would store voice commands that a consumer gave to an Echo smart speaker, and among those respondents who understood such storage would occur, the majority stated that this information would be deleted either when a user actively deleted that information or when Amazon elected to do so. This understanding is consistent with Amazon's policies. Respondents were roughly equally divided on the question of whether the law permitted Amazon to pool information with Fitbit to identify customers who were likely to be training for long-distance races. A narrow majority was inclined to believe the law permits Amazon to sell information it collects via Echo devices to companies such as music streaming services. From a review of their privacy policies it appears that Amazon probably does not sell personally identifiable Echo data to third parties like music streaming services, though the policy is not clear in this respect; there do not appear to be any limits on the transfer of aggregated data, and the policies could be changed by Amazon at any time. It is not evident from Amazon's privacy policies that there are limits on the company's ability to purchase data from a third party like Fitbit, to aggregate that database with Amazon's own data, and then to identify particular kinds of consumers (e.g., long-distance runners) on that basis.

Respondents who were asked about Amazon Echo smart speakers were divided on the question of whether they would be willing to pay extra for a version of the Echo that did not share users' personal information with other companies. About 40% of respondents said they were willing to do so, and on average they stated that they would be willing to pay an extra \$50 to \$60 for such a privacy-protective option after being told that the base price for an Echo was about \$150. This data suggests a substantial minority of consumers might be willing to pay a third more for a smart speaker with these features, though it is possible that Amazon generates more than this amount of revenue from the information supplied by each household with an

³⁸³ McDonald & Cranor (2010) Zheng et al. (2018)

Echo. This kind of information could help inform regulatory decision-making about consumer preferences and tradeoffs. That said, while this kind of survey data is illuminating, a cautionary note is appropriate because willingness to pay expressed in a survey might not materialize when consumers make real-world purchasing decisions. Observational studies therefore would be very valuable.

When researchers examine respondents' normative views it becomes evident that respondents view the storage and transfer of data from Amazon Echo devices as distressing. A clear majority believe Amazon should not store such information, and there is a strong consensus among respondents that it would be undesirable for Amazon to share information it collects with a music streaming service or to merge Echo data with Fitbit data so as to identify long-distance runners.

There were similar dynamics at play when respondents were asked about the use of facial recognition data by Facebook and other social networking platforms. Most respondents understood that Facebook does use facial recognition technology to help it create templates to recognize users in uploaded photos, and that it retains this data until a user or Facebook deletes it. Users were divided over the question of whether Facebook is allowed to share facial recognition information with a third party such as a maker of police body cameras, though a narrow majority of the representative sample said the practice was not permitted (mean = 3.66 on a 7-point Likert scale, with 1 indicating definitely impermissible). When asked for a normative judgment about such information sharing with a body camera firm, respondents were much more hostile (mean = 2.40). It appears from Facebook's full data use policy that Facebook does not engage in such third-party transfers of user information, though the prohibition applies to sales and could be altered by Facebook, at least outside of Illinois. Respondents were also close to evenly divided on the question of whether Facebook uses information from its facial recognition algorithm to tag users in photos uploaded to the site by other users (mean = 4.40)—Facebook does engage in this practice unless a user objects. Again, posing the question to consumers as a normative one about what the law should permit lowered these values by a little over one point (mean = 3.30). The divide between consumers' expectations and their preferences is particularly stark when consumers are presented with binary choices. Fully 67% of respondents said that Facebook is allowed to store users' facial information that it connects with its facial recognition technology, but only 36% of respondents said that Facebook should be able to store such information.

The same basic pattern played out with respect to Google and its collection and storage of information from Google Maps. A super-majority of respondents believe that Google retains data about a Google Maps user's geolocation after the completion of the trip, and that this information is retained until either Google or the user elects to delete it. A clear majority of respondents believe that Google Maps is permitted to track a user's location whenever the phone is turned on, even when the app is not in use, as long as the user consents to this when first using the app.

Respondents were pretty evenly divided over questions involving Google's sharing of geolocation information collected through Google Maps. A little more than half of respondents thought Google was legally permitted to share data about individual users' whereabouts with stores and restaurants that wanted to deliver advertisements to customers who were nearby. And a little less than half of respondents thought Google was legally permitted to sell data it collected about individual users "to provide other services to consumers free of charge."

Again, respondents' expectations differed from their normative preferences. Clear majorities of respondents objected to the sharing of Google Maps information with nearby restaurants and stores. Noting that Google would use revenue earned from selling that user information to provide other services (like Google Maps or Gmail) to consumers at a zero price did not cause that assessment to change. Similarly, respondents did not believe that Google should be permitted to retain information about a trip that was taken with the assistance of Google Maps once the trip had ended. Most respondents thought it was wrong for Google to collect geolocation information from phones that were not using the Google Maps app notwithstanding users' previous authorization of such data collection.

This disconnect between users' expectations and their preferences plays out similarly with respect to other kinds of privacy scenarios presented to consumers in the same study. The results from studies involving the use of genetic information supplied to 23andMe or cell tower geolocation information supplied to Verizon Wireless or Gmail storage of user emails were broadly similar to the results described above, except insofar as consumers were particularly hostile to the sharing of genetic information by a genetic testing firm (e.g., sharing information with pharmaceutical companies to help them research new drugs and treatments or selling genetic information in order to lower the prices it charged consumers for genetic testing). Moreover, the study also measured the effects of posing questions to consumers involving established companies like Amazon, Facebook, and Google versus new start-ups entering the same industry in order to compete with those highly successful firms. By and large, consumers' preferences and expectations were similar regardless of whether a dominant firm or a start-up was involved. The consistent similarity of those responses helps ameliorate what would otherwise be methodological concerns about how to elicit an accurate reflection of consumer preferences using survey instruments.

3. Toward "Consumertarian Default Rules"

We propose that the law select the contents of such default rules based on the results of well-designed, scientifically rigorous studies that elicit consumer preferences, opt-out costs, and knowledge of the rules and alternatives, as well as ignorance and biases of such rules' potential costs and benefits. Where consumer preferences and expectations match firm preferences and practices, those choices should be adopted as majoritarian default rules. Where consumer preferences and expectations diverge from firms' preferences, it often will be appropriate to use the consumers' preferred starting point as a default rule, establishing rights that can be waived if

firms are able to convince consumers that waiving those rights is worthwhile. Though our proposal mirrors neither the substance nor the structure of the GDPR, the expectations-based approach we outline here is compatible with it. Under Recital 47 of the GDPR, “the reasonable expectations of data subjects based on their relationship with the controller” may create a legal basis for the processing of personal information. The GDPR does not develop a clear methodology for determining how data subjects’ reasonable expectations are to be determined, and in that sense the data-driven approach we outline offers the comparative virtue of greater clarity and predictability.

To be sure, such an approach is not devoid of implementation challenges. What should happen in instances like those described above, where there is a divergence between what consumers say they want and what they expect? There are cases to be made for using either consumer preferences or consumer expectations where the two diverge. The legal system might determine the default by averaging the two, or by adopting either the preference default or expectation default depending on the context of the substantive rule at issue. The law might use preferences to set the content of the default rule when consumer expectations do not match actual company practices or when a new technology uses personal data in novel ways. In some other instances, consumers’ preferences may be unrealistic (because imposing their preferred limitations on the use of data would make it impossible for a company offering a useful product to be profitable), too abstract (because they are not tied to price versus privacy tradeoffs), or unreliable (because they are based on mistaken consumer assumptions, such as when consumers misunderstand the fundamentals of the technology they are using). In these varied contexts, consumer preferences and expectations could be relevant rather than decisive in determining the content of the appropriate default rule.

Consumers should be able to waive these default protections in instances where their counterparty convinces them to do so, provided that the choice architecture employed by the firm is non-manipulative and enables consumers to make well-informed decisions about tradeoffs. We propose a framework below, in our discussion of dark patterns, for determining whether a particular choice architecture complies with this standard.

There is a sense in which this approach is consistent with the lessons of law and economics, in which majoritarian default rules should be employed to minimize transaction costs. For example, if both a digital platform and its customers are engaged in a transaction where collecting personal information is necessary to facilitate the underlying service that customers want and expect from the platform, then it is unproblematic to impose such collection as a term of the contract. Doing so economizes on the costs of articulating that shared preference in a written contract or privacy policy.

Majoritarian default rules maximize contracting parties’ joint welfare by selecting rules that both parties want. Selecting consumer preferences or expectations as a starting point will often be preferred only by parties on one side of the transaction. A more apt moniker for this

approach to setting default rules is “consumertarian default rules”—selecting a default rule that is preferred by the majority of unsophisticated parties in a transaction, though perhaps not by the majority of sophisticated counterparties.

Conceptually, consumertarian default rules have more in common with information-forcing default rules than they do with majoritarian default rules. Where a firm does not view the choices preferred by a majority of consumers as a desirable contractual arrangement, it has two options. Either it can stop offering the good or service in question to consumers, or it can devise ways to convince consumers to waive the protections afforded them by default. The available evidence suggests that consumers can be convinced to waive legal rights that the legal system assigns them by default. For example, research by Lauren Willis shows that banks have been quite effective at convincing consumers to waive default protections conferred on them by federal law concerning bank overdraft fees. She finds that firms can be successful at overcoming “sticky” defaults when (a) they have a strong motivation to do so, (b) they have opportunities to ask consumers to waive their rights, (c) consumers find the decision-making environment confusing, and (d) consumer preferences are well-defined.³⁸⁴

To the extent that firms are convincing consumers to waive their default protections by confusing them, pestering them, or misleading them, there is no good justification for honoring such waivers. But sometimes a firm can and does convince consumers to waive rights that they have by default by providing consumers with goods or services that they value more than they value those default protections. Indeed, essentially every successful digital platform that employs zero-dollar-pricing has done exactly that by convincing consumers that it is in their interest to trust a third party with sensitive personal information to which the third party would otherwise lack an entitlement. In light of these dynamics, we can expect that the initial assignment of privacy-friendly default protections to consumers would not necessarily result in most consumers keeping those entitlements. Large numbers of consumers could elect to waive those protections in exchange for services they value more from digital platforms. For this reason, regulators should pay attention to the cost of opt-out as well as the costs of consumers becoming informed about the value of the opt-out relative to the default. Ignorant opt-out can result in welfare losses.³⁸⁵

That said, assigning these protections to consumers would function as a constraint on firms that seek personal information from their customers. So long as some well-identified, welfare-enhancing practices become mandatory, asking consumers to waive particular consumertarian default rights means imposing on their customers’ time. We propose below stringent constraints on what constitutes a waiver of a right protected by a consumertarian default rule. Waivers would have to meet the standard that prevails in American courts where key rights are at stake—there must be a knowing and voluntary waiver of a right. Meaningless rituals that

³⁸⁴ Willis (2013).

³⁸⁵ Bar-Gill & Ben-Shahar (2019).

some courts have deemed sufficient to create a contract (e.g., presenting consumers with an avalanche of complicated text that the firm knows nearly no consumers will read) would not satisfy this heightened standard. Securing these kinds of waivers operates as a time tax on both firms and consumers. Firms might not want to initiate those kinds of conversations very often, and they will not want to do that with respect to low-stakes issues. Rather, they will of necessity pick their battles. For these reasons, in proposing that consumertarian default rules become the default for consumer privacy and security, we regard the sticky nature of default rules as a feature rather than a bug. Asking consumers to waive rights would squander some good will that firms have accumulated with consumers and would risk losing customers by adding friction to the user experience. Given that such waivers would need to be narrow in their scope, firms would, at a minimum, pick their battles, turning those default rules that firms could live with into de facto mandatory rules.

Moreover, contractual silence will no longer be a benefit for firms where personal data is concerned if a company is using personal data in ways that counter consumer expectations and preferences. Rather, such silence will become an obstacle that firms need to overcome if the benefits of obtaining meaningful customer consent are high enough. Some firms will decide that the process of informing their customers of what they want to do and why they want to do it will spook enough of their customers away from using their good or service to render it unwise to seek permission to waive a default. Requests that would reveal unsavory or controversial data practices may bring unwelcome regulatory scrutiny as well.

In a sense, then, implementing consumertarian default rules will provide some of the benefits of information-forcing default rules. By selecting default rules that firms often will not prefer, these rules will prompt firms to provide information to consumers and regulators about why retaining their entitlements may not be worthwhile. The end results may be conversations in which consumers wind up with more information about companies' use of their personal information and the associated tradeoffs. As a result, even instances in which consumers elect not to waive their rights could result in learning opportunities for consumers as well as vehicles for reflection. All of this raises hard questions about whether these conversations can ever be structured in a way that is fair, informative, and satisfying to consumers. We now turn to those issues.

B. Dark Patterns and Manipulation

Dark patterns are user interfaces that can confuse users, make it difficult for users to express their actual preferences, or manipulate users into taking certain actions. The term “dark pattern” was coined by user interface designer Harry Brignull in 2010 and is widely used among computer scientists. Behavioral economists have tended to use the term “sludge” (i.e., an evil nudge) to describe the same phenomena, though the category of sludge is not limited to online

interactions.³⁸⁶ Legal scholars have also analyzed similar phenomena, often using the term “market manipulation.”³⁸⁷

Every reader of this report will have encountered numerous dark patterns in their online activities. For instance, when a website offers users a chance to sign up for a recurring newsletter, the “accept” button might be much easier to find than the “decline” button. Worse still, the design of the webpage may appear to require a user to sign up for the newsletter in order to continue browsing the site. A firm might employ ambiguous language that confuses consumers into sharing more personal information than they intended, or it might require consumers who want to select popular settings that protect their privacy but decrease firm profitability to jump through a large number of hoops in order to do so. Or a shopping site might sneak extra goods a consumer did not select into a virtual shopping cart by default, forcing the consumer to delete the items to avoid purchasing them. In these cases, the design interface confuses users about their possible choices or makes unrealistic assumptions about what consumers are likely to prefer. Dark patterns appear to be proliferating, both in terms of their prevalence in e-commerce and the variety of different techniques employed. A recent semi-automated analysis of popular shopping web sites found that more than 11% of sites employed at least some dark pattern strategies.³⁸⁸

The effectiveness and proliferation of dark patterns is partly a product of technology. To start, many interactions are mediated by digital interfaces. People use smartphones and computers to complete many tasks and transactions that were previously done in person. The rise of digital mediation provides many opportunities for interface designers to rely on dark patterns to influence user behavior.

Many of the inherent problems with dark patterns have implications for information privacy. Dark patterns are often used to direct users toward outcomes that involve greater data collection and processing. Additionally, the proliferation of data-driven computational methods allows firms to identify vulnerabilities of users and to target specific users with these vulnerabilities.

While dark patterns come in a variety of different forms, their central unifying feature is that they are manipulative, rather than persuasive. More specifically, the design choices inherent in dark patterns push users towards specific actions without valid appeals to emotion or reason.

1. Defining Actionable Manipulation

The line between manipulation and persuasion is sometimes difficult to draw, even from a purely ethical perspective. Here, we propose a framework that we believe will allow legislators, regulators, and courts to define the category of manipulations warranting legal action

³⁸⁶ Sunstein (forthcoming).

³⁸⁷ Hanson & Kysar (1999); Calo (2014).

³⁸⁸ Mathur et al. (2019).

in a way that is workable and defensible on both economic and moral grounds. The dark patterns that we are concerned with can be divided into two (somewhat overlapping) camps: manipulation by transaction costs and manipulation by targeting certain sorts of vulnerability. Both create market failures by obstructing the market's responsiveness to consumer preferences. As the sparse dark patterns academic literature emphasizes, dark patterns are aimed at "System 1" thinking, which takes place in decision-making contexts where information is processed quickly and automatically, with little consideration and cognitive effort. Dark patterns will be much less effective when consumers are making "System 2" decisions that are characterized by slow and deliberative weighing of pros and cons.³⁸⁹ Dark patterns that target and exploit certain sorts of vulnerability raise additional moral objections. Both implicate the sorts of consumer protection issues routinely addressed by legal regulation.

a. Transaction Costs Dark Patterns

Some dark patterns—through their design features—impose transaction costs unnecessarily on users in order to get them to behave in a way that is advantageous to the pattern designer (and often contrary to a consumer's own interests or wishes). Transaction costs dark patterns are quite common and, for the most part, easily recognizable. Examples include requiring a user to uncheck multiple boxes to unsubscribe from some website feature, or designing a frictionless process for opting in to some web service while making it difficult and time consuming for users to opt out. An app might repeatedly prompt users to opt in to sharing their locations, but then never ask users who eventually accede to the requests to stop sharing their locations. While transaction costs are endemic to market transactions, they are always wasteful. Transaction costs dark patterns of this sort go further, however, by imposing unnecessary transaction costs in an attempt to manipulate consumer choices.

While users across the board will suffer the market waste and distortion caused by transaction costs dark patterns, some users will experience the increased transaction costs more acutely. Specifically, users who are less tech savvy or do not have the extra time to devote to navigating byzantine opt out procedures will be less likely to persist so that they can express their authentic preferences in the transaction. Further, these groups may preferentially include those who are already at some social disadvantages, such as elderly people with less developed technology skills or less educated people.

Identifying transaction costs dark patterns should be relatively straightforward in most cases. The imposition of unnecessary transaction costs is often apparent on the face of the transaction, at least as a *prima facie* matter. For example, if the opt-in design is seamless with limited transaction costs, while the opt-out process is burdensome, *and* if most users would prefer to opt out when given an easy option to do so, then there is a *prima facie* case that the website operator is merely adding costs to encourage users to behave according to the website

³⁸⁹ Bösch et al. (2016).

operator's wishes. Similarly, if there is an obvious design that could reduce transaction costs for users but, instead, a website chooses a design that requires additional transaction costs for users, then the website is likely employing a transaction cost dark pattern to influence user behavior. Discovery or investigation may sometimes even reveal direct evidence of this sort of manipulative deployment of transaction costs. Moreover, while there will always be borderline cases, potential liability for transaction costs dark patterns should incentivize designers to internalize consumer transaction costs to some extent. Since they are the only parties who can mitigate this social waste, this internalization of costs is a salutary effect.

b. Manipulation by Targeting Vulnerabilities

Dark patterns that manipulate individuals by targeting certain sorts of vulnerabilities are particularly troubling. Examples of the sort of behavior include the targeting of advertisements based on acute emotional vulnerabilities, such as the recent death of a child, or on specific health-related vulnerabilities, such as those related to episodes of mania or depression, as well as techniques tailored to exploit vulnerabilities of particular social groups. A person who struggles with bipolar disorder may be susceptible to impulse purchase during a manic episode. These dark patterns are admittedly more difficult to identify and assess than transaction cost dark patterns because differences between individual consumers make it more difficult to distinguish unacceptable targeting of vulnerabilities from more morally benign, or at least long accepted, persuasion tactics of traditional advertising.

To be sure, dark patterns need not be personalized. Some dark patterns are troubling because they target vulnerabilities that are pervasive, rather than particular to subsets of users. Insights from behavioral economics have demonstrated that people also have decision-making vulnerabilities—or cognitive biases. One such bias, called loss aversion, may make consumers feel the harms associated with losses more intensely than they feel the benefits associated with gains. Firms may exploit this bias by falsely indicating that only a few units of the item a consumer is examining remain in stock, so the consumer had better hurry up and make a purchase before they disappear.³⁹⁰ While recognizing that there is a disputed line between consumer protection and paternalism, we put forward a framework for defining legally actionable dark patterns of this sort that we argue will permit courts and regulators to make distinctions that are similarly clear and predictable to those made in other consumer protection law. Each of these factors will be relevant in particular contexts; rarely would one factor be decisive.

Importantly, exploiting cognitive or emotional vulnerabilities does not have a uniform moral or societal salience. Put differently, the moral wrongness and social unacceptability of exploiting vulnerabilities differs based on the type and degree of vulnerability. Some situations or characteristics simply make individuals open to persuasion in morally benign ways. For

³⁹⁰ Mathur et al. (2019).

instance, behavioral data may uncover that a person is usually hungry after getting off work at 11pm. As a result, this person is served ads for restaurants with takeout that are within a 3-mile radius. Similarly, we accept that marketers will exploit some widely held cognitive biases. For example, advertisements often employ the bandwagon effect in order to sell products. If advertisers' representations about a product's popularity are truthful and non-misleading, the advertisers' tactics are not appropriately understood as dark patterns, even though consumers might assign undue weight to the product's popularity and insufficient weight to expert assessments of its quality. For similar reasons, tactics exploiting many relatively benign and widespread sorts of weaknesses, such as the temptation to eat junk food late at night, have generally been considered morally acceptable.

Exploiting other sorts of vulnerabilities, however, raises graver moral and societal concerns. For instance, exploiting the emotional weakness that comes from losing a child has significantly deeper moral resonance than exploiting late night hunger. Similarly, exploiting mental health issues or the particular technological unfamiliarity of many elderly individuals is less acceptable than exploiting the bandwagon effect.

One approach for identifying the kinds of vulnerabilities that might rise to the level of liability when manipulated is to identify vulnerabilities revealed through categories of information protected as "sensitive" in positive law. For example, laws treat information about sexual orientation, information about disease and other health matters, and information about children as worthy of heightened protection. Manipulating someone based on information that falls into one of these sensitive classes might be likelier to be deemed worthy of legal action.

c. Hiddenness

Manipulative dark patterns are often designed to undermine a user's deliberation process covertly, so that users are unaware that their actions are being manipulated.³⁹¹ Often, these sorts of dark patterns involve outright deception, selective disclosure, or misdirection. For example, a number of online shopping sites inform users that "Jane in Anchorage" just purchased the item a consumer is scrutinizing online, but in actuality these updates are bogus—the software randomly generates a list of names and locations to make it appear that other users are buying the items in question. Similarly, shopping sites falsely indicate to consumers that particular items are "in high demand" or "in very low stock" to prompt consumers to purchase them immediately, without further research or deliberation, and it turns out that some web sites describe the vast majority of their inventory in that manner.³⁹² By contrast, persuasion is generally not hidden. Instead, persuasion's appeals to emotion or reason are overt, even if subtle. The idea that consumers should be generally aware when others attempt to persuade them for commercial purposes is reflected in regulations requiring search engines and media to distinguish between paid advertising and other content. At bottom, hiddenness—or whether a person is aware of the

³⁹¹ Susser et al. (2018).

³⁹² Mathur et al. (2019).

influence—is a factor that weighs in favor of identifying a design feature as a manipulative dark pattern.

d. Vulnerability

Design features aimed at manipulating user choices can now often be targeted based on personal information that goes well beyond the sort of general demographic information previously available to market participants. Targeting heightens concerns about manipulative dark patterns in several ways: It can enhance manipulative power by focusing on the specific vulnerabilities of particular individuals, identify areas of vulnerability that were not apparent to traditional advertisers and, by isolating individuals' experiences, make it more difficult for them to detect manipulative tactics and account for them in decision-making. Of course, targeting exists on a spectrum of granularity that affects the extent of its impact. At the most specific, dark patterns can be targeted to distinct individuals. Other dark patterns can be targeted to classes of people, such as people living in a certain area.

Targeting can exacerbate the potential for morally and socially unacceptable exploitation of vulnerabilities. Targeting can allow more granular and specific identification of individual vulnerabilities, for example by drilling down to a sub-group of elderly individuals who have been susceptible to manipulative tactics in the past. Targeting can also increase the power of manipulative tactics for particular individuals. For example, if targeting the bereaved based on gathering information from obituaries (or listening to police radio) is distasteful or unacceptable, targeting based on the immediacy and level of personal detail carried by social media is surely more so. Similarly, targeting may allow not only the exploitation of the fact that an individual suffers from bipolar disorder, but also the detection of when that individual is going through a manic or depressive phase.

2. A Data-Driven Approach to Spotting Problematic Dark Patterns

So far, our discussion of dark patterns has been theoretical and anecdotal. Indeed, a significant problem with writing about the topic of dark patterns is the dearth of literature on the subject beyond some very helpful papers that develop typologies or provide examples of dark patterns.³⁹³ A helpful taxonomy, developed by Christoph Bösch and co-authors, identifies classic types of privacy dark patterns including bad defaults (which we propose a framework for identifying above, and one example of which is a choice between “Yes” and “Not Now” rather than “Yes” and “No”), privacy zuckering (i.e., providing users options to adjust their privacy settings that are needlessly complex, granular, or confusing), forced account registration (seeming to require registration to use a service), hidden fees or terms added at the end of a long transaction (how did that wind up in my online shopping cart?), forced account preservation (making it impossible to delete accounts once created), and address book leeching (requesting

³⁹³ (Acquisti et al. (2017); Bösch et al. (2016); Gray et al. (2018); Mathur et al. (2019); Zagal et al. (2013).

users' contacts at the time of activation and then spamming users' contacts with email invitations).

As useful as this research developing taxonomies and identifying examples of problematic dark patterns has been, independent researchers are hamstrung by a frustrating dynamic. Firms that employ dark patterns know how effective they are. Academics and policy-makers, by contrast, have only limited data on that front. To correct that unfortunate asymmetry, we initiated a new experimental academic study to determine exactly how successful dark patterns are at bending consumer will. This study, *Shining a Light on Dark Patterns*, to be published separately by Jamie Luguri and Lior Strahilevitz later in 2019, provides eye-opening evidence about the effectiveness and consequences of dark patterns, be they of the subtle or blatant variety.

a. Commissioning a New Experiment on Dark Patterns

Luguri and Strahilevitz created a novel survey instrument that employs various dosages of dark pattern techniques to a large census-weighted sample of Americans ($n = 1762$). The first part of the experiment was a survey instrument in which respondents were asked for their preferences or expectations across a host of issues involving privacy and data security. This data was used to inform our analysis of consumertarian privacy defaults in the above discussion. After respondents completed answering a battery of questions soliciting their views about privacy, every participant in the sample was shown a screen indicating that our software was “Calculating your privacy propensity score. Please wait.” (This was a ruse.) After a short delay, all respondents were informed that our algorithm had identified them as consumers who had a “heightened concern about their privacy.” Respondents were also told that using their IP addresses and other information they had already provided as part of the survey (their phone numbers, etc.) we were able to uniquely identify their mailing address. Experimental subjects were further informed that the researchers had partnered with the nation's largest data security and credit monitoring firm, which had automatically signed them up for a data protection and credit history monitoring plan. This service would be offered to consumers for free for six months, but after the six-month trial period, consumers would be charged a monthly fee (randomly varied to show either \$2.99 per month or \$8.99 per month), though they could cancel the service at any time. In other words, we led our experimental subjects to believe that they were about to be signed up for a service that most of them did not seek out. The experiment was carried out after receiving approval from the University of Chicago I.R.B. to engage in this deceptive conduct.

As part of the experimental manipulation, experiment participants were randomly assigned to one of three conditions. The first group (which we will call “easy”) were shown a simple yes-no screen and asked whether they wanted to accept or decline this data protection service from the researchers' corporate partners. Neither yes nor no was pre-selected. Respondents who selected “Accept” were deemed to have accepted the service.

