Approaches to Regulating Technology—From Privacy to A.I.

March 15th, 2019

AMERICAN ACTION FORUM

Edited by
Will Rinehart
Director of Technology and Innovation Policy

Published by the American Action Forum

Table of Contents

For citations, see americanactionforum.org/series/regulating-technology.

Introduction: How to Think About Tech Regulation

Will Rinehart, Director of Technology and Innovation Policy at the American Action Forum

If there were ever a year that changed the course of technology regulation, 2018 was it. In mid-March, the *New York Times* released a report detailing how the marketing firm Cambridge Analytica exploited Facebook's consumers leading up to the 2016 presidential campaign, sparking congressional hearings with Facebook CEO Mark Zuckerberg. At the end of May, the long-awaited European General Data Protection Regulation came into effect. On June 25th, the Supreme Court issued a major ruling in *Ohio v. American Express*, the first time the Court had directly addressed the economics of platforms, opening up the possibility for companies such as Facebook and Google to face scrutiny by courts in the future. On June 28th, the California Consumer Privacy Act (CCPA) was introduced and passed within a week.

The shifts in the legal and political environments could translate into this Congress legislating new federal regulatory policy for technology. Even the presidential campaigns could encourage action, as presidential hopefuls such as Senator Amy Klobuchar and Senator Elizabeth Warren have made tech regulation a central part of their platforms. But how are policymakers to tackle these difficult and broad topic areas?

The essays that follow are meant to help clarify the goals and implications of tech regulation. Since this topic is so broad, the essays are similarly expansive, but what unites them is their focus on innovation. If innovation is the primary challenge facing effective tech regulation, then it is critical to understand how legislative proposals either encourage or inhibit it.

The first essay comes from Ryan Hagemann of the Niskanen Center. Many assume that the best way for the government to address new technologies is through proactive and specific federal legislation. In contrast, Hagemann makes a case for a "soft law" approach. As he details, a number of agencies have worked alongside industry, advocates, and the research community to produce de facto governance through green papers, advisory circulars, guidance documents, and a range of other materials. This multi-stakeholder approach has been especially effective for fast-developing markets, such as autonomous vehicles, medical devices, and the Internet of Things. As policymakers consider legislation to deal with new and rapidly evolving technologies, a soft-law approach that supplements or even substitutes for statues should be the priority, especially if the goal is to ensure continued innovation.

In the second essay, Jennifer Huddleston of the Mercatus Center lays out the state of play in privacy regulation. California may have been the first to adopt a privacy law, passing one in 2018, but other states are following closely. As she explains, there will surely be problems if all 50 states pass some form of privacy legislation. The California legislation will not go into effect until 2020, giving Congress time to debate a federal privacy bill that preempts it. The only real solution to this legislative confusion will be a federal law that allows for innovation while also protecting consumers.

While many have turned their attention to privacy law, Ryan Radia of Lincoln Network explores how the broad consensus around antitrust regulation has similarly shifted. For decades, there was a bipartisan and empirically grounded agreement that antitrust laws should focus on consumer welfare. More recently, however, a "big is bad" approach to regulation is gaining ground as a solution to nearly every problem with tech companies. Radia argues that a change in course for antitrust would come at a steep price for consumers while not achieving the lofty goals of antitrust advocates.

But Congress and the agencies aren't the only places where the action is happening. The Supreme Court has tackled some tough tech issues as of late. As Ashley Baker of The Committee for Justice explains, the highest court has considered several tech-related issues, but some recent decisions leave many unanswered questions that Congress will need to address with legislation.

It is hard to have a conversation about technology today without mentioning artificial intelligence (A.I.), and in my own contribution to this volume, I highlight the areas where A.I. is having the biggest impact and provide a simple framework for understanding the technology. Countries and firms across the globe are racing to capitalize on this new tech, and policymakers should recognize that premature regulation is likely to stifle innovation and progress in A.I.

If it is the case that A.I. will revolutionize how we work, live, and connect, it will be imperative that government support this transition. In the final essay, Caleb Watney of the R Street Institute outlines some concrete steps policymakers can take to encourage the development and adoption of A.I. in the economy. To this end, he suggests passing immigration reform, encouraging the creation of open datasets, and avoiding political instability to international supply chains, among other proposals. Ensuring the United States remains a world leader in technology innovation and use may not require regulation, per se, but it will require proactive government engagement.

Watney's closing sentences reflect just how connected these seemingly disparate issues are. As he explains, "To ensure a competitive and innovative ecosystem going forward, policymakers should prioritize reducing the barriers to entry as our first line of defense." Indeed, the agenda Watney lays out to spark A.I. adoption could just as easily be an agenda to increase competition writ large. Technology policy in 2019 has become narrowly stuck on the old methods of regulation and enforcement, the sticks of government. But policy makers shouldn't deny how powerful carrots can be in ensuring better outcomes for all.

For years, the United States was the envy of the world because our light-touch regulatory regime created a dynamic and innovative ecosystem for today's tech companies, yet in the past year, the regulatory climate around tech and tech companies has notably changed. While shifting sentiments offer an opportunity to pass sensible laws, policymakers should be cautious about redesigning the entire system. The following essays should help to ground the discussion around how, and whether, to regulate technology in the United States.

What Policymakers Need to Know About "Soft Law"

Ryan Hagemann, Senior Fellow at the Niskanen Center

Many grade schoolers' first introduction to the law begins with Hammurabi, the Babylonian king renowned for his codification of the law in the Code of Hammurabi. The stele into which those legal tenets were carved was notably displayed in the public square so that all might know the law — and the punishments to be borne by those who violated it. But for all its notoriety in creating legal certainty, the *type* of certainty Hammurabi's laws established was harsh, exacting, and foundationally retaliatory. This is one of the core tensions of any social structure: providing for a legal regime that balances the need for certainty and stability against rules that offer a framework of reasonable and fair penalties for infractions.

Moving forward some 1500 years, those same tensions were on full display in the declining age of Republican Rome. Although the Romans are often celebrated as the great law-exporters to the wider Western world, the certainty with which those laws were applied was largely a function of informal adherence to traditional expectations. As Mike Duncan notes in his book *Storm Before the Storm*:

> What truly bound all Romans together, though, were unspoken rules of social and political conduct. The Romans never had a written constitution or extensive body of written law — they needed neither. Instead the Romans surrounded themselves with unwritten rules, traditions, and mutual expectations collectively known as mos maiorum, which meant "the way of the elders." Even as political rivals competed for wealth and power, their shared respect for the strength of the client-patron relationship, the sovereignty of the Assemblies, and wisdom of the Senate kept them from going too far. When the Republic began to break down in the late second century it was not the letter of Roman law that eroded, but respect for the mutually accepted bonds of mos maiorum.

Throughout history, legal regimes have been defined along this spectrum of certainty — from the informally enforceable and evolutionary rules of mos maiorum, to the unambiguous and punitively draconian commandments of Hammurabi's Code. For many decades, the American system of administrative law occupied a space somewhere between these two extremes. In recent years, however, the rapid pace of technological change has significantly outpaced the ability of regulators and bureaucrats to adjust to the changing expectations of an increasingly digital, interconnected world. The result has been a paradigmatic shift in administrative law practices in autonomous vehicles, commercial drones, the Internet of Things, and advanced medical technologies that has come to rely more and more on a modern variation of the Roman mos maiorum: soft law.

Collaborative Regulatory Governance

As legal scholars Gary Marchant and Braden Allenby describe the term, soft law is a set of "instruments or arrangements that create substantive expectations that are not directly

enforceable, unlike 'hard law' requirements such as treaties and statutes." In a forthcoming article in the *Colorado Technology Law Journal*, my co-authors — Mercatus Center senior research fellow Adam Thierer and legal fellow Jennifer Huddleston — and I expand on this definition by identifying and categorizing the specific outputs of the soft-law system, which we call "soft criteria." As we note in our law journal article:

> If soft law is generally defined as the implementation of those "arrangements that create substantive expectations that are not directly enforceable," then "soft criteria" refers to the corpus of "nonbinding norms and techniques" that serve as the instruments of soft law's implementation. In short, soft criteria are the means by which the soft law end is achieved — a skeletal structure that provides a governance foundation that can be built upon.

This "corpus of 'nonbinding norms and techniques'" includes things such as green papers, advisory circulars, guidance documents, interpretive rules, policy statements, opinion letters, voluntary standards, best practices, and much more. These deliverables can emanate from many different quarters of society, from industry consortia and trade associations to academic centers and think tanks. However, the essential feature of the soft-law system that legitimizes the "substantive expectations" established by these soft criteria is the multistakeholder process.

Multistakeholderism is a governance process by which a set of soft criteria (or other objectively measurable outcome or deliverable) is produced or reviewed and then legitimized (or discarded) via a deliberative, consensus-based dialogue involving a broad range of actors from government, civil society, industry, academia, and elsewhere. In developing governance responses to the challenges posed by emerging technologies over the past two decades, these proceedings have usually been convened by federal agencies — in particular, the National Telecommunications and Information Administration, which has become something of a general clearinghouse for multistakeholder meetings focusing on new technologies. This governance approach has recently come under fire from both right-of-center and left-of-center ideological quarters: Conservatives lament soft law's role in expanding the deference afforded to administrative agencies; Progressives decry its failure to both guard against self-regulatory excess and inability to actualize their preferred policy priorities. While the concerns regarding these informal approaches to rulemaking are not to be taken lightly, there are also many benefits of a more collaborative rulemaking process — especially for emerging technologies.

