

2021

From Automation to Autonomy: Legal and Ethical Responsibility Gaps in Artificial Intelligence Innovation

David Nersessian
Babson College

Ruben Mancha
Babson College

Follow this and additional works at: <https://repository.law.umich.edu/mtlr>



Part of the [Science and Technology Law Commons](#), and the [Torts Commons](#)

Recommended Citation

David Nersessian & Ruben Mancha, *From Automation to Autonomy: Legal and Ethical Responsibility Gaps in Artificial Intelligence Innovation*, 27 MICH. TECH. L. REV. 55 (2020).

Available at: <https://repository.law.umich.edu/mtlr/vol27/iss1/3>

This Article is brought to you for free and open access by the Journals at University of Michigan Law School Scholarship Repository. It has been accepted for inclusion in Michigan Technology Law Review by an authorized editor of University of Michigan Law School Scholarship Repository. For more information, please contact mlaw.repository@umich.edu.

FROM AUTOMATION TO AUTONOMY: LEGAL AND ETHICAL RESPONSIBILITY GAPS IN ARTIFICIAL INTELLIGENCE INNOVATION

David Nersessian, JD, PhD and Ruben Mancha, PhD***

ABSTRACT

The increasing prominence of artificial intelligence (AI) systems in daily life and the evolving capacity of these systems to process data and act without human input raise important legal and ethical concerns. This article identifies three primary AI actors in the value chain (innovators, providers, and users) and three primary types of AI (automation, augmentation, and autonomy). It then considers responsibility in AI innovation from two perspectives: (i) strict liability claims arising out of the development, commercialization, and use of products with built-in AI capabilities (designated herein as “AI artifacts”); and (ii) an original research study on the ethical practices of developers and managers creating AI systems and AI artifacts.

The ethical perspective is important because, at the moment, the law is poised to fall behind technological reality—if it hasn’t already. Consideration of the liability issues in tandem with ethical perspectives yields a more nuanced assessment of the likely consequences and adverse impacts of AI innovation. Companies should consider both legal and ethical strategies thinking about their own liability and ways to limit it, as well as policymakers considering AI regulation ex ante.

TABLE OF CONTENTS

I. INTRODUCTION.....	56
II. TYPES OF ARTIFICIAL INTELLIGENCE	63
III. LEGAL RESPONSIBILITY FRAMEWORKS	66
A. <i>Products Liability Law and AI Artifacts</i>	70
B. <i>AI Artifacts and the Internet of Things</i>	75
IV. IMPLICATIONS FOR THE DEVELOPMENT AND DEPLOYMENT OF AI ARTIFACTS	79

* Associate Professor of Law, Babson College.

** Assistant Professor of Information Technology, Babson College.

A. <i>Liability Implications—Impacts on AI Innovation</i>	79
B. <i>Responsibility Implications—Studying the Ethical Perceptions of AI Professionals</i>	87
V. CONCLUDING THOUGHTS.....	91
TECHNICAL APPENDIX.....	94

I. INTRODUCTION

Artificial intelligence (AI) is a powerful technological driver of innovation and change. As noted by the founder and chairperson of the World Economic Forum:

Already, artificial intelligence is all around us, from self-driving cars and drones to virtual assistants and software that translate or invest. Impressive progress has been made in AI in recent years, driven by exponential increases in computing power and by the availability of vast amounts of data, from software used to discover new drugs to algorithms used to predict our cultural interests. Digital fabrication technologies, meanwhile, are interacting with the biological world on a daily basis. Engineers, designers, and architects are combining computational design, additive manufacturing, materials engineering, and synthetic biology to pioneer a symbiosis between microorganisms, our bodies, the products we consume, and even the buildings we inhabit.¹

AI systems decide what information we see on social media,² brake our vehicles when obstacles suddenly appear in our path,³ and move money around with little human intervention.⁴ Apple’s Siri makes us laugh,⁵ and

1. Klaus Schwab, *The Fourth Industrial Revolution: What It Means, How to Respond*, WORLD ECON. FORUM (Jan. 14, 2016), <https://www.weforum.org/agenda/2016/01/the-fourth-industrial-revolution-what-it-means-and-how-to-respond>.

2. See, e.g., Jared Schroeder, *Marketplace Theory in the Age of AI Communicators*, 17 FIRST AMENDMENT L. REV. 22, 34 (2018) (noting that, for example, “Facebook’s news feed algorithm decides which items, out of countless possibilities, appear atop users’ apps and browsers when they use the social media outlet.”).

3. See, e.g., Adam Thierer & Ryan Hagemann, *Removing Roadblocks to Intelligent Vehicles and Driverless Cars*, 5 WAKE FOREST J.L. & POL’Y 339, 344 (2015) (discussing U.S. Department of Transportation’s five-part scheme on vehicle automation, noting that “Level 1” automation, already in use today in autonomous vehicles, “include[d] electronic stability control or pre-charged brakes, where the vehicle automatically assists with braking to enable the driver to regain control of the vehicle or stop faster than by acting alone.”) (internal citations omitted).

4. See, e.g., Elizabeth Boison & Leo Tsao, *Money Moves: Following the Money Beyond the Banking System*, 67 DEP’T OF JUST. J. FED. L. & PRAC. 95, 111 (2019) (noting that “Facebook has partnered with both MoneyGram and Western Union to integrate ‘chatbots’ into its messenger service, facilitating the initiation of international and domestic transfers by Facebook users directly from Facebook’s interface.”). AI also powers high speed trading on Wall Street, with mixed results. See, e.g., Thomas Belcastro, *Getting on Board with Robots:*

Google helps us to remember our lives.⁶ AI also plays a growing role in the physical world, facilitating human-digital interactions through IoT (Internet of Things)⁷ interfaces⁸ and other kinds of robotic systems.⁹ Trends suggest that the breadth and depth of AI utilization and integration into all aspects of daily life will multiply rapidly as technology develops and costs continue to drop.¹⁰

AI already plays a critical role in the commercial sector. It is the driving force behind the Fourth Industrial Revolution, which “is characterized by a fusion of technologies that is blurring the lines between the physical, digital, and biological spheres.”¹¹ The economic impact of AI in this sector is estimated to be between \$3.5–\$5.8 trillion annually across multiple industries and business functions.¹² New efficiencies and innovation grounded in

How the Business Judgment Rule Should Apply to Artificial Intelligence Devices Serving as Members of a Corporate Board, 4 GEO. L. TECH. REV. 263, 272 (2019) (noting positive returns from AI trading systems as well as the role of AI in “flash crashes” of the market).

5. See e.g., Nick Bilton, *Siri, Tell Me a Joke. No, a Funny One.*, N.Y. TIMES (Aug. 12, 2015), <https://www.nytimes.com/2015/08/13/fashion/siri-tell-me-a-joke-no-a-funny-one.html>.

6. See e.g., Joelle Renstrom, *The Sinister Realities of Google’s Tear-Jerking Super Bowl Commercial*, SLATE (Feb. 3, 2020, 6:08 PM), <https://slate.com/technology/2020/02/google-assistant-super-bowl-commercial-loretta.html> (criticizing Google’s “Loretta” Feb. 2, 2020 Super Bowl commercial).

7. The U.S. Federal Trade Commission has defined IoT as an “interconnected environment where all manner of objects have a digital presence and the ability to communicate with other objects and people.” *Prepared Statement of the Federal Trade Commission on Discussion Draft of H.R. __, Data Security and Breach Notification Act of 2015: Hearing on H.R. 1770 Before the Subcomm. on Com., Mfg. & Trade of the H. Comm. on Energy and Com.*, 114th Cong. 6 (2015) (statement of Jessica Rich, Director of the Bureau of Consumer Protection, Federal Trade Commission). Other countries take a similar approach. See, e.g., OFF. OF PRIV. COMM’R OF CANADA, *THE INTERNET OF THINGS: AN INTRODUCTION TO PRIVACY ISSUES WITH A FOCUS ON THE RETAIL AND HOME ENVIRONMENTS 1* (2016), https://www.priv.gc.ca/media/1808/iot_201602_e.pdf (defining IoT as “the networking of physical objects connecting through the Internet.”).

8. See, e.g., Bruno Zeller, Leon Trakman & Robert Walters, *The Internet of Things – The Internet of Things or of Human Objects? Mechanizing the New Social Order*, 47 RUTGERS L. REV. 15, 34 (2020).

9. See, e.g., Ugo Pagallo, *What Robots Want: Autonomous Machines, Codes and New Frontiers of Legal Responsibility*, in HUMAN LAW AND COMPUTER LAW: COMPARATIVE PERSPECTIVES, 25 IUS GENTIUM: COMPARATIVE PERSPECTIVES ON LAW AND JUSTICE 47, 47–65 (Mireille Hildebrandt & Jeanne Gaakeer eds., 2013) (considering the intersection of robotics technology with legal principles and systems).

10. See, e.g., RAYMOND PERRAULT ET AL., *ARTIFICIAL INTELLIGENCE INDEX: 2019 ANNUAL REPORT* (2019), https://hai.stanford.edu/sites/default/files/ai_index_2019_report.pdf (tracking the pace of AI developments in areas such as technical progress, financial investment, research citations and higher education impacts).

11. Schwab, *supra* note 1. For comparative purposes, “[t]he First Industrial Revolution used water and steam power to mechanize production. The Second used electric power to create mass production. The Third used electronics and information technology to automate production.” *Id.*

12. See JACQUES BUGHIN ET AL., MCKINSEY GLOBAL INSTITUTE, *NOTES FROM THE AI FRONTIER: MODELING THE IMPACT OF AI ON THE WORLD ECONOMY 6* (2018),

widespread AI adoption are predicted to grow the global economy by \$15.7 trillion (14%) by 2030.¹³ Billions of dollars in market value of top companies such as Alphabet, Microsoft, Amazon, and Apple are linked to their use of AI in products and services.¹⁴

Despite its widespread economic benefits, the new AI “arms race”¹⁵ presents unique challenges for the legal system. The law always has struggled to keep pace with technological innovation and often finds itself behind the curve.¹⁶ With AI innovation in particular, the legal system must “embrace change and innovation as an imperative in a journey towards an ever-shifting horizon.”¹⁷ There have been many proposals from a wide range of stakeholders—including executive¹⁸ and legislative¹⁹ policymakers

<https://www.mckinsey.com/~/media/McKinsey/Featured%20Insights/Artificial%20Intelligence/Notes%20from%20the%20frontier%20Modeling%20the%20impact%20of%20AI%20on%20the%20world%20economy/MGI-Notes-from-the-AI-frontier-Modeling-the-impact-of-AI-on-the-world-economy-September-2018.pdf?shouldIndex=false>. These include the core business practices of: (i) customer acquisition, retention and service, as well as pricing and promotions (~\$1.2–2.3 trillion); (ii) operations and supply chain management (~\$1.2–1.9 trillion); and (iii) business optimization, risk management, and automation tasks (~\$1.3 trillion). *Id.*

13. PRICEWATERHOUSECOOPERS, *SIZING THE PRIZE: WHAT’S THE REAL VALUE OF AI FOR YOUR BUSINESS AND HOW CAN YOU CAPITALISE?* 3–4 (2017), <https://www.pwc.com/gx/en/issues/analytics/assets/pwc-ai-analysis-sizing-the-prize-report.pdf>.

14. *See Google Leads in the Race to Dominate Artificial Intelligence*, *ECONOMIST* (Dec. 7, 2017), <https://www.economist.com/business/2017/12/07/google-leads-in-the-race-to-dominate-artificial-intelligence>.

15. *See generally* Peter Asaro, *What Is an “Artificial Intelligence Arms Race” Anyway?*, 15 *I/S: J.L. & POL’Y FOR INFO. SOC’Y* 45 (2019) (discussing the “global artificial intelligence (AI) arms race”).

16. *E.g.*, Edward A. Parson, *Social Control of Technological Risks: The Dilemma of Knowledge and Control in Practice, and Ways to Surmount It*, 64 *UCLA L. REV.* 464, 471 (2016) (noting that “[r]egulation . . . often lags behind innovation”).

17. Daryl Lim, *AI & IP: Innovation & Creativity in an Age of Accelerated Change*, 52 *AKRON L. REV.* 813, 874 (2019).

18. *See, e.g.*, Exec. Order No. 13,859, 84 *Fed. Reg.* 3967 (Feb. 11, 2019). *See also* Request for Comments on a Draft Memorandum to the Heads of Executive Departments and Agencies, “Guidance for Regulation of Artificial Intelligence Applications,” 85 *Fed. Reg.* 1825 (Jan. 13, 2020) (requesting comments on draft policy statement relating to legal, regulatory and non-regulatory oversight of the development and use of AI applications outside of the Federal government).

