
NISKANEN C E N T E R

White Paper

THE POLICYMAKER'S GUIDE TO EMERGING TECHNOLOGIES

Ryan Hagemann

Senior Director for Policy
Niskanen Center

Samuel Hammond

Director of Poverty and Welfare Policy
Niskanen Center

Dr. Joseph Majkut

Director of Climate Policy
Niskanen Center

Alec Stapp

Technology Policy Fellow
Niskanen Center

November 2018

EXECUTIVE SUMMARY

This guide is intended as an omnibus primer for policymakers interested in learning more about technology policy issues and recommendations for what they can do to effectively address emerging concerns.

It begins by discussing the broad contours of the emergent “soft law” regulatory governance order. As described in Part I, soft law is now the de facto system by which new rules are promulgated to regulate technologies that challenge the traditional command-and-control regulatory models of the 20th century. Part I examines the costs and benefits of this governance approach and lays the necessary foundation for understanding how the technology industry has outperformed other sectors, propelling American economic and productivity growth over the past quarter century.

Part II then looks at the many intersecting elements implicated in the ongoing socio-economic maturation of the Internet and digital communications technologies. In particular, it focuses attention on the purported problems related to antitrust, copyright, and privacy in the digital economy. It argues that the technology sector does not present unique challenges requiring a significant reevaluation of existing laws and standards. The law is remarkably resilient, and in many cases existing statutory authorities are sufficient to address problems related to competition, copyright, and privacy concerns. Policymakers should take care not to act too hastily in responding to perceived problems in these policy arenas; otherwise, they run the risk of impacting the future of the digital economy in profoundly unintended, and potentially disastrous, ways.

Part III goes on to examine a number of new technologies, the issues and concerns specifically related to their emergence, and a detailed set of recommendations that can help guide policymakers confronting the complicated questions that arise when considering whether or how to respond. Whether assessing the safety and effectiveness of new gene therapy treatments, the privacy and cybersecurity implications of the Internet of Things, or the possibility of a renaissance in airspace innovation, this section will provide actionable policy proposals that can maximize the benefits of emerging technologies while mitigating the most significant risks.

Finally, Part IV is dedicated to one of the most important technology policy issues of the day: artificial intelligence. Much like the Internet, artificial intelligence is a “nexus technology” — where advancements in its practical application impact developments in a wide range of other technologies — that holds the potential to upend large swaths of the American economy. Such disruption is perceived as a threat in many quarters of society, despite the many positive quality of life improvements presaged by ongoing developments and new innovative applications. Given the palpability of those concerns — and the rhetorical power they project in public dialogues on everything from privacy to autonomous vehicles — this section examines the actual, real-world impacts of artificial intelligence. It then offers a series of effective, innovation-friendly oversight mechanisms that are well suited to addressing the real, not merely perceived, harms that may emerge from the technology’s deployment.

The recommendations contained within this guide can help effectively guide policymakers towards policy decisions that will better balance the many complicated trade-offs associated with particular courses of legislative or regulatory action. The extent to which it is successful in helping strike a better balance and compromise in technology policy debates will ultimately be reflected in the degree to which leaders in government are committed to accelerating the rate of innovation in the world of atoms, bringing it a little closer to the brisk pace we observe in the world of bits.

ACKNOWLEDGMENTS

There are a great many Niskanen-affiliated scholars and fellows, past and present, whose research and scholarship helped lay the foundation for the background and recommendations contained within these pages. Without their tireless efforts, much of the work included here would not exist; for that, the authors are eternally grateful to Dr. Anastasia Greenberg, Dr. Joseph Gulfo, Dr. Jason Briggeman, Dr. Brandon Valeriano, Joshua Hampson, Regina Zernay, Adam Wong, Nicholas Ciuffo, Juno Zhang, and Eric Lieberman. Additionally, the authors wish to extend their gratitude to those who assisted in providing feedback and suggestions for improvements to this work: Virginia Postrel, Adam Thierer, Dr. Anastasia Greenberg, and Dr. Wayne Brough.


Finally, this compendium owes much of its substance to the research, writing, and ideas of numerous individuals who have labored and toiled in professions beyond the world of policy. Their evangelism on behalf of the future is the type of optimistic narrative that American political life desperately needs more of — now more than ever. To that end, the authors would like to thank those many luminary figures — many of whom are quoted within these pages — who have helped inspire a new generation to embrace the open-ended future: Marc Andreessen, Peter Thiel, Virginia Postrel, Joel Mokyr, Elon Musk, Ray Kurzweil, Matt Ridley, Benedict Evans, Saku Panditharatne, and Aubrey de Grey

Although this work is primarily intended for policymakers, the spirit of its contents are dedicated to them, and to all those committed to working for a better and brighter future for all humanity.

TABLE OF CONTENTS

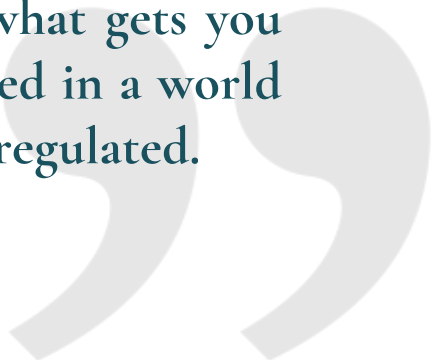
INTRODUCTION	I
PART I: “SOFT LAW” IS EATING THE WORLD	5
PART II: INNOVATION IN BITS	10
A. Antitrust	10
1. The American Model	11
2. The European Model.....	12
3. Policy Recommendations	12
B. Privacy	12
1. The American Model	13
2. The European Model.....	13
3. Policy Recommendations	13
C. Copyright	14
1. The American Model	14
2. The European Model.....	14
3. Policy Recommendations	14
PART III: INNOVATION IN ATOMS.....	16
A. Genetic Modification	16
1. Insurance Reimbursements for Gene Therapy Treatments	17
2. Improving the FDA Regulatory Approval Process	17
3. Policy Recommendations	20
B. Internet of Things	21
1. Encryption, Online Liability Protections, Cybersecurity Insurance, and the VEP	21
2. Quantum Computing: The Future of Cybersecurity.....	24
3. Policy Recommendations	24
C. Autonomous Vehicles	25
1. Addressing Privacy and Cybersecurity Concerns	25
2. Updating Existing Laws and Regulations	27
3. Policy Recommendations	27
D. Commercial Drones.....	27
1. Air Traffic Management and Airspace Auctions.....	28
2. Following Through on the Integration Pilot Program.....	28
3. Policy Recommendation.....	29
E. Supersonic Flight.....	29
1. Lifting the National Airspeed Limit.....	30
2. Noise Standards and Abatement	31
3. Policy Recommendations	32
F. Commercial Space.....	33
1. Licensing Commercial Activity in the Final Frontier	33
2. Promoting Certainty Through Default Approvals.....	34
3. Policy Recommendations	35

G. Climate Engineering.....	36
1. The Need for Further Research.....	36
2. OSTP Oversight and Multistakeholder Governance.....	36
3. Policy Recommendations.....	37
PART IV: ARTIFICIAL INTELLIGENCE.....	39
A. AI in Digital Advertising.....	39
1. The Economics of Online Advertising.....	40
2. Self-Regulatory Governance.....	41
3. Policy Recommendations.....	43
B. AI in Medical Devices.....	43
1. Soft Law at (and Beyond) the FDA.....	43
2. Voluntary Standards and the Software Precertification Pilot Program.....	44
3. Policy Recommendations.....	47
C. Algorithmic Accountability.....	47
1. Principles for Regulatory Oversight.....	48
2. Informational Injuries and AI-Specific Harms.....	49
3. Policy Recommendations.....	51
CONCLUSION.....	54
REFERENCES.....	55
“SOFT LAW” IS EATING THE WORLD.....	55
INNOVATION IN BITS.....	55
Antitrust.....	55
Privacy.....	56
Copyright.....	56
INNOVATION IN ATOMS.....	56
Genetic Modification.....	56
Internet of Things.....	57
Autonomous Vehicles.....	58
Commercial Drones.....	58
Supersonic Flight.....	59
Commercial Space.....	59
Climate Engineering.....	60
ARTIFICIAL INTELLIGENCE.....	60
AI in Digital Advertising.....	60
AI in Medical Devices.....	60
Algorithmic Accountability.....	61
APPENDIX A: FOUR-CATEGORY SAFETY AND EFFECTIVENESS PARADIGMS.....	62
A.1. Drugs and Biologics.....	62
A.2. Diagnostics.....	63
A.3. Devices.....	64
APPENDIX B: AI DEVELOPMENT BEST PRACTICES FOR MEDICAL PRACTICES.....	65



There's the question of stagnation, which I think has been a story of stagnation in the world of atoms, not bits. I think we've had a lot of innovation in computers, information technology, Internet, mobile Internet in the world of bits. Not so much in the world of atoms, supersonic travel, space travel, new forms of energy, new forms of medicine, new medical devices, etc. It's sort of been this two-track area of innovation.

There are a lot of questions of what has caused it and I think maybe that's a good part to start in terms of what gets you out of it. On a first cut, I would say that we lived in a world in which bits were unregulated and atoms were regulated.



— Peter Thiel

INTRODUCTION

Twenty years ago, in November 1998, Virginia Postrel released her seminal classic, *The Future and Its Enemies*. In that book, she eloquently articulated two competing visions of the future: *dynamism* and *stasism*.

Dynamists, she argued, not only recognized the reality that the world around them was in constant flux, but embraced it, understanding the “creation, discovery, and competition” that characterized the emerging future offered a vision of something better. At their core, dynamists are optimists who welcome the open-ended possibilities of the future.¹ In contrast, stasists place a premium on a world that is regulated and engineered for maximal order and stability. For them, the world of the present is something to be protected and sheltered, which ensures the known benefits of the present can be guaranteed in the face of an uncertain and unpredictable future. As Postrel herself summarized these differences:

*Stasists and dynamists are thus divided not just by simple, short-term policy issues but by fundamental disagreements about the way the world works. They clash over the nature of progress and over its desirability: Does it require a plan to reach a specified goal? Or is it an unbounded process of exploration and discovery? Does the quest for improvement express destructive, nihilistic discontent, or the highest human qualities? Does progress depend on puritanical repression or playful spirit?*²

As Americans confront the most profound and punctuated incidents of technological advancements in human history, these questions — and the worldviews the answers represent — are more relevant than ever. Is our society equipped to craft a stasist-minded blueprint for technocratically managing the emerging future, or is the unconstrained dynamism relished by techno-optimists the better — and perhaps only — means of charting a path through the murky unknown? Which path should our society and politics prioritize as

the ideal template for policymaking in the 21st century?

The history of China offers a lesson in what awaits societies that are skeptical and fearful of “verges” — those places where, as historian Daniel Boorstin noted, “something and something else” encounter one another.³ As Postrel noted, China was “for centuries the most innovative society on earth. In nearly every area of human endeavor, the Chinese developed technologies far ahead of their European counterparts. ... Then in the fifteenth century it all stopped. Even books disappeared.”⁴ She continues:

Voyages of exploration were forbidden, and the records of earlier voyages burned. By 1500, building a seagoing junk with more than two masts was punishable by death; a half-century later, going to sea to trade was considered a form of treason. Long-established technologies for mining, silk reeling, and telling time were forgotten. The state took over some foreign trades and shut down others, eliminating the unpredictable verges of the merchant cities. ...

*The details of how and why China abandoned its heritage of technological dynamism are murky. What is clear, however, is that reactionary ideals, technocratic administration, and monopoly power converged to enforce stability at the cost of stagnation. Even in its creative period, China’s dynamism was primarily technological, not social, economic, or political; the government was both highly bureaucratic and absolutist. ... The governing philosophy was one of order, subordination, and stasis. And like most bureaucracies, the mandarin state developed a strong interest in protecting the status quo. The system became reinforcing — a sterile verge between interests and ideology.*⁵

Every society characterized by dynamism is forever precariously situated on the edge of slipping into a stagnant abyss of stasism. “It is as if technological creativity is like youthful vitality,” observed economic historian Joel Mokyr. “As time

passes, the creative juices gradually dry up, and sclerosis sets in. Societies become increasingly risk-averse and conservative, and creative innovators are regarded as deviants and rebels.”⁶ The difficulty lies in keeping that sclerosis at bay, and ensuring that those “creative juices” that help drive technological progress and innovation don’t wither away. Touring the early American Republic, Alexis de Tocqueville expressed similar concerns, when he wrote of his fear “that men may reach a point where they look on every new theory as a danger, every innovation as a toilsome trouble, every social adventure as a first step toward revolution, and they may absolutely refuse to move at all.”⁷

While social and cultural pressures play no small part in safeguarding dynamism, the larger concern, and the one to which this guide focuses its attention, is the potential for, as Postrel noted, “reactionary ideals” and “technocratic administration” to become overly-protective vanguards of the status quo — forsaking the unknown possibilities of what *could be* for the contentment of what is. To that end, this guide is aimed at detailing how policymakers can embrace the fertile verge of progress, and the specific recommendations that will help the dynamic future unfold.

Part I begins with a discussion of dynamist-friendly approaches to administrative governance of the many emerging technologies already benefiting millions of Americans. This “soft law” governance order is currently the *de facto* system by which many regulators promulgate new rules to regulate those technologies that prove vexing for the traditional command-and-control regulatory machinery — machinery that was originally built to regulate 20th-century industries and which is now inadequate to keep pace with modern technological change. It goes on to examine the costs and benefits of this new governance approach, laying the foundation for an ongoing, interwoven discussion of soft law’s application to many of the emerging technologies discussed later.

The foundation of soft law is then followed in Part II with an examination of the many intersecting industries and policy issues implicated in

the ongoing maturation of the digital economy. This section will look at the issues of antitrust, privacy, and copyright: how they have (and have not) been changed by the post-Internet age, and what, if anything, policymakers should do to ensure the driving engine of American economic growth continues to hum along.

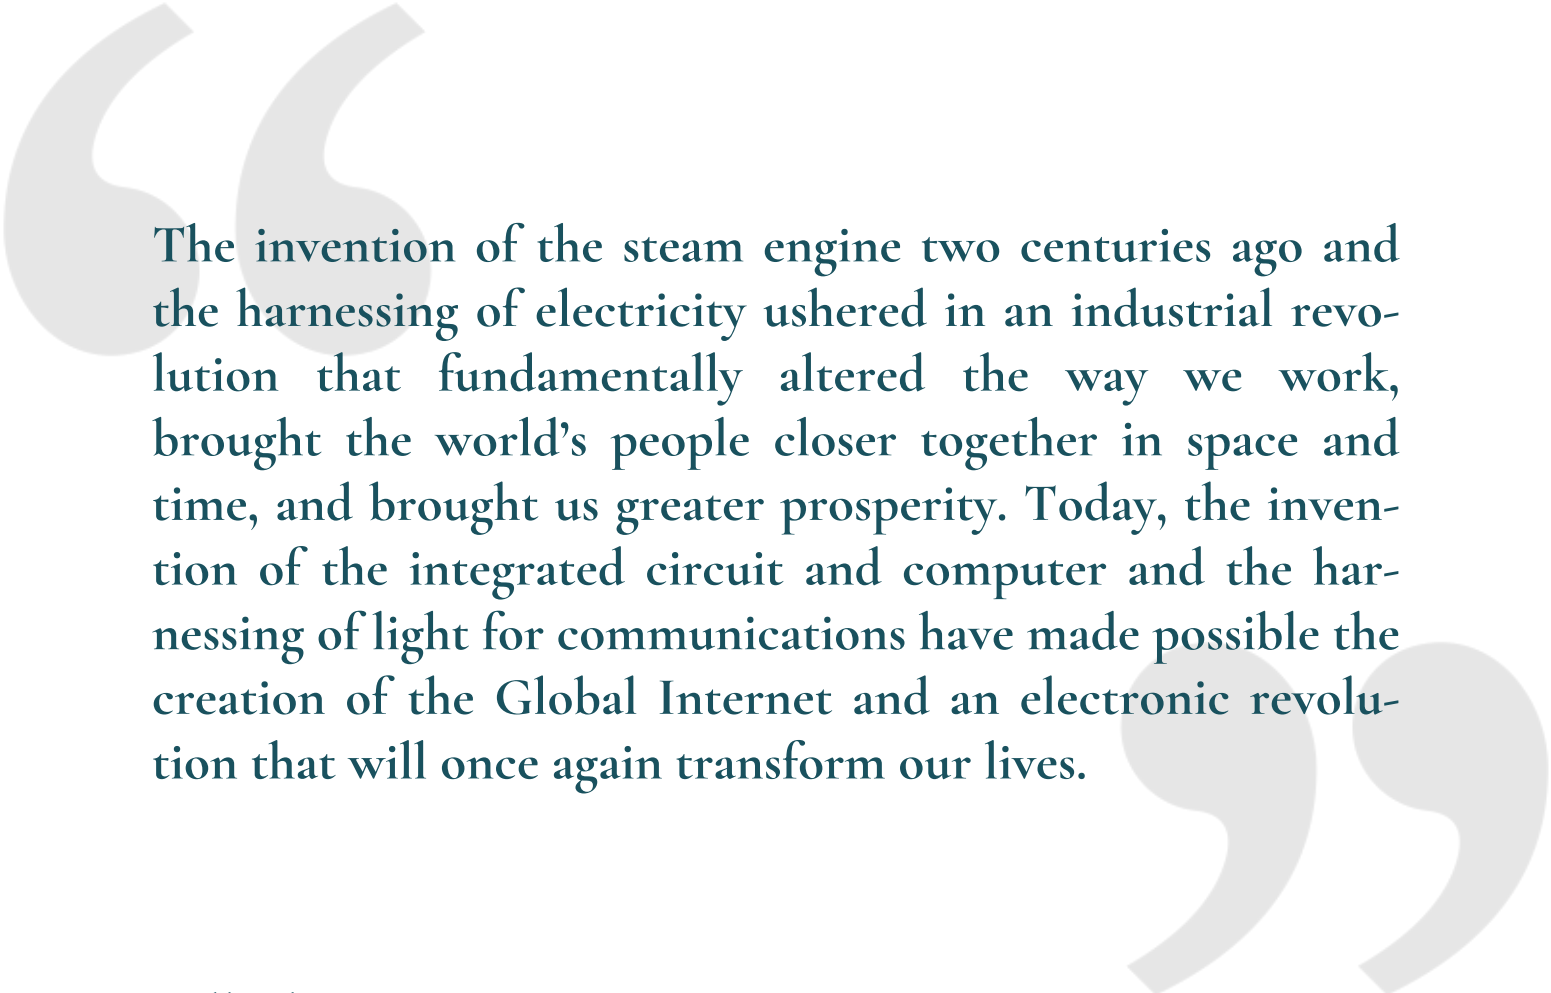
Part III details a number of specific emerging technologies, examining the issues and concerns related to their proliferation. Each section concludes by providing a set of detailed recommendations for policymakers confronting difficult issue-specific policy questions. This section is intended to provide actionable, politically pragmatic, and dynamist-friendly policy solutions to a wide range of new technologies.

Following on the heels of those specific recommendations, Part IV offers a more in-depth look at one of the most important technology policy issues of the day: artificial intelligence (AI). Like the Internet and digital communications technologies, AI is a “nexus technology” — that is, advancements in both basic R&D and AI’s practical application impact developments across a wide array of other industries and technologies. The disruption portended by AI is a flashpoint in the debate between dynamists and statists, with some perceiving its development as a potentially existential threat to human civilization, despite the many obvious quality-of-life improvements afforded by its proliferation. To address those purported concerns, this section will offer a sober analysis of the real-world impacts of AI, followed by a series of effective and innovation-friendly oversight mechanisms that can help address the real (not merely perceived) harms that may one day emerge as a result of its deployment.

Finally, the guide will conclude with a brief overview of why, in the face of the dynamic and unknown future, policymakers would do well to embrace the regulatory equivalent of the Hippocratic Oath: first, do no harm.

The set of recommendations provided in this guide is not a silver bullet for the problems plaguing technology governance. However, they can

nonetheless help effectively guide policymakers towards policies that will better balance the many complicated trade-offs associated with governing new technological realities. The extent to which it is successful in that endeavor will ultimately be reflected in how willing policymakers are to embrace the fertile verge of progress and pare back the regulatory burdens facing entrepreneurs, exporting the awesome pace of innovation in the world of bits to the world of atoms.



The invention of the steam engine two centuries ago and the harnessing of electricity ushered in an industrial revolution that fundamentally altered the way we work, brought the world's people closer together in space and time, and brought us greater prosperity. Today, the invention of the integrated circuit and computer and the harnessing of light for communications have made possible the creation of the Global Internet and an electronic revolution that will once again transform our lives.

— Bill Clinton

PART I: “SOFT LAW” IS EATING THE WORLD

Regulatory uncertainty in emerging technologies is a major problem confronting regulators and policymakers. The difficulties of attempting to provide certainty in an inherently uncertain landscape like technological development make old models of regulatory rulemaking nearly obsolete. Luckily, there exists a tried-and-true framework for agencies seeking to address concerns relating to new technologies. In 1997, the Clinton Administration released its *Framework for Global Electronic Commerce* (hereafter, *Framework*), which outlined a set of minimally burdensome policies for the emerging commercial Internet industry:

1. The private sector should lead;
2. Governments should avoid undue restrictions on new emerging technologies;
3. Where governmental involvement is necessary, it should support and enforce a predictable, minimalist, consistent, and simple legal environment for commerce; and
4. Government should recognize the unique qualities associated with new emerging technologies.⁸

In a 1997 interview, Ira Magaziner, the *Framework*'s lead architect, described the rationale for embracing a dynamist-friendly policy position on this still-emerging technology:

*Nobody really knows for sure where things are headed. On the other hand, government, by its nature, moves somewhat slowly and is somewhat inflexible when it does move. That may not be appropriate for the pace of the Internet and for the changeability of the Internet.*⁹

That prescience and foresight were, in retrospect, extremely well advised. Since then, the Internet has become one of the primary drivers of worldwide and domestic economic growth. The *Frame-*

work has been so successful that the Department of Commerce recently reaffirmed it during the course of its research into the Internet of Things (IoT).¹⁰ This success is largely because of policies that prioritized a flexible, adaptive, and light-touch approach to regulation — an approach that has come to be known as “soft law” governance.

Legal scholars Gary Marchant and Braden Allenby define soft law as “instruments or arrangements that create substantive expectations that are not directly enforceable, unlike ‘hard law’ requirements such as treaties and statutes.”¹¹ These “instruments or arrangements” can take many forms, from policy guidance documents to best practices and voluntary standards. They may begin as the result of proactive efforts to address potential concerns by industry consortia; multistakeholder engagements convened by government agencies; workshops organized by academic research centers; or policy briefs and white papers organized by think tanks. As shorthand, these outputs and deliverables can be characterized as “soft criteria” — the manifestation of soft law principles that aid in the creation of those “substantive expectations that are not directly enforceable.” Put another way:

*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.*¹²

The common thread that unites all of these informal soft criteria is that they rely on decentralizing regulatory governance responsibilities among institutions and actors outside traditional state regulators. At first glance, such an approach to regulation can seem chaotic and ineffective. After all, the soft law system lacks the procedural formality associated with more traditional rule-making processes, such as the notice-and-comment system outlined under the Administra-

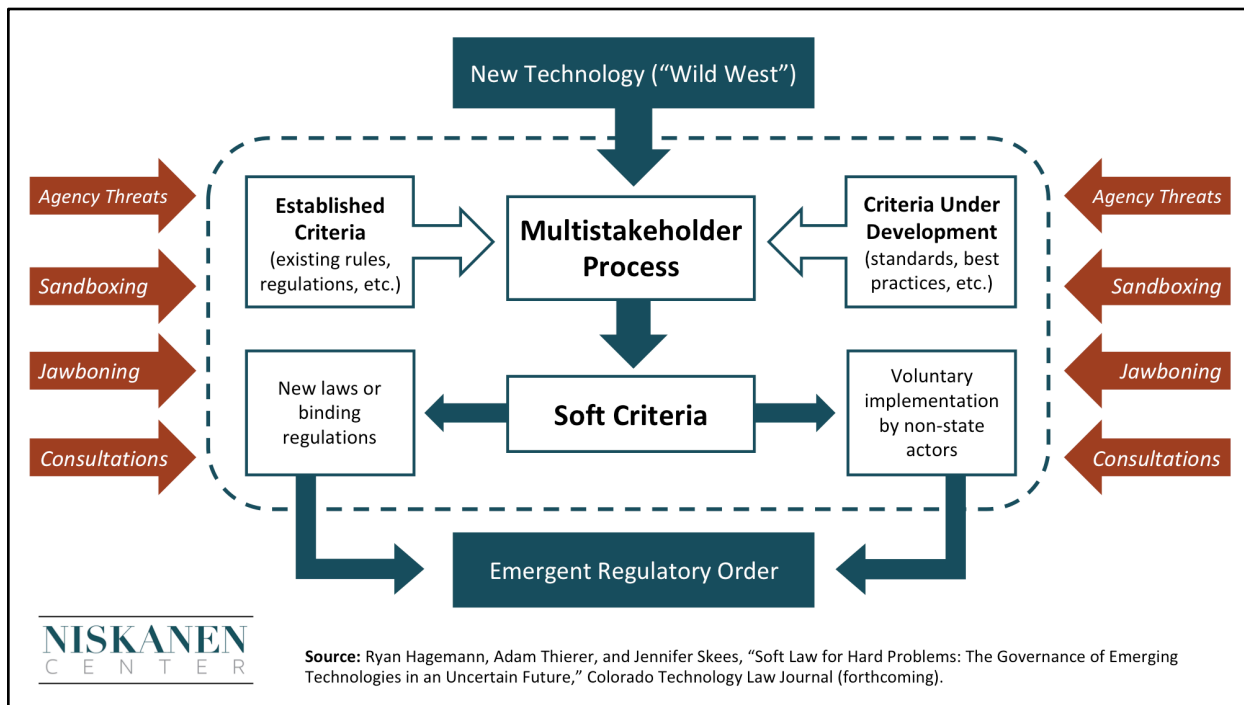


Figure 1: Providing a general overview of the soft law system for emerging technologies, and the potential pathways by which new rules emerge to regulate a new technology. “Agency threats,” “sandboxing,” “jawboning,” and “consultations” refer to specific means available to regulators that can be used to cajole industry and innovators into adjusting their behavior, but which seldom carry the force of law. This exogenous pressure from regulators constrains the contours of what transpires in soft law proceedings, like the multistakeholder process, and outputs, such as the promulgation of soft criteria. For more details, see Ryan Hagemann, Adam Thierer, and Jennifer Skees, “Soft Law for Hard Problems: The Governance of Emerging Technologies in an Uncertain Future.”

tive Procedure Act. However, there is a method to the apparent madness of soft law that roughly tracks along the paths detailed in Figure 1.

Soft criteria can gestate in any number of institutional settings — from the academy to industry trade associations. They do not begin to mature, however, until they pass through some collaborative process that lends legitimacy to their application. Those proceedings — multistakeholder processes — form the essential underpinning of the entire soft law system, serving as a gatekeeping forum for filtering out those soft criteria that cannot pass the muster of achieving some degree of consensus among a plurality of stakeholder participants. Multistakeholder efforts can be thought of as quasi-democratic assemblies or a type of “regulatory town hall,”¹³ where various actors from industry, government, civil society, trade associations, and other interested parties can air their perspectives and work together to achieve an outcome that is acceptable to all participants. The soft criteria that emerge from the multistakehold-

er process then go on to create a new emergent regulatory order based on one of two general pathways. Soft criteria either: (1) take the form of best practices or standards that inform voluntary self-regulatory frameworks, enforceable through a complex ecosystem of oversight and accountability housed both within and apart from the formal state regulatory apparatus; or (2) materialize as, or inform, new legislation or statutorily-binding rules. Over the past quarter-century, the general trend in the technology sector has followed the former path, and for good reason.

Soft law governance approaches provide significant benefits to the rapid innovation and progress that characterizes the technology industry. The soft criteria that emerge from collaborative, co-regulatory multistakeholder proceedings offer greater flexibility than rules produced through the notice-and-comment process of more traditional informal rulemaking. In an era where law and policy struggle to keep pace with rapid advancements in technology, standards and oversight mecha-

nisms that can better adapt to changing circumstances are necessary to maintain high levels of innovation and economic growth. Additionally, by crafting rules through a more collaborative, deliberative, and transparent process, soft law systems promote greater trust among stakeholders and incentivize compromise. That is, the multi-stakeholder process injects greater resiliency into the governance process by decentralizing oversight responsibilities among a wider set of interested parties; this also has the added benefit of providing a political steam valve for policymakers and regulators who may otherwise respond to knee-jerk public reactions calling for excessively precautionary government action in advance of known harms.

