COLUMBIA LAW REVIEW

VOL. 119 NOVEMBER 2019 NO. 7

KEYNOTE

A COMMON LAW FOR THE AGE OF ARTIFICIAL INTELLIGENCE: INCREMENTAL ADJUDICATION, INSTITUTIONS, AND RELATIONAL NON-ARBITRARINESS

*Mariano-Florentino Cuéllar**

INTRODUCTION

The majority of vehicles on California's vast network of roads make considerable use of information technology.[1] Although most are not yet capable of anything approaching fully autonomous driving, already it is possible to witness something like the following scene. A driver steering one vehicle spies a newer car's reflection in the rear-view mirror. The newer car appears to be driving itself. Whatever the official limits on that sleek vehicle's capability,[2] the person in its driver's seat seems to have no

* Justice, Supreme Court of California; Herman Phleger Visiting Professor and former Stanley Morrison Professor, Stanford Law School; affiliated faculty, Center for AI Safety and Freeman Spogli Institute for International Studies, Stanford University. I'm grateful for helpful feedback from Richard Rochman, invaluable conversations with Dario Amodei, Kate Crawford, Geoffrey Irving, Bob Kocher, Fei-Fei Li, and Meredith Whittaker, and superb research assistance from Natalie Heim and Derin McLeod. This Essay is adapted from my keynote address at the *Columbia Law Review*'s 2019 Symposium, "Common Law for the Age of AI."

1. To some extent, regulators have helped to drive the increasing importance of computing technology in the routine operation of automobiles. See Bill Canis, Cong. Research Serv., R44800, Issues with Federal Motor Vehicle Safety Standards 11–19 (2017), https://fas.org/sgp/crs/misc/R44800.pdf [https://perma.cc/TS69-SHYZ].

2. See generally David Welch & Elisabeth Behrmann, Who's Winning the Self-Driving Car Race?, Bloomberg (May 7, 2018), https://www.bloomberg.com/news/features/2018-05-07/who-s-winning-the-self-driving-car-race [https://perma.cc/HMC6-8LEN] (noting that the "road to autonomy is long and exceedingly complicated"); Tesla Autopilot—Review Including Full Self-Driving for 2019, AutoPilot Review, https://www.autopilotreview.com/tesla-autopilot-features-review [https://perma.cc/Z5CX-2CPS] (last visited July 29, 2019) (describing Tesla's self-driving capabilities). The extent to which a manufacturer appropriately represents to consumers or regulators the capacity of an autopilot function that falls short of full automation capability can raise plenty of legal issues—under contract law, tort law, and consumer protection statutes and regulations. See, e.g., Edvard Pettersson & Dana Hull, Tesla Sued over Fatal Crash Blamed on Autopilot Malfunction, Bloomberg (May 1, 2019), https://www.bloomberg.com/news/articles/2019-05-01/tesla-sued-over-fatal-crash-blamed-on-autopilot-navigation-error [https://perma.cc/VFX7-94UF]. Indeed, it's far from clear whether a concept such as "full" automation is even viable when functions that humans

interaction with the steering wheel when the driver of the older vehicle begins observing. Instead, the person in the driver's seat of that car is engaged in a mix of what seems like personal grooming, texting, and distracted glancing out the side window. Almost subconsciously, the driver of the older car realizes he is tweaking his own driving to test (within the limits of what's safe, of course) the way the algorithm appears to be driving the car behind him. If the older car slowed down or applied the brakes, the newer car behind would slow—gently if the front car decelerated slowly, and somewhat more suddenly if the driver of the older car applied the brakes more unexpectedly. Then the driver of the older vehicle realizes that if he stops for traffic and waits for the car in front to advance a bit before quickly accelerating, the autopiloted car stays behind and opens up a gap in traffic, tempting drivers in other lanes to switch into the opened-up spot. But if the driver of the older car speeds up more gradually, the newer vehicle stays close to the older car. So the older car's driver could effectively tighten the invisible coupling between his car and the more autonomous one or break it based on the rate of acceleration. Finally, when the lane next to the older car is clear, the driver realizes that a slight deviation in how centered his car is in the original lane achieves something significant—it seems to make the autopilot in the newer car behind disengage, forcing that driver to take over the steering wheel.

Even these few seconds of reciprocal steering and autopiloting on a California freeway tell a story: Simple choices can shape complex norms about how we rely on our machine infrastructure. More than simply emphasizing the importance of intricate algorithmic details affecting vehicular behavior, these stories also underscore how much humans are witnessing the steady integration of manufactured intelligence into everyday social life.[3] No doubt a human driver can feel like the Oscar Isaac character dancing with the robot in the film *Ex Machina*.[4] Sometimes this means that humans will be shaped in subtle but potentially enormously consequential ways by artificial intelligence (AI) techniques affecting the flow of information, the distance between cars, or the timing of persua-

---

colloquially bundle into a single category, such as driving, are easily disaggregated into distinct sub-functions that may call for different automation processes or degrees of human interaction, and when consumers routinely use available technologies in ways that fail to correspond to prescribed limits.

3. See Meredith Whittaker et al., AI Now Report 2018, at 10–11 (2018), https://ainowinstitute.org/AI_Now_2018_Report.pdf [https://perma.cc/JL95-7XKH] (describing the variety of settings where people routinely interact with systems displaying characteristics of artificial intelligence, and the broad range of functions performed); Ted Greenwald, What Exactly Is Artificial Intelligence, Anyway?, Wall St. J. (Apr. 30, 2018), https://www.wsj.com/articles/what-exactly-is-artificial-intelligence-anyway-1525053960 (on file with the *Columbia Law Review*) (same).

4. Ex Machina (Film4 & DNA Films 2014).

sive messages, for example.[5] Yet when we share the road, and indeed the world, with artificially intelligent systems, the direction of influence can also run in the opposite direction: Influencing the performance of an AI system need not be a very elaborate, high barrier-to-entry activity. The aforementioned driver's heavily analog, twentieth-century methods did fine in controlling, to some extent, a complex amalgam of software and hardware that is almost certainly also susceptible to—if surely somewhat tightly secured against—more sophisticated hacking.[6] Indeed, the co-evolution of human and artificial intelligence—what we could call our dance with machines—is well on its way to becoming routine. The dance continues as we navigate artificial chatbots, insurance transactions, court avatars, earnest advertising appeals, and borders.