For a recent analysis of current governmental use of AI, *see* DAVID FREEMAN ENGSTROM ET AL., *GOVERNMENT BY ALGORITHM: ARTIFICIAL INTELLIGENCE IN FEDERAL ADMINISTRATIVE AGENCIES* 6 (2020), <https://law.stanford.edu/education/only-at-sls/law-policy-lab/practicums-2018-2019/administering-by-algorithm-artificial-intelligence-in-the-regulatory-state/acus-report-for-administering-by-algorithm-artificial-intelligence-in-the-regulatory-state/#slsnav-report> (studying “AI use at the 142 most significant federal departments, agencies, and sub-agencies” and providing “cross-cutting analyses of the institutional, legal, and policy challenges raised by agency use of AI”).

19. *See, e.g.*, H.R. Res. 5356, 115th Cong. (2018) (proposing independent commission to consider artificial intelligence, machine learning, and related technologies from a national security standpoint); H.R. Res. 4625, 115th Cong. (2017) (proposing a coordinated national strategy for developing AI). For additional details on Congressional engagement with AI top-

in the United States and abroad,²⁰ state lawmakers,²¹ regulatory bodies,²² industry organizations,²³ and civil society groups²⁴—on how to address the unique legal implications presented by the creation and use of artificial intelligence. A wide body of scholarship is emerging around such varied digital innovation topics as virtual and augmented reality,²⁵ *sui generis* regulatory systems for AI implementation,²⁶ granting formal legal personality to AI systems,²⁷ and the role of AI in legal practice and the administration of justice.²⁸

Rather than add an incremental voice to an already-crowded chorus, this article takes an entirely different approach in order to arrive at a fresh perspective. It combines doctrinal analysis with an original research study, focusing on two complementary dimensions of AI innovation that have not

ics, see CONGRESSIONAL ARTIFICIAL INTELLIGENCE CAUCUS, <https://artificialintelligencecaucus-olson.house.gov/members> (last visited Feb. 15, 2020).

20. See, e.g., ACCESS NOW, MAPPING REGULATORY PROPOSALS FOR ARTIFICIAL INTELLIGENCE IN EUROPE (2018), https://www.accessnow.org/cms/assets/uploads/2018/11/mapping_regulatory_proposals_for_AI_in_EU.pdf.

21. States have taken different approaches to AI-powered autonomous vehicles, for example. See, e.g., Ben Husch & Anne Teigen, *Regulating Autonomous Vehicles*, NAT'L CONFERENCE OF STATE LEGISLATURES (Apr. 2017), <http://www.ncsl.org/research/transportation/regulating-autonomous-vehicles.aspx> (noting that 28 states had introduced legislation to regulate autonomous vehicles).

22. The use of AI in medical devices, for example, has drawn scrutiny from federal regulators. See, e.g., U.S. FOOD & DRUG ADMIN., *Artificial Intelligence and Machine Learning in Software as a Medical Device* (Jan. 28, 2020), <https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-and-machine-learning-software-medical-device> (implications of AI in medical devices).

23. Tom Simonite, *How Tech Companies Are Shaping the Rules Governing AI*, WIRED (May 16, 2019, 7:00 AM), <https://www.wired.com/story/how-tech-companies-shaping-rules-governing-ai> (discussing lobbying by industry group representing Microsoft, Facebook, and Apple around AI regulation). See also Yochai Benkler, *Don't Let Industry Write the Rules for AI*, 569 NATURE 161, 161 (2019), <https://www.nature.com/articles/d41586-019-01413-1> (noting that “[i]ndustry has mobilized to shape the science, morality and laws of artificial intelligence” and arguing that governments, consumers, and other stakeholders should take an active role in ensuring that industry concerns did not dominate the debate).

24. For consideration of the positive and negative human rights implications of AI, compare Salil Shetty, *Artificial Intelligence for Good*, AMNESTY INT'L (June 9, 2017), <https://www.amnesty.org/en/latest/news/2017/06/artificial-intelligence-for-good>, with Steven Melendez, *AI Could Bring “Nightmare Scenarios,” Warns Amnesty International*, FAST CO. (June 13, 2018), <https://www.fastcompany.com/40584711/ai-could-bring-nightmare-scenarios-warns-amnesty-international>.

25. See generally Mark A. Lemley & Eugene Volokh, *Law, Virtual Reality, and Augmented Reality*, 166 U. PA. L. REV. 1051 (2018).

26. See, e.g., Edmund Mokhtarian, *The Bot Legal Code: Developing a Legally Compliant Artificial Intelligence*, 21 VAND. J. ENT. & TECH. L. 145, 192 (2018).

27. See, e.g., Shawn Bayern, *The Implications of Modern Business-Entity Law for the Regulation of Autonomous Systems*, 19 STAN. TECH. L. REV. 93, 94 (2015).

28. See, e.g., Richard M. Re & Alicia Solow-Niederman, *Developing Artificially Intelligent Justice*, 22 STAN. TECH. L. REV. 242, 246 (2019); Harry Surden, *Artificial Intelligence and Law: An Overview*, 35 GA. ST. U. L. REV. 1305, 1326–35 (2019).

yet been considered vis-à-vis each other: the legal liability arising out of AI innovation and deployment and the ethical development of AI in the first place. It frames questions of products liability from a new perspective in two ways. First, it focuses on three key actors involved in bringing products with built-in AI capabilities (herein designated “AI artifacts”) to market: innovators, providers, and users. Second, it considers the relative sophistication of the technologies in question, noting important differences in AI automation, augmentation, and autonomy.

The reframing of liability is only one part of the picture. In business environments, principles of ethical responsibility work in conjunction with liability avoidance as powerful forces in shaping both corporate and individual behavior.²⁹ This article discusses findings from a new study on the ethical perspectives of AI developers and managers in the creation and deployment of AI technologies. Consideration of both perspectives yields a more nuanced assessment of the likely consequences and impacts of AI innovation, both for companies thinking about how to limit their own potential liability and for policymakers considering AI regulation *ex ante*.

Proposals for the regulation of AI tend to be highly context-dependent, varying considerably depending on what objectives the regulatory scheme is intended to achieve. Regulating AI for the purpose of personal safety,³⁰ for example, is very different from schemes aimed at protecting AI-generated creative works through copyright,³¹ which in turn differ markedly from broad regimes regulating AI as an aspect of wider social policies.³² This article focuses on questions of civil liability in tort for physical injury or

29. See, e.g., Milton C. Regan, Jr., *Risky Business*, 94 GEO. L.J. 1957, 1966 (2006) (“[C]onceptualizing ethics as a matter of avoiding liability can influence these dispositions, attitudes, and motives, and, therefore, how someone exercises her discretion Risk management conceives of ethical and legal provisions as a minefield of potential sources of liability.”).

30. See Matthew U. Scherer, *Regulating Artificial Intelligence Systems: Risks, Challenges, Competencies, and Strategies*, 29 HARV. J.L. & TECH. 353, 359–60 (2016) (discussing the challenges of defining AI in the context of a proposed regulatory regime to govern its use); see also *id.* at 388–98 (discussing tort liability and proposing a safety certification system for AI).

31. See Shlomit Yanisky-Ravid, *Generating Rembrandt: Artificial Intelligence, Copyright, and Accountability in the 3A Era—The Human-Like Authors Are Already Here—A New Model*, 2017 MICH. ST. L. REV. 659, 673 (2017) (“There are as many definitions as there are types of AI systems. John McCarthy, who coined the term ‘Artificial Intelligence,’ did not provide an independent definition, while scholars Stuart Russell and Peter Norvig suggested almost ten different definitions. Definitions generally vary according to the targeted subject, emphasizing different aspects of AI systems”); see also Kalin Hristov, *Artificial Intelligence and the Copyright Dilemma*, 57 IDEA 431, 431 (arguing that programmers should receive authorship rights for AI creations).

32. See, e.g., European Parliament Resolution of 16 February 2017 with Recommendations to the Commission on Civil Law Rules on Robotics, EUR. PARL. DOC. P8-TA(2017)0051 (2017) ¶¶ 49–65 (noting prospective need for liability regimes to cover next generation AI and robotic systems).

property damage caused by AI artifacts. Questions of criminal liability are not addressed here, though it certainly is possible for malicious innovators to use AI artifacts to violate criminal laws (e.g., using robots to rob a bank, drones to assault an enemy, advanced “peeping tom” technologies to violate privacy, etc.).³³

Another working assumption here is that the AI artifacts in question are not *intentionally* designed to cause injury to human beings. In such cases, AI-enabled devices that harm someone actually are not defective—they are operating as intended. Thus, our analysis excludes AI artifacts such as so-called “killer robots” designed for military purposes, even when those products make errors that end up killing the wrong army’s soldiers or civilian noncombatants.³⁴ Our discussion also excludes less direct forms of AI-enabled “injury,” such as the use of COMPAS, an evidence-based risk management system that predicts recidivism from a defendant’s interview and criminal file, the results of which are used by judges in deciding appropriate sentences of incarceration.³⁵

The article is organized in five parts. Part I sets the scene by describing the role of AI in society today, as well as the role that it increasingly will play in the future as technology advances at a rapid pace. Part II offers a more nuanced approach to thinking about AI systems by positing three primary classifications of AI artifacts according to their functioning:

- (i) *Automation AI*: Characterized by known pathways and defined characteristics, replacing known and repetitive human activities (e.g. sales chatbots, or repetitive tasks in manufacturing).
- (ii) *Augmentation AI*: Designs based on known interactions with human operators—helping workers to recall and analyze data but leaving judgment and strategizing to necessary human counterparts (e.g. surgical robots, or the augmented reality game Pokémon Go).
- (iii) *Autonomy AI*: Machine learning based on unknown interactions and environments, where the machine itself makes important, high stakes decisions—only primitive forms currently exist (e.g., today’s “self-driving” vehicles), but autonomy will be the

33. See, e.g., Rachel Charney, *Can Androids Plead Automatism? – A Review of When Robots Kill: Artificial Intelligence under the Criminal Law* by Gabriel Hallevy, 73 U. TORONTO FAC. L. REV. 69 (2015); see also Gabriel Hallevy, *The Basic Models of Criminal Liability of AI Systems and Outer Circles* (June 11, 2019), <https://ssrn.com/abstract=3402527>.

34. See, e.g., Bonnie Docherty, *Heed the Call: A Moral and Legal Imperative to Ban Killer Robots*, HUM. RIGHTS WATCH (Aug. 21, 2018), <https://www.hrw.org/report/2018/08/21/heed-call/moral-and-legal-imperative-ban-killer-robots>.

35. See, e.g., *State v. Loomis*, 881 N.W.2d 749 (Wis. 2016) (holding that a circuit court’s consideration of a COMPAS risk assessment at sentencing does not violate a defendant’s rights to due process).

inevitable result of AI increasingly gaining the ability to deal with unstructured data and complex settings.

Part III considers strict products liability law as applied to AI artifacts. When a product causes injury, normally all distribution chain participants can be held liable—including creators, designers, manufacturers, suppliers of component parts or designs, distributors, retailers, owners, and operators. For greater clarity, this writing simplifies the liability construct to focus more precisely on the three primary types of actors in the distribution chain for AI artifacts:

- (i) *Innovators*: Creators of either custom-designed or open-source AI algorithms and systems incorporated into AI artifacts. Innovators may be nonprofit actors (e.g., research universities) or individuals as well as corporate entities. The AI itself may be brand new or a refinement of existing AI (e.g., modifying code from an open-source AI library).
- (ii) *Providers*: Manufacturers who assemble various AI technologies and other components into products that may either be tactile (e.g., self-driving car) or virtual (e.g., high-frequency trading program), and the distributors and retailers who help to bring those products to the end-user market.
- (iii) *Users*: Owners who purchase AI-enabled products and those who operate them on behalf of the purchasers, which can be individual consumers with personal uses or business entities in a business-to-business setting.

Part III then discusses special liability issues associated with IoT devices, which incorporate AI artifacts into a variety of household and consumer products. This analysis reveals that the liability risks associated with certain types of AI for certain AI actors are predictable and understandable, while others are inherently unforeseeable and unknown, and thus are uninsurable. This has significant implications for the current and future development of AI artifacts, as discussed in Part IV.

Part IV considers how the absence of clear liability standards (a legal responsibility gap) might impact the development of AI in the private sector. It notes that the various actors in the AI value chain will seek to allocate liability and indemnity risks among themselves through purchasing contracts or intellectual property license regimes. This may mean that only the largest companies will be able to take on the potentially unlimited liability risk at issue. This in turn may stifle AI entrepreneurship at smaller scales or channel innovation to the types of AI artifacts most suitable to the needs of large corporate concerns.

Part IV also considers the role played by ethics in the use and development of AI. It discusses a new research study on the perceived responsibility of AI professionals in the development and use of AI innovations. This

study reveals a clear gap in ethical responsibility—namely, that although nearly everyone agrees that ethics and responsibility are critically important in AI innovation, nearly everyone also agrees that ethics and responsibility are someone else’s job. Part IV then discusses the implications of the legal and ethical responsibility gaps in AI innovation and offers recommendations to address them. Part V concludes with final thoughts and topics for future consideration.

