

Potential Policy Proposals for Regulation of Social Media and Technology Firms

Social media and wider digital communications technologies have changed our world in innumerable ways. They have transformed the way we do everything from shopping for groceries to growing our small businesses and have radically lowered the cost of, and barriers to, global communication. The American companies behind these products and services – Facebook, Google, Twitter, Amazon, and Apple, among others – have been some of the most successful and innovative in the world. As such, each of them deserves enormous recognition for the technological transformation they have engendered around the world. As their collective influence has grown, however, these tech giants now also deserve increased scrutiny.

In the course of investigating Russia’s unprecedented interference in the 2016 election, the extent to which many of these technologies have been exploited – and their providers caught repeatedly flat-footed – has been unmistakable. More than illuminating the capacity of these technologies to be exploited by bad actors, the revelations of the last year have revealed the dark underbelly of an entire ecosystem. The speed with which these products have grown and come to dominate nearly every aspect of our social, political and economic lives has in many ways obscured the shortcomings of their creators in anticipating the harmful effects of their use. Government has failed to adapt and has been incapable or unwilling to adequately address the impacts of these trends on privacy, competition, and public discourse.

Armed with this knowledge, it is time to begin to address these issues and work to adapt our regulations and laws. There are three areas that should be of particular focus for policymakers.

First, understanding the capacity for communications technologies to promote disinformation that undermines trust in our institutions, democracy, free press, and markets. In many ways, this threat is not new. For instance, Russians have been conducting information warfare for decades. During the Cold War, the Soviets tried to spread “fake news” denigrating Martin Luther King Jr. and alleging that the American military had manufactured the

AIDS virus.¹ Much like today, their aim was to undermine Americans' faith in democratic government. But what *is* new is the advent of social media tools with the power to magnify – and target – propaganda and fake news on a scale that was unimaginable back in the days of the Berlin Wall. As one witness noted during the March 2017 hearing on Russian disinformation efforts before the Senate Select Committee on Intelligence, today's tools seem almost *purpose-built* for Russian disinformation techniques.²

Just as we're trying to sort through the disinformation playbook used in the 2016 election and as we prepare for additional attacks in 2018, a new set of tools is being developed that are poised to exacerbate these problems. Aided in large part by advances in machine learning, tools like DeepFake allow a user to superimpose existing images and videos onto unrelated images or videos. In addition, we are seeing an increasing amount of evidence that bad actors are beginning to shift disinformation campaigns to encrypted messaging applications rather than using the relatively more open social media platforms. Closed applications like WhatsApp, Telegram, Viber, and others, present new challenges for identifying, rapidly responding to, and fact-checking misinformation and disinformation targeted to specific users.³

But it's also important to recognize that manipulation and exploitation of the tools and scale these platforms provide goes beyond just foreign disinformation efforts. In the same way that bots, trolls, click-farms, fake pages and groups, ads, and algorithm-gaming can be used to propagate political disinformation, these same tools can – and have – been used to assist financial frauds such as stock-pumping schemes, click fraud in digital advertising markets, schemes to sell counterfeit prescription drugs, and efforts to convince large numbers of users to download malicious apps on their phones.⁴ Addressing these diseconomies of scale – negative

¹ U.S. Department of State: *Soviet Influence Activities: A Report on Active Measures and Propaganda, 1986-1987* (August 1987), <https://www.globalsecurity.org/intell/library/reports/1987/soviet-influence-activities-1987.pdf>.

² U.S. Congress, Senate, Select Committee on Intelligence, *Open Hearing: Disinformation: A Primer in Russian Active Measures and Influence Campaigns*. 115th Cong., 1st sess., 2017.

³ See Elizabeth Dwoskin, & Annie Gowen, "On WhatsApp, fake news is fast – and can be fatal," Washington Post. July 23, 2018. https://www.washingtonpost.com/business/economy/on-whatsapp-fake-news-is-fast--and-can-be-fatal/2018/07/23/a2dd7112-8ebf-11e8-bcd5-9d911c784c38_story.html?utm_term=.ba2797f74d7d; Nic Dias, "The Era of WhatsApp Propaganda Is Upon Us," Foreign Policy. August 17, 2017. <https://foreignpolicy.com/2017/08/17/the-era-of-whatsapp-propaganda-is-upon-us/>.

⁴ See, e.g., Robert Gorwa, "Computational Propaganda in Poland: False Amplifier and the Digital Public Sphere," Working Paper No. 2017.4, Oxford Internet Institute, University of Oxford.

externalities borne by users and society as a result of the size of these platforms – represents a priority for technology policy in the 21st century.

A second dimension relates to consumer protection in the digital age. As online platforms have gained greater prominence in our lives, they have developed more advanced capabilities to track and model consumer behavior – typically across the multiple devices a consumer owns. This includes detailed information on viewing, window-shopping, and purchasing habits, but also more sensitive information. The prevailing business model involves offering nominally free services, but which results in consumers providing ever-more data in exchange for continued usage.

User tracking can have important consumer benefits, for instance by showing users more relevant ads and helping to optimize user experience across different apps. At the same time, these user profiles could provide opportunities for consumer harm – and in surreptitious, undetectable ways. Pervasive tracking may give platforms important behavioral information on a consumer’s willingness to pay or on behavioral tendencies that can be exploited to drive engagement with an app or service. These technologies might even be used to influence how we engage with our own democracy here at home, as we saw in recent months with the Cambridge Analytica scandal, where sensitive Facebook data from up to 87 million people may have been used to inappropriately target U.S. voters.