Whether the economic and social benefits of this system are a net positive to the broader economy is unclear, but as applied specifically to the technology sector, there is strong evidence to suggest the benefits outweigh the costs.¹⁴ Of course, while soft law provides many benefits for innovators and entrepreneurs, it is also open to the many criticisms associated with taking quasi-rulemaking functions away from federal agencies. Perhaps the most significant potential cost implicated by a soft law governance regime is the potential for lingering uncertainty to hinder investment and commercialization. As Robert Hoerr, co-founder of the nanotechnology pharmaceutical company Nanocropoeia, notes:

*Prolonged regulatory ambiguity is a cause for concern because markets place a high value on risk mitigation and predictability of outcomes. Developing innovation technology requires capital from venture capital investors who are comfortable with the risk of complete failure in exchange for the substantial rewards of success. Uncertainty in the regulatory environment has the potential to increase both the costs and time needed for development, thereby making the commercialization process unpredictable and, in the worst case, incapable of being financed.*¹⁵

Soft law-style regulatory systems are not a new creation of the digital age. Indeed, the Food and

Drug Administration (FDA) has been regulating according to non-binding policy guidance documents since the early years of the 20th century.¹⁶ Nor is the use of such governance approaches confined to the United States. The Netherlands, for example, has a long history of governing by cooperative tripartite negotiations to achieve consensus-based socio-economic policymaking. The “Polder Model,” as it has come to be known, emphasizes a distribution of shared governance responsibilities between the state and its social partners in civil society (in particular, corporations and labor unions).¹⁷

The old ways of regulating will not suffice in this new age of rapid technological progress, in which technological change consistently outpaces the capabilities of agencies — agencies that were designed to address different problems for a different era. The best way for the United States to avoid losing ground on developing advanced technologies is to retain its commitment to a relaxed regulatory system for the digital economy, and engage in a concerted effort to export soft law mechanisms to regulatory agencies managing non-digital sectors, such as health care and manufacturing.

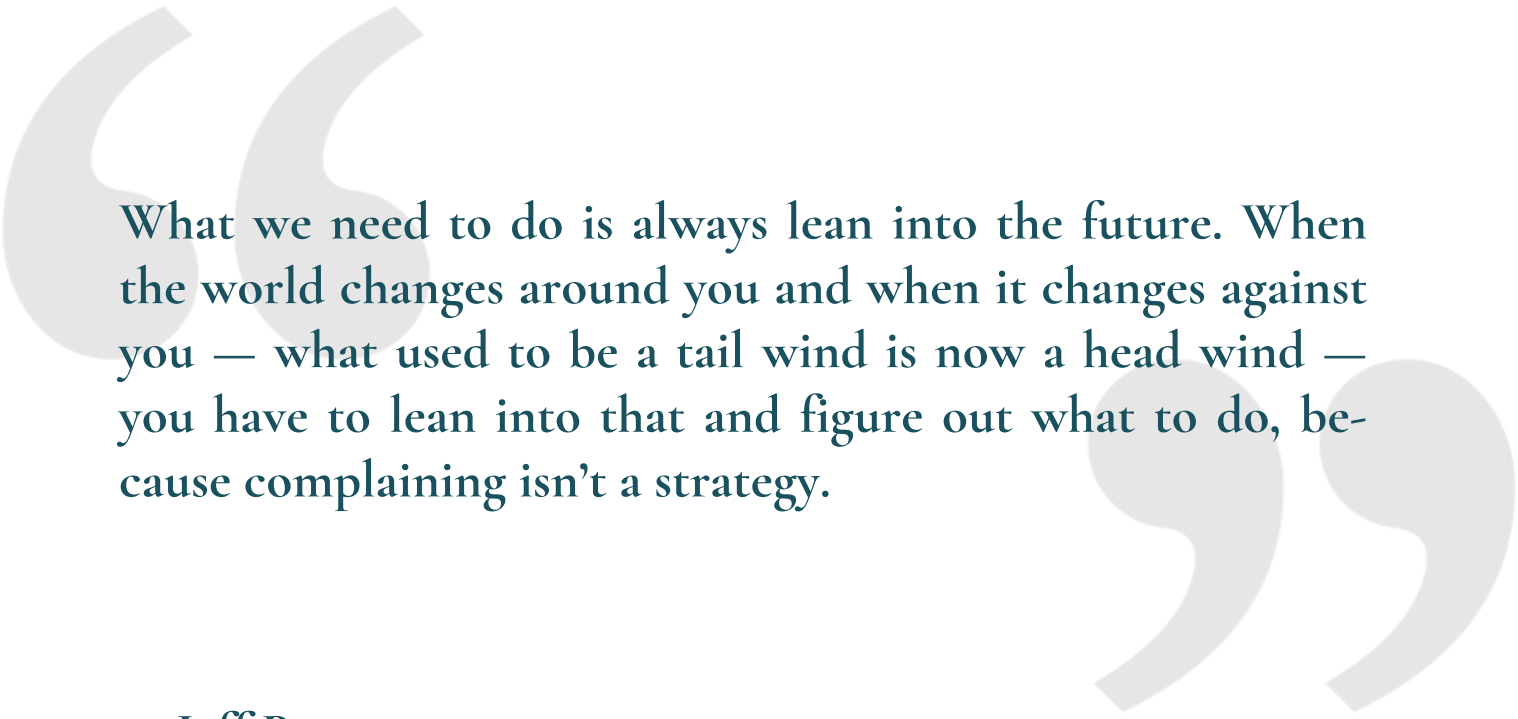
In particular, when considering new rules governing emerging technologies, policymakers and regulators should reaffirm their commitment to the principles outlined in the *Framework* and apply them to emerging technologies. However, embracing the *Framework*’s principles is merely the starting point in signaling a commitment to soft law governance.

In a 2011 essay for the *Wall Street Journal*, venture capitalist Marc Andreessen famously observed “software is eating the world.”¹⁸ His insight was spot-on, and many of his predictions have come true. Now, soft law is eating the world of technology regulatory governance. As described in a recent essay at *The Bridge*, this state of affairs was a direct outgrowth of Andreessen’s predictions surrounding software:

Andreessen’s point about software devouring everything was based on the realization that

the underlying drivers of the digital revolution — massive increases in processing power, exploding storage capacity, steady miniaturization of computing, ubiquitous communications and networking capabilities, and the digitization of all data — were rapidly spreading to other sectors of the economy. No longer was information technology or software its own isolated silo, but industries from farming to communications were being revolutionized by it.¹⁹

As software has spread to every corner of the economy and global society, it has become necessary to govern and regulate using a set of tools and norms better suited to a world in which the rate of innovation now sets the terms of governance realities. The future implications of this new world are profound, even as soft law has already created a set of de facto regulatory strategies for a wide swath of federal agencies confronting the challenges posed by innovation progress in new technologies. Part III will examine many of the individual technologies currently presenting difficult questions for regulators and policymakers and offers a set of tailored, soft law-inspired policy recommendations for addressing those concerns. First, however, it is worth examining how hard law regulatory models continue to impact the ongoing growth and development of the digital economy.



What we need to do is always lean into the future. When the world changes around you and when it changes against you — what used to be a tail wind is now a head wind — you have to lean into that and figure out what to do, because complaining isn't a strategy.

— Jeff Bezos

PART II: INNOVATION IN BITS

Since the dawn of the Internet, the digital economy has been increasing in importance relative to more traditional economic sectors. According to the Bureau of Economic Analysis, in the ten-year period between 2006 and 2016 “the digital economy grew at an average annual rate of 5.6 percent, outpacing overall U.S. economic growth of 1.5 percent per year.”²⁰ While this transition has led to vast increases in welfare, it has also raised critical questions about the role of antitrust, privacy, and copyright policies. None of these issues are new, but emerging technologies have prompted many to view them through a different lens.

In a 2014 essay for *Foreign Affairs*, Craig Mundie examined the current privacy fights in relation to the early days of the Internet, considering what might not have been if the government had embraced more prescriptive privacy rules:

*If, in 1995, comprehensive legislation to protect Internet privacy had been enacted, it would have utterly failed to anticipate the complexities that arose after the turn of the century with the growth of social networking and location-based wireless services. The Internet has proven useful and valuable in ways that were difficult to imagine over a decade and a half ago, and it has created privacy challenges that were equally difficult to imagine. Legislative initiatives in the mid-1990s to heavily regulate the Internet in the name of privacy would likely have impeded its growth while also failing to address the more complex privacy issues that arose years later.*²¹

Had the United States moved in a less-permissive direction, the innovative capacity of the digital economy may have been strangled before it had time to mature. The last 25 years have been a natural experiment — with the European Union as the treatment group and the United States as the control group — in learning how prescriptive rules for antitrust, privacy, and copyright regulation affect economic and technological outcomes.

Since the 1997 *Framework*, the United States has done an excellent job providing a predictable and minimalist regulatory environment for most emerging technologies in the world of bits. This philosophy runs counter to the strategy embraced in Europe, which prioritized prescriptive rules for data use, higher taxation, and broad precautionary approaches to addressing perceived risks. The best evidence for the success of the American model is that seven of the top 10 largest public or private Internet companies in the world — Apple, Amazon, Microsoft, Google, Facebook, Netflix, and eBay/PayPal — are in the United States. In contrast, not a single top 20 Internet firm is from Europe.²² Going forward, the United States faces a choice: continue on the path that has delivered prosperity and innovation, or follow Europe's lead in burdening industry with excessively burdensome regulations meant to prevent every potential harm. In other words, American policymakers can continue embracing dynamism or they can attempt to reengineer the socio-economic order to cater to stasis impulses. As the following sections will argue, the former is far preferable.

A. Antitrust

Rising concentration in some industries has sparked a renewed interest in antitrust policy in the United States.²³ But is big business bad *per se*? This line of reasoning is known as the “structuralist” school of thought and it was dominant in antitrust scholarship until being supplanted by the consumer welfare standard in the 1970s. Law and economics scholars, centered at the University of Chicago, argued that antitrust intervention should require evidence that consumers have been harmed by anticompetitive corporate behavior intended to increase profits. Now, structuralism is experiencing a revival in the antitrust debate with the rise of the Hipster Antitrust movement, which aims to use competition enforcement as a means of achieving progressive policy goals.²⁴ However, under the prevailing interpretation of American antitrust law, the consumer welfare standard still reigns supreme, which constrains the definition of anticompetitive behavior to situations in which a firm raises prices and reduces

output beyond the competitive level, thereby increasing profits at the expense of harms to consumers.

1. The American Model

The technology industry, led by the Big 5 (i.e., Apple, Google, Microsoft, Amazon, and Facebook), is no exception to the rising concentration observed in other sectors of the economy. However, this issue is not as simple as it might seem at first. In the digital economy, many companies price their services at zero and provide virtually unlimited output. This unique business model presents a challenge for the traditional regulatory policy toolkit. Thus far, policymakers have continued to review proposed mergers using a consumer welfare standard and have considered intervention on a case-by-case basis.

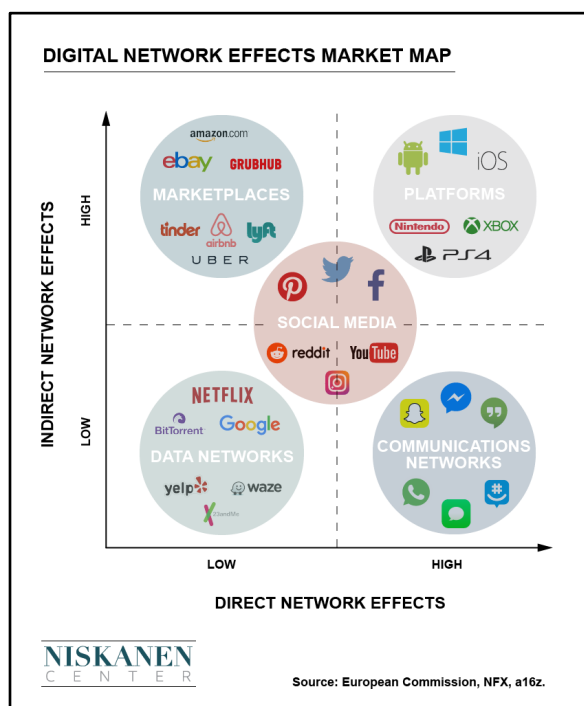


Figure 2: Digital Network Effects Market Map. *Source:* Alec Stapp, “You Can’t Understand Big Tech Without Understanding Network Effects.”

The unique nature of online network effects also plays a role in these analyses. Network effects refer to the phenomenon where the value of a good or service increases with the total number of users. These can take two forms: direct and indirect. Direct network effects occur when the addition of

a new user of a service provides a direct increase in the value of that service for other users. Examples in the digital economy include services like iMessage, Google Hangouts, WhatsApp, and Snapchat. The classic example of an analog era technology exhibiting direct network effects would be the telephone.

Indirect network effects are when the value of a service to one group of users increases as the size of a complementary group of users increases. This is one of the essential features of multi- or two-sided markets, such as operating systems and social media platforms. For example, as more users adopt a particular operating system — say, iOS versus Android — application developers might prioritize building products for that operating system, which in turn further fuels user demand for that platform over an alternative. An example of this in a pre-digital context would be brick-and-mortar marketplaces, such as shopping malls and flea markets.

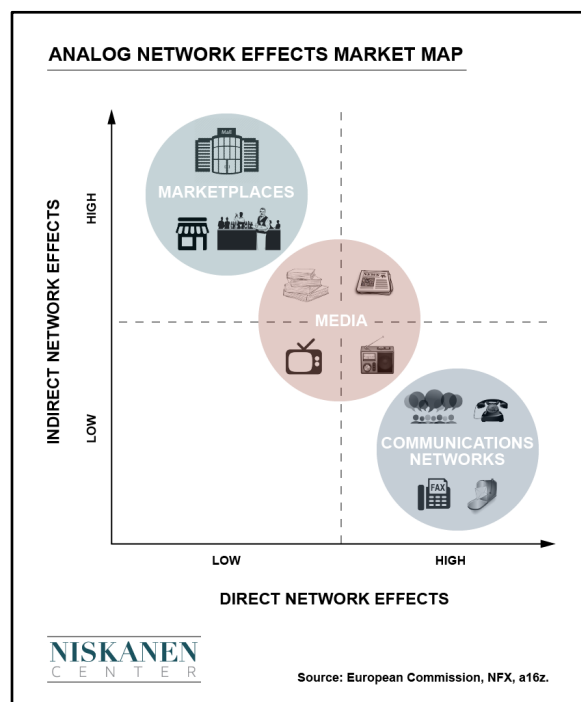


Figure 3: Analog Network Effects Market Map. *Source:* Alec Stapp, “You Can’t Understand Big Tech Without Understanding Network Effects.”

Although some have expressed concern regarding the potential negative consumer welfare implications of network effects locking incumbent firms

into monopoly positions, there are many reasons to believe network effects in the digital age are actually far weaker than they were in a pre-Internet world, including: (1) diminished switching costs due to the proliferation of physical devices and digital platforms, (2) more localized network effects that advantage smaller networks over their larger competitors; and (3) the incentives for users to go off-platform for repeat business or high-value transactions.²⁵

There are three reasons that a broader analysis of the competitive landscape in the technology industry militates in favor of fewer antitrust enforcement actions. First, the Big 5 were all among the top 20 companies *in the world* in total R&D expenditures in 2017, and four were in the top 10.²⁶ It seems that the leading tech giants are not growing complacent during their current period of prosperity. Second, while the incumbents' network effects and vast troves of consumer data may be barriers to entry for new firms, they are no more significant in scale than the substantial upfront capital investment required to build a factory or stock of inventory. Lastly, in general, entry into the software industry requires relatively little capital: anyone with a computer and Internet access can launch a startup and compete in the market.

2. The European Model

In Europe, given the size of recent fines against large American technology firms, securing a dominant market position at the expense of weaker rivals can be enough to warrant enforcement action. This framework presents two problems for sensible policy. First, it is not always the case that a firm with a large market share has gained market power. Small competitors and the threat of new entrants might constrain the dominant firm from profitably raising prices.

Second, regulators need to be extremely careful about how they define markets. For example, in the recent enforcement action against Google for the dominance of its Android operating system, the European Commission defined the market as only *licensable* mobile operating systems, which

excludes Apple's proprietary mobile operating system. But to anyone who has shopped for a smartphone can attest, Apple and Android phones are indeed competitors.

Also, consider how European regulators are laying the groundwork for a future antitrust case against Facebook. If and when that case comes, the key fact differentiating a judgment for or against the platform will be whether the market is defined as "social networking" or "entertainment." The debate over these companies is not going to quiet down any time soon, but policymakers should be wary of following the European model of punishing Big Tech for its success.

3. Policy Recommendations

- a. **Recommit to the consumer welfare standard for evaluating claims of anti-competitive behavior;**
- b. **For antitrust claims in two- or multi-sided markets, determine overall welfare effects across customer groups if and only if there are two distinct customer groups with a significant interdependence of demand; and**
- c. **Identify the negative and positive network effects for each platform individually and assess whether the effects are diminishing, increasing, or constant before determining market power.**

B. Privacy

As more economic and social activity moves online, people are increasingly worried about how much data technology companies collect and what they do with that data. But there is a paradox at the heart of privacy questions in the digital economy. On surveys, people will say that they highly value their privacy, but in practice they will give away their data in exchange for nominally free services. So is this just another example of cheap talk? In other words, do revealed preferences tell us more than stated preferences? Experimental evidence has shown that, indeed, people will trade

their privacy for relatively small amounts of money. They also exhibit predictable behavioral biases when it comes to valuing privacy. For example, people value privacy less in cases where the personal information is socially desirable (i.e., it reflects well on them) and more in cases where the data is socially undesirable (e.g., medical problems or criminal records).²⁷ How policymakers interpret this evidence and weigh privacy against other values will determine the future of the digital economy.

1. The American Model

The U.S. Constitution does not explicitly guarantee a right to privacy. However, Supreme Court cases over the years have cobbled together an implicit right to privacy against government intrusion from the First, Third, Fourth, and Fifth Amendments. Sector-based legislation, which addresses real harms to sensitive data, further protects consumers' right to privacy. Regulating privacy by industry — whether it is finance, health care, or education — is consistent with the empirical evidence regarding how consumers contextually value their privacy. The sector-based approach stands in contrast to an omnibus-style approach of setting national privacy standards that apply to all businesses, regardless of relative risks to consumers' privacy.

In the age of the Internet, encryption is the ultimate enabler of digital privacy. Encryption is what allows people to know their online banking activities, commercial transactions, and the exchange of other sensitive data is secure. This knowledge helped grease the wheels of online commerce and sparked the explosion of the digital economy over the last few decades. Without encryption, the online experience is fundamentally insecure. However, law enforcement has expressed valid concerns over the trend in digital communications moving beyond their reach — a phenomenon referred to as “going dark.” As more and more information becomes encrypted by default, law enforcement agencies will find it increasingly difficult to keep up. Legal battles like the 2016 Apple-FBI fight over the San Bernardino terror-

ist's encrypted iPhone are likely a sign of things to come in this debate.

2. The European Model

The European approach to privacy governance is most clearly demonstrated in its recent implementation of the General Data Protection Regulation (GDPR). This new package of regulations includes broad privacy mandates that apply to any business that deals with consumer data (which is to say all businesses). These rights include data portability, access, correction, erasure, and consent. While noble in intent, the new rights enumerated in this regulation are vague, ambiguous, and entail large compliance costs. For instance, the 500 largest companies in the world will spend roughly \$7.8 billion to ensure compliance with the new law.²⁸

Beyond the direct compliance costs and difficulty of interpreting ambiguous mandates, there is an inherent trade-off between privacy and prosperity in this domain that many are loath to accept. If people genuinely care about privacy more than other values, then it might be optimal from a welfare perspective to sacrifice some growth and innovation in exchange for more privacy. But policymakers should remember that the U.S. digital economy directly supports 5.9 million jobs and accounts for 6.5 percent of GDP, and omnibus privacy regulations have real costs for employment and incomes.²⁹

3. Policy Recommendations

- a. **Maintain sector-based privacy regulation for sensitive medical, educational, and financial information;**
- b. **Preempt state regulation of online privacy by passing baseline federal privacy regulation that is predictable, minimalist, consistent, and simple; and**
- c. **Firms suffering a data breach should be required to coordinate with law enforcement and notify affected users within a reasonable period of time.**

C. Copyright

Artists, content creators, and intellectual property holders currently have a right to exclude others from copying their work and selling it for a profit. Online, where the marginal cost of reproduction is functionally zero, this is difficult to achieve due to the ease of Internet piracy and file-sharing sites. The United States has a strong history of enforcing intellectual property rights to incentivize future innovators and artists. However, overzealous restrictions on intellectual property can also inhibit those same creators from building on previous work to make something new. The key to optimal copyright policy is finding the right balance between incentivizing new work and allowing old work to be remixed and reimaged.

1. The American Model

In the United States, the policy toward online service providers (OSPs) has struck a balance between protecting intellectual property rights and allowing innovative business models. According to Section 230 of the Communications Decency Act (CDA) of 1996, “No provider or user of an interactive computer service shall be treated as the publisher or speaker of any information provided by another information content provider”³⁰. In other words, platforms that allow users to upload media are not liable for content that violates copyright. The only duty of the website operator is to take down an infringing piece of media if the copyright owner submits a notice regarding the violation.

There is a separate longstanding copyright issue, commonly found in music, in which one artist will copy parts of another’s work without seeking permission or paying a licensing fee. Recent examples include Robin Thicke’s song *Blurred Lines* infringing on Marvin Gaye’s *Got to Give It Up*³¹ and Vanilla Ice’s *Ice Ice Baby* infringing on David Bowie and Queen’s *Under Pressure*.³² Technology could hold the answer for improving the baseline protections offered by Section 230 of the CDA. If a copyright lawsuit reaches trial, judges task jurors with using their subjective judgment to determine whether one artist has stolen from another. Con-

tent recognition systems (CRS) have the potential to automate this process and make it more objective.


For example, Google already uses algorithms on YouTube to detect whether an uploaded video violates copyright. The software can consistently match a new piece of content to any piece of work stored in its database. That leaves only two questions: (1) how accurate are the algorithms, and (2) do they provide reasonable results? When a CRS detects copyright infringement, the system could also automatically offer an option to pay the owner a predetermined licensing fee to use their work. These payments would be a new source of revenue for artists who currently miss out on them due to high transaction costs.

2. The European Model


The European Union recently made drastic changes to its copyright law.³³ The two most controversial changes in the law both augmented the rights of copyright holders at the expense of digital platforms and aggregators. The first is Article 11, which requires Internet companies to pay licensing fees to newspapers, magazines, and other media companies for using snippets of their work in aggregator services. The second is Article 13, which creates an inverse of the legal burden under Section 230 in the United States by making platforms liable for user-uploaded content that violates copyright. This burden-shifting makes Internet companies much more risk-averse and hampers the free exchange of ideas online. Compared to the American approach, the costs of the new changes to European law likely outweigh the benefits.

3. Policy Recommendations

- a. **Maintain safe harbor for digital platforms from copyright claims in any new copyright legislation; and**
- b. **Promote automated CRSs as a voluntary best practice, but do not make safe harbor protection contingent on its implementation.**



I'm interested in things that change the world or that affect the future, and wondrous new technology where you see it and you're like, 'Wow, how did that even happen? How is that possible?'



— Elon Musk

PART III: INNOVATION IN ATOMS

The *Framework's* general success in governing the world of bits was predicated on the lack of existing institutions and regulatory mechanisms that could provide an easy answer to the question of how the Internet and digital communications channels should be governed. An early stasis-styled proposal suggested the need for a National Information Infrastructure (NII) — a federal agency that would regulate and guide the development of the Internet.³⁴ Luckily, the more formalized, agency-centric approach proposed by the NII morphed into the *Framework*, eschewing technocratic regulatory oversight for more decentralized, market-oriented policies. Had proponents of the original NII won out in the early 1990s, it is unlikely the Internet would have developed into the global juggernaut that helped pave the way for the emergence of the modern digital economy. While the world of bits has benefited from a soft law governance alternative, the world of atoms has unfortunately lagged behind its digital counterpart.

This is not to suggest that progress in non-digital sectors has been completely stifled. Far from it. We live in a world where advancements in genetic modification technologies, autonomous vehicles and commercial drones, supersonic flight, and many other emerging technologies have been accelerated by the digital economy. But the breakneck pace of progress and technological advancement we've seen in the digital realm has not been matched in the physical world. A significant reason for this is the nature of legacy regulations in health care technology, ground and air transportation, and other industries, which make it far more difficult to disrupt entrenched incumbents built up over many decades. Developing a new app that can be replicated and distributed at low marginal cost across an interconnected and always-on global communications infrastructure is far easier than trying to enter the biopharmaceutical market, for example. And the difficulties of entry and disruption are only compounded by the

existence of regulatory gatekeepers who demand fidelity to a set of rules and standards crafted in — and designed for — a pre-digital world.

As discussed in Part I, however, there is a great deal of hope for innovation and progress in the world of atoms. This optimism is due in part to soft law's permeation throughout almost every precinct and quarter of the regulatory state. Even as agencies continue to make use of soft criteria and other non-binding regulatory pronouncements, there is much that policymakers can do to improve the governance of new technologies while more effectively balancing the trade-offs between protecting the public interest and accelerating the research, development, deployment, and ultimate adoption of new life-saving and life-improving technologies.

A. Genetic Modification

The history of medicine is a story of reaction and scarcity. Sick people enter our medical system, and a network of doctors, machines, and institutions react by providing the best care possible given limited time and resources. But what if we could prevent people from getting ill in the first place? And in unpreventable cases, what if we could use technology to stimulate the human body's regenerative process instead of relying on external therapies?

In the transition from reactive to proactive medicine, genetic modification technologies such as CRISPR offer the greatest hope yet of a healthier future. With the ability to directly edit the germline — the biological cells that pass on their genetic material to offspring — we might one day be able to render numerous genetic disorders to the dustbin of history. Researchers can accelerate efforts in genomics by using AI to connect genomic variation between individuals to phenomic variations (i.e., observable physical characteristics). To reach that future, policymakers need to consider the trade-offs of this new technology rationally. Fears about “designer babies” and superhuman soldiers are overblown, as the science is still in its early stages and we remain ignorant

about how editing DNA to select for one trait affects other traits.

Regenerative medicine and gene therapy treatments will be keystones in the foundation of a post-scarcity health care world. Therapies developed from this branch of research hold the promise of growing tissues or organs in a laboratory from a patient's tissues or cells. By further leveraging bioengineering technologies we might be able to solve the problems of organ transplant rejection and organ donation shortages forever.

1. Insurance Reimbursements for Gene Therapy Treatments

According to a recent study in the *Journal of Managed Care Medicine*, “There are now more than 930 clinical trials of regenerative and advanced therapies currently underway, with approximately 630 in Phase II to III and nearing market entry.”³⁵ Paying for gene therapy treatments, however, may prove to be the greatest challenge for widespread use and adoption, especially given the traditional insurance business model and reimbursement process.

Gene and cell therapies have curative potential because they target underlying biology. Cures have high clinical value for patients and high economic value for manufacturers. However, the potential adoption surges upon FDA approval are difficult to forecast and might blow up insurers' budgets. The insurance reimbursement process is currently designed to pay for drugs monthly rather than in one large upfront payment.³⁶ Without repeat administrations, the one-time price needs to be high, and there is a potential that patients will not be able to switch to alternative therapies easily. These products in this category are complex: They involve viral manufacturing, autologous cell processing, and a specific route-of-administration. For insurers, this level of complexity adds technical uncertainty to the reimbursement approval process, possibly requiring new and specialized codes and working with specialized centers of excellence.³⁷

Given the breakthrough potential of regenerative medicine, private and public insurance providers have an imperative to review their traditional reimbursement processes to accommodate these revolutionary therapies. For example, for a gene therapy treatment to qualify for new technology add-on payments (NTAPs) from Medicare, a manufacturer must prove that it is new, high cost, and provides a substantial clinical improvement in diagnosis or treatment.³⁸ However, even if a manufacturer meets these criteria, the NTAP designation still doesn't allow health care providers to recoup the full cost of the treatments: “Medicare pays a marginal cost factor of 50 percent for the costs of the new technology in excess of the [predetermined Medicare Severity Diagnosis Related Group (MS-DRG) payment rates].”³⁹

Fortunately, there are policy solutions to alleviate these problems. Manufacturers and insurers could agree to implement what are known as pay-for-performance, pay-for-outcomes, or value-based pricing systems. In practice, that means insurers would only pay for treatment if patients show improvement within a limited timeframe following administration. However, the Medicaid “best price” rule may prevent manufacturers from tying price to an outcome. If so, new legislation might be necessary to address this concern.⁴⁰

Aside from value-based pricing, manufacturers could work with insurers to set up installment payment plans to spread the cost over time.⁴¹ Furthermore, insurer value models currently only include the effect of a treatment on quality of life and direct medical costs. Given the transformative potential of gene therapies, models should be updated to include indirect costs such as unemployment and forgone caregiver wages.⁴²

2. Improving the FDA Regulatory Approval Process

Despite the unprecedented level of technological disruption causing headaches for many government regulators, the FDA and the health care sector more broadly have proven exceptionally resilient to the type of disruptive innovation that has infiltrated and upended many other quarters of

government and industry. While there are likely many explanations for this state of affairs, this is an arena that is nonetheless primed to benefit from many crosscutting innovations and technological developments, from the use of IoT devices analyzing patient responses to AI systems that can more accurately diagnose a wide variety of ailments. Before that can happen, however, the FDA needs broad updates to its approval process for new drugs, biologics, and gene therapy treatments.