II. TYPES OF ARTIFICIAL INTELLIGENCE

As discussed herein, AI refers to computational algorithms based on statistical and logical principles emulating human cognitive processes to engage in data-driven tasks: acquiring and processing data, actuating physical components, learning (including learning how to learn), and solving problems.³⁶ Put another way:

[T]o be considered [AI] . . . a computer system or robot must meet certain benchmarks: it must (1) communicate using natural language, (2) store information, (3) engage in automated reasoning (i.e., logic) to evaluate stored information to answer inquiries, (4) adapt to new situations and extrapolate patterns, (5) contain computer vision, and (6) include robotics functions.³⁷

AI creates value by emulating human cognitive processes, mainly those rooted in quantitative analysis and logic in narrow settings. All other cognitive processes, particularly those rooted in gut feeling, intuition, and emotion, are well beyond the capabilities of the most sophisticated AI systems. Common innovations in AI technology thus broadly fall into the categories of computer vision, virtual agents, natural language processing, autonomous vehicles, or smart robotics.³⁸ The most complex AI systems frequently incorporate multiple types of AI innovations and combine them with other technologies into comprehensive solutions. However, for analytic purposes here, we posit three primary classifications of AI according to its function: automation, augmentation, and autonomy.

36. See Scherer, *supra* note 30, at 360 (noting eight different functional definitions of AI “organized into four categories: thinking humanly, acting humanly, thinking rationally, and acting rationally. Over time, the importance of each of these definitional concepts has waxed and waned within the AI research community.”).

37. Nancy B. Talley, *Imagining the Use of Intelligent Agents and Artificial Intelligence in Academic Law Libraries*, 108 L. LIBR. J. 383, 387 (2016) (citing STUART J. RUSSELL & PETER NORVIG, *ARTIFICIAL INTELLIGENCE: A MODERN APPROACH* 2–3 (2d ed. 2003)).

38. See *Artificial Intelligence News*, BUS. INSIDER, <https://www.businessinsider.com/artificial-intelligence> (last visited Feb. 17, 2020) (discussing various developments and trends in AI).

Automation AI has found a home in production lines and warehouses, where many tasks are repetitive and fully identified in advance. It is also finding its way into white-collar settings, where its information processing capabilities help knowledge workers in decision-making processes.³⁹ As AI learns to deal with complexity and uncertainty, it becomes increasingly helpful in the completion of ambiguous tasks, which includes dealing with humans in unstructured contexts. Sales chatbots are an example of automation, as they replace known and repetitive activities traditionally completed by salespeople using a text-based interface.⁴⁰ Factory assembly and warehouse storage and retrieval robots are additional examples.⁴¹

Augmentation AI helps workers to recall and analyze data or perform precision tasks, leaving judgment and strategizing to a human counterpart.⁴² Workers in industrial settings, for example, find information aids on augmented reality helmets in support of their decisions.⁴³ So do players of the augmented reality game Pokémon Go. Often, augmentation AI is far more sophisticated and high stakes, as for example when surgery is assisted by the Da Vinci surgical robot.⁴⁴ Nevertheless, the surgeon remains squarely in control, albeit with the physician's capabilities augmented (enhanced) by the AI artifact.⁴⁵

Autonomy AI will be the inevitable result of AI gaining the ability to deal with unstructured data and complex environments. Emerging examples include delivery robots and "Level 4" (i.e., fully autonomous) self-driving vehicles,⁴⁶ although autonomous AI today remains far from the hypothetical

39. See, e.g., Arup Das, *There's a Bot for That*, A.B.A.: L. PRACT. TODAY (Dec. 14, 2018), <https://www.lawpracticetoday.org/article/lawyers-robotic-process-automation> (discussing use of robotic process automation in legal practice).

40. See, e.g., Steven Brykman, *Why We Desperately Need an AI Chatbot Law, Before We All Get Taken for a Ride*, CIO: THE BRYKMAN PREDICAMENT (June 13, 2018, 8:53 AM), <https://www.cio.com/article/3281375/why-we-desperately-need-an-ai-chatbot-law.html>.

41. See, e.g., Nick Wingfield, *As Amazon Pushes Forward with Robots, Workers Find New Roles*, N.Y. TIMES (Sept. 10, 2017), <https://www.nytimes.com/2017/09/10/technology/amazon-robots-workers.html>.

42. MIT TECH. REV. INSIGHTS, *Augmenting Human Intelligence*, MIT TECH. REV. (June 13, 2016), <https://www.technologyreview.com/s/601678/augmenting-human-intelligence>.

43. See Garrett Reim, *Augmented Reality Helmet Heads into Industrial*, L.A. BUS. J. (Sept. 2, 2016), <https://labusinessjournal.com/news/2016/sep/02/augmented-reality-helmet-heads-industrial>.

44. See Tim Lane, *A Short History of Robotic Surgery*, 100 ANNALS ROYAL COLL. SURGEONS ENG. 5–7 (Supp. 2018).

45. See Sumathi Reddy, *Robot-Assisted Surgery Costs More but May Not Be Better*, WALL ST. J.: YOUR HEALTH (Oct. 30, 2017, 1:11 PM), <https://www.wsj.com/articles/robot-assisted-surgery-costs-more-but-may-not-be-better-1509383463>.

46. See Thierer & Hagemann, *supra* note 3, at 344 ("Level 4: Full Self-Driving Automation. The vehicle is designed to perform all safety-critical driving functions and monitor roadway conditions for an entire trip. Such a design anticipates that the driver will provide

broad AI—the one capable of completing numerous cognitive tasks and called artificial general intelligence (AGI).

The differences between the three types of AI artifacts are summarized in Table 1:

TABLE I. COMPARISON OF AI AUTOMATION, AUGMENTATION AND AUTONOMY.

	Automation	Augmentation	Autonomy
Description	AI artifact completes well-defined and repetitive tasks following its programming.	AI artifact supports human worker in completing semi-structured and unstructured tasks.	AI artifact operates by following its programming, learning from its experiences, and adapting to new circumstances.
Scope of action	Replaces humans in known & understood activities. Can be applied to scenarios not yet contemplated but which are understood.	Supports humans in efficiently completing known & understood activities, discovering new information, patterns & connections; helping to find solutions to problems which are known but not understood.	Independently performs understood activities in anticipated or novel scenarios (known-unknowns). Can learn to act in scenarios not yet understood.
Human interaction	Direct supervision of operational process and outputs of the AI artifact.	Limited visibility into the operational processes completed by the AI artifact. Outputs of the AI artifact are considered by the human.	No visibility into the operational processes completed by the AI artifact. Outputs of the AI artifact impact the human as the anticipated or unanticipated result of the functioning.
Anticipation of failure	Error types and rates can be estimated.		Error types and rates cannot be anticipated in novel scenarios.

destination or navigation input, but is not expected to be available for control at any time during the trip. This includes both occupied and unoccupied vehicles.”).

III. LEGAL RESPONSIBILITY FRAMEWORKS

When it comes to “defective” products,⁴⁷ liability in any category (innovator, provider, user) will be based on three primary grounds: (i) when the product created deviates from its intended design (manufacturing defect), (ii) when the product should have been designed differently to avoid a foreseeable risk of harm (design defect), or (iii) when companies fail to provide instructions or warnings that could have avoided foreseeable risks of harm (failure to warn).⁴⁸ In all three cases, the plaintiff will assert that the product in question was “unreasonably dangerous”⁴⁹ due to a defect of some kind.⁵⁰ Each involves a fact-specific inquiry into the nature of the product and injury in question, as well as the circumstances in which the injury occurred.⁵¹

Manufacturing defects (for which companies are strictly liable)⁵² occur when products do not comport with the manufacturer’s intended design (e.g., if a data processing or coding error causes the AI to behave differently than it should).⁵³ The basis of liability is the product itself, which necessarily includes all of its component parts. Manufacturers thus can be held liable for product defects, even where the defect arises from a component manufactured by others.⁵⁴ However, in these cases the incorporated component itself must actually be defective. The original creator of the component is

47. There must be some cost or injury beyond the loss associated with the product itself not meeting specifications. *See, e.g.,* Sacramento Reg’l Transit Dist. v. Grumman Flexible, 158 Cal. App. 3d 289, 297 (Cal. Ct. App. 1984) (holding that a merchant could not assert a claim “in products liability for physical injury to its property where that injury consists of nothing more than the product defect upon which liability is founded”).

48. RESTATEMENT (THIRD) OF TORTS: PROD. LIAB. § 2 (AM. L. INST. 1998).

49. *Air & Liquid Sys. Corp. v. DeVries*, 139 S. Ct. 986, 1000 n.7 (2019) (Gorsuch, J., dissenting) (citing RESTATEMENT (SECOND) OF TORTS § 402A(1)(b) (AM. L. INST. 1965)) (noting that “[u]nder traditional tort principles, the seller of a defective, ‘unreasonably dangerous’ product may be liable to an injured user if the product ‘is expected to and does reach the user . . . without substantial change in the condition in which it is sold.’”).

50. “Whether a product is unreasonably dangerous is a distinct inquiry and must be established whether the claim is based on a manufacturing defect, a design defect, or a defective warning.” *Kaiser v. Johnson & Johnson*, 947 F.3d 996, 1008 (7th Cir. 2020).

51. *See, e.g., id.* (quoting *Koske v. Townsend Eng’g Co.*, 551 N.E.2d 437, 440–41 (Ind. 1990)) (“To decide whether a product is unreasonably dangerous, the fact-finder may consider several factors, including ‘the reasonably anticipated knowledge, perception, appreciation, circumstances, and behavior of expected users.’”).

52. Gregory C. Keating, *Products Liability as Enterprise Liability*, 10 J. TORT L. 41, 95 (2017) (citing RESTATEMENT (THIRD) OF TORTS: PROD. LIAB. § 2 cmt. a) (“Under both §402A and the Third Restatement, liability for manufacturing defects is strict enterprise liability.”).

53. Douglas A. Kysar, *The Expectations of Consumers*, 103 COLUM. L. REV. 1700, 1709 (2003).

54. *See, e.g.,* *Brocklesby v. United States*, 767 F.2d 1288, 1296 (9th Cir. 1985) (holding that “[a] seller is strictly liable for injuries caused by a defective product even though the defect originated from a component part manufactured by another party”).

not liable for injuries subsequently caused when another product manufacturer makes the poor choice to incorporate a component and task it to do something for which it is ill-suited.⁵⁵

Design defects, on the other hand, occur when there is no deviation from the manufacturing plan but where the design itself is flawed or the product could have been designed more carefully to lessen the risk of using it.⁵⁶ Design defects can be established in one of two ways. The first option allows plaintiffs to recover by showing that “the product failed to perform as safely as an ordinary consumer would expect when used in an intended or reasonably foreseeable manner” (known as the “consumer expectations test”).⁵⁷ Plaintiffs can prove their case using circumstantial evidence of how the product behaved in a particular situation to show that it must have been defective when causing their harm.⁵⁸

Alternatively, plaintiffs can establish a design defect by proving that a product should have been designed more safely in light of “the gravity of the danger posed by the challenged design, the likelihood that such danger would occur, the mechanical feasibility of a safer alternative design, the financial cost of an improved design, and the adverse consequences to the product and to the consumer that would result from an alternative design” (known as the “risk-benefit” test).⁵⁹ This test seeks to balance commercial and consumer interests. As one court described it:

A “risk-utility” analysis best protects both the manufacturer and the consumer. It does not create a duty on the manufacturer to create a completely safe product. Creating such a product is often impossible or prohibitively expensive. Instead, a manufacturer is charged with the duty to make its product reasonably safe, regardless of whether the plaintiff is aware of the product’s dangerousness.⁶⁰

Finally, manufacturers have a duty to warn about the risks associated with using their products. The duty arises “when the manufacturer ‘knows

55. See RESTATEMENT (THIRD) OF TORTS: PROD. LIAB. § 5 cmt. a (“If the component is not itself defective, it would be unjust and inefficient to impose liability solely on the ground” that others “utiliz[e] the component in a manner that renders the integrated product defective.”).

56. *Id.*; see also *Estate of Alex v. T-Mobile US, Inc.*, 313 F. Supp. 3d 723, 732 (N.D. Tex. 2018) (holding that plaintiff sufficiently alleged claim of design defect under Texas law by alleging that mobile device software was defectively designed because it failed to guard against software tampering that caused 9-1-1 telecommunications system to be congested and led to calls based on real emergencies, like babysitter’s call regarding injured child, to be placed on hold).

57. *Bookhamer v. Sunbeam Prods., Inc.*, 913 F. Supp. 2d 809, 818 (N.D. Cal. 2012) (citing *Barker v. Lull Eng’g Co.*, 573 P.2d 443, 457 (Cal. 1978)).

58. *Id.* (citing *Barker*, 573 P.2d at 454).