The allure of pervasive tracking also creates incentives to predicate services and credit on user behavior. Users have no reason to expect that certain browsing behavior could determine the interest they pay on an auto-loan, much less that what their friends post could be used to determine that. Further, numerous studies indicate users have no idea their information is being

<http://blogs.oii.ox.ac.uk/politicalbots/wp-content/uploads/sites/89/2017/06/Comprop-Poland.pdf>; Renae Merle, “Scheme created fake news stories to manipulate stock prices, SEC alleges,” *Los Angeles Times*. July 5, 2017. <http://www.latimes.com/business/la-fi-sec-fake-news-20170705-story.html>; Lauren Moss, “Xanax drug sold on social media found to be fake,” *BBC News*. March 26, 2018. <https://www.bbc.com/news/uk-england-43543519>; Danny Palmer, “Android malware found inside apps downloaded 500,000 times,” *ZDNet*. March 26, 2018. <https://www.zdnet.com/article/android-malware-found-inside-apps-downloaded-500000-times/>.

used in this manner, resulting in a massive informational asymmetry.⁵ Important policy mechanisms include requiring greater disclosure by platforms – and in clear, concise ways – about the types of information they collect, *and the specific ways they are utilizing it*.

Lastly, the rise of a few dominant platforms poses key problems for long-term competition and innovation across multiple markets, including digital advertising markets (which support much of the Internet economy), future markets driven by machine-learning and artificial intelligence, and communications technology markets. User data is increasingly the single most important economic input in information markets, allowing for more targeted and relevant advertisements, facilitating refinement of services to make them more engaging and efficient, and providing the basis for any machine-learning algorithms (which, for instance, develop decisional rules based on pattern-matching in large datasets) on which all industries will increasingly rely.

Unlike many other assets, which tend to illustrate declining marginal utility, the value of any piece of data *increases* in combination with additional data.⁶ Relatedly, data exhibits economies of scale, enabling more effective data analysis, computationally intensive pattern recognition and computational learning with greater collected data.⁷ As a consequence, firms with large preexisting data sets have potentially insuperable competitive advantages over new entrants and

⁵ Lee Raine, “Americans’ Complicated Feelings About Social Media in An Era of Privacy Concerns,” *Pew Research Center*. March 27, 2018. <http://www.pewresearch.org/fact-tank/2018/03/27/americans-complicated-feelings-about-social-media-in-an-era-of-privacy-concerns/> (noting that “people struggle to understand the nature and scope of the data collected about them”); Timothy Morey et al., “Customer Data: Designing for Transparency and Trust,” *Harvard Business Review*. May 2015. <https://hbr.org/2015/05/customer-data-designing-for-transparency-and-trust> (“While awareness varied by country...overall the survey revealed an astonishingly low recognition of the specific types of information tracked online. On average, only 25% of people knew that their data footprints included information on their location, and just 14% understood that they were sharing their web-surfing history too.”).

⁶ Maurice E. Stucke & Allen P. Grunes, *Big Data and Competition Policy* (Oxford University Press, 2016), 200-201.; OECD, “Data-Driven Innovation for Growth and Well-Being: Interim Synthesis Report,” (October 2014), 29. <https://www.oecd.org/sti/innodata-driven-innovation-interim-synthesis.pdf> (“The diversification of services leads to even better insights if data linkage is possible. This is because data linkage enables ‘super-additive’ insights, leading to increasing ‘returns to scope.’ Linked data is a means to contextualize data and thus a source for insights and value that are greater than the sum of its isolated parts (data silos).”).

⁷ OECD, “Exploring the Economics of Personal Data,” (2013), 34. (“The monetary, economic and social value of personal data is likely to be governed by non-linear, increasing returns to scale. The value of an individual record, alone, may be very low but the value and usability of the record increases as the number of records to compare it with increases.”); *see also* Frank Pasquale, “Paradoxes of Digital Antitrust.” *Harvard Journal of Law & Technology*. July 2013. <https://jolt.law.harvard.edu/assets/misc/Pasquale.pdf> (describing the “Matthew Effect” in digital markets).

nascent firms.⁸ Dominant platforms have also aggressively commercialized psychology research, identifying ways to exploit cognitive biases and psychological vulnerabilities to keep users on the site and addicted to their products, generating more behavior data to mine.⁹ As machine-learning and AI begin to animate a wider variety of fields – medicine, transportation, law, accounting/book-keeping, financial services – a handful of large platforms may be able to leverage large datasets to develop products faster and more efficiently than competitors. These advantages are especially pronounced because many machine-learning and AI techniques are openly-extensible: pattern recognition, decisional rules, and computational learning tools can be applied on a new dataset (like tumor images) even if they were developed from a completely dissimilar dataset (such as cat pictures).

Policy Options

The size and reach of these platforms demand that we ensure proper oversight, transparency and effective management of technologies that in large measure undergird our social lives, our economy, and our politics. Numerous opportunities exist to work with these companies, other stakeholders, and policymakers to make sure that we are adopting appropriate safeguards to ensure that this ecosystem no longer exists as the ‘Wild West’ – unmanaged and not accountable to users or broader society – and instead operates to the broader advantage of society, competition, and broad-based innovation.

The purpose of this document is to explore a suite of options Congress may consider to achieve these objectives. In many cases there may be flaws in each proposal that may undercut the goal the proposal is trying achieve, or pose a political problem that simply can’t be overcome at this

⁸ See Tom Simonite, “AI and ‘Enormous Data’ Could Make Tech Giants Harder to Topple,” *Wired*. July 17, 2017. <https://www.wired.com/story/ai-and-enormous-data-could-make-tech-giants-harder-to-topple/>; see also Alon Halevy, Peter Norvig, & Fernando Pereira, “The Unreasonable Effectiveness of Data,” *IEEE Intelligent Systems*. March/April 2009. (concluding that “invariably, simple models and a lot of data trump more elaborate models based on less data.”); Chen Sun, Abhinav Shrivastava, Saurabh Singh & Abhinav Gupta, “Revisiting Unreasonable Effectiveness of Data in Deep Learning Era,” *Google AI Blog*. July 11, 2017. <https://ai.googleblog.com/2017/07/revisiting-unreasonable-effectiveness.html> (finding that performance of computer vision models increases *logarithmically* based on the volume of training data).