The clinical trial approval process currently poses a significant regulatory cost for all but the most entrenched incumbent biopharmaceutical firms. The average length of a new drug approval — from application to final approval — takes between 10 and 12 years and can cost close to \$1 billion.⁴³ This reality poses an insurmountable hurdle for new biotechnology and genetic modification therapies startups, which seldom have the cash-on-hand to meet the rigorous, demanding, and often-unclear standards of approval set by the FDA's interpretation of what constitutes a safe and effective new drug or therapy. To the credit of Congress and FDA, much has been done in recent years to carve-out more expeditious pathways for the approval of new gene therapy treatments. The recent Regenerative Medicine Advanced Therapies (RMAT) designation, for example, wisely included new gene therapy applications under its definition.

For all the good, however, the FDA's mission has drifted considerably since the passage of the Federal Food, Drug, and Cosmetic (FD&C) Act, and no longer represents an appropriate balance between *resilience* and *adaptation* strategies — the FDA-specific variant of dynamism and stasis governance approaches, respectively. “Resilience” represents a favorable, optimistic disposition towards discovery, learning, and a decentralized trial-and-error process of experimentation. “Anticipation,” in contrast, is rooted in an aversion to risk and prefers utilizing centralized regulatory mechanisms to minimize errors, despite the impact on potential benefits. Indeed, the FDA has a long and unfortunate history of defaulting to anticipation strategies. Fortunately, the agency is

statutorily empowered to improve on this status quo — primarily by capitalizing on new technologies that can hasten the regulatory approval process, without sacrificing high standards for safety and effectiveness.

In a 2018 white paper, Niskanen Center adjunct fellows Dr. Joseph Gulfo and Jason Briggeman outlined a number of recommendations for how the FDA might achieve a reorientation back towards resilience strategies. In particular, they argued that it should be “widely understood that the FDA is to conduct reviews that assure safety and effectiveness when a product is used as labeled but do not require evidence of purported clinical utility.”⁴⁴ The reason, they argue, is that:

*Clinical utility is an elusive standard — it is tantamount to proving that there are, in some overall and ultimate sense, benefits to patient health from a product. Generally, even the best science cannot produce conclusive evidence on such a question, as attested by the many conflicting studies of the health effects of aspirin, for example. Aspirin is a safe and effective product, when used in accordance with its labeling — it generally delivers the promised effect to alleviate pain — but scientists continue even today to investigate whether taking aspirin is ultimately “good” when gauging different health risks, for different types of patients, over the long run, and so forth.*⁴⁵

The FDA should focus its approval efforts on drugs based on a more narrowly tailored interpretation of “safe and effective” and move away from determinations rooted in a misguided focus on “clinical utility” and “clinical outcomes.” This shift would not only provide valuable clarity for new startups but also minimize the costs associated with an excessively precautionary reliance on Phase II/III clinical trials. This change would also hand over the bulk of responsibility for determining clinical utility to the more appropriate realm of the medical marketplace: physicians and patients.

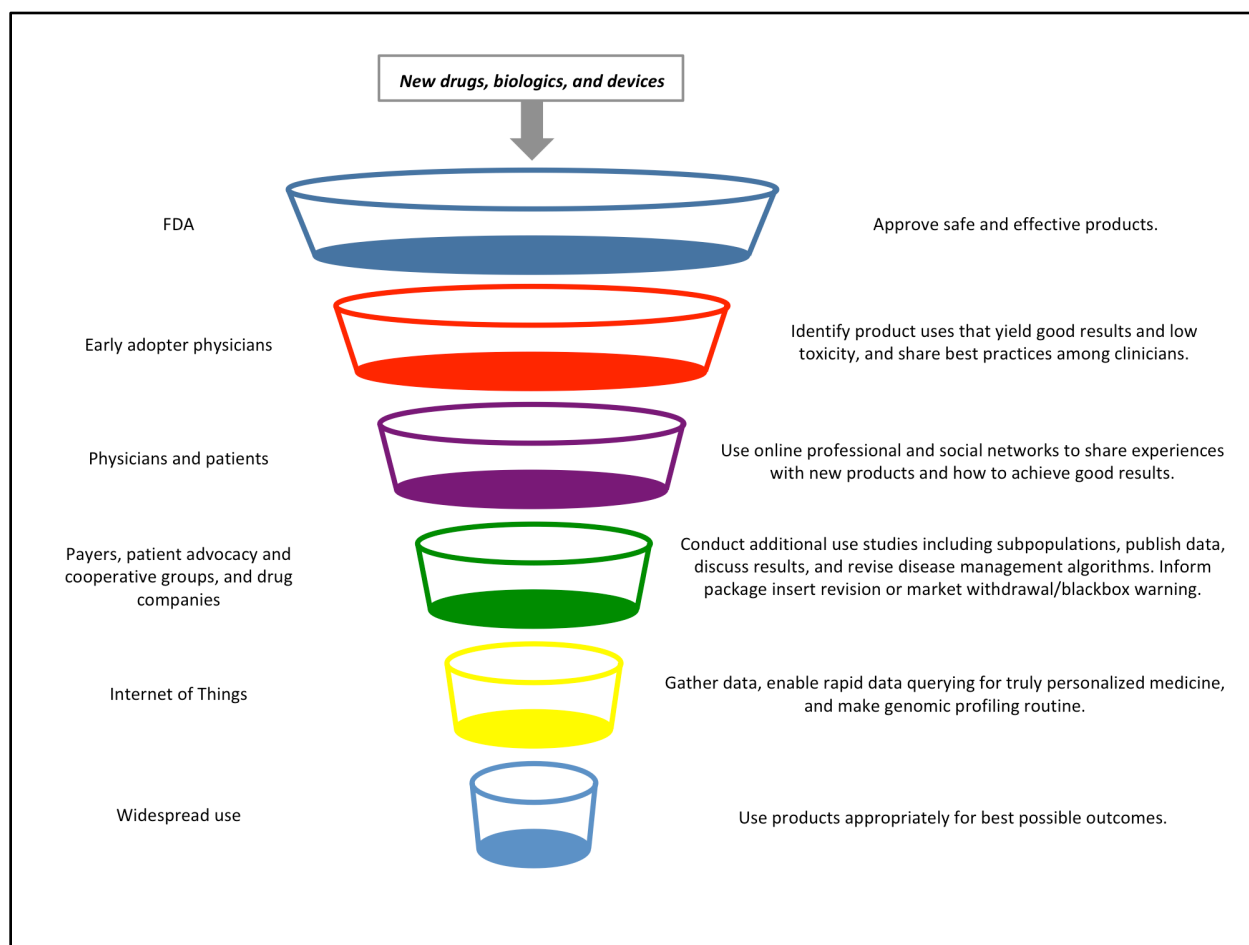


Figure 4: “The Medical Marketplace As It Should Be Today.” Source: Joseph V. Gulfo, Jason Briggeman, and Ethan C. Roberts, “The Proper Role of the FDA for the 21st Century,” p. 21.

The FDA can also improve its regulatory approval process by capitalizing on the emergence of new advanced technologies, such as the IoT and AI, which make postmarket monitoring of patient health outcomes far easier. As detailed in Figure 4, IoT devices and AI analytics tools can help expedite real-time and real-world data-gathering processes, allowing for more cost-effective clinical trials while also permitting the FDA to direct more resources away from premarket approvals and towards the postmarket surveillance ecosystem. This would have the dual benefit of: (1) promoting greater investment in, and competition within, the early-stage biotechnology and pharmaceutical startup space; and (2) better aligning the FDA’s resource allocations with real-world assessments of safety and effectiveness based on more individualized patient responses to new drugs, biologics, and gene therapy treatments.

Gulfo and Briggeman also recommend “that the FDA establish categories or orders of approval according to the nature of the evidence used to support effectiveness claims, a system that [they] believe will add substantial value atop the informational functions that premarket approval already serves.”⁴⁶ This would have the dual advantage of alleviating the FDA of the responsibility for dictating clinical endpoints to drug developers while leveraging the FDA’s comparative advantage and expertise in communicating relevant product knowledge to lower rungs of the medical marketplace, such as physicians, patients, and payers. They further argue that this approach:

Will also increase efficiency in the drug development process by enabling drug manufacturers to match the size, scope, duration, and goals of preapproval clinical studies with the

*claims being sought, in response to the demands of the medical marketplace — efficiency in addressing medical needs with new products, responding to the activities of competitors, and satisfying payer requirements.*⁴⁷

Thus, a regulatory approval process that prioritizes tiered categories of approval (described in further detail in Appendix A) would help save costs for innovators and regulators, expedite the regulatory review and approval processes, incentivize greater spending on new treatments that attempt to tackle diseases that afflict a wider proportion of the population (such as Alzheimer's and heart disease), and better position the FDA to capitalize on an emerging array of emerging technologies that are better-tailored to prioritizing precision medicine. More broadly, these reforms would also help more effectively constrain the FDA to its appropriate role in the hierarchy of the medical marketplace.

In a 2016 research paper for the Mercatus Center, Gulfo and Briggeman, joined by Ethan C. Roberts, also described the need for promoting a “more dynamic medical marketplace” in which the FDA is contained to “its proper role of permitting safe and effective drugs onto the market.”⁴⁸ Figure 4 shows how this reprioritization would be assisted by the use of advanced information technologies, such as the IoT, to better “enable rapid data querying for truly personalized medicine,” while helping to “make genomic profiling routine.” To achieve this new, more optimal ecosystem, they argue Congress needs to help guide the FDA back to its core statutory mission by, among other things:

1. Limiting the FDA's authority to consideration of a new drug's safety as it relates to an intended use, and according to its labeling. In particular, they note that “FDA reviews should not be permitted to speculate about the safety of off-label uses or of uses in populations beyond those the label indicates;”⁴⁹

2. Constraining the definition of “safety” to “the likelihood of causing death, debilitation, or severe harm;”⁵⁰ and
3. Defining “effectiveness” as a situation in which some “positive activity on the disease” can be shown.⁵¹

Whether the corpus of these reforms are enacted through soft law mechanisms or mandated through legislation will ultimately depend on the willingness of the FDA to achieve escape velocity from its institutional predisposition towards anticipation strategies. (Under the FD&C Act, the FDA is fully empowered to define “safety” and “effectiveness” as narrowly or broadly as it deems necessary.) Either way, the recommendations outlined in these two reports are an important first step in helping to improve the medical marketplace by placing greater emphasis on the post-market surveillance process, aided by emerging technologies, which can help expedite the clinical trial approval process, open the door to competition among a wider array of biotechnology and pharmaceutical firms (including early-stage startups), and ultimately help usher in an era of precision medicine for all Americans.

3. Policy Recommendations

- a. **Replace the FDA's bimodal approval process with a tiered approval process based on each therapy's risks and benefits;**
- b. **Narrow premarket approval standards to safety and effectiveness, the definitions of which should also be narrowed;**
- c. **Lift the statutory ban barring the FDA from considering clinical trials of heritable (germline) genome editing; and**
- d. **Embrace a multistakeholder approach to developing standards and regulations governing the application of genetic technologies for “enhancements” to the human genome.**

B. Internet of Things

The IoT generally refers to an ecosystem of devices using software and embedded sensors, which can be connected to the Internet or a local network. This allows previously unconnected and mundane physical systems to begin generating large amounts of data, which can then be collected and analyzed to achieve a wide array of objectives, including increasing supply chain efficiencies, expanding analytical forecasting power, and reducing the need for human intervention. Although these innovations hold the potential to disrupt a wide variety of industries, a lack of regulatory clarity means this burgeoning industry could be stifled before reaching maturation.

Recent events, however, show positive signs of government policymakers heading in the right direction. In January 2017, the Department of Commerce released a green paper detailing its policy perspective on the IoT and made the notable, and historic, decision to reaffirm the principles of the *Framework* — the first time a government agency has done so since they were first released over 20 years ago. As aptly noted in the green paper:

Over the past few decades in the United States, the role of government largely has been to establish and support an environment that allows technology to grow and thrive. Encouraging private sector leadership in technology and standards development, and using a multistakeholder approach to policy making, have been integral elements of the government's approach to technology development and growth. Following a review of public comments, meetings with stakeholders, and the public workshop, it is clear that while specific policies may need to be developed for certain vertical segments of IoT, the challenges and opportunities presented by IoT require a reaffirmation rather than a reevaluation of this well-established U.S. Government policy approach to emerging technologies.⁵²

While this is a heartening sign of things to come, reaffirming these principles is but the first step in

a long and difficult journey of ensuring the security of the IoT ecosystem.

1. Encryption, Online Liability Protections, Cybersecurity Insurance, and the VEP

While network cybersecurity for the IoT (and the Internet more broadly) is certainly a realm where government can potentially lend a helping hand, there are limits to how much it can do and how effectively it can do it. Fears of potential cybersecurity breaches of these systems are well-founded, but prescriptive regulatory mandates won't ameliorate those concerns. Instead, policymakers need to treat the security of the IoT as a feature to be improved, not a mandatory obligation to be imposed. There is no silver bullet in cybersecurity, but with each breach, attack, and intrusion we can react by improving best practices rather than by passing knee-jerk regulations, which often come with unintended consequences. Beyond best practices, there are a host of other interrelated policies and technologies that can help better secure the burgeoning IoT network while helping to foster, rather than inhibit, innovation.

One of the easiest and most effective means of ensuring the availability of strong network defense tools is to safeguard and support the use and proliferation of encryption protocols. While such tools can be a strong defensive bulwark against hackers and other nefarious actors, encryption is also a key tool in protecting the economic vibrancy of the digital landscape more broadly. As *The Atlantic* noted in its coverage of the Niskanen Center's 2015 research on the economic benefits of encryption, "The 40-plus trillion online banking industry ... would have been 'significantly stunted' without strong cryptography ... and the online purchases that in 2013 totaled more than \$3.3 trillion depended on encryption for trust and security."⁵³ As software has continued eating the world in the years since the release of Niskanen's research, the realities of encryption's value to the digital economy have only grown in scope and magnitude. Thus, the paper's conclusion still holds today: "The Internet is the lifeblood of the modern digital economy; encryption protocols are

the white blood cells. The health of the Internet ecosystem depends on the proliferation of strong encryption.”⁵⁴

Another piece of low-hanging policy fruit is ensuring open channels of communication and information sharing between government and the private sector. Codifying the vulnerabilities equities process (VEP) in law, as proposed in legislation like the PATCH Act, would be one such approach that could help stimulate and accelerate public-private cooperation in the fight against emerging online threats. Sharing information about “zero-day” exploits, as the VEP is intended to do, can help better secure the digital landscape by promoting greater collaboration and trust among industry, researchers, the intelligence community, and government more broadly.⁵⁵ Public-private coordination represents the ideal way to combat online insecurity while allowing government and industry to leverage their comparative strengths for the benefit of all. It also serves as yet another example of soft law systems plugging governance gaps that would otherwise be filled by ineffective government mandates.

In contrast, one of the worst approaches policymakers could take in attempting to secure IoT networks would be to attempt to codify technical baseline standards as legislative mandates. When considering ideal cybersecurity solutions, it is important for policymakers to remember that although mandatory technical standards may seem like an easy fix to a complicated problem, they always come with trade-offs and unintended consequences. Strong standards may be a brief hindrance for today’s bad actors, but the same mandated standard could expose broad swaths of the network to more serious penetration attacks tomorrow. The more centralized and standardized the approach to cybersecurity, the wider the attack surface for would-be cyber-assailants.

Additionally, this approach is often plagued by poor definitions that can create a host of unintended consequences for improvements to cybersecurity standards, such as limiting research into security testing software, as occurred after the signing of the Wassenaar Agreement. In an at-

tempt to limit the export and spread of surveillance tools, the Wassenaar Agreement inadvertently criminalized the very research necessary to fight back against hackers and other bad actors.⁵⁶ Export control agreements also played a significant role in the First Crypto War in the 1990s, assigning munitions-level control requirements to encryption that ultimately ran into clear First Amendment legal challenges. Code, after all, is a form of speech.

Of course, not all international agreements have been bad for cybersecurity. The recently-passed Clarifying Lawful Overseas Use of Data (CLOUD) Act, for example, was an important step in updating the law to accommodate the unique extraterritoriality issues raised by transnational digital networks, and has the added benefit of stipulating the need for any foreign government entering into a cross-border data-sharing agreement with the United States may not use any order issued under the terms of the law “to infringe freedom of speech.”⁵⁷ Most notably, the recently-updated North American Free Trade Agreement (NAFTA), colloquially dubbed NAFTA 2.0, embraces a market liberalism approach to digital trade, ensuring that the Internet economy does not succumb to digital protectionism and data localization requirements, both of which would undermine strong cybersecurity practices by cutting local providers off from the expertise of more global providers of such services.⁵⁸ Although international agreements for unrestricted cross-border data flows are important for cybersecurity, an even more pressing need is for policymakers to support the continued inclusion of strong online intermediary liability protections for online service providers in such agreements.

One example of the cybersecurity benefits resulting from strong online intermediary liability protections is the growth of content delivery networks (CDNs). CDNs, such as Cloudflare and Akamai, are services that operate linked servers, enabling faster and more secure content delivery services to users. While discussions of online intermediary liability protections often focus on content platforms like Facebook and Twitter, CDNs play an increasingly important role for cy-

bersecurity. As the Niskanen Center noted in 2016 comments to the United States Trade Representative, these services facilitate a safer and more secure online experience for individual users:

There are many benefits of utilizing CDNs, not least of which are the significant cost savings on storage and bandwidth when compared to central server streaming networks. Whatever benefits some actors participating in notorious markets may reap from CDN services, the mere possibility of a technological tool being used for ill is not justification enough for it to be held liable for the actions of users. As online content becomes more interactive and bandwidth-intensive, a more distributed network will increasingly become the most architecturally beneficial approach to optimizing user experience and services.⁵⁹

Cybersecurity threats continue to be a serious problem for the digital ecosystem. Luckily, the existence of balanced online intermediary liability protections has helped create a flourishing ecosystem of new innovative cybersecurity services, such as CDNs and other platforms that, without such protections, may never have found a pathway to commercialization. One such example is Cloudflare Orbit, which, as noted in 2017 Niskanen Center comments to the National Telecommunications and Information Administration (NTIA):

is a product that provides IoT vendors and users with the capability to push updates and patches at the network level, defraying the need to rely on updates to the physical devices themselves. Devices using Cloudflare Orbit will first pass through Cloudflare before connecting to the Internet, filtering out malicious activity. Essentially, Orbit operates as a “fog” between the physical device and the networked “Cloud.” The fog can filter out suspicious code and activity that would otherwise trickle down to user devices, all without necessitating additional security protocols that are hardcoded into the actual IoT device. Emerging technologies like this can act as additional layers of security for IoT devices, and

could help remedy concerns associated with relying on updates on the user-level.⁶⁰

Finally, policymakers can help encourage better cybersecurity by embracing policies that promote the adoption of cybersecurity insurance, and the growth of that still-budding market. The Niskanen Center argued for this approach in a separate set of 2017 comments to NTIA by noting that:

Cybersecurity insurance reduces financial risks associated with data breaches, network damage, distributed denial of service (DDoS) attacks, and other related cyber incidents. A growing number of companies are making use of cybersecurity insurance. According to PricewaterhouseCoopers, cybersecurity insurance annual gross written premiums are set to grow from around \$2.5 billion today to \$7.5 billion by 2020. That growth, however, is likely to face a number of challenges for actuaries attempting to underwrite cybersecurity policies.

Lack of actuarial data presents major challenges for insurance underwriters attempting to accurately quantify cyber risk, resulting in restricted coverage. The Department of Homeland Security National Protection and Programs Directorate is exploring the feasibility of creating a cyber incident data repository “that creates a trusted environment for enterprise risk owners to anonymously share sensitive cyber incident data.” A cyber incident data repository could enable insurers to more accurately assess the risks from online security threats and provide better coverage for the potential damages related to cyber incidents. Encouraging firms to purchase cybersecurity insurance can help lead to the adoption of better security practices and cybersecurity response plans tailored to the risks incurred by a variety of business models.⁶¹

As of this writing, there are also a number of ongoing soft law multistakeholder proceedings — within industry, trade associations, the academy, and government agencies — aimed at addressing

many of these thorny issues through the development of best practices for policymakers interested in the more technical details of this debate.⁶²

Just like the Internet, security-related concerns surrounding the IoT touch a broad spectrum of policies and issue spaces. Of course, many of the cybersecurity issues discussed above could quickly become irrelevant with the advent of quantum computing.

2. Quantum Computing: The Future of Cybersecurity

In classical computing, information is encoded as bits, which are expressed either a one or a zero. All integers, except zero and one, can be factored into prime numbers. Finding the prime factors for large numbers is computationally intensive (making it difficult to guess the right combination to “unlock,” for example, an encrypted email), which makes it appealing for cybersecurity. The most common cryptographic algorithms, such as Rivest-Shamir-Adleman (RSA), depend on the fact that it is trivially easy to find the product of two prime numbers — but practically impossible to do the reverse — to securely encrypt information. In these asymmetric cryptosystems a public key is generated, which is the product of two large primes, and anyone can use it to encrypt messages. An individual’s private key, which is used to decrypt those messages, consists of the two prime numbers used to generate the public key. It would take countless lifetimes to find the prime factors of a sufficiently large number using a classical computer.

Quantum computing, on the other hand, is much more powerful because it encodes information in quantum bits, or *qubits*, which can be a one, a zero, or any proportion of both states at once (known as a superposition). Quantum computing also leverages a unique property of quantum particles known as “entanglement,” which is when pairs or groups of particles interact in ways that make it impossible to describe their quantum states — position, momentum, spin, and polarization — independently of one another. In other words, the particles form a system, and if you

know the quantum state of one part of the system, you know the quantum state of the whole system.⁶³

This new type of computing has many potential practical applications, from searching databases more efficiently to simulating quantum systems in chemistry and nanotechnology. But its impact on the future of cybersecurity has generated the most discussion. The problem is that quantum computers in the future might be sufficiently powerful to cost-effectively compute prime factorizations for large numbers, thus breaking some of the most widely used methods of encryption, including RSA. While quantum computing will be able to bypass current best practices for encryption, the technology will also enable new methods of protecting data by generating quantum keys for encryption. When using quantum-secure communication, any attempted eavesdropping destroys the communication, and the eavesdropping is detected.

According to a recent report from the Congressional Research Service, “China designated [quantum information sciences (QIS)] research as one of four ‘megaprojects’ in its 15-year science and technology development plan for 2006-2020. Additionally, it designated quantum communications and computing one of six major goals for this period. China’s annual funding for QIS R&D has been estimated at \$244 million.”⁶⁴ To ensure the United States will be the world leader in quantum computing, policymakers need to increase funding for research and education in the field, remove roadblocks to knowledge and technology transfers out of universities, and coordinate research across institutional boundaries in the federal government.⁶⁵ On this latter point, the Office of Science and Technology Policy (OSTP) recently took an encouraging step by chartering a QIS subcommittee, which is tasked with coordinating a national agenda on this technology.⁶⁶

3. Policy Recommendations

- a. **Support the Department of Commerce’s reaffirmation of the Framework, and reiterate its value at fora,**

domestic and international, in which representatives of the U.S. government are involved;

- b. Continue to champion the ongoing NTIA multistakeholder working group strategy for addressing emerging IoT issues;
- c. Prioritize industry-led standards, not mandatory one-size-fits-all regulations, that can serve as the benchmark for cybersecurity best practices;
- d. Codify the VEP in law to build certainty, trust, and consistent communications channels between government and the private sector;
- e. Ensure ongoing innovation in the cybersecurity ecosystem by safeguarding the use and proliferation of encryption while refraining from placing undue regulatory restrictions on firms operating in this space;
- f. Promote the use of cybersecurity insurance in federal procurement standards for cybersecurity technologies and systems; and

C. Autonomous Vehicles

Every year, approximately 35,000 Americans die from automobile crashes as a result of human-operated error. The leading cause of death for children ages 4 and 11-13 are car accidents. Autonomous vehicles (AV) could help save tens of thousands of lives annually, and contribute to dramatic savings in commuting time, fuel costs, and carbon emissions. This future is only possible, however, if we get the rules of the road right.

Safety is likely to be of paramount concern with the deployment of AVs, especially in the early stages when these vehicles will be sharing the roadways with human-operated automobiles. The Department of Transportation currently possesses the necessary authorities to regulate the safety of vehicles on American roads, and this is as it

should be. Additionally, those authorities should preempt any states attempting to craft their own standards for AVs operating on the roadways; otherwise, the resulting patchwork of laws and regulations could inhibit more widespread deployment and adoption. This is one policy arena where harmonization of rules and authorities is imperative. However, issues relating to the testing and validation of software algorithms, remedying liability concerns, and setting standards and best practices are better handled by third party validators, the common law and tort system, and industry groups, respectively. Regulators should focus on areas in which they have experience and expertise, not expanding their authority into unknown waters where they may end up doing more harm than good.

1. Addressing Privacy and Cybersecurity Concerns

As detailed in Part II(B)(1), the problems with attempting to craft mandates for cybersecurity standards are legion. These problems are further punctuated by calls for technology-specific cybersecurity standards, which often fail to discern the nuance between various real-world applications of networked systems. AVs, in particular, have been the subject of numerous calls for government-mandated cybersecurity standards under the auspices of ensuring public safety. And indeed, there are unique safety concerns in the automotive space that are not replicated in, for example, online platforms providing social media, banking, and communications services. However, it is these unique considerations — that is, the complicated interactions between vehicle-to-vehicle (V2V) communications, transportation service networks, and ongoing developments in systems cybersecurity best practices — that should give policymakers pause before reactively endorsing a set of one-size-fits-all cybersecurity standards.

Although there have been considerable policy developments since the National Highway Traffic Safety Administration (NHTSA) released its *Preliminary Statement of Policy Concerning Automated Vehicles* back in 2013, the same general framework

of best practices for vehicular cybersecurity are still applicable today. Those include:

1. Ensuring a safe, simple, and timely process for any transition from a self-driving mode to human-operated mode;
2. Promoting systems that can detect, record, and apprise operators if or when automated systems malfunction;
3. Ensuring autonomous technologies do not interfere with or otherwise disable federally-mandated safety features; and
4. Ensuring event logs are maintained so that, in the event of a crash or loss of control, there is recorded information that can detail what occurred.⁶⁷

Those principles are still relevant in guiding ongoing developments in autonomous vehicle cybersecurity best practices. And as noted in a 2015 *Wake Forest Journal of Law & Policy* law journal article, it should not go unmentioned that:

*Manufacturers have powerful reputational incentives at stake here, which will encourage them to continuously improve the security of their systems. Companies like Chrysler and Ford are already looking into improving their telematics systems to better compartmentalize the ability of hackers to gain access to a car's controller-area-network bus. Engineers are also working to solve security vulnerabilities by utilizing two-way data-verification schemes (the same systems at work when purchasing items online with a credit card), routing software installs and updates through remote servers to check and double-check for malware, adopting routine security protocols like encrypting files with digital signatures, and other experimental treatments.*⁶⁸

These ongoing developments and reliance on existing security measures, such as encryption, are yet further reasons for policymakers to ensure fidelity with many of the recommendations already provided in Part III(B)(i).

Similarly, many of the concerns regarding privacy are not unique to autonomous vehicles. As a general matter, and as will be discussed at length throughout Part IV, there are good reasons to be wary of federal baseline privacy regulations that may inadvertently undermine potential future business models that could use ad-financing as a means of driving down the costs of shared autonomous vehicle services, or cost-sharing models that help subsidize the use of this technology for low-income riders. For the same reasons, policymakers should be skeptical of proposals that would treat autonomous vehicle privacy in a non-technologically neutral manner.

With regards to technological neutrality, there is an issue policymakers should keep their eyes on moving forward: mandates for V2V standards in general, and Dedicated Short Range Communications (DSRC) in particular. The DSRC standard has become something of a “zombie V2V” issue, at times thought to be moving full-steam ahead, at other times thought to be all but dead. While it has been billed as an ideal V2V standard that will promote safety and efficiency for communications between AVs, the technology is not only outdated and insecure but technically cumbersome as well. As detailed in a *Wired* op-ed from 2017, imposing a DSRC mandate on American vehicles would be the equivalent of the government mandating “that Henry Ford equip every one of his Model Ts with telegraph machines.”⁶⁹ It continues:

*A 19th century communications technology mandated for use in a 20th century innovation would have been a crushing blow to innovation and competition in the emerging automobile industry. That's precisely what is happening with the DSRC mandate, and the same potential for future innovation is at risk with its implementation.*⁷⁰

From a cost-perspective, DSRC carries a hefty price tag – upwards of \$5 billion per year — and doesn't accommodate interoperability with other standards. That means DSRC is only useful if installed on, and utilized by, every other vehicle on the road — a result that will take many decades

before achieving full implementation among the entire U.S. vehicle fleet.⁷¹

Most concerning, however, is that DSRC poses significant and serious risks to both cybersecurity and privacy. Cybersecurity expert Alex Kreilein noted many of these concerns in a comprehensive security review of the DSRC standard, highlighting “the critical gaps in the existing security architecture of vehicles and the inherent security and privacy concerns associated with the use of” DSRC, including “poor configuration management, a general lack of supply chain and code security,” and vulnerabilities that open the door to vehicular cyber intrusion through “at least six specific categories of attacks: deception attacks, denial of service attacks, cryptographic exploitation, malware exploitation, jamming and spoofing, and [vehicle-to-everything] exploitation.”⁷²

In short, DSRC is representative of what policymakers can expect if attempting to mandate a single technical standard to govern network security: minimal privacy protections, glaring security flaws, high adoption costs, and liability incentives for firms to become technologically locked-in, eschewing the need to continually innovate new cutting edge cybersecurity solutions.