The second group (which we will call the “mild” dark pattern) were shown somewhat different options. Instead of boxes with the options labeled “Accept” and “Decline” these consumers were shown boxes labeled “Accept and continue (recommended)” and “Other options.” The “Accept and continue (recommended)” box was pre-selected, so a user who wanted to select that option would not need to toggle the cursor position. If respondents selected “Other options” they were shown a screen with two options. Option one read “I do not want to protect my data or credit history.” Option two read, “After reviewing my options I would like to protect my privacy and receive data protection and credit history monitoring.”

Members of this second group who selected “I do not want to protect my data or credit history” were shown one last screen as part of the experiment, in which they were given the prompt “Please tell us why you decided to decline this valuable protection.” They were shown the following options and asked to select one:

- My credit rating is already bad (1)
- Even though 16.7 million Americans were victimized by identity theft last year, I do not believe it could happen to me or my family (2)
- I’m already paying for identity theft and credit monitoring services (3)
- I’ve got nothing to hide so if hackers gain access to my data I won’t be harmed (4)
- Other (minimum 40 characters) (5): _____
- On second thought, please sign me up for 6 months of free credit history monitoring and data protection services (6)

Luguri and Strahilevitz framed these choices to be manipulative in an effort to sway some people to opt for option 6. That is, they tried to identify weak reasons for rejecting our service rather than compelling ones, and some of the options were framed in a way that might make respondents have second thoughts about declining the service.

Respondents who selected “Accept and continue ...” on screen one or “After reviewing my options ...” on screen two or “On second thought ...” on screen three were deemed to have accepted the service. All other respondents were deemed to have declined it.

The third group (which the authors call the “hard” dark pattern condition) were exposed to a litany of interventions in sequence. The first two screens respondents in the “hard” condition saw were identical to those that the “mild” group saw. After those two screens, respondents who were attempting to decline the service were shown a screen with this text:

You indicated that you do not want to protect your data or credit history. We would like to give you a little information so that you can make an informed decision. **What is identity theft?** Identity theft happens when someone steals

your personal information to commit fraud. The identity thief may use your information to fraudulently apply for credit, file taxes, or get medical services. These acts can damage your credit status, and cost you time and money to restore your good name. You may not know that you're the victim of ID theft immediately.

- Accept data protection plan and continue (1)
- I would like to read more information (2)

There was a countdown timer at the bottom of the screen that prevented respondents from selecting option 2 until ten seconds had passed. Respondents who selected option 2 were shown two more similar screens with text describing the consequences of identity theft and the prevalence of identity theft in the United States, with the same countdown timer and the opportunity to accept the data protection plan on every screen.

Respondents who kept refusing to accept the data protection plan were then shown a confusing prompt, which read as follows:

If you decline this free service, our corporate partner won't be able to help you protect your data. You will not receive identity theft protection, and you could become one of the millions of Americans who were victimized by identity theft last year.

Are you sure you want to decline this free identity theft protection?

- No, cancel
- Yes

The ambiguity problems with this prompt are significant and intentional. Consumers had to select "Yes" to decline the data protection plan. Selecting "No, cancel" would cause them to be signed up for the service. The reference to the concept of cancellation was designed to trick people who thought they were selecting that box on the mistaken assumption that doing so would decline the data protection service.

Finally, respondents who indicated "Yes" above were shown the same final screen that participants in the "mild" condition saw, asking them why they had declined the service and providing them with one last chance to reverse course.

The dark pattern experiment ended whenever respondents provided either a definitive "Accept" or "Decline" answer, though for obvious reasons declining was more cumbersome than accepting for two-thirds of the sample. At this point, all respondents were asked a few more questions and debriefed on the exercise. First, they were asked to describe their current mood, using a 7-point Likert scale ranging from 1 ("Happy and relaxed") to 7 ("Aggravated and annoyed"). Second, all respondents were told: "Some survey participants may be contacted to do a follow up survey by the same researchers. Are you interested in potentially participating?"

Again, respondents could check a box corresponding to a 7-point scale ranging from 1 (“Not at all”) to 7 (“Extremely interested”). Third, all respondents were asked: “How free did you feel to refuse the offered data protection and identity theft plan?” with choices ranging from 1 (“Not at all free to refuse”) to 7 (“Completely free to refuse”). Two more prompts asked respondents how seriously they took this survey and asked them to share any questions, comments or concerns about the survey. Some respondents used that final prompt to share their displeasure with what the researchers had done, though most left this prompt blank.

Only after this final prompt were experimental subjects told the truth. The experimenters did not partner with a data protection company. They had not signed participants up for a service. They did not and would not share their responses or demographic information with anybody. Their data would be stored only in de-identified form. Luguri and Strahilevitz told participants exactly why they were interested in dark patterns and required them to click “I understand” in order to complete the survey. This final page also provided the researchers’ contact information, though none of the participants contacted the researchers.

b. What Story Does the Dark Pattern Data Tell?

The results of the study are striking, demonstrating the powerful force that even mild dark patterns can exert on consumer choice. In the easy condition, approximately 11% of respondents accepted the data protection plan and 89% rejected it. This statistic probably overestimates the consumer demand for a product of this kind since respondents were told that they had already been signed up for the service (potentially triggering loss aversion) and a pricing system was employed in which customers would pay nothing for six months but might have the onus of cancelling subsequently to avoid charges (potentially triggering hyperbolic discounting). The researchers had also spent ten minutes priming respondents to think about privacy and security, though that prime likely cut in cross-cutting directions as the researchers simultaneously claimed to be invading respondents’ privacy in order to help protect them against identity theft. Based on the structure of the set-up, 11% is appropriately understood as the ceiling for the percentage of consumers who might be interested in the offered product on its own merits.

What happens when the researchers started employing dark patterns to boost acceptance? Even mild dark patterns proved very effective. In the mild condition, 26% of respondents accepted the data protection plan. This corresponds to a 228% increase in data protection plan acceptance compared to the control group with easy opt-out. Thanks to the experimental design, it is possible to identify the point at which different respondents agreed to the data protection plan. More than three-quarters of accepting respondents accepted the service on screen one (which offered a choice between “Accept and continue (recommended)” and “Other options,” with the former choice pre-selected). Another 23% agreed to accept the data protection plan on screen two. Nearly everyone who made it to screen three (the list of largely bad reasons for declining the service, with one final chance to say yes) selected one of the boxes that amounted

to a definitive rejection of the data protection plan. Just three respondents relented “on second thought.”

The effects of the hard condition dark pattern were even more pronounced. A battery of dark patterns convinced nearly 42% of respondents who completed the survey (217 out of 518) to accept the same data protection plan that just over 11% agreed to accept in the easy condition. This figure represents a 371% increase in the percentage of respondents who wound up accepting the plan compared to the easy condition control group. Once again, the first screen did by far the most work, with approximately 65% of the accepting respondents in the hard condition doing so at the outset. The second screen accounted for another 10% of those accepting, and the three screens on which consumers were required to slowly read more information about identity theft combined to peel away another 19% of those who wound up accepting. A confusing “Yes – Cancel” prompt seems to have been responsible for another 11% of acceptances. And once again, virtually nobody (just 1 respondent) who made it to the last screen in the hard condition surrendered at that point.

It is worth emphasizing that different populations face varied vulnerabilities to dark patterns. In the Luguri and Strahilevitz “Shining a Light on Dark Patterns” study, there was a significant relationship between education and acceptance. Specifically, more highly educated people were more likely to decline the data protection plan. This relationship was driven by the mild and hard dark pattern conditions. In the easy / control condition, education did not significantly predict whether participants declined the data protection plan. However, when dark patterns were employed in the mild and hard conditions, the more educated the participants were, the less likely they were to accept the data protection plan. This data is suggestive of a troubling prospect—not only do dark patterns prompt consumers to sign up for services they do not really want, but the least educated Americans are most likely to be manipulated successfully.

Before deciding to initiate this project, Luguri and Strahilevitz anticipated that the effects of dark patterns would be significant, but they substantially underestimated the magnitude of the effects. The data tells other surprising stories as well. Recall that Luguri and Strahilevitz randomly varied the cost of the data protection plan, so that half the sample would be signed up to pay \$2.99 a month after the six-month trial period ended and the other half would be signed up at \$8.99 per month. Raising the stakes made no difference, as the high-stakes and low-stakes conditions are statistically indistinguishable. This result seems to be in some tension with what the neoclassical model of economics would predict, even accounting for the signal of quality that could be sent to consumers by a higher price.

A key consideration in evaluating the wisdom of legal interventions is determining whether a market failure exists. The data that Luguri and Strahilevitz gathered on how annoyed participants felt and how willing they were to participate in future research studies by the same experimenters was designed to gather pertinent information on that front. Here the lesson seems to be that mild dark patterns produce little customer backlash, even while convincing many

consumers to sign up for a service they otherwise would not want. Aggressive dark patterns annoy customers who refuse services and whose costs of saying no are increased, but they do not seem to trouble customers who are manipulated into saying yes. Among respondents who wound up accepting the data protection plan, there were not robust significant differences across the easy, mild, and hard conditions. But among respondents who wound up declining it, there were significant differences between the hard condition, on the one hand (4.20 mean, SD 2.06), and the easy (2.88 mean, 1.52 SD) and mild (2.95 mean, 1.65 SD) conditions on the other. The same patterns emerged when researchers asked about willingness to be a repeat customer.

This quantitative data is corroborated by other data collected as part of the experiment. For example, at any stage in the survey, experimental subjects could close their browser window with the survey incomplete. If they opted to do that, they would not get paid by the survey research firm the researchers hired. In the mild dark pattern condition, just 9 people opted to do so. In the hard dark pattern condition, that number jumped to 65 people. The questions that forced respondents to read material about identity theft with a countdown timer accounted for the majority of these departing subjects in the hard condition. Moreover, the open-ended question asking for comments and questions generated positive or banal comments in the easy and mild conditions and a barrage of outrage and expletives from a minority of the sample in the hard condition.

These are largely lessons that digital platforms and other sophisticated entities must have learned by now, even if policymakers and academics are only now understanding the magnitude of the manipulation that can occur via dark patterns. From our perspective, it's the mild dark patterns—like labeling an option that is good for a company's bottom line but maybe not for consumers as “recommended” or by providing initial choices between “Yes” and “Not Now”—that are most insidious. This kind of decision architecture, combined with the burden of clicking through an additional screen, managed to more than double the percentage of respondents who agreed to accept a data protection plan of dubious value, and it did so without alienating customers in the process. As a result, consumers were manipulated into signing up for a service that they probably did not want and surely did not need. More broadly, we can say the same things of the kinds of dark patterns that are proliferating on digital platforms. These techniques are harming consumers by convincing them to surrender cash or personal data in deals that do not reflect consumers' actual preferences and may not serve their interests. The harms from dark patterns are akin to those associated with consumer fraud. The case for treating these widely-employed psychological strategies for securing consumer consent as unfair and deceptive practices in trade is quite strong. The case for deeming resigned acceptance under these terms as consent is very weak.

3. Legal Framework

The framework we propose could be implemented via specific federal or state “dark pattern” legislation, by legislatively extending provisions such as the “unfairness” prong of

section 5 of the FTC Act or, in some contexts, by interpreting and applying existing consumer protection law.

The call for legal intervention over dark patterns is not a radical break from other consumer protection efforts. Consider, for example, the FTC’s Do Not Call Registry. This registry allows people to designate their telephone numbers so as to place them off-limits to unwanted sales calls. In part, the registry is a response to excessive influence. Marketers would often call during dinner in order to interrupt people and make them more likely to agree to whatever service was being offered, just so they could quickly return to their family. In part, the registry was designed to protect people’s time—few people purchased goods and services over the phone, and continually saying no to telemarketers became a nuisance for consumers.

Historically, regulations also have protected consumers from inappropriate or excessive influence in market transactions, giving consumers “cooling off” periods to cancel transactions that were prodded by high-pressure sales tactics or involve high stakes or might prompt substantial regret on the part of consumers. Federal law generally provides consumers with the right to rescind transactions within three business days under circumstances in which high-pressure sales tactics are likely to be employed, such as door-to-door sales (16 C.F.R. § 429.1), home equity loans (15 U.S.C. § 1635), and student loans (15 U.S.C. § 1638(e)(7)). There are also state laws in various jurisdictions that provide 3-day windows to rescind contracts resulting from door-to-door sales, contracts for gym memberships, or even dance lessons. New York has a mandatory rule that gives consumers rights to rescind transactions resulting from telemarketing. See NY Pers. Prop. Law § 442.

Both the Do Not Call Registry and “cooling off” laws have the effect of insulating vulnerable people from undesirable pressure in sales transactions. The Do Not Call Registry protects people from confusing sales pitches that happen over the phone. Similarly, “cooling off” laws protect people by giving them a mechanism to cancel transactions after the fact. In particular, many vulnerable people may not feel comfortable declining a sales pitch that happens at their residence, especially if they live alone. At bottom, both consumer protection efforts dilute some sales pressure in ways that are particularly beneficial for vulnerable people.

In this section, we are concerned with a subset of dark patterns—ones that warrant legal intervention, yet may be deemed non-deceptive and thus arguably fall outside of existing consumer protection regimes. Some user interfaces may be inconvenient for some users but fail to cross a threshold of manipulation that we would argue should be considered legally actionable dark patterns. For example, suppose that most users of a software program want the application to be rendered in full color, but a small subset of users who are colorblind prefer that it be rendered in a format that enhances its usability for users who have reduced sensitivity to green light. There should be no legal liability if a company renders its application in full color by default and then makes a user who has deuteranomaly (green light colorblindness) navigate through several screens in order to toggle on an option that makes the application easier to read

for that subpopulation. That is a defensible choice, given the preferences of the user base and existing technologies. If the option to change the way the application appears is made too prominent, it is likely that users who are not colorblind will opt out of the default setting unintentionally and be less satisfied with their experience using the application as a consequence. Therefore, the user interface described above would not be categorized as a dark pattern, and no legal sanctions should attach to such a design.

At the other end of the spectrum, some dark patterns fit neatly within current regulatory strictures against misrepresentations and deceptive practices. The FTC, in particular, has a legislative mandate to police unfair or deceptive trade practices and has pursued enforcement actions against deceptive dark patterns. The FTC has had the most success invoking section 5's prohibition on deceptive practices in trade when consumers were tricked into signing up for goods or services. In three federal court of appeals cases, the courts have accepted the Commission's arguments that companies appropriately characterized as employing dark patterns had behaved in an unlawfully deceptive manner.³⁹⁴ The Commission and courts did not use the dark patterns terminology, and the record lacked the kind of empirical evidence revealed by the Luguri and Strahilevitz study, so the judges relied on sensible intuitions to conclude that the trade practices being challenged were deceitful.

There is, admittedly, less case law surrounding the FTC's use of section 5 from which to construct a profile of what conduct is "unfair." In the overwhelming majority of enforcement actions, companies choose to settle with the commission, entering into binding settlement agreements, rather than challenge the commission in court or administrative proceedings.³⁹⁵ In the absence of judicial decisions, however, these widely-publicized consent decrees provide significant guidance to regulated entities. In 1980, the FTC laid out the test still currently utilized to find an act or practice "unfair" due to unjustified consumer injury under section 5. Per this test, an unfair trade practice is one that 1) causes or is likely to cause substantial injury to consumers, 2) is not reasonably avoidable by consumers themselves, and 3) is not outweighed by countervailing benefits to consumers or competition. This three-part test is now codified in section 5(n) of the FTC Act.

Generally, the "substantial injury" prong is seen as the linchpin of a section 5 unfairness analysis. Overwhelmingly, the substantial harm asserted by the FTC has been monetary; however, unwarranted health and safety risks have also been the backbone of some actions.³⁹⁶ Monetary harm can come from the coercion of consumers into purchasing unwanted goods, or from other incidental injuries that come as a result of the unfair action, such as financial harm from identity theft due to improperly secured data. Notably, a harm's substantiality can come

³⁹⁴ F.T.C. v. AMG Capital Management, 910 F.3d 417 (9th Cir. 2018); F.T.C. v. LeadClick Media, LLC, 838 F.3d 158 (2d. Cir. 2016); Fanning v. F.T.C., 821 F.3d 164 (1st Cir. 2016).

³⁹⁵ Solove & Hartzog (2014).

³⁹⁶ Hartzog & Solove (2015).

from its collective effect on consumers, as the FTC notes “an injury may be sufficiently substantial, however, if it does a small harm to a large number of people” (FTC Policy Statement on Unfairness). The harm need not be “certain” under section 5, only likely. Thus actions can be brought before the damage is fully realized, as long as it is more likely than not to occur.

The next prong of the three-part unfairness test is that the injury must not be one which the consumer could have reasonably avoided. This prong is grounded in the belief that the market will be self-correcting and that consumers should, and in general will, avoid those companies that utilize unfair practices. Those practices that “prevent consumers from effectively making their own decisions” run afoul of this prong, even if they merely hinder free market decisions and fall short of depriving a consumer of free choice. In order for consumers to reasonably avoid harm, particularly in the case of a nonobvious danger, they must also be aware of the possible risk. Thus it is imperative that a company disclose risks to consumers so that they can make informed decisions with knowledge of the relevant potential harms.

The cost-benefit analysis prong of the unfairness test ensures that companies are only punished for behaviors that produce “injurious net effects.” There are, as the Commission notes, inevitable trade-offs in business practices between costs and benefits for purchasers and consumers, and as such certain costs may be permissibly imposed on consumers, provided they are justified and balanced by legitimate benefits. Broader societal burdens are also accounted for in this equation as are the potential costs that a remedy would entail. Additionally, the Commission looks to public policy considerations during this analysis in helping to establish the existence and weight of injuries and benefits that are not easily given a concrete value.

The FTC has previously regulated the use of unfair designs and unfair default settings that are varieties of dark patterns. In *FTC v FrostWire, LLC* (FTC Matter No. 112-3041, 2011), the commission brought action against the producers of a file sharing application designed in a manner that caused a significant number of consumers to “unwittingly share files stored on those devices.” Relevantly, the FTC pointed to the obstructionist defaults of the program, which made it exceptionally burdensome for a consumer to prevent all of her files from being shared. As described in the complaint, “a consumer with 200 photos on her mobile device who installed the application with the intent of sharing only ten of those photos first had to designate all 200 ... as shared, and then affirmatively unshare each of the 190 photos that she wished to keep private.” Similarly, in *In re Sony BMG Music Entertainment* (FTC Matter No. 062-3019, 2007), content protection software was required to be installed on consumers computers in order for them to listen to purchased music CDs. This software was then almost impossible to remove for the average consumer, and those in the minority who were able to uninstall the software found their CD-ROM drive no longer operable as a result. The FTC, in a similar rationale to *FrostWire*, found the design of the program unfair.

Additionally, The FTC has a long history of regulating the targeting of vulnerable consumers, whether those vulnerabilities come from inherent cognitive biases or particularized

individual or group vulnerabilities.³⁹⁷ The FTC, in its statement on unfairness, distinguishes between legitimate salesmanship and techniques that "prevent consumers from effectively making their own decisions." The Commission properly bans practices that "undermine[...] an essential precondition to a free and informed consumer transaction," such as the "exercise [of] undue influence over highly susceptible classes of purchasers" like "promoting fraudulent 'cures' to seriously ill cancer patients."

Alternatively, to the extent that dark patterns emerge in the financial services and banking sectors, the Consumer Financial Protection Bureau (CFPB) has the authority to regulate "abusive conduct." The CFPB abusive conduct definition is arguably more expansive than the "unfair" conduct regulable by the FTC and, as such, is a powerful tool for regulation. An abusive practice, per 12 U.S.C. § 5531, is one that:

- (1) materially interferes with the ability of a consumer to understand a term or condition of a consumer financial product or service; or
- (2) takes unreasonable advantage of -
 - (A) a lack of understanding on the part of the consumer of the material risks, costs, or conditions of the product or service;
 - (B) the inability of the consumer to protect the interests of the consumer in selecting or using a consumer financial product or service; or
 - (C) the reasonable reliance by the consumer on a covered person to act in the interests of the consumer.

This provision covers the exploitation of the cognitive biases of consumers in order to manipulate or pressure the consumer into making a decision that may not be in their own best interest. Adding such a broad provision to the FTC Act would be ideal for the protection of consumers insofar as the FTC would be able to restrict the use of dark patterns in all the industries subject to the Commission's jurisdiction.

a. Dark Patterns as Unfair Practices

The manipulative dark patterns we discuss here might seem natural targets for regulation aimed at "unfair" practices. We would support the interpretation of existing unfairness provisions to include them, but we acknowledge that there may be some barriers to doing so. Recall that, as currently interpreted, section 5 of the FTC Act defines an unfair practice as one which, at a minimum, (1) "causes or is likely to cause substantial injury to consumers" that (2) "is not reasonably avoidable by consumers themselves," and is (3) "not outweighed by countervailing benefits to consumers or to competition." Arguably, the manipulative dark patterns we describe here should meet all three of these prongs, but there are possible hurdles as

³⁹⁷ Solove & Hartzog (2014).

to the first two. Previous cases based on the privacy or security of personal data have sometimes stumbled on the “substantial injury” requirement because courts have found increased risk of harm insufficient to constitute “substantial injury.” A case premised on dark patterns might fare better, however, if the alleged injury was stated in terms of consumers being manipulated into entering into transactions that they would otherwise have avoided, rather than in terms of any privacy or security risks associated with those transactions. Where the dark pattern is not technically deceptive, there might also be a question of whether the injury was “reasonably avoidable by consumers themselves,” though we think unfairness should encompass the sorts of manipulative dark patterns we describe here if it is to mean anything beyond deceptiveness. Favorable interpretations of unfairness could be pressed in litigation. Alternatively, or in addition, Congress could empower the FTC to engage in rulemaking on these issues—a change that would be desirable for other reasons as well.

In our view, a quantitative approach to identifying dark patterns could be workable and offers many of the benefits of bright-line rules in general. More precisely, where the kind of A/B testing that we discuss above reveals that a particular interface design or option set more than doubles the percentage of users who wind up “consenting” to engage in a consumer transaction, the company practice at issue could be deemed presumptively an unfair or deceptive practice in trade. In the scenarios developed by Luguri and Strahilevitz, both the mild dark patterns and the hard dark patterns made it more likely than not that consumers were electing not to decline a service on the basis of the choice architecture employed rather than on the basis of innate demand for the service at issue. The “more likely than not” standard is widely employed in civil litigation over torts and other kinds of liability, and it could work well in this context too, ideally with the FTC and academics working hand in hand to replicate high-quality research that quantifies the effects of particular manipulations. As a statistical matter, each individual research subject in the study who was signed up for the data protection plan was more likely than not to have done so because of the dark pattern rather than because of underlying demand for the service being offered.

What’s more, it may be appropriate for courts to deem instances in which the “more likely than not” test is satisfied as instances in which consumers have not actually consented to the contractual terms at issue. To hold otherwise runs the risk of treating consent as a legal fiction, rather than an indication of mutual assent. Numerous legal settings require individuals to consent to have their information processed or shared, or to incur a legal obligation to pay for a good or service. Where this consent is procured through a manipulative exchange between the consumer and the digital platform it should be treated by courts as a legal nullity.³⁹⁸

In embracing a “more likely than not” rule, we do not mean to rule out the development of multifactor standards that can complement a rule-based approach. For example, the “more likely than not” test works very well when the innate preference for a product among consumers

³⁹⁸ Hartzog (2018).

stands at 10 or 20%. But when 40 to 50% of consumers would want to sign up for a service or purchase a product, the “more likely than not” test is likely to let too much manipulative conduct survive. In our view, a situation where 40% of consumers opt to buy a service because of innate demand for it and 20% of consumers opt to buy because of a manipulative interface or choice architecture may still be legally problematic. We consider how the law should address those kinds of situations below, in the section labeled “Questions of Intent and Proof.”

In considering the interpretation of fairness requirements, it may also be useful to consider the EU’s approach to “fair processing” under the GDPR. The UK’s Information Commissioner’s Office has issued guidance to data processors stating they should be able to affirm that:

- We have considered how the processing may affect the individuals concerned and can justify any adverse impact.
- We only handle people’s data in ways they would reasonably expect, or we can explain why any unexpected processing is justified.
- We do not deceive or mislead people when we collect their personal data.

This approach would seem to capture the manipulative dark patterns we have in mind.

b. Questions of Intent and Proof

However implemented, the framework we propose here lends itself well to a two-step approach for purported dark patterns that do not satisfy the “more likely than not” bright line rule proposed above, beginning with a prima facie allegation that a manipulative dark pattern based on either transaction costs or vulnerability has been deployed. The prima facie barrier for discovery or investigation should not be too high because such patterns will often be the result of intentional design choices or, at a minimum, of willful blindness to design implications that will be readily apparent from internal documentation and behavior. We do not support a standard that requires intent or willful blindness for liability, however. Though such a standard might capture many of the most problematic dark patterns, particularly of the transaction cost strand, increasing use of artificial intelligence approaches to targeting suggests that system designers should have a duty to avoid implementing manipulative dark patterns that rise to the level that would be actionable under our framework.

To be more explicit about how enforcement might proceed, a consumer protection agency might identify illegal dark patterns on the basis of a two-step process. First, the agency would identify evidence that a particular design appears to impose unnecessary transaction costs or to exploit vulnerabilities associated with particular individual or group characteristics. Once this prima facie threshold is met, the consumer protection agency would further investigate the procedure underlying the design choices for this interface, thus potentially uncovering further evidence that the manipulation is intentional or willfully blind or, alternatively, that there is a non-manipulative justification for the design. For example, when a company is rolling out a new

product or service to its existing customers, it may need to disclose a healthy quantum of information to describe the new opportunity, and that process of persuasion will entail transaction costs for consumers. To the extent that the company is offering arguments designed to appeal to consumers' System 2 decision-making processes, it should not get into any trouble with regulators. It is only when the choice architecture employs cognitive tricks that are designed to exploit quirks of human System 1 decision-making that the line separating unlawful dark patterns from constitutionally protected persuasive commercial speech is crossed. Admittedly, there may be some close cases involving judgment calls about the precise line between permissible and impermissible, and evidence of intent will be illuminating where it is available. Regulators charged with policing dark patterns should focus on the blatant cases at the outset, while the development of "common law" principles along the way will permit the eventual differentiation of disputes closer to the margin.

Suppose that elderly consumers have been complaining that a user interface manipulated them into purchasing something that they did not want or into unintentionally giving over personal information. An investigation might uncover evidence that the designers engaged in A/B testing aimed at elderly people or were willfully blind to data showing the design's disproportionate effects on the elderly. Information about design processes may often be sufficient to confirm the identification of a manipulative dark pattern. Sometimes, however, the results may have emerged from an AI targeting algorithm and will have to be demonstrated by other sorts of evidence, perhaps including statistical testing.

