A Simpler World? On Pruning Risks and Harvesting Fruits in an Orchard of Whispering Algorithms

Mariano-Florentino Cuéllar[*]

This article explores some of the law-related challenges and trade-offs associated with the accelerating importance of "artificial intelligence" for organizations and the public. I consider the relevance of emergent phenomena that may arise from the interaction of humans and artificially intelligent systems (even when no one intended or expected such phenomena to arise), new problems in deliberation, and organizational choices with legal consequences that may be triggered by developments in artificial intelligence, and will in turn shape how such technology develops.

Although many of these developments may deliver distinctive benefits and in some cases have little, if any, precedent, in some respects they also evoke some of the experiences that societies confronted in the related realms of food, agriculture, and public health. In those contexts, advancing technology and scientific knowledge reduced the price of feeding ever-larger populations and contributed to substantial improvements in human health. But certain changes, such as large-scale industrialization of the alimentary sector and the rise of motorized transportation available on a massive scale, brought their own societal challenges and second-order regulatory problems that may illuminate trade-offs associated with artificial intelligence.

Ultimately, society's understanding of the emerging role of artificial intelligence and its associated legal dilemmas is best served — as has been

[*] Copyright © 2017 Mariano-Florentino Cuéllar. Justice, Supreme Court of California; former Stanley Morrison Professor of Law, Stanford Law School; and Scholar, Freeman Spogli Institute for International Studies. In developing and exploring the ideas described here, I have benefited from conversations with Fei-Fei Li, Ed Feigenbaum, Dan Boneh, John Seely Brown, Barbara Grosz, Jonathan Zittrain, Amy Zegart, Jon Mayer, Roberta Katz, Margaret Levi, Eli Sugarman, Larry Kramer, George Triantis, and Herbert Lin. I also appreciate helpful assistance from Elissa Winters, Leah Yaffe, and Shannon Galvin, and superb editorial assistance from Sean Yates and the UC Davis Law Review team. Errors are mine.

true in the food and public health contexts — by questioning reassuring assumptions about the self-correcting, social welfare-advancing nature of conventional markets or public deliberation. Societies will instead tend to make more sensible legal choices governing artificial intelligence by taking seriously the changing nature of human goals and values, the reality of time-inconsistent preferences, and subtle and potentially unintended organizational consequences of technologies that are reshaping our relationship to information, institutions, and each other.

TABLE OF CONTENTS

INTRODUCTION

During one of our recent oral arguments, our courtroom was graced with the presence of some sixty fourth graders that managed to stay remarkably engaged and quiet as the lawyers argued a case before us. What follows is a reflection on what strike me as some of the most interesting and potentially important legal questions our society is beginning to encounter as we prepare to turn over the world to these fourth graders. Adults today, in the early twenty-first century, are on one side of a major generational divide. Barring a cataclysm, theirs is the last generation that will be able to remember a world where computers were mostly on desks and in special rooms, and where most people on any given day were not walking around with little supercomputers in their pockets, or on their wrists. That generation of adults — our generation — is also the last one that will remember, however dimly, a world before almost everyone was tethered to a global network that keeps track of where we are and what we are doing.

Now imagine what life will be like for the grandchildren of today's adults — the children of those fourth graders. To do that, we need to take account of changes in the nature and use of the technologies surrounding us right now — technologies increasingly grouped under the rubric of "artificial intelligence." As with the ideas and technologies that once reshaped agriculture and health in decades past, the assortment of machine learning techniques and networked processes to harness data grouped under the artificial intelligence label will almost certainly reshape markets, politics, institutions, and societal norms. These changes, in turn, underscore the trade-offs and choices society confronts about how law structures our relationship to technology, information, and ultimately, each other.

My premise in what follows is that whether those trade-offs are ignored or confronted, they cannot be avoided. In order to better understand those trade-offs, I explore some of the most intriguing law-related challenges associated with the accelerating importance of artificial intelligence for organizations and the public. I consider the importance of emergent phenomena that may arise from the interaction of humans and artificially intelligent systems even when no one intended or expected such phenomena to arise, new problems in deliberation, and organizational choices with legal consequences that may be triggered by developments in artificial intelligence and will in turn help shape how such technologies further develop. Although many of these developments are distinctive and in some cases have little if any precedent, they also evoke some of the experiences that societies confronted in the related realms of food, agriculture, and

public health. In those contexts, advancing technology and scientific knowledge reduced the price of feeding ever-larger populations and contributed to substantial improvements in human health. But some changes, such as large-scale industrialization of the alimentary sector and the rise of motorized transportation on a massive scale, brought their own societal challenges and second-order regulatory problems that may illuminate trade-offs associated with artificial intelligence.

Ultimately, our understanding of the emerging role of artificial intelligence and its associated legal dilemmas is best served — as is true of food and public health contexts — by questioning reassuring assumptions about the self-correcting, social welfare-advancing nature of conventional markets or public deliberation. Societies will instead do well to make legal choices governing artificial intelligence by taking seriously the changing nature of human preferences, the reality of time-inconsistent preferences, and subtle and potentially unintended organizational consequences of technologies that are reshaping our relationship to information, institutions, and each other.

I. FOUR CONVERGING FACTORS RESHAPING THE SOCIETAL TERRAIN

First, we are getting strikingly better at amassing and analyzing huge amounts of data about human behavior — both in real time and over time. Some of this is a function of massively falling prices of storage technology and processing power over the last few decades.[1] There is change underway in the software and statistical techniques we use to analyze data.[2] And we have cheaper and more ubiquitous ways to gather data, from cameras, exercise bands, GPS, electronic tests, Internet searches, and so on.[3] According to one estimate, only about

[1] In the three decades between 1980 and 2011, computer storage costs dropped "by a factor of ten roughly every four years." For example, in 1984, a single gigabyte of data cost $85,000, whereas it cost about 5 cents in 2011. Orin S. Kerr, *The Next Generation Communications Privacy Act*, 162 U. PA. L. REV. 373, 391 (2014) (citing John Villasenor, *Recording Everything: Digital Storage as an Enabler of Authoritarian Governments*, BROOKINGS INST., Dec. 14, 2011, at 3, https://www.brookings.edu/wp-content/uploads/2016/06/1214_digital_storage_villasenor.pdf).

[2] *See generally* Jonathan Shaw, *Why "Big Data" Is a Big Deal: Information Science Promises to Change the World*, HARV. MAG., Mar.–Apr. 2014, at 30, http://harvardmagazine.com/2014/03/why-big-data-is-a-big-deal.

[3] One study comparing fitness tracker privacy and security found that most fitness tracking applications included in the study sent fitness and other data to company servers. For example, Jawbone "periodically transmits" geolocation information to its servers when the application is open. Additionally, many applications sent generated fitness data such as "steps taken, stairs climbed, stairs down, calories burned, speed, sleep depth, . . . heart rate and blood pressure."