2. Updating Existing Laws and Regulations

Updating the entire corpus of federal motor vehicle safety standards (FMVSS), while a daunting challenge, is a necessary step towards ensuring full legal and regulatory compliance with existing federal safety standards. Innovators and investors need market certainty to assuage fears that ongoing developments in commercializing AV technology will not run afoul of existing laws.

To that end, although many individual FMVSSs will require extremely technical updates to accommodate the unique, non-human-centric nature of AVs, a good place to start would be 49 C.F.R. § 571.3. This is the provision of the FMVSS code that defines a vehicle “operator” or “driver” to mean “the occupant of a motor vehicle seated immediately behind the steering control system.”⁷³

Obviously, this definition constrains the applicability of existing safety standards to autonomous vehicles, while having the added and unfortunate effect of prohibiting new vehicle designs that would be more appropriately suited to a vehicle where the driver is not, in fact, “seated immediately behind the steering control system,” but lives within the software of the vehicle itself.⁷⁴

3. Policy Recommendations

- a. **Establish a clear, reasonable, and well-defined baseline for what constitutes a “safe” AV;**
- b. **Prioritize updates to the FMVSSs using a “Clarity Through Simplicity” approach;**
- c. **Refrain from including any privacy provisions that are technology-specific, which could inadvertently stymie innovation in AV technology and new service offerings;**
- d. **Abstain from granting regulators pre-market approval authorities for autonomous-specific vehicles;**
- e. **Promote expanded exemptions for the testing and deployment of AVs; and**
- f. **Embrace third-party validators and industry-led standards over prescriptive mandates.**

D. Commercial Drones

Unmanned aerial systems (UAS), more commonly referred to as drones, have captured the imagination of a new breed of innovators. From would-be commercial entrepreneurs to more recreational users, these new systems have catapulted concerns over safety and privacy into the public limelight. Unfortunately, these concerns tend to overshadow the many potential benefits of UAS operations.

Airspace is one of the most underutilized resources in the United States, but it doesn’t need to be that way. When it comes to optimizing our

use of the national airspace for commercial innovations, the sky truly is the limit. Before we can transform the skies above into the next great platform for innovation, however, policymakers need to rethink how to better address air traffic control and management.

1. Air Traffic Management and Airspace Auctions

Establishing an effective air traffic management system that can handle low-altitude autonomous UASs is a necessary first step to ensuring wider integration of commercial drone operations. The development of a modern air traffic management system with this capability requires devolving governance resources and authorities to lower-tier municipal control. To accommodate this change, the Federal Aviation Administration (FAA) should spin off air traffic control responsibilities into a congressionally chartered, nonprofit user-coop (as is the case in most other countries).⁷⁵

Such a system will play a vital role in helping to accelerate innovations in the use of the national airspace, such as vertical take-off-and-landing (VTOL) drone transportation initiatives like those underway at Uber Elevate.⁷⁶ Although there are many technical challenges ahead for these skyward dreams, the biggest hurdles remain regulatory. The same is true with respect to Facebook's Aquila drone project, which could be a reasonably priced method of delivering Internet access to underserved areas of the United States.

The current open access approach has worked reasonably well for managing high-altitude airspace in the United States because there is a small number of players vying for access and their aircraft generally travel on fixed routes. As low-altitude airspace becomes more crowded with drones and small aircraft that travel on variable routes, this governance regime will need to be replaced with a new model for air traffic control.

As Brent Skorup and Melody Calkins point out, it would be wise for the FAA to learn the lessons from the Federal Communication Commission's (FCC) auctions of radio spectrum over the last

few decades.⁷⁷ For resources like airspace, allocating usage via property rights and the price mechanism has been shown time and again to be superior to a top-down regulatory approach. The auctions — and liquid secondary markets — ensure that a resource is put to its highest value use, which will be crucial as companies find new applications for UASs in the market.

2. Following Through on the Integration Pilot Program

In October 2017, the Trump administration issued an executive order directing the Department of Transportation to conduct a pilot program for UASs.⁷⁸ As the memorandum notes, it is indeed in the best interest of the United States “to promote the safe operation of [UAS]” and that “compared to manned aircraft, UAS provide novel, low cost capabilities for both public and private applications.”⁷⁹

In May of this year, the Department announced ten winning drone pilot projects, including from Google, Apple, and Microsoft.⁸⁰ Test flights for these projects are already underway and the fruits of innovation from this testing period will likely quickly follow.⁸¹ These first projects are a promising start, and more are likely to be authorized in the coming months, with some of the most exciting developments likely to come from Amazon, which plans on applying its proprietary drone technology to its distribution network in the near future.⁸²

As Deputy U.S. Chief Technology Officer Michael Kratsios noted in an op-ed for CNN, some safety risks need to be mitigated, but in general, the industry needs minimalist and consistent regulatory guidance to maximize investment and create new jobs:

The economic impact of the integration of drones into United States airspace is estimated to reach tens of billions of dollars. Providing the necessary legal authority to counter potential threats from drones will ensure that the United States benefits from this rapidly developing sector of the economy.

Realizing these potential benefits compels us to authorize safe, innovative drone operations. Responding to their potential risks should compel Congress to authorize our government to protect the American public. In both cases, the Trump administration will work to ensure the United States is the global leader in bringing the benefits of this emerging technology to our people.⁸³

This program, along with a reassessment of existing FAA regulations for commercial drone operations, will help innovators and consumers benefit from one of the country's most underutilized natural resources — low-altitude airspace — and open the door to a new age of airspace innovation for drone delivery services, new urban VTOL transportation networks, and a more market-friendly approach to airspace governance. To see that vision materialize, however, policymakers first need to free the skies for innovation to flourish.

3. Policy Recommendation

- a. **Modernize the air traffic control and management system, with particular attention paid to integrating systems that can accommodate autonomous operations;**
- b. **Apportion low-altitude airspace to be managed by public-private partnerships at a local level;**
- c. **Reform the more onerous provisions of existing UAS rules (such as restrictions on beyond-visual-line-of-sight and flights over non-operators, failures to permit the use of automated sense-and-avoid technologies, and daylight-only operations); and**
- d. **Support efforts that eliminate barriers to the emergence of new drone-related services, including easing the certification requirements on experimental aircraft designed for higher-altitude commercial operations.**

E. Supersonic Flight

The speed of sound is not constant. It is slowest in gases, faster in liquids, and faster still in solids. It rises and falls with the temperature. When a traveling object reaches the local speed of sound — Mach 1 — a sonic boom announces its arrival. The crack of a bullwhip or supersonic bullet is this same phenomenon in miniature. When an aircraft breaks the sound barrier, it is known as supersonic flight and, in a sufficiently moist environment, a vapor cone also appears around the fuselage to crown the achievement.

Since the retirement of the Concorde in 2003, there have been no sonic booms or vapor cones in commercial aviation. The Concorde was uneconomical to operate and public concerns about sonic booms damaging buildings, shattering windows, and creating noise pollution near airports effectively killed it off. Yet the death knell for a commercially viable supersonic passenger jet happened years before. In 1973, the FAA banned supersonic flight over the United States, confining it to international travel. This effectively cut off the natural entry point for a private sector (as opposed to government-backed) supersonic industry, which would have at least initially catered to overland business jet passengers. In turn, private research and development into “quiet” supersonic technologies that would serve to abate the sonic boom were chilled indefinitely.⁸⁴ Rescinding the ban and replacing it with a reasonable sonic boom noise standard would therefore significantly increase the addressable market for overland civil supersonic flights in the future.

The FAA Reauthorization Act of 2018 was a major step forward for supersonic aviation. In it, Congress directed the FAA to take leadership toward fostering the return of civil supersonic; to issue a new rulemaking for civil supersonic aircraft certification; and to review the practicality of amending the overland ban every two years, beginning in 2020. While short of repealing the ban outright, the current FAA reauthorization has made supersonic something that can no longer be neglected by regulators nor the broader technology policy community.




Figure 5: Showing the airspeed record peaking in 1976, with stagnating (and even regressing) speeds in the decades that followed.
 Source: Eli Dourado and Michael Kotrous, “Airplane Speeds Have Stagnated for 40 Years.”

1. Lifting the National Airspeed Limit

Despite the 1973 ban on supersonic overland (14 C.F.R. § 91.817), recent advances in materials science, aerospace engineering, and computer simulation techniques have enabled engineers to design conceptual models of supersonic planes with much quieter sonic booms. Technologically, we are light years ahead of the Concorde era. In fact, without necessarily optimizing for quieter booms, the boom generated by the next generation of transoceanic supersonic airliners is anticipated to be as much as 30 times quieter than Concorde’s boom, largely thanks to the lower mass of carbon fibre airframes. With a few changes to public policy, the coming age of supersonic flight will easily extend beyond transoceanic routes to routes overland. A two-hour flight from New York to Los

Angeles is only the most obvious application. Supersonic flights over the continental United States could open up many potential new, time-saving routes, like a significantly faster flight from London, England to Sydney, Australia travelling east-to-west.

For years the FAA’s official position on civil supersonic overland has been that more data is needed.⁸⁵ Critics argue that this has gotten the order of operations backwards, requiring the manufacturing and testing of a low boom supersonic airliner without the requisite regulatory certainty required for investors to know whether their target design will be quiet enough.⁸⁶ This is a departure from the FAA’s historical approach to noise regulation, which has focused on tracking noise standards against existing technological ca-

pabilities, rather than using tighter stringency rules to force new technology. Nominally, the FAA has stated that it is awaiting the results from the National Aeronautics and Space Administration's (NASA) low-boom flight demonstrator. Yet this is itself problematic, as the NASA-Lockheed "X-59 QueSST," once completed in 2022, will represent the bleeding-edge of low-boom technology, rather than a low-boom design that simultaneously optimizes for cost and passenger capacity. Setting a sonic boom noise standard based on the X-59 QueSST would be akin to setting subsonic noise standards based on the de minimis noise created by the quietest possible passenger jet.

A better approach would be to replace the supersonic speed limit with a sonic boom noise limit based on the wealth of existing data the FAA and others already have regarding the noise tolerances of populated areas. An initial noise limit would be based on what is technologically practical and economically reasonable, with a protocol for tightening noise stringency overtime as noise abatement technologies improve, in much the same way take-off and landing (T&L) noise standards have tightened overtime. Moreover, just as existing noise regulations are based on a measure of cumulative exposure rather than a single decibel level, an intelligent sonic boom standard would encourage supersonic flight paths that avoid densely populated areas, or which occurred during the daytime. Indeed, the cost-benefit literature on noise pollution is already incredible rich, and could be readily extrapolated to sonic booms, using factors like perceived decibels, boom duration, and cumulative population exposure as parameters.⁸⁷

Lifting the speed limit on civil aviation would be a major step toward reversing the stagnation in commercial airspeeds. Indeed, as shown in Figure 5, in the forty years since the Concorde's maiden flight, average commercial airspeeds have actually fallen. With major airlines focused on price competition, marginally slower flights make economic sense for fuel economy. Overland supersonic flights, on the other hand, would reintroduce a speed dimension to airline competition, redirecting research and development into genuinely

breakthrough technologies, rather than incremental improvements to existing models.

2. Noise Standards and Abatement

For years business jets represented the natural entry point for a viable commercial supersonic sector due to their rich clientele. Today, however, technologies have enabled several start-ups to leapfrog into the passenger jet category, bringing far superior affordability. The emerging supersonic leader, Boom Technology, is perhaps the best example of this.⁸⁸ Independent of overland capabilities, Boom is on track to begin manufacturing a 55 seat supersonic airliner for transoceanic routes in the mid 2020s. Although fares are ultimately determined by airlines, Boom projects initial fares could be as low as existing business-class tickets on a cost per passenger mile basis. Yet whether supersonic travel will be truly affordable for the average air traveller will depend a great deal on policy choices.

T&L noise standards are one of the most key policy parameters affecting the affordability of supersonic passenger travel. To understand why requires a brief detour into how aircraft noise standards work. Subsonic passenger jets are certified according to Aircraft Type and Airworthiness Certifications contained in 14 C.F.R. Part 36, which, among other things, regulates cumulative aircraft noise via the FAA's advisory circular, Noise Levels for U.S. Certificated and Foreign Aircraft. In the United States, aircraft noise standards are referred to terms of "stages" (Stage 1, Stage 2, Stage 3, and so on) with each stage representing a greater degree of noise stringency. Current jet and large turboprop aircraft are certified at Stage 4, however new type certificates have to be certified at the even more stringent Stage 5 after December 31, 2017 and December 31, 2020, depending on the weight of the aircraft.

Airport noise is a contentious issue in local politics. Residential associations neighboring airports are a well-organized interest group — so powerful that they have their own caucus in the U.S. House of Representatives known as the "Quiet Skies Caucus." Nonetheless, a study of airport noise

complaints revealed the vast majority of complaints are made by a small number of highly motivated individuals. In 2015, for instance, the Ronald Reagan Washington National Airport received 8,760 complaints about airport noise, 6,852 — or 78 percent — of which originated from one residence in North West D.C., a pattern that is consistent for airports across the country.⁸⁹

Moving from the local to the global, international aircraft noise standards are set by the International Civil Aviation Organization (ICAO), a specialized agency of the United Nations, based on the assistance of member states and industry groups on ICAO's Committee on Aviation Environmental Protection (CAEP). Given the dominance of the United States in aerospace technology, and the size of the U.S. economy more generally, the FAA has a significant (though by no means dispositive) influence over ICAO and CAEP proceedings. Conventionally, noise stringency standards are tightened with a significant lead time and in collaboration with industry to ensure standards reflect a dissemination of existing technologies and best practices, rather than serving to force new noise abatement technologies into existence.

By and large, subsonic passenger jets meet more stringent noise standards by increasing the bypass ratio of their turbofan engines. A larger bypass ratio means more air flows through the fan than the core turbine, resulting in less perceived noise. Indeed, the appearance of modern turbo fans is misleading, as the core turbine that provides the plane with thrust is substantially smaller than the fan in front of it. Given that a larger bypass ratio actually *improves* fuel efficiency at subsonic speeds, subsonic jet manufacturers have readily adopted more stringent noise standards overtime.

The story for supersonic aircraft is quite different. The next generation of supersonic passenger jets will realistically fly between two and three times the speed of sound. At these speeds, the drag coefficient of the aircraft becomes a make-or-break parameter from the point of view of fuel efficiency, and therefore affordability. To minimize drag, supersonic aircraft typically have elongated shapes and cruise at twice the altitude of subsonic planes

to take advantage of thinner atmosphere. Large bypass ratio turbofans are simply impractical at these speeds, akin to dragging one's arms off the sides of a canoe. As such, affordable supersonic passenger jets are likely infeasible under Stage 5 noise standards, at least given current technology. For instance, according to Boom Technology, complying with Stage 3 noise standards, rather than the current Stage 4 standard, would improve their fuel efficiency by over 20 percent.⁹⁰ This is worthwhile trade-off from the perspective of both affordability and the environment. While Stage 4 and 5 are undoubtedly quieter for communities that neighbor airports, they provide no noise advantage at supersonic's cruising altitude. Indeed, Stage 3 T&L noise standards prevailed for passenger jets throughout the 2000s without cause for alarm, suggesting a T&L noise standard of similar stringency (or a modified Stage 4 or 5 standard) would be a worthy compromise for enabling supersonic as an emerging aviation technology. (For comparison, Concorde was certified under much-less stringent Stage 2 standards). Additionally, CAEP's purview includes both aircraft noise and emissions, meaning they are obligated to consider fuel efficiency trade-offs when formulating international noise standards.

Fortunately, when the FAA was challenged on their interpretation of 14 C.F.R. Part 36, it was discovered that civil supersonic aircraft fall outside existing type certification regulations. As such, a supersonic aircraft certified in 2018 or thereafter will not necessarily be forced to adopt Stage 5 noise standards. Instead, the FAA Reauthorization of 2018 directs the FAA to issue a Notice of Proposed Rule Making for a supersonic-specific certification no later than December 2020. This in turn presents an opportunity for the FAA to develop noise standards that are crafted to supersonic's unique technological circumstances, with input from the public policy community.

3. Policy Recommendations

- a. **Make amending the ban on supersonic transport over the continental United States a key priority at the FAA;**

- b. Adopt operating standards for boom noise that are informed by noise levels tolerated elsewhere in society, and which properly account for the benefits of a robust supersonic market;
- c. Adopt a supersonic aircraft certification with T&L noise standards that (1) are practical to meet at supersonic speeds, (2) ensure affordable access to supersonic travel, and (3) appropriately consider fuel efficiency trade-offs; and
- d. Ensure the United States continues to exercise leadership at the international stage.

F. Commercial Space

Despite the remoteness of space, its impact on daily life is far more profound than most realize. In 2015 alone, the global space economy was valued at over \$320 billion. Of that total value, over 75 percent was attributable to the commercial infrastructure and associated systems that allow for everything from GPS-location services to satellite television.⁹¹ In the future, other still emerging — and yet-to-be imagined — space services could contribute even more substantially to both the global and American economies. Both Facebook and SpaceX, for example, are currently working on developing Internet-delivery systems using small satellite constellations, and companies like Bigelow Aerospace are creating new modular space habitation systems that could revolutionize the potential for “time-sharing” exo-atmospheric real estate, whether for space tourism or more practical research and scientific experimentation. That future is only possible, however, if innovators and policymakers can help drive down the cost of orbital payload launches, which currently average between \$27,000 and \$43,000 per pound of cargo.⁹²

Despite the cost challenges, a number of new space launch companies have emerged in recent years, capitalizing on the clear market demand for cheaper launch services. As more industry competitors enter this market, launch costs will con-

tinue declining, benefitting not only the commercial sector, but government as well. This will also bring a range of new innovations and technologies to Americans. To actualize this vision of a vibrant space economy, policymakers should promote policies that enable the development of a robust, commercialized, and accessible space sector. To do this, it will be necessary to embrace policies that support a strong domestic space economy while reforming government oversight of space activities to provide the regulatory certainty that is currently sorely lacking, especially for beyond-Earth orbit missions.

1. Licensing Commercial Activity in the Final Frontier

In a 2017 paper — “The Future of Space Commercialization” — Joshua Hampson detailed the many regulatory barriers facing commercialization of, and access to, space.⁹³ As he noted:

Current American regulations focus on systems leaving or entering Earth’s atmosphere and the capabilities of satellites in orbit. The current regulatory structure spans several government agencies, leading to a somewhat disjointed structure. While workable during an era of single-use outer space operations (placing satellites or space stations in orbit), it has become increasingly inadequate as more private actors enter the space economy and seek new opportunities.⁹⁴

While the Final Frontier is currently a vast expanse of unregulated commercial potential, getting there is anything but. First, in order for a satellite to escape Earth’s orbit, a company must undergo a payload review process for launches and reentries, overseen and approved by the FAA’s Office of Commercial Space Transportation (FAA AST). That checklist of requirements includes reviewing and approving a company’s:

1. Flight termination system design testing;
2. Operating techniques;
3. Launch and reentry site selections; and

4. Appropriate compliance with public health and safety regulation, international law and U.S. treaty obligations, and domestic national security requirements.

Then, assuming FAA AST signs off on the checklist of approval items, the purposes for which a satellite is to be used can run it into further regulatory hurdles. For example, most satellites make use of spectrum bandwidths, which require a license from the FCC. If the system is intended to serve in a remote-sensing capacity, it must acquire additional licensure through the National Oceanic and Atmospheric Administration (NOAA), which is tasked with ensuring all remote sensing undertaken by domestic firms adheres to both national security rules and is in compliance with all international treaty obligations to which the United States is a signatory.

These reviews also regularly implicate considerations of dual-use export controls, which are further complicated by controlled technologies being split between the International Traffic in Arms Regulations (ITAR) and the Export Administration Regulations (EAR) lists. Compliance is complicated even further as ITAR and EAR fall under the purview of separate departments — ITAR lives under the authority of the State Department, while EAR is housed at the Department of Commerce. When considering the number of individual components that make up each satellite, and the cacophonously disjointed and often ambiguous regulatory guidelines required for each commercial launch, it's a wonder any satellites are able to find their way to orbit at all. Streamlining this approval process should be a top priority for policymakers.

Regulatory uncertainty and complexity notwithstanding, there are also national security concerns that present hurdles to opening the door to greater commercialization of near- and beyond-Earth orbit. As Hampson pointed out in a 2017 *RealClearFuture* op-ed:

The military and intelligence communities rely on satellites for coordination, communication, and reconnaissance. For a long time, the

United States did not have to worry very much about satellites once they got into orbit. Today, however, there is a growing number of countries with active and burgeoning involvement in space operations. Potential rival powers are looking at America's satellite constellations as a vulnerability to exploit. There are serious concerns that, in a security crisis, a sophisticated country could destroy or degrade America's space capabilities.⁹⁵

Of course, limiting commercial access to, and use of, near-Earth orbit can also cut against national security and intelligence interests. Satellites are difficult to replace, given their unique operational domain (i.e., space) and national security institutions already rely heavily on commercial space infrastructure and launch services. Until the emergence of SpaceX, the only launch systems employed for transporting sensitive military and intelligence service satellites to orbit were provided by United Launch Alliance (ULA) — an industry cartel jointly-owned by Boeing and Lockheed-Martin. While its safety record is solid and offers a reliable service, ULA price tags come in at between \$164-350 million *per launch* for their Delta 4 system. ULA's sole alternative to the Delta 4 — the Atlas V — is cheaper, but it comes with additional national security concerns because it relies on Russian rockets.

Although ULA's new Vulcan Centaur will use rockets produced domestically by Blue Origin, there is no indication how long it will be before that system is ready for deployment.⁹⁶ In the meantime, the demand for both national security and commercial launch services only continues to grow as SpaceX continually improves its reusable rocket technology, driving down costs and promoting the entry of new firms looking to potentially make use of space for everything from tourism to scientific research.

2. Promoting Certainty Through Default Approvals

Even as commercial launch services help stoke interest in new uses of real estate in near-Earth orbit, more audacious private sector undertakings

have been proposed. In 2017, for example, Moon Express became the first private company cleared for a mission to the moon. Unfortunately, the ad hoc nature of the approval process means there has not been an opening in the regulatory aperture for other beyond-Earth orbit missions to follow suit. Although the Moon Express approval certainly sets a precedent, the path to future missions will necessitate a clear streamlined pathway for regulatory approval. As it currently stands, the approval process for beyond-Earth orbit missions is even more complicated and unwieldy than near-Earth orbit launch approvals:

While commercial activities beyond orbit have been established as legal in the United States, the current process relies on opaque, discretionary decision-making within multiple agencies. It's difficult to trace such decisions back to individual officials, who have to consider national security and foreign policy decisions. Without a formal process, firms have no way of knowing whether future missions will be permitted. With so many agency stakeholders involved and an international obligation to authorize and supervise all private space missions, the U.S. government might lapse into de facto non-approval. It's easy to understand, then, why commercial space companies are concerned about regulatory uncertainty. Industry concerns over the opacity and unpredictability of the mission approval process are likely to spur the government to consider new oversight mechanisms for the private exploration and use of outer space.⁹⁷

To assuage pervasive uncertainty in this emerging market, policymakers should place the burden of proof for denying future beyond-Earth orbit missions on government. Put differently, private space missions should be granted a default approval status, with the burden on government to show — clearly and unambiguously, based on a set of articulable standards — the proposed mission would have serious negative consequences for public safety, put national security at risk, or violate international treaties or agreements. Agency procedures for overruling that default should be

transparent and involve a clear, streamlined, and expeditious process for organizations appealing such decisions.

To help facilitate a policy that prioritizes permission-by-default, FAA AST should be elevated outside FAA, to a sub-cabinet level within the Department of Transportation. This would give the agency a larger mouthpiece in government, improve its budgetary position, and separate its mission — licensing commercial space operations and launches — from the FAA's mission to effectuate and police the safety of the domestic national airspace.

“The future success of American private space companies,” Hampson argues, “will rest on clear, coherent, and consistent oversight. Missions should be default-approved, with the burden of proof on the government to demonstrate why a particular initiative should not move forward.”⁹⁸ Unfortunately, given the number of interagency review processes and unclear standards currently required for beyond-Earth missions, it is unclear whether, and how, new missions can be approved.⁹⁹ To capitalize on the near-infinite expanse of space, that will have to change.

As Elon Musk noted in his September 2016 Mars colonization plans, “Life needs to be more than just solving problems every day. You need to wake up and be excited about the future.”¹⁰⁰ The prospect of humanity's future among the stars is one of the most exciting prospects of the future — one that is being driven largely by private sector investment and innovations in everything from advanced materials science to reusable rocket boosters. The key for policymakers is to ensure the regulatory environment is friendly and welcoming. Otherwise, we run the risk of deterring entrepreneurs from reaching for the stars.

3. Policy Recommendations

- a. **Promote competitive markets for commercial launch and space services by using consistent contracts across government purchases, reducing non-competitive government support for**

incumbents, and removing archaic regulations;

- b. Support and pass the Space Frontier Act of 2018 (S.3277) to streamline the certification process for commercial space activity;
- c. Give FAA AST the requisite authority and voice within government to promote commercial space by elevating it to a sub-cabinet position within the Department of Transportation; and
- d. Charter a non-governmental multi-stakeholder group to coordinate international space situational awareness.

G. Climate Engineering

While there is little scientific doubt that human activities are behind most of the recently-observed warming of the Earth — with more to come over this century — there is a great deal of uncertainty about the environmental and economic consequences of warming. As such, society should respond to climate change as a risk management problem and seek ways to reduce greenhouse gas (GHG) emissions, minimize societal vulnerability, and otherwise limit the potential costs of a warming planet.

1. The Need for Further Research

Geoengineering technologies are one prospective means of addressing climate change. As a class of technologies, they could potentially allow people to intervene in the Earth system at a large enough scale to deliberately alter the climate and decouple the total amount of warming from total emissions of greenhouse gases due to burning fossil fuels. Technological interventions to reduce the human influence on climate fall into two categories: Carbon dioxide (CO₂) removal and solar radiation management (SRM).

CO₂ removal, or negative emissions, describes technologies that would artificially remove CO₂ from the atmosphere and thereby limit the warming and chemical effects of excess CO₂. These are

interesting technologies for research that are already aligned with much of what is done at the Department of Energy (DOE) and other agencies. SRM refers to interventions that would decrease the amount of solar radiation that reaches the surface of the Earth by increasing the planet's reflectivity. While there are technical nuances and regional variances, the amount of cooling we could expect to see is roughly proportional to the decrease in radiation. These technologies could therefore be tuned to partially or fully offset the warming effects of increased CO₂ with a large enough intervention, while very small experiments would have no globally detectable signal. For SRM, the research agenda is more novel, and governance requirements more pressing, than CO₂ removal.

Not all of the considerations that will govern decisions to use or refrain from using SRM technologies are scientific. However, numerous scientific and engineering gaps prevent an informed understanding of the costs and benefits of potential unintended consequences. Reducing uncertainties and better characterizing those risks presents the scientific enterprise with the opportunity to add value for future policymakers. The potential scale of climate risks and the costs associated with transitioning to a low-carbon economy mean that the potential value of SRM could register in the trillions of dollars. Many organizations, like the Bipartisan Policy Center,¹⁰¹ the National Research Council,¹⁰² and the Governmental Accountability Office (GAO),¹⁰³ express some need to better understand the risks of solar geoengineering through research.

2. OSTP Oversight and Multistakeholder Governance

To maximize the potential gains from future research in this budding field, policymakers should grant jurisdictional authority over climate engineering policy efforts to OSTP and its subsidiary, the United States Global Change Research Program (USGCRP). As an interagency program focused on collaborative engagement with agency stakeholders, USGCRP is well positioned to address the near-term policy concerns of climate

engineering. It has the mission, expertise, and infrastructure to handle an influx of technical reporting data from ongoing experimentation, possesses the necessary working relationships with other agencies to serve as a platform for constructive interagency dialogue on overlapping issues, and lacks formal rulemaking authority that could otherwise lead to overly precautionary and scientifically stifling regulations.

Thus, USGCRP and OSTP can effectively further research and development of climate engineering science and technology through two primary activities: (1) convening a formal multistakeholder process to establish consensus-based policy recommendations for future research and experiments, and (2) make federal grants available for researchers. Any multistakeholder process should aim towards achieving a general consensus on the following outputs:

1. Research design standards or best practices for researchers;
2. A federal strategy and roadmap on funding priorities;
3. An OSTP/USGCRP green paper that establishes an official government policy position on climate engineering and SRM research (not deployment), describes the current state of climate engineering science, and issues a statement of principles to guide the federal government's involvement in research funding; and
4. Non-binding OSTP/USGCRP guidance on reporting protocols for researchers submitting updates on ongoing experimentation and research.

Longer-term issues that may warrant consideration include upper-bound restrictions on larger-scale experiments and more formally recognizing OSTP's role as the agency-of-record in international negotiations on climate engineering activities.

3. Policy Recommendations

- a. Affirm OSTP's role and authority in overseeing climate engineering research efforts;
- b. Convene an ongoing multistakeholder process under the purview of USGCRP that will prioritize the creation of voluntary, consensus-based policy guidelines and recommendations for future research priorities; and
- c. Apportion dedicated federal funding for climate engineering research and development.



As soon as it works, no one calls it AI anymore.