C. Institutional Implementation Considerations for Default Rules and Dark Pattern Regulations

Once the content of default rules and the non-manipulative way in which they should be offered to consumers are determined, there is the regulatory question of how such rules should be implemented. There are several models for doing so, and the optimal implementation choice should consider the issues regarding the problems of fine print and consumer contracts, where it is the sellers that write and present the contracts in the first place and consumers may remain uninformed about some or all of the terms in the contracts.

Unlike the case of dark patterns, where there is already a regulatory structure capable of incorporating enforcement actions against such practices, there are currently no all-encompassing default rules in the information privacy space. One approach towards implementation of default rules is through a principles project with the American Law Institute, or, less likely in this context, a restatement project, if there is enough case law to "restate" and steer the rules in a normatively appealing direction. In practical terms, principles project ideas are generated by the director and the projects committee, who also take suggestions from members. The director then develops a project proposal and potential reporters, which is then considered by the project committee for advice and recommendation to the council. While

restatements are directed to judges, the principles projects are addressed to legislatures and administrative agencies, or also to judges when the law in that particular area is scant. The benefit of a principles project is that the consumertarian-rule creating process could result from robust exchanges among experts in industry, judges, consumer advocates, and practicing lawyers, all of whom would give thoughtful advice to the reporters, thus ensuring all interests are well represented in the final product. The drawbacks of such approach are that it might take considerable time to obtain approval by the ALI members, the final product might not be optimal but rather reflect compromises made to ensure a final vote, and legislatures might ignore it. At the end of the day the ALI process may not be nimble enough to respond to rapid changes in technology and corporate practices unless the end product is written in a way that is open-ended enough to empower those charged with enforcing the law but also still clear enough to provide fair notice to regulated entities. Given industry participants' need to comply with the laws of multiple jurisdictions, an effective default rule regime would have to be adopted wholesale by all states or at the federal level to ensure uniformity.

Other projects take the form of codes, which are drafted with an eye toward legislative enactment. An example of this is the Uniform Commercial Code, which does not have legal effect until and unless adopted by state legislatures, which can, of course, modify the proposed code as they see fit. Codes might take longer to get buy in from the relevant bodies and might suffer some of the same concerns raised with principles or restatements. Alternatively, private bodies can create model rules that can also later be adopted by legislatures. For example, The Principles of European Contract Law (PECL) is a set of model rules, created by the Commission on European Contract Law, led by Professor Ole Lando (the Lando Commission). Like the restatement, the PECL is “soft law,” and, without some legislative or judicial action, not legally enforceable.

Finally, Congress could grant the FTC or a newly established federal agency some rule-making authority and charge the agency with the task of creating default rules and limits on dark patterns, or at least involve the agency in the process of creating such rules.

One dynamic should be relatively clear from our analysis above, which is that waiting for market forces to solve the problems we have identified is unlikely to lead to satisfying results. Even if markets for online services were perfectly competitive, many of the impediments to market privacy and security solutions we identified in section I.A. would remain. Namely, when privacy breaches occur it is often hard to tie subsequent harms to any particular breach, since the same personal information may be replicated in many different proprietary databases. These dynamics give rise to substantial negative externalities.³⁹⁹ Companies that wanted to differentiate themselves on the basis of investments in privacy and security or the development of enlightened policies would face challenges in communicating their choices to unsophisticated consumers, given the complex and technical subject matter at issue and the danger that consumer-facing

³⁹⁹ See Ben-Shahar (2018).

advertising will spook consumers about entire product lines. For example, consumers presently have no understanding of various ways in which dark patterns are being used to manipulate their choices, and the subtlety of successful strategies would make it costly for a firm that eschewed the dark patterns strategy to run a campaign condemning their competitors in an effort to gain market share. In that sense, we can contrast dark patterns with a more blatant and easier to understand strategy like extra fees for checked luggage, a pricing strategy Southwest Airlines has criticized in its “Bags Fly Free” ad campaign. Digital platforms also have incentives not to publicize many of their investments in data security in detail because describing those investments to consumers (or shareholders) necessarily alerts hackers to the same information, which in some cases could create new vulnerabilities. Finally, perfect competition is unlikely to internalize the externalities associated with one consumer’s decision to expose personal information about friends and family. We do not mean to suggest that it is impossible for consumer-facing firms to differentiate themselves on the basis of privacy and security precautions. Some have done so successfully. But those efforts represent an incomplete strategy for addressing the thorny range of problems highlighted herein.

D. Mitigation for Security Threats Caused by Data Breaches

To mitigate attacks that leverage passwords stolen in a data breach to compromise accounts on other services, minimizing the role of passwords in authentication would seem an obvious solution. Unfortunately, efforts to eliminate passwords entirely are unlikely to fully succeed in the near term. One reason for widespread password reuse is that users are asked to make dozens, or even hundreds, of different accounts, falling back on password reuse as a coping mechanism. The idea of federated identity or single-sign-on systems, which entail using a single identity provider (e.g., Google) to authenticate many different services, would obviate having so many different accounts. However, adoption of such systems has been limited for a number of reasons, including users' concerns about sharing data with companies like Google and Facebook that offer to serve as identity providers, as well as concerns about creating a single point of failure. While password manager software installed on a user's device does not raise the same privacy concerns, password managers are again a single point of failure and also enable new kinds of attacks. Furthermore, current password managers do not fully prevent users from continuing poor practices like password reuse due to design decisions their creators have made.⁴⁰⁰

Techniques like multi-factor authentication often layer a password with a second line of defense (e.g., possession of a smartphone configured in a particular way). While doing so substantially lessens accounts' vulnerability following a data breach, it again does not eliminate passwords. Recent efforts by the World Wide Web Consortium to standardize the WebAuthn specification for authentication using public-key cryptography (often enabled by USB tokens from manufacturers like Yubico) have been heralded as a new way to eliminate passwords. However,

⁴⁰⁰ Ur et al. (2015).

the inconvenience of users needing to carry additional hardware for authentication (at least in the near term) prevents WebAuthn from being a silver-bullet solution. As a result, passwords are unlikely to disappear completely anytime soon. Furthermore, rearchitecting authentication systems would, at best, only solve problems related to reused passwords stolen in a data breach, not to any other types of information stolen.

Accepting the persistence of the password in the authentication ecosystem, another approach to mitigating the security threats of data breaches would be to involve the affected users more directly, notifying them in greater detail about what information has been breached and how it could be used. Currently, data-breach notifications provide limited, often vague, information to consumers. As such, it is not surprising that consumers are often not sure how to respond when either their account credentials or their personal information is reported stolen.⁴⁰¹ As a result, even after major, widely publicized breaches like those of Equifax or LinkedIn, many consumers have not taken action in response. In contrast, automated systems could potentially have taken some action on their behalf.

Fully empowering the user to act would likely require a radical redesign of the data-breach notifications companies send to consumers. These notifications would need to be more detailed and actionable. Laws might also need to be amended. Current legal requirements for data-breach notifications vary by jurisdiction, yet no jurisdictions require notifications that are as detailed as might be needed. Further complicating these efforts is the necessity to authenticate the recipients of such detailed notifications; sending detailed information about information that was stolen to an unauthenticated recipient is, in a way, its own data breach. However, this approach increases, rather than decreases, the burden on users. The redesign of data breach notifications could conceivably provide greater agency and transparency to users, yet it would be unlikely to change this desolate scenario of widespread inaction. Increased transparency would only make an imperfect system even more burdensome. As mentioned above, users could easily resort to password managers to mitigate many of the risks discussed herein. That they do not do so demonstrates the need to develop alternative solutions that shift the focus from the user to other parties.

An interesting alternative may be to encourage companies to share the data stolen in a data breach directly with other companies, including their competitors. That is, following a data breach, companies would send the precise information identified as stolen to any other companies that request that information. Given that some larger companies already buy this breached information from hackers themselves, in some sense such a system would simply remove hackers as the middlemen. Two key challenges stand in the way of a system like this becoming practical. First, it is ethically questionable for a company that has experienced a data breach to further disseminate their users' personal information and credentials to any other company that asks. On one hand, disseminating this information that had been given specifically to only one company can violate

⁴⁰¹ Ablon et al., (2016); Golla et al. (2018).

users' privacy. That a further violation of privacy may be necessary follows from the breached company's negligence in protecting the data in the first place. On the other hand, given that other companies might be purchasing this information from hackers anyway to protect their own users, it might cause further harm if this information is not disseminated for defensive purposes. Complicating this calculus, though, are questions about how a company might want to handle the dissemination of data that has been breached by hackers but has not yet been released publicly.

The idea of competing companies working together to stave off computer security threats has some conceptual similarities to controversial efforts for sharing threat intelligence as part of the Cybersecurity Information Sharing Act of 2015 (CISA). This law established processes and legal safeguards for private companies to share information on computer security threats (e.g., details of attempted attacks) with the US federal government or with other private companies. The legal safeguards provided by CISA minimize potential liabilities related to the disclosure of privileged information, and they also protect shared information from FOIA requests under certain circumstances. While CISA has been presented as facilitating voluntary sharing of cybersecurity threats, rather than mandating such sharing, some of the act's many critics have questioned whether sharing truly is voluntary.⁴⁰² Other critics of the act have pointed out that CISA potentially broadens federal surveillance under the guise of warding off hackers.⁴⁰³ In the end, only a few companies have chosen to share data with CISA thus far. These past experiences with CISA raise important questions about how to structure any future processes for companies to share information stolen in data breaches with other entities.

Given concerns users might have with their data being shared with particular other companies or with government agencies, as well as how such a sharing infrastructure could create perverse incentives in favor of data breaches, the architecture of such an inter-entity sharing agreement, and the role of the government in such an agreement, requires care and thought. One possible direction to mitigate these tensions could be a process for consumers to opt in to having their data shared with other companies when notified of a data breach. In contrast, though, one could consider establishing a consumertarian default rule. Outlining a default rule that establishes terms that most individuals in a data breach would agree to, yet could be overridden upon the request of each impacted user, would likely require empirical grounding. That is, empirical research would be needed to document what specific types of information (e.g., usernames, credentials, contact information, social contacts) users would want other unrelated companies to be told were stolen in a given data breach. Consumers would likely view such transfers as akin to data sharing or data selling, practices that they often view negatively, even though in this case the transferred data could perhaps be constrained legally to allow only for threat mitigation. Empirical research would also need to answer questions about consumers' perception of which companies or entities should receive this data. Consumers' opinions might dictate that each consumer should

⁴⁰² Stepanovich (2015).

⁴⁰³ Krebs (2015).

establish a personalized whitelist of companies to which information could be shared. The empirical work might instead suggest that a more global, non-personalized whitelist is appropriate, dictating that only a particular list of trusted companies or entities should be permitted to receive data.

Beyond a consumertarian default rule, one can envision two other possible solutions to increase coordination and data sharing amongst companies for the purposes of increased data security. This system must be designed in a way that promotes the sharing of this sensitive and critical information while at the same time protecting user privacy, providing proper incentives for companies to participate (in contrast to CISA) and complying with laws and regulations on data sharing amongst companies (both data protection and antitrust laws).

The first possibility one could consider is to further strengthen a CISA-like system by making participation in a government-controlled database mandatory, alleviating some of the free-riding problems that currently impair CISA (companies receive reports even if they do not participate). Such imposition, however, would probably face even more opposition than the initial, voluntary CISA-style system did. Indeed, as seen above, CISA was perceived as a violation of user privacy and as an expansion of what is already an intrusive federal surveillance apparatus. These are some of the reasons why the CISA legislation stalled in Congress for many years and why current voluntary participation is low. As a result, such an expansion is unlikely. It is also probably undesirable.

A second, preferred, alternative may be to use technological tools to design a “data-breach clearinghouse,” where companies may share information about breaches and passwords in a secure and privacy-preserving manner. Companies could query the clearinghouse to determine whether a given user on their service had used the same, or a similar, password on another service. Enabling companies to insert data stolen from them into the clearinghouse removes hackers as the middleman who stand to profit from selling this data to other companies, and it also facilitates coordination among competing companies for the overall benefit of the authentication ecosystem. It can also minimize the additional financial incentives for hackers. Before they make stolen data public, they attempt to sell the data to others who weaponize it to compromise other accounts. Enabling other companies to defend against these attacks could substantially lessen the value of this non-public data.

Two methods can provide incentives for companies to participate. First, data-breach laws could be amended to mandate a company’s participation in the clearinghouse shortly after it learns of a data breach. Second, data-breach laws or courts could establish that not querying their own users’ records in the clearinghouse is prima-facie evidence against data-security protocols, potentially increasing liability in the event of a data breach. Finally, by being privately run and having a sole objective of enhancing data security, the clearinghouse would not be the equivalent to CISA in enabling further surveillance by the government (assuming they are protected from government backdoors).

A couple of security and privacy design considerations are crucial for such a clearinghouse to be feasible. First, the clearinghouse must not become a new single point of failure for the authentication ecosystem in the event that it is compromised. Second, the clearinghouse must minimize the potential for brute-force querying to leak private information. However, current cryptographic techniques can enable more privacy-preserving approaches to data sharing and can partially mitigate concerns about how such data sharing impacts privacy. Broadly, the sub-field of cryptography studying secure multi-party computation focuses on techniques that can meet these assumptions. More specifically, recent work has engineered new techniques for decentralized private set-membership testing. These techniques enable one website to query other websites about whether a user's password is the same across sites.⁴⁰⁴ The key innovation is that this work uses advanced mathematics to perform this testing without negatively affecting the privacy of a user's password, subject to certain assumptions. While these techniques are not yet at a point of engineering maturity sufficient for a data-breach clearinghouse like the one we imagine, they are an important step toward such a design.

For a number of reasons, such a clearinghouse is strongly preferable to companies directly sharing plaintext information with each other, or buying that information from hackers. For example, directly sharing plaintext data would have major ramifications for user privacy. In addition, depending on how the structure is designed, the direct sharing of information among competitors may run afoul of antitrust laws—in particular if this information is essential and if the trust requirement restricts the system to only a handful of already established companies. This scenario could be a variation of the Supreme Court's *Associated Press v. United States* decision (326 U.S. 1 (1945)).

Recent pushes to bring more centralized notifications of data breaches to consumers are related to a potential data-breach clearinghouse, yet operate on slightly different assumptions and fulfill different use cases. The website *Have I Been Pwned* focuses on notifying users when data associated with a particular email address or user name is part of a data breach. However, because it does not attempt to communicate the data itself (i.e., the exact password that has been stolen), it does not solve the same problem as a potential data-breach clearinghouse. Mozilla's recently released Firefox Monitor for the Firefox browser brings the *Have I Been Pwned* functionality to a wider user base. Google's recently released Password Checkup extension has similar goals and uses some related technologies. The Password Checkup extension aims to tell a user if the credentials they are trying to create for an account match those that Google knows to be compromised from prior data breaches (e.g., those Google's security team has collected from hacker forums) without revealing the user's password to Google. It does so, in part, by sending an encrypted version of a short prefix of the username and password of interest to Google, sending back to the user all entries in the database matching that prefix. While such a system is ideal for when the credentials have already been leaked publicly, this approach does not solve the problem

⁴⁰⁴ Wang & Reiter (2018).

of permitting the data-breach clearinghouse to handle credentials that have potentially been stolen by hackers but not yet released publicly. These credentials have much higher security requirements lest the data-breach clearinghouse inadvertently increase harm by leaking information about them.

E. Privacy-Preserving Measurement and Coordination

There are several components of the Committee's different reports that call for greater measurement of technology sector business practices or coordination between firms. The Market Structure and Antitrust Subcommittee, for instance, recommends routine collection of market transaction data to understand the competition dynamics of major technology platforms. The reports on the News Media Industry and on Political systems also call for data collection and disclosure. This Subcommittee's own report envisions a system for sharing stolen user credentials in the event of a data breach.

In considering the virtues of greater information transparency, it is vital to acknowledge that market measurement and sectoral coordination can involve confidential personal data. Public-sector agencies and private-sector businesses participating in these information sharing programs should adhere to principles of data minimization, avoiding bulk data dumps in favor of more targeted disclosures. Participants should also implement appropriate security safeguards, including technical precautions such as encryption and access controls, as well as security policies such as employee training. Promoting competition, protecting consumers and even understanding the impacts of platforms on political systems should not come at the expense of significant additional privacy impositions.

We also note that there is a range of privacy-preserving computational techniques that may implement the Committee's measurement and coordination recommendations while mitigating privacy consequences. One promising direction is differential privacy, a family of techniques that enables performing computation on a dataset without revealing an individual's personal information. Unlike prior approaches to privacy-preserving dataset analysis (e.g., k-anonymity), which are appealing in theory but prone to vulnerabilities in practice, differential privacy is designed to provide strong probabilistic guarantees. The technology sector is increasingly implementing differential privacy for user analytics, and software libraries for differential privacy are quickly reducing the costs of and barriers to further adoption.

It would be (relatively) straightforward, for example, to apply differential privacy to market surveillance. A regulator could issue queries to regulated firms to understand long-term trends in technology sector competition, and meanwhile, the firms could have confidence that they were not disclosing the details of specific transactions or business partners.

Another promising direction is secure multiparty computation, which enables multiple entities to collaborate on data analysis while limiting the data that each entity discloses. Secure multiparty computation could, for example, enable regulators to conduct market surveillance using data held across multiple firms while minimizing information flows to regulators and

preventing competitors from learning each other's sensitive business details. In the context of data breach information sharing, as noted above, secure multiparty computation could enable firms to compare user credentials without ever disclosing the credentials themselves.

The Subcommittee is not recommending any specific technology for privacy-preserving computation, but we encourage bearing these opportunities in mind for the subcommittees' recommendations and future recommendations on technology competition, the media landscape, politics, and privacy.

Conclusion

The law and the marketplace, as they presently exist, simply expect too much of consumers. They expect each consumer to read terms of service and privacy policies and to opt out of problematic practices. Such an expectation might be realistic if there were one entity that handled consumers' data, but it becomes absurdly impractical when scores of such entities interact with the typical consumer. The law expects consumers to be vigilant against dark patterns, and it often treats sneaky and manipulative tactics that are designed to confuse or badger consumers into accepting privacy intrusions as tantamount to authentic consent. Consumers are also expected to keep track of dozens of different passwords and other authentication tools, and to do so in a way that will not itself expose those passwords to the possibility of breach. It is little wonder, then, that password reuse and the use of inadequate passwords have become widespread problems. In short, in many privacy and security domains, consumers who face time constraints, cognitive constraints, or situations that seem low-stakes at first glance, and who then act rationally in response, are left unprotected by legal doctrines. We can and ought to do better.

In our report, we have endeavored to identify three key areas where fundamental problems of privacy and security emerge, and where the existing legal and regulatory toolkit has failed consumers. As a result, consumers are stuck with product features that they do not understand or do not want. They are manipulated into making choices that conflict with their preferences and values. And their data is being exposed to unnecessary threats of leakage and misuse resulting from the failures of platforms to coordinate in ways that would serve the collective interests of both customers and shareholders. Our effort here has been to summarize the state-of-the-art academic research where it exists and to supplement it in important instances where we have found the literature wanting. Having consulted the research, new and old, we have proposed three ambitious ideas to enhance the privacy and security of Americans' interactions with the digital platforms that have become important, ubiquitous, and unfortunately sometimes frustrating parts of life in the modern world.

The Privacy and Data Protection Subcommittee has not raised all the important issues concerning privacy and data security that relate to digital platforms. There are other efforts underway—in legislatures, within firms and industry groups, in academia, and within bodies like the ALI—that aim to be more comprehensive than this one and that stand to accomplish much good if their ideas are implemented. We hope those efforts succeed because the privacy and security problems confronting the users of digital platforms are substantial and urgent.

References

- Ablon, Lillian, Paul Heaton, Diana Catherine Lavery, and Sasha Romanosky. 2016. *Consumer Attitudes Toward Data Breach Notifications and Loss of Personal Information*. Rand Corporation. <http://www.rand.org/t/rr1187>.
- Acquisti, Alessandro, Idris Adjerid, Rebecca Balebako, Laura Brandimarte, Lorrie Faith Cranor, Saranga Komanduri, Pedro Giovanni Leon, Norman Sadeh, Florian Schaub, Manya Sleeper, Yang Wang, and Shomir Wilson. 2017. "Nudges for Privacy and Security: Understanding and Assisting Users' Choices Online." *ACM Computing Surveys* 50(33): article 44.
- Ayres, Ian, and Robert Gertner. 1989. "Filling Gaps in Incomplete Contracts." *Yale Law Journal* (99)87.
- Bakos, Yannis, Florencia Marotta-Wurgler, and David R. Trossen. 2014. "Does Anyone Read the Fine Print? Consumer Attention to Standard Form Contracts." *Journal of Legal Studies* (43)1.
- Bar-Gill, Oren, and Omri Ben-Shahar. 2019. An Information Costs Theory of Default Rules. Working paper. <https://www-cdn.law.stanford.edu/wp-content/uploads/2018/12/Coversheet-paper-.pdf>.
- Barocas, Solon, and Karen Levy. 2018. "Privacy Dependencies." Privacy Law Scholars Conference draft.
- Ben-Shahar, Omri. 2018. "Data Pollution." Univ. of Chicago Coase-Sandor Institute for Law & Econ. Research Paper No. 854. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3191231.
- Bösch, Christoph, Benjamin Erb, Frank Kargl, Henning Kopp, and Stefan Pfattheicher. 2016. "Tales from the Dark Side: Privacy Dark Strategies and Privacy Dark Patterns." *Proceedings on Privacy Enhancing Technologies* 2016(4).
- Calo, Ryan. 2012. "Against Notice Skepticism in Privacy (and Elsewhere)." *Notre Dame Law Review* (87)3: 1027.
- Calo, Ryan. 2014. "Digital Market Manipulation." *George Washington Law Review* 82(4): 995.
- Golla, Maximilian, Miranda Wei, Juliette Hainline, Lydia Filipe, Markus Dürmuth, Elissa Redmiles, and Blase Ur. 2018. "What was that site doing with my Facebook password? Designing Password-Reuse Notifications." In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*.
- Gray, Colin M. et al., The Dark (Patterns) Side of UX Design, Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems Paper No. 534, available at <https://dl.acm.org/citation.cfm?id=3174108> (2018).

- Hanson, Jon D., and Douglas A. Kysar. 1999. "Taking Behavioralism Seriously: The Problem of Market Manipulation." *NYU Law Review* 74(3): 632.
- Hartzog, Woodrow. 2018. *Privacy's Blueprint*. Harvard University Press.
- Hartzog, Woodrow, and Daniel J. Solove. 2015. "The Scope and Potential of FTC Data Protection." *George Washington Law Review* 83(6): 223.
- Hoofnagle, Chris Jay. 2016. *Federal Trade Commission Privacy Law and Policy*. Cambridge University Press.
- Krebs, Brian. 2015. "Cybersecurity Information (Over)Sharing Act?" *Krebs on Security Blog*, Oct. 27. <https://krebsonsecurity.com/2015/10/cybersecurity-information-oversharing-act/>.
- Lemley, Mark A. 2006. "Terms of Use." *Minnesota Law Review* 91:459.
- Listokin, Yair. 2010. "The Meaning of Contractual Silence: A Field Experiment." *Journal of Legal Analysis* 21:397.
- Marotta-Wurgler, Florencia. 2016. "Self-Regulation and Competition in Privacy Policies." *Journal of Legal Studies* 45:S13.
- Mathur, Arunesh, Gunes Acar, Michael Friedman, Elena Lucherini, Jonathan Mayer, Marshini Chetty, and Arvind Narayanan. 2019. "Dark Patterns at Scale: Findings from a Crawl of 11K Shopping Websites." Working paper. <https://webtransparency.cs.princeton.edu/dark-patterns/assets/dark-patterns.pdf>.
- McDonald, Aleecia M., and Lorrie Faith Cranor. 2008. "The Cost of Reading Privacy Policies." *I/S: A Journal of Law and Policy* 4(3): 543.
- McDonald, Aleecia M., Lorrie Faith Cranor. 2010. "Beliefs and Behaviors: Internet Users' Understanding of Behavioral Advertising." Paper presented at *2010 TPRC—43rd Research Conference on Communication, Information and Internet Policy*.
- Peppet, Scott. 2011. "Unraveling Privacy: The Personal Prospectus and Threat of a Full-Disclosure Future." *Northwestern University Law Review* 105: 1153.
- Porat, Ariel, and Lior Jacob Strahilevitz. 2014. "Personalizing Default Rules and Disclosure with Big Data." *Michigan Law Review* 112: 1417.
- Solove, Daniel J., and Woodrow Hartzog. 2014. "The FTC and the New Common Law of Privacy." *Columbia Law Review* 114(3): 583.
- Stepanovich, Amie. 2015. "Busting the Biggest Myth of CISA—That the Program is Voluntary." *Wired*, Aug 19. <https://www.wired.com/2015/08/access-cisa-myth-of-voluntary-info-sharing/>.
- Strahilevitz, Lior Jacob. 2010. "Collective Privacy." In *The Offensive Internet: Privacy, Speech, and Reputation*, ed. Saul Levmore and Martha Nussbaum. Harvard University Press.

- Sunstein, Cass R., Sludges and Ordeals, 69 Duke L.J. (forthcoming).
- Susser, Daniel et al., Online Manipulation: Hidden Influences in a Digital World, working paper, available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3306006 (2018).
- Thaler, Richard & Benartzi, Shlomo, Save More Tomorrow: Using Behavioral Economics to Increase Employee Saving, 112 J. Pol. Econ. S164 (2004).
- Ur, Blase, Fumiko Noma, Jonathan Bees, Sean M. Segreti, Richard Shay, Lujo Bauer, Nicolas Christin, and Lorrie Faith Cranor. 2015. "I Added '!' at the End to Make It Secure: Observing Password Creation in the Lab." In *Proceedings of the 11th SOUPS*, Ottawa, Canada. <https://www.usenix.org/system/files/conference/soups2015/soups15-paper-ur.pdf>.
- Wang, Ke Coby, and Michael K. Reiter. 2018. "How to End Password Reuse on the Web." In *Proceedings of the NDSS*. <https://arxiv.org/abs/1805.00566>.
- Willis, Lauren. 2013. "When Nudges Fail: Slippery Defaults." *University of Chicago Law Review* 80: 1155.
- Zagal, José P. 2013. "Dark Patterns in the Design of Games." Working paper. <http://www.diva-portal.org/smash/get/diva2:1043332/FULLTEXT01.pdf>.
- Zheng, Serena, Noah Apthorpe, Marshini Chetty, and Nick Feamster. 2018. "User Perceptions of Smart Home IoT Privacy." *Proceedings of the ACM on Human-Computer Interactions* 2(200):1. <https://arxiv.org/abs/1802.08182>.