25 percent of the world's stored information was digitized in 2000, whereas almost all of the world's stored information was digitized in 2012.[4] As of 2008, there were more devices on the Internet than people on the planet, and by 2020, an estimated 50 billion devices will be connected to the Internet.[5] The so-called "Internet of things"[6] will not only add to those means of gathering data; it is a vision that emphasizes the networked nature of these sensing technologies, allowing for more routine aggregation of such data so that people and organizations might better predict when behavior will stay constant and when it might change.[7] Which is why some observers are inclined to suggest that big data will make our world simpler, and its complex patterns more discernible.[8]

Second, our social norms seem to be changing when it comes to what we expect from computers. People today *routinely* accept that an algorithm is recommending not only music selections or movies to watch,[9] but whom to date[10] and how to invest.[11] The shift is from a

---

ANDREW HILTS, CHRISTOPHER PARSONS & JEFFREY KNOCKEL, OPEN EFFECT, EVERY STEP YOU FAKE: A COMPARATIVE ANALYSIS OF FITNESS TRACKER PRIVACY AND SECURITY 19-20, 22 (2016), https://openeffect.ca/reports/Every_Step_You_Fake.pdf.

[4] As of 2012, that was about 2.7 zettabytes of digitally stored information. Elizabeth E. Joh, *Policing by Numbers: Big Data and the Fourth Amendment*, 89 WASH. L. REV. 35, 39 (2014).

[5] PRESIDENT'S NAT'L SEC. TELECOMMS. ADVISORY COMM., NSTAC REPORT TO THE PRESIDENT ON THE INTERNET OF THINGS ES-2 (2014), https://www.dhs.gov/sites/default/files/publications/NSTAC%20Report%20to%20the%20President%20on%20the%20Internet%20of%20Things%20Nov%202014%20(updat%20%20%20.pdf.

[6] The "Internet of Things" ("IoT") refers to the "decentralized network" of everyday objects — from kitchen appliances to car tires — that are connected to the Internet and can process and communicate data about the physical world. *See id.* at ES-1; Scott R. Peppet, *Freedom of Contract in an Augmented Reality: The Case of Consumer Contracts*, 59 UCLA L. Rev. 676, 699 (2012).

[7] *See* Andrew Guthrie Ferguson, *The Internet of Things and the Fourth Amendment of Effects*, 104 CALIF. L. REV. 805, 813-18 (2016).

[8] *See* Shaw, *supra* note 2.

[9] Services that use algorithms to suggest music and movies such as Pandora, Spotify, and Netflix were among the most downloaded iPhone applications in 2016. *See* Lisa Eadicicco, *These Are the Most Popular iPhone Apps of 2016*, TIME (Dec. 6, 2016), http://time.com/4592864/most-popular-iphone-apps-2016.

[10] One study shows that in 2005, "few Americans" had tried online dating, whereas 15 percent of U.S. adults today have tried it. And attitudes toward online dating are becoming more positive: in 2005, only 44 percent of U.S. adults indicated that "online dating is a good way to meet people," whereas 59 percent agreed with this statement in 2015. Aaron Smith & Monica Anderson, *5 Facts About Online Dating*, PEW RES. CTR. (Feb. 29, 2016), http://www.pewresearch.org/fact-tank/2016/02/29/5-facts-about-online-dating.

[11] *See* John Melloy, *Algorithms Replacing Wall Street Analysts, Investors*, CNBC

world where it is almost quaint to accept that cars will drive themselves to one where it seems pretty straightforward to imagine that cars will not only choose the route but also the destination. "Take me somewhere fun," you might say. And we would not be too surprised if the car responded with something like: "Do you want someplace close to the water, or just a watered down version of the last place I chose for fun?"

We might not be surprised because, third, a renaissance of sorts is underway in the fields of machine learning and artificial intelligence ("AI"), both academic and commercial.[12] One can see this by following the venture capital and private sector funding dollars, where yearly disclosed investments of AI ventures have grown from $282 million in 2011 to $2.388 billion in 2015;[13] by reading about the AI research facilities Baidu, Samsung, Toyota, Google, Facebook, and

(Apr. 29, 2013, 4:29 PM), https://www.cnbc.com/id/100685958. Another striking example is the U.S. Army's chatbot, Sgt. Star, on GoArmy.com. Sgt. Star helps with recruiting efforts by answering questions online, and has answered 10.5 million questions in 2.8 million chat sessions, which equates to chatting with an average of 1,550 people each day. Dave Maass, *Answers and Questions About Military, Law Enforcement, and Intelligence Agency Chatbots*, ELECTRONIC FRONTIER FOUND. (Apr. 18, 2014), https://www.eff.org/deeplinks/2014/04/answers-questions-about-military-law-enforcement-and-intelligence-agency-chatbots.

[12] Precisely how to define "artificial intelligence" has been the subject of scholarly discussion for decades. In 1950, Alan Turing opened his influential paper on the subject with the words "I propose to consider the question, 'Can machines think?'" and suggested a test of machine intelligence — an "imitation game" — based on whether a human evaluator could distinguish genuine human natural language responses from those of a machine designed to produce similar responses. *See* Alan M. Turing, *Computing Machinery and Intelligence*, 59 MIND 433, 433 (1950), http://loebner.net/Prizef/TuringArticle.html. But in the more contemporary context relevant to the analysis here, it makes most sense to use the term "artificial intelligence" in a somewhat more colloquial sense, to mean the simulation of some aspect of human intelligence to undertake an activity or function that an informed layperson in the present would ordinarily consider human intelligence necessary or at least useful in performing a task. Instantiations of artificial intelligence may also involve a system or software program's relatively sophisticated capacity to communicate. Notice that this definition is not quite so demanding that it implies, say, the inability of a human to distinguish between communication with a human and with an application exhibiting some degree of artificial intelligence. And while certain activities such as real-time natural language-based communication likely belong within the scope of the concept long-term, precisely what else fits the definition will no doubt shift to some extent as broad public attitudes evolve regarding what functions or activities call for human intelligence.

[13] *Artificial Intelligence Explodes: New Deal Activity Record for AI Startups*, CB INSIGHTS (June 20, 2016), https://www.cbinsights.com/blog/artificial-intelligence-funding-trends.

others have built in San Mateo and Santa Clara counties;[14] and by talking to scholars who contrast our current times with the periodic "AI winters" where folks lost interest in the field (in the 1980s, for example).[15]

Traditional expert systems used law-like techniques to search through potential options when analyzing how to diagnose certain medical conditions, or how to categorize a certain kind of molecule — but they were cumbersome at best when it came to some of the seemingly simplest things that people could do almost "without thinking," like classifying visual objects, interpreting idiomatic expressions, or decoding nonverbal communication. As computing power gets cheaper and software improves, expert systems are ever more able to sift through millions of options quite quickly.