Lurking in the background is law, along with the assumptions and norms it helps sustain. That this dance is playing out in the world's most economically complex and geopolitically powerful common law jurisdiction—the United States, still the preeminent hub for innovation in AI[7]—makes it appropriate to explore what relevance the common law and AI hold for each other. In fact, even accounts of American law that foreground the administrative state retain a prominent if not starring role for the system of incremental adjudication associated with American common law. Indeed, the roads, buildings, and corners of cyberspace where humans are increasingly interacting with manufactured intelligence also reveal another development of considerable importance for lawyers and judges: AI is becoming an increasingly relevant development for the American system of incremental, common law adjudication. The design of a vehicle with some capacity for autonomous driving can spur contract and tort disputes with qualities both familiar and novel.[8] Even decades

---

5. See, e.g., Robert M. Bond et al., A 61-Million-Person Experiment in Social Influence and Political Mobilization, 489 Nature 295 (2012) (finding that randomly assigned political mobilization Facebook messages influenced Facebook users' offline political activity).

6. Lying somewhere in between sophisticated cybersecurity intrusions and easily deployed human-driven techniques to control AI systems is the use of adversarial attacks to disrupt the expected operations of machine learning systems. See, e.g., Alexey Kurakin, Ian J. Goodfellow & Samy Bengio, Adversarial Machine Learning at Scale 1–2 (2017), https://arxiv.org/pdf/1611.01236.pdf [https://perma.cc/2XBM-UVD2] ("[N]eural networks and many other categories of machine learning models are highly vulnerable to attacks based on small modifications of the input to the model at test time . . . .").

7. See Sarah O'Meara, China's Ambitious Quest to Lead the World in AI by 2030, 572 Sci. Am. 427, 428 (2019) ("Most of the world's leading AI-enabled semiconductor chips are made by US companies such as Nvidia, Intel, Apple, Google and Advanced Micro Devices.").

8. See, e.g., Mark A. Geistfeld, A Roadmap for Autonomous Vehicles: State Tort Liability, Automobile Insurance, and Federal Safety Regulation, 105 Calif. L. Rev. 1611, 1632–74 (2017) (discussing manufacturer liability for autonomous vehicle crashes and hacks); Bryant Walker Smith, Automated Driving and Product Liability, 2017 Mich. St. L. Rev. 1, 32–56 (discussing products liability and personal injury litigation in the context of autonomous vehicles); see also Jack Boeglin, The Costs of Self-Driving Cars: Reconciling

ago, American courts were sometimes already facing legal questions foreshadowing dilemmas one can reasonably expect the present century to serve up about the balance of human and machine decisionmaking. A court in Arizona, for example, was forced to resolve whether punitive damages could be imposed on a transportation company, which failed to use information technology to track the work of its drivers and limit them from working excessive hours.[9] Just as courts once had to translate common law concepts like chattel trespass to cyberspace,[10] new legal disputes—turning on subtle distinctions revealed by digital evidence of neural-network evolution that bear on a party's responsibility for causing harm, for example—will proliferate as reliance on AI becomes more common. The reasonableness of a driver's decision to rely on a vehicle's autonomous capacity, or an organization's choice to delegate a complicated health or safety question to a neural network, will almost certainly spur a new crop of disputes in American courtrooms.

Given the speed and importance of these developments, my purpose here is to begin surveying the fertile terrain where the American system of common law adjudication intersects with AI. American society depends both on technology and the role of incremental common law adjudication in the legal system. The growing importance of AI gives us reason to consider how AI, common law, and society may affect each other. In particular, such an exploration should take account of the common law's role as a default backstop in social and economic life in the United States and a number of other major economies. Even beyond the strict doctrinal limits of torts, property, and contracts, common law ideas tend to set the terms for conversations among elites and even the larger public about the way social and economic interactions ordinarily occur, and how public agencies should analyze the problems—ranging from financial regulation to occupational safety—they are designed to mitigate.[11] Beyond serving as a default means of structuring interactions and a framework for analyzing social and economic life, the common law also offers an apt metaphor for how law, society, and technological change affect each other over the drawn-out process of applying broad social commitments to specific fact patterns. So it is no surprise that any

---

Freedom and Privacy with Tort Liability in Autonomous Vehicle Regulation, 17 Yale J.L. & Tech. 171, 174–75, 185–201 (2015) (noting the "uncertainty surrounding the complex liability issues for crashes involving [autonomous vehicles], which, in many ways defy the traditional conceptions of fault and agency at play in automobile accidents").

9. See Torres v. N. Am. Van Lines, Inc., 658 P.2d 835, 838–39 (Ariz. Ct. App. 1982).

10. See Intel Corp. v. Hamidi, 71 P.3d 296, 308 (Cal. 2003) (declining to find that emails from a former employee to numerous current employees criticizing the company's employment practices could, despite their unauthorized nature, constitute trespass to chattels).

11. See, e.g., Mariano-Florentino Cuéllar, Administrative War, 82 Geo. Wash. L. Rev. 1343, 1439 (2014) (discussing how ideological norms and the common law appeared to buttress each other and reinforced concerns about government ownership of industry during wartime mobilization in the early 1940s).

intellectually candid conversation about law and AI—particularly in the United States—must be to a considerable extent a conversation about the relationship between AI and the common law.

After defining some terms and setting the stage, I offer three preliminary ideas. First, our society already regulates AI through a backstop arising from the common law—and rightly so. Second, some degree of explainability that is well-calibrated to foster societal deliberation about consequential decisions is foundational to making any AI involved in human decisionmaking compatible with tort and other common law doctrines. At least one version of this ideal that merits attention could be termed "relational non-arbitrariness" to foreground the importance of buttressing—through both the common law and public law—society's capacity to deliberate about, and revise, the process through which it makes the choices that matter most. Finally, common law doctrines have room to integrate societal considerations involving organizational realities and institutional capacity, and concerns about matters such as the erosion of human knowledge that would be risky to ignore.

I. THE SCOPE OF A SHARED CONVERSATION ABOUT LAW AND AI: POINTS OF DEPARTURE

One can think of society as the aggregate of people who live together in a more or less ordered community. The very idea of society implies at least some degree of fairly constant change. Just as people across generations are defined by their evolving relationships to different groups or formal organizations,[12] those same people and the organizations with which they are affiliated are defining, through their behavior, their bonds with the increasingly adaptive technologies that surround them. For two reasons, close observation of the legal system and its common law component proves a revealing method to discern how some of that change happens. For one, the lawyers, clients, judges, and policymakers working through or within the legal system often play a part—sometimes a pivotal one—in the struggles over how society evolves.[13] But

12. See Max Weber, The Theory of Social and Economic Organization 118–20 (Talcott Parsons ed., A.M. Henderson & Talcott Parsons trans., 1947) (explaining the shifting nature of various kinds of social relationships); see also Raymond Geuss, History and Illusion in Politics 14–20 (2001) (explaining how a single "political association" can shift over time); Sheldon S. Wolin, Max Weber: Legitimation, Method, and the Politics of Theory, 9 Pol. Theory 401, 409–10 (1981) (explaining that Weber's definition of "culture" was concerned with social "meaning" and "patterns").