As Ian Ayres and John Braithwaite articulated in their 1992 book *Responsive Regulation: Transcending the Deregulation Debate*, these types of deliberative proceedings possess three distinct benefits over more traditional regulatory rulemaking processes:

> First, it grants the [public interest group] and all its members access to all the information that is available to the regulator. Second, it gives the [public interest group] a seat at the negotiating table with the firm and the agency when deals are

> done. Third, the policy grants the [public interest group] the same standing to sue or prosecute under the regulatory statute as the regulator.

Opening the door to more substantive engagement with both civil society and industry, agency-convened multistakeholder processes, Ayres and Braithwaite note, can "produce more efficient regulatory outcomes because bad arguments and bad solutions are less likely to go unchallenged. And genuine communication means that when challenges are advanced, they are listened to." These processes are especially adept for dealing with fast-paced technologies.

All Roads Lead to Soft Law

In many ways, these soft law pathways are an inevitable byproduct of the modern bureaucratic state. As the German sociologist Max Weber noted in *The Theory of Social and Economic Organization*, all socio-economic systems of political organization inevitably trend toward some degree of bureaucratization. "The question," Weber notes, "is always who controls the existing bureaucratic machinery. And such control is possible only in a very limited degree to persons who are not technical specialists."

However, nothing preordains that a society's "bureaucratic machinery" must be centralized under the sole control and oversight of a state authority. Weber argues that bureaucratic administration plays the "crucial role in our society as the central element in any kind of large-scale administration," conceding only "by reversion in every field — political, religious, economic, etc. — to small-scale organization would it be possible to any considerable extent to escape [the bureaucratic machinery's] influence." But in many ways, soft law is creating the space within the administrative state for these small-scale organizational structures (i.e., multistakeholder governance proceedings) to flourish. Although he didn't explicitly foretell the emergence of soft-law systems, Weber did recognize that one class of individuals within society was more likely than others to escape the gravitational pull of a centralized bureaucracy:

> The capitalistic entrepreneur is … the only type who has been able to maintain at least relative immunity from subjection to the control of rational bureaucratic knowledge. All the rest of the population have tended to be organized in large-scale corporate groups which are inevitably subject to bureaucratic control. This is as inevitable as the dominance of precision machinery in the mass production of goods.

This makes sense, as bureaucratic cultures tend to focus on rules and processes that address known problems created by emergent industries. In contrast, emerging innovations that lack historical precedent confound administrative agencies with missions and objectives predicated on the application of institutional knowledge to observed market failures. In other words, bureaucrats, like generals, are always fighting the last war.

In an age of rapid advancements in technology, the limitations of regulators' knowledge are perhaps more apparent than ever. As such, soft law offers an ideal compromise: it allows agencies to more effectively balance their statutory missions to protect the public interest without imposing undue hardships on the entrepreneurs whose work helps drive economic growth and societal well-being. By leaning more heavily on the use of informal rules and exporting authority to self-regulatory governance mechanisms, regulators can capture the benefits of a more responsive, flexible, and adaptive governance culture without sacrificing their statutory oversight functions. In *The Promise and Pitfalls of Co-Regulation: How Government Can Draw on Private Governance for Public Purpose*, Edward J. Balleisen and Marc Eisner echo similar sentiments:

> Legislators and administrative agencies should view nongovernmental regulation as a policy instrument that can make sense in many, if by no means, regulatory contexts. The key challenge is to design systems that provide that benefits of self-governance without sacrificing the high levels of accountability that one expects from public regulation.

Autonomous Vehicles: An Application

The development of the Department of Transportation's (DOT) guidance on self-driving cars, *Preparing for the Future of Transportation: Automated Vehicles 3.0* (commonly known as AV Guidelines 3.0), best exemplifies the soft-law approach in action. In just a few short years, autonomous vehicles have been developed and deployed. Large car manufacturers have entered the market, often working in conjunction with tech upstarts to bring this technology to market. While there have been setbacks, development quickly outpaced the ability of legislators to make rules. The DOT has managed this change by relying on multistakeholder processes and agency workshops that include working groups, researchers, state and local actors, industry specialists, advocates, and policy experts. At the core of this process has been the creation of a flexible regulatory framework, the Guidelines, that encourages entrepreneurship while still maintaining regulatory clarity for everyone involved. In the process, the Guidelines have helped to avoid a state or federal patchwork of regulations while also promoting safety.

On the whole, soft law tends to strike a reasonable balance between the flexibility and adaptability of a self-governance regime with the legitimizing power of administrative oversight ensuring a backstop against egregious excesses that self-regulation may fail to effectively address. Policymakers confronting the realities of rapid technological progress should take notice of soft law, and seriously consider how they can take advantage of its benefits.

Preventing Privacy Policy From Becoming a Series of Unfortunate Events

Jennifer Huddleston, Research Fellow at the Mercatus Center

It seems that every day new headlines are calling our attention to the "growing problem" of data privacy. From the Equifax breach to a series of Facebook scandals to new laws in Europe and California, there seems to be a growing chorus advocating that "something must be done." Unfortunately, these calls often neglect the full spillover effects of increasing compliance burdens on innovation and free speech. Further, they also often neglect the existing tools available both to consumers and the government regarding data privacy.

The Present Position of Public Policy on Privacy

To begin, let's examine the current privacy laws in the United States. Contrary to what critics may claim, the United States is not completely a "Wild West" when it comes to data privacy. The lack of comprehensive privacy policy rather has recognized that choices regarding privacy and transactions involving data primarily exist as an individual preference between consumers and companies. Nevertheless, there are still specific rules when necessary to stop harm or protect those individuals or data that are particularly vulnerable and/or significantly more likely to result in real, cognizable harm if breached.

For example, Congress recognized that children under the age of 13 needed special protections regarding data collection. The Children's Online Privacy Protection Act (COPPA) requires verifiable parental consent for the collection of personal information collected online and is enforced and interpreted by the Federal Trade Commission (FTC). Other laws have carved out additional data privacy protections for particularly sensitive information such as medical, financial, and credit data held by certain entities, as well as information that is more likely to reveal information an individual would likely consider not to be public such as data gathered while driving and personal video rental history. Rather than taking a restrictive but comprehensive approach to data privacy that might have prevented many of the most popular uses of the Internet, such as tailored song recommendations or social media, these laws have served to carve out particular types of data that would most likely result in real, cognizable harm to individuals if a breach were to occur.

But actions not covered by these narrow categories are not entirely without recourse either. If the privacy harm in question is harming consumer welfare, the FTC has broad authority to go after the companies using its unfair and deceptive practices authority. As Neil Chilson points out, in this role the FTC has brought privacy actions against companies regarding data privacy "where consumers are substantially injured, could not reasonably avoid the injury, and this injury isn't outweighed by benefits to consumers or competition." This approach focused on consumer welfare has tended to preserve a broad array of options to reflect individual preferences regarding data privacy and allow individual consumers and companies to select standards that fit these preferences. Many states

have given similar broad authority to their consumer protection authorities that would allow them to go after bad actors if there was harm to consumer welfare. The consumer welfare standard used by the FTC limits the ability of this broad power to be abused for just a sense of "creepiness" and rather limits its pursuit of such incidents ideally to only those where harm has occurred or is most likely to occur.

It is important to remember that most if not all of these interactions are voluntary and in exchange for services and not mandated by government authorities.

Privacy Paradox Problems

When news of a data breach or a perceived abuse breaks, there is an uptick in popular concern about the state of online privacy. This interest, however, appears to be immediate reactions to specific events and not an actual desire for change or willingness to make tradeoffs. For example, a 2018 NetChoice survey found less than 1 percent of younger adults and teenagers chose to leave a platform over changes in a privacy policy.

Much of this type of data has been labeled the "privacy paradox." The privacy paradox is the distinction between expressed and revealed preferences for data privacy online. Individuals may state that they value data privacy, but their actions and further preferences reveal that they are willing tolerate little that would decrease efficiency or increase costs to achieve it. Many complain or are unwilling to deal with privacy measures that would increase friction of usage or change the user experience and even fewer are willing to pay for actual measures. It is important to note that in many cases additional products and solutions are available to those whose expressed preferences for data privacy do align with such willingness to take action.

Furthermore, many consumers willingly participate in the data-driven marketplace and appreciate its benefits. For example, a 2017 Accenture study found the majority of consumers would be willing to share financial data in exchange for service benefits such as investment advice, insurance, and banking. When surveyed, few adults say that concerns about privacy are significant enough to make them willing to pay for access to online platforms rather than exchange data for a free, ad-supported version. Policymakers should carefully consider that in choosing to prioritize privacy over consumer preferences they may be making choices that consumers aren't actually looking for. As Chris Koopman points out, "In truth, it seems that consumers have a pretty good idea that their personal information is being collected but seem to value it less than what they're getting in exchange from the platforms."