⁹ Ian Leslie, “The Scientists Who Make Apps Addictive,” *The Economist: 1843 Magazine*. October/November 2016. <https://www.1843magazine.com/features/the-scientists-who-make-apps-addictive>.

time. This list does not represent every idea, and it certainly doesn't purport to answer all of the complex and challenging questions that are out there. The hope is that the ideas enclosed here stir the pot and spark a wider discussion – among policymakers, stakeholders, and civil society groups – on the appropriate trajectory of technology policy in the coming years.

Disinformation and Misinformation/Exploitation of Technology

Duty to clearly and conspicuously label bots – Bots play a significant role in the amplification and dissemination of disinformation. Bot-enabled amplification and dissemination have also been utilized for promoting scams and financial frauds.¹⁰ New technologies, such as Google Assistant's AI-enabled Duplex, will increasingly make bots indistinguishable from humans (even in voice interfaces). To protect consumers, and to inhibit the use of bots for amplification of both disinformation and misinformation, platforms should be under an obligation to label bots – both those they provide (like Google's Duplex) and those used on the platforms they maintain (e.g. bot-enabled accounts on Twitter). California lawmakers have proposed something like it – colloquially referred to as a 'Blade Runner law' after the 1980s movie – to do just this.¹¹

Duty to determine origin of posts and/or accounts – Anonymity and pseudo-anonymity on social media platforms have enabled bad actors to assume false identities (and associated locations) allowing them to participate and influence political debate on social media platforms. We saw this during the 2016 election, as IRA-affiliated actors pretended to be real Americans, flooding Facebook and Twitter newsfeeds with propaganda and disinformation. Forcing the platform companies to determine and/or authenticate the origin of accounts or posts would go far in limiting the influence of bad actors outside the United States. Facebook appears to have trialed an approach similar to this in May 2018:

¹⁰ Samuel C. Woolley & Philip N. Howard, "Computational Propaganda Worldwide: Executive Summary," Working Paper No. 2017.11. (Oxford Internet Institute, University of Oxford), <https://blogs.cranfield.ac.uk/is/i-reference-working-paper>.

¹¹ SB-1001, "Bolstering Online Transparency (BOT) Act of 2018," CA Legislature 2017-2018.



However, due to the widespread use of VPN's and other methods for masking IP addresses, determining the true origin of posts or accounts can be technically challenging. Such a scheme could result in a large number of false positives, potentially undermining its value. Facebook's trial, for instance, apparently associated pages with particular locations simply because a page admin had logged into their Facebook account from that country while traveling.

A duty on the part of service providers to identify the origin of posts or accounts raises a number of privacy concerns. For one, it may incentivize online service providers to adopt identity verification policies at the cost of user privacy. Facebook has, for instance, come under criticism from a variety of groups and advocates – LGBT, Native American, and human rights groups – for its real name policy. It may also better enable online platforms to track users. Lastly, location identification could potentially enable oppressive regimes to undermine and attack freedom of expression and privacy – particularly for those most vulnerable, including religious and ethnic minorities, dissidents, human rights defenders, journalists, and others. Any effort on this front must address the real safety and security concerns of these types of at-risk individuals.

Duty to identify inauthentic accounts – A major enabler of disinformation is the ease of creating and maintaining inauthentic accounts (not just bots, but in general, accounts that are

based on false identities). Inauthentic accounts not only pose threats to our democratic process (with inauthentic accounts disseminating disinformation or harassing other users), but also undermine the integrity of digital markets (such as digital advertising). Platforms have perverse incentives not to take inauthentic account creation seriously: the steady creation of new accounts allows them to show continued user growth to financial markets, and generates additional digital advertising money (both in the form of inauthentic views and from additional – often highly sensational – content to run ads against). A law could be crafted imposing an affirmative, ongoing duty on platforms to identify and curtail inauthentic accounts, with an SEC reporting duty to disclose to the public (and advertisers) the number of identified inauthentic accounts and the percentage of the platform’s user base that represented. Legislation could also direct the FTC to investigate lapses in addressing inauthentic accounts under its authority to address unfair and deceptive trade practices. Failure to appropriately address inauthentic account activity – or misrepresentation of the extent of the problem – could be considered a violation of both SEC disclosure rules and/or Section 5 of the FTC Act.

Like a duty to determine the origin of accounts or posts, however, a duty on the part of online service providers to identify inauthentic accounts may have the effect of incentivizing providers to adopt identity verification policies, at the cost of user privacy. Mandatory identity verification is likely to arouse significant opposition from digital privacy groups and potentially from civil rights and human rights organizations who fear that such policies will harm at-risk populations. In addition, any effort in this area needs to consider the distinction between inauthentic accounts created in order to mislead or spread disinformation from accounts clearly set up for satire and other legitimate forms of entertainment or parody.

Make platforms liable for state-law torts (defamation, false light, public disclosure of private facts) for failure to take down deep fake or other manipulated audio/video content

– Due to Section 230 of the Communications Decency Act, internet intermediaries like social media platforms are immunized from state tort and criminal liability. However, the rise of technology like DeepFakes – sophisticated image and audio synthesis tools that can generate fake audio or video files falsely depicting someone saying or doing something – is poised to usher in an unprecedented wave of false and defamatory content, with state law-based torts

(dignitary torts) potentially offering the only effective redress to victims. Dignitary torts such as defamation, invasion of privacy, false light, and public disclosure of private facts represent key mechanisms for victims to enjoin – and deter – sharing of this kind of content.