13. For examples of how key actors within the legal system affect policy outcomes through an alchemy of discretionary choices, legal interpretations, and strategies for reform of institutions and their legal authority, see Mariano-Florentino Cuéllar & Keith Humphreys, The Political Economy of the Opioid Epidemic, 37 Yale L. & Pol'y Rev. (forthcoming 2019) (manuscript at 48–53) (on file with the *Columbia Law Review*) (analyzing these dynamics in policymaking and litigation associated with opioid abuse); Mariano-Florentino Cuéllar, Refugee Security and the Organizational Logic of Legal Mandates, 37 Geo. J. Int'l L. 583, 587 (2006) (exploring this process in international law); Mariano-

whether any particular actors playing their part in the legal system fail or succeed at their goals, the legal opinions, legislative enactments, and administrative materials that memorialize much of the legal system's work also tend to reveal a story of society's public narrative of justifications and rationales for action: its compromises and aspirations.[14] That disputes about property interests, contract formation, products liability, and other aspects of the legal system can reveal so much is not only reason to take seriously the ideas and internal dynamics that define its work as a description of societal conflict and change; it's also a reason to retain humility about what deeper normative insights that may be persuasive across different segments of society—or even across cultures—can be gleaned from it. Bearing in mind that spirit of humility, we can at least observe in the American common law tradition and its related statutory or regulatory developments some insights about ideas that may be valuable amidst the transitions currently under way.

The enormous changes in the alchemy of algorithms and data, in social norms about computing, and in the resources available for technological development suggest two scenarios that may arise in the next few decades with respect to those transitions. One scenario takes as its point of departure the still-substantial limitations that bedevil many aspects of AI technology well into the twenty-first century—including in domains such as natural language processing and complex motor functions in robotics. Under this scenario, steady but gradual change occurs in AI as well as in the norms, institutions, and financial arrangements affecting its use. As we further leverage reinforcement learning and its variants, alongside conventional uses of supervised and unsupervised learning, the broad outlines of our world could remain much the same as they are now. Lawyers and their clients will continue navigating familiar disagreements about domestic and international politics. Whether they graduate from college or struggle to even finish high school, young people will navigate a labor market in which unemployment is a problem—but a manageable one in advanced countries—and relationships are primarily among humans. The autonomous vehicles in that scenario change only slowly from the one behind me in the recollection I shared; the rates at which its successors improve are limited by constraints of money, physics, lack of human imagination, and familiar global developments like recessions and climate change. AI is in that picture, but the difference compared to how it works now isn't categorical: It's still mostly a technology

---

Florentino Cuéllar, The Tenuous Relationship Between the Fight Against Money Laundering and the Disruption of Criminal Finance, 93 J. Crim. L. & Criminology 311, 403–04 (2003) (discussing this process in the criminal and national security context).

14. See Gerald E. Frug, The Ideology of Bureaucracy in American Law, 97 Harv. L. Rev. 1276, 1279–86 (1984) (analyzing corporate and administrative law "as a series of stories that assure us about the acceptability of bureaucratic organizations"); Duncan Kennedy, The Structure of Blackstone's *Commentaries*, 28 Buff. L. Rev. 205, 210, 214–16 (1979) (describing categories of legal reasoning as "social construction[s]").

of massive data analyzed using some kind of artificial neural network deploying enormous computing, where some tasks like automated translation and clustering get faster and become more ubiquitous, but—to channel Richard Haass—the world of the future is much like the world of the present.[15]

But in another scenario, the next one or two decades are quite discontinuous relative to our present. Here some mix of cheaper and greater computing power, innovation in designing algorithms, and our understanding of intelligence prove far more transformative over the next two decades or so. It becomes possible to imagine a world where some material subset of the population has deep emotional attachments to AI systems; where far more of the language we respond to or learn from is artificially generated; where some forms of friendship and work attachments are commodified through AI; where many major decisions about resources, entertainment, coercion, or innovation are routinely made with almost no human intervention; and where labor markets bear little resemblance to present ones. The difference between these scenarios is not the main subject of my talk, though it lurks in the background.

I suspect the distinction between these two scenarios turns heavily on several almost certainly interrelated questions: two technical and one social. The two technical questions that loom large are whether (1) enough progress occurs in natural language processing to simulate routine human communication of medium-to-high complexity (whether written or spoken); and (2) whether we scale the availability of reliable autonomous transportation. The social question is whether norms about the value of human decisionmaking, and the propriety of quite complex, emotionally meaningful communication and relationships with AI relative to humans, shift in favor of even more robust acceptance of AI-driven decisions and interaction.

That these scenarios are distinct in important ways should not obscure a crucial point of convergence. Whether drastic changes in employment or social norms about our relationship to machines occur in the next two decades or take longer, the legal system in general—and the common law in particular—will be a major focal point for certain pronounced societal dilemmas associated with AI. We can better understand those choices not only by recognizing the common law's role as a regulatory backstop but also by focusing attention on the centrality of reasoned deliberation across people and institutions, at least in the American legal tradition. And because the aspirations associated with the legal system inevitably run the risk of encountering detours and roadblocks, we must also acknowledge how much institutions matter—both as the targets for

15. See generally Richard N. Haass, Where to Go from Here: Rebooting American Foreign Policy, Foreign Aff., July/Aug. 2017, at 2, 9 (arguing that "the old challenges have not gone away," despite technological advances and increased globalization).

much of legal doctrine and also as the means through which the legal system implements much of its work.[16]

In this context, the term artificial intelligence means information technology that learns in some way, can perform some functions we consider to require intelligence if performed by humans, and provides functions or outputs that non-experts would consider trusting. Framed this way, the term encompasses both the domain-specific applications performing specific functions involving financial analysis or autonomous driving, for example, as well as systems aiming to simulate general intelligence through conversation or analytic capacity across domains. This description pivots, too, on the presence of a distinction between AI and conventional statistical inference—though obviously specific machine learning techniques at the heart of certain AI applications implicate both computer science and statistics.[17]

II. INCREMENTAL COMMON LAW ADJUDICATION AS THE DEFAULT FORM OF SOCIAL REGULATION

In a market economy with our historical tradition, the common law is the default framework for making sense of social and economic interaction.[18] Its conventions have informed, and indeed predated, the rise of the modern administrative state. The common law's influence is therefore powerful not only in its direct consequences for discrete transactions, such as the buying or selling of land, but in the ideas it's buttressed about who owes what to whom and for what reason—what duty of care, for example, two people owe each other, and thus what features of social life call for some judge-made or administrative remedy. Which means that in some sense, the pervasive common law backstop to social life provides a kind of first-draft regulatory framework—however imperfect—for managing new technologies ranging from aviation to email. Despite sometimes strong protestations to the contrary (especially from

16. See Mariano-Florentino Cuéllar, Beyond Weber: Law and Leadership in an Institutionally Fragile World, 69 Stan. L. Rev. 1781, 1785–87 (2017).