— John McCarthy

PART IV: ARTIFICIAL INTELLIGENCE

AI and machine learning (ML) are perhaps the least well-understood emerging technologies confronting policymakers. Although headlines will often speak of the newest breakthroughs with extreme exaggerations, the reality is that AI is already all around us. From Google's search algorithm to advanced translation systems, this technology is already deeply embedded in our everyday lives. This is because, as the famous AI researcher John McCarthy once said, "as soon as it works, no one calls it AI anymore."¹⁰⁴

Rapid advancements in AI/ML hold the potential to significantly disrupt, and possibly entirely upend, a wide range of industries. It will contribute to ongoing developments in robotics and automation, greatly increasing manufacturing efficiencies; it will help analyze massive troves of digital data, yielding insights and breakthroughs in everything from medicine and materials science to driverless cars and video optimization. Of course, much of the talk surrounding AI is focused on far-flung hyperboles that have no basis in reality — from the less-flashy forecasts of mass robotic-induced unemployment to visions of *Deus Ex*- and *Terminator*-style apocalyptic futures.

Though some will argue to the contrary, AI does not pose an existential threat to civilization. Portents of dystopian doom and destruction stemming from AI make for good clickbait and better blockbuster films, but they should not serve as the basis for public policy-making. When discussing serious AI policy, policymakers should avoid focusing on worst-case scenarios and instead embrace language that speaks to the promising benefits this technology heralds. That messaging, framed around the need for regulatory restraint, is imperative to help provide the market certainty necessary to continue fueling ongoing investments in AI/ML research and development.

The transformative potential of the many emerging technologies discussed in Part III are all, to some degree, reliant on AI to actualize their full

potential. Whether it's advanced IoT sensor suites providing real-time data tracking for postmarket surveillance and patient oversight of new gene therapy treatments or the software powering an autonomous vehicle, each of these technologies makes use of AI/ML and automated decision-making systems. AI is a "nexus technology," powering advancements in other commercial applications. It is because of the uniquely significant role AI/ML plays in the broader technology ecosystem that policymakers should take special care to approach any new broad-based regulatory proposals with skepticism.

A. AI in Digital Advertising

Discussions of online advertising often focus on amorphous and non-quantifiable concerns surrounding consumer privacy and the latter's incompatibility with a digital economy premised on the monetization of individuals' data. But contrary to many assertions, these values — economic growth and privacy — are not necessarily in conflict. As one economic journal article put it:

*Extracting economic value from data and protecting privacy do not need to be antithetical goals. The economic literature we have examined clearly suggests that the extent to which personal information should be protected or shared to maximize individual or societal welfare is not a one-size-fits-all problem: the optimal balancing of privacy and disclosure is very much context dependent, and it changes from scenario to scenario. ... Thus, it stands to reason that, case by case, diverse combinations of regulatory interventions, technological solutions, and economic incentives, could ensure the balancing of protection and sharing that increases individual and societal welfare.*¹⁰⁵

Despite the simplistic narrative often highlighted in news media reports of technology's impact on privacy, the general consensus among researchers is that individual assessments of privacy are complicated, contextual, occasionally contradictory relative to stated preferences, and above all highly subjective and atomistic. Nonetheless, privacy

debates continue to be heated flashpoints characterized by highly-charged emotional rhetoric that is often unmoored from evidence and empirical analysis — rhetoric that is unfortunately all-too commonly inflamed by advocacy groups dedicated to a dogmatic and myopic focus on privacy as the one human right to rule all others.

Those who advocate for this perspective do an immense disservice to discussions of technological governance by pushing the belief that privacy “occupies so profound a position in the constellation of human values that it is innately inalienable and should never be subject to commodification, regardless of the potential social or economic benefits. ... Although these voices tend to be louder and more emotionally forceful than others, the shrillness with which a conviction is proclaimed is not dispositive of some manifest truth; merely disputing the absolute sacrosanctity of privacy does not imply a cavalier indifference to its value.”¹⁰⁶

The unfortunate casualty of this focus on hypothetical privacy harms is the lack of discussion about the very real, quantifiable, and observable benefits of digital advertising — increasingly buttressed by the use of automated programmatic ad purchases — on consumer welfare, and the contributions of AI/ML to those welfare gains. A 2010 analysis from McKinsey & Company, for example, noted the significant proportion of surplus value from ad-funded Internet services that went to consumers (approximately 85 percent, with the remaining 15 percent accruing to producers).¹⁰⁷ The report further explained that over “80 percent of current Internet users generate significantly more value from using the Web than what they would be willing to pay to eliminate [advertising] disturbances.”¹⁰⁸ In short, the evidence of individual consumer valuations of privacy in relation to the trade-off of zero-priced online services weighs heavily in favor of the latter.

1. The Economics of Online Advertising

Before discussing the economic benefits of AI/ML in online advertising, understanding *how* online

advertising markets actually work will help dispel some of the common mischaracterizations often portrayed by Privacy Fundamentalist advocates and the media. In a 2013 Cato Institute policy brief, Larry Downes described the process, and exorcised the myths and misnomers, thus:

It’s important to dispel right from the start some persistent myths about how advertising actually works. The marketing of transaction data is far more complex than advocates for more government regulation of privacy would have us believe. It’s not “your” information that’s being sold. First, the information is collected and stored by the service with your permission. If the data ever was “yours,” you freely traded your interests in it. For most of the Internet’s billion users, the exchange is a good one, generating far more value for users than the costs of supplying and parting with the information.

Data being processed for advertising isn’t “yours” in a second sense: It doesn’t identify you as the source. Search engines such as Google don’t sell information about what individuals searched for, and advertisers don’t then turn around and advertise to those individuals based on what they have learned about them. Google doesn’t even know “who” is doing the searching, only that the searches originated from the same computer. Google keeps track of the activities of that computer (which could be used by one person, a family, the patrons of a library, or a bot), and it does so only by storing cookies on the computer that maintains the connection.

But the cookie doesn’t contain identifiable information about the user — the name, address, and so on. And once you delete a cookie, the data collection has to start all over again. (Your searches will get less streamlined if you do, as Google’s software will make worse guesses about what you’re actually looking for.)

More to the point, ads you see on Google search results or other applications only ap-

*pear to be optimized as personal messages. In most cases, the services and their sponsors don't make use of the individual cookie data, or at least not on its own. Say you searched for "carpet cleaners in Berkeley, CA." Google doesn't sell that fact to carpet cleaners in Oakland, who then pass along an advertisement to the computer of whoever typed that search. The actual science of advertising is both more and less sophisticated than that.*¹⁰⁹

While 2013 is practically an eon in the lifecycle of developments in the digital economy, the reality is that the fundamentals of digital advertising have not actually changed all that much — companies have just gotten much better at optimizing and tailoring those ads, largely as a result of advancements in AI/ML.

The digital advertising market is a leader in applying AI/ML to real-world problems. According to one estimate, approximately 80 percent of digital display ad spending in the United States in 2018 was programmatic, meaning it was auctioned off in real-time across a network of publishers and platforms using AI technology.¹¹⁰ Digital ads are used to fund the free or heavily subsidized online services that form the backbone of the digital economy. In fact, a 2010 report by McKinsey & Company found that the value users obtain from using Internet services was six times greater than the amount they were willing to pay to avoid advertising disturbances and privacy risks.¹¹¹ AI/ML also helps increase satisfaction in the advertising ecosystem — among users, advertisers, and publishers — by improving the relevance of targeted ads. As a 2013 paper in the *Journal of Machine Learning Research* aptly recognized, “the publisher must balance short term interests, namely the immediate revenue brought by the ads displayed on each web page, and long term interests, namely the future revenues resulting from the continued satisfaction of both users and advertisers.”¹¹²

AI/ML is an increasingly valuable tool for helping publishers balance those interests. Among its many benefits in the online advertising economy, AI/ML optimizes ad placements, personally-tailors content to individual consumer expecta-

tions, and allows for greater ease of scalability — providing a wider audience of both users and firms with the beneficial capabilities, and outcomes, of targeted advertising. These are not trivial benefits. For consumers, the gains of AI/ML uses in programmatic ad buys come in the form of targeted and relevant advertisements, improving the online user experience without diminishing the quality of the zero-priced services many Internet denizens have come to expect as a given of online life. While many will decry the monetization of user data as somehow fueling an already overly-commercialized culture that feeds our worst purchasing impulses, the reality is that online ads deliver very real welfare-enhancing outcomes. Downes aptly summarized the reality that many Privacy Fundamentalists often discount:

*Advertisements are offers. Those that are perceived as “ads” are offers that are at least slightly off. But an ad for the right product or service, offered at the right time to the right person at the right price, isn't an ad at all. It's a deal.*¹¹³

These benefits have already proliferated across the digital ecosystem, with some estimates placing the number of online auction market bids using narrow AI for programmatic ad buys at 80 percent.¹¹⁴ That number will almost certainly continue to rise in coming years. For consumers in particular, these benefits come in the form of better-targeted and more relevant online advertisements, improving the online user experience while maintaining the competitive race to ensure quality for even zero-priced online services remains higher than ever. For advertisers, the benefits of deploying AI/ML for these programmatic ad-buys are perhaps even more profound, when accounting for the amount of revenue losses incurred as a result of digital ad fraud.¹¹⁵

2. Self-Regulatory Governance

Where existing laws fail to address consumer harms, private industry has been remarkably successful in its commitment and adherence to an ecosystem of self-reinforcing and self-regulating

governance norms and frameworks. In the realm of consumer data and privacy, there is an unheralded story of how digital advertising consortia have effectively stepped in to help govern the thorny questions that regulators have had difficulty addressing. Third-party certification organizations like the National Advertising Initiative (NAI) and Digital Advertising Alliance (DAA) are one component of this self-regulatory landscape, holding participating companies accountable to a complex set of frameworks, standards, and best practices, with regular compliance audits.¹¹⁶

Alternatively, less-decentralized regulatory approaches, such as omnibus federal privacy legislation, are very likely to have a chilling effect on research and development of AI/ML technologies. In particular, broad-based privacy rules — such as those included in both GDPR and the recently passed California Consumer Privacy Act (CCPA) — have tended to prioritize greater user control at the cost of innovation by codifying stringent limitations on both the collection and use of user data. As the Niskanen Center noted in August 2018 comments to the Federal Trade Commission (FTC), this imbalanced trade-off poses significant problems for advancements in AI/ML:

*Greater restrictions on data collection and limitations on their use naturally obstruct more sizable gains in AI/ML innovation, through increased development costs and disincentivizing greater levels of investment. In the same way that enhancing individuals' control over their information restricts the free flow of information, so too does restricting the collection of that information inhibit potentially beneficial, and entirely unforeseeable, future applications of that data for AI/ML purposes. The end result is a net decrease in consumer welfare, as consumers lose out on products and services that may otherwise have benefited them.*¹¹⁷

In addition to this concern, there are a great many reasons to be skeptical of shifting away from self-regulatory governance in the use of AI/ML in the online advertising industry. The foreseeable prob-

lems that could result from embracing omnibus privacy regulation over the current framework include:

1. Creating an online free speech crisis, as increasing individuals' control over greater aspects of their information flows (such as through a Right to be Forgotten) increases the level of a posteriori censorship in society;
2. Misallocating resources towards compliance and away from innovation, as every dollar a firm must spend on new legal mandates is a dollar not spent on R&D;
3. Diminishing future innovation due to more limited access to, or increased costs associated with acquiring, data;
4. Providing structural advantages for incumbent firms that curtail competition by demanding greater investments in privacy counsels, data protection officers, and regulatory compliance staff;
5. Disincentivizing research into data protection tools and systems — such as native browser controls for cookies and advertising ID opt-out settings — as firms can default to newly-heightened legal requirements as a liability shield; and
6. Shifting priorities and funding away from self-regulatory governance approaches, which provide greater responsiveness and flexibility to changing technological circumstances and consumer expectations.

Thus, rather than attempting to craft one-size-fits-all rules for the many different types of data that firms collect, policymakers should instead continue to embrace the sectoral-based status quo, which prioritizes privacy rules based on the sensitivity of particular types of data according to industry use (e.g., health care vs. finance vs. social media use cases of data).¹¹⁸ This approach has served American innovators and consumers well over the past quarter-century, as described in a

2001 working paper from the AEI-Brookings Joint Center for Regulatory Studies:

Complaints about the patchwork of regulations governing information privacy in the U.S. notwithstanding, there are valid reasons for supporting a selective approach to information privacy protection. Passing laws one at a time, for specific areas, allows for a more careful evaluation of issues. In principle, such laws are less prone to — although certainly not immune from — unintended consequences. This approach also allows (at least in theory) for lawmakers to consider the costs and benefits that each proposed act is likely to entail. ... With the federal government, enacting legislation in an incremental fashion is easier than eliminating bad policy once it is on the books.¹¹⁹

While there are certainly drawbacks to a less uniform approach to privacy regulation, the superiority of a less centralized approach to governing privacy is striking when looking at the disparities between the American and European digital economies. Of the top 25 technology companies worldwide, 15 are U.S. firms, with a combined market capitalization value of over \$5 trillion. European firms, in contrast, occupy only three of those spots, with a comparatively meager \$285 billion in total market capitalization. Putting that in perspective, Intel alone is as valuable as the entire European technology industry.¹²⁰

3. Policy Recommendations

- a. **When considering “purpose specification” rules for data, data retention mandates, or default opt-in requirements, the FTC should give considerable weight to evidence detailing the economic costs associated with such proposals;**
- b. **Continue embracing the ecosystem of self-regulatory mechanisms that can effectively balance innovation and consumer welfare with the need to address specific harms; and**

- c. **Refrain from adopting rules that specifically target, and that may inhibit, the use and deployment of AI/ML systems in online advertising**

B. AI in Medical Devices

The rise of big data and ongoing developments in AI/ML offers the potential for breakthroughs in medicine that will deliver enormous improvements in patient outcomes. This future of precision medicine, however, requires a regulatory environment conducive to new innovations in order to deliver on its promise. For example, as the Niskanen Center discussed in 2018 comments submitted to the FDA, software-based medical diagnostic devices are relatively low-risk to patient safety and should bear a proportionately small regulatory burden before being cleared for clinical use. AI-enabled devices in particular present a major opportunity to improve the accuracy, enhance the personalization, and lower the cost of medical diagnoses. Clear guidance and minimal regulatory hurdles will increase investment by the biotech industry in this space and accelerate the pace at which patients and doctors reap the benefits of this technology.

1. Soft Law at (and Beyond) the FDA

As noted in Part III(A)(2), the FDA is in desperate need of a regulatory update. It has remained significantly behind the curve in adopting a 21st century approach to regulatory governance, most notably in its failure to make broader use of the many possibilities afforded by the Internet and IoT devices. One solution to this problem, of course, is for the agency to continue its reliance on soft law guidance while vigorously committing to a program that incorporates advanced technologies into the regulatory approval process. Updating regulatory approvals in this manner could have a particularly significant impact on expediting the approval of new AI-based medical devices. Such changes are long overdue at the FDA.

“In the 20th century,” write Robert Graboyes and Sara Rogers, “medical device regulation was designed for large, heavy, relatively immobile, easily

visible, expensive technologies with easy-to-track money trails. That works well for a multi-ton, multimillion-dollar MRI machine. It's far less effective for an inexpensive artificial pancreas system produced by the user or passed hand to hand through informal channels."¹²¹ They conclude:

*The FDA would be well-advised to expand its use of soft-law regulation. Taking full advantage of the dizzying stream of innovations coming our way requires a faster, defter, more predictable regulatory regime than those procedures which dominated technology in the 20th century. This likely means more voluntary standards, co-regulation (tasks shared by multiple agencies, along with self-regulation), and guiding principles rather than hard-and-fast mandates.*¹²²

Ironically, despite the FDA's long history of reliance on soft criteria and soft law, the most notable soft law proceedings for health care technologies are actually currently underway at the Department of Commerce.

NTIA recently convened an ongoing series of multistakeholder working group meetings focused on defining Software Bill of Materials (SBoM) and proposing methodologies for their promulgation.¹²³ An SBoM is intended to promote transparency in software design to solve numerous problems associated with cybersecurity risk, including the need to develop systems for sharing relevant data on individual software components. Although the process is intended to apply to the broader software ecosystem, there is a working group specifically dedicated to creating a "Health Care Proof of Concept" for SBoMs. The goal of this working group is to establish "a collaborative effort between health care delivery organizations and medical device manufacturers ... to demonstrate successful use of [SBoMs] and relate to the overall cross-sector effort to establish standardized formats and processes."¹²⁴

While soft law has been criticized for its inability to promote long-term clarity and regulatory certainty, in the realm of medical devices utilizing AI/ML it may actually — and perhaps counter

intuitively — help promote a more predictable regulatory regime through the use of voluntary standards and pilot programs. One such pilot program currently under development at the FDA is a perfect example of how the agency might be taking important, albeit small, steps in the right direction for promoting the use of AI/ML in medical devices.

2. Voluntary Standards and the Software Precertification Pilot Program

Despite repeated statements committing to the creation of a clear pathway for regulatory approval of AI-based medical devices, the FDA's corpus of official policy guidance documents has remained notably silent on this issue. Indeed, a cursory query for "AI" or "artificial intelligence" will yield almost nothing on the agency's website. That may be changing, however.

The agency recently released its second iteration of "Developing Software Precertification Program: A Working Model vo.2" (hereafter, *A Working Model vo.2*). This program is intended as an experimental approach to regulating the emerging Software as a Medical Device (SaMD) industry based on vetting software development organizations producing these devices, rather than a specific premarket review of each SaMD product up for regulatory consideration. Unfortunately, *A Working Model vo.2* — despite being touted as the means by which the FDA will begin assessing regulatory approval processes for AI-based medical devices — provides no specific AI/ML recommendations, voluntary standards, best practices, or almost any mention at all. To add insult to injury, the current draft of the pilot program, while technically an improvement on the previous version, still lacks basic clarity and in many cases actually creates more confusion.

At a high-level, the pilot program describes a set of "Excellence Principles" and Key Performance Indicators (KPI) that an organization can use to determine their eligibility for participation in the program. On the one hand, the program lays out a very broad scope of eligibility, potentially opening

the door for smaller health care and AI startups to compete with large biotechnology, pharmaceutical, and even other more established technology firms.¹²⁵ In practice, however, the vague program guidelines offer little guidance as to what, precisely, an organization must demonstrate to meet a given KPI, and the FDA does not deign to offer examples of what might constitute a qualifying KPI.

In an attempt to glean greater insight into what could be an important program in developing the FDA's approach to future AI medical device regulation, in July 2018 the Niskanen Center submitted comments to the agency. These focused on 12 recommendations for improvements to the current iteration of the pilot program.¹²⁶ Those recommendations included:

1. Ensure grant of precertification status to a diverse portfolio of organizations for the pilot program and beyond;
2. Create a scoring sheet and minimum thresholds for identifying whether each of the excellence principles are met, based on demonstrated elements within each organizational domain;
3. Include example KPIs for each element used to appraise organizational excellence;
4. Include an additional element under the Deployment and Maintenance domain of the excellence appraisal chart to evaluate an organization's ability to collect and analyze post-launch, real-world performance data;
5. Include a new organizational domain in the excellence appraisal chart called "Artificial Intelligence Development Best Practices" with appropriate elements and KPIs;
6. Explicitly allow for outsourcing of core organizational activities in the excellence appraisal process;

7. Develop an excellence appraisal process template and example including estimated timelines for each application stage;
8. Abolish two levels of precertification and adopt a single precertification standard;
9. Simplify the SaMD risk categorization framework by adopting a single formula yielding a single number that is categorized into three risk levels: low, moderate, and high;
10. Require precertified organizations to proceed to Streamlined Review only for high-risk SaMDs.
11. Clarify the Streamlined Premarket Review process requirements; and
12. Include a domain that measure dynamic physician reliance on SaMD tools under User Experience Analytics for Real-World Performance Analytics assessments.

These suggestions are primarily aimed at clarifying the lack of certainty involved in both the selection process and criteria for approving participating organizations, as well as focusing the FDA's attention on the need for embracing "Least-Burdensome" regulatory principles.

Recommendations 1-7 are aimed at clarifying the appraisal process for eligibility and particular Excellence Principles that, as currently constructed, leave far too much breadth in interpretation. This uncertainty also unfairly favors large entrenched incumbents at the expense of potentially more innovative, flexible, and financially lean startups. In particular, the comments emphasized the paramount need for the FDA to institute recommendation 5 (see Appendix B for further details and a specific case study application of sample KPIs) to actually move the agency towards a competent and transparent treatment of SaMDs that incorporates AI/ML technologies capable of improving their capabilities through learning systems. In conjunction with the tiered categories of approval reform proposed by Gulfo and Briggeman in Part

III(A)(2), these reporting standards for the use of AI in medical diagnostic devices could help significantly minimize the existing regulatory hurdles — manifested primarily as pervasive uncertainty regarding what is and is not permitted — for innovators and entrepreneurs.

Recommendation 8 is intended to further level the playing field between large and small firms by streamlining the precertification pathway to include only one precertification level. Currently, the pilot program bifurcates these precertification pathways, applying two sets of organizational requirements based on whether an organization had a demonstrated track record of SaMD products that met the amorphous standards of the program's Excellence Principles. This despite the FDA noting that the agency's "current thinking reflects the belief that an organization of any size without a medical device or SaMD currently on the market should have the opportunity to deliver products for medical purposes as a precertified organization."¹²⁷ While it is certainly true the "opportunity" to become precertified is present, the segmentation of these separate approval levels — differentiated by a standard that appears intended to discriminate based on, among other things, a company's size — clearly suggests the program will inevitably favor large incumbents over smaller startups. Further, the agency offers no explanation as to the need for, or benefit of, distinct precertification levels.

Finally, recommendations 9-12 offer a further critique of the lack of clarity associated with the proposed Streamlined Premarket Review process, which would require "high-risk SaMDs" to go through an additional approval process, while "low-risk SaMDs" would require no additional FDA scrutiny for approval. As with the rest of the pilot program, the devil is in the definitions. As the Niskanen Center comments noted:

The current risk categorization framework is unnecessarily complex and based on arbitrary and theoretical conjectures of how risky an SaMD product is expected to be. From a practical standpoint, many AI-based SaMDs would not easily fit into such a framework.

For example, if an AI SaMD is developed to process and analyze skin biopsy images, would the device be considered to diagnose a skin condition such as melanoma or to drive clinical management? Such an SaMD would most certainly be trained on data that was labeled as either positive or negative for melanoma. When providing an output as to whether melanoma is likely present or not, it could be seen as driving clinical management by allowing a physician to examine the biopsy further and to send the patient for extra testing, or it could be seen as diagnosing the condition outright. In either case, such an approach is highly subjective, and fails to provide clarity for applicants seeking an understanding of what distinguishes diagnosing from driving in the context of AI-based SaMDs.¹²⁸

To help clarify what constitutes "low" versus "high" risk in such products, Niskanen's comments proposed a simplified risk categorization framework that would be more generalizable to all SaMD product types covered by the pilot program, using a simple formula to calculate a quantifiable risk number:

$$\text{SaMD risk} = (\text{Severity of most serious health condition} + \text{Proportion in user population} + \text{Error rate}) / 3$$

Those three variables, in turn, would be defined as follows:

1. **Severity of most serious health condition:** Value between 0 and 1, with 0 representing no health consequence and 1 representing a potentially lethal condition.
2. **Proportion in user population:** Value between 0 and 1, with 0 representing the non-existence of the most severe health condition in the user population and 1 representing all users with the condition in question.
3. **Error rate:** Value between 0 and 1, with 0 representing perfect accuracy of a SaMD product in detection and classification of the existing health condition and 1 repre-

senting a complete mismatch in detection and classification.

Such an approach is a far more powerful means of promoting transparency and certainty in the approval process. In contrast, the more abstract language currently used to define “low” versus “high” risk categorizations offers no possibility of anything *but* ad hoc, non-standardized, and opaque interpretations.

While *A Working Model vo.2* is a good first step towards better addressing the issues of SaMDs, the program would benefit from a great deal of added simplicity. As currently construed, the program poses an especially acute barrier for smaller companies with less experience navigating the difficult bureaucratic requirements imposed by the FDA, and which have far more limited resources to devote to unnecessarily complicated precertification requirements for a mere pilot program.

3. Policy Recommendations

- a. **Embrace a technologically-neutral approach to regulating software in medical devices;**
- b. **Emphasize flexibility and adaptability in development standards;**
- c. **Prioritize new guidance describing and clarifying FDA’s thinking on AI medical and diagnostic devices; and**
- d. **Expand the Software Precertification Pilot Program to include AI software developers.**

C. Algorithmic Accountability

AI has a “black box” problem. Like the human brain, the decision-making processes behind ML systems and neural networks are immensely complex and opaque. The algorithms that underlie these technologies can learn and improve without explicit instruction from a human, which makes explaining the exact reasons for a particular decision functionally impossible. But some techno-

phobes in the policy community are demanding a level of transparency and explainability from AI decisions that we do not require of human decisions. Why is there a double standard of rigor when AI has already demonstrated the ability to make better decisions than humans across a wide variety of domains?

Part of the reason is likely an emergent narrative that attributes the prevalence of bias in the use of automated decision-making tools. Of course, these discussions are fraught with misnomers, and media stories surrounding bias in algorithms often confuse the nature of *statistical* bias with a more *colloquial* understanding of the term, which is often conflated with “fairness” (although it is true that at times these two distinct understandings of “bias” do indeed converge).¹²⁹ The problem with this approach, as Chris Stucchio and Lisa Mahapatra have pointed out, is that it leads to a popular discourse that increasingly blurs reality with the ideal. “Bias,” Stucchio and Mahapatra note, “is defined as the difference between an AI’s ‘typical’ output and reality. Bias has a magnitude and a direction, and is a systematic tendency to get the same kind of wrong answers.”¹³⁰ They continue:

Very importantly, “wrong” is defined as being relative to reality. One problem with reality is that it often fails to live up to our desires and expectations. There is often a very large gap between the two. When a journalist discusses bias, they typically do not mean it in the same manner that statisticians do. As described in the examples above, a journalist typically uses the term “bias” when an algorithm’s output fails to live up to the journalist’s ideal reality.¹³¹

As a result, the standard narrative surrounding algorithmic bias has become immensely misleading, as many people increasingly come to presume these systems are inclined towards “incorrect” decisions. It is certainly true, of course, that the training data used in AI — itself the aggregate result of many billions of individual human decisions — will be the product of decisions predicated on biases. There are, however, a corpus of best

practices for machine learning practitioners and developers, and obtaining representative data is immensely critical in building effective models — practices that developers strive to improve on a daily basis. If we abstain from making use of real-world data to improve these AI systems until all the biases of humankind are corrected, the many benefits heralded by the use of AI will almost certainly never materialize.

There are means of improving problem solving in ML, including prioritizing diversification of the AI workforce itself through investments in promoting general STEM education among underserved communities. This is not, however, a silver bullet to righting the many prejudices intrinsic to segments of AI training data; and indeed, the only true path to rectifying these historical and structurally systemic social wrongs is broader reforms to the education, social welfare, and criminal justice systems, among many others. In the meantime, there is much that can be done to promote the many effective, ethical, and socially beneficial uses of AI. A good starting point would be promoting rules that prioritize technological neutrality in regulation and focusing on verifiable consumer harms.

In a very real way, the outcomes we unearth as a result of AI decision-making are actually uncovering the oftentimes hidden biases in society. Far from systematizing and entrenching bias in society, AI/ML can be used to help shed light on these realities. The algorithms, however, cannot make the policy decisions for ameliorating those problems on our behalf: those solutions require a very human response from policymakers.

1. Principles for Regulatory Oversight

As detailed at length in Parts III(B)(2) and III(C)(1), the general principle that policymakers should always prioritize when considering new proposed regulations for emerging technologies is technological neutrality. This principle is similarly applicable to AI, which is simply a technological means of automating tasks and supplementing human decision-making. Thus, focusing on the *outcomes* of those decisions is a far more ideal

starting point for assessing the need for regulatory intervention or penalization. The ultimate goal of any legislation or regulation, therefore, should be aimed at minimizing actual, observable harms — not attempting to preemptively address hypothetical concerns. Legislation like the AI in Government Act, proposed by Senators Cory Gardner (R-CO) and Brian Schatz (D-HI), which would establish an “Emerging Technology Policy Lab” under the oversight of the General Services Administration, can serve as an effective platform for convening future multistakeholder proceedings dedicated to examining many of the issues related to AI/ML — effectively embracing a soft law approach to AI governance.

Some advocates, however, have argued the best way of achieving these ends is through “algorithmic transparency” — a term lacking in precise meaning, but which generally holds the need to peer through the black box of neural networks in an attempt to understand *why* particular automated decisions might have been made. Although well-intentioned, such an approach is all but impossible given current technologies; and even if it were possible, rules that subject the inner workings of an AI system and its code to pre-review and/or pre-approval would almost certainly be incapable of communicating any information of actionable value to consumers — to say nothing of the clear First Amendment implications.