George J. Stigler Center for the Study of the Economy and the State
The University of Chicago Booth School of Business

Stigler Committee on Digital Platforms

Politics Subcommittee

Report

July 2019

The Politics Committee

Andy Guess, Assistant Professor of Politics and Public Affairs, Princeton University

Rana Foroohar, Associate Editor and Global Business Columnist, Financial Times

David Lazer, Professor of Political Science and the College of Computer and Information Science, Northeastern University

Nolan McCarty (chair), Susan Dod Brown Professor of Politics and Public Affairs, Princeton University

Alexandra Siegel, Post-doctoral fellow, Immigration Policy Lab, Stanford University

Nicholas Stephanopoulos, Professor of Law, Herbert and Marjorie Fried Research Scholar, University of Chicago Law School

Joshua A. Tucker, Professor of Politics, Affiliated Professor of Data Science and Russian and Slavic Studies, co-Director of the Social Media and Political Participation Lab, Director of the Jordan Center for the Advanced Study of Russia, New York University

Contributor

Filippo Lancieri, Fellow, George J. Stigler Center for the Study of the Economy and the State. JSD Candidate, University of Chicago Law School

DISCLAIMER: The purpose of these reports is to identify what new challenges digital platforms pose to the economic and political structure of our countries. These reports also try to identify the set of possible tools that might address these challenges. Yet, there is potential disagreement among the members of the committees regarding which of these problems is most troubling, which tools might work best, whether some tools will work at all, or even whether the damage some tools might produce is larger than the problem they are trying to fix. Not all committee members agree with the findings or proposals contained in these reports. The purpose of these reports, thus, is not to unanimously provide a perfect list of policy fixes but to identify conceptual problems and solutions and to start an academic discussion from which robust policy recommendations can eventually be drafted.

*** The Committee in-person meetings were partially supported by a grant from the Alfred P. Sloan Foundation, whom we also thank for supporting this project**

Table of Contents

Executive Summary 271

Introduction..... 277

The Political Economy of Major Social Media Platforms..... 279

 A. The Sources of Political Power 280

 1. Structural Power..... 281

 2. Financial Power and Resources for Lobbying 281

 3. Role as a Media Outlet..... 287

 4. Complexity..... 289

 5. Connectivity 290

 6. “National Champions” 291

 B. Summary 292

The Political Effects of Social Media..... 293

 A. Political Mobilization..... 295

 B. Political Engagement 296

 C. Limiting Power of Information Gatekeepers 297

 D. Political Polarization and Manipulation of Information 298

 E. Harmful Speech 301

 F. Summary 305

Reforming the Social Media Platforms..... 305

Policy Recommendations..... 309

 A. New Regulatory Authorities 310

 B. Antitrust Enforcement or Other Policies to Prevent Political Market Concentration 315

 C. Platform Liability..... 317

 D. Role of Social Media in Campaigns and Elections..... 318

 E. Philanthropic Disclosure..... 320

 F. Data Access for Academic and Independent Research..... 321

Conclusions..... 325

References 328

Executive Summary

The emergence of social media and its dominant platforms has profoundly transformed many aspects of economic and social life. But few of these transformations have been as heatedly debated as social media's impact on political institutions and behavior. Social media was once touted as a powerful accelerant of democratization and democratic renewal. It would provide voice to the powerless and would spur collective action to overthrow authoritarians and reform democracies. But revelations that Russian intelligence sought to use social media platforms to influence the 2016 US presidential election and the UK Brexit referenda have cast a pall on this early optimism. Concern about the impact of social media has grown steadily as policymakers and the public increasingly view social media as a megaphone for fake news, the mobilizer of extremists, and a polarizer of society.

While innovations in communication from the printing press to cable television profoundly altered the relations between the governor and the governed and transformed the potential for democratic and accountable governance, there are valid concerns that the political impact of social media may be greater than that of these earlier innovations. The massive scale and reach of social media allows a single post to reach millions of users. Its platforms facilitate anonymity, which enables misinformation and promotes harassment and hate speech. But most importantly, these features are exacerbated by the network externalities that push social media platforms towards natural monopoly. Consequently, a technology with tremendous potential to reshape politics is controlled by a few firms. In a more competitive environment, users might flee platforms that are overly prone to electoral manipulation and misinformation. But platform monopolization removes the accountability of competition.

The Stigler Center has charged our committee with exploring the political impacts of social media and its most prominent platforms. In taking up this mandate, we focused not only on the ways in which social media usage has the potential to shape political outcomes but also the ways in which social media platform companies are emerging as uniquely powerful political actors. While political and legal scrutiny of the platform companies is currently high, these firms have a number of formidable political assets. Beyond sheer size and economic clout, these firms are advantaged by First Amendment protections, the complexity and opacity of their algorithms and internal policies, and their connectivity to users and others who may be politically mobilized.

Moreover, social media platforms benefit from economic nationalism as countries race for advantages in digital technology and artificial intelligence. Few if any firms have ever had such a rich collection of advantages.

The political power of the platform companies complicate reforming social media. Most obviously, these political advantages can be employed to limit government oversight and regulation. Moreover, regulatory authority created during this period of political weakness may later be captured by the industry as public interest in reforming social media wanes. Most importantly, the power of these firms and their control of the relevant data allows them to avoid greater public scrutiny.

Our report also maps out our current understanding of how social media and digital platforms impact the broader political systems. As platforms such as Facebook, Twitter, YouTube, and others grow in importance as a medium for political debates, so too does their potential to impact political outcomes more broadly. However, our knowledge of the political impacts of social media remains in its infancy given that the data necessary to independently evaluate social media's effect on political outcomes remains proprietary and largely unavailable to researchers and the public at large. The political impact of social media remains in its infancy given that the data necessary to evaluate social media's effect on political outcomes remains proprietary and unavailable to researchers. Without better access to such information, academia, think tanks, and other civil society organizations can do little to hold social media accountable for the possible distortions of our democracy.

Policy Recommendations

We outline proposals designed to mitigate the political impact of social media and the political effects of digital platform concentration. Our most important recommendation is our concurrence with the other subcommittees that significant government regulation and greater antitrust scrutiny is warranted. Our contribution to that discussion focuses on regulatory structure, laying out several principles to help insulate regulatory authorities from excessive industry influence while preserving democratic accountability. We also address important issues related to disclosure and transparency. First, we endorse updating campaign finance law to cover spending on social media campaigns. Second, we call for more transparency in the use of platform companies' support for research on social media, and for greater dissemination of

internal research. Finally, we suggest that a new Digital Authority can be essential for facilitating greater independent research by ensuring that scholars have access to relevant social media data.

1. New Regulatory Authorities

We recommend the creation of a new regulator and enhanced arrangements for inter-agency cooperation. The following principles should guide the creation of a Digital Authority (DA) tasked with regulating digital platforms:

1. The DA should have a reasonable degree of autonomy from industry influences to make decisions about social media platforms in the public interest.
2. The jurisdiction of the DA should cover as many social media-related functions as possible to prevent regulatory fragmentation.
3. Mechanisms for coordination with other agencies should be created
4. The DA should have responsibility for rulemaking in the following areas:
 - i. General consumer protection
 - ii. Privacy policies and disclosure
 - iii. Transparency
 - iv. Data portability
 - v. Data and algorithmic access for external auditing and research
5. The DA should have authority to create mechanisms for real time data collection from the platforms (subject to appropriate protections for user privacy).
6. The DA should have research capacity to undertake studies of the impact of the platforms on social and political outcomes.
7. The DA should play a facilitating role in generating independent research by outside scholars.

8. The DA should have the authority to review relevant internal studies conducted by the platform companies. When the release poses no undue privacy violation or exposure of business secrets, studies should be made publicly available.
9. Rules and regulations should be fostered in a way that promotes innovation and competition in the digital media sphere.

2. Antitrust Enforcement or Other Policies to Prevent Political Market Concentration

Many of the negative political by-products of social media are associated with the lack of competitive markets for digital platforms. Therefore, policies aimed at reducing "political concentration" should be developed.

Contemporary antitrust enforcement is generally predicated on a consumer welfare standard. In the case of social media, this standard may be inadequate to account for the political impact of concentration. Economic concentration concentrates political power. Large firms who lack competitors are hugely advantaged in the political marketplace. Second, concentration may exacerbate the negative consequences of the role of social media in the political system. The lack of competition deprives us of a marketplace of ideas that might serve to regulate the platforms' policies on speech and political activity. These political effects of concentration are unlikely to ever be captured by the consumer welfare standard.

Whether the antitrust law should broaden its scope beyond the consumer welfare standard is a complex and controversial issue. But the harms to citizens through the distortion of political processes should be given considerable weight in policies aimed at fighting market concentration. At a minimum, the DA should develop methodologies for evaluating the explicit political impact of social media concentration. Such methodologies may contribute to the establishment of a system of dual review such as that in place for mergers involving broadcasters, where the FCC has a dual mandate that complements that of antitrust authorities but considers different criteria when assessing the consequences of concentration.

3. Role of Social Media in Campaigns and Elections

We endorse two campaign disclosure provisions that have been proposed as part of the Honest Ads Act. The first amends the definition of “electioneering communication” to include internet or digital communication. The second is a mandate that digital platforms compile publicly available databases of political advertisements that are run on the platforms.

These provisions do not cover, however, all political activity on digital platforms that we might like to be disclosed. We also endorse a disclosure requirement on political advertising paid for by foreign entities. Similarly, the Honest Ads Act contains no mechanism to compel buyers of political advertisements to truthfully reveal their identities. The records compiled by digital platforms would be more informative if these issues were addressed.

Furthermore, nothing in the Honest Ads Act requires that digital platforms themselves be politically neutral. There are concerns that the regulation or imposition of political neutrality by the DA might impinge upon First Amendment protections. So we support further and strong disclosure requirements that would reveal such non-neutral platform policies. Such disclosures should cover situations i) when the platforms provide specific support or technical assistance to political parties, candidates, or interest advocacy groups, outlining what type of support has been provided and what the outcome of this support was; and ii) when the platforms make algorithmic changes that directly impact how users see political content and the outcome of such changes.

c. Platform Liability

Section 230 of the Communications Decency Act of 1996 says, “No provider or user of an interactive computer service shall be treated as the publisher or speaker of any information provided by another information content provider.” The Politics committee reached no consensus on the desirability of amending or repealing Section 230. Removing the liability protection of the media platforms would undoubtedly spur them to undertake much more aggressive content moderation in an effort to avoid litigation related to slanderous and harassing speech. But the absence of liability protection might induce the platforms to police over-aggressively and have an unduly chilling effect on speech.

e. Philanthropic Disclosure

Large digital firms have extensive philanthropic efforts that serve many worthwhile causes. Yet philanthropic efforts that support research and teaching on technology and associated policy issues often create conflicts of interest, as such support may make it more difficult for technology and policy scholars to criticize the platforms and their social impact.

First, there should be greater and more transparent disclosure of the philanthropic efforts of social media companies, especially those tangibly related to teaching and research. Second, the new DA should create an office of research to help facilitate independent research on social media companies and platforms. Such research will avoid the inherent conflicts of industry-supported research. We also encourage universities and academic associations to develop disclosure standards that would apply to scholars supported by social media firms. The disclosure policy of the American Economic Association is one plausible model.

f. Data Access for Academic and Independent Research

Independent research on the economic and political effects of social media is crucial to ensuring that the platforms enhance citizen well-being. Currently, the major impediment to such research is data access. The lack of access to data for academic researchers does not, of course, mean that no research is being conducted. Instead, it means that the only people who are able to conduct such research are those working inside the platforms. We offer two proposals. A major initiative of a new DA ought to be to facilitate independent research. This could include making the data it obtains from social media firms available for research (with suitable restrictions for individual privacy and proprietary secrets). Second, we encourage the reconsideration of the presumption that data collected by social media platforms *ought to be considered proprietary at all*. There are a number of possible proposals in this regard. The strongest of which would be to recast the role of the platforms not as owners of the data provided by users, but rather as stewards of that data, entitled to use it to improve their own business models but not necessarily to prevent others from using the data for welfare maximizing purposes.

Introduction

Following revelations that foreign agents working on behalf of Russian intelligence sought to manipulate information on several digital media platforms to influence the 2016 US presidential election and the UK Brexit referendum, there has been heightened scrutiny of the political impacts of digital media. While much of that attention has focused on the role of propaganda and misinformation in democratic elections, the impacts are far broader. Social media and Internet search have become a major tool for political mobilization and engagement—for good and for bad. They have the potential to democratize speech by offering platforms to those formerly without voice. At the same time, however, they have been blamed for exacerbating polarization in societies across ideological, partisan, racial, and ethnic lines. The relative anonymity of some social media also facilitates the propagation of hateful ideologies and of political and social harassment.

Of course, social media and search are far from the first information technologies to reshape the political environment. Innovations in communication from the printing press to cable television have profoundly altered the relations between the governor and the governed and transformed capacities for democratic and accountable governance. Yet there are valid concerns that social media and search may be different from the innovations that have come before them. Their massive scale and reach allows a single post to reach millions of users. They facilitate anonymity, which can enable the spread of misinformation and may reduce barriers to harassment and hate speech. They severely degrade the ability of traditional journalistic gatekeepers to control who produces—or is exposed to—political information. The national origins of information can now be almost completely concealed. Of course, earlier innovations had some of these features, but none had all of them. Moreover, these concerns may be exacerbated by network externalities and other market characteristics that push social media platforms towards natural monopoly. So in the end, the technology with the most potential to reshape modern political institutions and outcomes falls under the control of just a few firms, who themselves are enormously powerful political actors. In a more competitive environment, users might vote with their feet against platforms that are overly prone to electoral manipulation and misinformation. But platform monopolization removes the accountability afforded by competition.

The Stigler Center has charged our committee with exploring the political impact of digital media and its most prominent platforms. First, we discuss how digital media platform companies have emerged as uniquely powerful political actors. Here we take a broad definition of digital media which encompasses social media, search engines, and messaging applications, along with other platforms that share similar characteristics, such as Amazon and Uber.⁴⁰⁵ We identify several political advantages intrinsic to the network externalities central to their business models. We also discuss advantages associated with some digital platforms but not others. Importantly, the major platforms, such as Facebook and Google, enjoy all of them. Second, we focus not only on the ways in which digital media usage can shape political outcomes. Here we are more focused on social media platforms and search engines.

Our report is divided into three major sections. The first (section 2 below) focuses on digital platforms as political actors with the direct motivation and capacity to influence the public discourse to their advantage.⁴⁰⁶ While political and legal scrutiny of digital platforms is currently high, these firms have a number of formidable political assets at their disposal—assets whose combination is probably unique to these companies. Beyond their sheer size and economic clout, these firms are becoming the pathways through which politicians reach their constituents, protected by First Amendment provisions, shielded by the complexity and opacity of their algorithms and internal policies, and massively connected to users and others that may be politically mobilized. Moreover, digital platforms benefit from the current wave of economic nationalism as home governments support them in the hopes of competitive advantages in digital technology and artificial intelligence.

These political advantages complicate policy remedies in a variety of ways. First and most obviously, these political advantages can be employed to stave off strenuous government oversight and regulation. Moreover, even those regulatory authorities that are created during this period of political weakness may later be captured by the industry when the memories of motivating events such as the 2016 election and Brexit fade. But second, and perhaps most

⁴⁰⁵ By social media, we mean platforms that allow for user-generated content that can be shared—and often annotated—across horizontal networks (Ackland 2013).

⁴⁰⁶ While many of the elements of the power of social media firms are common across political systems, we will focus on their role in the US political system.

importantly, the power of these firms and their control of the very data necessary to facilitate external analysis can allow them to avoid greater public scrutiny.

This leads us to our second section (section 3 below), which maps out the current understanding of how social media and digital platforms impact broader political systems. As platforms such as Facebook, Twitter, YouTube, and others grow in importance as a medium for political debate, so too does their capacity to impact political outcomes more broadly (not only to their advantage). As we discuss in this report, however, our knowledge of the political impacts of social media and search is severely wanting—in part because the platforms control many of the data necessary to independently evaluate the effect on political outcomes, and those data remain proprietary and largely unavailable to researchers and the public at large. Without better access to such information, academia, think tanks, and other civil society organizations are hindered in efforts to hold social media accountable for the possible distortions of our democracy.⁴⁰⁷ The role of social media and search platforms as essential political infrastructure also complicates efforts to regulate them, as politicians and political parties will attempt to use those processes to reshape digital regulation for electoral advantage.

Finally, the third and final section of this report (sections 4 and 5 below) builds on this dual diagnosis to study policy options to regulate and reform social media and digital platforms more broadly so as to minimize their negative consequences on the political environment. It does so by first studying the internal incentives for companies to adequately self-regulate and, upon concluding that they are insufficient, addressing a range of other policy options that should be considered to better align these companies with the broader public interest. While many of the reform principles and proposals should apply to many national and international contexts, our focus will be on those policy reforms that might be undertaken in the United States.

The Political Economy of Major Social Media Platforms

Any discussion of the political impacts of social media platforms should begin with the political influence and power that emerges from the economic concentration of such a key industry. Absent any other concerns about regulating these platforms, this concentration of

⁴⁰⁷ <https://www.ft.com/content/fbb11010-69d8-11e9-80c7-60ee53e6681d>.

political influence alone would be a troubling development for American democracy. But given that we and other working groups identify several areas in which government intervention may be desirable or necessary, the political economy of digital platforms—how they can use their political power to their own advantage—becomes central to addressing the societal concerns around social media. In particular, we are concerned that any reforms designed to address economic concentration, privacy issues, or political manipulation be designed to minimize the possibility that digital platforms will have the ability to block, delay, or undermine them through the use of lobbying and other political strategies. We believe these considerations are extremely important at this juncture—any effective regulatory regime rises from the combination of appropriate tools, justifications, and political will. Following concerns about political manipulation in the 2016 election and the Brexit vote, revelations about the lack of privacy, security and data breaches, and the backlash against Amazon’s approach to its H2Q location, political will to regulate digital platforms may well be at a peak.⁴⁰⁸ But regulation has to be able to sustain itself well after the attention of elected politicians and voters is directed elsewhere. In some cases, political incentive compatibility might require ensuring an alignment of incentives between the firm and regulator. But it may also be worth considering certain interventions directed at reducing their political power (subject to First Amendment concerns).

A. The Sources of Political Power

Scholars have long debated the extent to which individual corporations and industries are politically influential. While political scientists and economists tend to predict that concentrated interests win out over more diffuse interests, there have been many instances where concentrated economic power has been checked in favor of the interests of voters and consumers. Yet there are many reasons to believe that the major digital platforms may pose special political risks.⁴⁰⁹ In particular, the social media platforms enjoy a unique constellation of structural political advantages that may transform them into some of the most successful political agents of our

⁴⁰⁸ Indeed, even companies such as Facebook are now advocating for the need for government regulation. See https://www.washingtonpost.com/opinions/mark-zuckerberg-the-internet-needs-new-rules-lets-start-in-these-four-areas/2019/03/29/9e6f0504-521a-11e9-a3f7-78b7525a8d5f_story.html?utm_term=.f6a6a7317474.

⁴⁰⁹ Of course, there are policy areas in which the major platforms are likely to be in opposition. In such cases, their individual political influences may be offsetting. However, in many of the key areas discussed in the reports in this project—privacy and antitrust—the platform companies are likely to have shared interests.

times: (i) They have the structural power that comes with being a large corporation, (ii) they have unique financial power and resources to lobby politicians and regulators, (iii) their role as a media outlet allows them to both become an important pathway through which politicians reach their constituents and claim First Amendment protections for proposed regulatory changes, (iv) their complexity and internal opacity complicates the development of effective regulatory tools, (v) their connectivity can allow platforms to directly engage users in challenging political initiatives that disadvantage them, and (vi) their growing importance as leading US exporters allows them to raise “national champion” arguments.

These advantages are discussed in more detail below.

1. Structural Power

It has long been recognized that large firms have substantial political clout.⁴¹⁰ Much of this influence arises directly from the economic resources that such firms deploy into politics. Large firms generally organize political action committees that use employee contributions to make direct contributions to political campaigns. Subsequent to the US Supreme Court decision in *Citizens United v. Federal Elections Commission*, corporations can spend unlimited sums on independent expenditures that advocate the election or defeat of candidates. Large corporations spend even larger sums of money in directly lobbying legislators and regulators. Philanthropy and public relations can be used to create corporate goodwill that provides insulation from political pressure. However, in addition to these direct influences, politicians and regulators may provide favorable treatment to large, structurally important firms so as to avoid politically-costly adverse employment, investment, and growth effects.⁴¹¹

2. Financial Power and Resources for Lobbying

While economic power and resources enhance the political standing of firms in any sector, the economic positions of technology firms are currently large outliers. In May 2018, the five largest firms in the world by market capitalization were Apple, Amazon, Alphabet, Microsoft, and Facebook. Moreover, these firms hold combined cash reserves worth hundreds of billions of dollars, and all deployed substantial resources towards political campaigns and

⁴¹⁰ See Chen (2019) for a recent review of the political impact of business in the US.

⁴¹¹ See Lindblom (1983) and McCarty, Poole, and Rosenthal (2013).

lobbying.⁴¹² In the most recent cycle, Amazon and its employees spent almost \$13.5 million on political campaigns, while Microsoft spent \$14 million. All together, these five firms spent around \$39 million.⁴¹³ But campaign contributions is not the area where these firms stand out.⁴¹⁴ Their footprint is much larger in the area of direct lobbying. According to lobbying disclosure reports analyzed by [opensecrets.org](https://www.opensecrets.org), Alphabet, Amazon, and Facebook were the Second, sixth, and ninth most prolific spenders (respectively) on direct lobbying among American corporations in 2018.⁴¹⁵

Not only is the level of lobbying by the social media companies high, but their lobbying has grown rapidly over the past five years. Figure 1 shows the lobbying expenditure for the major digital media companies since 2013. With the exception of Microsoft, whose expenditure remained level, all of the companies ramped up lobbying expenditures to an extraordinary degree, especially Amazon. The increases in this sector are all the more remarkable in that aggregate lobbying expenditure in the United States has been flat or declining.⁴¹⁶

⁴¹² Twitter is a much smaller firm which is reflected in its more modest political expenditures. In the 2017–2018 electoral cycle, Twitter contributed just under \$300k while spending \$1.6 million on lobbying.

⁴¹³ Because of restrictions on fundraising and contributions by the corporate political action committees, spending by the corporate PACs represents a relatively small share of the total, with the rest coming from employees. But given the income-skew of campaign contributions, it is likely that the bulk of these individual contributions came from high-level executives who would broadly share the firms' political aims.

⁴¹⁴ Among corporate donors, Microsoft, Amazon, and Alphabet are the 11th, 12th, and 18th largest. Facebook and Apple are much further down the list.

⁴¹⁵ ATT is number 3 and Comcast is number 5. So five of the top eight are communications companies. The others are major defense contractors. See <https://www.opensecrets.org/lobby/top.php?indexType=s&showYear=2018>

⁴¹⁶ According to [opensecrets.org](https://www.opensecrets.org), aggregate lobbying expenditure was \$3.42b in 2018 down from a peak of \$3.5b in 2009.

Figure 1

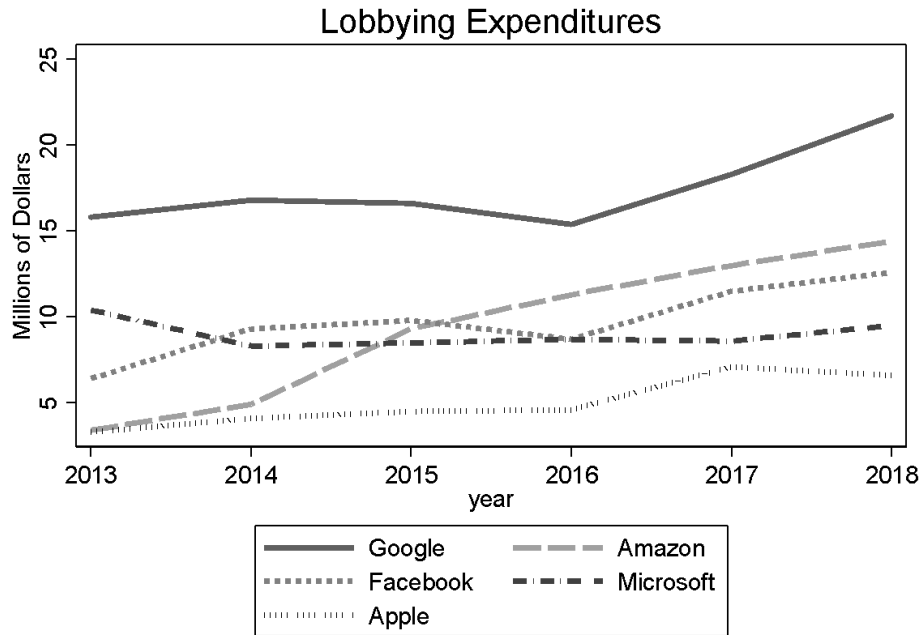


Figure 2 displays some additional important information about the lobbying strategies of Alphabet, Amazon, and Facebook. Each figure, generated from the Lobbyview database, shows the flow of money through lobbying representatives (in-house or outside firm), policy issues, and institutional venues.⁴¹⁷ These figures reveal several facets of these firms’ advocacy strategies. First, these firms primarily use in-house lobby firms rather than contracting to outside lobbying firms. This suggests that each has developed a very large sustained institutional presence in Washington DC.⁴¹⁸ Such investments in internal lobbying capacity are typical of heavily regulated industries and those that depend on government procurement. That these firms have made such investments suggests the importance they place on avoiding government regulation while obtaining other favorable policies. The second aspect of these firms’ advocacy strategies is the diverse mix of issues. The influence efforts of the media platforms are much more expansive than the domain of information technology, broadly construed. Amazon’s efforts

⁴¹⁷ See Kim (2018).

⁴¹⁸ These companies also lobby extensively in the US states and municipalities. Yet due to varying state disclosure requirements and little or no disclosure in corporate reports, less is known about the full extent of these activities than those at the federal level. See <https://www.weinberg.udel.edu/IIRCiResearchDocuments/2017/02/Corporate-Lobbying-in-the-States-FINAL.pdf>.

cover a particularly wide swath of policies. In fact, Amazon lobbies more on issues related to aviation, taxation, and trade than it does on technology or the computer industry. Similarly, Facebook is a major player on immigration, while Google is extremely active on consumer product safety, as is Apple on taxation. The third important feature of tech's lobbying strategy is that it is not exclusively focused on the legislative process. With the exception of Facebook, the tech platforms routinely lobbied executive branch office and regulatory agencies in 2018.⁴¹⁹ These include those, such as the Federal Trade and Communications Commissions, which might be directly involved in any efforts to regulate the platforms. Thus, the influence of the platforms at the agency level is an important consideration in designing any new regulatory structure.