17. See Michael Jordan, Artificial Intelligence—The Revolution Hasn't Happened Yet, Medium (Apr. 19, 2018), https://medium.com/@mijordan3/artificial-intelligence-the-revolution-hasnt-happened-yet-5e1d5812e1e7 [https://perma.cc/56CM-MLPR] ("The developments which are now being called 'AI' arose mostly in the engineering fields associated with low-level pattern recognition and movement control, and in the field of statistics—the discipline focused on finding patterns in data and on making well-founded predictions, tests of hypotheses and decisions.").

18. See Grant Gilmore, The Ages of American Law 3–4 (2d ed. 2014); see also Francesco Parisi, The Efficiency of the Common Law Hypothesis, *in* The Encyclopedia of Public Choice 519, 519 (Charles K. Rowley & Friedrich Schneider eds., 2004); Rafael La Porta, Florencio Lopez-de-Silanes, Andrei Shleifer & Robert W. Vishn, Law and Finance, 106 J. Pol. Econ. 1113, 1115–17 (1998).

my home community in Silicon Valley[19]), society already "regulates" AI. That's true even in the absence of statutes and regulatory rules governing AI (though some of those exist, too, particularly in the autonomous vehicle context[20]). But the ultimate regulatory backstop here is the common law.

For thoughtful lawyers practicing in the American tradition, this may seem like a commonplace or even banal observation. But consider the implications: Through tort, property, contract, and related domains, society will shape how people use AI and will define what it means to abuse AI. We may draw on property law to address whether an annoyingly chatty AI system programmed to seek what it defines as autonomy, or a close affiliation with another person it claims has influenced its reasoning, can be "owned" and to what end. Contract law will be one setting in which we work on resolving whether the perfect, AI-spiked pitch for a bargain tailored and timed to overwhelm someone's judgment is unconscionable.[21] All this will indirectly shape norms and semiformal alternative dispute resolution systems, and it will play out even—indeed, especially—if statutes or regulatory rules meant to comprehensively regulate AI "sectors" remain on the shelf gathering dust. Sometimes the risk of disparate doctrinal decisions across many states or doctrines that cut against economic interests can prompt efforts to preempt the common law. But given the context of a relatively robust American federalism, it's not always easy to preempt state-level common law decisionmaking. Consider, for example, how even sweeping federal autonomous vehicle legislation carves out a robust domain (presumably on "laboratories of democracy" type grounds) for states to make common law decisions through the courts about autonomous vehicles.[22] We may also expect criminal law—that ever demanding stepchild of the common law and our statutory present—to deliver a share of dilemmas with interesting trade-offs.

---

19. See, e.g., Tom Simonite, Google Says It Wants Rules for the Use of AI—Kinda, Sorta, WIRED (Feb. 2, 2019), https://www.wired.com/story/google-says-wants-rules-ai-kinda-sorta/ [https://perma.cc/W2KZ-FTHW].

20. See Jack Karsten and Darrell West, The State of Self-Driving Car Laws Across the U.S., Brookings Inst.: Techtank (May 1, 2018), https://www.brookings.edu/blog/techtank/2018/05/01/the-state-of-self-driving-car-laws-across-the-u-s/ [https://perma.cc/P2TE-F48V].

21. See, e.g., Maxwell v. Fid. Fin. Servs., Inc., 907 P.2d 51, 53–54 (Ariz. 1995); Perdue v. Crocker Nat'l Bank, 702 P.2d 503, 512–13 (Cal. 1985).

22. See, e.g., David H. Coburn, Dane Jaques & Anthony J. LaRocca, Senate Commerce, Science, and Transportation Committee's AV START Act Advances, Steptoe (Oct. 11, 2017), https://www.steptoe.com/print/content/20768/Senate-Commerce-Science-and-Transportation-Committees-AV-START-Act-Advances.pdf?q= [https://perma.cc/LU8K-KC3V] (noting that "the AV START Act also preserves the existing rule that compliance with a federal safety standard does not exempt a person from common law liability under state law"); Aarian Marshall, Congress Races to Pass a Self-Driving Car Law by Year's End, WIRED (Dec. 5, 2018), https://www.wired.com/story/av-start-act-senate-congress-new-language-self-driving/ [https://perma.cc/88E9-YWHF] (noting that the bill clarifies "which level of government controls what part of self-driving car testing and operations").

It's worth reflecting on what society's frequent inability to recognize the common law as a form of "regulation" may suggest. Obviously important distinctions lurk in the difference between the common law and statutes or regulatory rules. And as Frank Pasquale's essay for this Symposium reminds us, interaction between regulatory rules and the common law can have a powerful effect on specific substantive domains ranging from bank finance to insurance to family relationships.[23] Yet we may consider what this blind spot about how our legal system operates says about how even informed laypeople ordinarily perceive the market economy: as having a nonregulatory default. This is not only deeply misleading; it misses the status quo and empowers certain institutions—in this case state courts, for example—that may be in some ways less and in other ways more prepared for the challenge than even they realize.

III. GIVING MEANING TO OPEN-ENDED COMMON LAW CONCEPTS BY REDUCING OPACITY: RELATIONAL NON-ARBITRARINESS

Whether the common law exerts an influence on AI primarily through courts or through common law-inspired norms shaping the public's understanding of social and economic relations, we can glean some insight into the evolving place of AI in society by training attention on some of the common law's most prominent doctrinal building blocks. Some of its most crucial components, such as the concept of reasonable care in tort law, presuppose an ability for a decisionmaker (such as a court) to observe sufficient contextual details to gauge objective reasonableness. As the use of AI becomes more common to design content and display information, and to inform, and in some cases essentially make, socially important decisions such as whom to hire and how to interact with the market, the capacity to assess context by understanding how AI systems make decisions in order to determine reasonableness becomes particularly important. Let me illustrate the point by delving more deeply into tort law.

Anglo American tort law relies on concepts of proximate causation, foreseeability, and duty, which together provide a more adaptive alternative to many forms of less flexible regulation meant to force parties to internalize the social costs of their actions and decisions.[24] Because the nature of the proximate causation and often-related foreseeability inquiries are flexible enough to take account of changing social, technological, and economic conditions, resolution of cases in this area often involves an interplay between application of well-settled principles and flexibility for the doctrine to adapt to the concerns raised by new situa-

---

23. See Frank Pasquale, Data Informed Duties in AI Development, 119 Colum. L. Rev. 1917, 1928–31 (2019) (describing the interaction between judge-made common law and expert-led agencies).