The Risk of a Privacy Patchwork Problem

Perhaps the more pressing privacy public policy problem is not the lack of privacy regulation, but the potential pitfalls of the growing patchwork of policies aimed at privacy issues. The United States has traditionally embraced a permissionless approach to the Internet, including privacy policies,

and has allowed individual online communities and individual users to develop their own preferences via a system of notice and choice. Any legislative intervention into the existing paradigm should be as narrowly tailored as possible to preserve allowing consumers a wide array of options to select personal preferences, and only intervening when consumer welfare is being harmed. Data do not obey borders, and increasingly state and global regulations risk changing the innovation-friendly framework that allowed the Internet to succeed and creating a quagmire of conflicting regulations.

Notably, the European Union (EU) has taken a very different approach to data privacy, first in establishing a "right to be forgotten" and more recently through the General Data Privacy Regulation (GDPR). These requirements have come at substantial costs to the firms operating in both the financial expenses of compliance as well as hiring additional employees to focus on the issue and dedicating time to insuring compliance rather than innovation. In some cases, faced with such burdens, small- and mid-size firms have chosen to exit the market rather than undertake the compliance burden and costs. As a result, the largest market players that could afford to comply, such as Google, have increased their market share and consumers have fewer choices.

GDPR and privacy laws do not just impact those traditional data collectors such as social media and search engines, but also many online entities like personal blogs or small businesses that do not have large compliance teams at their disposal. Several U.S. newspapers are still not available in the EU due to not being GDPR compliant, and other websites and services from video games to a church website collecting prayer requests have geofenced EU users or stopped offering certain features. Even large tech companies may be rethinking potentially beneficial additions in light of these new standards. For example, GDPR likely prevented Facebook from launching a suicide prevention alert in Europe.

Given the amount of time and money many American tech companies invested into GDPR compliance, there is concern that it — rather than the more permissionless American approach — could emerge as a global default. GDPR style laws value privacy over innovation and put privacy choices in the hands of regulators rather than individual consumers. A shift to such a default could prevent the development of better options including those focused on data privacy due to the increased cost of compliance and the regulatory burdens. As a result, small players may never be able to grow big enough to offer alternatives to current giants.

Not only internationally are new laws such as the GDPR creating a complicated and potentially innovation-limiting framework for data privacy, domestically state laws are creating problems for data privacy policy. Notably the California Consumer Privacy Act (CCPA) was passed in August 2018 and would create a wide variety of requirements that would increase compliance costs and dissuade new innovators online in a very short period of time. While it may be an option to segment off and provide different defaults and services to comply with GDPR, the burden of doing so on a state-by-state level would not be sustainable. Additionally, the California law's definition of covered entities could be interpreted broadly enough to cover any data if the company has even a single California resident access the website — whether from California or elsewhere. Others may see this

as a way to follow suit creating a problem in which it is difficult if not impossible to launch a universal product available in all communities. New York City recently stated in a comment filing that it intended to similarly pass its own data privacy law regardless of if a federal regulation occurred. Much like the GDPR, these state and local laws may further consolidate business in the large players who can afford to devote resources to complying with them and — because of the relatively lower number of users at which such requirements kick in — discourage current small players from trying to grow large enough to challenge existing giants. As a result, such state laws would deter innovation and competition.

With the rise of the Internet as a tool for commerce and the borderless nature of data flows, it seems the dormant commerce clause would likely render state or local privacy laws such as the CCPA unconstitutional. Similar issues are arising in the courts regarding state-level net neutrality laws. Yet because of the length of time it would take businesses to engage in necessary steps for compliance and the funds that would need to be allocated to such, some degree of damage could be done even if such laws are later struck down by the courts. If such laws were not struck down, it could create a scenario of conflicting requirements or one where the most restrictive approaches or largest states win for fear of noncompliance or losing the most lucrative markets.

While the California law is the most expansive, it is not the first state to carve out specific privacy regulations for its citizens, although prior attempts have been much more limited in both the data and entities covered. Perhaps the other most relevant example of such deviations from the general notice-and-choice process is Illinois' Biometric Information Privacy Act (BIPA). But even this relatively specific policy had consequences that prevented Illinois residents from using certain services. For example, many Illinoisans were surprised to find they could not use Google's Arts & Culture app due to the law. Social media companies such as Facebook that use facial recognition to help users identify themselves or friends in photos have found themselves subject to class action lawsuits and hefty fines under the law. Offline it could prevent the use of biometric security because of consent requirements such as the recent lawsuit involving Six Flags amusement park and its collection of annual passholders' fingerprints. If left unchecked, such litigation could redefine harm so broadly as to undermine the sharing of information at the heart of many online communities or limit the technologies that companies are able to offer consumers. Texas and Washington passed similar but less restrictive laws regarding biometric data, and as a result have not experienced the same degree of loss of opportunity.

State data breach laws provide an example of what a 50-state patchwork could look like and even this relatively narrow privacy issue shows the potential compliance quagmire that could emerge if state policies pollute the current privacy frameworks. These laws vary in the amount of time in which users whose data have been breached must be notified and even what data are covered. As a result, nationwide companies may find the most restrictive laws to be de facto national regulations rather than risk noncompliance and must determine how best to deal with seemingly contradictory requirements.

It quickly becomes apparent that state governments are not the proper actor for data privacy policymaking, and if any reforms are necessary they must occur at the federal level. Recent actions suggest that when they are not preempted from doing so, states may try to pursue regulatory actions that could eliminate certain consumer choices and, potentially, fundamentally change the Internet.

Preemption and Potential Privacy Public Policy Solutions

Perhaps the most obvious solution to this perplexing problem is federal preemption that retains the current framework of notice-and-choice for most privacy issues, continues case-by-case examination of consumer harms, and resolves current and potential data-privacy public policy patchworks caused by state and local interventions.

The United States has recognized since the mid-1990s that the Internet is an important tool for innovation, communication, and commerce. Now over 20 years after this initial framework, this overall permissionless framework that allowed the Internet to evolve faces significant challenges from a variety of policies on the state and federal level.

Policymakers should carefully consider the tradeoffs associated with any data privacy policies and if such actions will eliminate benefits and increase the burdens to providing consumers better and more innovative choices. One way to prevent or eliminate the growing patchwork problem would be to expressly preempt state laws regarding data. This move would prevent laws such as the CCPA from becoming de facto national regulations for millions and likely undo existing patchwork problems related to conflicting data breach laws or policies such as the BIPA.

Preemption of state policies that would disrupt the current system would not mean that a federal policy such as a U.S. GDPR is necessary. Such preemption could provide the opportunity to reiterate the U.S. emphasis on a permissionless approach to Internet policy, including notice-and-choice regarding data privacy and delegation of authority on a case-by-case basis regarding consumer harm to the FTC. Any policy should maintain the current neutrality in the U.S. system, which allows new entrants to emerge that may offer consumers more and better options and allows individuals to make choices regarding data privacy that reflect their own preferences.

Perhaps the best policy solution is not to change policy regarding data privacy at all, but rather for the government to increase efforts regarding consumer education around privacy choices and to continue to have the FTC examine the specifics of instances where consumer welfare is truly harmed. The government's role as educator in such debates regarding possible data privacy solutions is often neglected. By informing consumers of steps to take if they seek to improve data privacy and what options are available to them in the case of specific harms, the government can empower individual citizens and their families to make the choices that are right for them and would also encourage new products to enter the market and provide consumers the options they desire.

The True Privacy Problem

The presumption that data privacy is broken is based on individual incidents such as breaches rather than real harm to consumers or competition. This presumption could lead to tradeoffs that impact innovation and remove choices that consumers enjoy and value above their privacy. Many of the proposed solutions do not align either with consumer actions, expressed choices, or the reality of innovation. Perhaps the real privacy problem is knowing how not to start down a slippery slope that could lead to diminishing or dismantling many of the technologies that have revolutionized our lives, and that could prevent future innovation that brings even more benefits.

New Antitrust Thinking Isn't a Return to the Good Old Days

Ryan Radia, Senior Policy Counsel at Lincoln Network

Amazon, Apple, Facebook, Google, Microsoft. These are America's five most valuable tech companies, and they have a public policy challenge in common: antitrust. When antitrust makes headlines, it's often because one of these five companies is looking to buy a smaller firm or facing some sort of investigation from the Federal Trade Commission or Department of Justice. Although antitrust usually comes up in the context of a specific transaction or probe, a growing movement of activists and commentators is urging policymakers to rethink the framework by which antitrust regulators scrutinize how big businesses behave and when they can enter into mergers or acquisitions. This movement, deemed "neo-Brandeisian" for its adherence to Supreme Court Justice Louis Brandeis's fear of the "curse of bigness," seeks to restore the antitrust approach that prevailed in the mid-20th century of condemning market concentration and large businesses as illegal and harmful. America's top tech companies are squarely in the crosshairs of this proposed policy shift.

This effort doesn't seem particularly close to overtaking the long-prevailing view among policymakers that competition law should chiefly serve consumers, but it has won over some influential allies. When leading congressional Democrats unveiled their legislative agenda called "A Better Deal" in July 2017, a section that drew considerable attention called for reforming U.S. antitrust laws to crack down on "corporate monopolies." Several bills were soon introduced — unsuccessfully — to implement this proposal. Even some prominent right-leaning pundits, such as Fox News commentator Tucker Carlson, have emerged in favor of stricter antitrust intervention — though this stance might be driven less by a newfound affection for antitrust enforcement and more by ideological disagreement with the left-wing political views often espoused by senior leadership at big American tech firms.