An even more important change is underway in the realm of so-called "machine learning." This branch of computer science involves the application of computing capacity and analytical techniques to enable computers to learn without being programmed explicitly. Such applications involve the use of statistical methods and neural networks to recognize patterns in large aggregations of data. In particular, the software architecture used in many successful "artificial intelligence" applications is increasingly relying on convolutional neural networks and reinforcement learning — architectures that extend the purview of AI beyond the stultified domain of experts to encompass those difficult-to-replicate abilities of young children to recognize a cat or a horse, or learn how to grasp an object.[16] The intricate pattern

[14] *See, e.g.*, *Advanced Applied Research*, SAMSUNG, http://www.sra.samsung.com/research/advanced-applied-research (last visited June 4, 2017); *Facebook AI Research (FAIR)*, FACEBOOK, https://research.fb.com/category/facebook-ai-research-fair (last visited June 4, 2017); *SAIL-Toyota Center for AI Research at Stanford*, http://aicenter.stanford.edu (last visited June 4, 2017); Sam Shead, *DeepMind's Army of AI Researchers in Mountain View Is Now over 20 People Strong*, BUS. INSIDER (May 31, 2017, 11:44 AM), http://www.businessinsider.com/deepminds-small-army-of-ai-researchers-in-mountain-view-is-growing-2017-5 (describing the growth at Google's subsidiary, DeepMind, at its headquarters in Mountain View, California); *Silicon Valley AI Lab*, BAIDU RES., http://research.baidu.com/silicon-valley-ai-lab (last visited June 4, 2017).

[15] *See* DANIEL CREVIER, AI: THE TUMULTUOUS HISTORY OF THE SEARCH FOR ARTIFICIAL INTELLIGENCE 203 (1993).

[16] Two machine learning techniques are generating particular interest. First are so called "deep learning" neural networks inspired by the layout of the human brain to spot patterns and leverage "big data." "Deep learning" systems embody a particular architecture for neural networks that avoid some persistent problems neural networks have had in developing adaptive responses to new data, and have sparked interest because of their capacity to solve pattern-recognition problems in computer vision

recognition made possible by these techniques comes at an analytical price, though, as in some cases it is far from clear, even to the designers of the systems, precisely how they have arrived at their conclusions.[17] These machine learning techniques are a great way to write a nearly security-proof, pesky computer virus (we will get back to that). It is through machine learning, too, that new progress is underway on many of those apparently simple but devilishly hard technical problems like vision and speech recognition.[18]

Fourth, in large measure because of the application of machine learning techniques leveraging large amounts of computing power to vast datasets, society will likely see even more change in the realm of user interfaces. Such innovation may involve changes to visual interfaces, as with augmented or virtual reality; further use of natural language and speech recognition like Siri, Alexa, or Cortana; or, eventually, new ways of integrating computers and the brain. The key development in this domain is not necessarily sophisticated data visualization, but rather the application of machine learning techniques to simulate intelligent conversation, for example, without some stark philosophical milestone involving conscious thought by machines.

But we are nonetheless approaching a milestone of sorts, as there is good reason to conclude that these four trends will drive stark, and potentially transformative, change in what society colloquially understands as artificial intelligence. It is not one factor, but all of them together, that can create a game-changer: huge amounts of information, leveraged in fluent, carefully intoned English with extraordinary visuals to make a persuasive appeal to people whose own norms increasingly accept persuasion by computer. Whether or when that becomes self-awareness is a different conversation, and one that — while interesting — may prove almost distracting.

---

and other fields. Second are reinforcement learning techniques or closely-related "genetic algorithms" that emerge by developing simple algorithms — computer programs capable of evolution — to solve a problem like spotting suspicious financial transactions, allowing those algorithms to mutate slightly over time, and then applying feedback to select for those algorithms that beat the others on a given metric. *See generally* Mariano-Florentino Cuéllar, *Cyberdelegation and the Administrative State*, *in* ADMINISTRATIVE LAW FROM THE INSIDE OUT: ESSAYS ON THEMES IN THE WORK OF JERRY L. MASHAW 134 (Nicholas R. Parrillo ed., 2017).

[17] *See generally* Joshua A. Kroll et al., *Accountable Algorithms*, 165 U. PA. L. REV. 633 (2017) (discussing the regulatory difficulties posed by advancements in machine learning).

[18] *See generally* Robert D. Hof, *Deep Learning*, MIT TECH. REV., https://www.technologyreview.com/s/513696/deep-learning (last visited June 4, 2017).

Instead, my goal in what follows is to advance a conversation about how legal doctrines and institutional arrangements may be affected by the resulting opportunities and challenges — ones so close on the horizon that we are all but fated to live with their consequences. And I want to explore what choices lawyers, judges, policymakers, and the public may need to make as we come to terms with these changes. The backdrop for this assessment is an expectation that law is at least one relatively reliable means through which society makes trade-offs and manages change. It is true enough that societies rely on other means to make such trade-offs as well — including changes in informal societal norms — and that in some settings, one can reliably expect a lack of consistency with or even outright disregard for legal principles. Nonetheless, in the United States and many other countries with advanced economies, policymakers and the public frequently express a desire for consistency between societal norms and practices, on the one hand, and legal rules and standards, on the other — whether embodied in constitutions, statutory provisions, regulatory principles, or case law. In such societies, law is about resolving disputes and about defining and implementing societal values. Which means people tend to expect a lot from the legal process and from law itself.

II. HOW ARTIFICIAL INTELLIGENCE COULD AFFECT LAW AND THE LEGAL SYSTEM

That societies in advanced democracies with developed economies have relatively high expectations of the legal system is evident from just a few examples. In California and elsewhere in the United States, if someone is a witness in a criminal trial but does not speak English, we want them to have an interpreter.[19] If a tenant who is about to be kicked out of her home wants to appeal, she should be able to do so.[20] When an agency decides to implement a rule that regulates certain molecules in pesticides or pharmaceutical products, there should be a justification that courts and other parts of government can follow. If a big new development subject to the California Environmental Quality Act ("CEQA") requirements allegedly merits a "negative declaration," there should be evidence to support that conclusion.[21] When an army or a police organization uses force, it should be justified. The devil may be in the details of government tort doctrines or section 1983 actions, but the premise is as simple as it is important.

[19] *See* 28 U.S.C. § 1827 (2012).

[20] *See* CAL. CIV. PROC. CODE § 1176 (2017).

[21] *See* 14 CAL. CODE REGS. tit. 14, § 15070 (2017).