24. See, e.g., S. Cal. Gas Leak Cases, 441 P.3d 881, 887–88 (Cal. 2019).

tions. Evaluating these decisions will raise some intricate flexibility–fidelity trade-offs for our profession.

Those trade-offs will arise whether decisions involve routine matters such as driving, or more specialized activity such as detecting suspicious transactions. As a number of you have anticipated in papers and comments, one concern is that evaluating proximate causation requires a certain understanding (and ultimately, explainability) of the rationales for AI-driven decisions, without which it's difficult to complete in any defensible way the conventional doctrinal inquiry (because, at a minimum, it's not clear how justifiable it is for a person to rely on a particular kind of decisionmaking technology). So making concepts like reasonableness relevant in a world more reliant on AI systems will depend to some extent on the intricate design choices affecting how AI systems exchange information with humans, and how decisionmaking "justifications" are somehow extracted from artificial neural networks and similar systems. The challenge here is in some ways not so different from what happens when AI technology informs judgment about open-ended statutes like the Administrative Procedure Act or constitutional concepts like reasonable suspicion.[25] The most likely way to make an infrastructure of machine decisionmaking (or at least decision support) conform to a system with the aspirations of our own is to expect machine answers to conform to what I call "relational non-arbitrariness"—a concept not unrelated to what Ashley Deeks calls xAI,[26] and perhaps an example of Kate Strandburg's point about how human decisionmaking tends to be shared decisionmaking.[27]

Rooted to some extent in familiar concerns about shared deliberation[28] as well as decision costs,[29] relational non-arbitrariness calls for

---

25. See, e.g., Cary Coglianese & David Lehr, Regulating by Robot: Administrative Decision Making in the Machine-Learning Era, 105 Geo. L.J. 1147, 1170–76 (2017) (discussing "adjudicating by algorithm" and "rulemaking by robot"); Andrew Guthrie Ferguson, Big Data and Predictive Reasonable Suspicion, 163 U. Pa. L. Rev. 327, 351 (2015) (describing how big data will change policing and suspicion of criminal activity).

26. See Ashley S. Deeks, The Judicial Demand for Explainable Artificial Intelligence, 119 Colum. L. Rev. 1829, 1829–30 (2019); Ashley S. Deeks, Predicting Enemies, 104 Va. L. Rev. 1529, 1569 (2018); see also Matt Turek, Explainable Artificial Intelligence (XAI), Program Information, DARPA, https://www.darpa.mil/program/explainable-artificial-intelligence [https://perma.cc/N23J-KSWS] (discussing explainable AI and explainable machine learning as a concept) (last visited Aug. 26, 2019).

27. Katherine J. Strandburg, Rulemaking and Inscrutable Automated Decision Tools, 119 Colum. L. Rev. 1851, 1854–55 (2019).

28. See, e.g., John Dewey, The Public and Its Problems: An Essay in Political Inquiry 174–91 (Melvin L. Rogers ed., Ohio Univ. Press 2016) (1927) ("[Democracy] is the idea of community life itself"); Jurgen Habermas, Between Facts and Norms 126–28 (1996) ("The citizens themselves become those who deliberate and, acting as a constitutional assembly, decide how they must fashion the rights that give the disclosure principle legal shape as a principle of democracy."); Alexander Meiklejohn, Political Freedom: The Constitutional Powers of the People 24–28 (1965) (examining the freedom of speech and discussion in the context of "the traditional American town meeting").

the evaluation of how not only private organizations but public institutions make decisions. It takes seriously that decision costs should not routinely swamp the benefits of policies or legal rules, and depends on first considering whether some basis exists for a decision made by a human in close consultation with an AI system or by the system itself, in principle, such that we can defend the decision as non-arbitrary. Second, it calls for asking whether the relationship between the machine or analytical tool and the human conveys some of the complexity involved in the analysis and competing values at stake in the decision. And third, it calls for asking whether the process for making decisions supports further deliberation about the decision among some members of the community of people involved in or affected by it.

Admittedly, relational non-arbitrariness is a mouthful, but then so are terms like reasonable foreseeability and joint and several liability. The point of using a broader term is to think across fields like torts and civil procedure, regulation and constitutional law—and to ask whether an explanation is sufficient to let lawyers and informed laypeople engage in meaningful conversations about how a decision was made and for what reason. A touch of explanation will help you see how much it's already part of our law's fabric: It's taking seriously tort cases like *Biakanja v. Irving* that call for consideration of factors like prevention of future harm and moral blame to decide on the existence of a duty of care,[30] and seeking clarity about the assumptions underlying the arguments about these factors. It's considering whether the justifications offered by an agency arguing its conduct was supported by substantial evidence generalize to other contexts, or at least disclose the extent of the fit between an agency's argument and the relevant legal standard. It's expecting that a hearing to satisfy due process will involve enough transparency to know whether a decisionmaker has effectively delegated all power to an algorithm.[31]

A focus on reasoned explanation and public deliberation is perhaps especially prominent in public law. One can readily discern a judicial concern not only with reason-giving but also with how reason-giving permits deliberation and decisionmaking about accountability in the Supreme Court's recent majority opinion in *Department of Commerce v.*

---

29. See Mariano-Florentino Cuéllar, Auditing Executive Discretion, 82 Notre Dame L. Rev. 227, 251–52, 256–57 (2006).

30. 320 P.2d 16, 19 (Cal. 1958).

31. See State v. Loomis, 881 N.W.2d 749, 774–76 (Wis. 2016) (Abrahamson, J., concurring) (arguing that courts using risk assessment tools to inform sentencing decisions "must set forth on the record a meaningful process of reasoning addressing the relevance, strengths, and weaknesses" of the tool); see also Michael T. v. Crouch, No. 2:15-CV-09655, 2018 WL 1513295, at *3, *9–10 (S.D. W. Va. Mar. 26, 2018) (analyzing the use of algorithms to determine individualized budgets for state disability benefit recipients).

New York[32]—the case arising from the Commerce Department's decision to include a citizenship question in the nationwide decennial census:

> We are presented, in other words, with an explanation for agency action that is incongruent with what the record reveals about the agency's priorities and decisionmaking process. It is rare to review a record as extensive as the one before us when evaluating informal agency action—and it should be. But having done so for the sufficient reasons we have explained, we cannot ignore the disconnect between the decision made and the explanation given. Our review is deferential, but we are "not required to exhibit a naiveté from which ordinary citizens are free." The reasoned explanation requirement of administrative law, after all, is meant to ensure that agencies offer genuine justifications for important decisions, reasons that can be scrutinized by courts and the interested public. Accepting contrived reasons would defeat the purpose of the enterprise. If judicial review is to be more than an empty ritual, it must demand something better than the explanation offered for the action taken in this case.[33]

Perhaps in an even more pointed and explicit way than many run-of-the-mill administrative law opinions, this opinion emphasizes the need for further proof in a situation involving complex organizations with a variety of difficult-to-observe procedures and internal routines.[34] Yet in some respects, the situation is not dissimilar from what might arise when an AI system both seeks to reduce an underlying cost function while separately optimizing the likelihood that the information presented will entice the reviewing authority to find the relevant justification acceptable. Given the relevant agency problems and the understandable assumptions people make about their legal institutions in a society that values judicial independence and integrity, the goal must not be merely to generate justifications for public or private action that, on their face, are acceptable. The goal must instead extend to permitting review or at least some form of dialogue about reasoning and justification. That discussions of "reasonableness" arise in a different doctrinal context in tort law doesn't change at least one key aspect of the concept's function: to permit assessment of how a member of our civic community—whether

---

32. 139 S. Ct. 2551 (2019).