The neo-Brandeisians reject the idea that the antitrust laws exist primarily to serve consumer well-being by empowering the government to block mergers, acquisitions, and business practices that tend to push up prices, reduce output, and undermine competition itself. Instead, these commentators want judges and antitrust regulators to dust off antitrust principles long ago abandoned by the courts, such as condemning bigness in American business as presumptively unhealthy and restricting economic concentration for the sake of "social and political goals." This approach rejects the approach prevalent since the 1970s among scholars and judges that a market's concentration doesn't tell us how the market will perform, and that big firms can sometimes deliver efficiencies that smaller firms cannot.

A Reality Check

Outside the world of punditry and policymaking, the American public appears relatively unconcerned about the scale of the country's leading tech companies. A summer 2018 survey

conducted by Georgetown and NYU researchers found that among 20 top public and private U.S. institutions, Google and Amazon "universally inspire a great deal of confidence." One notable outlier: Facebook, which placed near the bottom of the list among Democrats and Republicans alike. But this skepticism about Facebook likely has more to do with the social media platform's highly publicized recent snafus involving user privacy and foreign election interference than the firm's competitive practices or acquisition history. (Ironically, Facebook's efforts to address privacy fears by greatly restricting third-party applications' access to its application programming interface (API) have made it harder for users to seamlessly share content across social media platforms.)

America's leading technology firms enjoy impressive market valuations, with the nation's five most valuable tech firms worth a combined $3.6 trillion — or roughly 17 percent of the 500 companies listed on the S&P 500 index. This market signal indicates that investors are, by and large, confident that Alphabet (Google's parent company), Amazon, and Microsoft will grow more profitable in coming years, while Facebook and Apple will maintain their enviable profits. But U.S. tech firms' record-breaking performance in the stock market should not obscure the immense size and scope of the rest of the nation's economy, which includes numerous companies with well-established brands, deep pockets, and a thirst for success in the digital world.

For the time being, American consumers and businesses still spend far more offline than online, even when it comes to retail purchases and ad buys. In the United States, for instance, Amazon's direct retail sales plus its third-party sellers' revenue amounted to about $200 billion in 2017 — an impressive figure, but still far behind Walmart's $308 billion in U.S. sales last year. Similarly, although Google and Facebook are expected to sell about $62 billion combined in U.S. digital ads this year, that still adds up to under one-third of the overall domestic advertising market.

These tech firms haven't stopped growing, to be sure: Amazon and its affiliates may outsell Walmart domestically in 2019. And Google and Facebook may end up raking in most U.S. advertising dollars within the next few years. But portraying companies like Amazon or Google as behemoths swallowing up the competition at every turn, monopolizing market after market, is a gross exaggeration.

Jet.com, an e-commerce site launched in 2015 with aspirations to take on Amazon by ditching annual fees, was acquired by Walmart for $3.3 billion in late 2016. Microsoft's Bing may not be taken too seriously among power users or industry analysts, but Microsoft has the incentive and ability to capitalize on any shortcomings at Google if it doesn't keep improving its search product. And despite Uber's new leadership and revamped marketing strategy, Lyft remains a popular — and innovative — alternative.

For critics of U.S. tech firms who want greater antitrust intervention, the rationale for such regulation seems to transcend the question of what's best for consumers. Instead, subsumed in the case for an antitrust crackdown is a laundry list of complaints about tech companies that spans a broad array of policy areas. From accusations of overrepresentation of Asian employees to insufficient responsiveness to sexual harassment complaints to too much (or too little) moderation

of user-generated content, the clear takeaway is that America's tech firms are engaged in lots of bad behavior.

Policymakers could address each of these issues in context, critically examining the evidence and the legal frameworks that underlie each area. Or, officials could avoid this complexity and instead take up antitrust intervention as a powerful cudgel to save the day for the many constituencies supposedly harmed by tech giants — and perhaps mete out some well-deserved punishment at the same time. To the populist progressives pining for a pound of flesh from big tech, the appeal of the latter approach is obvious. For consumers, however, the benefits of antitrust interventionism are dubious — while the downside is real.

Undoing decades of bipartisan, rigorous, empirically grounded efforts to rationalize U.S. antitrust laws and rekindling the "big is bad" approach to antitrust will undoubtedly alter how big tech firms behave. Indeed, this is the point of antitrust regulation. But this behavioral change could come at a steep price for consumers. To understand why, it's worth revisiting the tumultuous decade and a half in the smartphone marketplace.

Google and the Smartphone

Just 15 years ago, the smartphone was a niche product popular among busy executives, but its transformation into a must-have consumer device was imminent. Against this backdrop, Apple's decision to develop the iPhone wasn't all that surprising of a business decision, given the success of the iPod. More surprising, however, was Google's decision to purchase Android for a reported $50 million in 2005. Given that Nokia, BlackBerry, Microsoft, and Apple, among others, were investing heavily in the mobile marketplace, Google's foray into consumer devices was hardly a foregone conclusion. Following its Android acquisition, nearly five years would elapse before the platform even made a dent in the smartphone market. Fast forward to 2019, and over 80 percent of the world's smartphones are running Android.

Google's extraordinary success with Android is just one example of how consumers benefit when tech firms make risky bets entering adjacent markets. Even when these vertical gambits don't topple incumbent firms, they can make a difference. Consider Google Fiber, announced in 2010 as an effort to jumpstart the U.S. broadband sector by building fiber-to-the-home networks in cities across the nation. Nine years later, Google Fiber is available in over a dozen cities, but Google has paused network buildout in new cities. Still, although only a tiny fraction of Americans are served by Google Fiber, the initiative likely contributed to the vast improvement of broadband in America. Over 80 percent of U.S. households are served by at least one provider offering service with downstream speeds of a gigabit or more. Just a few years ago, only a lucky few Americans could get gigabit broadband service at home.

Google continues to invest in emerging markets, including, perhaps most notably, automated vehicles (also known as self-driving cars). Alphabet has poured well over a billion dollars into its

Waymo division, which is widely viewed as a global leader in automated vehicle technology. If the technology is ultimately successful, it could revolutionize the safety, efficiency, and affordability of surface transportation to an extent greater than perhaps any innovation since the automobile itself. But it will almost certainly take decades for automated vehicles to proliferate across the country, even once the technology itself has been perfected.

Private sector investment in potentially game-changing innovations can have positive, far-reaching societal implications. Whatever one thinks about the proper role of government in fueling research and development, market-driven efforts to develop revolutionary technologies in hopes of achieving commensurate rewards are an essential ingredient in human progress. But we shouldn't take for granted the willingness of brilliant and creative minds to sweat it out despite the high risk of failure that comes with trying to change the world. Just as governments can help establish the conditions in which these creative efforts thrive, governments can also stymie such efforts.

Had Google known from the start that Android's extraordinary success would mean regulatory headaches down the road, would it have still entered the mobile ecosystem? Quite possibly. But imagine if those headaches were more like throbbing migraines. Even if Google were to still have pursued Android, it might have followed Apple's path, targeting the world's most affluent users to generate solid profits while keeping a lower profile in terms of market share.

How will the government react if Waymo ends up first to market with a safe, affordable, automated vehicle? If such success is met with exacting bureaucratic oversight of every decision Waymo makes when it comes to pricing, strategy, and acquisitions, at what point do the diminished rewards of success make the risk-taking no longer worth it? Like any company, Google will throw in the towel if the prospect of success is sufficiently slim relative to the benefits. Government should foster an economic environment in which efficient risk-taking can thrive. Scrutiny will scale with success, but antitrust intervention should never amount to a hefty marginal tax on innovation.

Rewriting U.S. antitrust laws won't just affect existing firms; it will also influence the evolution of markets that have yet to exist. Although Google's earliest days were financed by personal credit cards, like most of America's tech leaders, the company's formative years were financed by angel investors and venture capital funding. Even the 42-year old Apple got off the ground in the 1970s thanks to wealthy funders willing to bet on a long shot.

But it's astronomically rare for a startup to enjoy the success of Apple or Google. Indeed, investors have poured countless sums into seemingly promising firms that ended up failing spectacularly. Pets.com spent $300 million in under two years before shutting down in 2000. Groupon raised $1.4 billion in its first four years, then another $700 million in its initial public offering — yet it's now worth under $2 billion. Juicero raised almost $120 million in three years before shutting down in 2017. Even MoviePass, whose business model boils down to losing money on every customer in hopes of making up the difference with volume, garnered nearly $70 million in seed funding. These are just a few examples of how, each year, investors pour billions of dollars into startups that never even come close to breaking even, let alone justifying their seed funding.

Why do investors make such risky bets? Because of the potentially huge upside they might enjoy if just one of their investments turns into a multi-billion dollar unicorn. Even a few moderately successful startups that end up being acquired by larger players is often enough to make a venture capital fund worth its risk premium. But as antitrust regulation makes it harder for the rare success story to develop into a large, sustainably profitable operation — or for leading firms to acquire promising startups — it shifts the risk-reward proposition of angel investment and venture capital funding. The upshot? Fewer bold ideas make it to market, fewer creative thinkers quit their day jobs to become entrepreneurs, and more investable assets end up in the pockets of the very incumbents that advocates of antitrust regulation aspire to weaken.