59. *Id.* (citing *Barker*, 573 P.2d at 455).

60. *Sperry-New Holland v. Prestage*, 617 So. 2d 248, 256 (Miss. 1993) (footnote omitted).

or has reason to know' that its product 'is or is likely to be dangerous for the use for which it is supplied' and the manufacturer 'has no reason to believe' that the product's users will realize that danger."⁶¹ This duty to warn extends only to the known aspects of the product itself—not to unanticipated situations where one company's product is incorporated into another's product in an unexpected way.⁶²

Tort liability for defective products also can be predicated on the doctrine of negligence. "In plain English, a person suing for negligence alleges that the defendant owed her a duty of reasonable care and injured her by breaching that duty."⁶³ Although negligence claims likely would play a role in lawsuits arising out of defective AI artifacts,⁶⁴ in order to facilitate a sharper focus on the product itself, negligence theories and contractual warranties under the Uniform Commercial Code⁶⁵ are not discussed further here. This is not to say that such alternative claims are unimportant. Indeed, it is a virtual certainty that claims arising under all three theories would be asserted simultaneously in any lawsuit arising out of an AI artifact.⁶⁶ How-

61. *Air & Liquid Sys. Corp. v. DeVries*, 139 S. Ct. 986, 993 (2019) (quoting RESTATEMENT (SECOND) OF TORTS § 388). "[W]arnings also may be needed to inform users and consumers of nonobvious and not generally known risks that unavoidably inhere in using or consuming the product." RESTATEMENT (THIRD) OF TORTS § 2, cmt. I. Users also may have to be advised when there are safer ways to use a product. *See, e.g., Liriano v. Hobart Corp.*, 170 F.3d 264, 270–71 (2d Cir. 1999) (discussing two categories of required warnings—the first covering the ways in which a product may be dangerous and the second instructing on safer ways to use that dangerous product).

62. *Air & Liquid Sys. Corp.*, 139 S. Ct. at 993–94 (2019) (noting "foreseeability that the product may be used with another product or part that is likely to be dangerous is not enough to trigger a duty to warn. But a manufacturer does have a duty to warn when its product requires incorporation of a part and the manufacturer knows or has reason to know that the integrated product is likely to be dangerous for its intended uses.").

63. John C.P. Goldberg & Benjamin C. Zipursky, *The Restatement (Third) and the Place of Duty in Negligence Law*, 54 VAND. L. REV. 657, 658 (2001).

64. Broadly speaking, "[t]ort law imposes 'a duty to exercise reasonable care' on those whose conduct presents a risk of harm to others." *Air & Liquid Sys. Corp.*, 139 S. Ct. at 993 (citing RESTATEMENT (THIRD) OF TORTS: PHYSICAL & EMOTIONAL HARM § 7 (AM. L. INST. 2010)).

65. *See* U.C.C. § 2-318 (AM. LAW INST. & UNIF. LAW COMM'N 1977) [hereinafter U.C.C.] (on third party beneficiaries of warranties express or implied); *see also* U.C.C. §§ 2-313 (express warranties), 2-314 (implied warranties of merchantability), 2-315 (implied warranty of fitness for a particular purpose), 2-316 (governing how warranties can be modified and disclaimed). For more details on contractual liability, *see* Timothy Davis, *UCC Breach of Warranty and Contract Claims: Clarifying the Distinction*, 61 BAYLOR L. REV. 783 (2009) (addressing "the distinction between breach of warranty and breach of contract claims" under the UCC); *see also* William L. Stallworth, *An Analysis of Warranty Claims Instituted by Non-Privity Plaintiffs in Jurisdictions That Have Adopted Uniform Commercial Code Section 2-318 (Alternatives B & C)*, 27 AKRON L. REV. 197 (1993) (discussing warranty claims and defenses under Article 2 of the UCC).

66. Claims seeking redress for injuries allegedly caused by sophisticated medical devices are a good analogy. *See, e.g., Tyree v. Bos. Sci. Corp.*, 54 F. Supp. 3d 501, 515 (S.D. W. Va. 2014) (plaintiff's complaint alleging negligence, strict liability for design and manufactur-

ever, there are a number of analytic benefits to focusing more precisely here on strict liability and leaving it to future scholarship to address the application of other compensation doctrines to AI artifacts.

First, the focus on strict products liability law makes it unnecessary to contemplate the separate but critical question of whether a legal “duty” exists for AI-related activities under common law or legislation. Duty is a prerequisite for all negligence claims,⁶⁷ and the scope of that duty is critical in determining whether the compensation gateway remains open or closed.⁶⁸ But it is not necessary here to get bogged down in the theoretical quagmire of where a legal duty might come from in relation to AI artifacts (meaning, whether the duty is an outgrowth of traditional tort principles or some other form of *sui generis* evolution).⁶⁹

Second, the focus on strict liability eliminates the challenges associated with contractual disclaimers⁷⁰ as well as the complications that can arise when multiple causes of action “become entangled with the structure of products liability actions and with different limitation periods and accrual rules applying to warranty, strict products liability and negligence actions.”⁷¹ This allows for more analytic consistency across different types of AI artifacts.

Finally, the focus on strict liability alone presents the clearest opportunity for policy consideration of the potential benefits and drawbacks of widespread AI innovation and adoption. Strict products liability focuses on

ing defects, strict liability for failure to warn, and breach of express warranty); *McPhee v. DePuy Orthopedics, Inc.*, 989 F. Supp. 2d 451, 455 (W.D. Pa. 2012) (alternative claims including strict liability, negligence, and breach of express and implied warranties).

67. See, e.g., *Phila. Indem. Ins. Co. v. Amazon.com, Inc.*, 425 F. Supp. 3d 158, 164–65 (E.D.N.Y. 2019) (holding that it was necessary for a plaintiff to establish that the defendant owed the plaintiff a legal duty before the defendant could be held liable for negligence).

68. See, e.g., David G. Owen, *The Five Elements of Negligence*, 35 HOFSTRA L. REV. 1671, 1672 (2007) (summarizing the traditional formulation of the elements of a negligence claim—duty, breach, causation, and damages).

69. For example, whether Autonomous AI can have its own legal obligations. See, e.g., Lawrence B. Solum, *Legal Personhood for Artificial Intelligences*, 70 N.C. L. REV. 1231, 1248–53 (1992) (in the context of capacity to uphold legal duties associated with trust administration, noting that an AI artifact could not be deemed to sufficiently capable of exercising judgment and discretion for purposes of legal liability unless it was proven capable of responding capably to unexpected changes in circumstance, exercising proper moral judgment and fairness, and making good decisions in the resolution of legal issues and disputes). For updated consideration of Solum’s early theories and predictions, see generally Dina Moussa & Garrett Windle, *From Deep Blue to Deep Learning: A Quarter Century of Progress for Artificial Minds*, 1 GEO. L. TECH. REV. 72 (2016).

70. See, e.g., Jay M. Feinman, *Implied Warranty, Products Liability, and the Boundary Between Contract and Tort*, 75 WASH. U. L.Q. 469 (1997) (noting distinctions between implied warranty and products liability standards).

71. Samuel J. M. Donnelly & Mary Ann Donnelly, *Commercial Law*, 47 SYRACUSE L. REV. 379, 398 (1997).

whether a product—negligently designed or not⁷²—is “unreasonably dangerous” to those with whom it subsequently comes into contact.⁷³ As one author put it,

When products are involved, negligence liability is liability for harms that would not have happened given reasonably safe product design and reasonable product warnings. By contrast, enterprise liability is liability for harms that flow from an activity’s (or an enterprise’s) characteristic risks, whether or not those risks should have been eliminated through the exercise of reasonable care.⁷⁴

This places the focus squarely on the risks associated with the inherent nature of AI artifacts themselves, together with the behavior that the law should require innovators, providers, and users to engage in in order to mitigate those risks.

A. Products Liability Law and AI Artifacts

The question of assessing legal responsibility for AI innovations is important because the law is currently poised to fall behind technological reality, if it hasn’t done so already. As discussed below, the liability risks associated with certain types of AI will be inherently unforeseeable. The various participants in the AI value chain undoubtedly will seek to allocate potential liability risks among themselves through purchasing contracts or intellectual property licensing regimes. They cannot, however, limit their collective liability when an AI artifact causes an indivisible harm to a third party outside the value chain.

Fundamental differences in the three primary categories of AI artifacts discussed above will lead to different types of legal responsibility for firms involved in bringing them to market. When a defective product causes harm, the law generally seeks to apportion liability based upon the relative fault of the actors involved (which includes injured parties who contribute

72. See RESTATEMENT (SECOND) OF TORTS § 402A(2) (providing for liability even where “the seller has exercised all possible care in the preparation and sale of his product”).

73. See, e.g., *Greenman v. Yuba Power Products, Inc.*, 377 P.2d 897, 900 (Cal. 1963) (holding that a party need not prove negligence to recover in products liability; rather, a plaintiff need only establish that a product was “defective” and that the defect caused the injuries in question). See also Keating, *supra* note 52, at 80 (discussing liability for selling a “product in a defective condition unreasonably dangerous to the user or consumer”) (citing RESTATEMENT (SECOND) OF TORTS § 402A(1)).

74. Keating, *supra* note 52, at 78. Absolute safety is not required. See, e.g., *Miller v. Dvornik*, 501 N.E.2d 160, 165 (Ill. App. Ct. 1986) (refusing to hold motorcycle manufacturer liable for injuries from vehicle accident because the injuries in question arose from the inherent risks in operating a motorcycle, which were obvious to the user).

to their own injuries—say, by misusing a product).⁷⁵ An important consideration is which parties were best positioned to prevent the harm in the first place—typically the product’s creators and sellers.⁷⁶ All of the actors in a product’s distribution chain can be held liable for injuries caused.⁷⁷ This usually includes:

- Creators/designers of a product or suppliers of component parts provided to manufacturers;
- Product manufacturers, who select and assemble a variety of component parts according to specific designs in order to create a product that meets a specific market need; and
- Distributors and retailers, ranging from mere sellers of “boxed” products from manufacturers (e.g., retailers), to players with a specific role in “prepping” the product for market and selling end users on its benefits (e.g., car dealers), to purveyors of sophisticated equipment who are heavily involved in helping users with product selection and in educating them on proper use (e.g., medical devices).⁷⁸

Products liability theories cannot be used to secure compensation from end-users or the occasional seller.⁷⁹ In cases where third parties are injured, claims against owners, who held either title or a leasehold interest in the product at the time it causes injury, and operators, who actually were using the product at the time the injury occurs, most likely would be brought under a negligence theory. This makes sense because any wrongful conduct that owners and operators engage in relates to how the product was maintained and/or used at the time of the accident, rather than how the product was designed or manufactured. In a similar vein, negligence would play some role in the apportionment of liability when human error combines

75. See, e.g., *Sperry-New Holland v. Prestage*, 617 So. 2d 248, 256 (Miss. 1993) (“In balancing the utility of the product against the risk it creates, an ordinary person’s ability to avoid the danger by exercising care is also weighed.”).

76. See, e.g., *Greenman*, 377 P.2d at 901 (noting that the purpose of imposing strict liability is “to insure [sic] that the costs of injuries resulting from defective products are borne by the manufacturer that put such products on the market rather than by the injured persons who are powerless to protect themselves.”).

77. See RESTATEMENT (THIRD) OF TORTS: PROD. LIAB. § 1 cmt. c (“One engaged in the business of selling or otherwise distributing products who sells or distributes a defective product is subject to liability for harm to persons or property caused by the defect.”). While the focus here is on U.S. law, similar rules apply in other countries—throughout Europe, for example. See Council Directive 85/374, art. 6, 1985 O.J. (L 210) 29 (EC) (“A product is defective when it does not provide the safety a person is entitled to expect.”).

78. See RESTATEMENT (SECOND) OF TORTS § 402A cmt. f (noting that liability “applies to any person engaged in the business of selling products for use or consumption. It therefore applies to any manufacturer of such a product, to any wholesaler or retail dealer or distributor”).

79. See *id.*

with a defective AI artifact to cause harm (as with the distracted driver in the Uber case discussed below,⁸⁰ where the driver was using her mobile phone at the time of the crash).⁸¹

For purposes of analytic clarity, a simplified liability chain is offered here to focus more precisely on three primary types of actors in connection with products constituting AI artifacts: innovators, providers, and users.⁸² Under this scheme, products liability theories would apply only to innovators and providers in relation to defective AI artifacts. Although separated here for analytical purposes, it is important to bear in mind that the development of AI artifacts can be achieved through either centralized or highly distributed processes. In practice, a single company even could play all three roles—innovating, producing, and using AI artifacts. Google parent company Alphabet’s “moonshot factory,”⁸³ for example, has internal laboratories designing algorithms for products and services including household robots,⁸⁴ self-driving cars,⁸⁵ the detection of kidney disease,⁸⁶ quantum computing,⁸⁷ combating threats to geopolitical instability,⁸⁸ and even creating human-level artificial intelligence.⁸⁹ All of these various projects provide feedback, learning, and user information that Alphabet then uses to innovate further in many different realms.