Currently the onus is on victims to exhaustively search for, and report, this content to platforms – who frequently take months to respond and who are under no obligation thereafter to proactively prevent the same content from being re-uploaded in the future.¹² Many victims describe a ‘whack-a-mole’ situation.¹³ Even if a victim has successfully secured a judgment against the user who created the offending content, the content in question in many cases will be re-uploaded by other users. In economic terms, platforms represent “least-cost avoiders” of these harms; they are in the best place to identify and prevent this kind of content from being propagated on their platforms. Thus, a revision to Section 230 could provide the ability for users who have successfully proved that sharing of particular content by another user constituted a dignitary tort to give notice of this judgement to a platform; with this notice, platforms would be liable in instances where they did not prevent the content in question from being re-uploaded in the future – a process made possible by existing perceptual hashing technology (e.g. the technology they use to identify and automatically take down child pornography). Any effort on this front would need to address the challenge of distinguishing true DeepFakes aimed at spreading disinformation from satire or other legitimate forms of entertainment and parody.

Reforms to Section 230 are bound to elicit vigorous opposition, including from digital liberties groups and online technology providers. Opponents of revisions to Section 230 have claimed that the threat of liability will encourage online service providers to err on the side of content takedown, even in non-meritorious instances. Attempting to distinguish between true disinformation and legitimate satire could prove difficult. However, the requirement that plaintiffs successfully obtain court judgements that the content in question constitutes a dignitary tort – which provides significantly more process than something like the Digital Millennium

¹² Chris Silver Smith, “Paradigm Shift: Has Google Suspended Defamation Removals?” *Search Engine Land*. December 30, 2016. <https://searchengineland.com/paradigm-shift-google-suspends-defamation-removals-266222>.

¹³ Kari Paul, “Reddit’s Revenge Porn Policy Still Puts the Onus on Victims, Advocates Say,” *Motherboard*. February 26, 2015. https://motherboard.vice.com/en_us/article/8qzkz3/anti-revenge-porn-advocates-are-skeptical-of-reddits-new-policy.

Copyright Act (DMCA) notice and takedown regime for copyright-infringing works – may limit the potential for frivolous or adversarial reporting. Further, courts already must make distinctions between satire and defamation/libel.

Public Interest Data Access Bill – One of the gravest problems identified by people like Tristan Harris, Wael Ghonim, and Tom Wheeler is that regulators, users, and relevant NGOs lack the ability to identify potential problems (public health/addiction effects, anticompetitive behavior, radicalization) and misuses (scams, targeted disinformation, user-propagated misinformation, harassment) on the platforms because access to data is zealously guarded by the platforms.¹⁴ Under this view, we could propose legislation that guarantees that platforms above a certain size provide independent, public interest researchers with access to anonymized activity data, at scale, via a secure API. The goal would be to allow researchers to measure and audit social trends on platforms. This would ensure that problems on, and misuse of, the platforms were being evaluated by researchers and academics, helping generate data and analysis that could help inform actions by regulators or Congress.

While at first glance this might seem drastic, the upshot is that the platforms have already developed methods by which researchers can gain anonymized activity data, at scale; the current problem is that much of this research is proprietary and platforms typically condition access to it on a researcher signing an NDA (compromising their independence). Further, as Bloomberg has reported, platforms have typically sought collaborations with researchers whose projects comport with their business goals, while excluding researchers whose work may be adverse to their interests.¹⁵ Under immense public and political pressure, Facebook has proposed a system

¹⁴ Wael Ghonim & Jake Rashbass, “It’s Time to End the Secrecy and Opacity of Social Media,” *Washington Post*. October 31, 2017. https://www.washingtonpost.com/news/democracy-post/wp/2017/10/31/its-time-to-end-the-secrecy-and-opacity-of-social-media/?noredirect=on&utm_term=.f66cb6adce18; Stefan Verhulst & Andrew Young, “How the Data that Internet Companies Collect Can Be Used for the Public Good,” *Harvard Business Review*. January 23, 2018. <https://hbr.org/2018/01/how-the-data-that-internet-companies-collect-can-be-used-for-the-public-good>; Tom Wheeler, “How to Monitor Fake News,” *New York Times*. February 20, 2018. <https://www.nytimes.com/2018/02/20/opinion/monitor-fake-news.html>.

¹⁵ Karen Weise & Sarah Frier, “If You’re A Facebook User, You’re Also a Research Subject,” *Bloomberg*. June 14, 2018. <https://www.bloomberg.com/news/articles/2018-06-14/if-you-re-a-facebook-user-you-re-also-a-research-subject>.

somewhat similar to a public interest research access regime, in collaboration with the Social Science Research Council.

Large-scale implementation of such an initiative does present a number of practical challenges, however. To protect user privacy, a number of controls would need to be required – including contractual controls, technical controls, criminal penalties for misuse of data by researchers, extensive auditing, compliance checks, and institutional review boards (IRBs). At the same time, extensive privacy protections may simultaneously inhibit the ability of researchers to effectively use platform data for research.

Further, experts point out that as important as ensuring researcher access to platform data is regulating the commercial use of behavior data by platforms. Experts have pointed to a need to regulate the use of corporate behavioral science, focusing on research controls (such as requiring companies to run research through an IRB) and the implications of behavior research on their business models. Commercial behavioral science may provide large platforms with unfair competitive advantages, allowing platforms to use behavior data to model new features that drive higher levels of user engagements. These practices even extend to *conditioning* user behavior – designing (and refining) products to be intentionally habit-forming. These practices raise important questions related to consumer protection, competition, and privacy.

Require Interagency Task Force for Countering Asymmetric Threats to Democratic

Institutions – After multiple briefings and discussions, it is evident that the intelligence and national security communities are not as well-positioned to detect, track, attribute, or counter malicious asymmetric threats to our political system as they should be. From information operations to cyber-attacks to illicit finance and money laundering, our democratic institutions face a wide array of new threats that don't fit easily into our current national security authorities and responsibilities. As just one example, programs to detect and protect against information operations are disparately positioned with unclear reporting chains and lack metrics for measuring success. Standing up a congressionally-required task force would help bring about a whole-of-government approach to counter asymmetric attacks against our election infrastructure and would reduce gaps that currently exist in tracking and addressing the threat. This typically

could be done by the President without legislation; however, President Trump seems unwilling to touch the issue, and as such, Congress could force the issue as they did with the creation of the State Department Global Engagement Center. However, as the GEC has proven, without engaged leadership these types of legislated entities can easily be starved of resources or authorities.