Instead of requiring “algorithmic transparency,” policymakers should insist on “algorithmic accountability.” In practice, this would mean that the operator of an AI decision-making algorithm — not software engineers, developers, or researchers — would be responsible for the consequences of that decision and expected to provide a non-technical, high-level explanation for a particular decision. As a recent Center for Data Innovation report from Joshua New and Daniel Castro argues:

Rather than establish a master regulatory framework for all algorithms, policymakers should do what they have always done with regard to technology regulation: enact regulation only where it is required, targeting specif-

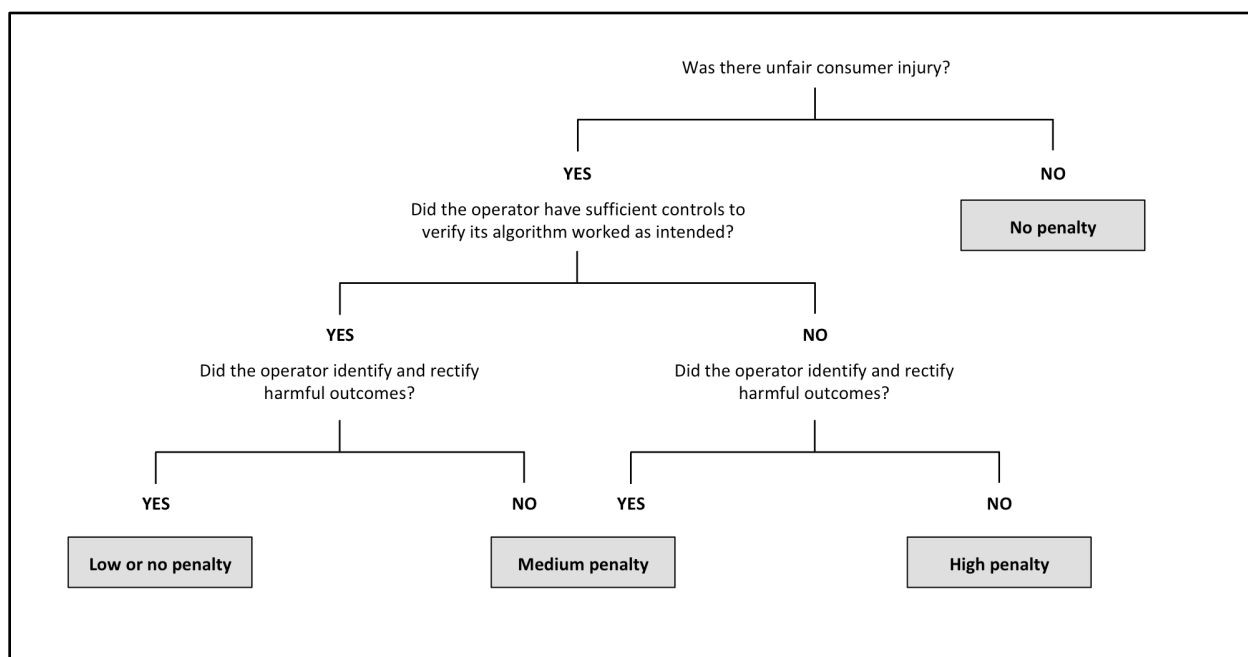


Figure 6: “The Regulator’s Neural Network.” Describing the optimal decision-making process when considering whether a harm has resulted from the use of AI. Source: New and Castro, “How Policymakers Can Foster Algorithmic Accountability.”

ic harms in particular application areas through dedicated regulatory bodies that are already charged with oversight of that particular sector. To accomplish this, regulators should pursue algorithmic accountability — the principle that an algorithmic system should employ a variety of controls to ensure the operator (i.e., the party responsible for deploying the algorithm) can verify it acts in accordance with its intentions, as well as identify and rectify harmful outcomes.¹³²

How would such an oversight process work in practice? As shown in Figure 6, New and Castro detail a prospective decision-making flowchart — “The Regulator’s Neural Network”¹³³ — for guiding regulators through the process of determining the existence of, and penalties for, consumer harms resulting from the use of an AI system. This flowchart is not intended to serve as a template for new legislation, or as a means to justify expanding FTC authority to create new categories of rules. (Indeed, it would be far preferable to simply allow the common law to address the emergence of such harms.) Rather, it is meant to assist regulators and policymakers in assessing whether an actual harm has materialized in the

context of an automated decision. But how does one make a determination as to the existence of an “unfair consumer injury” resulting from an AI decision-making process? Answering that question will be the focus of the next section.

2. Informational Injuries and AI-Specific Harms

In 2017 comments to the FTC, Castro and Alan McQuinn, writing on behalf of the Information Technology and Innovation Foundation (ITIF), outlined a framework that defined different types of personally identifiable information (PII), each of which would implicate different types of consumer harms. Those harms, in turn, would necessitate different regulatory responses.¹³⁴ Thus, in answering the initial question posed by the “The Regulator’s Neural Network” — that is, “was there unfair consumer injury” resulting from an operator’s deployment of an algorithm? — Table 1 below offers a helpful taxonomy in separating out the specific types of PII and the corresponding types of consumer harms that might result from that information’s exposure.

The ITIF comments also identify different “levels of PII collection and use,” each of which can also

result in different levels of risk for information injuries (see Table 2), For the purposes of algorithmic decision-making, levels 2 and 3 are the only situations that would conceivably be implicated by the use of such systems. While informational injuries resulting from the mere collection of PII is possible, AI systems that are used for re-

al-world decision-making purposes require that PII be both collected *and* used, whether humans are involved or not. Thus, for the purposes of AI, levels 2 and 3 are the only levels of PII collection and use likely to implicate any specific consumer harms necessitating regulatory responses.

Type of Information	Definition	Informational Injury	Definition
Observable Information	“Personal information that can be perceived first-hand by other individuals.”	Autonomy Violations	“Result in harm for consumers when information they consider sensitive and would prefer to keep private becomes public through involuntary means and tend to be harms that are ‘reputational or interpersonal’ in nature.”
Computed Information	“Information inferred or derived from observable or observed information” that “is produced when observable or observed information is manipulated through computation to produce new information that describes an individual in some way.”	Discrimination	“Occurs when personal information is used to deny a person access to something, such as employment, housing, loans, or basic goods and services.”
Observed Information	“Information collected about an individual based on a third party’s observation or provided by the individual, but does not allow someone else to replicate the observation.”	Autonomy Violations or Discrimination	(See above)
Associated Information	“Information that a third party associates with an individual” that “does not provide any descriptive information about an individual.”	Economic Harm	“Results when a consumer suffers a financial loss or damage as a result of the misuse of PII,” such as in the case of “identity theft, fraud, or larceny.”

Table 1: Examining the potential for informational injuries, corresponding to the type of information collected. *Source:* Castro and McQuinn, “Comments submitted to the Federal Trade Commission RE: Informational Injury Workshop.”

Taken together, the Castro-McQuinn taxonomy of informational injuries and PII use and collection can be superimposed over the initial question posed by “The Regulator’s Neural Network.” The result, as shown in Figure 7, is a more detailed means by which regulators can address the basic question of whether an “unfair consumer injury” exists that requires regulatory redress.

Level	Description	Potential for Information Injury
Level 0	No collection and use	None
Level 1	Collection and no use	Low
Level 2	Collection and use (no human)	Low
Level 3	Collection and use (with human)	High

Table 2: Describing the levels of PII collection and use and how each level has the potential for different levels of informational injury. *Source:* Castro and McQuinn, “Comments submitted to the Federal Trade Commission RE: Informational Injury Workshop.”

Using these taxonomies and flowcharts, policymakers will be far-better positioned to make more informed decisions regarding the appropriate means by which they might address clear consumer harms resulting from the misapplication of algorithmic systems.

When confronting calls for greater transparency in AI architectural design, policymakers need to weigh the purported benefits of such approaches against the likely costs of their implementation. As discussed previously, algorithmic transparency itself is a disputed concept, often conflated with the essential features of algorithmic accountability and seldom cognizant of the technical limitations of AI systems. In contrast, the process of oversight and accountability described in this section makes use of evidence-based frameworks that provide for remedying known harms, while pains-

takingly detailing the very transparency that detractors often opine for. Policymakers have a clear roadmap here for how to strike a winning balance between ensuring the public is safeguarded against unfair and deceptive automated decisions and protecting the dynamism necessary for innovation and entrepreneurship to continue flourishing.

Unlike frameworks touting the value of algorithmic transparency, these operational rules for algorithmic accountability can actually be implemented while ensuring America’s AI industry continues to lead the world in research, development, and commercialization.

3. Policy Recommendations

- a. Promote industry self-regulation as a means of developing best practices and standards in AI research;
- b. Regulations, if necessary, should be based on tangible and known costs and benefits, not hypotheticals;
- c. The FTC should recognize a framework for “algorithmic accountability” as the ideal approach to regulating AI/ML, while continuing to promote self-regulatory governance mechanisms in lieu of broader omnibus privacy rules;
- d. Policymakers should direct resources towards examining how implementing specific rules for “algorithmic accountability” could address potential harms in sector-specific contexts;
- e. Support and pass the AI in Government Act; and
- f. Abstain from granting the FTC rule-making authority.

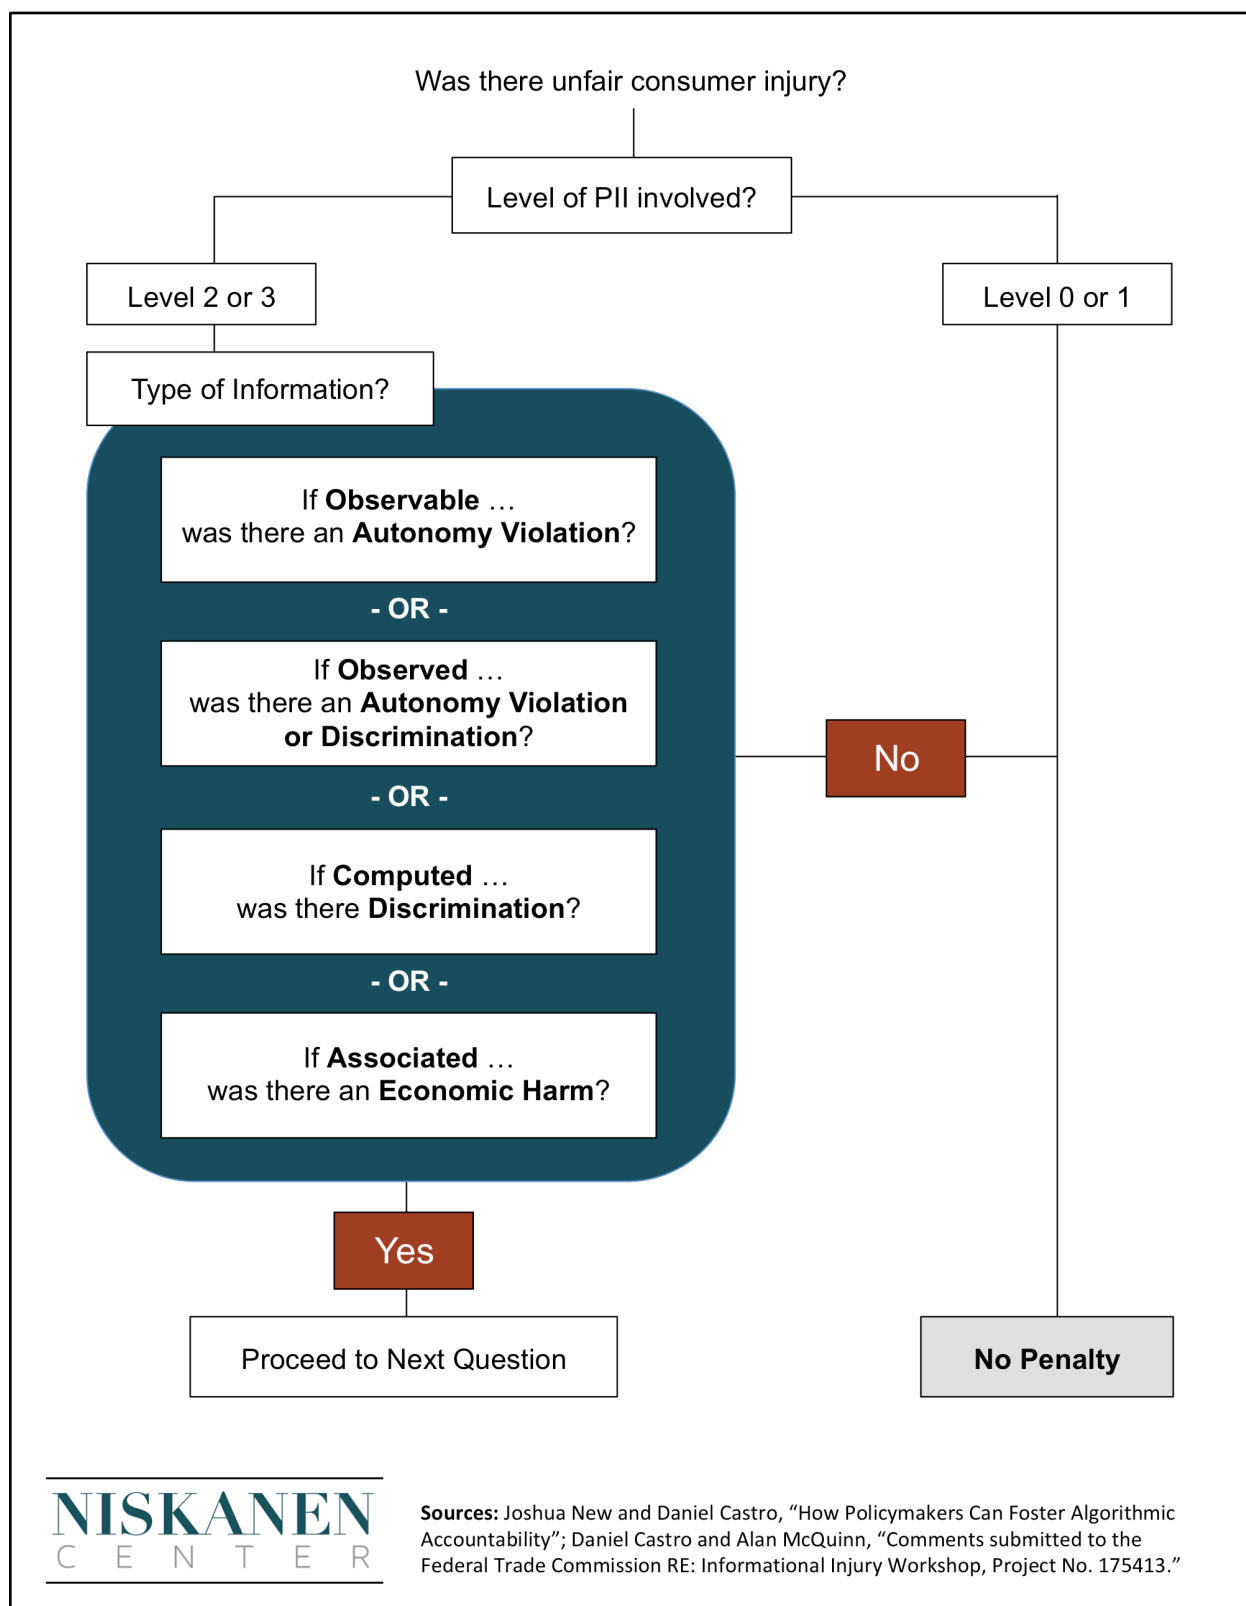
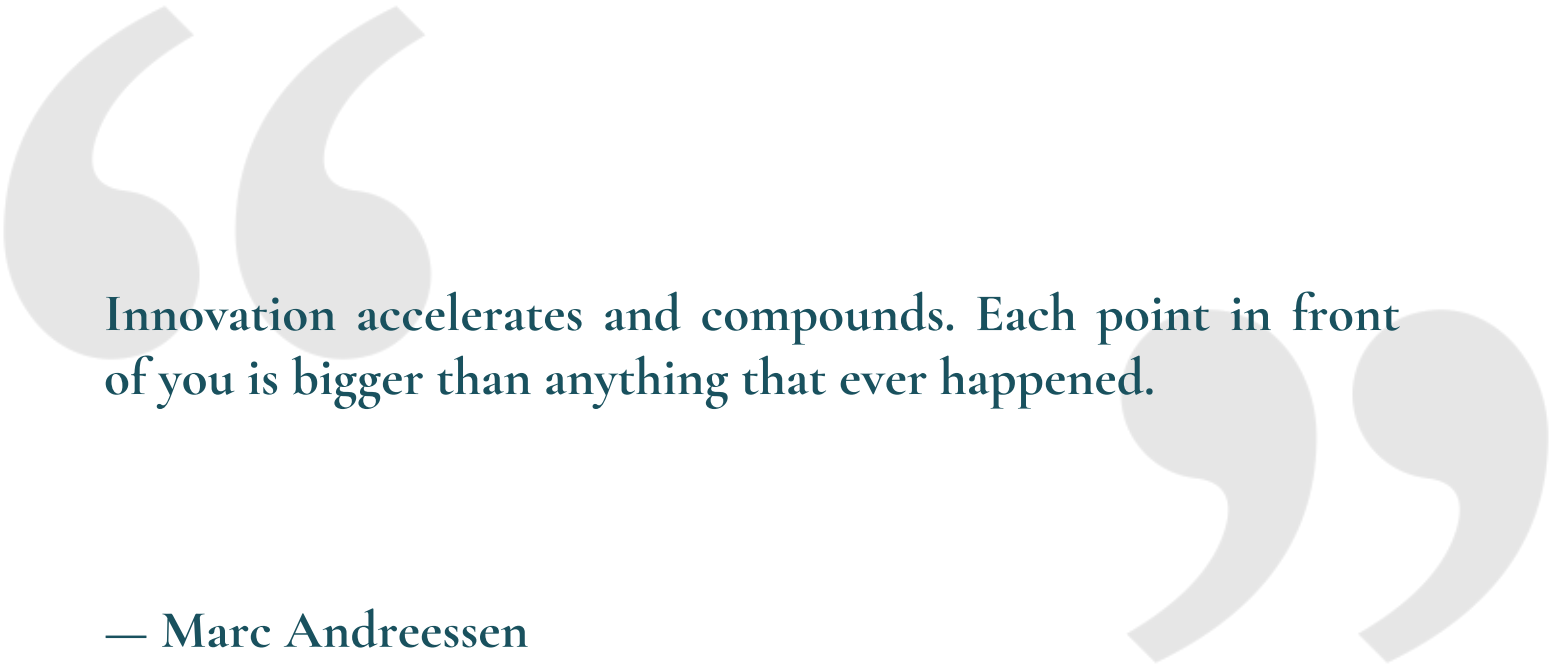


Figure 7: Applying the Castro-McQuinn taxonomy of informational injury to the New-Castro "Regulator's Neural Network" flowchart to answer whether an "unfair consumer injury" occurred as a result of an algorithmic decision-making system. Assessing the existence of such an injury can be determined using the Level of PII collection and use and the type of information collected and used. *Sources:* Niskanen Center; Castro and McQuinn, "Comments submitted to the Federal Trade Commission RE: Informational Injury Workshop"; New and Castro, "How Policymakers Can Foster Algorithmic Accountability."



Innovation accelerates and compounds. Each point in front of you is bigger than anything that ever happened.

— Marc Andreessen

CONCLUSION

The recommendations contained in this analysis can serve as signposts on the road to a brighter future. Like any sign, however, it can only *guide* policymakers towards the better path; it cannot ease the burdens of the journey. Policymakers should be skeptical of proposals that would purport to provide easy solutions to these complex problems. It is notoriously difficult to chart a clear and easy course to a better future without expecting to experience some difficulties, and not every decision that seems wise at any given moment is necessarily paving the way to something better — that’s why rules that allow for maximum flexibility and adaptability are far more ideal than those that propose a One True Golden Path. The end of the 2007 movie *Charlie Wilson’s War* provides a pitch-perfect summation of this lesson.

Having successfully helped the Afghan mujahedeen drive the Soviet army out of Afghanistan, Gust, the gruff CIA officer played by Philip Seymour Hoffman, asks Congressman Wilson, played by Tom Hanks, if he’s ever heard the story of the Zen Master and the Little Boy.

Gust: There was a little boy, and on his 14th birthday he gets a horse. And everybody in the village says, “How wonderful! The boy got a horse!” And the Zen master says, “We’ll see.” Two years later the boy falls off the horse, breaks his leg, and everybody in the village says, “How terrible!” And the Zen master says, “We’ll see.” Then, a war breaks out and all the young men have to go off and fight, except the boy can’t ‘cause his leg’s all messed up, and everybody in the village says, “How wonderful!”

Charlie: And the Zen master says, “We’ll see.”

Gust: So you get it

Charlies: No ... [laughing] No, ‘cause I’m stupid.

Gust: You’re not stupid, you’re just in Congress.¹³⁵

The story is intended to showcase the dangers of excessive triumphalism in the face of an ever-changing, dynamic world. But it also provides a lesson in the hazards of defining a path towards the future based on the assumptions of the moment — a lesson that is particularly relevant in the field of technology policy and regulatory governance.

Returning once more to *The Future and Its Enemies*, Postrel echoes similar cautions for those policymakers who would attempt to craft a “particular, carefully outlined future” designed according to a single technocratic vision. “The future,” she writes, “will be as grand, and as particular, as we are. We cannot build a single bridge from here to there, for neither here nor there is a single point. And there is no abyss to cross.”¹³⁶

REFERENCES

“SOFT LAW” IS EATING THE WORLD

Ryan Hagemann, Jennifer Skees, and Adam Thierer, “Soft Law for Hard Problems: The Governance of Emerging Technologies in an Uncertain Future,” *Colorado Technology Law Journal* (forthcoming), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3118539.

Ryan Hagemann, “New Rules for New Frontiers: Regulating Emerging Technologies in an Era of Soft Law,” *Washburn Law Journal*, Vol. 57, No. 2: 235-263 (Spring 2018), <http://washburnlaw.edu/publications/wlj/issues/57-2.html>.

Ryan Hagemann, Jennifer Huddleston Skees, and Adam Thierer, “‘Soft Law’ Is Eating the World,” *The Bridge*, 11 Oct. 2018, <https://www.mercatus.org/bridge/commentary/soft-law-eating-world-driverless-car>.

Ryan Hagemann, “All Roads Lead to Soft Law,” Niskanen Center, 11 June 2018, <https://niskanencenter.org/blog/all-roads-lead-to-soft-law/>.

Ryan Hagemann and Joshua Hampson, *Comments submitted to the Bureau of Industry and Security in the Matter of: Emerging Technology and Research Advisory Committee Meeting*, submitted 14 Mar. 2017, https://niskanencenter.org/wp-content/uploads/2017/03/NiskanenCenter_CommentsETRACMeetingBIS.pdf.

Ryan Hagemann, “New Rules for New Frontiers: A Regulatory Manifesto for Emerging Technologies,” Niskanen Center, 30 Jan. 2017, <https://niskanencenter.org/blog/new-rules-new-frontiers-regulatory-manifesto-emerging-technologies/>.

Ryan Hagemann and Alec Stapp, *Comments submitted to the National Telecommunications and Information Administration in the Matter of: International Internet Policy Priorities*, Docket No. 180124068-8068-01, 17 July 2018, <https://niskanencenter.org/wp-content/uploads/2018/07/Comments-International-Internet-Policy-NTIA.pdf>

INNOVATION IN BITS

Antitrust

Alec Stapp, “Reports of Antitrust’s Death Have Been Greatly Exaggerated: Economics, Law, and Technology in the Supreme Court’s *Amex* Decision,” Niskanen Center (Washington, D.C.: July 2018), <https://niskanencenter.org/wp-content/uploads/2018/07/Reports-of-Antitrusts-Death-Have-Been-Greatly-Exaggerated.pdf>.

Ryan Hagemann, “Data Price Gouging: A Stalking Horse for a Neo-Brandeisian Antitrust Doctrine?,” Niskanen Center (Washington, D.C.: 8 May 2018), https://niskanencenter.org/wp-content/uploads/2018/05/Brief-Data-Price-Gouging_-A-Stalking-Horse-for-a-Neo-Brandeisian-Antitrust-Docctrine_.pdf.

Ryan Hagemann, “The Theory of the (Tech) Firm,” Niskanen Center, 28 Nov. 2017, <https://niskanencenter.org/blog/theory-tech-firm/>.

Should Washington break up Big Tech?, American Enterprise Institute panel with Ryan Hagemann, Andrew McAfee, Michael R. Strain, Luigi Zingales, and Jim Pethokoukis, 27 Nov. 2017, <https://www.aei.org/events/should-washington-break-up-big-tech/>.

Privacy

Ryan Hagemann and Joshua Hampson, “Encryption, Trust, and the Online Economy: An Assessment of the Economic Benefits Associated with Encryption,” Niskanen Center (Washington, D.C.: 9 Nov. 2015), https://niskanencenter.org/wp-content/uploads/2015/11/RESEARCH-PAPER_EncryptionEconomicBenefits.pdf.

Ryan Hagemann, “The Tech Side of the Encryption Debate,” 1776, 18 Dec. 2015, <https://www.1776.vc/insights/encryption-privacy-security-terrorists-terrorism-global-economy-san-bernardino-paris-surveillance-snow-den-revelations/>.

Ryan Hagemann and Andrew Chang, “Encryption and Lack of Trust in Big Brother,” *Wall Street Journal*, 5 May 2016, <https://www.wsj.com/articles/encryption-and-lack-of-trust-in-big-brother-1462385212>.

Ryan Hagemann, “The Path Forward on Encryption: The McCaul-Warner Commission,” *Lawfare*, 24 June 2016, <https://www.lawfareblog.com/path-forward-encryption-mccaul-warner-commission>.

Decoding the Encryption Dilemma: A Conversation on Backdoors, Going Dark, and Cybersecurity, Information Technology and Innovation panel presentation with David Bitkower, Chris Calabrese, Daniel Castro, Ryan Hagemann, Bruce J. Heiman, Jules Polonetsky, and Morgan Reed, 31 Mar. 2016, <https://itif.org/events/2016/03/31/decoding-encryption-dilemma-conversation-backdoors-going-dark-and-cybersecurity>.

Copyright

Regina Zernay, “The Quest for Automated Copyright Enforcement,” Niskanen Center, 26 Oct. 2016, <https://niskanencenter.org/blog/quest-automated-copyright-enforcement/>.

Regina Zernay and Ryan Hagemann, “ACES in the Hole? Automated Copyright Enforcement Systems and the Future of Copyright Law,” Niskanen Center (Washington, D.C.: 6 June 2017), https://niskanencenter.org/wp-content/uploads/2017/06/AutomatedCopyrightSystems_Final.pdf.

Ryan Hagemann, *Comments submitted to the United States Trade Representative in the Matter of: A Rebuttal to “A Request for Comment on the 2016 Special 301 Out-of-Cycle Review of Notorious Markets,”* Docket No. USTR-2016-2013, submitted 20 Oct. 2016, https://niskanencenter.org/wp-content/uploads/2016/10/NiskanenCenter_USTRCommentsNotoriousMarketsRebuttal.pdf.

Ryan Hagemann, “2017 Policy Priorities: A Copyright Regime for the Digital Age,” Niskanen Center, 14 Dec. 2016, <https://niskanencenter.org/blog/2017-policy-priorities-copyright-regime-digital-age/>.

INNOVATION IN ATOMS

Genetic Modification

Dr. Anastasia Greenberg, “Lighting a Path to Precision Medicine: Regulatory and Policy Implications of Optogenetic Technology,” Niskanen Center (Washington, D.C.: Sep. 2018), <https://niskanencenter.org/wp-content/uploads/2018/09/Research-Paper-Policy-Implications-of-Optogenetics-Greenberg.pdf>.

Joseph V. Gulfo and Jason Briggeman, “Fostering Resilience in the Medical Marketplace: A Plan for Reform of Pharmaceutical Regulation,” Niskanen Center (Washington, D.C.: June 2018), <https://niskanencenter.org/wp-content/uploads/2018/07/Fostering-Resilience-in-the-Medical-Marketplace.pdf>.

Joseph V. Gulfo, Jason Briggeman, and Ethan C. Roberts, “The Proper Role of the FDA for the 21st Century,” Mercatus Center at George Mason University (Arlington, VA: 2016), <https://www.mercatus.org/system/files/Gulfo-Proper-Role-FDA-vi.pdf>.

Ryan Hagemann, Elena Milova, and Keith Comito, *Comments submitted to the Food and Drug Administration in the Matter of: Expedited Programs for Regenerative Medicine Therapies for Serious Conditions*, Docket No. FDA-2017-D-6159, submitted 15 Feb. 2018, <https://niskanencenter.org/wp-content/uploads/2018/02/Comment-Regenerative-Medicine-FDA.pdf>.

Ryan Hagemann, *Comments submitted to the Food and Drug Administration in the Matter of: Benefit-Risk Assessments in Drug Regulatory Decision-Making*, Docket No. FDA-2017-N-4076-0001, submitted 19 Oct. 2017, <https://niskanencenter.org/wp-content/uploads/2017/10/Comments-Benefit-Risk-Assessments-FDA.pdf>.

Ryan Hagemann, “The Coming Age of Genetic Modification, Part I,” Niskanen Center, 8 Aug. 2017, <https://niskanencenter.org/blog/coming-age-genetic-modification-part-i/>.

Ryan Hagemann, “The Coming Age of Genetic Modification, Part II,” Niskanen Center, 10 Aug. 2017, <https://niskanencenter.org/blog/coming-age-genetic-modification-part-ii/>.

Ryan Hagemann, “The Parallel Fears Driving Perceptions of AI and Genomics,” Niskanen Center, 30 Aug. 2017, <https://niskanencenter.org/blog/parallel-fears-driving-perceptions-ai-genomics/>.