33. Id. at 2575–76 (quoting United States v. Stanchich, 550 F.2d 1294, 1300 (2d Cir. 1977)).

34. Cf. Banco Multiple Santa Cruz, S.A. v. Moreno, 888 F. Supp. 2d 356, 376–80 (E.D.N.Y. 2012) (denying summary judgment to an issuer of an annuity who honored a fraudulent withdrawal request because it failed to perform basic due diligence and ignored various factors that should have triggered greater inquiry). Compare DiLieto v. Cty. Obstetrics & Gynecology Grp., 998 A.2d 730, 751–52 (Conn. 2010) (allowing for evidence of a hospital's procedures and expert explanation of those procedures to establish standard of care), with Blankenship v. Collier, 302 S.W.3d 665, 670–72 (Ky. 2010) (holding that the procedures in a hospital's training manual were insufficient to establish standard of care).

relying on an AI system or not—justifies her actions relative to a more broadly applicable standard of conduct, and to permit reflection on how such a standard should be adjusted over time.[35]

Yet it's fair to ask whether courts are taking all that into account now, even beyond the common law. In *State v. Loomis*, the Wisconsin Supreme Court held that a trial court's use of an algorithm for risk assessment in sentencing didn't violate due process, even though the methodology used to produce the assessment remained undisclosed to either the defendant or the court.[36] The *Loomis* court insisted on a mild procedural safeguard instead: a "written advisement" accompanying presentencing reports.[37] Irrespective of how one weighs the court's understandable concerns about practical constraints and avoiding excessive discovery, it's far from clear that its holding promotes the kind of meaningful deliberation a reasonable observer would naturally associate with relational non-arbitrariness—about the design of the algorithm, the data used, or even the user interface. That this ideal of practically informed reason-giving, shared deliberation, and manageable decision costs is difficult to achieve even without AI in the picture should be obvious.[38] We should at least recognize a need to protect deliberation by making thoughtful use of our technologies to protect human-centered deliberation in the search for more sensible, less arbitrary choices about rules and statutes, constitutions, and the common law.

The tight bond between serious concern about deliberation and discussions of institutional structure underscore why it was far from a fluke—and instead the kind of pattern that even a simple, appropriately trained neural network could spot—that the inimitable Nobel laureate in economics, political scientist Herbert Simon, would move so naturally from studying organizations to becoming an AI pioneer.[39] There's probably good reason to think of at least functional formal organizations as a form of AI—arrangements that display a kind of intelligence yet work quite differently from individual human minds. Charles Stross develops

---

35. See Christopher H. Schroeder, Rights Against Risks, 86 Colum. L. Rev. 495, 551–58 (1986).

36. 881 N.W.2d at 753, 760.

37. Id. at 769.

38. See Jerry L. Mashaw, Organizing Adjudication: Reflections on the Prospect for Artisans in the Age of Robots, 39 UCLA L. Rev. 1055, 1064–65 (1992); Jerry L. Mashaw, Reasoned Administration: The European Union, the United States, and the Project of Democratic Governance, 76 Geo. Wash. L. Rev. 99, 117–20 (2007); Jerry L. Mashaw, Small Things Like Reasons Are Put in a Jar: Reason and Legitimacy in the Administrative State, 70 Fordham L. Rev. 17, 21–23 (2001); Robert Post, Managing Deliberation: The Quandary of Democratic Dialogue, 103 Ethics 654, 661–63 (1993).

39. See Hunter Crowther-Heyck, Herbert A. Simon: The Bounds of Reason in Modern America 275–90 (2005).

this point with respect to private corporations.[40] If Max Weber were here he'd probably agree with me that the point generalizes to the bureaucratic arrangements networking human brains to achieve somewhat common goals. Plainly, the turn-of-the-century social theorist Max Weber and pioneering cyberlaw scholar Jonathan Zittrain[41] (for example) speak a somewhat different language, and some of what a highly functional AI system can accomplish is different in speed and even substance relative to what a high-performing organization can accomplish. But I'm confident enough about this aspect of my argument to treat it as a rebuttable presumption that—at least for purposes of any conversation about ethics and governance—the similarities are more relevant than the differences. Maximizing any one goal, for example—whether it's shareholders' wealth or sharecroppers' health—can yield a better harvest.[42] AI and organizations can both serve to dilute responsibility, making it harder to know where action comes from and what justifies it. AI and organizations can deaden initiative, too, or spur creativity. And as with the organizational form, the more contemporary and future versions of AI will both help solve and more fundamentally continually raise core questions about governance that we tend to solve best with as much humility and awareness of competing values as we do with technical precision.

In short, as an ideal to guide our evaluation of reliance on AI for consequential decisions, I would have us emphasize not what an individual decisionmaker thinks, but rather buttress the often-implicit legal concern to focus on what networks of decisionmakers can reasonably consider. What is likely most consistent with the pronounced interest in reason-giving and justification in both the common law and public law traditions is to treat as pivotal the centrality of forms of justification that can be defended in human networks incorporating at least a material balance of principled, reasonable deliberations—networks designed to weigh whether certain reasons are sound enough to justify the use of coercive power, or the rejection of a presumed duty of care (for example) that members of a civic community owe each other.[43] There is little

---

40. Charlie Stross, Dude, You Broke the Future!, Charlie's Diary (Jan. 2, 2018), http://www.antipope.org/charlie/blog-static/2018/01/dude-you-broke-the-future.html [https://perma.cc/WC6A-594M].

41. See, e.g., About, Jonathan Zittrain, https://blogs.harvard.edu/jzwrites/about/ [https://perma.cc/SE3Z-YLU7] (last visited Sept. 2, 2019).

42. See, e.g., Nick Bostrom, Superintelligence: Paths, Dangers, Strategies 127–43 (2014) (describing AI and utility maximization); see also Daniel J. Phaneuf, Catherine L. Kling & Joseph A. Herriges, Estimation and Welfare Calculations in a Generalized Corner Solution Model with an Application to Recreation Demand, 82 Rev. Econ. & Stat. 83, 89–91 (2000) (providing an economic analysis of utility maximization).