The development of AI artifacts proceeds through a variety of highly discrete actors who may range from high school students to professionals, with coding taking place in settings as diverse as corporate plants, garages, and even dorm rooms. These actors may participate in the development pro-

80. See *infra* notes 98–100 and accompanying text.

81. See, e.g., Tom Krisher, *Safety Agency Says Distracted Driver Caused Fatal Uber Crash*, U.S. NEWS (Nov. 18, 2019, 11:51 PM), <https://www.usnews.com/news/business/articles/2019-11-19/official-safety-lacking-before-uber-self-driving-car-crash>.

82. These terms were defined above. See *supra* p. 62.

83. See Dereck Thompson, *Google X and the Science of Radical Creativity*, ATLANTIC (Nov. 2017), <https://www.theatlantic.com/magazine/archive/2017/11/x-google-moonshot-factory/540648> (discussing Alphabet’s “moonshot factory”).

84. E.g., Tom Simonite, *Alphabet’s Dream of an ‘Everyday Robot’ Is Just Out of Reach*, WIRED (Nov. 21, 2019, 1:00 PM), <https://www.wired.com/story/alphabets-dream-everyday-robot-out-reach> (innovation on general purpose robots who could be tasked anything from helping the elderly to sorting trash).

85. E.g., Andrew J. Hawkins, *Inside Waymo’s Strategy to Grow the Best Brains for Self-Driving Cars*, VERGE (May 9, 2018, 8:00 AM), <https://www.theverge.com/2018/5/9/17307156/google-waymo-driverless-cars-deep-learning-neural-net-interview>.

86. E.g., Tom Simonite & Gregory Barber, *Alphabet’s AI Might Be Able to Predict Kidney Disease*, WIRED (July 31, 2019, 1:00 PM), <https://www.wired.com/story/alphabets-ai-predict-kidney-disease>.

87. E.g., *Quantum*, GOOGLE RSCH., <https://research.google/teams/applied-science/quantum> (last visited Feb. 15, 2020).

88. E.g., *Jigsaw*, GOOGLE, <https://jigsaw.google.com> (last visited Feb. 15, 2020).

89. E.g., Jeremy Kahn, *Inside Big Tech’s Quest for Human-Level A.I.*, FORTUNE (Jan. 20, 2020, 3:30 AM), <https://fortune.com/longform/ai-artificial-intelligence-big-tech-microsoft-alphabet-openai>.

cess through a broad range of ventures, belonging to different organizations or no organizations at all—even contributing anonymously. The contributions may involve discrete work and custom designed code or modifications of AI components developed by others contained in open-source libraries. All of this builds considerable opacity into the operations of complex systems powered by AI artifacts, especially where the interactions of the various AI components are not fully understood or transparent to users or regulators, or even to the innovators creating them or the providers bringing them to market.⁹⁰

The three AI value chain categories mapped against the traditional products liability categories appear in Table 2.

90. See Scherer, *supra* note 30, at 369–72 (discussing diversity, diffusion, discretion and opacity in AI development).

TABLE 2. AI VALUE CHAIN CATEGORIES MAPPED AGAINST TRADITIONAL LIABILITY CATEGORIES.

Value Chain for AI Artifacts	Products Liability Distribution Chain	Function	Example
Innovator	Designer/ Creator of component parts	Creates algorithms and AI systems	AI flight navigation system
Provider	Product Manufacturer	Assembles products incorporating AI systems & algorithms as components	Flying package-delivery drone
	Product Distributor	Final “product prep” and assembly, incl. instructions & warnings; training users	Drone dealer—helps buyers select needed features & instructs on safe operation
User	Product Owner	Uses the product to fulfill a certain business or personal need/desire	Amazon (or a similar company)
	Product Operator	Operating/controlling product at the time injury occurs	Drone delivery “pilot” hired by Amazon—employee or sub-contractor

Some have suggested an additional player for products liability claims specific to AI: allowing claims against the AI artifact itself or against a product incorporating that artifact. For example, if my friend’s butler robot crushes my fingers when it returns my car keys after dinner, I could sue the robot itself—just as I would a human being. Such proposals include things like giving AI artifacts formal legal personality, similar to the type given to corporations that enable them to be sued in courts,⁹¹ or even more broadly conceiving of AI algorithms as “citizens,” which would create duties for in-

91. David C. Vladeck, *Machines Without Principals: Liability Rules and Artificial Intelligence*, 89 WASH. L. REV. 117, 124 (2014) (“there is no a priori reason why truly autonomous machines should not be accorded some formal legal status, making them, like corporations and certain trusts, ‘persons’ in the eyes of the law and thus subject to suit.”).

novators and providers to “‘raise’ their AIs to act as responsible, productive members of society.”⁹²

This category is not considered here for two reasons. First, while theoretically intriguing, this category of liability is not yet recognized in law⁹³ and thus is of little help to innovators and providers making decisions in today’s uncertain environment. Second, these categorizations are conceptually dissimilar to the other participants in the value chain. They often amount to semantic shorthand for a risk-shifting policy mandate (e.g., requiring insurance on self-driving cars that moves with the vehicle itself, thus enabling the crash victims to “sue the car” in order to access the insurance).⁹⁴ Even the “algorithm as citizen” notion really amounts to a mandate for companies to design their AI responsibly and to ensure that their creations do not breach important social covenants.

B. *AI Artifacts and the Internet of Things*

The linkage of AI artifacts and the kinetic world through the IoT has the potential to drastically change our conception of what exactly the “product” in “product liability” is. The IoT takes the virtual world and makes it tangible. Consider just the problem presented by the more than 278 million passenger cars on the road today in the United States.⁹⁵ Modern cars contain 50–100 sensors constantly monitoring and recording everything from speed and fuel consumption to tire pressure.⁹⁶ The failure of any number of these instruments, or certain combinations of failures, could cause an accident and personal injury to the driver, passengers, or other motorists or pedestrians. That said, when there is an instrument failure, the root cause of an accident relating to the device failure in question often is traceable back to a particular component part or that part’s manufacturer.⁹⁷

The harder challenge arises not when a sophisticated component is physically defective, but rather when there is either a problem with the AI artifact operating that component or when AI software interacts with that

92. ACCENTURE, TECHNOLOGY VISION FOR ORACLE 5 (2018), https://www.accenture.com/_acnmedia/PDF-77/Accenture-Technology-Vision-Oracle-2018.

93. As the court noted in *United States v. Athlone Industries, Inc.*, 746 F.2d 977, 979 (3d Cir. 1984), “robots cannot be sued, but they can cause devastating damage,” such that “the defendant . . . was twice sued as the ultimate responsible distributor for various violations of the Consumer Product Safety Act.”

94. Vladeck, *supra* note 91, at 124 n.27.

95. Nathan Bomey, *Old Cars Everywhere: Average Vehicle Age Hits All-Time High*, USA TODAY (June 28, 2019, 1:06 PM), <https://www.usatoday.com/story/money/cars/2019/06/28/average-vehicle-age-ihs-markit/1593764001>.

96. J. Murgóitio & J. I. Fernández, *Car Driver Monitoring by Networking Vital Data*, in *ADVANCED MICROSYSTEMS FOR AUTOMOTIVE APPLICATIONS* 37, 38 (Jürgen Valldorf & Wolfgang Gessner eds., 2008).

97. See, e.g., *in re Takata Airbag Prods. Liab. Litig.*, 396 F. Supp. 3d 1101 (S.D. Fla. 2019) (consolidated class action claims arising out of defective airbags installed in vehicles).

component in a way that makes it behave in an unexpected manner. In such complex systems, “[t]he various component parts and their respective roles in causing a malfunction may be hard to discern and separate for the purpose of assigning responsibility.”⁹⁸

This is not to suggest that AI-related causation—and the interaction between AI artifacts and physical components—can never be proven. For example, a self-driving Uber SUV that struck and killed a pedestrian in Tempe, Arizona, in March 2018 was equipped with a camera as well as radar and LIDAR systems.⁹⁹ These components worked properly at the time of the accident, but the algorithm running them did not recognize a woman walking across the street with her bicycle outside of a crosswalk as a jaywalking pedestrian, and thus did not activate the braking system until it was too late to stop.¹⁰⁰ In a similar situation, a man was killed in a Tesla vehicle operating in auto-pilot mode when the vehicle struck an eighteen-wheeler after the system failed to distinguish the white truck from the bright sky.¹⁰¹ Neither of these vehicles was truly autonomous, and thus both would be classified as augmentation AI as they both provided driver assistance, rather than being trusted to do all of the actual driving.

In both of these instances, the AI operating the self-driving vehicle was found to be at fault for the crash. The problems with the AI artifacts in question¹⁰² were identified by government agencies tasked with post-accident

98. Gary E. Marchant & Rachel A. Lindor, *The Coming Collision Between Autonomous Vehicles and the Liability System*, 52 SANTA CLARA L. REV. 1321, 1328 (2012) (discussing autonomous vehicle systems).

99. Richard Gonzales, *Feds Say Self-Driving Uber SUV Did Not Recognize Jaywalking Pedestrian in Fatal Crash*, NPR (Nov. 7, 2019, 10:57 PM), <https://www.npr.org/2019/11/07/777438412/feds-say-self-driving-uber-suv-did-not-recognize-jaywalking-pedestrian-in-fatal->.

100. NAT’L TRANSP. SAFETY BD., VEHICLE AUTOMATION REPORT, TEMPE, AZ, NTSB NO. AZ-HWY-18-MH-010, 11–12 (2019), <https://dms.nts.gov/public/62500-62999/62978/629713.pdf>; see also Hannah Knowles, *Uber’s Self-Driving Cars Had a Major Flaw: They Weren’t Programmed to Stop for Jaywalkers*, WASH. POST (Nov. 6, 2019, 10:06 PM), <https://www.washingtonpost.com/transportation/2019/11/06/ubers-self-driving-cars-had-major-flaw-they-werent-programmed-stop-jaywalkers>. A separate lawsuit was brought against the city for allegedly creating the appearance of a safe crossing area that led directly into traffic. Patrick O’Grady, *Tempe Faces \$10M Suit in Uber Self-Driving Death*, PHX. BUS. J. (Feb 3, 2019, 9:00 PM), <https://www.bizjournals.com/phoenix/news/2019/02/03/tempe-faces-10m-suit-in-uber-self-driving-death.html>.

101. Danny Yadron & Dan Tynan, *Tesla Driver Dies in First Fatal Crash While Using Autopilot Mode*, GUARDIAN (June 30, 2016, 7:14 PM), <https://www.theguardian.com/technology/2016/jun/30/tesla-autopilot-death-self-driving-car-elon-musk>. Two other accidents involving Tesla’s auto-drive feature are under investigation at the time of writing. Alexander B. Lemann, *Autonomous Vehicles, Technological Progress, and the Scope Problem in Products Liability*, 12 J. TORT L. 157, 169–70 (2019) (providing details of two other Tesla crashes).

102. Major IoT risks—both of which seem to apply to the Uber and Tesla crashes—involve “sensor perception and decision-making under conditions of uncertainty.” See Mat-

investigation. The full power of the U.S. government—acting through the National Transportation Safety Board’s investigatory branches —was brought to bear in sorting out causation. Establishing the root causes of alleged failures by AI artifacts in cases that do not command the attention of powerful government agencies will be a much harder and more costly proposition.¹⁰³

Although the AI artifacts often will be involved in IoT-powered device failures,¹⁰⁴ proving exactly what went wrong with an AI artifact will become increasingly difficult as its complexity increases. Traditional concepts of foreseeability already are hard to apply because “an AI system’s solution may deviate substantially from the solution typically produced by human cognitive processes. The AI’s solution thus may not have been foreseeable to a human — even the human that designed the AI.”¹⁰⁵ And these are just the problems that exist with machine learning around narrowly defined but still complex tasks (e.g., processing data from a combination of radar, LIDAR, and camera inputs and ascertaining whether a vehicle should suddenly brake because a pedestrian or large truck lies in the way). The problem will become exponentially worse as AI transitions from automation and augmentation and AI autonomy becomes a reality:

[T]he possibility that an autonomous system will make choices other than those predicted and encouraged by its programmers is inherent in the claim that it is autonomous. If it has sufficient autonomy that it learns from its experience and surroundings then it may make decisions which reflect these as much, or more than, its initial programming. The more the system is autonomous then the more it has the capacity to make choices other than those predicted or encouraged by its programmers.¹⁰⁶

thew Michaels Moore & Beverly Lu, *Autonomous Vehicles for Personal Transport: A Technology Assessment* (June 2, 2011), <https://ssrn.com/abstract=1865047>.