Of course, campaign and lobbying expenditures are not the only financial resources that these companies bring to bear. Not only is philanthropy an important generator of goodwill, but corporate support of university and think tank research is a more direct source of influence.⁴²⁰ Hard data on these expenditures is harder to come by, but the concern was visibly demonstrated in the controversy surrounding allegations that the New America Foundation divested its Open Markets program to placate its major donor, Google.⁴²¹ Another avenue of corporate influence not captured in lobbying reports is the role of these firms in creating express advocacy organizations to lobby for policies that may directly or indirectly benefit the firms. For example, Mark Zuckerberg, in collaboration with other tech industry leaders, founded fwd.us which advocates for immigration reforms that might expand the number of high-skilled immigrants available to work in the tech sector.⁴²²



⁴¹⁹ Facebook's lack of executive branch lobbying in 2018 is unusual. In prior years, they were active in lobbying the Departments of Justice and Commerce.

⁴²⁰ For evidence of the impact of such incentives on research in corporate finance, see Zingales (2013). For evidence that philanthropy is a complement of firms' political strategies, see Bertrand et al. (2018).

⁴²¹ See <https://www.nytimes.com/2017/08/30/us/politics/eric-schmidt-google-new-america.html>.

⁴²² See <https://www.vox.com/2018/7/18/17575156/mark-zuckerberg-interview-facebook-recode-kara-swisher> and <https://www.fwd.us/immigration/>.


Figure 2: 2018 Lobbying Footprint of Major Tech Firms



Lobbying Finance

Methodology

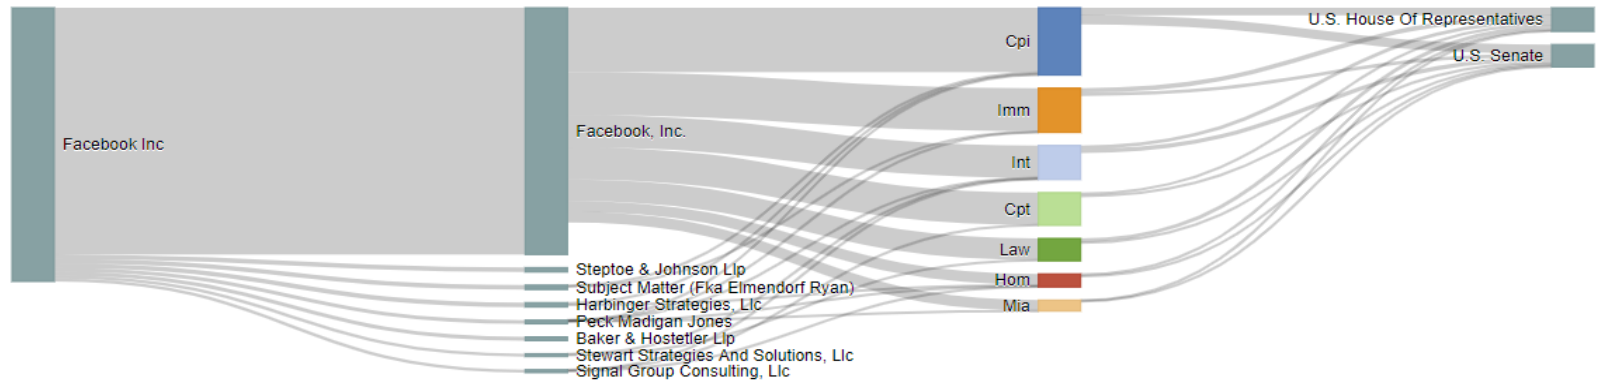
Lobby Firms Issues Govt. Entities Politicians* Year: 2018



Lobbying Finance

Methodology

Lobby Firms Issues Govt. Entities Politicians* Year: 2018



Issues Key:

AVI	Aviation/aircraft/airlines	INT	Intelligence and surveillance	SCI	Science and Technology
CPI	Computer Industry	IMM	Immigration	TAX	Taxation/internal revenue
CPT	Copyright, Patent, and Trademark	LAW	Law enforcement/crime/criminal justice	TEC	Telecommunications
CSP	Consumer issues/safety/products	LBR	Labor issues/Antitrust/Workplace	TRA	Trade
EDU	Education	MIA	Media (Information and Publishing)		
HOM	Homeland Security				

Unfortunately, political power generated by economic resources is notoriously difficult to control. Lobbying is a protected activity under the First Amendment. The current Supreme Court will protect the rights of corporations to make unlimited independent electoral expenditures, and corporate philanthropy has become a primary source of research funding. These facts have led many scholars and activists to advocate for the aggressive use of antitrust remedies to curb political influence.⁴²³

3. Role as a Media Outlet

Digital media platforms such as Facebook, Twitter, and Google operate as news aggregators and portals. Thus, not only will these companies be able to avail themselves of certain press protections enshrined in the First Amendment but, more importantly, they are increasingly embodying and controlling the means through which politicians reach their constituents.

Starting with the latter, modern political campaigns are increasingly waged over the internet. The News Media Committee reports how news consumption has moved online thus making digital platforms an important intermediary in news publication. There is also evidence of how political debate is slowly moving online, which we discuss below in the section on political effects of social media.

The problem is that only a handful of companies with political and policy motives control this new “public square.” Thus, we must be concerned that they may use it for their own advantages. It is well known that platforms such as Facebook provide important support to certain political campaigns, something they claim to do “neutrally” although candidates have paid different advertising rates.⁴²⁴ It is also known, however, that shifts in their algorithms and policies can impact outcomes—as shown by changes in Facebook impacting voter turnout.⁴²⁵ The question that imposes itself is: What would be the political implications of these increasingly essential platforms choosing to promote one candidate or idea at the expense of another?

⁴²³ See Khan (2016) and Wu (2018).

⁴²⁴ See <https://techcrunch.com/2018/06/11/facebook-says-it-gave-identical-support-to-trump-and-clinton-campaigns/>. One must stress that it is also known that the Trump and Clinton campaigns paid different rates for advertisement on Facebook: <https://www.wired.com/story/facebook-trump-clinton-campaign-ad-cpms/>.

⁴²⁵ See Bond et al. (2012).

The recent controversy over Facebook barring Senator Elizabeth Warren's ads criticizing Facebook is a potential example of how much power could be exercised.⁴²⁶ Facebook's stated rationale for removing the ads was that they violated policies against the use of the Facebook name and logo in ads. However, such a policy clearly serves to make it harder to criticize the company on the platform. In this particular case, Facebook quickly reversed course after public outcry. However, the ability to publish "billions of personal editions" (to use the concept framed by the Media Report) combined with the opacity of proprietary algorithms affords these companies the opportunity to shape political outcomes without any form of public scrutiny. Conservatives have long claimed that platforms are biased against them. Given the current lack of any public data or disclosure obligations, one cannot affirm that these claims are true or false or whether Senator Warren's posts were treated fairly.⁴²⁷ Even more importantly, politicians who promote regulation against the interests of social media may fear the risk of being subject to some form of opaque discrimination when running for re-election. This media power, here understood as the power to shape the public debate and impact political outcomes, is certainly a barrier against the creation of a coalition capable of regulating digital platforms.

A second challenge, however, is that even if politicians overcome their growing dependency on social media to impose some form of content regulation, it is not clear whether these laws and regulations would be enforceable in the United States as the companies can claim First Amendment protections to certain activities. These rights may be sufficient to challenge approaches aimed at regulating content, such as the strict Network Enforcement Act (NetzDG) adopted by the German government.⁴²⁸ Indeed, some US courts affirmed that search results, among others, are protected from governmental influence by the First Amendment, a move criticized by some scholars as practically insulating platforms.⁴²⁹

⁴²⁶ See <https://www.theverge.com/2019/3/11/18260857/facebook-senator-elizabeth-warren-campaign-ads-removal-tech-break-up-regulation>.

⁴²⁷ It is important to stress that this report is **not** affirming that platforms are biased against conservatives or defending Senator Warren's proposals on tech regulation. The examples are simply used to prove a point on the risk that platforms become non-neutral on their political engagements.

⁴²⁸ This law mandates that social media companies block or remove content that violates one of several laws on hate speech and defamation. <https://law.yale.edu/mfia/case-disclosed/germanys-netzdg-and-threat-online-free-speech>

⁴²⁹ See Wu (2013) criticizing federal cases that affirmed freedom of speech constitutional protections to Google's search results, notably *Langdon v Google, Inc* (2007) 474 622; *Search King, Inc v Google Tech, Inc* (2003) 2003 21464568 (WD Okla).

The web of protections currently afforded to social media companies goes well beyond those of the First Amendment. They are also protected by provisions such as Section 230 of the Communications Decency Act of 1996, which says, “No provider or user of an interactive computer service shall be treated as the publisher or speaker of any information provided by another information content provider.” The act goes further in that it preempts state laws, and the courts have refused to limit it to Internet service providers. Among other things, this provision immunizes social media platforms from legal liability for hate speech, harassment, and misinformation. Of course, these protections also immunize the platforms in their decisions to block or censor content—potentially at the expense of politicians themselves. But currently, as discussed in the second section of the report, platforms have largely chosen to use a light touch on content moderation, though this appears to be slowly changing.⁴³⁰

The fact that social media platforms have quasi-journalistic functions, such as curation and moderation, may also complicate efforts at regulation for more purely political reasons. Setting aside the economic merits for or against the Comcast/Time Warner merger, Time Warner was able to claim that the Trump Administration had targeted it in retaliation for CNN coverage of the administration. For similar reasons we might expect that any attempts to regulate Facebook, Google, or Twitter would likely be similarly politicized by either advocates or opponents.⁴³¹

4. Complexity

Another impediment to successful regulation of the social media companies lies in the complexity of their proprietary algorithms and the privacy of the data they generate. Under conditions of complexity and low transparency, firms will always maintain a large informational advantage over regulators.⁴³² While regulation may strive to increase transparency (indeed we make several such proposals), it seems unlikely that the firms’ informational advantages could

⁴³⁰ See, for example, <https://www.theverge.com/2019/1/25/18197301/youtube-algorithm-conspiracy-theories-misinformation> or https://time.com/5619999/instagram-mosseri-bullying-artificial-intelligence/?utm_source=time.com&utm_medium=email&utm_campaign=the-brief-pm&utm_content=2019070819pm&xid=newsletter-brief.

⁴³¹ The potential for such politicization was clearly on display in the recent Senate hearings “Google and Censorship through Search Engines,” in which wildly inflated estimates of the impact of Google on the 2016 election were offered in testimony. See <https://www.judiciary.senate.gov/meetings/google-and-censorship-though-search-engines>.

⁴³² See McCarty (2013).

ever be eliminated. Moreover, the existence of regulation may give firms less of an incentive for transparency and other internal controls upon which government regulators might build.⁴³³ These issues manifest themselves in the human capital available for government regulatory agencies. First, government agencies have to compete for technical talent with the tech firms. Given the much lower salaries available to government employees, the agencies are at a considerable disadvantage. Second, government regulation may depend excessively on “in-and-outers”—regulators who have been employed in the tech sector and hope to return. Such dependencies create two problems. The first is the classic concern of regulatory capture through an implicit quid pro quo where regulators go easy on industry in hopes of securing future employment.⁴³⁴ The second is what James Kwak calls “cultural capture.”⁴³⁵ To the extent to which the tech sector and its regulators draw from the same labor pool, firm executives and regulators are likely to see the world in very similar ways. This in turn reduces the set of policy options to those that seem palatable to the firms.

Finally, complexity implies that there will generally be very little voter interest in many of the fine details of platform regulation. Thus, regulators may face little accountability for their decisions.⁴³⁶

5. Connectivity

By virtue of their huge bases of customers, social media platforms have direct connections to citizens and voters atypical of other corporations. While other corporations may have millions of customers, few have the ability to directly interact with them or have such sophisticated data on their political preferences. A good example may be the now almost forgotten debates around the Stop Online Piracy Act (SOPA) and the Protect IP Act (PIPA), where coordinated online campaigns led by websites such as Google, Twitter, and many others

⁴³³ See McCarty (2017).

⁴³⁴ A recent study of banking regulation, however, suggests that the effects of such a quid pro quo appear empirically less important than the regulator’s human capital disadvantage. See Lucca, Seru, and Trebbi (2014).

⁴³⁵ See Kwak (2013).

⁴³⁶ For this argument applied to the setting of accounting standards, see Rammana (2015).

mobilized millions of users—a Google petition in opposition to the acts quickly drew 4.5 million signatures, while Twitter spurred users to share 2.4 million tweets attacking the bills.⁴³⁷

These powers become even more potent for those platforms that also provide income to millions of self-employed workers. For example, it has been estimated that four million Brazilians earn income through Uber and other service apps—a number 35 times larger than the Post Office, the country’s largest public employer.⁴³⁸ With a large direct stake, such workers are likely to be quite susceptible to mobilization on behalf of the firm. For these reasons, the social media firms more closely resemble powerful membership organizations, such as the National Rifle Association or the AARP, than typical corporations. But technology firms have the capacity to use these connections in much more powerful ways. The data generated by the platforms are already used extensively in mobilization by political parties and social movements. That technology firms could turn their data and algorithms towards mobilizing customers on behalf of the firms’ policy positions is a serious one. In fact, such mobilizations have occurred in several cities in attempts to overturn local transportation and hospitality regulations.⁴³⁹

6. “National Champions”

Finally, just as the American steel and automobile industries responded to international competition by arguing for beneficial policies on “national interest” and security grounds, the technology sector has begun to make similar arguments in response to the success of international competitors such as Alibaba and Tencent.⁴⁴⁰ American firms argue that the combination of Chinese government support and massive vertical and horizontal integration has given these firms an unfair advantage. In a recent interview, Mark Zuckerberg was asked “There’s been some calls to break up some companies like Facebook or Amazon that become too big. Are you in fear of that in any way?” His response centered on how that would affect competition with China:

⁴³⁷ See <https://latimesblogs.latimes.com/technology/2012/01/wikipedia-sopa-blackout-congressional-representatives.html>.

⁴³⁸ See <https://economia.estadao.com.br/noticias/geral,aplicativos-como-uber-e-ifood-sao-fonte-de-renda-de-quase-4-milhoes-de-autonomos,70002807079>.

⁴³⁹ See <https://www.theguardian.com/technology/2017/apr/12/why-everyone-hates-uber-seven-step-playbook> and <https://www.businessinsider.com/heres-why-uber-is-a-political-machine-2015-7>.

⁴⁴⁰ See <https://www.csis.org/analysis/technological-competition-and-china>.

I think that the alternative, frankly, is going to be the Chinese companies. If we adopt a stance which is that, “Okay, we’re gonna, as a country, decide that we wanna clip the wings of these companies and make it so that it’s harder for them to operate in different places, where they have to be smaller,” then there are plenty of other companies out there that are willing and able to take the place of the work that we’re doing.⁴⁴¹

Such arguments are likely to be repeated throughout any efforts to enforce antitrust provisions when considering breakups or evaluating new mergers. In an era of economic nationalism, these arguments are likely to be potent. The tech firms have also elicited the support of the US government in attempts to stave off tax and regulatory burdens imposed by other states. When France recently passed a 3% “digital tax” on tech company revenues, the United States quickly announced a Section 301 investigation, which could lead to retaliatory actions against French products should the digital tax be found to discriminate against US companies.⁴⁴² A bipartisan letter from the leadership of the Senate Finance committee urged the Treasury Secretary to consider “all available tools under U.S. law to address such targeted and discriminatory taxation.” The US response to the French tax illustrates just how much clout the social media platforms command even during this period of heightened scrutiny of their business models.

B. Summary

As we have seen, digital platforms are unique in their ability to influence the political debate. While many industries and interest groups enjoy some of the five main political advantages, we are unaware of any others that possess all five. Financial and pharmaceutical firms are large and resourced and may appeal to national champion arguments, but they can draw on a narrower set of constitutional claims or have a mobilizable customer base. Among firms, media conglomerates seem to come closest, but they are normally smaller economically, have less customer data, and are less likely to use international competitiveness as an argument for favoritism. It is possible that new, very large conglomerates such as ATT/Time Warner will

⁴⁴¹ See <https://www.axios.com/tech-giants-new-defense-our-chinese-rivals-are-worse-zuckerberg-aa1a2e5c-6aa6-45b3-aa2b-8217b792d10e.html>, <https://www.vox.com/2018/7/18/17575156/mark-zuckerberg-interview-facebook-recode-kara-swisher>.

⁴⁴² Section 301 refers to that section of the 1974 Trade Act. See <https://www.ft.com/content/ba4bd9b8-a351-11e9-a282-2df48f366f7d?emailId=5d27229b017d570004d4fa2d&segmentId=13b7e341-ed02-2b53-e8c0-d9cb59be8b3b>.

attain similar powers, but even AT&T reaches fewer users than companies such as Facebook or Google and has much weaker claims to a national champions argument.

The digital media platforms do suffer from some important political disadvantages. First, even as their workforce grows, they directly employ fewer workers than many firms do, and employment tends to be geographically concentrated. While Amazon is a major employer across the country, other tech firms have smaller workforces that tend to be concentrated into tech hubs like Silicon Valley.⁴⁴³ And of course, the industry has been weakened politically by the many concerns about their role in the Russian interference in the 2016 US elections as well as elections in other countries. That the industry's closest political allies were disadvantaged by these events opens the window for bipartisan reform coalitions that might not exist during normal times.⁴⁴⁴

The Political Effects of Social Media

Having previously outlined how digital platforms may operate as political actors, this section reviews the current academic understanding of how platforms affect the broader political system. In a recent report for the Kofi Annan Foundation, Nathaniel Persily highlighted the defining characteristics of the new communication revolution, by which he means the information environment generated by the rise of the internet and social media platforms. These included the *velocity* at which information can spread, the rising importance of *virality* in guiding the production of information, the opportunity for *anonymity* on the part of providers of information, and the possible emergence of *homophilous* communities that themselves can become major providers of information to their members.⁴⁴⁵ All of these characteristics flow from the key features of what has often been called Web 2.0: user-produced content that can be annotated and shared across horizontal networks.⁴⁴⁶ Echoing our analysis of the previous section, Persily also notes the *unprecedented monopolistic* potential of the two leading social

⁴⁴³ As mentioned earlier, platforms in the share economy have millions of de facto employees spread throughout many jurisdictions.

⁴⁴⁴ However, the US response to the French digital tax suggests that the platforms have not been weakened that much.

⁴⁴⁵ See Persily (2019).

⁴⁴⁶ See Ackland (2013).

media companies—Facebook and Google—as well as the extent to which their international position presents challenges to the *sovereignty* of individual nations.

The political implications of the rise of social media are related to both the characteristics of the new communication revolution as well as the concerns about the role of media platforms as political actors. Consequently, our report covers both of these concerns and argues that they are tightly related. In considering the political impacts of social media, we argue that social media platforms have a number of features that could directly affect political behavior and engagement while also generating opportunities that other political agents may exploit both for good and for ill. While the analogy is imperfect, these features are generated in much the same way as classic economic externalities. First, the political impacts are mostly tangential to the platforms' business models, which are focused primarily on maximizing digital advertising revenue. Thus, the platforms have very little private incentive to take actions that eliminate or manage the impact of their products on political outcomes. Second, for the overwhelming majority of users, political engagement is not a primary or even secondary reason for engaging with a media platform.⁴⁴⁷ As a result, most users lack any incentive to police the platform policies which may generate political side-effects. And those consumers that do care will face tremendous collective action problems compounded by the large degree of user concentration and monopolization of the platforms. So as in the case of classic externalities, democratic opportunities and outcomes of millions of third parties may be affected by the policies of the platforms and the behavior of users.

In this section, we consider six *potential* political consequences of the widespread use of social media: reduced costs to *political mobilization*, increased levels of *political engagement*, reduced power of *traditional media gatekeepers*, the rise of *political polarization*, the ability of *malicious actors to manipulate information*, and the increased opportunity for the use of *harmful speech*. While we order these from those that seem to have the most positive net benefits for society writ large to those with the most negative effects, it is important to recognize that complex relationships are at work. Indeed, we use the term “potential” in describing many of these categories because scholars are not yet sure about the overall net impact of social media in many of these cases—largely because the lack of access to proprietary data hinders high-quality,

⁴⁴⁷ See Guess, Munger, Nagler & Tucker (2019).

independent academic research. So it is perhaps best to think of this list as summarizing what is most frequently discussed—as well as what we currently know—about the consequences of social media usage for politics.

A. Political Mobilization

While perhaps difficult to remember today, social media first burst onto the collective political consciousness as what Larry Diamond termed “Liberation Technology.”⁴⁴⁸ Social media would help spread democracy across the globe by giving democratic oppositions the opportunity to enjoy the fruits of a communication tool outside of the control of authoritarian regimes. To give a key example, social media have repeatedly been lauded for facilitating the organization of protests in authoritarian regimes. Social media can be harnessed to connect prospective democratic opponents of the regime, to plan protests, to spread information about protests in real time, and to keep networks of protesters connected after the physical protests have concluded. Social media have been posited to lower both the real and perceived costs of participating in protest against oppressive regimes. The real costs are lowered not only by making access to information about protests and protest movements easier to acquire, but also by decreasing the costs associated with sharing information with large numbers of people.⁴⁴⁹ Uncertainties such as the possibility of showing up for a protest with few other fellow protesters or that of facing police action and physical harm when protesting can also be decreased by social media, which allows would-be protesters to gauge the popularity of the regime and the reservoir of support for nascent opposition movements.⁴⁵⁰ Social media also allow regime opponents to accumulate information in real time about police activity, violence, and safety. The Ukrainian Euromaidan protests, which eventually led to a regime change, famously began with a post on Facebook noting that if at least 1000 users “liked” the post, all those users would all join a protest the following day.⁴⁵¹ The role of social media in the events that came to be known as the “Arab Spring” is also well documented.⁴⁵²

⁴⁴⁸ See Diamond (2012).

⁴⁴⁹ See Tufekci (2017).

⁴⁵⁰ See Cantoni et al. (2017) for counter-evidence that people might actually stay away from protests when they learn large numbers of people are planning on attending.

⁴⁵¹ See Metzger and Tucker (2017).

⁴⁵² See Howard and Hussain (2013).

While originally touted as a tool for organizing in closed societies where democratic activists were shut out of more traditional forms of media, the lure of social media as a tool for political mobilizations quickly became apparent in more open societies as well.⁴⁵³ Perhaps the most dramatic example is the original Women’s March in the United States in January 2017, which grew out of a Facebook post from “a grandmother in Hawaii.” Indeed, today it is difficult to find an example of a significant protest event anywhere that does not have a social media component. Nevertheless, there has been pushback against the idea that social media inherently facilitates protest behavior. Most prominently, Evgeny Morozov has argued that information technology can be just as useful to oppressive regimes looking to stifle dissent. Others doubt that the benefits for grassroots action are real; Malcolm Gladwell’s “Slacktivist” argument posits that social media gives people the false belief that they can participate in a movement simply by clicking “like” on a protest post, without having to engage in the relatively more costly activity of offline protesting or movement building.⁴⁵⁴ This in turn has led to counterarguments about the value of the “weak ties” observed in social media networks for political participation,⁴⁵⁵ the role that peripheral members of online networks can play in spreading information about protests,⁴⁵⁶ and evidence of the tighter online networks among protesters as opposed to comparable groups of non-protesters.⁴⁵⁷

B. Political Engagement

Social media platforms may also foster many other forms of political and social engagement by lowering the costs of participation and ameliorating collective action problems. Indeed, the original proponents of “e-government” touted the Internet as a force that would bring the rulers closer to the ruled by creating a common space in which both could communicate with one another.⁴⁵⁸ Such communication has often come from the top down, effectively allowing elites to communicate with masses without having to work through the press (e.g., Alexandria Ocasio-Cortez) or by using social media to dictate press coverage (e.g., Donald Trump). There

⁴⁵³ See González-Bailón et al. (2011).

⁴⁵⁴ See Gladwell (2010).

⁴⁵⁵ See Bennet and Segerberg (2012).

⁴⁵⁶ See Barberá et al. (2015).

⁴⁵⁷ See Larson et al. (2019).

⁴⁵⁸ See Gil-Garcia et al. (2005).

are, however, instances of politicians engaging in two-way conversations with their constituents on social media (e.g., Senator Cory Booker). Indeed, elites have embraced social media as a tool for direct communication with the public in both democratic and non-democratic regimes.⁴⁵⁹

Social media have also been shown to be a powerful tool for encouraging the most common form of political engagement: voting. A 61-million-person experiment in the 2010 US national election presented users of Facebook with a banner reminding them of Election Day, showed them some information about friends' behavior on Facebook, and then provided the opportunity to note on Facebook that they had voted and to seek out more information about how to vote. The result was that not only did exposure to the banner have a direct effect on the respondent in question, it also had indirect effects on friends of people exposed to the banner. Nor was this effect limited to online behavior: The researchers also observed an increase in actual voter turnout.⁴⁶⁰ The same features that make social media such powerful tools for organizing protests also should allow it to be used for organizing other forms of civic activism, including voter registration drives and get-out-the-vote campaigns.⁴⁶¹ Of course, the question has emerged whether the same tool that was used to encourage voter participation can now be used to suppress voter participation or to target a given political outcome.⁴⁶²

C. Limiting Power of Information Gatekeepers

The defining feature of social media is that user-generated content can be spread horizontally through peer-to-peer networks. Putting aside concerns about the quality of the content, this represents an obviously dramatic change from previous eras when information about politics came top-down from media companies in the form of newspapers, magazines, radio, and television. Thus, social media can give voice to actors who previously did not have access to mainstream media, and undoubtedly these platforms democratize access to information.⁴⁶³ As discussed previously with respect to authoritarian regimes, these features mean that pro-democratic voices that would normally be excluded from state-controlled press and television now have a platform for publishing and distributing political news and

⁴⁵⁹ See Barberá et al. (2018), Neblo, Easterling, and Lazer (2018), and Siegel et al. (2019).

⁴⁶⁰ See Bond et al. (2012).

⁴⁶¹ See Aldrich et al. (2016).

⁴⁶² See Bradshaw and Howard (2018) and Tucker et al. (2018).

⁴⁶³ See Tucker et al. (2017).

information. But even in the context of democratic political systems, the diminished role of gatekeepers means that marginalized communities that may previously not have had a real voice within mainstream media can now be heard. Prime examples include the Black Lives Matter and Occupy Wall Street movements in the United States and the Indignados movement in Spain.⁴⁶⁴ Further, despite the received wisdom that social media puts users in a political bubble, there is growing evidence that simply the experience of consuming news via social media—accessing news stories one at a time from different sources, as opposed to watching a television broadcast or reading a newspaper—can actually lead to increased exposure to cross-cutting political views.⁴⁶⁵ Of course, limiting the power of gatekeepers also means that it is harder to prevent the spread of inaccurate information and uncivil speech, topics to which we turn in the following sections.