To Understand Antitrust, First Understand Markets

Markets occupied by two or three major players are often demonized as overly concentrated, but the actual number of firms needed to make a market competitive is just *one*. One reason for this is that markets are inherently contestable; even where entry costs may seem formidable, no incumbent can ever be sure that some innovation will eliminate its apparent dominance. Even in markets where entry barriers and other factors render multiple competitors infeasible, artificially introducing competition by forcing incumbents to break up into smaller, less efficient firms may increase higher prices for consumers. No matter how vigorous the price competition, it cannot overcome the economic reality that firms must charge enough to cover their costs in the long run.

It turns out that so-called "natural monopolies"—i.e., markets that the government has deemed incapable of sustaining multiple competitors — rarely, if ever, exist in nature. Instead, they tend to emerge due to government regulation that protects incumbents, thwarting entry. Assume for the sake of argument, however, that Google's dominance in search — or Amazon's in online retail — will continue to grow and endure. Even then, the case for antitrust intervention is hardly a slam dunk. For well over a century, government agencies have regulated monopolies in sectors including telecommunications, electricity, rail, trucking, and aviation. Despite immense efforts by regulators tasked with advancing the public interest to make these markets function well, the results have been abysmal. Failures have been frequent, and merely maintaining a market's mediocrity is considered a regulatory success.

The neo-Brandeisian movement is right about one thing: lawmakers should revisit America's antitrust laws. Their vague wording and susceptibility to wildly different judicial interpretations is problematic by itself. More fundamentally, though, lawmakers should rethink the notion that the government can make consumers better off by banning entire categories of voluntary transactions among market participants. The only obvious beneficiaries of the merger review process are lawyers, economists, and lobbyists. But even price coordination among rival firms, often considered the most obviously problematic form of supposedly anti-competitive conduct — and a criminal offense in many situations — can generate real efficiencies.

Even those who don't share this skepticism of antitrust should recognize the harms of giving regulators and judges far greater powers to shape the evolution of the digital marketplace. Especially with America's tech sector continuing its trajectory of remarkable progress, intervening in this marketplace is a recipe for denying consumers the unknowable rewards of innovations that no one creates. Antitrust may be far from perfect in its current state, but it could be far worse.

New Technology, Same Principles: The Supreme Court and Tech

Ashley Baker, Director of Public Policy at The Committee for Justice

During the confirmation hearings for then-Judge Neil Gorsuch, Senator and former Judiciary Committee Chairman Orin Hatch asks a question about interpreting constitutional provisions in the digital age. How, he asks, can a two-century old document apply to technologies that were not even imagined by the Founders?

Gorsuch responds, "So, the technology changes, but the principles do not. And it cannot be the case that the United States Constitution is any less protective of the people's liberties today than it was the day it was drafted."

Will the Supreme Court — with the recent addition of Justices Gorsuch and Kavanaugh — adhere to the dictum "new technology, same principles?" Only time will tell, but a few recent cases may serve as good indicators. Still, the law governing emerging technologies is predominantly statutory, which means that despite decisions by the courts Congress will inevitably have many questions to address.

Between the current cases before the Supreme Court and sensible legislation from Congress, there is still hope that our institutions will succeed in protecting both innovation and the principles of our founders in this brave new high-tech world.

Privacy

What's Left of the Fourth Amendment?

In what was one of the most important technology cases in Supreme Court history, the justices ruled 5-4 last June that the historical cell phone location data used to convict Timothy Carpenter of armed robbery is subject to the protection of the Fourth Amendment. Acknowledging that Fourth Amendment doctrine must evolve to account for "seismic shifts in digital technology," the high court said that the government was required to obtain a search warrant for the data.

The question before the justices in *Carpenter v. United States* was whether the third-party doctrine and the lower standards of the Stored Communications Act (SCA) allow law enforcement authorities to obtain such data from an individual's cell phone provider without the finding of "probable cause" required for a search warrant. The third-party doctrine, first articulated by the Supreme Court in *United States v. Miller* (1976) and *Smith v. Maryland* (1979), holds that people have no reasonable expectation of privacy under the Fourth Amendment when they voluntarily convey information to a third party, such as a bank or a telephone service provider.

In a world where digital information can be transmitted without affirmative consent, the doctrine has created a gaping hole in the Fourth Amendment. Cell phones are so ubiquitous that, as Chief Justice Roberts wrote in *Riley v. California* (2013), "nearly three-quarters of smart phone users report being within five feet of their phones most of the time, with 12 percent admitting that they even use their phones in the shower."

With *Carpenter*, the Supreme Court concluded that the *voluntary* conveyance assumption behind the third-party doctrine doesn't hold up in light of the precise, retrospective nature of cell phone location data. As Roberts explained in the majority opinion, "a cell phone logs a cell-site record by dint of its operation, without any affirmative act on the part of the user beyond powering up."

As the Court said in the ruling: "Our decision today is a narrow one. We do not express a view on (scenarios) not before us." Because the high court's reasoning logically applies to a variety of current and future technologies, however, the ramifications of the *Carpenter* decision are likely to be anything but narrow.

One practical, short-term implication of the decision is that law enforcement, defense attorneys, and courts will begin to think differently about metadata. Because *Carpenter* was so vague, attorneys may apply the logic of *Carpenter* to a wide array of situations involving other technologies. *Carpenter* may well cause law enforcement to adjust their practices and seek a warrant when searching the digital data that is so pervasive in our lives today. That alone makes the ruling a major victory for privacy rights that will reign in the tendency of the government to use subpoena power to circumvent the higher standards of a search warrant.

Possible Paths to Digital Property Rights

The Court's opinion in *Carpenter* included four dissents — the most written in a single case since *Obergefell*, the 2015 same-sex marriage decision. One of those dissents came from Justice Neil Gorsuch. His dissent was more like a concurrence on other grounds, but it reads as if he was laying the groundwork for future changes to Fourth Amendment jurisprudence, which may ultimately be in his hands.

Justice Gorsuch's dissent was partially based on the Court's failure to confront the third-party doctrine head-on. Gorsuch's dissent argues that rather than employ a Fourth Amendment analysis focusing on the reasonable expectation of privacy, the court should follow a property rights-based theory of the Fourth Amendment. Under that theory, Carpenter had a property interest in his cell phone data.

Although the Court is unlikely to completely discard the third-party doctrine anytime soon, it will likely slowly chip away at the doctrine, and Justice Gorsuch's suggestion of a positive-law approach may be the best path forward. When the legislative branch acts to create specific property rights, as Justice Gorsuch noted, "that may supply a sounder basis for judicial decision-making than judicial guesswork about societal expectations." And since the Supreme Court's ruling in *Carpenter* did not

include more invasive and increasingly prevalent technologies such as facial recognition software, stingray devices, DNA collection, and drone surveillance, it may be time to discuss another approach to vindicate the full protections of the Fourth Amendment.

Meanwhile, Congress should act to strengthen statutory requirements for searches enabled by new technologies. Now that the third-party doctrine is no longer the bright-line rule it once was, confusion is inevitable, and the branch closest to the people is best-equipped to weigh in on societal privacy expectations.

Whether or not Congress picks up its glacial pace of legislation, the Supreme Court and lower courts will continue to settle important technology issues. Policymakers need to be especially cognizant of how legal responses to innovation will impact technologies' application and implementation, as well as expose the need for congressional action.

Antitrust

Will the Supreme Court Take a Bite Out of Apple?

This term's big antitrust case is *Apple v. Pepper*, in which users of Apple's iPhone claim that the company is violating federal antitrust laws by requiring them to buy apps exclusively from Apple's App Store. The question before the Supreme Court is whether iPhone users have standing to bring this suit for damages. The answer hinges on the Court's application of its 1977 *Illinois Brick Co. v. Illinois* precedent, which holds that only the direct purchaser of a product or service may sue for antitrust damages.

Apple argues that it is merely acting as an agent or middleman for app developers, as evidenced by the fact that it sells the apps at the prices set by the developers. The only direct purchasers, Apple says, are the developers themselves, who pay the company a 30 percent commission for use of the App Store. The plaintiffs disagree, noting that iPhone users pay Apple directly and are thus direct purchasers.

If oral argument is any indication, the plaintiffs are likely to prevail. All four members of the Court's liberal bloc and at least a couple of the more conservative justices voiced skepticism of Apple's argument that the instant facts are analogous to *Illinois Brick*. In that case, contractors purchased bricks from the defendant company and used them in buildings the plaintiffs purchased, but the plaintiffs had no direct contact with the defendant. Unlike those plaintiffs, the plaintiffs here buy apps directly from Apple.

Some of the justices seemed to hint that the Court should overrule *Illinois Brick*. Justice Brett Kavanaugh noted that, at very least, it is not clear how *Illinois Brick* should apply to the facts here and suggested that the Court therefore look to the broad language of the governing Clayton Antitrust Act under which "any person injured" by an antitrust violation can sue. If that language is taken literally, iPhone users would have standing.