43. See David Dyzenhaus & Michael Taggart, Reasoned Decisions and Legal Theory, *in* Common Law Theory 134, 145–50 (Douglas E. Edlin ed., 2007); Frederick Schauer, Giving Reasons, 47 Stan. L. Rev. 633, 656–59 (1995); cf. Judith N. Shklar, The Liberalism of Fear, *in* Liberalism and the Moral Life 21, 28–30, 36–38 (Nancy L. Rosenblum ed., 1989) (arguing that "[w]ithout the institutions of representative democracy and an acces-

doubt about the seriousness of agency and information economics problems implicit in heavy reliance on opaque AI systems—meant in principle to spot relationships eluding human judgment—for decisionmaking. The difficulty of choosing a principle (often among many viable ones) to turn the stuff of reinforcement-learning techniques or other uses for artificial neural networks into a neatly organized set of persuasive bullet points will likely spawn a second-order body of doctrine about aligning explanation and phenomenon explained. Surely one payoff of our shared conversation is to travel even a modest distance toward that destination. Crucial to that journey—down both the common law and statutory roads—is a recognition of the distinction between optimizing for perceptions of a decision's legitimacy from the audience, and instead seeking the right level of distrust and skepticism from the audience.

IV. PRESERVING THE COMMON LAW'S CAPACITY TO CONSIDER SOCIETAL IMPERATIVES ABOUT INSTITUTIONS AND ORGANIZATIONS

This discussion of tort law and relational non-arbitrariness also serves to introduce the third point—which is about the importance of integrating case-specific considerations with broader social imperatives. What's individually reasonable, such as reliance on automated building-design tools, may not scale in a benign way. As currently implemented, virtually all effective AI technologies depend to some extent on human collective action to produce data (for example, some interpreters, interpreting without an algorithm, essentially do the work for the algorithm). At least under the aforementioned scenario involving more stark change, excessive, organization-wide, or societal reliance on AI for entire classes of decisions may introduce systemic safety and security problems.[44] Neither consideration cuts decisively against allowing AI to play some role in consequential decisionmaking. But they do suggest that organizations and society may have reasons to seek an optimal degree of avoidance of the use of AI, to continue generation of unmediated data, and to hedge on safety and security concerns. No doubt the public, through market behavior and democratic responses (at least in some countries), will have some impact on user interfaces, natural language processing, deep learning architectures, and trade-offs about security. Yet there's no good reason to think these concerns can be safely ignored because of some kind of self-correcting market mechanism or reliable calibration process built into political democracy. At a minimum, analyses of

---

sible, fair, and independent judiciary open to appeals, and in the absence of a multiplicity of politically active groups, liberalism is in jeopardy").

44. See Jennifer M. Bernstein, Are We Literally Losing Our Way by Relying on GPS Devices?, Wash. Post (Dec. 2, 2018), https://www.washingtonpost.com/national/health-science/by-relying-on-gps-devices-are-we-literally-losing-our-way/2018/11/30/dd9eb6ae-e9bd-11e8-bbdb-72fdbf9d4fed_story.html (on file with the *Columbia Law Review*) (suggesting that increased dependence on navigational devices, like GPS, "have been linked to lower spatial cognition, poorer wayfinding skills and reduced environmental awareness").

how markets and political pressures affect the evolution of AI technologies must take account of the collective action and common pool problems playing a starring role in climate change perils;[45] intertemporal utility conflicts that complicate reasoned decisionmaking about addictive products;[46] transaction costs that complicate bargaining and coordination among consumers;[47] and path-dependent dynamics that can lock in certain practices, institutions, and products bearing little if any relationship to long-term social welfare.

Crucial to any sensible analysis of society-wide concerns about AI is recognizing the extent to which human endeavors occur within organizations. At least in countries with complex economies and societies, most of us do our work in organizations. As Charles Perrow's work has shown, the density and power of large organizations has grown massively over the last 150 years, especially in the United States, but also in other advanced industrialized countries.[48] Of course virtually everything we do occurs against the backdrop of institutions and institutional rules, but "organization" implies something more specific—a formal or semiformal entity with some internal rules or procedures and almost always, a common culture. Not surprisingly, problems of governance, compliance with law, and ethics are in some sense problems of organization and not just decisionmaking. We worry not only about the optimal use of force in the abstract, but about how police departments decide to use force and who's accountable for that. How the military ensures soldiers understand international law, and how it promotes unit cohesion. How our court system ensures that a family law judge facing two pro se litigants who don't speak English behaves when she realizes these litigants need an interpreter.

Many people already face these problems of compliance, policy, and ethics in mixed human–machine settings—in a dance with the machines.

---

45. See, e.g., Intergovernmental Panel on Climate Change, Global Warming of 1.5°C, at 71–72 (Valérie Masson-Delmotte et al. eds., 2018), https://www.ipcc.ch/site/assets/uploads/sites/2/2019/06/SR15_Full_Report_High_Res.pdf [https://perma.cc/6R7C-EXQV]; Scott Barrett, Collective Action to Avoid Catastrophe: When Countries Succeed, When They Fail, and Why, 7 Global Pol. 45, 46–50 (2016); Paul G. Harris, Collective Action on Climate Change, 47 Nat. Resources J. 195, 210–20 (2007); Manfred Milinski, Ralf D. Sommerfeld, Hans-Jürgen Krambeck, Floyd A. Reed & Jochem Marotzke, The Collective-Risk Social Dilemma and the Prevention of Simulated Dangerous Climate Change, 105 Proc. Nat'l Acad. Sci. 2291, 2292–94 (2008).

46. See Jon D. Hanson & Kyle D. Logue, The Costs of Cigarettes: The Economic Case for Ex Post Incentive-Based Regulation, 107 Yale L.J. 1163, 1193–1209 (1998).

47. See, e.g., Jeff John Roberts, Big Tech vs. Big Privacy Lawsuits, Fortune (Feb. 23, 2019), https://fortune.com/2019/02/23/big-tech-vs-big-privacy-lawsuits/ [https://perma.cc/DX49-QZVZ]; see also Elizabeth J. Cabraser & Samuel Issacharoff, The Participatory Class Action, 92 N.Y.U. L. Rev. 846, 852–60 (2017) (describing the effect of technology on class action participation and communication); Saul Levmore & Frank Fagan, The End of Bargaining in the Digital Age, 103 Cornell L. Rev. 1469, 1472–87 (2018) (describing the ways in which bargaining can be inefficient).