Disclosure Requirements for Online Political Advertisements – As the Senate Select Committee on Intelligence helped to uncover during its investigation into Russian interference in the 2016 elections, the ease by which our foreign adversaries purchased and targeted politically oriented ads during the campaign exposed an obvious threat to the integrity of our democracy. Because outdated election laws have failed to keep up with evolving technology, online political ads have had very little accountability or transparency, as compared to ads sold on TV, radio, and satellite. Improving disclosure requirements for online political advertisements and requiring online platforms to make all reasonable efforts to ensure that foreign individuals and entities are not purchasing political ads seem like a good first step in bringing more transparency online. The Honest Ads Act (S.1989) is one potential path, but there are other reasonable ways to increase disclosure requirements in this space.

Public Initiative for Media Literacy – Addressing the challenge of misinformation and disinformation in the long-term will ultimately need to be tackled by an informed and discerning population of citizens who are both alert to the threat but also armed with the critical thinking skills necessary to protect against malicious influence. A public initiative – propelled by federal funding but led in large part by state and local education institutions – focused on building media literacy from an early age would help build long-term resilience to foreign manipulation of our democracy. Such an effort could benefit from the resources and knowledge of private sector tech companies, as well as the expertise and training of some of the country’s most credible and trustworthy media entities. One particularly difficult challenge in any long-term effort like this, however, is establishing and tracking metrics for real success. It is not enough for social media companies or the tech community to simply give lip service to building long-term resiliency and media literacy without taking some much more significant short-term steps in addressing the

threat we face in the here and now. A public effort like this should be seen as augmenting or supporting more assertive and more aggressive policy steps.

At the same time, technology scholars such as danah boyd have argued that emphasis on media literacy obscures the real problems around online consumption of misinformation: distrust of media sources and a proclivity of users to *deploy* online information in service of strongly-held ideological or identity-based claims or beliefs.¹⁶ A recent study by Gallup and the Knight Foundation found that “People rate [p]artisan news stories as more or less trustworthy depending on whether the source is viewed as sympathetic or hostile to their political preferences” rather than on the content of the story.¹⁷ Under this view, empowering individuals as fact-checkers and critics may exacerbate distrust of institutions and information intermediaries. More important than building capacity for individuals to scrutinize sources is cultivating a recognition that information can (and will) be weaponized in novel ways, along with an understanding of the pathways by which misinformation spreads.

Increasing Deterrence Against Foreign Manipulation – The U.S. government needs to do more strengthen our security against these types of asymmetric threats. We have to admit that our strategies and our resources have not shifted to aggressively address these new threats in cyberspace and on social media that target our democratic institutions. Russia spends about \$70 billion a year on their military. We spend ten times that. But we’re spending it mostly on physical weapons designed to win wars that take place in the air, on land, and on sea. While we need to have these conventional capabilities, we must also expand our capabilities so that we can win on the expanded battlefields of the 21st century. Until we do that, Russia is going to continue getting a lot more bang for its buck.

The consequences of this problem are magnified because we lack a deterrence strategy that would discourage cyberattacks or information warfare targeting our democracy. In the absence

¹⁶ danah boyd. “You Think You Want Media Literacy...Do You?” *Medium*. March 9, 2018. <https://points.datasociety.net/you-think-you-want-media-literacy-do-you-7cad6af18ec2>.

¹⁷ “An Online Experimental Platform to Assess Trust in the Media,” *Gallup Inc.* and the *John S. and James L. Knight Foundation*. July 18, 2018. <https://www.knightfoundation.org/reports/an-online-experimental-platform-to-assess-trust-in-the-media>.

of a credible deterrent, there is nothing preventing Russia or another adversary from just continuing to use a tool that, frankly, has been working. It is not even clear which of the numerous agencies and departments tasked with responding to the cyber threat is supposed to be in charge.

We must spell out a deterrence doctrine, so that our adversaries don't see information warfare or cyberattacks against us as a "free lunch." The U.S. has often done too little to respond to these attacks against us or our allies. When we do respond, it has often been done quietly, and on a one-off basis. That's not been enough to deter future action. We need to make clear to Russia and other nations, that if you go after us in the cyber realm, we're going to punch back using our own cyber capabilities. And we need to increase the costs of this activity with robust sanctions and other tools.

Privacy and Data Protection

Information fiduciary – Yale law professor Jack Balkin has formulated a concept of "information fiduciaries" – service providers who, because of the nature of their relationship with users, assume special duties to respect and protect the information they obtain in the course of the relationships. Balkin has proposed that certain types of online service providers – including search engines, social networks, ISPs, and cloud computing providers – be deemed information fiduciaries because of the extent of user dependence on them, as well as the extent to which they are entrusted with sensitive information.¹⁸ A fiduciary duty extends beyond a mere tort duty (that is, a duty to take appropriate care): a fiduciary duty would stipulate not only that providers had to zealously protect user data, but also pledge not to utilize or manipulate the data for the benefit of the platform or third parties (rather than the user). This duty could be

¹⁸ Jack M. Balkin, "Information Fiduciaries and the First Amendment," *UC Davis Law Review*, Vol. 39, No. 4. April 2016. https://lawreview.law.ucdavis.edu/issues/49/4/Lecture/49-4_Balkin.pdf (noting that in addition to performing professional services, "fiduciaries also handle sensitive personal information. That is because, at their core, fiduciary relationships are relationships of trust and confidence that involve the use and exchange of information."); Jack M. Balkin & Jonathan Zittrain, "A Grand Bargain to Make Tech Companies Trustworthy," *The Atlantic*. October 3, 2016. <https://www.theatlantic.com/technology/archive/2016/10/information-fiduciary/502346/>.

established statutorily, with defined functions/services qualifying for classification as an information fiduciary.