103. See Kyle Graham, *Of Frightened Horses and Autonomous Vehicles: Tort Law and Its Assimilation of Innovations*, 52 SANTA CLARA L. REV. 1241, 1270 (2012) (“For a plaintiff to reach a jury on a design-defect claim, she may have to engage in a searching review of the computer code that directs the movement of these vehicles. This project may be difficult, and expensive.”).

104. See, e.g., Marchant & Lindor, *supra* note 98, at 1328 (noting, for example, that “the malfunction in an autonomous vehicle will usually be a programming error or system failure.”).

105. Scherer, *supra* note 30, at 365.

106. Robert Sparrow, *Killer Robots*, 24 J. APPLIED PHIL. 62, 70 (2007) (discussing AI in the context of LAWS—Lethal Autonomous Weapons Systems). See also Scherer, *supra* note 30, at 365 (noting that “the development of more versatile AI systems combined with advances in machine learning make it all but certain that issues pertaining to unforeseeable AI behavior will crop up with increasing frequency and that the unexpectedness of AI behavior will rise significantly.”).

Causation may be even more difficult to establish in the not-too-distant future where AI artifacts become intimately involved in creating other AI artifacts along with designing and building other products that cause injury after release into the stream of commerce.¹⁰⁷

Although cases involving AI artifacts certainly will be novel—particularly with autonomous AI—the use of traditional strict products liability principles to determine liability and set compensation in such cases is sound as a matter of public policy. Innovators and providers utilizing AI artifacts have special duties to society as a whole.¹⁰⁸ The major justification for imposing strict liability is that these actors are in the best position to avoid the potential for harm in the first place and subsequently to pass on and spread the cost of compensating victims for their injuries.¹⁰⁹

That said, it is important to note that the theoretical linkage between fault, prevention, and payments for injuries is at best imperfect. When multiple tortfeasors create a single indivisible harm (which usually is the case when a defective product hurts someone), the law generally holds each of them “jointly and severally liable.” This means that the injured party can recover the full amount of damages from any of them, regardless of that individual defendant’s actual percentage of fault.¹¹⁰ Plaintiffs’ lawyers thus will seek to assert claims against *all* companies or individuals involved—however slightly—in the creation and delivery chain that brought the defective product into contact with the injured party.

The plaintiff’s lawyer’s target zone thus will include all identifiable innovators, providers, and users that can be linked to the injury in question. Note also that with sophisticated products, such as self-driving cars, there may be multiple parties in each category—numerous AI innovators whose algorithms and systems are incorporated into an AI artifact created by a provider that in turn may be aggregated with other products in the activities of given owners or operators.

107. In a statement more rhetorical flourish than realistic description, Tesla’s Elon Musk has called his largely automated manufacturing facility (the Gigafactory) a “machine that builds the machine.” Elon Musk, CEO of Tesla, Inc., *quoted in* Sean O’Kane, *Tesla Will Live and Die by the Gigafactory*, VERGE (Nov. 30, 2018, 10:01 AM), <https://www.theverge.com/transportation/2018/11/30/18118451/tesla-gigafactory-nevada-video-elon-musk-jobs-model-3>.

108. “The basis for the rule is the ancient one of the special responsibility for the safety of the public undertaken by one who enters into the business of supplying human beings with products which may endanger the safety of their persons and property, and the forced reliance upon that undertaking on the part of those who purchase such goods.” RESTATEMENT (SECOND) OF TORTS § 402A(1)(b) cmt. f.

109. *See, e.g.,* *Daly v. Gen. Motors Corp.*, 575 P.2d 1162, 1165–69 (Cal. 1978) (discussing policy justifications for strict liability).

110. *See, e.g.,* *Chase v. Roy*, 363 N.E.2d 402, 408 (Mass. 1973) (“[I]f two or more wrongdoers negligently contribute to the personal injury of another by their several acts, which operate concurrently, so that in effect the damages suffered are rendered inseparable, they are jointly and severally liable.”).

In any case, once a presumably-blameless¹¹¹ injured plaintiff is fully compensated, all of the various defendants then will assert legal claims against each other to apportion damages based on each defendant's relative contribution to the overall harm.¹¹² Such contribution and indemnity claims have significant limitations. A party that had a minimal role in the product containing the defective AI artifact but which is financially well-heeled can be left holding the proverbial bag for the entire group if the other defendants lack insurance, go bankrupt, or benefit from a statutory liability cap—as in the case of many universities¹¹³—that make them unable to contribute their pro rata share.

IV. IMPLICATIONS FOR THE DEVELOPMENT AND DEPLOYMENT OF AI ARTIFACTS

A. Liability Implications—Impacts on AI Innovation

As discussed above, in addition to the three actors (innovators, providers, and users) in the value chain for AI artifacts, AI also can be characterized based on its level of sophistication. AI autonomy is markedly different than automation and augmentation. The key value proposition of autonomy is the machine's ability to learn from its environment and make its own decisions based on a wide variety of inputs. But its inherent unpredictability works against the players in the AI value chain because its autonomy also means that unexpected things can go wrong at unexpected times and in unexpected ways. Put more colloquially, and quoting former U.S. Defense Secretary Donald Rumsfeld: "There are known knowns. There are things we know that we know. There are known unknowns. That is to say, there are things that we now know we don't know. But there are also unknown unknowns. There are things we do not know we don't know."¹¹⁴

The liability matrix appears as set out in Table 3.

111. A plaintiff's recovery may be reduced or even barred outright if the plaintiff is at fault for their own harm. *See, e.g.*, Linda J. Rusch, *Products Liability Trapped by History: Our Choice of Rules Rules Our Choices*, 76 TEMP. L. REV. 739, 752 (2003) ("The defenses of contributory or comparative responsibility, assumption of the risk and product misuse are also used in strict liability cases.") (internal citations omitted). Product alteration also may bar recovery. *See, e.g.*, *Davis v. Berwind Corp.*, 690 A.2d 186, 190–91 (Pa. 1997) (holding that the removal of a safety device sufficiently altered the product to act as supervening cause of injury that relieved manufacturer of liability). But these kinds of misuse themselves are subject to a foreseeability test. *See, e.g.*, *Hall v. E. I. Du Pont De Nemours & Co.*, 345 F. Supp. 353, 363 (E.D.N.Y. 1972) ("A manufacturer cannot ignore a probable 'misuse' of his product.").

112. *See, e.g.*, MASS. GEN. LAWS ch. 231B, § 1 (West 2017).

113. *See, e.g.*, MASS. GEN. LAWS ch. 231B, § 85K (West 2017) (limiting liability of non-health care charitable organizations to \$20,000).

114. Donald Rumsfeld, Sec'y of Def., News Briefing (Feb. 12, 2002), *quoted in* David C. Logan, *Known Knowns, Known Unknowns, Unknown Unknowns and the Propagation of Scientific Enquiry*, 60 J. EXPER. BOTANY 712, 712 (2009).

TABLE 3. LIABILITY MATRIX FOR AI INNOVATORS, PROVIDERS, AND USERS.

Value Chain for AI Artifacts	Legal Duties	Automation (Known Unknowns)	Augmentation (Known Unknowns)	Autonomy (Unknown Unknowns)
Innovator	Design non-defective AI (“state of the art” quality); “manufacture” algorithms to design standards.	Known pathways and defined characteristics; can predict error rates & types.	Design based on known interactions with human operators (e.g., surgical robot).	Machine learning based on unknown interactions & environments.
Provider	Manufacture AI artifacts to design standards; instruct in the proper use; warn against foreseeable misuses—may be duty to make product unmodifiable.	Predictable operations and intersections between AI and tactile world.	Defined uses—types of injuries can be predicted based on known purpose of the augmentation.	End uses may be undefined and unknown at the time of manufacture and sale and may change over time.
User	Operate according to instructions; make authorized modifications only.	Known risk of modification (e.g., removal of safety device).	Known interactions between human and technology; training & instructions to mitigate harms.	Unknown multitude of potential users and types of interactions; unclear authorized versus unauthorized modifications.

The “known unknowns” are relatively easy cases because the role of the AI is limited and error rates and types can be predicted. A factory robot that punches three holes in a metal plate, rotates it 90 degrees, and moves it 12.5 centimeters to the left behaves in a fairly predictable way. If it moves 12.5 centimeters to the right, causing a machine to jam and subsequently break apart, injuring the operator, liability can be established in a fairly straightforward manner. Indeed, statistics from the Occupational Safety and Health Administration indicate that “‘dumb robots,’ designed for repetitive tasks

that are dirty, dangerous or dull,” are known to kill one to two factory workers per year in the United States.¹¹⁵

Likewise, for example, with surgical AI augmentation for a human doctor. If a robot that is supposed to move only two millimeters instead moves five, severing an artery instead of a ligament, we can determine what went wrong with relative ease after we factor in the doctor’s conduct and whether she met the standards of a reasonable medical professional¹¹⁶ in conducting the robotic-assisted surgery.¹¹⁷ And, as noted above, investigators have been able to determine what happened in those accidents involving the semi-autonomous Uber and Tesla vehicles.¹¹⁸

The real problem rests with AI autonomy. When an AI artifact causes physical harm that can be traced back to an unknown unknown, or even an unknowable one, assuming that accurate forensic analysis is even possible, who should bear the costs? In the autonomy zone, the AI artifact teaches itself and learns from its unique environment. It therefore may prove nearly impossible for innovators, providers, and users to predict the types of errors that could occur, the frequency of those errors, whether one error might cascade in a complex system and cause other types of errors, and the ultimate type and magnitude of harm that might arise. This means that the usual legal paradigms—focused on ensuring that innovators, providers, and users guard against foreseeable risks in design and manufacture and warn against foreseeable misuses when those risks cannot be avoided—are of little help in determining legal responsibility. This is a problem for all three actors, as the harm often will prove indivisible, and thus very difficult—as a forensic matter—to trace back to one particular source.

It is extremely difficult to discover whether software, as opposed to hardware, is responsible for the glitch that led to an accident. If the software is responsible, it would be hard to determine whether the precise cause was the operating system or the application (and, if the latter, which application). This analysis is all the more difficult

115. John Markoff & Claire Cain Miller, *As Robotics Advances, Worries of Killer Robots Rise*, N.Y. TIMES (June 16, 2014), <https://www.nytimes.com/2014/06/17/upshot/danger-robots-working.html>.

116. See, e.g., *Sargis v. Donahue*, 65 A.3d 20, 25 (Conn. App. Ct. 2013) (quoting *Mcchietto v. Keggi*, 930 A.2d 817, 821 (Conn. App. Ct. 2007)) (“[T]o prevail in a medical malpractice action, the plaintiff must prove (1) the requisite standard of care for treatment, (2) a deviation from that standard of care, and (3) a causal connection between the deviation and the claimed injury”).

117. See, e.g., *Balding v. Tarter*, No. 4-12-1030, 2013 WL 4711723, at *1 (Ill. App. Ct. Aug. 29, 2013) (denying appeal of summary judgment awarded to physician in medical malpractice case alleging that plaintiff suffered nerve damage to patient during laparoscopic robotic-assisted prostatectomy due to the surgeon’s lack of familiarity with robotic procedures, and hence excessive time necessary to perform the procedure).

118. See *supra* notes 99–101 and accompanying text.

where the software is open source (since no single author is responsible) and the hardware can be easily modified.¹¹⁹

Further complications arise with foreseeability when an AI artifact is modified, customized, or “taught” to behave in certain ways “because the manufacturer could not necessarily anticipate the universe of potential problems that might stem from third-party innovation and provide warnings or modify the platform design in response.”¹²⁰

Several important consequences follow. First, it is fanciful to suggest that simply because harm is unforeseeable, innovators, providers, and users would not be held liable for injuries caused by autonomous AI artifacts. From a policy perspective, AI value chain participants will be perceived—rightly or wrongly—as having created the problem in the first place, as best positioned to prevent accidents (or at least reduce their frequency and costs), and as having the most resources available to compensate injured parties.¹²¹

As a legal matter, all innovators and providers involved most likely will be held liable in the absence of a legislative mandate or regulatory scheme that creates immunity.¹²² The doctrine of *res ipsa loquitur* (Latin for “the thing speaks for itself”) aids injured plaintiffs in situations where the precise source of their harm is unclear but control of the harm-causing mechanisms rests exclusively in the hands of others.¹²³ In these cases, courts “infer a defect of some kind on the theory that the accident itself is proof of defect, even if there is compelling evidence that cuts against a defect theory.”¹²⁴

A related consequence is that the risks associated with autonomous AI artifacts may prove impossible—or prohibitively expensive—to insure. After all, insurance is predicated on actuarial tables that predict the likelihood of events based in part on historical occurrences. But with autonomous AI, there are no precedent scenarios, and there is much opacity around causa-

119. M. Ryan Calo, *Open Robotics*, 70 MD. L. REV. 571, 597 (2011) (internal citations omitted).

120. *Id.* at 596.