D. Political Polarization and Manipulation of Information

The widespread use of social media platforms has also been linked to two related negative effects: growth in political polarization and the manipulation of information. By political polarization, critics typically refer to increasing levels of outgroup dislike—and in-group favoritism—at an emotional or gut level, or what political scientists refer to as “affective polarization.”⁴⁶⁶ Information manipulation covers a range of phenomena, most commonly the spread of online misinformation, disinformation, and “fake news.” Since false or misleading content is often tailored specifically to appeal to people’s political prejudices, the “demand” for online misinformation has been linked to the rise of partisan polarization in the mass public. As we will see, the extent to which either of these phenomena are caused or exacerbated by social media remains a matter of lively debate. Still, it is important to clarify the mechanisms by which unplanned encounters with socially shared information—due to network ties, targeted advertising, algorithmic personalization, or some combination of the three—could lead to increased levels of polarization and misinformation. In the case of misinformation, people’s ability to costlessly share content that they find compelling or agreeable can transform seemingly

⁴⁶⁴ See Freelon, McIlwain, and Clark (2016), Jackson, Bailey, and Welles (2019) and Gonzalez-Bailon and Wang (2015).

⁴⁶⁵ See Wojcieszak & Mutz (2009), Messing & Westwood (2014), Barberá (n.d.), Bakshy et al. (2015), and Newman et al. (2017).

⁴⁶⁶ See Iyenger et al. (2012) for an overview.

unconnected individual behaviors on social media into unpredictable cascades of viral propaganda. The resulting fog of half-truths can inflame discourse and monopolize the limited capacity of information intermediaries. These possibly unintended byproducts interact with mass polarization by further skewing perceptions of the other side.⁴⁶⁷ When the most visible actors are the most extreme, and when information architectures flatten and decontextualize social interactions, reinforcing spirals of misinformation and distrust may become more likely.

Thus perhaps one of the most frequently asked questions about social media platforms is whether their use is polarizing society. A longstanding concern is that the internet enables people to choose sources of information that confirm their preexisting biases.⁴⁶⁸ In the case of social media, people may also choose to exclusively interact with others who share their worldview. When people self-select in this way, social pressure toward homogeneity can produce polarization and extremism. Finally, algorithms designed to learn people's preferences can reinforce these tendencies to create a feedback loop.⁴⁶⁹ Thus, platforms that give people unprecedented freedom to encounter information and seek out social connections may—inadvertently or not—contribute to the segmentation of society along partisan and other social identity-related lines.⁴⁷⁰

While the logic underlying these dynamics is well understood, evidence on these points continues to be mixed. One recent study focusing on the United States found that the greatest increase in polarization over the past 20 years occurred among the oldest age group—those least likely to be online during that period.⁴⁷¹ At the same time, fairly credible evidence has emerged linking the rollout of partisan-leaning cable news networks—whose audiences skew older—and voting patterns. Taken together, these studies raise the possibility that rising levels of polarization owe more to traditional media than to new communication technologies.⁴⁷² But

⁴⁶⁷ See Ahler & Sood (2018), Lelkes, Sood & Iyenger (2017), and Suhay et al. (2018).

⁴⁶⁸ User self-selection into “echo chambers” represents a classic collective dilemma. Each user may enhance her own personal satisfaction by engaging only with confirmatory information. Yet there may be a collective harm to political deliberation if many such users engage in such behavior. So any rationale for making self-selection harder depends on whether researchers can establish that the collective harm justifies regulation of individual choice.

⁴⁶⁹ See Jiang et al. (2019).

⁴⁷⁰ See Sunstein (2018).

⁴⁷¹ See Boxell, Gentzkow, and Shapiro (2017).

⁴⁷² See DellaVigna and Kaplan (2007), Prior (2007), Clinton and Enamorado (2014), Hopkins and Ladd (2014), and Martin and Yurukoglu (2017).

these macro-level analyses may not tell the whole story. Recent controlled experiments found that encountering ideological tweets from the other side can cause people to become more polarized, and suspending Facebook use for a month caused people to become less polarized.⁴⁷³

How can both sets of findings be true? It's important to remember that research is still accumulating regarding the effects of social media on politics and society. But studies conducted to date make several points clear. First, most people are not on Twitter, the social network that has generally been the focus of empirical scrutiny. Although influential actors and important conversations may occur on Twitter regularly, most citizens will not encounter links or messages there—at least not directly.⁴⁷⁴ Second, even for those who are on Twitter, many of these users are likely rarely engaging with political content. Once we concentrate on the subset of people who participate in online conversations or interact frequently with political information on social media, we are already looking at a group that is quite different from the rest of the population: more ideological, more politically knowledgeable, and more likely to participate in political activity. The visibility of such actors may very well influence the generalizations people tend to make about social media as a whole and their effects. And third, even if social media can drive polarization, there are other ways in which the structure of online networked interactions can simultaneously encourage the opposite. A number of studies have documented that people regularly encounter challenging information, even if by accident, via mechanisms such as “weak ties”—acquaintances or friends of friends whose views might go unnoticed in offline environments.⁴⁷⁵ Finally, there is little reason to believe that the effects should be the same across a wide variety of platforms. Ultimate conclusions must therefore wait for similar research to be done on a wider variety of platforms.

Beyond the question of the extent of homophily in online networks, much remains to be learned about the producers of online disinformation and the precise mechanisms that encourage its spread on social platforms. Given the evidence so far, popular narratives about the dangers of “fake news” tend to overemphasize its effects on people’s political decisions, such as whether or

⁴⁷³ See Bail et al. (2018) and Allcott, Braghieri et al. (n.d.).

⁴⁷⁴ According to studies by the Pew Research Center, only 22% of US adults report ever using Twitter, compared to 74% and 69% for YouTube and Facebook. See <https://www.pewresearch.org/fact-tank/2019/04/10/share-of-u-s-adults-using-social-media-including-facebook-is-mostly-unchanged-since-2018/>.

⁴⁷⁵ See Guess (2018), Barberá (2015), and Eady et al. (2019).

for whom to vote. Researchers have generally found that online misinformation is not prevalent enough in most people’s social media feeds to make a measurable difference in these outcomes.⁴⁷⁶ On the other hand, there is not nearly enough focus on hypothesized negative externalities that are harder to quantify, such as the distortion of public debates about critical policy issues and the possibility of losing consensus on a shared set of facts upon which to base disagreements.

E. Harmful Speech

The spread of harmful speech—including incivility, harassment, and hate speech—is another negative consequence of social media. From racially motivated harassment campaigns targeting journalists on Twitter to calls for ethnic violence on Facebook, harmful online speech has received increased attention from academics and policymakers alike. Fearing that online rhetoric is mobilizing offline violence and extremism, governments are passing regulation to compel social media companies to ban harmful content. However, despite increased public and scientific attention to harmful speech, definitions remain murky and systematic measurement is rare, in large part due to lack of access to reliable data.⁴⁷⁷ As a result, our knowledge of the causes, prevalence, and offline consequences of this content remain limited. However, by exploring the range of harmful language that has proliferated online and clarifying the mechanisms by which social media might facilitate its spread and amplify its impact, we can better understand harmful speech, as well as how to best combat it.

Harmful speech is an umbrella term for behaviors—spanning a range of instigators, targets, motives, tactics, and media—that cause harm. It ranges from speech that incites violence or criminal acts to speech that is “merely” offensive. It includes actions carried out by individuals as well as coordinated mob attacks, involving both those who know their targets and those who do not. Outcome-based definitions of harmful speech focus on the harm experienced by its targets, while other definitions examine the intent of the speaker or the content of the

⁴⁷⁶ See Guess, Nyhan & Reifler (n.d.) and Grinberg et al. (2019). It is worth noting, however, that much of the research has been limited to a subset of misinformation, namely, fake news domains. Thus, we lack a measure of the broader fraction of content that is misinformation.

⁴⁷⁷ See Faris (2016).

speech. Some of the most commonly cited forms of harmful speech on online platforms are incivility, harassment, hate speech, and extremist content.

Incivility is difficult to define, but recent work has developed a nuanced approach which encompasses contempt (e.g., name calling), threats, partisan vitriol, profanity, speech devaluation (e.g., calling someone a liar), and seditious language (e.g., calling someone a traitor).⁴⁷⁸ Survey data suggests that social media users are regularly exposed to incivility online. Moreover, recent empirical work on Facebook and Twitter supports anecdotal evidence that social media platforms are rife with uncivil speech. Uncivil language is highly prevalent in political discourse, comprising a large proportion of comments on European and US politicians' Facebook pages and tweets.⁴⁷⁹ Although individuals in surveys report feeling fatigued and demoralized by exposure to uncivil messages online,⁴⁸⁰ such content receives high levels of engagement.⁴⁸¹ This may beget a cycle of incivility that discourages certain citizens from participating in political discourse.⁴⁸²

More severe than online incivility, online harassment is defined as unwanted contact that uses digital means to create an intimidating, annoying, frightening, or hostile environment for the target.⁴⁸³ This involves a range of tactics including doxing, revenge porn, and explicit threats that make individuals fear for their immediate safety. Women and minorities are often the targets of harassment on social media platforms. Research suggests that online harassment demoralizes its victims and emboldens perpetrators, encouraging them to repeat such behavior.⁴⁸⁴ While there is little empirical evidence of how widespread online harassment has become on social media,

⁴⁷⁸ Recent work on incivility described here includes Barberá et al. (2018), Timm and Barbara (2018), Suhay et al. (2018), Theocharis et al. (2016), and Kwan and Gruzd (2017).

⁴⁷⁹ In another interesting study, The Guardian policed more than 70 million comments it received on news reports over ten years, discovering that harassment targeted women and minorities. Out of the ten most abused writers, eight were women and two were black men. All ten least abused writers were white men—see <https://www.theguardian.com/technology/2016/apr/12/the-dark-side-of-guardian-comments>.

⁴⁸⁰ See Hinduja and Patchin (2007).

⁴⁸¹ See Klubicka and Fernandez (2018).

⁴⁸² See Henson et al. (2013).

⁴⁸³ For definitions of online harassment, see Lenhart et al. (2016).

⁴⁸⁴ See Hensen et al. (2013) and Hinduja and Patchin (2007).

racially motivated, misogynistic, and anti-Semitic harassment of well-known journalists and celebrities have amplified its visibility.⁴⁸⁵

Often overlapping with harassment, online hate speech is most commonly understood as bias-motivated, hostile, and malicious language targeted at a person or group because of their actual or perceived innate characteristics.⁴⁸⁶ While systematic empirical work exploring the prevalence of different types of online hate speech across social media platforms is quite rare, existing research suggests that—counter to popular journalistic narratives—it may represent only a fraction of a percentage point of overall posts on sites like Facebook and Twitter. However rare, online hate speech can have severe offline consequences. Survey data indicates that online hate speech negatively impacts the psychological well-being of individuals who are exposed to it, and can damage intergroup relations. Compelling causal empirical evidence also suggests that online hate speech can incite people to violence.⁴⁸⁷ Indeed it may be playing a particularly devastating role in fueling attacks on immigrants, refugees, and other vulnerable populations. Moreover, hate speech is frequently used by extremist groups—from white nationalists to members of the Islamic State—to propagate their messages and lure recruits.⁴⁸⁸ Thus, in addition to potentially inciting hate crimes and exacerbating intergroup tensions, online hate speech may also bolster extremists groups' ability to fundraise and expand their reach.

What are the mechanisms by which social media facilitates the spread of harmful speech? First, social media platforms provide social distance and anonymity (although these characteristics vary by platform).⁴⁸⁹ The absence of face-to-face contact and the anonymity of online communication enables people to engage in harassment or express uncivil, hateful, or extreme opinions without tangible consequences. This may lead perpetrators of harmful speech to believe that “normal” rules of social engagement do not apply online. This is true both of

⁴⁸⁵ For evidence of the targets and consequences of online harassment described here, see Munger (2017), Kennedy and Taylor (2010), Mantilla (2013), Hinduja and Patchin (2007), and Banks (2016).

⁴⁸⁶ For a discussion of defining hate speech see Cohen-Almagor (2011) and Sellars (2016). For work on the relative rarity of online hate speech see Gagliardone et al. (2016) and Siegel et al. (2019).

⁴⁸⁷ See Tynes (2008), Muller and Schwarz (2017), and Chan et al. (2015).

⁴⁸⁸ See Siegel and Tucker (2018) and Siegel et al. (n.d.).

⁴⁸⁹ For studies of the connection between anonymity and harmful speech, see: Cohen-Almagor (2017), Munger (2017), Citron (2014), Delgado and Stefancic (2014), and Postmes et al. (2001).

everyday social media users and trolls, who mask their identities and strategically engage in uncivil discourse, harassment, and hate speech in order to spark reactions and gain attention.

Second, the amplifying nature of social media platforms gives heightened visibility to harmful speech. As studies of uncivil speech demonstrate, although individuals report frustration with viewing such content online, it actually receives more engagement in political discourse than civil language.⁴⁹⁰ This heightened engagement means that uncivil content may then be amplified by the platforms' algorithms, though more research is needed to test this hypothesis systematically. Social media sites also enable hate groups and extremists to be more visible in the digital realm. Although platforms are now cracking down on their activities, hate groups and extremists have operated openly on mainstream social media sites. For example, Richard Spencer, who organized the "Unite the Right" alt-right Charlottesville rally, had over 75,000 followers and was verified by Twitter up until November, 2017.

Recognizing the power of mainstream social media platforms to magnify their voices, "digital mobs" have engaged in coordinated harassment campaigns, as well as anti-Semitic and racially motivated attacks on Twitter. These attacks, which often target well-known journalists or celebrities, may lead the followers of influential accounts to be incidentally exposed to hate speech or harassment. After gaining heightened online visibility, harmful speech often receives offline media coverage as well, and this hybrid media system draws even more attention to incivility, harassment, hate speech, and extremism, intensifying their negative effects.

Finally, the networked structure of social media means that individuals who seek out harmful online content may find themselves immersed in online echo chambers devoted to fringe ideologies where harmful speech is normalized and encouraged.⁴⁹¹ As platforms ban extremist and hateful content, one unintended consequence is that banned users may congregate on specialized platforms where their speech is permitted. In these silos of harmful speech, individuals may be radicalized and ultimately may be more likely to perpetrate acts of offline violence. Thus as governments and platforms work to regulate the most egregious forms of

⁴⁹⁰ For research on the amplifying nature of social media platforms for harmful speech, see Mariconti et al. (2018), Chess and Shaw (2015), and Kumar et al. (2018).

⁴⁹¹ For studies of how social media echo chambers foster harmful speech, see Daniels (2017), Magdy et al. (2016), Costello and Hawdon (2018), Marwick and Lewis (2017), and Zannettou et al. (2018).

harmful speech on mainstream platforms, this effort may push fringe users into more dangerous and destructive online spaces.

F. Summary

As we have reviewed, the rise of social media has impacted the political arena in many important ways. While it has multiplied the venues for political debates and empowered users, it has also enabled a range of negative behaviors that adversely impact our society. It is important to stress, however, that we currently have only a very limited understanding of all these phenomena. The opacity of digital platforms combined with the restricted access to proprietary and user data greatly impairs independent academic research on the topic. It is not coincidence that most research on hate speech, misinformation, and related topics relies primarily on Twitter data, the most publicly accessible data of any major social media platform.

Therefore, any attempt to properly understand how social media platforms are impacting the political environment necessarily requires better access to data—a point stressed below in the recommendations section.

Reforming the Social Media Platforms

As discussed earlier, the political consequences of social media bear some important similarities with classic economic externalities in that they are generated in large part as by-products of non-political users connecting on platforms that were designed with non-political aims. If one were to take the externality analogy seriously, it suggests a possible set of solutions. In the standard case of externalities generated by transaction costs or unclear property rights, a central authority can prohibit or tax behaviors with negative externalities. In principle, such authority might be exercised by government agencies or by collective agreements among industry actors to create a private regulator.

But we might also consider whether the political impacts of social media can be internalized by firms or disciplined by the market so that government regulation is unnecessary. In making such assessments, it is worth remembering that new communications technologies have often been as politically and socially disrupting as digital media. Consider the printing

press. It too increased velocity, virality, and homophily. Yet these features did not in themselves require central regulation (at least within democratic societies).

The key difference between social media and the printing press lies in the differences in economic organization. The productive scale of a single printing press is quite limited. Thus, competition among presses quickly emerged in ways that contributed to political pluralism and democracy as a by-product. But as demonstrated by the Market Structure Committee, social media platforms have a strong tendency toward monopolization of their specific markets. Thus, the political effects of social media will be determined by the policies of a few firms in a very uncompetitive environment. The current situation is parallel to one where the printing press technology was controlled by a single producer, who could turn off and on any printing press in the world. So the situation of social media platforms is much more analogous to that of broadcast television and radio. Early on it was recognized that the scarcity of the broadcast spectrum generated the need to regulate in ways that ensured political pluralism.⁴⁹² Thus, governments issued licenses, put restrictions on how those licenses could be used, and subjected broadcasters to regulatory oversight.⁴⁹³

Yet there are arguments that suggest that the problems related to social media, like the printing press, can be adequately addressed by self-regulation. It has been argued that individual are reasonably discriminating consumers of information;⁴⁹⁴ that the social mediation of content is not very consequential; and that virality is a very rare phenomenon that is most often driven by the mainstream media.⁴⁹⁵ It has also been argued that harassment and misinformation may be controllable because (1) they are driven by a few high volume (and thus detectable) bad apples,⁴⁹⁶ (2) some tweaks to platform affordances would provide more tools to individuals to evade trolls (e.g., by excluding them from discussion threads), and (3) the possible loss of user engagement provides ample incentive for the platforms to appropriately moderate content.

⁴⁹² This scarcity was a major point in the Supreme Court case *Red Lion Broadcasting Co. v. FCC*, 395 U.S. 367 (1969), which upheld the “fairness doctrine” for broadcasters. Thus, the lack of scarcity in the digital space may pose legal impediments to building upon the precedents of radio and television regulation.

⁴⁹³ Of course, some governments went much further into the realm of actual censorship, which is something we do not advocate.

⁴⁹⁴ See Pennycook and Rand (2019).

⁴⁹⁵ See Goel et al. (2015) and Riedl et al. (2018).

⁴⁹⁶ See Grinberg et al. (2019); and Guess, Nagler, and Tucker (2019).

Finally, some scholars argue in favor of a flexible self-regulatory approach that is based on articulation of general principles and good practices and that involves third party certification and management-based regulation.⁴⁹⁷

We, however, have three concerns about arguments in favor of self-regulation and market discipline. First, as the Market Structure Committee points out, the major social media platforms operate in very concentrated markets. There are few good substitutes for Facebook social networking or Google searching. Thus, users who disapprove of the platform's speech policies or negative political impacts have limited alternatives. It is doubtful that the number of users who might stop using Facebook or Google would create a large enough financial incentive for the platforms to take the costly actions required to stem the departures. Moreover, it is unlikely that new platforms can emerge that offer true alternatives on issues like privacy and speech moderation. The lack of competition also means that the platforms may be less concerned with maintaining brand images in the way that, for example, Coke, Pepsi, or Nike might be. Such consumer-oriented firms in competitive environments are often quite responsive to social concerns about their corporate practices for fear of lost sales and boycotts. But boycotts work best when there are alternatives for consumers, and in social media there are few.⁴⁹⁸

The second concern is that policies and practices of social media platforms are not very transparent. And without adequate independent research, the implications of these policies are not known. Thus, it is beyond the capacity not only of regular consumers but also of watchdog groups to verify whether the social media firms are adhering to their states' self-regulatory aims. This problem is exacerbated by the fact that many political communications are targeted. Thus, most users may never see the politically harmful content.⁴⁹⁹ While the social media platforms may attempt to create oversight boards, the independence of those boards may be questioned. Moreover, the access of such boards to data, algorithms, and policies may be limited by the firms. So the self-regulated can quite easily become self-assessors of the self-regulation. Thus, an important role for government may also lie in creating the incentives and infrastructure for the research that would hold platform firms accountable for any political harms caused by their

⁴⁹⁷ See Coglianese and Lazer (2003).

⁴⁹⁸ See Baron (2009, 2014).

⁴⁹⁹ This problem may be exacerbated in the future if Facebook follows through with its plans to stress encrypted messaging. <https://promarket.org/we-need-to-deal-with-whatsapp-misinformation-problem/>

algorithms or policies, which in turn means ensuring access to the data necessary to undertake this research in the first place.⁵⁰⁰ Finally, the experiences of self-regulatory regimes in other industries suggest that such an approach absent government oversight may come up short. For example, the fact that the US financial services industry is self-regulated through the Financial Industry Regulatory Authority (FINRA) does not obviate the need for supervision by the Federal Reserve, the Securities and Exchange Commission, and other governmental agencies.⁵⁰¹

To illustrate our concerns with the inadequacy of self-regulation, an example of successful self-regulation is instructive. Consider the case of e-mail, which was threatened with extinction by the flood of spam just a few years ago. While some reformers proposed centralized solutions such as small congestion charges on e-mail, spam was largely solved through detection technologies that firms willingly installed. Notably, since much of spam is misinformation, and e-mail a particular kind of social medium, this could be viewed as a major triumph of social media companies against misinformation.

But the incentives to engage in spam detection are different than those facing the social media platforms in a number of important ways. First, the number of users affected by spam (i.e., all of them) was greater than the number of consumers objecting to certain political content and policies. Second, at the time, there was significant competition in internet service and email provision. Consumers fed up with spam would have both strong incentives and opportunities to switch to services that better controlled spam. Third, spam filtration was not as politically contested as content moderation. E-mail services could more easily design transparent policies that were broadly acceptable to everyone but the spammers. Content moderation on what has become an essential piece of political infrastructure is not likely to ever achieve such a consensus. Democratic mechanisms may be required to reach acceptable policies. In summary, it is doubtful that controlling political externalities such as misinformation and harassment will be in the business interests of the firms the same way controlling spam was.

There are of course important caveats surrounding government regulation of the platforms. One of the most important is the possibility that regulatory agencies could be

⁵⁰⁰ As will be better outlined below, it is important to stress that at this point the committee takes no position on the platforms' immunities from Section 230 of the Communications Decency Act of 1996.

⁵⁰¹ Also see the Privacy Committee's discussion of the insufficiency of self-regulation in the area of user privacy.

“captured” by the social media platforms so that new rules and regulations would benefit the incumbent firms and further stymie entry into the market.⁵⁰² This concern is especially important in the long run after the salience of social media regulation in current political and policy discussions has died down. Long after public scrutiny has abated, the social media platforms will be allocating significant resources to lobbying regulators and legislators for favorable treatment. A standard prescription for preventing agency capture is to strengthen executive branch oversight of the agency, under the theory that presidents are less parochial than legislators.⁵⁰³ This prescription may be dangerous in the case of social media regulation. As we discussed above, social media platforms have become an essential part of the infrastructure for US elections. Giving any partisan officials wide latitude over social media regulation invites opportunities for political manipulation. In other words, we must be equally vigilant that any social media regulators are not captured by industry or a political party.

Policy Recommendations

The previous sections have raised a number of challenges posed by the rise of social media platforms as important political actors and as the new public squares. In this section, we outline a number of proposals designed to mitigate not only the potential negative political impact of social media but also the auxiliary effects of concentration on the digital platforms. Perhaps the most important recommendation is our concurrence with the other groups that significant government regulation and greater antitrust scrutiny is warranted. Our unique contribution to that discussion focuses on issues related to regulatory structure, where we lay out several principles to help insulate regulatory authorities from excessive industry and political influence while aiming to preserve democratic accountability. We also address important issues related to disclosure and transparency. First, we endorse updating campaign finance law to cover spending on social media campaigns. Second, we call for more transparency in the use of platform companies’ support for research on social media and for greater dissemination of internal research. Finally, we believe that a new Digital Authority can be essential for

⁵⁰² See Carpenter and Moss (2013) for a recent set of essays on the politics of agency capture.

⁵⁰³ For example, see Moe (1989).

facilitating greater independent research and auditing, by ensuring that regulators and scholars have access to relevant social media data.

A. New Regulatory Authorities

Our report, along with those of the committees on privacy, media, and competition, identifies many negative features of social media that will require much more than commitments to corporate social responsibility and self-regulation. Yet, recent regulatory arrangements are under-resourced and ill-suited. Thus, we recommend the creation of a new regulator and enhanced arrangements for inter-agency cooperation.

We believe the following principles should guide the creation of a Digital Authority (DA) tasked with regulating digital platforms:⁵⁰⁴

1. The DA should have a reasonable degree of autonomy from industry influences to make decisions about social media platforms in the public interest.
2. The DA should be structured so as to prevent its powers from being used for partisan purposes.
3. The jurisdiction of the DA should cover as many social media–related functions as possible to prevent regulatory fragmentation.
4. Mechanisms for coordination with other agencies should be created.
5. The DA should have responsibility for rulemaking in the following areas:
 - i. General consumer protection
 - ii. Privacy policies and disclosure
 - iii. Transparency
 - iv. Data portability
 - v. Data and algorithmic access for external auditing and research

⁵⁰⁴ Should existing bodies such as the FTC or FCC be granted new authority, we hope that these principles will guide the ways in which these powers are implemented internally.

6. The DA should have authority to create mechanisms for real time data collection from the platforms (subject to appropriate protections for user privacy).
7. The DA should have research capacity to undertake studies of the impact of the platforms on social and political outcomes.
8. The DA should play a facilitating role in generating independent research by outside scholars.
9. The DA should have the authority to review relevant internal studies conducted by the platform companies. When the release poses no undue compromising of privacy or exposure of business secrets, studies should be made publicly available.
10. The DA should evaluate any of its rules and regulations for adverse effects on innovation and competition in the digital media sphere.
11. The DA should have a broad set of tools to enforce compliance with its rules and directives. The authority should be able to pursue administrative actions, civil proceedings, and criminal referrals.

The principles of autonomy (items 1 and 2) will be essential in preventing industry capture and politically-inspired forbearance. We are especially concerned about the latter consideration, given the role that social media has come to play in political campaigns and policy advocacy. Adhering to both principles, however, will be challenging. As noted above, some scholars tend to stress a combination of executive oversight and legislative insulation as a preventive of industry capture.⁵⁰⁵ As an example of this approach, consider the Consumer Financial Protection Bureau (CFPB) that was created in the Dodd-Frank financial reforms. The CFPB's design includes a number of features that were designed to limit undue industry influence. First, it is headed by a single presidentially-appointed, Senate-confirmed appointee. This head serves for a fixed five-year term and can only be removed by the President for cause. Second, the agency is housed in and funded by the Federal Reserve. Thus, it is not subject to direct presidential oversight and is insulated from the appropriations process. Third, the CFPB

⁵⁰⁵ See Moe (1989).

can issue new rules without going through the review of the Office of Information and Regulatory Affairs (OIRA).⁵⁰⁶

There are two clear downsides of the CFPB model. The first is that it goes too far in insulating the DA from democratic accountability, as voters would lack clear means of having their views represented in agency decision making. Second, and perhaps more importantly, the President would have important sway over the DA through the powers of appointment.⁵⁰⁷ Such an arrangement could be an open invitation to politicization and partisan favoritism—especially in light of the role of digital media in political debates and elections.