Some kinds of software applications leveraging artificial intelligence — whether expert systems, genetic algorithms, or neural nets — can help address all the problems we currently face living up to these ideals. We do not have enough interpreters or lawyers to help poor people.[22] Identifying the right molecules to regulate is difficult. And deciding to use force is fraught and complex — just ask Chicago Mayor Rahm Emanuel.[23]

If society enhances the artificially intelligent tools available for addressing challenges of such enormous legal consequence, we will gain new opportunities to close the considerable gap between legal aspirations and reality that currently bedevils aspirations for justice. Convolutional neural networks and certain kinds of expert systems with natural language user interfaces grafted on can help with interpreting, facilitate legal advice,[24] and enhance the capacity of agencies to discern what they should regulate. So-called deep learning systems could analyze data from military or police body cameras in real time — indeed, perhaps even helping society make smarter decisions about when to use force.

These opportunities are important but they also raise quite profound questions. For starters, we do not have neat formulas to answer the question of what government should maximize, because we often lack consensus about how best to define individual rationality (especially across time), and aggregate social welfare. That is what we spend a lot of our time thinking about. Translating powerful, complex ideas into the language of algorithms and machine learning protocols is the mother of all statutory drafting and interpretation problems. I assure you this is more than simply a naked curricular power grab by an administrative law and legislation scholar.

---

[22] *See generally* Muneer I. Ahmad, *Interpreting Communities: Lawyering Across Language Difference*, 54 UCLA L. REV. 999 (2007) (calling for an integration of community interpreters in the lawyering process).

[23] *See generally* Mitch Smith & Timothy Williams, *Chicago Police Adopt New Limits on Use of Force*, N.Y. TIMES (May 17, 2017), https://www.nytimes.com/2017/05/17/us/chicago-police-force-shooting.html (describing the process behind Chicago's adoption of new restrictions on the use of force).

[24] For example, ROSS Intelligence, which was built on IBM's Watson cognitive computing platform, is a legal research tool that responds to questions an individual would ordinarily expect to ask a human, and for this reason, does not rely on keyword searches. ROSS was designed primarily to reduce the cost of legal research, and in its present and future iterations, its design and user interface implicate choices that can shape how a human understands legal concepts and their importance. Anthony Sills, *ROSS and Watson Tackle the Law*, IBM (Jan. 14, 2016), https://www.ibm.com/blogs/watson/2016/01/ross-and-watson-tackle-the-law. For more information about ROSS, see generally ROSS, http://www.rossintelligence.com.

Second, thoughtful observers of technological change could spend weeks or more talking about cybersecurity — for example, the risks to democracy, the transnational dimension, or the legal complexities associated with differing legal norms across jurisdictions. For present purposes, our discussion should highlight three challenges: (1) that the Internet is the backbone of the emerging world of machine learning; (2) that the Internet was designed to be adaptable and relatively resilient rather than secure; and (3) that societies appear to be rushing towards a world where Internet-connected computers have more control over infrastructure and human decisions. We are doing this while security is an unresolved problem, making people vulnerable to subtle manipulation, and organizations as well as countries vulnerable to attack.[25] In a world where election-related data breaches feature prominently in American politics, and where the Stuxnet computer virus covertly circled the world to wreck uranium enrichment centrifuges in Iran,[26] it borders on madness not to expect more security issues. And as illustrated by the clever architecture of the Stuxnet virus, the risk is not simply that an algorithm will fail to perform as promised, but that it will fail to do so even as the computer reiterates the promise that no problem is afoot.[27]

Third, heavy reliance on computer programs — particularly adaptive ones that modify themselves over time — may further complicate public deliberation about administrative decisions, because few if any observers will be entirely capable of understanding how a given decision was reached. Much of our law is about managing "principal-agent" problems or resolving disputes — whether between voters and politicians, between shareholders and managers, or between two people who cross-complain in tort. Yet neither in public nor private law have we designed a great legal framework to manage principal-agent problems or resolve disputes between people and very complex, self-modifying machines.

---

[25] In October 2016, for example, attackers reportedly relied on "hundreds of thousands of internet-connected devices like cameras, baby monitors and home routers" to conduct a cyber-attack against Dyn, a Domain Name System provider. Nicole Perlroth, *Hackers Used New Weapons to Disrupt Major Websites Across the U.S.*, N.Y. TIMES (Oct. 21, 2016), http://www.nytimes.com/2016/10/22/business/internet-problems-attack.html. As a result, people across the country temporarily had difficulty accessing a host of major websites, including Twitter, Netflix, Reddit, and *The New York Times*. *Id.*

[26] David Weissbrodt, *Cyber-Conflict, Cyber-Crime, and Cyber-Espionage*, 22 MINN. J. INT'L L. 347, 351-52 (2013).

[27] *See generally* KIM ZETTER, COUNTDOWN TO ZERO DAY: STUXNET AND THE LAUNCH OF THE WORLD'S FIRST DIGITAL WEAPON (2015).

Agencies have long relied on software to analyze data and inform assessments of the costs and benefits of particular actions. What is different is how the scope of reliance on software may be expanding to encompass new kinds of functions — such as policing decisions — and activities that had been previously treated primarily as the domain of human judgment.[28] The policing tool PredPol — which was reportedly used in almost 60 police departments across the country as of 2015[29] — uses machine learning to enhance predictions about the locations of future crimes, including locations other than those of past crimes.[30] In the military context, *The New York Times* reported that "the Pentagon has put artificial intelligence at the center of its strategy . . . spending billions of dollars to develop . . . autonomous and semiautonomous weapons . . . ."[31]

True: the human mind itself — what a neural network imperfectly imitates[32] — sometimes forges opaque and unreliable attachments to certain justifications for legal or governmental decisions. Indeed, the mind may come to accept as true wildly inaccurate reports of why a decision was reached. But the more interesting comparison is not between a neural network's conclusion and the faux certainty of a person's subjective account of what drove her decision. It is between reliance on an analytically opaque form of AI, such as a convolutional neural network, and the common language of justification and accountability that institutions force on an audience. The resulting

[28] *See generally* Kroll et al., *supra* note 17.

[29] Ellen Huet, *Server and Protect: Predictive Policing Firm PredPol Promises to Map Crime Before It Happens*, FORBES (Feb. 11, 2015, 6:00 AM), http://www.forbes.com/sites/ellenhuet/2015/02/11/predpol-predictive-policing/#2ec979fa407f.

[30] DAVID ROBINSON & LOGAN KOEPKE, STUCK IN A PATTERN 2-3 (2016), https://www.teamupturn.com/static/reports/2016/predictive-policing/files/Upturn_-_Stuck_In_a_Pattern_v.1.01.pdf.