Only Chief Justice John Roberts seemed to lean toward Apple, noting that the plaintiffs' arguments had evolved from what they said in their complaint. Corey Andrews of Washington Legal Foundation explains that on the way "to the Supreme Court, the plaintiffs' theory of monopolization went from the claim that Apple monopolized the distribution market for apps by charging app developers a supracompetitive 30 percent commission to the claim that Apple's App Store is itself a monopoly because consumers can't buy an app from anywhere else."

Roberts also voiced concern about the possibility of duplicative recoveries in situations such as this if, say, both iPhone users and app developers sue for the same antitrust violation. Justice Gorsuch, however, pointed out that even in states that allow indirect purchasers to file antitrust suits under state law, duplicative recoveries have not been a problem: "Shouldn't we question *Illinois Brick* perhaps," Gorsuch asked, "given the fact that so many states have done so? They've repealed it."

Apple v. Pepper Is About Who Is Entitled to Sue

If the plaintiffs do prevail on standing, the case will go back to the lower courts where the iPhone users will still need to demonstrate that Apple is in violation of antitrust laws. Similarly, any precedent set by the Court's decision later this year will be limited to standing and will not address the central question of whether the App Store violates antitrust law. In other words, *Apple v. Pepper* will decide who is allowed to sue.

The reach of any decision will also be limited by the practical fact that Apple has a unique business model regarding app sales. Since the plaintiffs were unable to point to other distributors with a similar business model, the justices may not worry, as some amici do, that a ruling against Apple would open the floodgates to similar e-commerce lawsuits.

That said, if the Court overturns *Illinois Brick*, giving indirect purchasers standing, the impact on antitrust law will not be limited to the tech industry. For the Court to overturn *Illinois Brick*, however, would be an extraordinary step, especially since the plaintiffs have not asked it to do so and the Court did not grant certiorari on the question of the precedent's viability.

Moreover, Congress has rejected 17 bills over the past four decades to overrule *Illinois Brick* and repeal the direct purchaser rule. Nonetheless, given Congress's recent heightened interest in enforcing the antitrust laws against the tech giants, this case — especially a decision for Apple — could motivate Congress to consider legislation that negates or modifies the rule.

Free Speech

Will a Free Speech Case "Break the Internet"?

A relatively low-profile but important free speech case this term is *Manhattan Community Access Corporation v. Halleck*, in which two content contributors are suing a public-access television channel, claiming it is violating their First Amendment rights by barring them from appearing on

the channel for harassing an employee in 2012. The question before the Supreme Court is whether the private operator of this channel (known as MNN) is a "state actor" — that is, a person or entity acting on behalf of the government — and is therefore subject to the First Amendment, which normally applies only to the government.

The potentially broader issue before the high court is when private property can be a "public forum" — a place like a public street or park, where free speech is protected — and the owner of the property can be deemed a state actor and thus be subject to the First Amendment. This question is increasingly relevant as Congress and the legal community debate whether privately owned social media platforms such as Facebook and YouTube have any obligation to respect the requirements of the First Amendment.

If the justices conclude that Manhattan's designation of a private company to operate a public-access channel turns the company into a state actor, could the precedent apply to online media providers as well? That's an unlikely outcome, but the case will undoubtedly impact the larger ongoing debate over content moderation.

While acknowledging that on its face this case has little to do with online platforms, a brief filed by the Internet Association urges the Court to issue an "exceedingly narrow" ruling because of concerns "that any decision that deems MNN a state actor will be misinterpreted in ways that are highly damaging to the Internet."

Another amicus brief warns the justices that "Internet service providers and platforms are directly affected by the uncertainty" generated by this case, because "these companies have developed business models that rely on providing open platforms" and some have even "sought to partner with municipalities to provide communities with access to the Internet." Therefore, it is important to clarify the line between state and private actors.

The concerns of Internet companies may have been alleviated somewhat at oral argument, when the plaintiffs' attorney emphasized that the Court was not asked to consider the broader public-forum question. Furthermore, plaintiffs explain, "only a government can create a public forum," so Internet platforms can't be analogized to government-mandated public-access channels.

At least two of the Court's liberal justices, Ruther Bader Ginsburg and Sonia Sotomayor, appeared sympathetic to the plaintiffs' argument that Manhattan's designation of MNN to run the channel and its operational requirements that limit MNN's discretion render the company little more than an agent of the government. Justice Brett Kavanaugh, on the other hand, was skeptical of that argument. He noted that MNN was a private company "not operating on government property" and was thus very much like public utilities, which have been held to not be state actors.

Kavanaugh and Congress Could Matter Most

For the past three decades, a substantial portion of opinions in important free speech cases were authored by Justice Anthony Kennedy, giving him substantial influence on First Amendment

jurisprudence. Although he, like Justice Sandra Day O'Connor, was often a swing vote in these cases, some would argue that Justice Kennedy's central legacy was his insistence on putting First Amendment cases first.

Following his departure from the Court last June, it remains to be seen which justice will fill his shoes in prioritizing free speech cases. Although Chief Justice Roberts recently declared himself the First Amendment's "most aggressive defender," Justice Kavanaugh's confirmation to the Court is likely to be more consequential, as he is replacing the swing vote. And if lower court judicial records are any indication, Kavanaugh is likely to weigh in on free speech cases, and may even place an emphasis on online speech. The role of the D.C. Circuit combined with the increasing number of tech-related cases over the past several years has given him ample exposure to the issue.

There is no shortage of evidence supporting this prediction. In an opinion dissenting from the denial of rehearing en banc in *United States Telecom Association v. FCC* (2016), Kavanaugh wrote that the government may not "regulate the editorial decisions of Facebook and Google" or "impose forced-carriage or equal-access obligations on YouTube and Twitter." In *Cablevision Systems Corp. v. FCC* (2010), Kavanaugh also wrote separately to discuss the First Amendment problems raised by carrier restrictions, noting that "[t]he First Amendment endures, and it applies to modern means of communication as it did to the publishers, pamphleteers, and newspapers of the founding era."

Although the high court's decision in *Manhattan Community Access Corporation v. Halleck* is likely to be a narrow one, policymakers should pay extra attention to Justice Kavanaugh for potential discussion of Internet platforms.

As for Congress, while this case won't directly impact social media platforms, conversation surrounding it is certain to spill over into the debate over Section 230, the law that protects online speech by providing limited immunity from third-party speech liability. Lawmakers are upset about social media companies' content moderation practices, and — rightly or wrongly — want to do something about it. That Section 230 immunity is politically vulnerable was demonstrated when Congress passed SESTA-FOSTA last year, ending immunity for content related to sex trafficking.

Since the Court won't grant a more directly relevant case in the immediate future, it may be that Congress decides the fate of their limited immunity provision that inadvertently created the Internet as we know it. It is therefore important for policymakers to consider carefully each aspect of this issue rather viewing it through the lens of politics alone.

Conclusion

In the areas of the law most important to the future of America's tech industry — privacy, antitrust, and free speech — both Congress and the Supreme Court can contribute to the continued vibrancy of America's tech industry and the economic growth it spurs.

The Supreme Court holds tremendous power, but as these three cases demonstrate, it both moves incrementally and is constrained by the law. Sometimes the Supreme Court is the wrong forum for updating American law to reflect technological advancement, making the need for congressional action greater. Protecting the Internet could require wise legislation on current issues.

Understanding Calls for Regulating Artificial Intelligence

Will Rinehart, Director of Technology and Innovation Policy at the American Action Forum

Elon Musk, Bill Gates, Mark Cuban, and the late Stephen Hawking have been among the most vocal luminaries calling for the regulation of artificial intelligence (A.I.), but they are hardly alone. Countless papers, conferences, and talks dedicated to algorithms and artificial intelligence call for the same. Without detailing the harms, or explaining how the market has failed, many tend to focus on proposals to tax, regulate, and limit artificial intelligence.

Embedded in these calls for new government power are countless uncertainties about the direction of technology. Yet, the track record of technology forecasts is far from stellar. One of the largest retrospective reviews of technology forecasts found that predictions beyond a decade were hardly better than a coin flip. In an analysis that focused specifically on A.I. predictions, the authors warned of "the general overconfidence of experts, the superiority of models over expert judgement, and the need for greater uncertainty in all types of predictions." Predictions that general A.I. is just around the corner have failed countless times across several decades.

This uncertainty indicates a fundamental reality about A.I. It is a developing collection of technologies with a tremendous variety of applications. As a result, the goal for policymakers should not be a singular A.I. policy or strategy, but a regulatory and policy approach that is sensitive to developments within society, leaving room for innovation and change.

Terms and Origins

To understand artificial intelligence, it is helpful first to define terms, especially "narrow" A.I. and "general" A.I. Narrow A.I. references models built using real-world data to achieve narrow, specific objectives such as translating languages, predicting the weather, spotting tumors in chest scans and mammograms, and helping people identify caloric information just from pictures of food.

Narrow A.I. can be contrasted with general A.I., which refers to decision-making systems able to cope with any generalized task like a human. Arnold Schwarzenegger's early 1990's movies and, more recently, Samantha from the movie *Her* represent this kind of A.I. While some fret over the risks posed by super-intelligent agents with unclear objectives, task-specific A.I. holds immediate promise while general A.I. is still far from full realization.