121. *E.g.*, Amar Kumar Moolayil, *The Modern Trolley Problem: Ethical and Economically-Sound Liability Schemes for Autonomous Vehicles*, 9 CASE W. RES. J. L. & TECH. 1, 20 (2018) (discussing these issues in the context of self-driving cars).

122. For example, Internet service providers benefit from an immunity from liability for defamation and related claims arising out of material posted on their websites. Communications Decency Act, 47 U.S.C. § 230(c)(1) (2012) (safe harbor provision immunizing Internet service providers from liability for third-party content posted on their websites). That said, it is unlikely that the public would stand for similar limitations on liability for innovators or providers of AI artifacts if physical injuries and property damage were at issue.

123. *See, e.g.*, Moussa & Windle, *supra* note 69, at 78 (discussing application of *res ipsa loquitur* to AI failures).

124. Vladeck, *supra* note 91, at 128; *see also* RESTATEMENT (THIRD) OF TORTS: PROD. LIAB. § 3(a) (*res ipsa loquitur* may apply where a product’s failure “was of a kind that ordinarily occurs as a product defect”).

tion, applicable theories of liability, event frequency, and overall damages. This likely will deter all but the most adventurous insurers from entering the market, at least until some of these parameters become known and thus predictable for assessing casualty rates, experience histories, and premium calculations. This in turn requires innovators, providers, and users of AI artifacts to retain considerable risk in their operations. Two market-facing consequences are likely to follow.

First, the risks may discourage smaller players, such as entrepreneurs, from entering the market or force existing participants to leave. Things may change once the liability rules are clarified, but for now, the potential for catastrophic liability may leave the field dominated by only one type of company: well-heeled corporations that can absorb considerable losses and self-insure against such risks, such as Google, Apple, Facebook, IBM, or large manufacturing companies like the “Big 3” automakers in the United States.¹²⁵ This, in turn, may limit the types of AI developed to those that suit the interests of large technology companies, instead of the greater good—or at least other types of AI suited to smaller, more entrepreneurial ventures.

Second, the enhanced risk may discourage investment in AI ventures or require a higher return. This will deter entrepreneurship and limit the spectrum of AI investment to financiers capable of withstanding greater losses. It also may impact the valuation of existing AI companies and the market for buying and selling them. Again, larger incumbent players will likely have a considerable advantage.

In the business-to-business context, where consumer protection concerns do not limit liability disclaimers,¹²⁶ it is likely that companies facing uncertain sources and degrees of liability will seek to protect themselves by allocating liability elsewhere in the value chain through the contracting process. The parties may say, in effect, “I don’t think this will happen, as I have a good product, but I don’t know for sure, so if something unexpected does happen, I want you to pay for it.” For example, a provider may seek in its

125. See Parker O’very, *3 Ways Self-Driving Cars Will Affect the Insurance Industry*, VENTURE BEAT (Jan. 26, 2018), <https://venturebeat.com/2018/01/26/3-ways-self-driving-cars-will-affect-the-insurance-industry> (noting that “Google, Volvo, and Mercedes-Benz already accept liability in cases where a vehicle’s self-driving system is at fault for a crash” and that “Tesla is taking things a step further by extending an insurance program to purchasers of Tesla vehicles.”). The larger market impacts on the insurance industry itself are likely to be significant. See, e.g., Paul Tullis, *Self-Driving Cars Might Kill Auto Insurance as We Know It*, BLOOMBERG (Feb. 19, 2019), <https://www.bloomberg.com/news/articles/2019-02-19/autonomous-vehicles-may-one-day-kill-car-insurance-as-we-know-it>.

126. See, e.g., Sarah Denis, *Using the Class Action Fairness Act as a Loophole Around the Magnuson Moss’s Jurisdictional Requirements*, 19 J. CONSUMER & COM. L. 124, 125 (2016) (noting that “Congress enacted the Magnuson-Moss Warranty Act, 15 U.S.C. §§ 2301 et seq. (2015), with the purpose of protecting consumers from deceptive warranty practices, specifically, narrow consumer product warranties that were often too convoluted for a layperson to understand.”).

purchase of an AI algorithm (perhaps structured as an IP license) to allocate all risks of harm caused by products containing that algorithm to the innovator who created it.¹²⁷ The innovator likely would prefer things the other way. In automation and augmentation settings, the issue can be resolved with insurance because a carrier can predict and model potential losses. The inability to do so in autonomy seems to suggest that the risk itself will be uninsurable, or perhaps only insurable at an unaffordable rate. Contracting will be further complicated by the inherent difficulty in valuing these defense and indemnity allocations, in that calculating the risk premium for these contract terms presents the same problems faced by insurers noted above.

This process will tend to favor players with market dominance, although unlike situations involving third-party liability, this does not automatically mean the largest and most financially powerful companies. A small innovator holding patents and other exclusive rights to a badly-needed technology can extract considerable monopoly-like concessions from providers seeking to create new products for fickle customers. The weakest party in the chain—in relative terms—could end up being stuck with one hundred percent of the liability if things go wrong. Again, this could impact entrepreneurship and the types of players willing to participate in the AI market.

The situation may also transform the very nature of the relationship between value chain participants. With autonomous AI, it may be impossible for parties to conclude bargained-for one-off transactions (or a series of them) for AI technology for at least two reasons. First, an algorithm never wears out or expires in the way that a mechanical device does. Errant AI incorporated into products may be blamed for unpredicted injuries many years or even decades later, particularly as other innovators and providers in the AI artifact's distribution chain exhaust any available insurance or go out of

127. An important part of software contracting involves controlling costs by limiting the liability of software developers for consequential damages arising from software errors. See, e.g., David R. Collins, *Shrinkwrap, Clickwrap, and Other Software License Agreements: Liti-gating a Digital Pig in a Poke in West Virginia*, 111 W. VA. L. REV. 531, 539 (2009) (“Because an entire company may rely on only one software program to conduct business, permitting recovery for products liability without limitation can potentially hold a developer liable for consequential damages amounting to millions or even billions of dollars for a single copy of a software product.”). Subsequent users of algorithms or software (e.g., those incorporating AI artifacts into their products and services) become bound by any liability provisions in the original distribution end-user license because any other use exceeds the license itself. See also Christian H. Nadan, *Software Licensing in the 21st Century: Are Software “Licenses” Really Sales, and How Will the Software Industry Respond?*, 32 AIPLA Q.J. 555, 588 (2004) (noting that “[i]f the distributor does pass on the end user license, the end user becomes bound by the liability limiting contract terms in the end user license. In this way, liability-limiting contract terms—even though not themselves license limitations—can pass down the distribution channel with the software copy.”).

business.¹²⁸ This may transform the nature of the dealings between the various parties in the chain into much more of an ongoing relationship with a longer-term perspective.

Second, the specter of indivisible liability will encourage innovators, providers, and users to work collaboratively on an ongoing basis to prevent harm in the first place. The roles create a feedback loop from innovation down to end-user and back again as the AI-enabled product learns with experience. Even where defense and indemnity are shiftable by contract, it often will prove far less expensive for the players to collaboratively repair problems discovered in AI artifacts than to incur the litigation costs of defending cases brought by injured third parties. Bad press has its own costs, and as a practical matter an indemnity clause in an IP licensing agreement may do little to restore a company's damaged reputation.

The need to cooperate also may force the parties to engage in different types of future transactions aimed at minimizing third-party liability. These could include duties to gather and retain information that will enable the parties to sort out responsibility amongst themselves if an unexpected event does occur. This may be harder than it sounds in practice, potentially creating additional challenges in complying with consumer privacy laws, government reporting obligations, maintaining trade secrets, and the like, not to mention adherence to contractual provisions that require it.

Cooperation also may require companies to make different design choices in the first place—for example, building reversibility into their AI that will allow them to undo undesirable learning by their algorithms or otherwise accounts for unexpected events,¹²⁹ or designing processes that enable querying the AI system and increase accountability without revealing the inner workings of the system itself.¹³⁰ This might increase costs, slow down the development pipeline, and make complex systems even more cumbersome. Again, the transformation of the dealings between the various parties

128. The nature and magnitude of liability certainly may be impacted by statutes of limitation (generally premised on the date that an injury occurs or upon which it is discovered). See, e.g., Richard C. Ausness, *Replacing Strict Liability with a Contract-Based Products Liability Regime*, 71 TEMP. L. REV. 171, 173 (1998) (noting that statutes of limitation in tort cases “begin to run when the plaintiff’s injury occurs, or in some cases, when the injury is discovered”). Certain industries also may lobby for special statutes of repose for their products, which time-limit claims to a period that begins to run on the date of first sale or transfer. See, e.g., *Montgomery v. Wyeth*, 580 F.3d 455, 467 (6th Cir. 2009) (holding that claims were barred under statute of repose that required claims against pharmaceutical companies to be brought within the shorter of ten years from the date of first sale or within one year after the drug’s expiration date).

129. See Lucas D. Introna, *Maintaining the Reversibility of Foldings: Making the Ethics (Politics) of Information Technology Visible*, 9 ETHICS & INFO. TECH. 11, 21 (2007).

130. See Finale Doshi-Velez et al., *Accountability of AI Under the Law: The Role of Explanation* (Dec. 20, 2019), <https://arxiv.org/abs/1711.01134>.

into an ongoing relationship has implications for the AI value chain as a whole.

In the absence of an option to spread risk to others, who “wins” in a contract negotiation depends on a lot of factors—including relative bargaining power, tolerance for risk, the depth of negotiators’ pockets, and ability to satisfy a judgment if one is assessed. Still, the negotiation process itself also has its own intrinsic value. It forces the parties to think long and hard about what might go wrong. Most contracts amount to bargained-for exchange, at least in theory, so the conversation will help to clarify at least some of the liability risks in a way that regulatory choices categorically allocating liability to one party versus another would not.

A more nuanced understanding of the nature of AI technology allows managers and executives to identify their organizations’ interests and objectives working with AI solutions as well as their legal liability exposure. This is particularly important if they are considering transitioning from automation and augmentation AI to autonomous AI. By studying their role(s) in the AI value chain, organizations can anticipate the nature of their legal exposure. Reframing their risk calculus away from a sole focus on contractual or tort liability in favor of promoting comprehensive responsibility in organizational behavior may be the most effective approach to mitigate future legal risks in the face of unknowns. Various factors support this argument, including the evolving nature of AI towards autonomy, the long shelf-life of digital inventions such as AI artifacts, and the modular nature of digital systems, with AI components that may find their way into a large number of AI solutions.

To limit their legal liability, organizations should actively foster the responsible development and use of AI innovation. This tenet, while applicable to any technology, is particularly important for AI for several reasons. First, as noted above, an AI artifact’s lifespan is potentially indefinite. The digital aspects of AI artifacts do not have the limited shelf-lives associated with physical decay, and an AI algorithm can be reused and fully or partially incorporated into novel AI systems. Second, AI is not subject to the replicability constraints of physical products; AI code can be replicated a large number of times at limited marginal cost and find its way into a great many types of AI systems.¹³¹ This greatly magnifies the consequences of error. Third, the impending evolution of AI systems towards autonomy presents unforeseeable outcomes. This increases the imperatives on designers, providers, and users to identify and control as much of the downside as possible in the AI context.

131. As noted above, modification of the AI artifact before incorporation into another product or process may break the liability chain, provided that the modification itself was not foreseeable.

All of this would suggest that companies are incentivized to take great care in minimizing all possible risks in the development and deployment of AI artifacts, controlling as much as possible on their own in order to minimize liability. But despite the importance of risk management in AI innovation, a new study of the actual behavior of AI professionals counterintuitively identifies considerable gaps in responsible behavior around AI systems, which companies should immediately begin to take steps to address.

B. Responsibility Implications—Studying the Ethical Perceptions of AI Professionals

Individuals with managerial responsibilities for the implementation, design, development, and distribution of AI systems within their organizations, or who are involved in offering such products and services to other organizations, also fall into the categories of AI innovators, distributors, and users. Given their understanding of AI technologies, as well as system capabilities and the risks associated with their use, AI professionals serve as critical participants in the value chain of AI development.

In order to preliminarily explore ethical perceptions of those making decisions about AI and the level of responsibility in the development of AI innovations,¹³² we conducted a brief survey of technology managers and executives who play some role in AI development and who have expertise in the field of artificial intelligence (designated herein as “AI professionals”).¹³³ The respondent AI professionals answered ten questions about personal and corporate responsibility and ethical behavior within their organizations in the development and deployment of AI technologies.

Eight questions that dealt with the importance of ethical principles and organizational practices in the ethical development or implementation of technologies are presented in Table 4, along with the percentages of positive responses. Overall, we found strong agreement among AI professionals on these questions. The respondents support the importance of ethical behavior in the development and implementation of new technologies, contemplate the potential social harm of their work, and follow ethical principles in their own development and implementation efforts.