48. See Charles Perrow, Organizing America: Wealth, Power, and the Origins of Corporate Capitalism 19–21, 31–47 (2005).

The plaintiffs, defendants, witnesses, and jurors in our courtrooms work in hospitals or companies or agencies that rely on software and computer systems to assess the environment and support decisionmaking. Some of these aspects of our work are beneficial and some raise challenges. But it means innovations can scale quickly because people are somewhat used to interacting with machines to make decisions. And it means, as Max Weber would have appreciated,[49] that many and perhaps most of the beneficial uses of AI systems we can imagine—including the use of automated systems to support decisions in the justice system or in medicine—depend on particular assumptions about organizations. That they learn from their mistakes, for example, or provide minimally adequate cybersecurity. Some may argue that with the AI systems on the horizon we can just do away with many such organizations eventually, and surely our institutions will evolve in response to some of what you build.[50] But I'm skeptical we can do without them entirely, so I ask you to consider what organizational assumptions—about competence, adaptation, leadership, efficiency, or whatever—are built into any technically oriented definition of success we want to apply to a particular AI system or robotics technology.

The common law's relevance to our collective societal deliberations about the place of AI in a (still eminently) human-led world depends heavily on working into the analysis of doctrinal questions, such as the existence of a duty of care or proximate causation, these organizational realities and assumptions. Although assessing these institutional concerns raised by AI may sometimes justify some form of administrative regulation, properly interpreted and applied, tort law is at least one setting where judges and lawyers can take seriously the risks of eroding knowledge and other institutional concerns when performing the requisite social calculus necessary to resolve questions about reasonable foreseeability,[51] or the existence of a duty.[52] Both individual and society-wide safety benefits deserve an important place in the tort analysis, but so do offsetting considerations. Indeed, appropriately reasoned organizational decisions to prudently restrict reliance on some decisionmaking technologies incorporating AI that arguably represent the current norm of practice may be well justified under tort law. And of course, the use of

---

49. See Max Weber, The Profession and Vocation of Politics, *in* Weber: Political Writings 309, 313–15 (Peter Lassman & Ronald Speirs eds., 1994).

50. See Clay Shirky, Here Comes Everybody: The Power of Organizing Without Organizations 260–92 (2008); see also Ajay Agrawal, Joshua Gans & Avi Goldfarb, Prediction Machines: The Simple Economics of Artificial Intelligence 7–20 (2018); Mark Muro, Robert Maxim & Jacob Whiton, Metro. Policy Program, Brookings Inst., Automation and Artificial Intelligence: How Machines Are Affecting People and Places 29–46 (2019).

51. See, e.g., Restatement (Second) of Torts § 435 (Am. Law. Inst. 1965).

52. See, e.g., Regents of the Univ. of Cal. v. Superior Court, 413 P.3d 656, 663–74 (Cal. 2018) (finding that universities have a duty to protect students from reasonably foreseeable harms).

technologies that take explainability and legibility seriously and justify decisions in terms that can be shown to be consistent with a duty of care (subject to auditing that can confirm this) may merit some degree of recognition—perhaps through a rebuttable presumption that a cluster of related tort law responsibilities was taken seriously by the party whose behavior is in question.[53]

CONCLUSION

We should retain some humility in any conversation about the interdependent effects of law, AI, and society. History and the common law remind us that understanding change in law and society is a subtle enterprise, replete with episodes of profound consequence, such as the climate change spurred by generations of growing greenhouse gas emissions,[54] that were difficult to understand fully at an earlier point in the historical slipstream. But past experience also offers a reminder that certain patterns rhyme even if they don't recur precisely: As with technologies ranging from aviation to the internet, AI traces some of its roots not only to the industrial economy but to geopolitical competition. And because the common law has long affected both prevailing assumptions about who owes what to whom as well as society's day-to-day responses to emerging disputes and trends, it's a mistake to assume that AI is so exotic that the common law has nothing to contribute to its responsible development. Yet it's also worth acknowledging that the common law's relevance in this context persists in no small measure because its prevailing methodology is capacious enough to permit—in ways distinct but not wholly unrelated to what's possible in organs of the administrative state—sustained deliberation about society-wide consequences that should rightly inform how society assigns responsibility for the use of the evermore elaborate tools that human ingenuity has forged.

Without slipping into common law romanticism, it's fair to discern in the common law something far more interesting and consequential than a mere recipe for sensibly resolving disputes. In the courtroom arguments, judicial opinions, and public presumptions that define the common law one can also see a means of contending with different values and rationales—one relying on systems for argument using human networks, rather than by identifying a single value or goal to maximize.[55]

---

53. Cf. Wright v. Ford Motor Co., 508 F.3d 263, 268–74 (5th Cir. 2007) (applying Texas law, which in this context establishes a rebuttable presumption that a manufacturer is not liable for a design defect when it complied with applicable federal safety regulations). See generally Dan B. Dobbs, Paul T. Hayden & Ellen M. Bublick, The Law of Torts § 167 (2d ed. 2019) (explaining the use and effects of presumptions in tort law).

54. See, e.g., Intergovernmental Panel on Climate Change, supra note 45, at 53; 1 U.S. Glob. Change Research Program, Climate Science Special Report: Fourth National Climate Assessment 35–36, 38 fig.1.1 (D.J. Wuebbles et al. eds., 2017), https://science2017.globalchange.gov/downloads/CSSR2017_FullReport.pdf [https://perma.cc/3CBY-M7WU].

55. See Gilmore, supra note 18, at 4–10.

Indeed, the quest to build ethically aligned AI systems may go wrong if the noble intentions behind it turn into a presumption that we can realistically solve the most difficult ethical dilemmas by entrusting any single decisionmaker or ethical framework. And it is just as wrong to presume that some self-explanatory, easily defended concept of innovation or the market either provides straightforward guidance on this front or justifies preempting tort law while judges and lawyers calibrate their assumptions as they endeavor to steer a fast and reasonably safe course down the winding road that awaits and realize the need to change their assumptions about how decisions in society are made.

Let's review our progress on this brief leg of that remarkable road trip. First, we already regulate AI through the common law—and rightly so. We also regulate it through statutory and regulatory obligations on organizations, such as emerging standards governing autonomous vehicles, and we may yet do so through more AI-specific variants. None of this changes the fact that judges ruling on common law-type claims will likely play a quite central role in how our society governs AI, just as judges have at times set the default principles for how much we analogize cyberspace to physical space, or how far into the air the property rights go that are associated with a piece of land held in fee simple. Second, some degree of "explainability" is foundational to making any AI involved in substantially important human decisionmaking—about what products to design or sell, for example, or what promises to make or honor—compatible with tort and other common law doctrines. And third, common law doctrines have room to integrate societal considerations involving organizational realities and institutional capacity, security, and concerns about the erosion of human knowledge that would be risky to ignore. Reflect on that as your eyes go back for a few moments longer to the curving road while your vehicle dances with the (other) machines in your midst. I trust you've fastened your seat belt.