[31] Matthew Rosenberg & John Markoff, *The Pentagon's 'Terminator Conundrum': Robots that Could Kill on Their Own*, N.Y. TIMES (Oct. 25, 2016), http://www.nytimes.com/2016/10/26/us/pentagon-artificial-intelligence-terminator.html?_r=0. Moreover, a 2009 U.S. Air Force report acknowledged the possibility that advances in technology could create autonomous machines, and stated that ethical and policy discussions are needed to figure out how to use that technology. The report states, "Authorizing a machine to make lethal combat decisions is contingent upon political and military leaders resolving legal and ethical questions." U.S. AIR FORCE, UNITED STATES AIR FORCE UNMANNED AIRCRAFT SYSTEMS FLIGHT PLAN 2009-2047, at 41 (May 18, 2009), https://fas.org/irp/program/collect/uas_2009.pdf; *see also* Peter Finn, *A Future for Drones: Automated Killing*, WASH. POST (Sept. 19, 2011), https://www.washingtonpost.com/national/national-security/a-future-for-drones-automated-killing/2011/09/15/gIQAVy9mgK_story.html.

[32] *See* Peter Margulies, *Surveillance by Algorithm: The NSA, Computerized Intelligence Collection, and Human Rights*, 68 FLA. L. REV. 1045, 1065-66 (2016).

conversation, however imperfect, allows citizens, organizations, and groups to engage in a common conversation about how to draw the civil-criminal liability line in a particular case or whether to regulate a molecule. The question is how that conversation changes or the scope of who can participate in it begins to constrict as reliance on the most opaque AI software — convolutional neural networks, for example, rather than expert systems capable of pinpointing the symbolic logic underlying their conclusions — gradually rises even as norms make reliance on their analyses more routine.

Fourth, existing statutes and legal doctrines may be poorly suited to the task of carefully policing the balance of decision-making responsibility between humans and computer systems in situations where — on reflection — society might prefer a human in the loop. Consider as an example how courts, at least when policing public sector decision-making, tend to rely on presumptions of regularity to avoid awkward questions about whether an official legally responsible for making a decision has actually made that decision rather than relying on advisers. Not only does litigation-related discovery have its costs, but it is also substantively difficult to draw a line between laudable reliance on other minds to supplement a decision-maker's judgment and inappropriately allowing someone else to decide. It is far from obvious that these problems would be easier to handle when the line must be drawn between the computer as a mere decision-support tool and the human as the implementer of machine-mediated decision.

And in many settings, such as those raising tort law questions, existing law may encourage automation without some careful weighing of aggregate risks or consequences.[33] We are then faced with an exceedingly blurry line between computer-assisted human choice and human-ratified computer choice. We can begin to see the complexity of this question by looking to older cases examining liability for both excessive reliance and insufficient reliance on computing systems. For example, in 1986 the Supreme Court of

[33] The Supreme Court's decision in *Kyllo v. United States* could provide one way of looking at new technologies. The Court determined that police use of a thermal imaging device to take a thermal image of the defendant's home (from outside the home) was an unreasonable search in violation of the Fourth Amendment, in part because of the "sense-enhancing technology" used by the police. 533 U.S. 27, 34 (2001). Moreover, in a concurring opinion in *United States v. Jones*, Justice Sotomayor acknowledged how technology may impact our understanding of a "reasonable expectation of privacy" (one possible component in the determination of whether an unreasonable search occurred), noting that "even short-term [GPS] monitoring . . . will require particular attention." 565 U.S. 400, 414-15 (2012).

Georgia found that it was a question for the jury "whether the plaintiff was negligent in relying solely upon its computer, considering the current widespread use of computers for the purpose of keeping business records, and that the plaintiff's computer . . . may not have been known to be inaccurate."[34] In contrast, in 1983 the Court of Appeals of Arizona concluded the jury was reasonable in finding the defendant trucking company liable for a wrongful death because the company could have quickly set up a computer program to verify driver log books and better regulate driver fatigue.[35] However we resolve these over- and under-reliance questions, make no mistake: a human relying entirely on a computer program — however carefully-designed or slickly marketed — to decide when to use force is analytically indistinguishable from an autonomous weapon.

Fifth, there is good reason to think that the labor and employment consequences of artificial intelligence — like those in the past associated with trade policy and immigration — will prove contentious and place a strain on public institutions already weakened by the depletion of public trust. Barring major changes in existing trends, automation is likely to further polarize our politics.[36] Unless something rather stark changes in our labor markets, our understanding of education and cognition, or the course of technological progress, the role of machines in our economy will become more contentious in the years to come. Think about the arguments over labor and law in the 1920s that were as contentious as anything we have seen in law and politics, or the contemporary arguments about trade. Political schisms over disappearing jobs will fuel fights over cybersecurity, human dignity, and basic income amidst the possibility that millions of drivers, middle-class white collar workers, and (as robotics takes off) security guards and medical care workers may lose their jobs in the decades to come.[37]

Finally, we confront around us the question of how organizations adapt over time to domains where human knowledge may erode — or mass cognition may change — because of delegation to AIs. The ability of people or organizations to persuade others through the use

[34] 49 AM. JUR. *Trials* § 2 n.31 (1994).

[35] *Id.* §§ 2 n.62, 6.

[36] *See* President Barack Obama, Remarks by the President in Farewell Address (Jan. 10, 2017), https://obamawhitehouse.archives.gov/the-press-office/2017/01/10/remarks-president-farewell-address.

[37] *See* Claire Cain Miller, *The Long-Term Jobs Killer Is Not China. It's Automation.*, N.Y. TIMES (Dec. 21, 2016), https://www.nytimes.com/2016/12/21/upshot/the-long-term-jobs-killer-is-not-china-its-automation.html.

of artificial intelligence will be heavily bolstered by the convergent capacity to leverage data about an individual or organization, knowledge of human cognition, and machine learning techniques. In a variety of settings ranging from policing to operational decisionmaking in the private sector, such persuasion will likely enhance a willingness to delegate decisions to an AI infrastructure itself dependent on leveraging past human knowledge acquired over time. Long term, such routine delegation will almost certainly disrupt the process through which humans ordinarily acquire institutional knowledge — through a mix of practical decisions, close observations, routinized conduct, and social interactions. That disruption would affect many organizations' ability to enhance performance over time, particularly if their reliance on AI were ever disrupted, and to reflect carefully on what is and is not working well. The sooner we realize that we are often already essentially working in organizations of mixed machine and human intelligence, the better we will be to think about the right uses for emerging AI innovations.