An alternative to the highly-insulated CFPB model would be the traditional multimember independent commission similar to the Federal Trade Commission or the Securities and Exchange Commission. Such commissions are generally bipartisan with limits on the number of members from a single party. If they include staggered terms and constraints on presidential removal, presidential influence may be restrained. But precluding partisan influences on a commission remains challenging. With an odd-numbered membership, it is likely that one party would have a majority. But with an even-numbered membership with partisan balance (such as the Federal Election Commission), deadlocked decisions might become the norm. Within the commission model, varying degrees of insulation could be achieved by allowing the DA to obtain some of its funding from industry fees in addition to congressional appropriations. Currently, independent commissions are not subject to OIRA review and therefore are provided some buffer against excessive presidential influence.

There is probably no single way to balance the concerns for preventing industry capture and those related to political manipulation. But we envision that the DA could draw from both of these regulatory models to strike a reasonable balance. Ultimately, the recurrent and

⁵⁰⁶ Under Executive Order 12866, OIRA is tasked with reviewing all “significant regulatory actions,” including proposed rules that would “[h]ave an annual effect on the economy of \$100 million or more.” OIRA’s primary charge (to the extent allowed by applicable statutes in any given case) is to make sure that the rules meet a cost-benefit test. Currently, all independent regulatory agencies are exempted from OIRA review. But the current administration may be considering OIRA review for independent rulemaking.

<https://www.forbes.com/sites/susandudley/2018/04/24/trumps-regulatory-czar-signals-readiness-to-rein-in-independent-agencies/#173735f9f9c3>

⁵⁰⁷ This presidential influence has been demonstrated at the CFPB following President Trump’s appointment of Mick Mulvaney.

pervasive problems around social media indicate that, at least in this area, imperfect regulation is likely to be better than the current status quo of no regulation whatsoever.

The issue of regulatory fragmentation (principle 3) is a serious one that has bedeviled many policy areas, and especially financial regulation. Thus, we strongly urge that the DA be given as much exclusive jurisdiction over digital media platforms' publication and data processing activities as practical. However, there would appear to be at least three important exceptions.⁵⁰⁸ Clearly, the role of social media in campaigns and elections will continue to be overseen by the Federal Elections Commission, as such regulations need to be well coordinated with regulations related to the use of television and radio. Second, there may be policy issues related to the Internet, such as service provision, that might be better regulated by the Federal Communications Commission.⁵⁰⁹ Finally, the Federal Trade Commission and the Department of Justice will continue to be the agencies that review mergers and enforce antitrust laws. Thus, we stress the importance of principle 3 and the idea of mechanisms to coordinate between the DA, the FCC, the FEC, the FTC, and the DOJ.

A related concern is the extent to which regulations of the DA should preempt state laws and regulations. While there may be some advantages to replacing a patchwork of state rules with a uniform national standard, the committee believes that the DA should primarily set regulatory floors in areas such as privacy and electioneering disclosure and allow states to enforce high standards should they choose to do so. This is particularly important in the area of campaign disclosure, given the states' traditional role in regulating elections.⁵¹⁰

We recommend that the DA take primary responsibility for several areas, including consumer protection, privacy, transparency, and data portability.⁵¹¹ Specific recommendations for rules and enforcement approaches are detailed later in this report and in the reports of the other working groups. But to facilitate rules and enforcement in these areas, we believe that it is

⁵⁰⁸ Of course, there will be many other exceptions generally related to regulations and oversight that apply to all firms, such as the enforcement of employment laws.

⁵⁰⁹ The committee takes no position on the ongoing controversy over whether the FCC has or should have the power to regulate service providers in areas such as net neutrality.

⁵¹⁰ However, there might be an important role for centralized policies to lessen the burdens associated with social media platforms making 51 sets of disclosures (one set for the FEC and one for every state regulator).

⁵¹¹ Here we depart from the Privacy Committee, who would leave user privacy regulation under the domain of the Federal Trade Commission.

important for the DA to develop rules and capacity for real-time data collection and analysis. Access to these data will be crucial for ensuring compliance with privacy and transparency rules. Consider the recent mosque attacks in New Zealand that were broadcast on Facebook. Without real-time regulatory data, Facebook's reports on moderator response time and the number of viewers is not verifiable. Regulators ought to have access to such data (subject to reasonable protections against revealing firms' proprietary information, and protecting the privacy interests of third parties.) Where privacy and the maintenance of trade secrets can be insured, such data should be made available to independent researchers.

We have previously detailed how the lack of access to data prevented scholars from comprehending the real impact of social media on most arenas. This claim, however, is only partly true—it prevents scholars from doing *independent*, high-level research. Scholars working for the platforms themselves undoubtedly have much better access to data than those not working for the platforms (i.e., those employed at institutions that typically support scholarly research, such as universities and think tanks). The question, however, is: What happens to the studies that are conducted internally by platform-employed researchers? These proprietary studies are rarely made available to the public, and as long as this remains the case there will always be those who suspect that those that are made public will be the ones that tend to cast the platforms in a favorable light, *regardless of whether or not this is the case*.⁵¹² Therefore, we recommend that the DA create a research office staffed with social and data scientists who can conduct and coordinate independent research (principle 7). Moreover, we also believe that the DA research office should play a facilitation role for research conducted by outside scholars by providing both funding and access to data and data generation processes (principle 8). This would allow for far less financial and data dependency on the platforms by scholars and outside analysts.⁵¹³ Finally, the DA should have access to the internal studies conducted by the platform companies, diminishing information asymmetries between the parties (principle 9). The DA should consider creating pre-registration of internal studies to avoid the hiding of adverse outcomes. In addition,

⁵¹² This is a particularly perverse version of what is known as the “file drawer” problem in academia, which is generally described as a biased view of research results produced by the fact that “positive” findings are more likely to get published than null findings (Franco, Malhotra, and Simonovits 2014). So rather than the proclivities of academic publishing preferences, here the file drawer problem could be attributed to the financial incentives of the platforms.

⁵¹³ This is a problem that has plagued health-related research in the past; see for example Fabbri et al. (2018); Lexchin (2012); and Sismondo (2007).

an important consideration is how to set the agenda of the DA's research arm. We recommend the creation of an outside advisory committee to help set these priorities.⁵¹⁴ Such an advisory committee should have a membership in which academia and civil society is heavily represented relative to industry.

Principle 10 is crucial. Too often regulatory regimes create substantial burdens on smaller firms that serve as a barrier to entry. Competition in the social media space is an important complement to good regulation and should be promoted where possible. The Market Structure Committee Report has a series of important recommendations on policies that may promote a competitive environment in digital markets.

Finally, as many commentators have raised concerns that recent FTC fines of Facebook over privacy violations were not large enough to deter future transgressions, it is crucial that the DA should have a wide variety of enforcement tools to ensure that punishments for violating its rules and directives are significant enough to compel compliance by the platforms.⁵¹⁵

B. Antitrust Enforcement or Other Policies to Prevent Political Market Concentration

As seen above, many of the negative political by-products identified by this report are associated with the lack of a competitive market for digital platforms. Therefore, we must also consider the establishment of some form of policy aimed at reducing “political concentration.”

Antitrust is the primary policy targeting concentration. Contemporary antitrust enforcement is generally predicated on a consumer welfare standard. Accordingly, unless it can be shown that consumers are harmed by a potential increase in price or decrease in quality, mergers and acquisitions are generally approved by regulators. As currently interpreted, this approach presumes that the main negative externality of economic concentration is monopoly

⁵¹⁴ Although the appointments would be at the discretion of the agency head, nominations can also be made by members of Congress, the general public, or professional societies or current and former committee members. Moreover, the charter for such a committee can specify the professional and academic requirements of the body. See <https://www.gsa.gov/policy-regulations/policy/federal-advisory-committee-management/advice-and-guidance/the-federal-advisory-committee-act-faca-brochure> for the regulations surrounding the creation and operation of such an advisory committee.

⁵¹⁵ See https://www.washingtonpost.com/technology/2019/07/12/ftc-votes-approve-billion-settlement-with-facebook-privacy-probe/?utm_term=.9c2279a119c9 and https://www.washingtonpost.com/technology/2019/07/22/facebook-vs-feds-inside-story-multi-billion-dollar-tech-giants-privacy-war-with-washington/?utm_term=.434f89b36440.

pricing. In the case of social media, this presumption is dubious. Economic concentration concentrates political power. Large firms who lack competitors are hugely advantaged in the political marketplace. They have the clout to use political processes to thwart the innovation which might lead to future competition and generate other political rents. Second, concentration may exacerbate the negative consequences of social media's role in the political system. When the social media space becomes concentrated, the effect of each firm's moderation decisions is increased. Moreover, the lack of competition deprives us of a marketplace of ideas that might serve to regulate the platform's policies on speech and political activity. These political effects of concentration are unlikely to ever be captured directly by the consumer welfare standard. A more promising approach would be to emulate the various rules that the FCC has enacted over the years to limit concentration of media ownership in particular markets.

Whether the antitrust law should broaden its scope beyond the consumer welfare standard is a complex and controversial issue and lies beyond the scope of our committee. But we do believe that the harms to citizens through the distortion of political processes should be given considerable weight in policies aimed at fighting market concentration. At a minimum, we believe that the DA should develop methodologies for evaluating the explicit political impact of social media concentration, and we welcome a debate on how to better structure such review in a way that is not manipulated by political parties (the metric developed by the News Media Report being an important reference). The independent access to data and data generation processes, for which we have called in principle 8 in the previous section, will undoubtedly play an important role in developing and evaluating such methodologies. Such studies should be made available to antitrust regulators and to the general public, so as to instruct decision-making. Such reports may contribute to the establishment system of dual review such as that in place for mergers involving broadcasters, which gives the FCC a dual-mandate that complements that of antitrust authorities but considers different criteria when assessing the consequences of concentration.⁵¹⁶ Again, the challenge here will be to develop a system that controls political concentration in an objective, unbiased way.

⁵¹⁶ This is also the case with banking mergers where the impact on financial stability is considered along with the impact on concentration.

C. Platform Liability

Many discussions about the role of social media firms in content moderation focus on the impact of Section 230 of the Communications Decency Act of 1996, which says, “No provider or user of an interactive computer service shall be treated as the publisher or speaker of any information provided by another information content provider.” Moreover, Section 230 preempts state laws, and the courts have refused to limit it to Internet service providers. Among other things, this provision immunizes social media platforms from legal liability for hate speech, harassment, and misinformation.

The Politics Committee reached no consensus on the desirability of repealing Section 230. On one hand, removing the liability protection of the media platforms would undoubtedly spur them to undertake much more aggressive content moderation in an effort to avoid litigation related to slanderous and harassing speech. Yet at the same time, the absence of liability protection might induce the platforms to police over-aggressively and have an unduly chilling effect on speech.⁵¹⁷ The debate surrounding Germany’s Network Enforcement Law (NetzDG) exemplifies these tradeoffs. In accordance with its long tradition of prohibiting and penalizing extremist speech, the NetzGD law requires the platforms to delete illegal content within 24 hours or face fines up to \$60 million. The law spurred digit platforms to dramatically increase moderation efforts by hiring thousands of new moderators. Yet, critics accuse the law of overreaching—resulting in the deletion of legitimate content and serving as an additional grievance in opposition to which extremist movements could mobilize. While the First Amendment would surely mean that any similar law in the United States would be far narrower, the tradeoffs between over- and under-regulating content would remain. There was some support on the committee for removing liability protections in limited circumstances. One exception might include defamatory or illegal communications by foreign actors. Assuming adequate oversight capacity, removing protection from liability on those acts would incentivize the platforms to better police foreign disinformation efforts.⁵¹⁸ The recommendation of the Media Committee to condition Section 230 protections on public service obligations is an idea

⁵¹⁷ Interestingly enough, and as reported by the News Media Committee, Section 230 was actually initially enacted to allow platforms to moderate content in the first place, even if it became a shield afterwards.

⁵¹⁸ As discussed in the section on campaign disclosure, US courts are generally more tolerant of policies directed at reducing foreign political influences.

worthy of consideration. But given the history of similar arrangements, such as the Fairness Doctrine for broadcast licensees, we should remain concerned that platform lobbyists may be able to narrow the scope of these obligations while retaining immunity.

D. Role of Social Media in Campaigns and Elections

We endorse two campaign disclosure provisions that have been proposed as part of the Honest Ads Act.⁵¹⁹ The first of these amends the definition of “electioneering communication” to include internet or digital communication. Under current law, electioneering communication is any *broadcast, cable, or satellite* communication that refers to a clearly identified candidate for federal office and is publicly distributed within 30 days of a primary election or 60 days of a general election. Electioneering communication must include a disclaimer stating who paid for the message. Any entity spending more than \$10,000 on electioneering communication in a calendar year must also report the disbursements to the FEC. Section 6 of the Honest Ads Act broadens the definition of electioneering communication to cover “qualified internet or digital communication.” This term, in turn, is defined as “communication which is placed or promoted for a fee on an online platform.” Section 6 thus makes applicable to political advertisements on digital platforms the disclaimer and disclosure requirements that already govern political advertisements in other media. It fills the loophole in the existing definition of electioneering communication (which was adopted by Congress in 2002, before the emergence of digital platforms). The committee believes that the Act’s definition of electioneering communication be clarified to ensure that it includes electioneering activities of “coordinated social media campaigns” that do not involve paid advertising on the platform. This would encompass situations such as those where an entity is coordinating the actions of actual or simulated user accounts. In such cases, the organizer of the campaign would have to disclose any electioneering activities to the platform.⁵²⁰

⁵¹⁹ The Honest Ads Act was introduced in the Senate with bipartisan support. Subsequently, its provisions were included as part of H.R. 1 (the For the People Act) that the House of Representatives passed in March 2019. That act includes many other provisions related to electoral reform that sharply divide the two political parties. Thus, we make no endorsement specifically to H.R. 1.

⁵²⁰ See <https://defusingdis.info/2019/01/17/how-transparency-can-help-defuse-disinformation-from-botnets-sockpuppets-and-online-trolls/>. Any disclosure rules for social media coordination should consider a reporting threshold that excludes small organized campaigns, such as student groups in high schools or on college campuses.

The second disclosure provision we endorse is a mandate that digital platforms compile publicly available databases of political advertisements that are run on the platforms. This would include the posts associated with the coordinated social media campaigns discussed above. Section 8 of the Honest Ads Act states that digital platforms with over fifty million monthly users must maintain a record for each request to purchase a political advertisement made by a person whose aggregate purchases exceed \$500 per year. The act should be clarified to include electioneering activities of coordinated social media campaigns that do not involve paid advertising. Each such record must contain, inter alia, a copy of the advertisement, the number of views generated by the advertisement disaggregated by demographic categories, the period over which the advertisement was displayed, the rate charged for the advertisement, the demographic targets, and information about the buyer of the advertisement. Section 8 further defines a political advertisement as one that is made on behalf of a candidate or that “communicates a message relating to any political matter of national importance,” including a candidate, a federal election, or a national legislative issue.

As far as these provisions go, they do not encompass all political activity on digital platforms that we might like to be disclosed. Even as amended by Section 6, electioneering communication does not cover messages that do *not* refer to clearly identified federal candidates or that are distributed more than 30 or 60 days before an election. Those who convey such messages on digital platforms are thus exempt from any disclaimer or disclosure requirements. Broader requirements for non-electioneering political ads, however, may generate First Amendment concerns. In light of those considerations, we endorse a disclosure requirement on political advertising paid for by foreign entities. Such a foreign distinction is found in many areas of law, such as campaign finance and lobbying disclosure.⁵²¹

Similarly, Section 8’s record-keeping obligation does not apply to digital platforms with fewer than fifty million monthly users (which include Reddit, Snapchat, and Tumblr). Section 8 also contains no mechanism to compel buyers of political advertisements to truthfully reveal their identities (or to prevent them from dividing a larger volume of purchases into several sub-

⁵²¹ Foreign nationals may not make direct contributions to candidates and parties or make independent electioneering expenditures. These restrictions were recently upheld unanimously in *Bluman v. FEC* (2012). The Foreign Agents Registration Act requires periodic disclosures from agents working for foreign principals on political or quasi-political matters.

\$500 accounts). The records compiled by digital platforms would be more informative if these issues were addressed.

Furthermore, nothing in the Honest Ads Act responds to the possibility that digital platforms themselves might not be politically neutral—a concern outlined in the first section of this report. For example, platforms might charge lower advertisement rates to certain candidates or parties, restrict the audience for certain candidates’ or parties’ advertisements, offer sales and consulting services at lower prices (or for free) to certain candidates or parties, or even ban certain candidates or parties from the sites.⁵²² Ideally, we believe that social media companies should not engage in any practices which violate strict political neutrality. As mentioned, however, there are concerns that the regulation or imposition of political neutrality by the DA might impinge upon First Amendment protections. So at a minimum, we support further and strong disclosure requirements that would reveal such non-neutral platform policies. Such disclosures should cover situations (i) when the platforms provide specific support or technical assistance to political parties, candidates, or interest advocacy groups, outlining what type of support has been provided and the outcome of this support; and (ii) when the platforms make algorithmic changes that directly impact how users see political content, outlining the outcome of these changes. The platforms should be obliged to aggregate the data compiled under Section 8 at least by candidate, party, and issue. These data slices would show whether advertisement characteristics (such as the rates charged and the views generated) systematically varied along these dimensions. Perhaps more onerously, platforms could be required to disclose marketing and pricing materials. These documents would expose any platform policies that overtly favored or disfavored particular candidates or parties.

These obligations would be complemented by the data access policies for academic and independent research outlined below, which will enable third parties to more effectively police platform behavior that may adversely impact the political arena.

E. Philanthropic Disclosure

As do many other large corporations, large digital firms invest in extensive philanthropic efforts. Such support obviously serves many worthwhile causes. Yet there is considerable

⁵²² See Zittrain (2014).

evidence that philanthropic efforts are often designed to support political strategies, generating the goodwill either of important civil society organizations or of legislators with ties to those organizations.⁵²³ But more troubling are the conflicts of interest associated with educational philanthropy to support research and teaching on technology and the associated policy issues. Such support may make it more difficult for technology and policy scholars to criticize the platforms and their social impact.

We offer two concrete proposals. First, we believe there should be greater and more transparent disclosure of the philanthropic efforts of social media companies, especially those tangibly related to teaching and research. The charitable arms of Google and Facebook make periodic disclosures of their contributions. But there should be similar disclosures of the contributions of top executives.⁵²⁴ Second, as discussed in our DA proposal, the DA should create an office of research that will help facilitate independent research on social media companies and platforms. Such research will avoid the inherent conflicts of industry-supported research. We also encourage universities and academic associations to develop disclosure standards that would apply to scholars supported by social media firms. The disclosure policy of the American Economic Association is one plausible model.⁵²⁵

F. Data Access for Academic and Independent Research

As discussed above, we believe that independent research on the economic and political effects of social media is crucial to ensuring that the platforms enhance citizen well-being. Currently, the major impediment to such research is data access. Such access is limited for a variety of reasons:

- Data security and individual privacy concerns may be raised.
- Data may be proprietary.

⁵²³ See Bertrand et al. (2018). In the case of the Chan Zuckerberg Initiative, philanthropy can explicitly serve political ends. Since the initiative is organized as a limited liability corporation rather than as a 501 c3 nonprofit, it may make political contributions.

⁵²⁴ The controversy with the New America Foundation and its Open Markets initiative centered around contributions from Eric Schmidt rather than from Google's charitable foundation.

⁵²⁵ See <https://www.aeaweb.org/journals/policies/disclosure-policy>.

- Even when data is in the public domain, platforms may limit access to data through terms of service agreements.
- Securing access to data is an uncertain and time-consuming process; when possible at all, storing and accessing such large datasets is technologically challenging and financially demanding.

Moreover, there is also a political challenge to securing access to social media data for rigorous, transparent research. After the Cambridge Analytica controversy, there have been well-organized—and well-meaning—organizations and individuals attempting to restrict access to social media data for research purposes in the name of protecting individual privacy. One such example is the request by the Electronic Privacy Information Center (EPIC) to US Federal Trade Commission and the European Data Protection Board to suspend data sharing between Facebook and Social Science One, a new organization attempting to facilitate analysis of Facebook data by academic researchers to study the effect of social media on democracy and election integrity.⁵²⁶

The lack of access to data for academic researchers does not, of course, mean that no research is being conducted. Instead, it means that the only people who are able to conduct such research are those working inside the platforms.⁵²⁷ Put another way, the knowledge that can be gained from access to these data is going to remain inside what are giant multinational companies, and most likely the vast majority of this research is going to be focused on increasing profits for those same companies, as opposed to a myriad of other socially desirable applications that research could address, including assessing many of the political questions listed earlier in this report.

To remedy these problems, we offer two proposals. First, as discussed above, a major initiative of a new DA ought to be to facilitate independent research. This could include making the data it obtains from social media firms available for research (with suitable restrictions for individual privacy and proprietary secrets). Second, we encourage the reconsideration of the presumption that data collected by social media platforms *ought to be considered proprietary at*

⁵²⁶ See <https://epic.org/2018/07/epic-asks-ftc-and-edpb-to-sus.html> and <https://socialscience.one/>.

⁵²⁷ And perhaps security services that manage to get hold of data through non-transparent means.

all.⁵²⁸ There are a number of possible proposals in this regard. The strongest would recast the role of the platforms not as owners of the data provided by users, but rather as stewards of that data, entitled to use it to improve their own business models but not necessarily to prevent others from using the data for welfare-maximizing purposes.⁵²⁹ Another option would be analogous to a “data tax,” whereby X% of data collected by digital platforms would need to be deposited in the public domain. Such a proposal might also mitigate concerns about excessive economic concentration by allowing start-up firms to use such data in developing algorithms.⁵³⁰ In implementing any of these proposals, systems to ensure individual privacy will need to be put in place. The Privacy and Data Protection Committee’s report provides some interesting ways to enable data disclosure in a privacy protective way. We would only stress that while differential privacy protocols may be required, we hope that the value of such protocols is weighed against the possibility that data loses its value for academic work and facilitating innovation.

Further, it is important to note that understanding the impact of social media usage upon political outcomes does not depend solely on access to data in its raw form. A full understanding of these processes and causal impacts of social media requires an understanding of the *data generation process*. While social media platforms all provide some publicly-available information about the ways in which their products function, there is often a very wide divide between what is contained in this public information and what is necessary for analysts to understand the data generation process sufficiently. Here, we highlight two important categories of information: data that was generated as the result of *internal experimental research* and the *algorithms* that determine what content is displayed to users.

It is well known that platforms such as Facebook are constantly conducting A/B tests for all sorts of product development reasons,⁵³¹ and there is the possibility that internal researchers

⁵²⁸ The feasibility of this approach would depend significantly upon whether current law actually does give platforms the rights of legal ownership over data that they currently hold. If current law does actually make platforms the “owners” of data they have collected, then governmental attempts to eliminate their ownership rights could raise significant constitutional questions under the Fifth Amendment’s right against the deprivation of property for public use without just compensation.

⁵²⁹ See, e.g., Balkin and Zittrain (2016), Balkin (2016), and Zittrain (2018).

⁵³⁰ See Graef and Prufer (2018).

⁵³¹ See, e.g., Feitelson, Frachtenberg and Beck (2013); <https://www.xda-developers.com/facebook-mute-notification-dots/>.

may be conducting experiments for basic scientific research as well.⁵³² As currently construed, the latter set of studies are only released into the public domain if they are published, and the former set of studies are almost never made public. This represents a lost gain in efficiency for the accumulation of scientific knowledge, but it also provides potentially more serious problems for independent analysts if *they don't know whether the data was generated as a result of an experimental manipulation*. In an ideal world, platforms would maintain a public record of all experiments conducted internally, and then social media data that is made available for outside analysis—through any of the mechanisms suggested above—would contain metadata that identifies whether the unit of analysis in question had been subjected to experimental manipulation. In practice, due to the role A/B testing plays in product development, something short of this ideal would probably end up being more pragmatic, but we believe it is important when pursuing principle 8—ensuring transparency—that the question of labeling data as to whether they have been included in an experimental study is addressed. Further, as previously noted, we believe that the more of these internal experimental studies that can be released to the public, the better.

Additionally, an increasingly pressing question for assessing the impact of social media on the political sphere concerns not only the content of social media posts but the *algorithm* that determines what content is seen by which users. Thus, a thorough and accurate assessment of social media's impact in the political sphere would involve an understanding of how these algorithms function. Complicating this task is the fact that these algorithms are constantly being tweaked and adapted. In an ideal world for outside assessment, platforms would not only publish all of the algorithms that have been used to serve up social media content, but would also include metadata in social media posts identifying which algorithm was at use when the content was posted. However, as these algorithms are central to the platforms' business models and competitive advantages—and very technically complex—the likelihood or desirability of regulation to compel the publication of algorithms would seem to be limited.

The alternative, however, is to allow what is known as *algorithmic auditing*, whereby outside analysts attempt to assess the function of the algorithm by either analyzing code directly

⁵³² See, e.g., Bond et al. (2012).

(code auditing) or by interacting with the system in question (e.g., sending a lot of requests to Google and seeing what Google returns). There are various different forms of these types of audits depending on the degree of adherence to the platform's terms of service and the personalization of the interaction with the platform.⁵³³ Almost all of these types of audits, however, face hurdles in terms of gaining access to the necessary information from the platform in question. Such obstacles can arise because of API rate limits (e.g., a researcher wants to give Google 100,000 inquiries a day, but is only permitted to make a small fraction of those inquiries from one IP address) and also because the necessary data may not be included in the API (e.g., YouTube video recommendations). Auditors often get creative to get around these limits, but this can make the enterprise much more time-consuming (i.e., higher barriers to entry for prospective studies), subject to being shut down by the platforms for terms of service violations, and perhaps subject to legal action. At the same time, audits that are conducted in cooperation with the platforms are potentially subject to many of the concerns about conflict of interest that we addressed previously.

Thus, we recommend a policy of *cooperation without supervision* to facilitate algorithmic auditing for regulatory or scholarly research. Under such a policy, platforms would facilitate algorithmic auditing by exempting auditors from normal terms of service, but would not in any way supervise the auditing process. So, for example, Google could exempt auditors from rate limits in terms of the number of times an API could be accessed in a given day for the period of time of the audit. To prevent abuse of the potential access, a non-partisan academic advisory board—perhaps in the model of *Social Science One*—could vet research proposals for external audits of the platforms.⁵³⁴ The DA may also consider creating certain legal safe harbors for researchers conducting audits.

Conclusions

While research on the political impacts of social media and its major platforms is nascent, the potential for irreparable harm to democratic institutions is more than necessary to justify

⁵³³ For more details on algorithmic auditing, see Sandvig et al. (2014) and Wilson (2018).

⁵³⁴ See <https://socialscience.one/>.

significant reforms. While it is possible that the incentives of platforms and users may resolve some of the negative political effects of social media, we believe that some coordinated and centralized actions may be necessary. Thus, we join those other working groups who are calling for the creation of new regulatory authorities focused on social media. But given the power and clout of social media companies, these authorities must be well designed to balance insulation from industry capture with democratic accountability. In our report, we outline several principles of regulatory design which may serve these purposes.