Concretely defining what responsibilities a fiduciary relationship entails presents a more difficult challenge. Appropriate responsibilities may vary based on a number of factors, including the value that consumers derive from the service, whether consumers are paying monetarily for the service, and the extent of data collection by the service provider. Applying a one-size-fits-all set of fiduciary duties may inhibit the range of services consumers can access, while driving online business models towards more uniform offerings.

Privacy rulemaking authority at FTC – Many attribute the FTC’s failure to adequately police data protection and unfair competition in digital markets to its lack of genuine rulemaking authority (which it has lacked since 1980). Efforts to endow the FTC with rulemaking authority – most recently in the context of Dodd-Frank – have been defeated. If the FTC had genuine rulemaking authority, many claim, it would be able to respond to changes in technology and business practices. In addition, many have suggested that Congress should provide the FTC with additional resources. The FTC’s funding since 2010 has fallen by 5%. Significantly more funding is necessary in order for the FTC to develop tools necessary to evaluate complex algorithmic systems for unfairness, deception, or competition concerns.

Comprehensive (GDPR-like) data protection legislation – The US could adopt rules mirroring GDPR, with key features like data portability, the right to be forgotten, 72-hour data breach notification, 1st party consent, and other major data protections. Business processes that handle personal data would be built with data protection by design and by default, meaning personal data must be stored using pseudonymisation or full anonymization. Under a regime similar to GDPR, no personal data could be processed unless it is done under a lawful basis specified by the regulation, or if the data processor has received an unambiguous and individualized consent from the data subject (1st party consent). In addition, data subjects have the right to request a portable copy of the data collected by a processor and the right to have their data erased. Businesses must report any data breaches within 72 hours if they have an adverse effect on user privacy. One major tenet of the GDPR (that the US could or could not adopt) is the potential of

high penalties for non-compliance in which a company or organization can be fined (in the EU, penalties are up to 4% of its annual global turnover or €20 million - whichever is higher).

U.S. firms have voiced several concerns about the GDPR, including how it will be implemented and the scale of potential fines. In addition, if GDPR- like legislation were to be proposed, a central authority would need to be created to enforce these regulations. E.U. member states have their own data privacy authorities to enforce the GDPR, but this does not exist in the U.S. Delegating this responsibility to states could result in a patchwork of data protection and privacy regulations.

In some respects, there are also indications GDPR may take too extreme a view of what constitutes personal data. For instance, domain registration information – the historically public information about the individual who has registered a given domain, which operates much like a phonebook – is treated as personal data under GDPR. This poses serious problems to operation of the WHOIS database – a vital repository of domain registry information for those investigating online scammers – and many have suggested it will undermine cybersecurity investigations.

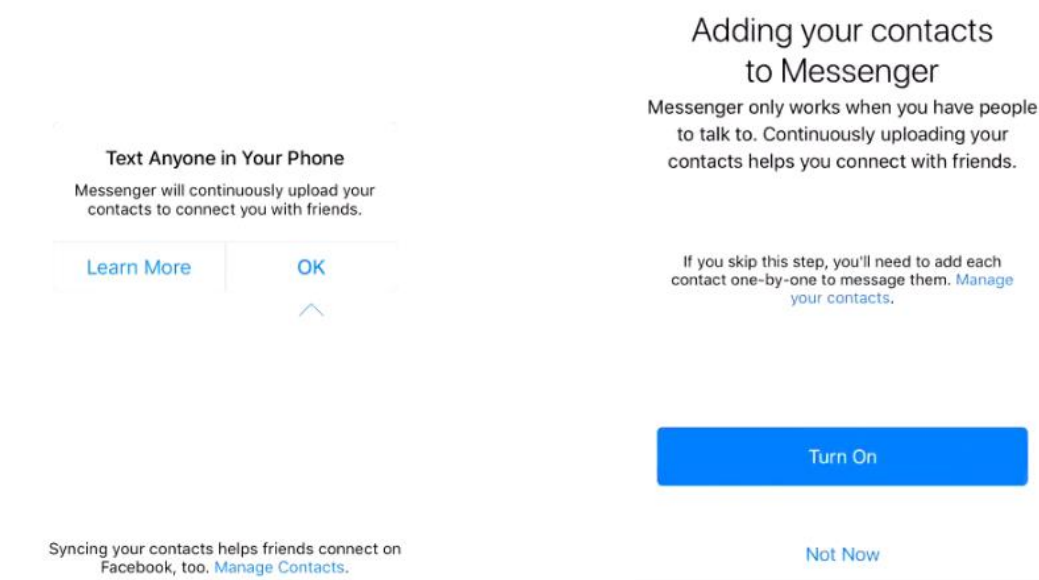
1st Party Consent for Data Collection – The US could adopt one specific element of GDPR: requiring 1st party consent for any data collection and use. This would prevent third-parties from collecting or processing a user’s data without their explicit and informed consent. Because third-party data is a practice reliant on consent that isn’t explicit, GDPR renders all third-party activity obsolete. Critics have acknowledged the need to remove some of the more salacious practices that go on with third-party access to data, but have also called for more clarity on the explicit consent side due to the negative connotations that can result from removing the third party data market in its entirety. Under GDPR, the supply of first-party data will likely decrease. A 2010 study by the European Commission (EC) found that “89% of respondents agreed that they avoid disclosing their personal information online.”

Critics have noted, however, that a focus on user consent tends to mask greater imbalances in the bargaining power between users and service providers. The strong network effects of certain

online services – and the costs to users of foregoing those services – may undermine the extent to which consent is ‘freely’ given. Further, absent restrictions on practices like ‘dark patterns’ (which manipulate user interfaces to steer users towards consenting to settings and practices advantageous to the platform), an emphasis on user consent may be naïve.