The diversity of what one thinks of as A.I. extends beyond these categories, too. Machine learning is another commonly referenced term, which denotes a process whereby a machine analyzes data and learns without supervision. Yet the barriers between A.I. and machine learning and more standard computer programming are blurry. In practice, there often isn't much difference between narrow A.I. and complex computer programming like machine learning. But there is an upside to this

diversity, as Stanford's "One Hundred Year Study on Artificial Intelligence" noted: "[T]he lack of a precise, universally accepted definition of A.I. probably has helped the field to grow, blossom, and advance at an ever-accelerating pace."

This diversity stems from the technology's democratic origins. As the Obama White House noted in its "Preparing For the Future of Artificial Intelligence" report,

> The current wave of progress and enthusiasm for A.I. began around 2010, driven by three factors that built upon each other: the availability of big data from sources including e-commerce, businesses, social media, science, and government; which provided raw material for dramatically improved machine learning approaches and algorithms; which in turn relied on the capabilities of more powerful computers.

In January of 2010, the machine learning library, scikit-learn, was released to the public, democratizing the tools of A.I. and sparking the current rush. This program finds its genesis in Google's Summer of Code programs, and many different companies and entrepreneurs have applied these tools in manifold ways. As Representative Will Hurd said in June, "The United States boasts a creative, risk-taking culture that is inextricably linked to its free enterprise system."

Google, Facebook, Microsoft, and other large tech companies have played a large part in the development of A.I. While it has been popular of late to criticize the largest tech companies, policymakers should be comfortable with large firms such as Google, Facebook, and Microsoft taking the lead on A.I. implementation. Even though these companies have been lambasted for their size, they shouldn't be penalized for adopting advanced technologies.

The democratic nature of A.I. development over the last decade means that there are a variety of experiments in the ecosystem, and shifting to A.I.-embedded processes will not be frictionless for firms or social institutions. As the American Action Forum has noted before, firms face significant practical hurdles in implementing A.I.-driven systems, as they aren't cheap and most automation schemes fail to achieve any positive results. The same kind of implementation problems exist in government institutions as well. In the most comprehensive study of its kind, George Mason University law professor Megan Stevenson tracked Kentucky's state-wide implementation of an algorithm meant to automate bail decisions by judges. While there were significant changes in bail-setting practices, over time these changes eroded as judges returned to their previous habits.

Regulation and A.I.

The shifting landscape and unclear implications of A.I. mean that policymakers should adopt three outlooks regarding narrow A.I. regulation.

First, A.I. is a general purpose technology, like electricity, the automobile, the steam engine, and the railroad, that will have a variety of regulatory impacts. A.I. isn't going to converge industry

regulation but make it more variegated. Thus, calls to impose a singular regulatory framework on A.I. are misplaced. Some industries might need clarity, others might need a shift in liability rules, and yet others might need additional consumer safeguards. Still, we are a long way from those deep societal impacts. In the near term then, policymakers should be on alert to the potential barriers that could hobble growth in A.I. application, which might necessitate the liberalization of rules.

Second, premature action is likely to be deleterious to A.I. innovation and progress, as privacy regulation in Europe has shown. A rush to legislate A.I. applications, and thus constrain and narrow them, would signal to investors and innovators that their time, money, and talents should be put elsewhere. Such a shift would be a real loss, as the opportunities for A.I. applications are enormous. The United Kingdom's National Grid has turned to A.I. to reduce service outages. Facebook and MIT are using A.I. to give addresses to people throughout the world without them. And even the *New York Times* is getting into the game, by installing a recommendation feed for its users through A.I.

Regulatory restraint does not mean consumers are exposed to harms. Consumers can be protected if policymakers choose the route of soft law. As Ryan Hagemann, Jennifer Huddleston, and Adam Thierer explained, "soft law represents a set of informal norms, multi-stakeholder arrangements, and non-binding guidance standards that provide an adaptable alternative to more traditional regulations or legislation." These approaches have been successfully applied to autonomous vehicles, Internet of Things, advanced medical technologies, FinTech, and electric scooters. Relying on these strategies would be a smart strategy for A.I. regulation.

As a final matter, policymakers should temper concerns about the ethical implications of A.I. The terminator scenario of A.I. might be well known, but it is not indicative of the current hurdles that A.I. researchers face. Instead, practitioners tend to be concerned with more concrete obstacles, such as avoiding side effects that reward hacking, ensuring that there is scalable supervision, and stopping undesirable behavior during the learning process.

Moreover, countless organizations are dedicated to these ethical problems, such as Data&Society, the Ethics and Governance of A.I. Initiative, and the A.I. Now Institute, just to name a few. Companies are beginning to hire researchers focused in A.I. ethics and are creating internal ethics boards for A.I. Moreover, educational facilities are beginning to implement ethics within curriculum. As Computer Science Professor Yevgeniy Vorobeychik explained in a filing, "the vast majority of A.I. researchers already set public good, broadly construed, as their aim." Policymakers should be optimistic about society's ability to consider and act on A.I.'s ethical implications with both speed and nuance.

In short, policymakers should embrace regulatory restraint, although there are opportunities for policy to strengthen A.I. deployment.

A Framework for Increasing Competition and Diffusion in Artificial Intelligence

Caleb Watney, Fellow at R Street Institute

Artificial Intelligence (A.I.) is developing rapidly, and countries from around the globe are beginning to articulate national strategies for handling the political ramifications. With A.I. powering innovations such as driverless cars, autonomous drones, full-sequence genetic analytics, and powerful voice-assistant technology, the future certainly looks full of potential. Unsettled questions, however, about who will reap these benefits and when they will be achieved leave storm clouds on the political horizon.

Amid questions of industrial concentration and economic inequality on one side, and concerns about lagging U.S. productivity and the slow pace of A.I. diffusion on the other, there is an underexamined overlap that connects these questions to the same set of policies: high barriers to entry due to supply-side constraints.

There are significant barriers to entry in A.I. development and application, many of which stem directly from government policies. These barriers have inadvertently boosted the market power of incumbent firms and in reducing them, we may enable new firms to compete better, while also removing some of the bottlenecks that slow down research and integration of A.I. systems across the entire economy.

Supply of Skilled A.I. Analysts

Perhaps the single biggest bottleneck in A.I. development and application today is the supply of skilled data scientists and machine-learning engineers. Typical A.I. specialists can expect to earn between $300,000 and $500,000 at top tech firms, numbers that are significantly higher than their peers in other computer-science-related subfields. In addition to these ballooning salaries, industry experts such as Hal Varian have pointed to the scarcity of adequate A.I. talent as the largest factor behind slow application in the economy.

Reform Our Immigration System to Allow More High-Skill A.I. Talent

The policy lever with perhaps the highest degree of leverage to begin immediately alleviating this talent shortage is our immigration system and, more specifically, reforming visas for international graduate students.

In 2015, the United States had 58,000 graduate students in computer science fields, the overwhelming majority of which (79 percent) were international. This influx of talent represents a significant portion of the overall A.I. talent supply being cultivated each year, as students from all over the world are attracted to the nation's top education system. In particular, the United States

attracts large numbers of students from China and India. Due to a limited number of visa slots, however, only a fraction of these students are allowed to work in the country long term.

The primary pathway for these highly skilled immigrants to stay in the country is through the H-1B visa program. For the past 16 years, however, the H-1B limit has been exhausted and, in more recent years, the number of applications filed has consistently been twice as high as the number of available spots. This discrepancy is almost certainly understating the scope of the problem, as it does not account for the ways in which foreknowledge about the difficulty of acquiring a work visa may deter students from applying in the first place.

Although it also limits the talent pool available to large tech firms, the status quo is especially daunting for startups, as they do not have the specialized Human Resources personnel to handle the bureaucracy of the immigration visa application process. Including application and attorney fees, to sponsor a work visa typically costs around $5,000 per employee, and the paperwork burdens appear to be increasing. Both the financial and bureaucratic costs are easier for established firms to bear, given their larger size and increased resources.

In turn, this cost impacts the types of firms high-skill immigrants will apply to work for in the first place. Even when attracted to work at startups, foreign workers may ultimately privilege their applications to incumbents because they will likely have a better chance of obtaining work visas at established firms. Additionally, since startups face high failure rates, job loss could mean termination of work authorization as well — which would mean that the entire visa application process would have to be approached anew.

Accordingly, to allow more international students to live and work in the United States upon completion of their degree — either through an expansion and simplification of the H-1B visa program or through the creation of a new technical worker visa program — would be a relatively straightforward and effective method to alleviate the country's talent shortage around A.I. In particular, this reform would benefit smaller firms and startups that are unable to access existing foreign-born talent to the same degree as established firms.

Allow Companies to Deduct the Cost of Training A.I. Talent

In addition to reforming our immigration pathways for high-skilled A.I. talent, it would be wise for the United States to extend more effort toward building up domestic talent. One way to achieve this end would be to better align incentives for companies to develop A.I. talent internally.