Internet of Things

Ryan Hagemann, *Comments submitted to the National Telecommunications and Information Administration in the Matter of: Green Paper: Fostering the Advancement of the Internet of Things*, Docket No. 170105023-7023-01, submitted 9 Feb. 2017, https://www.ntia.doc.gov/files/ntia/publications/niskanencenter_commentsiotgreenpaperntia.pdf.

Ryan Hagemann, *Comments submitted to the National Telecommunications and Information Administration in the Matter of: The Benefits, Challenges, and Potential Roles for the Government in Fostering the Internet of Things*, Docket No. 160331306-6306-01, submitted 23 May 2016, https://www.ntia.doc.gov/files/ntia/publications/niskanencenter_ntia_iot_comments.pdf.

Managing the Opportunities and Risks of the Internet of Things and Big Data, 2016 Internet Governance Forum panel presentation with Dan Caprio, Jeff Brueggeman, Alan Davidson, Michelle De Mooy, Ryan Hagemann, and Dean Garfield, 22 June 2016, <https://www.igf-usa.org/panel-emerging-technologies-that-affect-internet-governance-managing-the-opportunities-and-risks-of-the-internet-of-things-and-big-data/>.

Ryan Hagemann, “2017 Policy Priorities: Making Way for the Internet of Everything,” Niskanen Center, 30 Nov. 2016, <https://niskanencenter.org/blog/2017-policy-priorities-making-way-internet-everything/>.

Ryan Hagemann, *Comments submitted to the National Telecommunications and Information Administration in the Matter of: Promoting Stakeholder Action Against Botnets and Other Automated Threats*, Docket No. 170602536-7536-01, submitted 28 July 2017, https://www.ntia.doc.gov/files/ntia/publications/niskanencenter_comments_botnets_ntia.pdf.

Autonomous Vehicles

Adam Thierer and Ryan Hagemann, “Removing Roadblocks to Intelligent Vehicles and Driverless Cars,” *Wake Forest Journal of Law & Policy*, Vol. 5, No. 2: 339-391 (2015), https://wfulawpolicyjournal.com.files.wordpress.com/2016/05/5-thierer-hagemann_final.pdf.

Ryan Hagemann, *Comments submitted to the National Highway Traffic Safety Administration in the Matter of: Federal Automated Vehicle Policy*, Docket No. NHTSA-2016-0090, submitted 21 Nov. 2016, <https://niskanencenter.org/wp-content/uploads/2016/11/CommentsAutonomousVehicleStandardsNHTSA.pdf>.

Ryan Hagemann, *Comments submitted to the Department of Transportation in the Matter of: Automated Vehicle Policy Summit*, Docket No. DOT-OST-2018-0017, submitted 5 Mar. 2018, <https://niskanencenter.org/wp-content/uploads/2018/03/Comment-Automated-Vehicle-Policy-Summit-DOT.pdf>.

Ryan Hagemann, *Comments submitted to the National Highway Traffic Safety Administration in the Matter of: Removing Regulatory Barriers for Vehicles with Automated Driving Systems*, Docket No. NHTSA-2018-0009, submitted 14 Mar. 2018, <https://niskanencenter.org/wp-content/uploads/2018/03/Comment-Barriers-to-Autonomous-Vehicles-NHTSA.pdf>.

Ryan Hagemann, *Comments submitted to the National Highway Traffic Safety Administration in the Matter of: Automated Driving Systems: A Vision for Safety*, Docket No. NHTSA-2017-0082, submitted 3 Oct. 2017, <https://niskanencenter.org/wp-content/uploads/2017/10/Comments-Autonomous-Vehicle-Guidance-NHTSA.pdf>.

Ryan Hagemann, “Senseless Government Rules Could Cripple the Robo-car Revolution,” *Wired*, 1 May 2017, <https://www.wired.com/2017/05/senseless-government-rules-cripple-robo-car-revolution/>.

Ryan Hagemann, “DSRC Isn’t the Path Forward for Connected Vehicles,” Niskanen Center, 13 June 2017, <https://niskanencenter.org/blog/dsrc-isnt-path-forward-connected-vehicles/>.

Ryan Hagemann, “How Congress Can Accelerate Us Towards the Autonomous Roadway,” Niskanen Center, 27 Feb. 2017, <https://niskanencenter.org/blog/congress-can-accelerate-us-towards-autonomous-roadway/>.

Commercial Drones

Ryan Hagemann, “Consumer Privacy in an Age of Commercial Unmanned Aircraft Systems,” *The Independent Review*, Vol. 23, No. 1: 1-14 (Summer 2018), available at <http://www.independent.org/publications/tir/article.asp?id=1304>.

Ryan Hagemann, “Legislative Analysis of the Federal Aviation Administration Reauthorization Act of 2017, Title II, Subtitle A — Unmanned Aircraft Systems Reform,” Niskanen Center (Washington, D.C.: 27 June 2017), <https://niskanencenter.org/wp-content/uploads/2017/06/FAAReauthorizationActLegislativeAnalysis.pdf>.

Ryan Hagemann, *Comments submitted to the Federal Trade Commission in the Matter of: The Privacy Implications of Commercial Drone Operations*, prepared for the FTC’s seminar addressing drones, 13 Nov. 2016, https://www.ftc.gov/system/files/documents/public_comments/2016/11/00016-129462.pdf.

Eli Dourado, Ryan Hagemann, and Adam Thierer, *Comments submitted to the Federal Aviation Administration in the Matter of: Operation and Certification of Small Unmanned Aircraft Systems*, Docket No. FAA-2015-0150, submitted 24 April 2015, <https://www.mercatus.org/system/files/Dourado-UAS-PIC.pdf>.

Ryan Hagemann, “FAA Drone Rules Could Kill Innovation,” Niskanen Center, 21 June 2016, <https://niskanencenter.org/blog/faa-drone-rules-kill-innovation/>.

Supersonic Flight

Supersonic Myths, Niskanen Center, <http://www.supersonicmyths.com/>.

Eli Dourado and Samuel Hammond, “Drop the Supersonic Aircraft Ban, Watch Business Boom,” *Wall Street Journal*, 13 June 2016, <https://www.wsj.com/articles/drop-the-supersonic-aircraft-ban-watch-business-boom-1465769638>

Eli Dourado and Samuel Hammond, “Make America Boom Again: How to Bring Back Supersonic Transport” Mercatus Center at George Mason University (Arlington, VA: Oct. 2016), <https://www.mercatus.org/system/files/mercatus-dourado-supersonic-transport-vi.pdf>.

Samuel Hammond, “The Return of Supersonic,” Niskanen Center, 19 June 2017, <https://niskanencenter.org/blog/return-of-supersonic/>.

Samuel Hammond, “The Business Case for Supersonic Overland,” Niskanen Center, 28 Feb. 2017, <https://niskanencenter.org/blog/supersonic-overland/>.

Eli Dourado and Michael Kotrous, “Airplane Speeds Have Stagnated for 40 Years,” Mercatus Center at George Mason University, 20 July 2016, <https://www.mercatus.org/publication/airplane-speeds-have-stagnated-40-years>.

Commercial Space

“Commercial Space Actors: Disruptors or Solid Partners for National Security,” NSI Virtual Think Tank Report produced in support of the Strategic Multilayer Assessment Office, Feb. 2018, <https://goo.gl/5hHMD1>.

Joshua Hampson, “The Future of Space Commercialization,” Niskanen Center (Washington, D.C.: 25 Jan. 2017), <https://science.house.gov/sites/republicans.science.house.gov/files/documents/TheFutureofSpaceCommercializationFinal.pdf>.

Joshua Hampson, “National Security Needs Robust Commercial Space,” Niskanen Center, 21 June 2017, <https://niskanencenter.org/blog/national-security-needs-robust-commercial-space/>.

Ryan Hagemann, “Humanity’s Future Among the Stars,” Niskanen Center, 29 Sep. 2016, <https://niskanencenter.org/blog/humanitys-future-among-stars/>.

Joshua Hampson, “Committing to Commercial Space Launch,” Niskanen Center, 6 Jan. 2017, <https://niskanencenter.org/blog/committing-commercial-space-launch/>.

Joshua Hampson, “Making Commercial Space Great,” *The Hill*, 31 Jan. 2017, <http://thehill.com/blogs/congress-blog/technology/317120-making-commercial-space-great>.

Climate Engineering

Dr. Joseph Majkut, “Essay for the Forum on U.S. Solar Geoengineering Research,” Harvard’s Solar Geoengineering Research Program, 24 Mar. 2017, https://geoengineering.environment.harvard.edu/files/sgrp/files/forum_report.pdf.

Dr. Joseph Majkut, Adam Wong, and Ryan Hagemann, *Comments submitted to the American Geophysical Union in the Matter of: Geoengineering Responses to Climate Change Require Enhanced Research, Consideration of Societal Impacts, and Policy Development*, submitted 25 Sep. 2017, <https://niskanencenter.org/wp-content/uploads/2017/10/Niskanen-Center-Comments-Climate-Engineering-AGU-2.pdf>.

Dr. Joseph Majkut, *Statement concerning geoengineering research to the House Committee on Science, Space, and Technology*, 8 Nov. 2017, <https://science.house.gov/sites/republicans.science.house.gov/files/documents/HHRG-115-SY18-WState-JMajkut-20171108.pdf>.

ARTIFICIAL INTELLIGENCE

AI in Digital Advertising

Ryan Hagemann, *Comments submitted to the Federal Trade Commission in the Matter of: Hearing on Competition and Consumer Protection in the 21st Century: The Consumer Welfare Implications Associated With the Use of Algorithmic Decision Tools, Artificial Intelligence, and Predictive Analytics*, Docket Number FTC-2018-0056, Project Number P181201, submitted 20 Aug. 2018, <https://niskanencenter.org/wp-content/uploads/2018/08/Comments-Consumer-Welfare-Implications-of-AI-FTC.pdf>.

Alec Stapp and Ryan Hagemann, *Comments submitted to the Federal Trade Commission in the Matter of: Hearing on Competition and Consumer Protection in the 21st Century: The Identification and Measurement of Market Power and Entry Barriers, and the Evaluation of Collusive, Exclusionary, or Predatory Conduct That Violates the Consumer Protection Statutes Enforced by the FTC, in Markets Featuring “Platform” Businesses*, Docket Number FTC-2018-0050, Project Number P181201, submitted 20 Aug. 2018, <https://niskanencenter.org/wp-content/uploads/2018/08/Comments-Market-Power-of-Platforms-FTC.pdf>.

Alec Stapp and Ryan Hagemann, *Comments submitted to the Federal Trade Commission in the Matter of: Hearing on Competition and Consumer Protection in the 21st Century: The Intersection Between Privacy, Big Data, and Competition*, Docket Number FTC-2018-0051, Project Number P181201, submitted 20 Aug. 2018, <https://niskanencenter.org/wp-content/uploads/2018/08/Comments-Privacy-Big-Data-and-Competition-FTC.pdf>.

AI in Medical Devices

Dr. Anastasia Greenberg and Ryan Hagemann, *Comments submitted to the Food and Drug Administration in the Matter of: Software Precertification Program (vo.2)*, Docket No. FDA-2017-N-4301-0001, submitted 18 July 2018, <https://niskanencenter.org/wp-content/uploads/2018/07/Comments-Software-Precert-FDA.pdf>.

Dr. Anastasia Greenberg and Ryan Hagemann, *Comments submitted to the Food and Drug Administration in the Matter of: The Benefits and Risks of Software as a Medical Device: A Response to a Request for Input RE: “Development of 21st Century Cures Act Section 3060 Required Report,”* Docket No. FDA-2018-N-1910, submitted 27 June 2018, <https://niskanencenter.org/wp-content/uploads/2018/06/Comments-Section-3060-Software-Device-Guidance-FDA-FINAL.pdf>.

Algorithmic Accountability

Curt Levey and Ryan Hagemann, “Algorithms With Minds of Their Own,” *Wall Street Journal*, 12 Nov. 2017, <https://www.wsj.com/articles/algorithms-with-minds-of-their-own-1510521093>.

Joshua New and Daniel Castro, “How Policymakers Can Foster Algorithmic Accountability,” Center for Data Innovation (Washington, D.C.: 21 May 2018), <http://www2.datainnovation.org/2018-algorithmic-accountability.pdf>.

Ryan Hagemann, *Comments submitted to the Office of Science and Technology Policy in the Matter of: A Request for Information on Artificial Intelligence*, Docket No. 2016-15082, submitted 22 July 2016, <https://niskanencenter.org/wp-content/uploads/2016/07/CommentsArtificialIntelligencePolicyOSTP.pdf>.

Ryan Hagemann, “2017 Policy Priorities: Embracing the Ghost in the Machine,” Niskanen Center, 17 Nov. 2016, <https://niskanencenter.org/blog/2017-policy-priorities-embracing-ghost-machine/>.

Ryan Hagemann, “The Creeping Hysteria Over Artificial Intelligence,” Niskanen Center, 17 Apr. 2017, <https://niskanencenter.org/blog/creeping-hysteria-artificial-intelligence/>.

Ryan Hagemann, “Why, Robot?,” Niskanen Center, 7 Sep. 2016, <https://niskanencenter.org/blog/why-robot/>.

Daniel Castro and Alan McQuinn, *Comments submitted to the Federal Trade Commission RE: Informational Injury Workshop, Project No. 175413*, Information Technology and Innovation Foundation (Washington, D.C.: 27 Oct. 2017), <http://www2.itif.org/2017-informational-injury-comments.pdf>.

APPENDIX A: FOUR-CATEGORY SAFETY AND EFFECTIVENESS PARADIGMS

A.1. Drugs and Biologics

Proposed safety and effectiveness paradigm based on type of evidence provided			
Safety: determination of safety is to be made relative to the conditions of use specified by the sponsor, per the FD&C Act—“safe for use under the conditions prescribed, recommended, or suggested in the proposed labeling”—not in anticipation of potential uses or abuses of the product outside the claim sought in the approval application. Special emphasis is placed on the likelihood of use causing death, debilitation, or severe harm and on ways to mitigate these risks.			
Effectiveness: categories consistent with the nature of the endpoints used to demonstrate substantial evidence of effectiveness. Labeling will be color coded to facilitate communications to physicians, patients, and other groups.			
Category 1: Biomarkers	Category 2: Clinical Signs and Symptoms	Category 3: Disease Modulation / Modification	Category 4: Clinical Outcomes and Survival
Improvement in a biomarker known to be elevated or decreased in patients with specific diseases—for example, fasting blood glucose, hemoglobin A1c, carcinoembryonic antigen (CEA), CD4/CD8 ratio, prostate specific antigen (PSA), blood clotting (INR), LDL cholesterol, HDL cholesterol, etc.)	Reduction in pain; improvement in activities of daily living; tumor response (size, local control, improved progression-free interval); improvement in forced expiratory volume; improved walking distance; improved bone mineral density; improved treadmill performance and EKG findings (atrial fibrillation, premature ventricular contractions); patient-reported outcomes; etc.	Reduction in flares of diarrhea, arthritis, or headache; reduction in suicidal ideation; fewer heart failure readmissions; reduction in joint space narrowing; reduction in use of other medications (steroids); reduction in development of deep vein thrombosis or pulmonary embolism; reduction in unstable angina; etc.	Improvement in survival; reduction on major cardiac events (myocardial infarction, heart failure, rehospitalization); etc.

A.2. Diagnostics

Proposed safety and effectiveness paradigm based on type of evidence provided			
Safety: determination of safety is to be made relative to the conditions of use specified by the sponsor, per the FD&C Act—“safe for use under the conditions prescribed, recommended, or suggested in the proposed labeling”—not in anticipation of potential uses or abuses of the product outside the claim sought in the approval application. Special emphasis is placed on the likelihood of use causing death, debilitation, or severe harm and on ways to mitigate these risks.			
Effectiveness: categories consistent with the nature of the data used to demonstrate substantial evidence of effectiveness. Labeling will be color coded to facilitate communications to physicians, patients, and other groups.			
Category 1: Associated with disease or current state of disease in patients with an established diagnosis when used alone or when considered with other diagnostic tests and clinical information	Category 2: Predicts safety and effectiveness in patients receiving drug/biologic therapy	Category 3: Predicts for disease presence or progression	Category 4: Information provided by the test induces interventions that favorably alter the natural history of the disease
<u>Examples:</u> <ol style="list-style-type: none"> 1. Measurement above a threshold is associated with disease recurrence. 2. Rising level is associated with progression of disease. 	<u>Examples:</u> <ol style="list-style-type: none"> 1. Companion diagnostics. 2. Test correlates with drug/biologic effect, taken during drug therapy to determine whether: <ol style="list-style-type: none"> a. continued treatment is likely to be safe; or b. clinical response is likely (clinical signs and symptoms, disease modulation, clinical outcomes and survival). 	<u>Examples:</u> <ol style="list-style-type: none"> 1. Screening test that enables diagnosis earlier than currently available methods. 2. Test in patients with established diagnosis identifies those at higher risk for progression and other poor outcomes (clinical measures: disease burden or severity, survival, progression, or quality of life, etc.). 	<u>Examples:</u> <ol style="list-style-type: none"> 1. Screening test leads to initiation of therapy (surgery, drug, device) that results in improved survival or quality of life. 2. Test in patients at high risk or with established diagnosis leads to initiation of therapy (surgery, drug, device) that results in improved survival or quality of life.

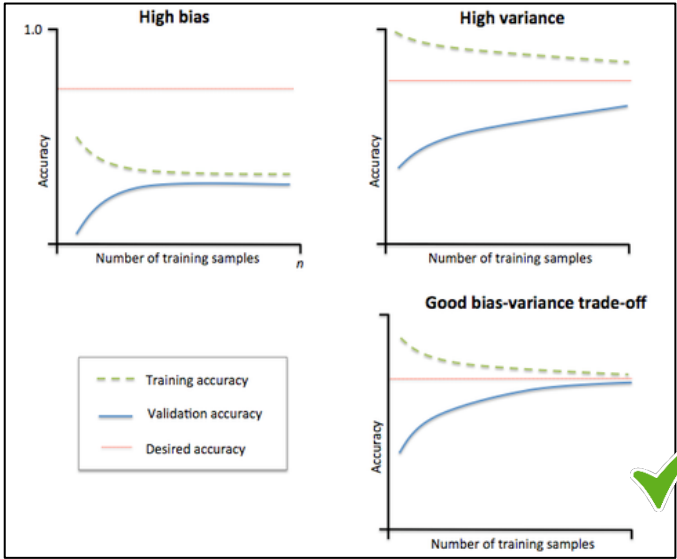
A.3. Devices

Proposed safety and effectiveness paradigm based on type of evidence provided			
Safety: determination of safety is to be made relative to the conditions of use specified by the sponsor, per the FD&C Act—“safe for use under the conditions prescribed, recommended, or suggested in the proposed labeling”—not in anticipation of potential uses or abuses of the product outside the claim sought in the approval application. Special emphasis is placed on the likelihood of use causing death, debilitation, or severe harm and on ways to mitigate these risks.			
Effectiveness: categories consistent with the nature of the data used to demonstrate substantial evidence of effectiveness. Labeling will be color coded to facilitate communications to physicians, patients, and other groups.			
Physical Action		Clinical Sequelae	
Category 1: Tools - Used in conjunction with diagnostic or therapeutic intervention (surgery or drug delivery)	Category 2: Used to diagnose disease, to provide treatment, or to repair or replace damaged or nonfunctional or dysfunctioning tissues	Category 3: Clinical Improvement	Category 4: Improved Clinical Outcomes
<u>Example:</u> 1. Used to help conduct or facilitate diagnostic or therapeutic procedures.	<u>Examples:</u> 1. Diagnostic equipment. 2. Eradicate, ablate, or destroy tissue. 3. Enhance, augment, or substitute functioning of tissues or organs.	<u>Examples:</u> 1. Improve or stabilize clinical signs and symptoms of disease. 2. Reduce complications from surgery or drug therapy. 3. Disease modulation or modification.	<u>Examples:</u> 1. Disease progression; progression-free survival. 2. Reduce major cardiovascular events (MACE). 3. Survival.

APPENDIX B: AI DEVELOPMENT BEST PRACTICES FOR MEDICAL PRACTICES

Organizational Domain	Elements	Excellence Principles					KPIs
		PS	PQ	ClinR	CybR	PC	
Artificial Intelligence Development Best Practices	Assessment of potential problematic machine learning (ML) features/variables that include biased or non-representative data.	X	X			X	<p>Descriptive statistics of all features (input variables) and output variables that were used to build the ML classifier/algorithm.</p> <p><i>Melanoma Example:</i> Proportion of male/female in dataset; average and standard deviation of variables such as: age, number of previous health conditions, family history of melanoma; representative example of skin biopsy image and averaged skin biopsy image; number of positive and negative melanoma examples in dataset; etc.</p>
	A strong theoretical justification for how various patient data are related to the diagnosis/treatment of a given disease/condition to help to prevent unintended consequences of complex ML models.	X	X	X		X	<p>Explanation of reasons for including each feature as a variable for training the algorithm.</p> <p><i>Melanoma Example:</i> Age was included as a variable because likelihood of melanoma is known to increase with age. Family history of melanoma was included as a variable because there is an established heritable component of melanoma; etc.</p>

	Ensuring generalizability of ML model to new data.	X	X	X		X	<p>Description of how data were split into independent training, validation, and test sets to avoid biasing the final model.</p> <p><i>Melanoma Example:</i> Training examples with skin biopsies are labeled as positive or negative for melanoma. The dataset is obtained across different patients and randomized to split 60 percent into the training set, while saving 20 percent for the validation set and the remaining 20 percent for the test set. During model development, the model/training parameters are only modified based on results from the validation set, and never from the test set.</p>
	Ensuring that the model is measuring what it claims to measure.	X	X	X		X	<p>Description of how the initial dataset was labeled for training. This includes an explanation of how a “ground truth” was established for labeling the data in supervised models. For example, was the ground truth established based on clinical data? Was multiple observer averaging used to make a final labeling determination?</p> <p><i>Melanoma Example:</i> The skin biopsy images were evaluated by three radiologists for presence or absence of melanoma. Only images with 100 percent agreement across radiologists were included in the model development dataset.</p>
	Assessment for unintended data alterations.		X			X	<p>Description of any preprocessing and/or data cleaning steps and verifications that these steps did not alter the data in unintended ways.</p> <p><i>Melanoma Example:</i> Principle Component Analysis (PCA) was performed on the skin biopsy images to reduce data dimensionality and a Gaussian filter was applied to smooth the images. Following these processing steps, a random subset of the processed images were re-evaluated by two radiologists for melanoma with no disagreement between the unprocessed and processed images, indicating that no</p>

						unintended alterations have occurred.
	Identification of model building flaws.	X	X	X	X	<p>An inclusion of the learning curves during model training to allow for identification of major model building issues, such as models that show high bias (i.e., underfitting) or high variance (i.e., overfitting). High bias or high variance is an indication that the model will not generalize well to new data.</p> <p><i>Melanoma Example:</i> Training and validation curves were plotted for melanoma detection model development and the curves are within an acceptable bias/variance trade-off range:</p>  <p>Figure 8: Various possible learning curves obtained during machine learning algorithm development. Source: https://sebastianraschka.com/faq/docs/ml-solvable.html.</p>

	Assessment of model accuracy.	X	X	X			<p>Reporting of model accuracy scores based on calculations that are appropriate for a given ML situation. For example, F1 scores should be included for models trained on input data with rare events (i.e., low positive-to-negative ratio in the input data -presence of disease is very rare) to avoid misleading accuracy scores.</p> <p><i>Melanoma Example:</i> The dataset of skin biopsy images includes 1 percent positive melanoma cases. After model development was completed, the F1 score was calculated to be 0.95, showing a good level of model prediction accuracy.</p>
	Commitment to model updates based on real-world performance	X	X		X		<p>Evidence of plans and mechanisms for obtaining new data after SaMD release for use in improvement to model accuracy and generalizability.</p> <p><i>Melanoma Example:</i> Mechanisms are put in place to continuously obtain skin biopsy images from new patients with radiologist establishment of presence/absence of Melanoma. Following substantial amount of new data collection, plans are in place to re-train the model and to retest improvement to model accuracy, as well as update the SaMD when accuracy and generalizability are shown to have improved.</p>

¹ Virginia Postrel, *The Future and Its Enemies: The Growing Conflict Over Creativity, Enterprise, and Progress* (Simon & Schuster, Inc.: New York, NY, 1998), ISBN 978-0-684-86269-9, p. xiv.

² *Id.*

³ Daniel J. Boorstin, “The Fertile Verge: Creativity in the United States,” address to the Carnegie Symposium on Creativity, Inaugural Meeting of the Council of Scholars of the Library of Congress (No. 19-20, 1980), pp. 3-4. (“America was a land of verges — all sorts of verges, between kinds of landscapes or seascapes, between stages of civilization, between ways of thought and ways of life. During our first centuries we experienced more different kinds of verges, and more extensive and more vivid verges, than any other great modern nation. ... The creativity — the hope — of the nation was in its verges, in its new mixtures and new confusions.”)

⁴ Postrel, *supra* note 1, at 208.

⁵ *Id.* at 208-209.

⁶ Joel Mokyr, “Creative Forces,” *Reason*, May 1993, p. 65.

⁷ Alexis de Tocqueville, *Democracy in America*, Vol. 2 (New York, Alfred A. Knopf, 1945), p. 279.

⁸ White House, *A Framework for Global Electronic Commerce*, 1 July 1997, (hereafter, the *Framework*), <https://clintonwhitehouse4.archives.gov/WH/New/Commerce/summary.html>.

⁹ Blake Harris, “Ira Magaziner: A Framework for Global Electronic Commerce,” *GovTech*, 20 Sep. 1997, <http://www.govtech.com/magazines/gt/Ira-Magaziner-A-Framework-for-Global.html>.

¹⁰ *Fostering the Advancement of the Internet of Things*, Department of Commerce, Internet Policy Task Force and Digital Economy Leadership Team, 12 Jan. 2017, https://www.ntia.doc.gov/files/ntia/publications/iot_green_paper_01122017.pdf.

¹¹ Gary E. Marchant and Braden Allenby, “New Tools for Governing Emerging Technologies,” *Bulletin of the Atomic Scientists*, Vol. 73, No. 108 (2017), p. 112.

¹² Ryan Hagemann, Jennifer Skees, and Adam Thierer, “Soft Law for Hard Problems: The Governance of Emerging Technologies in an Uncertain Future,” *Colorado Technology Law Journal* (forthcoming), p. 10, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3118539.

¹³ Ryan Hagemann, “New Rules for New Frontiers: Regulating Emerging Technologies in an Era of Soft Law,” *Washburn Law Journal*, Vol. 57, No. 2: 235-263 (Spring 2018), <http://washburnlaw.edu/publications/wlj/issues/57-2.html>.

¹⁴ *Id.* at 244-255. (These costs can potentially include, among other things: “(1) Diminished long-term legal clarity, given the lack of common law adjudication of soft criteria and the soft law systems they engender; (2) Subjecting agencies to criticism that such approaches ignore the rule of law and provide new avenues by which industry interests can engage in regulatory capture, thereby undercutting institutional legitimacy; and (3) Opening the door to policy entrepreneurs motivated by hostility to technology and progress, ideologically dogmatic policy preferences, or nefarious intentions.”)

¹⁵ Robert A. Hoerr, “Regulatory uncertainty and the associated business risk for emerging technologies,” *Journal of Nanoparticle Research*, Vol. 13: 1513-1520 (2011), p. 1514.

¹⁶ Hagemann, et. al., *supra* note 12, at 11-13.

¹⁷ *Cutting Red Tape: Administrative Simplification in the Netherlands*, OECD Directorate for Public Governance, ISBN 9789264037496, (2007), available at <http://www.oecd.org/gov/regulatory-policy/cuttingredtapeadministrativesimplificationinthenetherlands.htm>; N.M.J. Kurstjens, “The Dutch ‘Polder Model’: Social Pacts,” Nijmegen School of Management at Radboud University (Nijmegen, Netherlands: Aug. 2015), p. 6, https://theses.uibn.ru.nl/bitstream/handle/123456789/1179/Kurstjens%2C_Nanda_1.pdf?sequence=1. (“Policymakers view the Dutch system of cooperation between government, unions and employers as a guiding example of successfully spreading and sharing power within the field of political economy. This ‘Polder Model’ is the frequently praised way of policymaking due to its ability to bring antagonists into a mode of dialogue, for developing agreements based on consultation, compromise and consensus and for organising a support base for its pacts and policies. Not without reason, the Netherlands is characterised as a consensus democracy and a traditional corporatist country.”)

¹⁸ Marc Andreessen, “Why Software Is Eating the World,” *Wall Street Journal*, 20 Aug. 2011, <https://www.wsj.com/articles/SB10001424053111903480904576512250915629460>.

¹⁹ Ryan Hagemann, Jennifer Huddleston Skees, and Adam Thierer, “‘Soft Law’ Is Eating the World,” *The Bridge*, 11 Oct. 2018, <https://www.mercatus.org/bridge/commentary/soft-law-eating-world-driverless-car>.