Statutory determination that so-called ‘dark patterns’ are unfair and deceptive trade

practices – Dark patterns are user interfaces that have been intentionally designed to sway (or trick) users towards taking actions they would otherwise not take under effective, informed consent. Often, these interfaces exploit the power of defaults – framing a user *choice* as agreeing with a skewed default option (which benefits the service provider) and minimizing alternative options available to the user. A vivid example of this practice is below, where Facebook deceptively prods users into consenting to upload their phone contacts to Facebook (something highly lucrative to Facebook in tracking a user’s ‘social graph’):



(First screen presented to users)

(Second screen presented to users upon clicking ‘Learn More’)

The first screen gives the false impression that there is only one option (“OK”), with a bouncing arrow below the “OK” option pushing users towards consent. If users click “Learn More” (which is the path towards *declining* consent) they’re presented with yet another deceptively-designed

interface – where the opt-*in* is highlighted (despite getting to this screen by not opting-in on the first screen), and the opt-out option is in smaller font, positioned at the bottom of the screen, and not highlighted with a blue button. The FTC Act could be updated to define these kinds of practices – which are based on design tricks to exploit human psychology – as *per se* unfair and deceptive.

One drawback of codifying this prohibition in statute is that the law may be slow to address novel forms of these practices not anticipated by drafters. To address this, the FTC could be given rulemaking authority to ensure that the law keeps pace with business practices.

Algorithmic auditability/fairness – The federal government could set mandatory standards for algorithms to be auditable – both so that the outputs of algorithms are evaluated for efficacy/fairness (i.e. were you justifiably rejected for a mortgage based on the defined factors?) as well as for potential hidden bias. This could be established for algorithms and AI-based systems used for *specific functions* (like eligibility for credit, employment, and housing opportunities). For instance, Finland recently passed a law prohibiting the “discriminatory use” of artificial intelligence in decisions about financial credit. Or it could be established based on magnitude (in other words, an algorithmic system that covers over 200M people). Under GDPR, users have numerous rights related to automated decision-making, particularly if those processes have legal or significant effects. These include furnishing individuals with information about the automated decision-making process, providing ways for the consumer to request human intervention in the process (or a challenge of the automated process that is adjudicated by a human), as well as regular audits of the automated decision-making process to ensure it is working as intended. A first step towards this (something that could, for instance, be inserted into the annual National Defense Authorization Act) would be to require that any algorithmic decision-making product the government buys must satisfy algorithmic auditability standards delegated to NIST to develop.

More broadly, a mandate would require service providers to provide consumers with the sources of data used to make algorithmic determinations or classifications. Service providers would also need to furnish consumers with information on the recipients of that data or

determinations/classifications, while also establishing processes by which consumers can correct or amend erroneous data.

Critics of this approach will likely argue that many methods of machine learning produce outputs that cannot be precisely explained and that requiring explainability will come at the cost of computational efficiency – or, because the outputs of machine learning-based systems are not strictly deterministic, explainability is not feasible. Particularly in the context of employment, credit, and housing opportunities, however, a degree of computational inefficiency seems an acceptable cost to promote greater fairness, auditability, and transparency. Moreover, while complete algorithmic transparency may not be feasible or preferable, a range of tools and techniques exist to determine whether algorithms align with key values, objectives, and legal rules.¹⁹

Competition

Data Transparency Bill – The opacity of the platforms’ collection and use of personal data serves as a major obstacle to agencies like the FTC addressing competitive (or consumer) harms. This lack of transparency is also an impediment to consumers ‘voting with their wallets’ and moving to competing services that either protect their privacy better or better compensate them for uses of their data.²⁰ One of the major problems identified among dominant platforms, for instance, is that the terms of the bargain – free services in exchange for access to consumer data – continue to be amended in favor of the platform. Google’s free email service, for instance, was once predicated on the idea that users received a free email service in exchange for Google using the email data for more targeted ads. Increasingly, however, Google has found *other* uses for this data, beyond the terms of the original deal. Similarly, Facebook has made consumers agree to give up additional data as a condition for using its messaging service on their smartphones: whereas previously they could use the messaging feature through their web browsers, Facebook later made them download a dedicated ‘Messenger’ app that collects considerably more data.

¹⁹ Joshua A. Kroll et al., *Accountable Algorithms*, 165 U. Pa. L. Rev. 633 (2017).

²⁰ Maurice E. Stucke & Allen P. Grunes, *Big Data and Competition Policy* (Oxford: Oxford University Press, 2016), 333.

Many observers have noted that the collection and use of data is at odds with consumer expectations. Legislation could require companies to more granularly (and continuously) alert consumers to the ways in which their data was being used, counterparties it was being shared with, and (perhaps most importantly) what each user's data was *worth* to the platform. Requiring that 'free' platforms provide users with an annual estimate of what their data was worth to the platform would provide significant 'price' transparency, educating consumers on the true value of their data – and potentially attracting new competitors whose services (and data collection/use policies) the consumer could evaluate against existing services. Lastly, data transparency would also assist antitrust enforcement agencies like the FTC and DOJ by providing concrete and granular metrics on how much value data provides a given company, allowing enforcement agencies to identify (particularly retrospectively) anticompetitive transactions as ones that significantly *increase* the value a company extracts from users (which in the data-centric markets is equivalent to a price increase).

Data Portability Bill – As platforms grow in size and scope, network effects and lock-in effects increase; consumers face diminished incentives to contract with new providers, particularly if they have to once again provide a full set of data to access desired functions. The goal of data portability is to reduce consumer switching costs between digital services (whose efficiency and customization depends on user data). The adoption of a local number portability requirement by Congress in the Telecommunications Act of 1996 had a substantial procompetitive effect, particularly in the mobile market by facilitating competitive switching by customers. A data portability requirement would be predicated on a legal recognition that data supplied by (or generated from) users (or user activity) is the *users'* – not the service provider's. In other words, users would be endowed with property rights to their data. This approach is already taken in Europe (under GDPR, service providers must provide data, free of charge, in structured, commonly-used, machine-readable format) but a robust data ownership proposal might garner pushback in the U.S. More modestly, a requirement that consumers be permitted to port/transfer their data – in structured data, machine-readable format – without addressing the underlying ownership issue, would be more feasible.