As the number of newly minted machine learning (ML) Ph.D. students continues to dwindle, some companies are looking at training employees internally to essentially create new supply. Such training, however, requires significant investment on the company's part, both in time and resources, to train new A.I. specialists this way, and the gains from this training are mostly captured by the newly trained worker in the form of higher wages. Since workers can jump ship from the companies that train them at any time for a higher salary at a competitor, employers have few

opportunities to recoup the costs of worker training. It thus seems likely that employers are generally underinvesting in worker training when compared to the amount that might otherwise be efficient. We should therefore look more closely at incentivizing this socially desirable behavior through the tax code.

Employers may currently deduct a portion of the costs of worker training as long as it is to improve productivity in a role they already occupy, but this credit is fairly small and employers may not deduct the costs if it would qualify them for a new trade or business. Expanding this deduction — both in size and scope — so that the full cost of worker training for new trades could be deducted, would incentivize more investment in building the A.I. workforce that is needed to fuel our economy. Given the pre-existing level of interest by employers in this strategy, it seems likely this could become a fruitful part of our domestic A.I. pipeline, if given more support.

Supply of Data

In many ways, the supply of high-quality machine-readable training data is the key enabler of ML. Without access to some underexplored dataset, a team of talented A.I. specialists can be left twiddling their thumbs. Consumer data in the United States is particularly valuable, but here again large incumbents have significant (though not unsurmountable) data advantages over startups.

But we can potentially create high-leverage opportunities for startups to compete against established firms if we can increase the supply of high-quality datasets available to the public. As with increasing the supply of A.I. talent, this reform will help both incumbents and startups, but on the margin it will be the smaller firms with less access to consumer data who benefit most.

Encourage the Creation of Open Datasets and Data Sharing

One of the easiest ways to begin this process would be a more thorough examination of existing government datasets that are not public. As an example of previous projects that were broadly successful, consider the U.S. National Oceanic and Atmospheric Administration and Landsat projects, both of which made weather-satellite data available to the public and, in turn, developed into a multi-billion-dollar industry, creating more accurate forecasts of extreme weather and crop patterns.

There appears to be even more potential from datasets the government owns but has not made public. For example, many cities and municipalities have useful data around traffic patterns, electricity usage and business development that, if made accessible, could lead to reduced-cost service provision and better analytics.

There is also the matter of industries in which open data might become the norm if existing regulations are relaxed or streamlined. The healthcare industry seems a particularly promising target in this respect, as the Health Insurance Portability and Accountability Act (HIPAA) has long

been considered a barrier to the development of data sharing between medical professionals and companies. Allowing consumer health data to be more easily shared with the proper privacy safeguards could enable a renaissance in drug development and personalized medicine, as recent ML advances have proven quite promising when appropriate data have been available.

Each new dataset that can be easily shared or, when appropriate, made public, increases the odds both that a new startup will be able to leverage it for success, and also that a new industry can thrive around the increased predictive analysis the released data has enabled. For recent advances in A.I. to diffuse throughout the economy, we must make sure the underlying data is accessible.

Clarify the Fair-Use Exemption for Training Data

In addition to making more government datasets open source, we should also take a second look at some of the intellectual property laws that intersect and interact with the ML process, specifically copyright law.

Imagine a hypothetical startup focused on the creation of a natural-language-processing application. One readily available source of human dialogue the company might consider learning from would be the last 50 years of Hollywood scripts, many of which are scrapable from various online databases. Such an endeavor, however, would stand on legally dubious grounds, as these scripts remain copyrighted works and there have not been clear legal guidelines established to delineate what is allowable as fair use in ML training data. Given this uncertainty, it is more likely that such a startup would avoid this potential legal minefield and consider what other datasets might be available with less risk.

Such is the ambiguous state of copyright enforcement in ML today. And it may also have important and underexplored applications for the state of competition in A.I.

There are an enormous number of copyrighted works that are scrapable from the Internet, the data of which is currently underexploited in part because of its legally dubious standing if used as training data. This reform could represent, then, a significant lever to create new arbitrage opportunities for scrappy startups willing to find and leverage interesting datasets.

Given the existing ambiguity around the issue and the large potential benefits to be reaped, further study and clarification of the legal status of training data in copyright law should be a top priority when considering new ways to boost the prospects of competition and innovation in the A.I. space.

Access to Specialized Hardware

Underlying the data being used to train ML models and the data scientists who are building them is the physical infrastructure of the A.I. world. This primarily takes the form of the computer servers

and chipsets that ML models are trained and operated on. In recent years, this hardware has become increasingly specialized to keep up with the pace of A.I. development.

While a natural and necessary part of the A.I. development process, such a trend toward specialized hardware does increase the fixed costs required to be competitive. This cost manifests not only in the expense of these systems, but in the elaborate supply chains that have been built up to support them. While the policy recommendations that flow out of this insight are less clear cut than those for the supply of A.I. analysts or datasets, maintaining access to valuable A.I. hardware is a key policy consideration.

Avoid Political Instability in International Supply Chains

As A.I. hardware becomes more specialized, the supply chains for very specific chips become a critical ingredient for cutting-edge ML research. While the United States maintains advanced manufacturing facilities that are vital to the supply chain, much of the production for particular parts (like back-end semiconductor fabrication) have been outsourced. Given the importance of chip foundries in Taiwan and China in particular, the perceived stability of trade in the region will alter investment patterns and domestic access to these sophisticated chips.

To ensure access in spite of political tensions, large companies such as Apple, Google, and Nvidia are beginning to re-shore production of especially valuable chips. Smaller competitors and startups, however, are much more limited in this capacity and thus are more reliant on existing international supply chains.

Insofar as recent U.S. trade tensions with China have increased the perceived instability of regional trade, the disparate impact this instability will have on smaller firms should be recognized. Ultimately, new foundries and semiconductor manufacturing plants will shift wherever they are most profitable. Accordingly, in the event of a long-term trade war, production could eventually shift elsewhere. Trade tensions, however, will certainly shape short- and medium-term access to specialized hardware.

Maintain a Healthy Ecosystem Around Distributed Platforms

The other significant trend in A.I. hardware utilization is the growth of cloud-computing platforms such as Amazon Web Services (AWS) and the Google Cloud platform. Cloud computing has notable pro-competitive effects in that it transforms what is normally a fixed cost in server capacity into a variable one. Allowing a startup to buy only the discrete server space they will need for that month significantly reduces the amount of venture capital needed to get a company off the ground.

This becomes even more important as A.I. hardware becomes more specialized. Requiring a startup to buy different chips for the various life cycles of training and operating an ML algorithm would be a significant financial outlay and almost certainly hurt the ability of startups to compete. Fortunately, both AWS and Google Cloud have been competing with one another by adding

specialized A.I. hardware as a part of their platform offerings. This offering essentially allows startups to spread out the increased fixed costs of specialized hardware over a longer time horizon, which makes it more manageable.

In addition to the physical servers themselves, cloud computing companies are increasingly offering ML services such as voice recognition, translation, and image recognition to save startups the hassle of building their own software tools for each discrete task. Again, it is difficult to understate how much easier this makes the process of launching a startup, and it is a very positive development for the overall health of the A.I. ecosystem.

As this portion of the ecosystem largely seems to be developing in a healthy manner, the United States should be careful to avoid data-localization laws, excessive privacy laws, and other legislative efforts that might disrupt the careful balance. On the whole, recommendations for this area should largely follow the Hippocratic Oath and "First, do no harm."

What About Antitrust?

It is worth contrasting this general approach of reducing barriers to entry with another commonly cited remedy: stronger antitrust enforcement. While concern over the level of domestic competition faced by large tech firms is, of course, not unique to A.I., it has certainly raised the stakes given how central the technology is to their current and future business models.

Traditional antitrust measures, however, may prove to be both fairly difficult to implement and high risk for dealing with this perceived problem. After all, there are many plausible arguments supporting the current consolidated structure of the A.I. industry, particularly those that emphasize the importance of cross-cutting technical expertise, and the ability to leverage data and services from one business application to another.

If critics are right, breaking up or actively restricting the merger activities of large tech firms could lead to more innovation in the long run. If these companies are indeed leveraging their significant market power to make it harder for startups to compete with them, breaking them up or constraining them could be a remedy.

If critics are wrong about the optimal market structure of A.I. development and strong antitrust action is pursued, however, the consequences could be dire. An increasing amount of evidence suggests that a small sliver of firms on the technological frontier have been responsible for the lion's share of productivity gains in the economy. Breaking up these large tech firms potentially risks killing the goose that lays the golden egg.

By contrast, focusing on lower barriers to entry is a fairly low-risk strategy for injecting more competition into the A.I. landscape. If the United States can make it easier for startups to compete against large, established incumbents, it increases the likelihood of achieving the boosts to

dynamism and innovation that antitrust advocates champion. Further, it would do so without risking the destruction of the current market equilibrium that is producing significant gains for consumers and for the broader economy. If incumbents can withstand the Schumpeterian winds of increased competition from startups, it is all the better for them.

As this essay suggests, there are significant barriers to entry in A.I. development that have boosted the market power of incumbent firms. If, in the absence of these barriers, new startups can successfully compete, it will be a win for innovation, consumers and for the dynamism of the economy as a whole. To ensure a competitive and innovative ecosystem going forward, policymakers should prioritize reducing the barriers to entry as our first line of defense.