III. BUILDING INSTITUTIONAL CAPACITY AND ENHANCING PUBLIC KNOWLEDGE

What it means for so many of us to be working in these mixed human-machine arrangements already is that we are living through a moment rich in irony. On the one hand, the ubiquity of society's growing dependence on algorithms of various flavors — and the subtlety with which these systems are designed — makes it easier for many people to treat our interactions with persuasive technologies as familiar, and not to question how human decisions are methodically nudged in particular directions. On the other hand, it becomes simpler to imagine how AI-catalyzed changes in labor markets or relationships will affect so many livelihoods. Civil servants in policing, national security, or public health could have their judgment augmented or their expertise displaced — and the organizations where they work may risk losing their subtle organizational knowledge over time through reliance on AIs for making decisions. Doctors may reach more patients even as they become ever more dependent on AI,[38] and more vulnerable to cybersecurity breaches.[39]

[38] *See* Siddhartha Mukherjee, *A.I. Versus M.D.*, NEW YORKER (Apr. 3, 2017), http://www.newyorker.com/magazine/2017/04/03/ai-versus-md.

[39] On April 8, 2016, Senator Barbara Boxer (D-CA) sent a letter to the FBI about ransomware attacks on hospitals operated by MedStar Health in Washington, D.C., as well as attacks on the Alvarado Hospital Medical Center in San Diego and the

Meanwhile, consumers of media and services may face subtle manipulation by large players with market power using data about them to achieve persuasive optimization and shape their tastes long-term.[40] Some people will be considered suspicious by AIs because of their associations, speech, and/or past actions that seem to have no special connection to the matter being investigated.[41] As parents unwrap a snazzy new AI doll to interact with their child, their ability to shape their kids' values and learning may diminish (which of course assumes that we have at least a little sway with our kids!). Highly persuasive AIs — fed by copious amounts of data about their families and kids — may seamlessly weave together the doll, immersive video games, educational settings, and commercial/marketing activities.

As these dolls are unwrapped and given to eager children keen to share their secrets, society is left to ponder the fit between its needs, its laws, and its technologies. By asking the right questions, societies could find that they value organizations like police departments, schools, or drug companies that develop explicit plans to safeguard their organizational capacity from erosion, or that work to ensure they avoid relying too much on any one particular AI technology. Citizens and consumers may approach their own relationships — and society's — with AI differently if lasting changes occur in how society educates children from an early age to recognize their vulnerability to AI-fueled persuasion, and even to cybersecurity risks. Those citizens may be more keen to support laws and policies that force entities marketing AIs to more thoroughly and thoughtfully internalize cybersecurity risks. They may prefer that society should ask seriously — and often

Hollywood Presbyterian Medical Center. The letter indicates that the Hollywood hospital had to "divert critical 911 patients to other medical centers and to administer care without access to the important information contained in electronic medical records," before paying a ransom to regain access to the system. Letter from Barbara Boxer, U.S. Senator, to James Comey, Dir., Fed. Bureau of Investigations (Apr. 8, 2016), https://votesmart.org/public-statement/1085536/letter-to-the-honorable-james-comey-director-of-the-federal-bureau-of-investigation-information-on-growing-threat-of-ransomware-attacks#.WXv9ZdPytUQ.

[40] Indeed, many companies are already tracking and collecting data from Internet users through means such as third-party cookies, which may be used to develop a "detailed history of the types of sites you frequent" so as to "deliver ads tailored to your interests." *Online Tracking*, FED. TRADE COMM'N, https://www.consumer.ftc.gov/articles/0042-online-tracking (last visited Dec. 21, 2016).

[41] Chicago's "Strategic Subjects List" is a predictive policing system that uses a computer program to determine who is most likely to be involved in a shooting (either as the person committing the crime or the victim). The program is based in part on past offenses, possible gang affiliations, and past behavior. ROBINSON & KOEPKE, *supra* note 30, at 3.

— how much concentration of power is too much in the market for AIs. Whether achieved through common law doctrines of tort and contract law or carefully-tailored regulatory frameworks, greater transparency regarding a user interface's implicit objectives and how these are pursued through leveraging user data can help individuals understand how they are being swayed and affected by their interactions, and when particular representations or standards of reasonable care are being violated.[42]

These possibilities underscore how the problems and opportunities faced by an AI-embracing society are distinctive in some respects. Never have societies so massively centralized data and computing power, or used the results so persistently to predict or persuade. And it is likely that during the lives of our children or their children, societies may seek from our laws a means of deciding when certain AI systems might be best treated as having agency or interests worth protecting. That question may arise not necessarily because they meet some widely-accepted definition of being alive, but because reasonable people whose own interests are protected by law say of an AI: "this is my girlfriend" or "this is my child."

IV. LEARNING FROM SOCIETY'S EXPERIENCES WITH PUBLIC HEALTH AND THE FOOD ECONOMY

Distinctive though our future history will be, we can also learn from the past. That history teaches how societies that were increasingly urban and interconnected managed problems involving public health and food. Occasionally, advanced industrialized countries still face epidemics, food safety risks, and the threat of pandemics. But even in developing countries, public health practices have reduced disease and increased lifespans.[43] The world has benefited from improvements in basic hygiene such as hand-washing that have become commonplace, disclosure requirements and rules to reduce the risk of health problems from certain food additives and pathogens, and agencies capable of responding to infectious disease emergencies.[44] Although

[42] *See generally* Jon Hanson & Douglas Kysar, *Taking Behavioralism Seriously: The Problem of Market Manipulation*, 74 N.Y.U. L. REV. 630, 634 (1999) (seeking reevaluation in light of the impact of certain persuasive techniques of "any legal concept that relies in some sense on a notion of reasonableness or that is premised on the existence of a reasonable or rational decisionmaker").

[43] *See, e.g.*, Dr. Sheila Isanaka et al., *Efficacy of a Low-Cost, Heat-Stable Oral Rotavirus Vaccine in Niger*, 376 NEW ENG. J. MED. 1121 (2017) (reporting results of a new vaccine to help prevent severe gastroenteritis in young children).

[44] On improvements in basic hygiene, see ANGUS DEATON, THE GREAT ESCAPE:

far too many people still struggle to feed themselves and their families well into the twenty-first century, the number of people around the world who are under-nourished dropped by roughly 216 million between 1992 and 2015, to approximately 795 million.[45] As with urbanization, division of labor, advances in agricultural technology, and global interdependence, our reliance on machines and software leveraging artificial intelligence may continue to deliver something of a bounty for many aspects of human welfare.