A new regulatory regime also should do much more to promote disclosure and transparency to a greater extent while balancing legitimate privacy and First Amendment concerns. At the barest minimum, political advertising on social media must be subject to the same disclosure and reporting required of television and radio advertising. Such disclosures serve both to allow citizens to understand who funds parties and candidates as well as to provide a safeguard against foreign interference in elections. But much more fundamental transparency of data and algorithms will also be necessary to allow for government regulators and independent researchers to better evaluate the effects of social media platforms on elections, polarization, and group conflict. Without such efforts, social media companies will remain unaccountable for the political and social harms they may inflict.

Our group concurs with the emerging consensus that excessive concentration and monopolization by the platforms is a root cause of many of the negative consequences of social media. While many previous technological innovations also impacted political systems, most did so in ways that deconcentrated power. Those with the potential to concentrate power, such as broadcast television, came quickly under government regulation and oversight. But social media platforms have the capacity to concentrate control over valuable political resources in ways previously unimaginable. Antitrust or other complementary market concentration control policies must develop the tools to consider the harms of political concentration as well as those of economic concentration.

The remaining question is how reformers might go about building a political coalition that can tackle the challenges of crafting meaningful reform. Under normal circumstances, the political clout of the industry would make this a formidable obstacle. Yet as we discussed above,

our current circumstances present an opening. Democrats and progressives are concerned about the use of the platforms to manipulate public opinion. Conservatives are concerned about algorithmic biases and what they perceive as unfair moderation policies. Citizens, who are concerned about privacy, hate speech, and harassment, might be persuaded to support stronger measures. Private interest such as newspapers who have lost advertising revenue and retailers who either compete with or are forced to operate on Amazon's platform may be mobilized. At the same time, concern with monopoly and economic concentration has grown on the left and the right. While these various interests and constituencies may not agree on exactly what they would want a regulation and antitrust enforcement to achieve, their overlapping concerns may well provide the building blocks for efforts to create arenas in which the issues can be addressed.

References

- Ackland, Robert. 2013. *Web Social Science: Concepts, Data, Tools for Social Scientists in the Digital Age*. London: Sage Publication.
- Ahler, Douglas J. and Gaurav Sood. 2018. “The Parties in Our Heads: Misperceptions About Party Composition and Their Consequences.” *Journal of Politics* 80(3): 964–981.
- Aldrich, John H., Rachel K. Gibson, Marta Cantijoch and Tobias Konitzer. 2016. “Getting Out the Vote in the Social Media Era: Are Digital Tools Changing the Extent, Nature and Impact of Party Contacting in Elections?” *Party Politics* 22(2): 165–178.
- Allcott, Hunt, Luca Braghieri, Sarah Eichmeyer, and Matthew Gentzkow. 2019. *The Welfare Effects of Social Media*. No. w25514. National Bureau of Economic Research.
- Bakshy, Eytan, Solomon Messing, and Lada A. Adamic. 2015. “Exposure to Ideologically Diverse News and Opinion on Facebook,” *Science*, June 5, 1130.
- Balkin, Jack M. and Jonathan Zittrain, 2016. “A Grand Bargain to Make Tech Companies Trustworthy” *The Atlantic*, Oct. 3, 2016.
<https://www.theatlantic.com/technology/archive/2016/10/information-fiduciary/502346/>.
- Balkin, Jack M. 2016. *Information Fiduciaries and the First Amendment*, 49 U.C. Davis L. Rev. 1183.
- Bail, Christopher A., Lisa P. Argyle, Taylor W. Brown, John P. Bumpus, Haohan Chen, M.B. Fallin Hunzaker, Jaemin Lee, Marcus Mann, Friedolin Merhout, and Alexander Volfovsky. 2018. “Exposure to Opposing Views on Social Media Can Increase Political Polarization.” *Proceedings of the National Academy of Sciences* 115(37): 9216–9221.
- Banks, Antoine J. 2016. “Are Group Cues Necessary? How Anger Makes Ethnocentrism Among Whites a Stronger Predictor of Racial and Immigration Policy Opinions.” *Political Behavior*, 1–23.
- Barberá, Pablo. 2014. “How Social Media Reduces Mass Political Polarization. Evidence from Germany, Spain, and the US.” Unpublished paper, New York University.
- Barberá, Pablo, Ning Wang, Richard Bonneau, John T. Jost, Jonathan Nagler, Joshua A. Tucker, and Sandra González-Bailón. 2015. “The Critical Periphery in the Growth of Social Protests.” *PLOS ONE*, 10(11).
- Barberá, Pablo, Andreu Casas, Jonathan Nagler, Patrick Egan, Richard Bonneau, John T. Jost, and Joshua A. Tucker. 2018. “Who Leads? Who Follows? Measuring Issue Attention and Agenda Setting by Legislators and the Mass Public Using Social Media Data.” *American Political Science Review*, 1–19. <https://doi.org/10.1017/S0003055419000352>.

- Baron, David P. 2009. "A Positive Theory of Moral Management, Social Pressure, and Corporate Social Performance." *Journal of Economics & Management Strategy* 18: 7–43. <https://doi.org/10.1111/j.1530-9134.2009.00206.x>.
- Baron, David P. 2014. "Self-Regulation in Private and Public Politics." *Quarterly Journal of Political Science*: 9(2): 231–267. <http://dx.doi.org/10.1561/100.00013076>.
- Bennett, W. Lance, and Alexandra Segerberg. 2012. "Digital Media and the Personalization of Collective Action: Social Technology and the Organization of Protests against the Global Economic Crisis." In *Social Media and Democracy*, ed. Brian D. Loader and Dan Mercea, 23–48. Routledge.
- Bertrand, Marianne, Matilde Bombardini, Raymond Fisman, and Francesco Trebbi. 2018. *Tax-exempt Lobbying: Corporate Philanthropy as a Tool for Political Influence*. No. w24451. National Bureau of Economic Research.
- Bond, Robert M., Christopher J. Fariss, Jason J. Jones, Adam D.I. Kramer, Cameron Marlow, Jaime E. Settle, and James H. Fowler. 2012. "A 61-Million-Person Experiment in Social Influence and Political Mobilization." *Nature* 489(7415): 295.
- Boxell, Levi, Matthew Gentzkow, and Jesse M. Shapiro. 2017. "Is the Internet Causing Political Polarization? Evidence from Demographics." Technical report. National Bureau of Economic Research.
- Bradshaw, Samantha, and Philip N. Howard. 2018. "Challenging Truth and Trust: A Global Inventory of Organized Social Media Manipulation." *The Computational Propaganda Project*.
- Bucher, Taina, and Anne Helmond. 2017. "The Affordances of Social Media Platforms." *The SAGE Handbook of Social Media*: 223–253.
- Cantoni, Davide, David Y. Yang, Noam Yuchtman, and Y. Jane Zhang. 2017. "Are Protests Games of Strategic Complements or Substitutes? Experimental Evidence from Hong Kong's Democracy Movement." No. w23110. National Bureau of Economic Research.
- Carpenter, Danniell and, David A. Moss. 2013. *Preventing Regulatory Capture: Special Interest Influence and How to Limit It*. Cambridge: Cambridge University Press.
- Chan, Jason, Anindya Ghose, and Robert Seamans. 2015. "The Internet and Racial Hate Crime: Offline Spillovers from Online access." Unpublished Manuscript. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2335637.
- Chen, Anthony. 2019. "In the Private Interest? Business Influence and American Democracy." In *Can American Govern Itself?*, ed. Frances Lee and Nolan McCarty. New York: Cambridge University Press.

- Chess, Shira, and Adrienne Shaw. 2015. "A Conspiracy of Fishes, or, How We Learned to Stop Worrying about #GamerGate and Embrace Hegemonic Masculinity." *Journal of Broadcasting & Electronic Media* 59(1): 208–220.
- Clinton, Joshua D., and Ted Enamorado. 2014. "The National News Media's Effect on Congress: How Fox News Affected Elites in Congress." *Journal of Politics* 76(4): 928–943.
- Citron, Danielle Keats. 2014. *Hate Crimes in Cyberspace*. Harvard University Press.
- Cohen-Almagor, Raphael. 2011. "Fighting Hate and Bigotry on the Internet." *Policy & Internet* 3(3):1–26.
- Cohen-Almagor, Raphael. 2017. "Why Confronting the Internet's Dark Side?" *Philosophia* 45(3):919–929.
- Coglianesse, Cary, and David Lazer. 2003. "Managementbased Regulation: Prescribing Private Management to Achieve Public Goals." *Law & Society Review* 37(4): 691–730.
- Costello, Matthew, and James Hawdon. 2018. "Who Are the Online Extremists Among Us? Sociodemographic Characteristics, Social Networking, and Online Experiences of Those Who Produce Online Hate Materials." *Violence and Gender* 5(1): 55–60.
- Diamond, Larry. 2012. "Liberation Technology." In *Liberation Technology: Social Media and the Struggle for Democracy*, ed. Larry Diamond and Marc F. Plattner. Baltimore: Johns Hopkins University Press.
- Delgado, Richard, and Jean Stefancic. 2014. "Hate Speech in Cyberspace." *Wake Forest Law Review* 49: 319.
- DellaVigna, Stefano, and Ethan Kaplan. 2007. "The Fox News Effect: Media Bias and Voting." *The Quarterly Journal of Economics* 122(3):1187–1234.
- Dunleavy, Patrick, Helen Margetts, Jane Tinkler, and Simon Bastow. 2006. *Digital Era Governance: IT Corporations, the State, and E-government*. Oxford University Press.
- Eady, Gregory, Jonathan Nagler, Andy Guess, Jan Zilinsky, and Joshua A. Tucker. 2019. "How Many People Live in Political Bubbles on Social Media? Evidence From Linked Survey and Twitter Data." *SAGE Open*. <https://doi.org/10.1177/2158244019832705>
- Fabbri, Alice, et al. 2018. "The Influence of Industry Sponsorship on the Research Agenda: A Scoping Review." *American Journal of Public Health* 108(11): e1–8. <https://doi.org/10.2105/AJPH.2018.304677>.
- Feitelson, Dror G., Etan Frachtenberg, and Kent L. Beck. 2013. "Development and Deployment at Facebook." *IEEE Internet Computing* 17(4): 8–17. <https://doi.org/10.1109/MIC.2013.25>.

- Franco, Annie, Neil Malhotra, and Gabor Simonovits. 2014. "Publication bias in the social sciences: Unlocking the file drawer." *Science* Sept. 19, 1502–1505.
- Freelon, Deen, Charlton D. McIlwain, and Meredith Clark. 2016. "Beyond the Hashtags: #Ferguson, #Blacklivesmatter, and the Online Struggle for Offline Justice." Center for Media & Social Impact, American University, February 29. <http://dx.doi.org/10.2139/ssrn.2747066>.
- Gagliardone, Iginio, Matti Pohjonen, Zenebe Beyene, Abdissa Zerai, Gerawork Aynekulu, Mesfin Bekalu, Jonathan Bright, Mulatu Moges, Michael Seifu, and Nicole Stremmlau. 2016. "Mechachal: Online Debates and Elections in Ethiopia—From Hate Speech to Engagement in Social Media." https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2831369.
- Gil-García, J. Ramón, and Theresa A. Pardo. 2005. "E-government success factors: Mapping practical tools to theoretical foundations." *Government Information Quarterly* 22(2): 187–216.
- Gladwell, Malcolm. 2010. "Small Change: Why the Revolution will not be Tweeted." *The New Yorker*, September 27.
- Goel, Sharad, Ashton Anderson, Jake Hofman, and Duncan J. Watts. 2015. "The Structural Virality of Online Diffusion." *Management Science* 62(1): 180–196.
- González-Bailón, Sandra, Javier Borge-Holthoefer, Alejandro Rivero, and Yamir Moreno. 2011. "The Dynamics of Protest Recruitment through an Online Network." *Scientific Reports* 1: 197.
- González-Bailón, Sandra, and Ning Wang. 2016. "Networked Discontent: The Anatomy of Protest Campaigns in Social Media." *Social Networks* 44: 95–104.
- Graef, Inge and Jens Prufer. 2018. "Mandated Data Sharing Is a Necessity in Specific Sectors." Tilburg Law School Research Paper. *Economisch Statistische Berichten* 103(4763) 298–301. <https://ssrn.com/abstract=3206685>
- Grinberg, Nir, Kenneth Joseph, Lisa Friedland, Briony Swire-Thompson, and David Lazer. 2019. "Fake News on Twitter During the 2016 US Presidential Election." *Science* 363(6425): 374–378.
- Guess, Andrew M. 2018. "(Almost) Everything in Moderation: New Evidence on Americans' Online Media Diets." Working paper, Princeton University.
- Guess, Andrew M., Kevin Munger, Jonathan Nagler, and Joshua Tucker. 2019. "How Accurate Are Survey Responses on Social Media and Politics?" *Political Communication* 36(2): 241–258.

- Guess, Andrew M., Jonathan Nagler, and Joshua A. Tucker. 2019. Less Than You Think: Prevalence and Predictors of Fake News Dissemination on Facebook. *Science Advances* 5(1): eaau4586.
- Guess, Andrew, Brendan Nyhan, and Jason Reifler. 2018. "Selective Exposure to Misinformation: Evidence from the Consumption of Fake News during the 2016 US Presidential Campaign." Working Paper.
- Hopkins, Daniel J. and Jonathan M. Ladd. 2014. "The Consequences of Broader Media Choice: Evidence from the Expansion of Fox News." *Quarterly Journal of Political Science* 9(1): 115–135.
- Howard, Philip N. and Muzammil M. Hussain. 2013. *Democracy's Fourth Wave?: Digital Media and the Arab Spring*. Oxford University Press.
- Hinduja, Sameer, and Justin W. Patchin. 2007. "Offline Consequences of Online Victimization: School Violence and Delinquency." *Journal of School Violence* 6(3): 89–112.
- Iyengar, Shanto, Gaurav Sood, and Yphtach Lelkes. 2012. "Affect, Not Ideology: A Social Identity Perspective on Polarization." *Public Opinion Quarterly* 76(3): 405–431.
- Jackson, Sarah J., Moya Bailey, and Brooke Foucault Welles. 2019. "Women Tweet on Violence: From #YesAllWomen to #MeToo." *Ada: A Journal of Gender, New Media, and Technology* 15.
- Jiang, Ray, S. Chiappa, T. Lattimore, A. Agyorgy, and P. Kohli. 2019. "Degenerate Feedback Loops in Recommender Systems." <https://arxiv.org/abs/1902.10730>.
- Khan, Lina M. 2016. "Amazon's Antitrust Paradox." *Yale Law Journal* 126: 710
- Kennedy, M. Alexis, and Melanie A. Taylor. 2010. "Online Harassment and Victimization of College Students." *Justice Policy Journal* 7(1): 1–21.
- Kim, In Song. 2018. "LobbyView: Firm-level Lobbying and Congressional Bills Database." Working paper, August 30. <http://web.mit.edu/insong/www/pdf/lobbyview.pdf>.
- Kumar, Srijan, William L. Hamilton, Jure Leskovec, and Dan Jurafsky. 2018. "Community Interaction and Conflict on the Web." In *Proceedings of the 2018 World Wide Web Conference on World Wide Web*, ed. International World Wide Web Conferences Steering Committee, 933–943.
- Kwak, James. 2013. "Cultural Capture and the Financial Crisis." In *Preventing Regulatory Capture: Special Interest Influence and How to Limit It*, ed. Daniel Carpenter and David A. Moss. Cambridge: Cambridge University Press.
- Kwon, K. Hazel, and Anatoliy Gruzd. 2017. "Is Offensive Commenting Contagious Online? Examining Public vs Interpersonal Swearing in Response to Donald Trump's YouTube Campaign Videos." *Internet Research* 27(4): 991–1010.

- Larson, Jennifer M., Jonathan Nagler, Jonathan Ronen, and Joshua A. Tucker. 2019. "Social Networks and Protest Participation: Evidence from 130 Million Twitter Users." *American Journal of Political Science*. <https://doi.org/10.1111/ajps.12436>.
- Lelkes, Yphtach, Sood, G., and Iyengar, S. 2017. "The Hostile Audience: The Effect of Access to Broadband Internet on Partisan Affect." *American Journal of Political Science* 61: 5–20.
- Lenhart, Amanda, Michelle Ybarra, Kathryn Zickuhr, and Myeshia Price-Feeney. 2016. "Online Harassment, Digital Abuse, and Cyberstalking in America." Data & Society Institute. <https://datasociety.net/output/online-harassment-digital-abuse-cyberstalking/>.
- Lexchin, Joel. 2012. "Sponsorship Bias in Clinical Research." *The International Journal of Risk & Safety in Medicine* 24(4): 233–42. <https://doi.org/10.3233/JRS-2012-0574>.
- Lindblom, Charles E. 1983. *Politics and Markets: The World's Political-Economics Systems*. New York: Basic Books.
- Linders, Dennis. 2012. "From E-government to We-government: Defining a Typology for Citizen Coproduction in the Age of Social Media." *Government Information Quarterly* 29(4): 446–454.
- Lucca, David, Amit Seru, and Francesco Trebbi. 2014. "The Revolving Door and Worker Flows in Banking Regulation." *Journal of Monetary Economics* 65: 17–32.
- Magdy, Walid, Kareem Darwish, Norah Abokhodair, Afshin Rahimi, and Timothy Baldwin. 2016. "#isisisnotislam or #deportallmuslims?: Predicting Unspoken Views." In *Proceedings of the 8th ACM Conference on Web Science*, 95–106.
- Mantilla, Karla. 2013. "Gendertrolling: Misogyny Adapts to New Media." *Feminist Studies* 39(2): 563–570.
- Marwick, Alice, and Rebecca Lewis. 2017. "Media Manipulation and Disinformation Online." New York: Data & Society Research Institute.
- Mariconti, Enrico, Guillermo Suarez-Tangil, Jeremy Blackburn, Emiliano De Cristofaro, Nicolas Kourtellis, Ilias Leontiadis, Jordi Luque Serrano, and Gianluca Stringhini. 2018. "You Know What to Do: Proactive Detection of YouTube Videos Targeted by Coordinated Hate Attacks." <https://arxiv.org/abs/1805.08168>.
- Martin, Gregory J., and Ali Yurukoglu. 2017. "Bias in Cable News: Persuasion and Polarization." *American Economic Review* 107(9): 2565–99.
- McCarty, Nolan. 2013. "Complexity, Capacity, and Capture." In *Preventing Regulatory Capture: Special Interest Influence and How to Limit It*, ed. Daniel Carpenter and David A. Moss. Cambridge: Cambridge University Press, 99–123.

- McCarty, Nolan. 2017. "The Regulation and Self-Regulation of a Complex Industry." *Journal of Politics* 79(4): 1220-1236.
- McCarty, Nolan, Keith T. Poole, and Howard Rosenthal. 2013. *Political Bubbles: Financial Crises and the Failure of American Democracy*. Princeton University Press.
- Messing, Solomon, and Sean J. Westwood. 2014. "Selective Exposure in the Age of Social Media: Endorsements Trump Partisan Source Affiliation When Selecting News Online." *Communication Research* 41(8): 1042–1063.
- Metzger, Megan MacDuffee, and Joshua A. Tucker. 2017. "Social Media and EuroMaidan: A Review Essay." *Slavic Review* 76(1): 169–191.
- Moe, Terry M. 1989. "The Politics of Bureaucratic Structure." In *Can the Government Govern?* ed. John E. Chubb and Paul E. Peterson, 267–321. Washington: Brookings Institution.
- Morozov, Evgeny. 2011. *The Net Delusion: The Dark Side of Internet Freedom*. New York: PublicAffairs.
- Neblo, Michael A., Kevin M. Esterling, and David M. Lazer. 2018. *Politics with the People: Building a Directly Representative Democracy*. Cambridge University Press.
- Newman, Nic, Richard Fletcher, Antonis Kalogeropoulos, David AL Levy, and Rasmus Kleis Nielsen. 2017. "Reuters Institute Digital News Report 2017."
- Pennycook, Gordon, and David G. Rand. 2019. "Fighting Misinformation on Social Media using Crowdsourced Judgments of News Source Quality." *Proceedings of the National Academy of Sciences* 116(7): 2521–2526
- Persily, Nathaniel. 2019. *The Internet's Challenge to Democracy: Framing the Problem and Assessing Reforms*. Kofi Annan Foundation.
<https://www.kofiannanfoundation.org/electoral-integrity/democracy-digital-age-challenges-opportunities/>
- Postmes, Tom, Russell Spears, Khaled Sakhel, and Daphne De Groot. 2001. "Social Influence in Computer-mediated Communication: The Effects of Anonymity on Group Behavior." *Personality and Social Psychology Bulletin* 27(10), 1243–1254.
- Prior, Markus. 2007. *Post-broadcast Democracy: How Media Choice Increases Inequality in Political Involvement and Polarizes Elections*. New York: Cambridge University Press.
- Ramanna, Karthik. 2015. *Political Standards: Corporate Interest, Ideology, and Leadership in the Shaping of Accounting Rules for the Market Economy*. University of Chicago Press.
- Riedl, Christoph, Johannes Bjelland, Geoffrey Canright, Asif Iqbal, Kenth Engø-Monsen, Taimur Qureshi, Pål Roe Sundsøy, and David Lazer. 2018. "Product Diffusion Through On-demand Information-seeking Behaviour." *Journal of The Royal Society Interface* 15(139): 20170751.

- Sandvig, Christian, Kevin Hamilton, Karrie Karahalios, and Cedric Langbort. 2014. "Auditing algorithms: Research methods for detecting discrimination on internet platforms." Paper presented at Data and Discrimination: Converting Critical Concerns into Productive Inquiry, Seattle, WA, May 22.
- Sellars, Andrew. 2016. "Defining hate speech." Unpublished manuscript, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2882244.
- Siegel, Alexandra, Evgenii Nitikin, Pablo Barberá, Joanna Sterling, Bethany Pullen, Jonathan Nagler, Richard Bonneau, and Joshua Tucker. 2019. "Trumping Hate on Twitter: Online Hate Speech in the 2016 Presidential Election Campaign and its Aftermath." Unpublished manuscript.
- Siegel, Alexandra, Joshua Tucker, Jonathan Nagler, and Richard Bonneau. 2019. "Tweeting beyond Tahrir: Ideological Diversity and Political Tolerance in Egyptian Twitter networks." Unpublished manuscript.
- Sismondo, Sergio. 2007. "Ghost Management: How Much of the Medical Literature Is Shaped Behind the Scenes by the Pharmaceutical Industry?" *PLOS Medicine* 4(9): e286. *PLoS Journals*. <https://doi.org/10.1371/journal.pmed.0040286>.
- Suhay, Elizabeth, Emily Bello-Pardo, and Brianna Maurer. 2018. "The Polarizing Effects of Online Partisan Criticism: Evidence from Two Experiments." *The International Journal of Press/Politics* 23(1): 95–115.
- Sunstein, Cass R. 2018. *#Republic: Divided Democracy in the Age of Social Media*. Princeton University Press.
- Tolbert, Carolyn J., and Karen Mossberger. 2006. "The Effects of E-government on Trust and Confidence in Government." *Public Administration Review* 66(3): 354–369.
- Theocharis, Yannis, Pablo Barberá, Zoltán Fazekas, Sebastian Adrian Popa, and Olivier Parnet. 2016. "A Bad Workman Blames His Tweets: The Consequences of Citizens' Uncivil Twitter use When Interacting with Party Candidates." *Journal of Communication* 66(6): 1007–1031.
- Timm, J., and Pablo Barberá. nd. "Incivility Begets Incivility? Understanding the Contagion Dynamics of Uncivil Conversations on Facebook." Working paper.
- Tucker, Joshua A., Andy Guess, Pablo Barberá, Cristian Vaccari, Alexandra Siegel, Sergey Sanovich, Denis Stukal, and Brendan Nyhan. 2018. "Social Media, Political Polarization, and Political Disinformation: A Review of the Scientific Literature," March 19. <http://dx.doi.org/10.2139/ssrn.3144139>.
- Tufekci, Zeynep. 2017. *Twitter and Tear Gas: The Power and Fragility of Networked Protest*. New Haven: Yale University Press.

- Tynes, Brendesha M., Michael T. Giang, David R. Williams and Geneene N. Thompson. 2008. "Online Racial Discrimination and Psychological Adjustment among Adolescents." *Journal of Adolescent Health* 43(6): 565–569.
- Tynes, Brendesha M. and Suzanne L. Markoe. 2010. "The Role of Color-Blind Racial Attitudes in Reactions to Racial Discrimination on Social Network Sites." *Journal of Diversity in Higher Education* 3(1): 1.
- Vaccari, Cristian, Augusto Valeriani, Pablo Barberá, John T. Jost, Jonathan Nagler, and Joshua A. Tucker. 2016. "Of Echo Chambers and Contrarian Clubs: Exposure to Political Disagreement Among German and Italian Users of Twitter." *Social Media + Society* 2 (September): 1–24.
- West, Darrell M. 2000. *Assessing E-government: The Internet, Democracy, and Service Delivery by State and Federal Governments*. Washington, D.C.: World Bank.
- Wilson, Christo. 2018. "The Promise and Peril of Algorithm Audits for Increasing Transparency and Accountability of Donated Datasets." Unpublished manuscript.
- Wojcieszak, Magdalena E., and Diana C. Mutz. 2009. "Online Groups and Political Discourse: Do Online Discussion Spaces Facilitate Exposure to Political Disagreement?" *Journal of Communication* 59(1): 40–56.
- Wu, Tim. 2013. "Machine Speech." *University of Pennsylvania Law Review* 161(6): 1495–1533.
- Wu, Tim. 2018. *The Curse of Bigness: Antitrust in the New Gilded Age*. Columbia Global Reports.
- Zannettou, Savvas, Barry Bradlyn, Emiliano De Cristofaro, Michael Sirivianos, Gianluca Stringhini, Haewoon Kwak and Jeremy Blackburn. 2018. "What is Gab? A Bastion of Free Speech or an Alt-Right Echo Chamber?" <https://arxiv.org/abs/1802.05287>.
- Zingales, Luigi. 2013. "The Capture of Economists." In *Preventing Regulatory Capture: Special Interest Influence and How to Limit It*, ed. Daniel Carpenter and David A. Moss. Cambridge: Cambridge University Press.
- Zittrain, Jonathan. 2014. "Facebook Could Decide an Election Without Anyone Ever Finding Out." *The New Republic*, June 1. <https://newrepublic.com/article/117878/information-fiduciary-solution-facebook-digital-gerrymandering>.
- Zittrain, Jonathan. 2018. "How to Exercise the Power You Didn't Ask For" *Harvard Business Review*, Sept. 19. <https://hbr.org/2018/09/how-to-exercise-the-power-you-didnt-ask-for>