-
- ²⁰ “Initial Estimates Show Digital Economy Accounted for 6.5 Percent of GDP in 2016,” National Telecommunications and Information Administration, accessed 23 Oct. 2018, <https://www.ntia.doc.gov/blog/2018/initial-estimates-show-digital-economy-accounted-65-percent-gdp-2016>.
- ²¹ Craig Mundie, “Privacy Pragmatism,” *Foreign Affairs*, Vol. 93, No. 2 (March/April 2014), p. 517, <http://www.foreignaffairs.com/articles/140741/craig-mundie/privacy-pragmatism>.
- ²² Mary Meeker, “Internet Trends Report 2018,” Kleiner Perkins, 30 May 2018, <https://www.kleinerperkins.com/perspectives/internet-trends-report-2018>.
- ²³ Jay Shambaugh, et. al., “The State of Competition and Dynamism: Facts about Concentration, Start-Ups, and Related Policies,” Hamilton Project at the Brookings Institution, June 2018, https://www.brookings.edu/wp-content/uploads/2018/06/ES_THP_20180611_CompetitionFacts_20180611.pdf.
- ²⁴ Elyse Dorsey, et. al., “Hipster Antitrust Meets Public Choice Economics: The Consumer Welfare Standard, Rule of Law, and Rent-Seeking,” *CPI Antitrust Chronicle*, 18 Apr. 2018, <https://www.competitionpolicyinternational.com/wp-content/uploads/2018/04/CPI-Dorsey-Rybnick-Wright.pdf>.
- ²⁵ Alec Stapp, “You Can’t Understand Big Tech Without Understanding Network Effects. Here’s a Road Map,” Niskanen Center, 13 Sep. 2018, <https://niskanencenter.org/blog/you-cant-understand-big-tech-without-understanding-network-effects-heres-a-road-map/>.
- ²⁶ “2017 Global Innovation 1000,” PriceWaterhouseCooper, 27 Oct. 2017, <https://www.strategyand.pwc.com/media/file/2017-Global-Innovation-1000-Fact-Pack.pdf>.
- ²⁷ Bernardo Huberman, et. al., “Valuating privacy,” *IEEE security & privacy*, Vol. 3, Issue 5: 22-25 (2005), https://papers.ssrn.com/sol3/papers.cfm?Abstract_id=488324.
- ²⁸ “Global 500 Companies to Spend \$7.8B on GDPR Compliance,” International Association of Privacy Professionals, accessed 24 Oct. 2018, <https://iapp.org/news/a/survey-fortune-500-companies-to-spend-7-8b-on-gdpr-compliance/>.
- ²⁹ NTIA *supra* note 20.
- ³⁰ 47 U.S.C § 230, <https://www.law.cornell.edu/uscode/text/47/230>.
- ³¹ Ben Sisario, “Blurred Lines’ Verdict Upheld by Appeals Court,” *The New York Times*, 21 Mar. 2018, <https://www.nytimes.com/2018/03/21/business/media/blurred-lines-marvin-gaye-copyright.html>.
- ³² Jordan Runtagh, “Songs on Trial: 12 Landmark Music Copyright Cases,” *Rolling Stone*, 8 June 2016, <https://www.rollingstone.com/politics/politics-lists/songs-on-trial-12-landmark-music-copyright-cases-166396/>.
- ³³ Ryan Browne, “Catastrophic: EU passes controversial copyright law that could hit the likes of Google and Facebook,” *CNBC*, 12 Sep. 2018, <https://www.cnbc.com/2018/09/12/eu-lawmakers-pass-controversial-digital-copyright-law.html>.
- ³⁴ White House, *The National Information Infrastructure: Agenda for Action*, 15 Sep. 1993, <https://clintonwhitehouse6.archives.gov/1993/09/1993-09-15-the-national-information-infrastructureagenda-for-action.html>.
- ³⁵ Eric Faulkner, et al., “Ensuring Patient Access to Regenerative and Advanced Therapies in Managed Care: How Do We Get There?” *Journal of Managed Care Medicine*, 2018, <https://alliancerm.org/wp-content/uploads/2018/05/JMCMArm.pdf>.
- ³⁶ Ben Hirschler, “Gene Therapy Is Now Available, but Who Will Pay for It?” *Scientific American*, 8 Aug. 2017, <https://www.scientificamerican.com/article/gene-therapy-is-now-available-but-who-will-pay-for-it/>.
- ³⁷ Francesca Cook, et al., “Regenerative Medicine Is Here: New Payment Models Key To Patient Access,” *In Vivo*, July/Aug. 2018, available at <https://invivo.pharmaintelligence.informa.com/IV005371/Regenerative-Medicine-Is-Here-New-Payment-Models-Key-To-Patient-Access>.
- ³⁸ “New Medical Services and New Technologies,” Centers for Medicare & Medicaid Services, 17 Oct. 2018, <https://www.cms.gov/Medicare/Medicare-Fee-for-Service-Payment/AcuteInpatientPPS/newtech.html>.
- ³⁹ Ana Stojanovska, “How to Navigate Cell and Gene Therapy Reimbursement: Part One,” AmerisourceBergen Corporation, 29 June 2018, <https://www.amerisourcebergen.com/abcnew/insights/how-to-navigate-cell-and-gene-therapy-reimbursement-part-one>.
- ⁴⁰ *Id.*
- ⁴¹ *Id.*

-
- ⁴² Arlene Weintraub, “Gene Therapy Is Booming, But How Will We Manage The Costs?” *Forbes*, 1 Dec. 2017, <https://www.forbes.com/sites/arleneweintraub/2017/12/01/gene-therapy-is-booming-but-how-will-we-manage-the-costs/>.
- ⁴³ Joseph A. DiMasi, Ronald W. Hansen, and Henry G. Grabowski, “The Price of Innovation: New Estimates of Drug Development Costs,” *Journal of Health Economics*, Vol. 22, No. 2: 151-185 (2003), available at <https://www.ncbi.nlm.nih.gov/pubmed/12606142>.
- ⁴⁴ Joseph V. Gulfo and Jason Briggeman, “Fostering Resilience in the Medical Marketplace: A Plan for Reform of Pharmaceutical Regulation,” Niskanen Center (Washington, D.C.: July 2018), p. 3, <https://niskanencenter.org/wp-content/uploads/2018/07/Fostering-Resilience-in-the-Medical-Marketplace.pdf>.
- ⁴⁵ *Id.*
- ⁴⁶ *Id.* at 4.
- ⁴⁷ *Id.*
- ⁴⁸ Joseph V. Gulfo, Jason Briggeman, and Ethan C. Roberts, “The Proper Role of the FDA for the 21st Century,” Mercatus Center at George Mason University (Arlington, VA: 2016), p. 20, <https://www.mercatus.org/system/files/Gulfo-Proper-Role-FDA-v1.pdf>.
- ⁴⁹ *Id.* at 29.
- ⁵⁰ *Id.*
- ⁵¹ *Id.*
- ⁵² *Fostering the Advancement of the Internet of Things*, National Telecommunications and Information Administration, 12 Jan. 2017, p. 40, https://www.ntia.doc.gov/files/ntia/publications/iot_green_paper_01122017.pdf. (“The Department reaffirms its commitment to the policy approach that has made the United States the leading innovation economy. This approach is reflected in the 1997 Framework for Global Electronic Commerce, and has been maintained across all subsequent Presidential administrations.”); see also Ryan Hagemann, *Comments submitted to the National Telecommunications and Information Administration in the Matter of: Green Paper: Fostering the Advancement of the Internet of Things*, Docket No. 170105023-7023-01, submitted 8 Feb. 2017, https://www.ntia.doc.gov/files/ntia/publications/niskanencenter_commentsiotgreenpaperntia.pdf.
- ⁵³ Kaveh Waddell, “How Much Is Encryption Worth to the Economy?” *The Atlantic*, 9 Nov. 2015, <https://www.theatlantic.com/politics/archive/2015/11/how-much-is-encryption-worth-to-the-economy/458466/>.
- ⁵⁴ Ryan Hagemann and Josh Hampson, “Encryption, Trust, and the Online Economy: An Assessment of the Economic Benefits Associated With Encryption,” Niskanen Center (Washington, D.C.: 9 Nov. 2015), https://niskanencenter.org/wpcontent/uploads/2015/11/RESEARCH-PAPER_EncryptionEconomicBenefits.pdf.
- ⁵⁵ Ryan Hagemann, “How to Cure What Ails American Cybersecurity,” *RealClearPolicy*, 18 May 2017, http://www.realclearpolicy.com/articles/2017/05/18/how_to_cure_what_ails_american_cybersecurity.html. (“By turning the current VEP policy into a law, the PATCH Act ensures a continuity of policy across administrations. That certainty is necessary to ensure a stable and consistent approach to a key facet of federal cybersecurity policy. It also embraces the necessary TAO (transparency, accountability, and oversight) of surveillance reform, carefully balancing intelligence-gathering operations and the individual security needs of consumers.”); Ryan Hagemann, “Balancing Cybersecurity and National Security,” Niskanen Center, 28 July 2016, <https://niskanencenter.org/blog/balancing-cybersecurity-national-security/>. (“Congress should pass legislation that would formalize the process for VEP. Executive orders can be rescinded at any time, leading to real concern that any headway made towards a better, more transparent disclosure process could be lost following the Obama Administration’s exit. Implementation of the VEP policy should be enshrined in law. Congress can play a leading role in helping to shape better, more transparent, and more stable cybersecurity practices by taking the lead on legislation.”)
- ⁵⁶ Joe Uchill, “When governments turn spyware on citizens,” *Axios*, 20 Sep. 2018, <https://www.axios.com/governments-spyware-technology-israel-fb950f76-1b3b-4d60-82db-6e8ff188ee44.html>.
- ⁵⁷ 18 U.S.C. § 2523(b)(4)(E) (2018).
- ⁵⁸ Dr. Anastasia Greenberg, “NAFTA 2.0 Will Benefit the Digital Economy,” Niskanen Center, 9 Oct. 2018, <https://niskanencenter.org/blog/nafta-2-0-will-benefit-the-digital-economy/>.
- ⁵⁹ Ryan Hagemann, *Comments submitted to the United States Trade Representative in the Matter of: A Rebuttal to “A Request for Comment on the 2016 Special 301 Out-of-Cycle Review of Notorious Markets,”* Docket No. USTR-2016-2013, submitted 20

Oct. 2016,

https://niskanencenter.org/wpcontent/uploads/2016/10/NiskanenCenter_USTRCommentsNotoriousMarketsRebuttal.pdf.

⁶⁰ Ryan Hagemann, *Comments submitted to the National Telecommunications and Information Administration in the Matter of: Promoting Stakeholder Action Against Botnets and Other Automated Threats*, Docket No. 170602536-7536-01, submitted 28 July 2017, https://www.ntia.doc.gov/files/ntia/publications/niskanencenter_comments_botnets_ntia.pdf.

⁶¹ *Id.*

⁶² In particular, we would direct policymakers' attention to this 2017 technical white paper written in response to the Trump Administration's Executive Order 13800: Industry Technical White Paper, Communications Sector Coordinating Council, 17 July 2017, https://docs.wixstatic.com/ugd/0a1552_18ae07afc1bo4aa1bd13258087a9c77b.pdf.

⁶³ Patricia Moloney Figliola, "Federal Quantum Information Service," Congressional Research Service, 2 July 2018, <https://fas.org/sgp/crs/misc/IF10872.pdf>.

⁶⁴ *Id.*

⁶⁵ *Id.*

⁶⁶ Carten Cordell, "OSTP forms new subcommittee to focus on quantum technology," *FedScoop*, 22 June 2018, <https://www.fedscoop.com/ostp-forms-new-subcommittee-focus-quantum-technology/>.

⁶⁷ National Highway Traffic Safety Administration, *Preliminary Statement of Policy Concerning Automated Vehicles*, 30 May 2013, pp. 12–13, http://www.nhtsa.gov/staticfiles/rulemaking/pdf/Automated_Vehicles_Policy.pdf.

⁶⁸ Adam Thierer and Ryan Hagemann, "Removing Roadblocks to Intelligent Vehicles and Driverless Cars," *Wake Forest Journal of Law & Policy*, Vol. 5, No. 2: 339–391 (2015), p. 377, https://wfulawpolicyjournal.com.files.wordpress.com/2016/05/5-thierer-hagemann_final.pdf.

⁶⁹ Ryan Hagemann, "Senseless Government Rules Could Cripple the Robo-car Revolution," *Wired*, 1 May 2017, <https://www.wired.com/2017/05/senseless-government-rules-cripple-robo-car-revolution/>.

⁷⁰ *Id.*

⁷¹ Michaela Ross, "Regulatory Chill May Pivot Connected Vehicle Tech's Course," *Bloomberg Law*, 8 Feb. 2017, <https://www.bna.com/regulatory-chill-may-n57982083525/>.

⁷² Alex Kreilein, "Dedicated Short Range Communications (DSRC) Expose Critical Gaps in Security and Privacy," *SecureSet*, 29 Mar. 2017, <https://glenechogroup.isebox.net/securesetaccelerator/dedicated-short-range-communications-dsrc-expose-critical-gaps-in-security-and-privacy>.

⁷³ 49 C.F.R. § 571.3.

⁷⁴ Ryan Hagemann, *Comments submitted to the National Highway Traffic Safety Administration in the Matter of: Removing Regulatory Barriers for Vehicles with Automated Driving Systems*, Docket No. NHTSA-2018-0009, submitted 14 Mar. 2018, <https://niskanencenter.org/wp-content/uploads/2018/03/Comment-Barriers-to-Autonomous-Vehicles-NHTSA.pdf>.

⁷⁵ Marc Scribner, "Reforming Air Traffic Control with the 21st Century AIRR Act," *Competitive Enterprise Institute*, 21 June 2017, <https://cei.org/blog/reforming-air-traffic-control-21st-century-airr-act>.

⁷⁶ *Fast-Forwarding to a Future of On-Demand Urban Air Transportation*, Uber Elevate, 27 Oct. 2016, <https://www.uber.com/elevate.pdf>.

⁷⁷ Brent Skorup and Melody Calkins, "Why not auction off low-altitude airspace for exclusive use?," *Technology Liberation Front*, 27 June 2017, <https://techliberation.com/2017/06/27/why-not-auction-off-low-altitude-airspace-for-exclusive-use/>.

⁷⁸ "Presidential Memorandum for the Secretary of Transportation," White House, 25 Oct. 2017, <https://www.whitehouse.gov/presidential-actions/presidential-memorandum-secretary-transportation/>.

⁷⁹ *Id.*

⁸⁰ "Alphabet, Apple and Microsoft will be part of government drone pilots, but Amazon was left out," *Reuters*, 10 May 2018, <https://www.cnn.com/2018/05/09/alphabet-apple-microsoft-part-of-usdot-drone-pilots-amazon-bypassed.html>.

⁸¹ Miriam McNabb, "The UAS Integration Pilot Program is Underway: North Dakota Gets Started with Parazero Parachutes for Drones," *Drone Life*, 16 Aug. 2018, <https://dronelife.com/2018/08/15/the-uas-integration-pilot-program-is-underway-north-dakota-gets-started-with-parazero-parachutes-for-drones/>.

⁸² Charlie Rose, "Amazon's Jeff Bezos Looks to the Future," *CBS News*, 1 Dec. 2013, <https://www.cbsnews.com/news/amazons-jeff-bezos-looks-to-the-future/>.

-
- ⁸³ Michael Kratsios, “Trump aide: The future of American aviation is all about drones,” *CNN*, 6 June 2018, <https://www.cnn.com/2018/06/06/opinions/trump-administration-drone-technology-kratsios/index.html>.
- ⁸⁴ Samuel Hammond, “The Business Case for Supersonic Overland,” Niskanen Center, 28 Feb. 2017, <https://niskanencenter.org/blog/supersonic-overland/>.
- ⁸⁵ See the FAA’s 2008 policy statement on “Civil Supersonic Airplane Noise Type Certification Standards and Operating,” available at https://www.faa.gov/about/office_org/headquarters_offices/apl/noise_emissions/supersonic_aircraft_noise/media/noise_policy_on_supersonics.pdf.
- ⁸⁶ Eli Dourado and Samuel Hammond, “Drop the Supersonic Aircraft Ban, Watch Business Boom,” *Wall Street Journal*, 13 June 2016, <https://www.wsj.com/articles/drop-the-supersonic-aircraft-ban-watch-business-boom-1465769638>.
- ⁸⁷ For a longer discussion of this issue, see Eli Dourado and Samuel Hammond, “Make America Boom Again: How to Bring Back Supersonic Transport,” Mercatus Center at George Mason University (Arlington, VA: Oct. 2016), <https://www.mercatus.org/system/files/mercatus-dourado-supersonic-transport-v1.pdf>.
- ⁸⁸ See Boom Supersonic, <https://boomsupersonic.com/>.
- ⁸⁹ Eli Dourado and Raymond Russell, “Airport Noise NIMBYism: An Empirical Investigation,” Mercatus Center at George Mason University (Arlington, VA: 17 Oct. 2016), <https://www.mercatus.org/publications/airport-noise-nimbyism>.
- ⁹⁰ Eli Dourado, “Lee-Gardner Amendment would reduce supersonic fuel burn by 20 percent or more,” Boom Supersonic, 27 June 2017, <https://blog.boomsupersonic.com/lee-gardner-amendment-would-reduce-supersonic-fuel-burn-by-20-percent-or-more-34a189465af4>.
- ⁹¹ The Space Foundation, “The Space Report: 2016,” 2016, http://www.spacefoundation.org/sites/default/files/downloads/The_Space_Report_2016_OVERVIEW.pdf.
- ⁹² Sarah Kramer and David Mosher, “Here’s how much money it actually costs to launch stuff into space,” *Business Insider*, 20 July 2016, <http://www.businessinsider.com/spacex-rocket-cargo-price-by-weight-2016-6/#does-this-sound-ridiculously-expensive-10>.
- ⁹³ See Joshua Hampson, “The Future of Space Commercialization,” Niskanen Center (Washington, D.C.: 25 Jan. 2017), <https://science.house.gov/sites/republicans.science.house.gov/files/documents/TheFutureofSpaceCommercializationFinal.pdf>.
- ⁹⁴ *Id.* at 14.
- ⁹⁵ Joshua Hampson, “Why National Security Needs Commercial Outer Space,” *RealClearFuture*, 7 Oct. 2016, http://www.realclearfuture.com/articles/2016/10/07/why_national_security_needs_commercial_outer_space_111944.html.
- ⁹⁶ Avery Thompson, “ULA’s New Rocket Will Be Powered by Blue Origin’s Engine,” *Popular Mechanics*, 27 Sep. 2018, <https://www.popularmechanics.com/space/rockets/a23492986/ulas-new-rocket-will-be-powered-by-blue-origins-engine/>.
- ⁹⁷ Hampson, *supra* note 93, at 16.
- ⁹⁸ Joshua Hampson, “Making commercial space great,” *The Hill*, 31 Jan. 2017, <https://thehill.com/blogs/congress-blog/technology/317120-making-commercial-space-great>.
- ⁹⁹ *Id.* (“Companies are often left uninformed about why a mission was denied in the review process — a process often seen as arbitrary. Within the remote imaging license process, for example, rejections or revocations of licenses are often done without an explanation. If this secretive approach becomes the norm for the space economy, space companies may be unable to raise financing. This can lead to companies that are initially interested in entering space markets to ultimately avoid doing so altogether. Certainty will be key to ensuring the space economy flourishes.”)
- ¹⁰⁰ Ryan Hagemann, “Humanity’s Future Among the Stars,” Niskanen Center, 29 Sep. 2016, <https://niskanencenter.org/blog/humanitys-future-among-stars/>.
- ¹⁰¹ Jane Long, et. al., “Geoengineering: A national strategic plan for research on the potential effectiveness, feasibility, and consequences of climate remediation technologies,” Task Force on Climate Remediation Research, Bipartisan Policy Center (Washington, D.C.: 2011), <https://bipartisanpolicy.org/wp-content/uploads/sites/default/files/BPC%20Climate%20Remediation%20Final%20Report.pdf>.
- ¹⁰² National Research Council, “Climate Intervention: Reflecting Sunlight to Cool Earth,” National Academies Press (Washington, DC: 2015), <https://www.nap.edu/catalog/18988/climate-intervention-reflecting-sunlight-to-cool-earth>.

¹⁰³ “Climate Change: A Coordinated Strategy Could Focus Federal Geoengineering Research and Inform Governance Efforts,” Report to the Chairman, Committee on Science and Technology, House of Representatives, United States Government Accountability Office, GAO-10-903, Sep. 2010, <http://www.gao.gov/new.items/d10903.pdf>.

¹⁰⁴ Moshe Y. Vardi, “Artificial Intelligence: Past and Future,” *Communications of the ACM*, Vol. 55, No. 1: 5 (Jan. 2012), <https://cacm.acm.org/magazines/2012/1/144824-artificial-intelligence-past-and-future/fulltext>.

¹⁰⁵ Alessandro Acquisti, et. al, “The Economics of Privacy,” *Journal of Economic Literature*, Vol. 52, No. 2 (Mar. 2016), p. 48, <https://dx.doi.org/10.2139/ssrn.2580411>.

¹⁰⁶ Ryan Hagemann, *Comments submitted to the Federal Trade Commission in the Matter of: Hearing on Competition and Consumer Protection in the 21st Century: The Consumer Welfare Implications Associated With the Use of Algorithmic Decision Tools, Artificial Intelligence, and Predictive Analytics*, Docket Number FTC-2018-0056, Project Number P181201, submitted 20 Aug. 2018, p. 2, <https://niskanencenter.org/wp-content/uploads/2018/08/Comments-Consumer-Welfare-Implications-of-AI-FTC.pdf>.

¹⁰⁷ *Consumers driving the digital uptake: The economic value of online advertising-based services for consumers*, McKinsey & Company, September 2010, p. 5, http://www.youronlinechoices.com/white_paper_consumers_driving_the_digital_uptake.pdf.

¹⁰⁸ *Id.*

¹⁰⁹ Larry Downes, “A Rational Response to the Privacy ‘Crisis,’” Policy Analysis No. 716, Cato Institute (Washington, D.C.: 7 Jan. 2013), pp. 4-5, <https://object.cato.org/sites/cato.org/files/pubs/pdf/pa716.pdf>.

¹¹⁰ Sami Main, “Programmatic Digital Display Ads Now Account for Nearly 80% of US Display Spending,” *AdWeek*, 18 Apr. 2017, <https://www.adweek.com/tv-video/programmatic-digital-display-ads-now-account-for-nearly-80-of-usdisplay-spending/>.

¹¹¹ McKinsey & Company, *supra* note 107.

¹¹² Leon Bottou, et. al., “Counterfactual Reasoning and Learning Systems: The Example of Computational Advertising,” *Journal of Machine Learning Research*, Vol. 14 (2013), <https://www.microsoft.com/en-us/research/wp-content/uploads/2013/11/bottou13a.pdf>.

¹¹³ Downes, *supra* note 109, at 6.

¹¹⁴ Main, *supra* note 110.

¹¹⁵ For a more thorough accounting of the particular costs incurred by digital advertisers, see Hagemann, *supra* note 106 at 5. (“Every year, digital ad fraud — in particular, “invalid traffic” automated systems that artificially inflate the number of clicks, impressions, views, etc. with the aim of generating revenue for the perpetrators — costs online advertisers billions of dollars, and continues to skyrocket. In 2015, a report commissioned by the Interactive Advertising Bureau found that digital ad fraud cost the industry \$8.2 billion, with invalid traffic comprising the largest segment of those costs (\$4.6 billion). A 2017 report from Juniper Research forecasted that these costs could jump to \$19 billion in 2018. While AI/ML technologies are likely being employed in the commission of ad fraud, they are also being harnessed as a potential solution. Existing systems, such as PPC Protect, and forthcoming innovations, such as the NOIZ decentralized digital advertising platform, promise to significantly curtail these costs, minimizing the deadweight losses for advertisers and diminished online experience for users.”)

¹¹⁶ *Id.* at 11.

¹¹⁷ *Id.* at 13.

¹¹⁸ *Id.* at 15. (“Strict, one-size-fits-all data protection and privacy regulations are fundamentally anti-consumer and anti-innovation — anti-consumer because they drive up costs and diminish competition, and anti-innovation because they stifle the development of new products and services while incentivizing firms to prioritize regulatory compliance over maximizing consumer welfare. Additionally, by assuming the existence of a problem requiring a solution, such rules fail to adequately assess the costs and benefits of alternative governance strategies. Finally, when compared to a light-touch, sector-based regulatory approach to addressing privacy harms, it’s not at all clear that broad data protection mandates would do a better job of promoting consumer privacy interests, and may actually have the opposite effect.”)

¹¹⁹ Robert W. Hahn and Anne Layne-Farrar, “The Benefits and Costs of Online Privacy Legislation,” AEI-Brookings Joint Center for Regulatory Studies, Working Paper 01-14 (Washington, D.C.: Oct. 2001), p. 43, http://papers.ssrn.com/abstract_id=292649.

¹²⁰ The top technology firms in the United States (Apple, Amazon, Google, Microsoft, Facebook, Intel, Cisco, Oracle, IBM, NVIDIA, Adobe, Texas Instruments, Salesforce, Qualcomm, Micron Technologies) account for over \$5 trillion in total market capitalization; the Asia Pacific region's top technology firms (Tencent, Samsung, Taiwan Semiconductor, Broadcom, Container Store, Sony, Nintendo) account for almost \$1.4 trillion; the top European technology firms (SAP, Accenture, ASML Holding) together account for a mere \$285 billion in total market capitalization. See Meeker, *supra* note 22.

¹²¹ Robert Graboyes and Sara Rogers, "As free innovation encounters health care regulation, thing 'soft laws'," *STAT*, 12 Sep. 2018, <https://www.statnews.com/2018/09/12/free-innovation-health-care-regulation/>.

¹²² *Id.*

¹²³ *Multistakeholder Process on Promoting Software Component Transparency*, National Telecommunications and Information Administration, Notice of Open Meeting, Federal Register Vol. 83, No. 110, 7 June 2018, <https://www.ntia.doc.gov/files/ntia/publications/fr-notice-07192018-meeting-software-component-transparency.pdf>. ("NTIA encourages discussion of approaches and considerations from diverse sectors such as the medical device community, where the applicability of a "bill of materials" has garnered increased discussion and interest.")

¹²⁴ "NTIA Software Component Transparency," National Telecommunications and Information Administration, 2 Oct. 2018, <https://www.ntia.doc.gov/SoftwareTransparency>.

¹²⁵ *Developing Software Precertification Program: A Working Model vo.2*, U.S. Food and Drug Administration, June 2018, p. 13. ("Any organization that intends to develop or market regulated software in the United States would be considered in-scope for the Software Precertification Program.")

¹²⁶ Dr. Anastasia Greenberg and Ryan Hagemann, *Comments submitted to the Food and Drug Administration in the Matter of: Software Precertification Program (vo.2)*, Docket No. FDA-2017-N-4301-0001, submitted 18 July 2018, <https://niskanencenter.org/wp-content/uploads/2018/07/Comments-Software-Precert-FDA.pdf>.

¹²⁷ *A Working Model vo.2*, *supra* note 125, at 19.

¹²⁸ Greenberg and Hagemann, *supra* note 126, at 6.

¹²⁹ Chris Stucchio and Lisa Mahapatra, "A.I. 'Bias' Doesn't Mean What Journalists Say It Means," *Jacobite Magazine*, 29 Aug. 2017, <https://jacobitemag.com/2017/08/29/a-i-bias-doesnt-mean-what-journalists-want-you-to-think-it-means/>. ("Algorithms are nothing more than a repeatable mathematical procedure, executed by a computer. Think of the step-by-step directions you were given in grade school for solving systems of linear equations or the quadratic formula. How could such a formula suddenly become biased against, say, football players? The media is misleading people. When an ordinary person uses the term 'biased,' they think this means that incorrect decisions are made — that a lender systematically refuses loans to blacks who would otherwise repay them. When the media uses the term 'bias', they mean something very different — a lender systematically failing to issue loans to black people regardless of whether or not they would pay them back.")

¹³⁰ *Id.*

¹³¹ *Id.*

¹³² Joshua New and Daniel Castro, "How Policymakers Can Foster Algorithmic Accountability," Center for Data Innovation (Washington, D.C.: 21 May 2018), pp. 1-2, <http://www2.datainnovation.org/2018-algorithmic-accountability.pdf>.

¹³³ Alec Stapp, "Whither Algorithmic Accountability?," Niskanen Center, 24 May 2018, <https://niskanencenter.org/blog/whither-algorithmic-accountability/>. (New and Castro "introduce a novel flowchart to show how a regulator could impose penalties on algorithm operators. In a sense, they've constructed what a *Regulator's Neural Network* might look like when determining the existence and magnitude of an algorithmic harm. While this diagram is excellent at explaining how regulators should discipline operators after a consumer injury has been established, it is also worthwhile to consider how regulators and industry, acting separately or in collaboration, might incorporate best practices for algorithmic accountability in advance of potential harms.")

¹³⁴ Daniel Castro and Alan McQuinn, *Comments submitted to the Federal Trade Commission RE: Informational Injury Workshop*, Project No. 175413, Information Technology and Innovation Foundation (Washington, D.C.: 27 Oct. 2017), <http://www2.itif.org/2017-informational-injury-comments.pdf>.

¹³⁵ *Charlie Wilson's War*, Dir. Mike Nichols, Universal Pictures, 2007, Perf. Tom Hanks and Philip Seymour Hoffman.

¹³⁶ Postrel, *supra* note 1, at 218.