Such reliance also forces us to think about problems analogous to what the world has confronted in public health — including what we might call the "digital" hygiene of disclosure governing an AI's subtle biases, regulation of security vulnerabilities that imperil others you will never meet, and agency capacity to protect the public when our reliance on machines has undesired or unexpected consequences.[46] It

HEALTH, WEALTH, AND THE ORIGINS OF INEQUALITY 98-100 (2013). Regarding the regulation of food additives and pathogens, see, e.g., Barbara Wilhelm et al., *The Effect of Hazard Analysis Critical Control Point Programs on Microbial Contamination of Carcasses in Abattoirs: A Systematic Review of Published Data*, 8 FOODBORNE PATHOGENS & DISEASE 949 (2011). For an example of the interpretive questions and societal trade-offs arising when strict limits are imposed on the use or concentration of carcinogenic additives, including those found in pesticides, see *Les v. Reilly*, 968 F.2d 985, 988-90 (9th Cir. 1992). The importance of techniques deployed in wealthy countries to manage infectious disease, including medical research infrastructures, disease monitoring, and high-capacity public health agencies capable of responding to infectious disease emergencies unlikely to be addressed effectively locally or through private action, see generally PATRICE BOURDELAIS, EPIDEMICS LAID LOW: A HISTORY OF WHAT HAPPENED IN RICH COUNTRIES (Bart K. Holland trans., Johns Hopkins Univ. Press 2006) (2003).

[45] FOOD & AGRIC. ORG. OF THE UNITED NATIONS, THE STATE OF FOOD SECURITY IN THE WORLD 8 (2015), http://www.fao.org/3/a-i4646e.pdf.

[46] One could, for example, conceivably compare immunization efforts with the future of autonomous vehicles. U.S. vaccination coverage is at record high levels and vaccines have saved numerous lives. The Centers for Disease Control and Prevention ("CDC") acknowledges, however, that advances in medicine come with increasingly complex infrastructure challenges, and that the successful provision of vaccines requires the cooperation of people and administrative systems. Similarly, a transition to autonomous vehicles could provide considerable benefits (e.g., road safety and reduced pollution) but will also require legal, administrative, and physical infrastructures to support their use — arrangements that will be likely experienced as significant constraints by individuals and communities. *See Achievements in Public Health, 1900–1999: Impact of Vaccines Universally Recommended for Children — United States, 1900–1998*, 48 MORBIDITY & MORTALITY WKLY. REP. 243, 247 (1999), https://www.cdc.gov/mmwr/PDF/wk/mm4812.pdf.

Moreover, efforts to reduce the infant and maternal mortality rates in the United States highlight the importance of understanding — and making more salient to the public — the decision-making power of sophisticated computer systems and how such power may affect different groups in different ways. The infant mortality rate has

is no surprise that marketing for tobacco or certain food products may downplay the harms and make it more difficult for individuals to switch away from such products later, shaping not only their information about the product but their capacity to resist its allure by — among other techniques — leveraging people's tendency to discount the future and neglect time-inconsistent preferences.[47]

Trade-offs arise not only for individuals across time, but for communities relying on vehicles and infrastructure that reduced walking. Consider, as an analogy, how the existence of motorized transportation improved human welfare even as its easy availability eventually introduced new dilemmas. To the extent societies valued both ease of transportation and other goals such as public health, they would need to take account of the risk that the built environment — by being so carefully calibrated to facilitate motorized transportation — implicitly discourages basic exercise through walking.[48] A similar dilemma may arise as intricate AI applications proliferate, allowing people to save time and gain capacity to achieve their goals when they analyze problems or interact with other people. Such applications might further substantially diminish the time and effort individuals expend in answering communications from others, deciding how to structure and prioritize tasks at work, and even how to woo someone into a romantic relationship. That tort law as well as conventional statutory and regulatory law will impact these applications should be taken as a given. The more fundamental question is to what end. Of course, part of the answer emerges when we consider how it is in no small measure by analyzing problems and interacting with others that we decide on our goals and the values underlying them in the first place.

As people rely increasingly on these systems, society encounters dilemmas similar to those involving motor transportation and urban design to preserve physical activity: even if such reliance falls well short of atrophying entirely a person's capacity for unassisted analysis

declined more than 90 percent in the past century, but there continue to be "significant disparities by race and ethnicity," such that black infants are more than twice as likely to pass away than white infants; it will be necessary to ensure that discrimination is not intentionally or unintentionally systematized through new technologies. *See Achievements in Public Health, 1900–1999: Healthier Mothers and Babies*, 48 MORBIDITY & MORTALITY WKLY. REP. 849, 849 (1999), https://www.cdc.gov/mmwr/PDF/wk/mm4838.pdf.

[47] *See, e.g.*, Kris N. Kirby & Richard J. Herrnstein, *Preference Reversals Due to Myopic Discounting of Delayed Rewards*, 6 PSYCHOL. SCI. 83, 84 (1995).

[48] *See, e.g.*, Jenna Panter et al., *Impact of New Transport Infrastructure on Walking, Cycling, and Physical Activity*, 50 AM. J. PREVENTIVE MED. e45 (2016), https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4712020.

or human interaction, his or her acuity in these areas may suffer even while (as with motorized transportation) such reliance appears justified in the short term. By carefully considering the cost of this second order effect arising from reliance on robust AI applications, societies can counterbalance the pressures to create legal, organizational, and technological infrastructures — such as legal and physical arrangements in cities strongly encouraging people in certain areas to use autonomous vehicles. Such infrastructures are akin to busy parkways designed only to speed vehicular traffic, with no sidewalks or other elements to catalyze physical activity. To the extent policymakers, lawyers, or the public are concerned about preserving cognitive or organizational capacity, perhaps they can learn from deliberate choices made in some regions to design infrastructure encouraging residents to exert themselves as a means of enhancing their health.

In California's distinctive agricultural heartland, the ongoing story of humanity's relationship with food as it impacts our health and culture is especially resonant. The practices and technologies that enabled the Green Revolution made it possible to feed a hungry planet once seen as on its way to famine. Access to food has also expanded because of the development of processed foods designed to last without spoiling. They combine long lists of multisyllabic ingredients and leverage precise scientific knowledge to combine salt, sugar, and fat with a perfect crunch that keeps people reaching for the next morsel. They are prepared in facilities that simultaneously reduce the risks of certain food safety problems — because part of the processing may reduce risks that adulterants such as E. coli O157:H7 will be present — even as these facilities also introduce new risks that a single security problem, whether intended or not, could affect large groups of people.[49] Whatever else is true of the AIs on which private sector decisionmakers, policymakers, and the public will increasingly rely, they tend to reflect in certain respects characteristics associated with the domain of highly processed food and processed information. Both are liable to be chock-full of intricate ingredients — whether the chemical compounds used to process foods or the specification of the neural networks relied on to drive outputs — designed for allure and

[49] *See* RENÉE JOHNSON, CONG. RESEARCH SERV., R43358, FOOD FRAUD AND "ECONOMICALLY MOTIVATED ADULTERATION" OF FOOD AND FOOD INGREDIENTS 1 (2014), https://fas.org/sgp/crs/misc/R43358.pdf (describing the risks of tampering associated with processed foods, and citing reports that melamine-contaminated baby formula had sickened 300,000 children in China).

conveying to consumers a measure of risk even as the food or information satisfies a craving or meets a short-term need.