One potential complexity in establishing a data portability requirement is whether to extend it to *observed* data collected by a service provider. In one sense, this data is data *about* the user, derived from the activity of the user. Service providers, however, are likely to claim that observed data – for instance, classifications or generalizations about a user based on observed activity – belong to the service provider. Service providers may even invoke 1st amendment protections against sharing that data – which they may characterize as compelled commercial speech – with third parties.

Additionally, data portability can pose a number of cybersecurity risks if not implemented correctly. Specifically, it increases attack surface by enlarging the number of sources for attackers to siphon user data; further, if the mechanism by which data is ported (typically an API) is not implemented correctly, unauthorized parties could use it to access data under the guise of portability requests.

There is also a risk that, if not devised appropriately, data portability could be used by dominant providers *to the detriment* of smaller, emerging providers. Large providers are best-positioned to offer incentives to users to submit portability requests to new entrants who may pose a competitive threat to them. Smaller providers also may have less ability to process portability requests, and less ability to implement portability mechanisms securely. For this reason, any portability mandate should ideally be imposed on providers *above* a certain size, or who have been determined to hold dominant positions in particular markets.

Interoperability – Imposing an interoperability requirement on dominant platforms to blunt their ability to leverage their dominance over one market or feature into complementary or adjacent markets or products could be a powerful catalyst of competition in digital markets. More importantly, an interoperability requirement acknowledges that in some contexts – for instance, where network effects are so pronounced, or where it would be uneconomical for a new platform to radically reinvent key functions provided by a dominant incumbent– data portability alone will not produce procompetitive outcomes. For instance, allowing messaging or photo-sharing startups access to the ‘social graph’ of Facebook would allow users to communicate more broadly without a new startup having to (unfeasibly and uneconomically) recreate an

entirely new Facebook. A prominent template for this was in the AOL/Time Warner merger, where the FCC identified instant messaging as the ‘killer app’ – the app so popular and dominant that it would drive consumers to continue to pay for AOL service despite the existence of more innovative and efficient email and internet connectivity services. To address this, the FCC required AOL to make its instant messaging service (AIM, which also included a social graph) interoperable with at least one rival immediately and with two other rivals within 6 months. Another example was the FTC’s interoperability decrees with respect to Intel’s treatment of NVIDIA.

Interoperability is seen as falling within the “existing toolkit” regulators have to address a dominant platform; observers have noted that “Regional Bell Operating Company” (RBOC) interoperability with long distance carriers actually worked quite well. Some experts have expressed concern with the managed interoperability approach, suggesting it might create too cozy a relationship between regulatory agencies and the platforms. However, a tailored interoperability requirement may not pose the same regulatory capture concerns. Interoperability could be achieved by mandating that dominant platforms maintain APIs for third party access. Anticipating platforms’ counter-arguments that fully open APIs could invite abuse, the requirement could be that platforms maintain transparent, third-party accessible APIs under terms that are fair, reasonable, and non-discriminatory (FRAND).

As with data portability, security experts have observed that interoperability could increase the attack surface of any given platform. Implementing APIs securely can be difficult for even mature providers; for instance, it was a weakness in Apple’s iCloud API (allowing attackers to make unlimited attempts at guessing victims’ passwords) that contributed to the 2014 hacks of major celebrities’ photos.

Opening federal datasets to university researchers and qualified small businesses/startups

– Structured data is increasingly the single most important economic input in information markets, allowing for more targeted and relevant advertisements, facilitating refinement of services to make them more engaging and efficient, and providing the basis for any machine-learning algorithms (which develop decisional rules based on pattern-matching in large sets of

training data) on which all industries will increasingly rely. Large platforms have successfully built lucrative datasets by mining consumer data over significant timescales, and separately through buying smaller companies that have unique datasets. For startups and researchers, however, access to large datasets increasingly represents the largest barrier to innovation – so much so that university researchers are steadily leaving academia not only for higher *salaries* but also for access to unrivalled or unique datasets to continue their work. The federal government, across many different agencies, maintains some of the most sought-after data in many different fields such that even the largest platforms are pushing the Trump Administration to open this data to them. To catalyze and sustain long-term competition, however, Congress could ensure that this data be provided *only* to university researchers and qualified small businesses, with contractual prohibitions on sharing this data with companies above a certain size. Numerous precedents already exist for government contractual agreements only with smaller or non-commercial entities (e.g. procurement).

Essential Facilities Determinations – Certain technologies serve as critical, enabling inputs to wider technology ecosystems, such that control over them can be leveraged by a dominant provider to extract unfair terms from, or otherwise disadvantage, third parties. For instance, Google Maps maintains a dominant position in digital mapping (enhanced by its purchase of Waze), serving as the key mapping technology behind millions of third party applications (mobile and desktop) and enabling Google to extract preferential terms and conditions (such as getting lucrative in-app user data from the third-party apps as a condition of using the Maps function). Legislation could define thresholds – for instance, user base size, market share, or level of dependence of wider ecosystems – beyond which certain core functions/platforms/apps would constitute ‘essential facilities’, requiring a platform to provide third party access on fair, reasonable and non-discriminatory (FRAND) terms and preventing platforms from engaging in self-dealing or preferential conduct. In other words, the law would not mandate that a dominant provider offer the service for *free*; rather, it would be required to offer it on reasonable and non-discriminatory terms (including, potentially, requiring that the platform not give itself better terms than it gives third parties). Examples of this kind of condition are rife in areas such as telecommunication regulation, where similar conditions have been imposed on how Comcast’s NBC-Universal subsidiary engages with Comcast and Comcast rivals.