Admittedly, the food economy is shaped not only by its distinctive politics but by a variety of factors rooted in culture and social norms.[50] Yet an analogy of sorts between food and information may shed light on some of the legal and policy dilemmas society may face in addressing aspects of artificial intelligence. The processing of food on a massive scale has created both opportunities and risks for society — with more people fed at a lower price, but new risks ranging from food safety problems at large processing facilities to heavy public consumption of tasty and convenient foods that nonetheless risk diminishing nutrition relative to the alternatives and in some cases diminish public understanding of the provenance of their nutrition. The neural networks at the heart of some of the most significant recent advances in artificial intelligence are, at core, means of processing information in a particular fashion to deliver a variety of benefits, from translation of languages to medical diagnostics.[51] The risks could encompass security breaches distorting the output of sophisticated software, as well as people's willingness to rely on the output of systems they do not understand especially well — a situation diminishing users' ability to understand the nature of the trade-offs involved in producing the information on which they relied.

The analogy has its limits. Even the scale and complexity of the food economy may not match the greater likelihood of emergent, unplanned properties as highly adaptive neural networks increasingly affect human behavior and engage in recursive interactions with not only people but with other such networks.[52] Along with other emerging machine learning techniques, these neural networks would likely lie at the heart of software performing functions in domains

[50] *See, e.g.*, Mariano-Florentino Cuéllar et al., *Institutions, Interests, and Incentives in American Food and Agriculture Policy*, *in* THE EVOLVING SPHERE OF FOOD SECURITY 87 (Rosamond L. Naylor ed., 2014).

[51] *See* Jeremy Hsu, *Can a Crowdsourced AI Medical Diagnosis App Outperform Your Doctor?*, SCI. AM. (Aug. 11, 2017), https://www.scientificamerican.com/article/can-a-crowdsourced-ai-medical-diagnosis-app-outperform-your-doctor; Catherine Matacic, *Google's New Translation Software is Powered by Brainlike Artificial Intelligence*, SCIENCE (Sept. 27, 2016), http://www.sciencemag.org/news/2016/09/google-s-new-translation-software-powered-brainlike-artificial-intelligence; Roger Parloff, *Why Deep Learning is Suddenly Changing Your Life*, FORTUNE (Sept. 28, 2016), http://fortune.com/ai-artificial-intelligence-deep-machine-learning.

[52] For a helpful overview of emergence underscoring its relevance in this context, see generally JOHN F. PADGETT & WALTER W. POWELL, THE EMERGENCE OF ORGANIZATIONS AND MARKETS (2012).

ranging from energy generation to education to public health. Nonetheless, these echoes between the food and information economies suggest that policymakers and the public may glean potentially interesting insights from understanding how certain legal arrangements structure the relationship between the public's interest in population health, private sector actors, and the food economy in advanced democracies. In many countries, statutory and regulatory rules require disclosure of basic nutrition information — and in the United States, the Food and Drug Administration ("FDA") recently finalized regulations requiring restaurants and similar food establishments to list calories associated with particular menu items.[53] Food safety measures, including risk analyses, audits of processing facilities, limitations on certain additives, and requirements imposed on third parties, help reduce the spread of foodborne illness and long-term health risks.[54] Tort liability forces participants in the food economy that can avoid societal costs by complying with reasonable standards of care to internalize those costs in their own actions.[55] Public sector purchases and public awareness campaigns may help the public better understand the relationship between nutrition and health, as well as the risks of ingredients that may seem tasty or convenient but pose longer-term health threats.

While there is no guarantee these analogies will fit all the legal dilemmas associated with artificial intelligence, our increasing reliance on sophisticated convolutional neural networks processing enormous amounts of data evoke some of the same questions and trade-offs society has encountered in managing its relationship to processed food — regarding food safety rules, tort liability, the role of public sector purchases, and public awareness efforts. Even if the analogy is not always a perfect fit, to contend that societies can handily manage those trade-offs without even reflecting on the history of our relationship to physical nutrition is a bold claim that casts aside the possibility of time-inconsistent preferences, troubling concentrations of political

[53] Food Labeling, 79 Fed. Reg. 71156 (Dec. 1, 2014) (to be codified at 21 C.F.R. pts. 11, 101).

[54] *See* RENÉE JOHNSON, CONG. RESEARCH SERV., RS22600, THE FEDERAL FOOD SAFETY SYSTEM: A PRIMER (2016), https://fas.org/sgp/crs/misc/RS22600.pdf. For a discussion of why it makes little sense simply to assume that these measures are the product of agency capture, see generally Mariano-Florentino Cuéllar, *Modeling Partial Agency Autonomy in Public Health Policymaking*, 15 THEORETICAL INQUIRIES L. 471 (2014).

[55] *See* Timothy D. Lytton & Lesley K. McAllister, *Oversight in Private Food Safety Auditing: Addressing Auditor Conflict of Interest*, 2014 WIS. L. REV. 289, 309 (2014).

and economic power, and other reasons why market failures even as these are conventionally understood may arise.

CONCLUSION

The fourth graders I described at the outset will inhabit a world that will benefit from being aware of, and pruning, the risks lurking in its midst even as it harvests the bounty emerging from its fertile soil. Their emerging legal world will prove in some respects as unfamiliar (at least to their parents) as it is familiar in other ways. They will confront a variety of new legal complexities and practical opportunities arising from converging developments involving computing power and data, evolving societal norms, intricate interdependence between machines and human organizational systems, and "deep learning" neural networks of enhanced, almost intuitive capacity.

Yet ironically, most of the reasons rendering these challenges both analytically interesting and practically significant are rooted in the continuing relevance to law of certain fundamental qualities of human existence and behavior. Among these are the frequent presence of time-inconsistent preferences, the extent of malleability in human values and goals, the difficulties living up to ideals of deliberation and self-government, and the possibility of emergent forms of organization amidst interconnected organizations and technological systems. As with many legal and policy questions involving food, agriculture, and public health — from nutrition information disclosure laws to urban design choices encouraging people to walk rather than drive — these qualities of human existence highlight why it is it less realistic or intellectually honest to expect that only lightly-regulated markets functioning against the backdrop of traditional common law doctrines will persistently deliver what the public most values over time.

Instead of being able to deploy a simple optimization algorithm, society will be left to use law as a means of both pruning risks and harvesting the fruits of this fertile age. Perhaps the most important step in both pruning and harvesting is to question whether any technology that promises a simpler world only makes the world *look* simpler. It is likely the trade-offs are pervasively happening just beyond what a pair of sleek new augmented reality glasses show. Or perhaps they are buried deep in the underbrush of fine print in a contract your AI assistant conveniently decided that you needn't read.

So much for big data making the world simpler.