With respect to ethical behavior at the enterprise level, the AI professionals surveyed also consistently report that their organization has written policies for the socially responsible development and implementation of new technologies, that these policies are followed, and that the declared values of their organization agrees with their own. Most also agree that

132. There is some dispute whether such questions even are possible to answer. See, e.g., Cade Metz, *Is Ethical A.I. Even Possible?*, N.Y. TIMES (Mar. 1, 2019), <https://www.nytimes.com/2019/03/01/business/ethics-artificial-intelligence.html>.

133. See *infra* Technical Appendix for details about sample collection, preparation, and other descriptive statistics, including demographic variables.

company leadership follows the organization's stated ethics and values. Finally, a large number of these AI professionals disclose that their organization offers mechanisms to report improper behavior in the development and implementation of new technologies.

TABLE 4. RESPONSES TO QUESTIONS ABOUT ETHICAL PRINCIPLES AND THE RESPONSIBLE DEVELOPMENT OF TECHNOLOGIES.

Question	Agree
Q1. Being ethical is highly important in the development and/or implementation of new technologies.	98.8%
Q2. My organization ensures that its development and/or implementation of new technologies is done in a socially responsible manner.	97.56%
Q3. I agree with the declared values of my organization in relation to the socially responsible development and/or implementation of new technologies.	95.12%
Q4. I feel a personal obligation to ensure the socially responsible development and/or implementation of new technologies.	95.12%
Q5. My organization has written policies relating to the socially responsible development and/or implementation of new technologies.	93.90%
Q6. The behavior of my organization's leaders in relation to the socially responsible development and/or implementation of new technologies is consistent with the stated ethics and values of my organization.	93.90%
Q7. I have the opportunity in my work to consider the wider social implications of what I am working on and the potential social harm(s) it could create.	93.90%
Q8. My organization has procedures for reporting improper behavior in relation to the socially responsible development and/or implementation of new technologies.	87.80%

Two other questions for AI professionals focused on the scope of their personal responsibility and their actions to date in terms of reporting any

concerns relating to the socially responsible use of technologies. These findings are summarized in Table 5.

TABLE 5. RESPONSES TO QUESTIONS ABOUT THEIR ETHICAL RESPONSIBILITY AND PERSONAL CONCERNS ABOUT THE RESPONSIBLE USE OF TECHNOLOGIES.

Question	Agree
Q9. I see my role as creating, implementing, and/or using the best technology possible; it is the responsibility of others to determine how such technology should be used and what limits should be imposed on it.	90.24%
Q10. Report having personal concerns relating to the socially responsible use of technology but did not communicate that concern to supervisors	74.39% ¹³⁴ (n=61)

The survey results reveal a significant responsibility gap between the beliefs and the actions of the AI professionals. A large majority of respondents expressed great awareness and concern about the responsible development and use of technology. They also agreed solidly with the values of their companies in this regard and reaffirmed the importance of ethical behavior, including their own actions, in the development and implementation of new technologies. Yet less than ten percent of them felt that it was their personal responsibility to determine the scope of responsible use for the technologies they build or deploy. Ninety percent of them thought that was someone else's job. And it's clear that these are not theoretical issues; nearly three quarters of the respondents confirmed not reporting to supervisors actual concerns relating to the socially responsible use of technology in their organizations. The most frequent justifications for not reporting are captured in Table 6.

134. Percentage reflects survey respondents who identified at least one reason for not communicating ethical concerns to their supervisors. On average, those responding positively identified 2.32 reasons for not doing so.

TABLE 6. TOP REASONS FOR NOT REPORTING TO A SUPERVISOR
CONCERNS ABOUT THE SOCIALLY RESPONSIBLE USE OF TECHNOLOGY

Reason	Percent (count) ¹³⁵
The belief that reporting would not be anonymous	39.34% (24)
The fear of retaliation	36.07% (22)
The belief that someone else would do it	31.15% (19)
The belief that corrective action would not be taken	31.15% (19)
To avoid getting colleague(s) into trouble	26.23% (16)
Out of loyalty to the department or business unit	18.03% (11)
The belief that the issue was not important to immediate supervisors	16.39% (10)
The belief that it was not part of their job	13.11% (8)

As discussed above, in order to minimize legal liability for AI innovations, it would seem that decision-makers leading the innovation and production of AI solutions in modern enterprises have many increased incentives to act *responsibly*. At a minimum, this means operating under a set of known organizational values and using clearly defined procedures to identify and address ethical violations. Yet based on these findings, this does not appear to be happening—at least amongst the AI professionals who participated in the survey.

This study of AI professionals reveals an apparent gap between stated responsibilities and values and behaviors when managing technology. They report clear personal and organizational values to minimize negative social impacts and act ethically in AI development. Yet they do not feel directly responsible for the actual execution of such responsible behaviors and consistently fail to report ethical concerns. This AI responsibility gap is of great

135. See *infra* Technical Appendix for details on the collection and recoding of survey responses.

consequence to organizations working or relying on AI innovations, as it could greatly increase future legal exposure in automation and augmentation as well as autonomy. Given the liability risks inherent in the creation, deployment, and utilization of AI systems, it is important for organizations developing and using AI to address this gap—and quickly.

A few clear recommendations and challenges emerge from this study. First, companies must not only identify clear values and hire individuals who share those principles, but also reframe responsibility in the use and development of AI innovations to favor a deeper understanding of the social and legal consequences and a willingness to act based on that understanding.

Second, companies should look beyond short-term economic considerations and immediate outcomes when acting on their AI responsibilities, and they should train their employees to do so as well. In other words, companies should avoid short-term thinking when evaluating AI systems' impact and revenue-oriented metrics of performance that discourage employees from taking responsible action. Many challenges will arise from short-term organizational perspectives, particularly in digital companies seeking to market innovations and generate revenue quickly, especially given the unforeseeable legal liabilities that their AI innovations may create in the long term. While this change in perspective may be difficult to achieve and requires further research, it is one that companies should seek to encode in their organizations' actions as well as their vision statements and operational guidelines. This is especially important because algorithms do not "wear out" like tangible objects do, such that the applicable time horizon in which harm can arise is far longer. AI managers and executives thus should adopt a long-term approach in evaluating the consequences of their AI innovations and products.

Finally, it is important to create an opportunity to fix problems when they become apparent, especially those problems that could not have been anticipated beforehand (the "unknown unknowns"). As a matter of company policy, AI innovators and providers should aim to build reversibility¹³⁶ and visibility into AI artifacts so that actions can be undone and to ensure that the unexpected behavior of the AI system can be explored and better understood. As noted above, all innovators and providers in the AI artifact supply chain have significant incentives to collaborate with one another in the long term to investigate and repair "defective" AI artifacts.

V. CONCLUDING THOUGHTS

The goal of this article was to begin a conversation about legal, ethical, and managerial responsibility in creating and using AI technologies. Many

136. See Inrona, *supra* note 129, at 15–23.

more questions require careful consideration. One of these is how to balance the various stakeholder interests identified herein from a policy perspective. Liability for AI participants can inhibit innovation and product development, reduce investment, limit the willingness of owner/operators to purchase AI artifacts in the first place, and tie up resources in legal defense and cross-claims for indemnity that could be used for other things.

Too much liability will create a drag on the entire AI value chain, which is the lifeblood of AI advancement. This suggests that some limitations on liability may be appropriate in some circumstances.¹³⁷ Yet innovators and producers also clearly need incentives to be as careful as possible in their design, manufacture, and operational choices, and imposing liability on them clearly advances this goal.¹³⁸ The right balance between these competing interests to some extent involves larger questions of social welfare maximization. Although the degree to which tort law promotes efficient resource allocation is an open question,¹³⁹ some suggest that the future benefits of AI will far outweigh its costs, even when human life is concerned.

Take the case of autonomous vehicles. How the legal system should attribute responsibility for the crashes that autonomous vehicles cause is an open and hotly debated question.¹⁴⁰ Yet given the prediction that automated vehicles will save thousands of lives per year by eliminating human error,¹⁴¹ one author hypothesized “[w]ho cares if Tesla’s Autopilot could have been improved in some incremental way that would have prevented three fatalities if it can save tens of thousands every year? To the welfare economist, the appropriate question is not whether Autopilot was “defective,” but rather whether imposing liability will help encourage an efficient allocation of resources.”¹⁴²

These kinds of considerations ultimately may determine the extent to which individuals will be permitted to recover at all, as well as the measure of damages that will be available.¹⁴³ No-fault regimes, such as those pro-

137. See, e.g., Calo, *supra* note 119, 601–12 (discussing selective immunity for robots).

138. See, e.g., Weston Kowert, Note, *The Foreseeability of Human-Artificial Intelligence Interactions*, 96 TEX. L. REV. 181, 199 (2017) (noting that if vaccine and automobile “manufacturers aren’t going to be held liable, then they lose much of their incentive to improve their product.”).

139. See, e.g., John C. P. Goldberg, *Twentieth-Century Tort Theory*, 91 GEO. L.J. 513, 544–60 (2003) (discussing various facets of the interpretive and prescriptive economic deterrence theories and critiques of the doctrines).

140. See Alexander B. Lemann, *Autonomous Vehicles, Technological Progress, and the Scope Problem in Products Liability*, 12 J. TORT L. 157, 175–76 (2019).

141. *Id.* at 157 (“Autonomous vehicles are widely expected to save tens of thousands of lives each year by making car crashes attributable to human error – currently the overwhelming majority of fatal crashes – a thing of the past.”).

142. *Id.* at 192–93.

143. It is not uncommon to impose damage limiting schemes in circumstances where the public interest outweighs that of individual claimants—as in the case of medical malpractice. See Michael J. Cetra, *Damage Control: Statutory Caps on Medical Malpractice Claims*, *State*

posed in Europe for certain AI products,¹⁴⁴ may prove useful, though their application to various players within the AI value chain has yet to be considered. Insurance pools funded by manufacturing groups also may be worth exploring.¹⁴⁵

Another option is to consider the pros and cons of leaving the legal gray area untouched and simply letting courts innovate through the common law.¹⁴⁶ While this would make the AI development landscape harder for companies to navigate, it has some appeal. This is particularly true because it can be extraordinarily difficult to achieve correct regulatory solutions *ex ante*, especially for rapidly-developing technologies. As Judge Easterbrook noted in connection with emerging issues of cyberspace in the mid-1990s, a market-based solution might be best:

“Error in legislation is common, and never more so than when the technology is galloping forward. Let us not struggle to match an imperfect legal system to an evolving world that we understand poorly. Let us instead do what is essential to permit the participants in this evolving world to make their own decisions.”¹⁴⁷

Much useful dialogue will come simply from helping to ask the right questions, particularly because autonomous AI will undoubtedly raise unique issues that the law has not yet encountered.¹⁴⁸

Constitutional Challenges, and Texas' Proposition 12, 42 DUQ. L. REV. 537, 542–43 (2004) (discussing state constitutional amendments authorizing state legislatures to cap damages in medical malpractice cases).

144. *E.g.*, Roeland de Bruin, *Autonomous Intelligent Cars on the European Intersection of Liability and Privacy – Regulatory Challenges and the Road Ahead*, 7 EUR. J. RISK REG. 485, 490 (2016).

145. Insurance pools may be especially useful in the face of potentially unlimited liability that could leave a company bankrupt. *See, e.g.*, Leonard J. Long, *Bankruptcy Lesson of Future Mass Tort Claims: Potential Mass Tort Victims Should Have Catastrophic Injury Insurance*, 16 QUINNIAC L. REV. 357, 367 (1997) (“[I]f the remedy (principally compensatory damages for injuries suffered), available to tort victims who suffer catastrophic injury or substantial loss of income is a separately funded insurance pool, then such tort victims are not dependent on the continued solvency of their tortfeasor and the precarious fortunes of the bankruptcy process.”).

146. *See, e.g.*, Marta Katarzyna Kolacz et al., *Who Should Regulate Disruptive Technology?*, 10 EUR. J. RISK REGUL. 4, 4 (2019) (arguing that the judiciary is best suited to handling cases presenting risky new technologies whereas regulation was required for technologies whose risks are not foreseeable at the time of the technological innovation); Brandon W. Jackson, *Artificial Intelligence and the Fog of Innovation: A Deep-Dive on Governance and the Liability of Autonomous Systems*, 35 SANTA CLARA HIGH TECH. L.J. 35, 35 (2019) (discussing a significant role for the judiciary in shaping legal regulation of AI).

147. Frank H. Easterbrook, *Cyberspace and the Law of the Horse*, 1996 U. CHI. LEG. F. 207, 215–16 (1996).

148. *See* Jack M. Balkin, *The Path of Robotics Law*, 6 CAL. L. REV. CIR. 45, 49–51 (2015).

TECHNICAL APPENDIX

The data (n=82) obtained in this study¹⁴⁹ through an online Qualtrics panel focused exclusively on managers and executive-level employees who supervised other employees and for whom artificial intelligence technologies were a significant part of their role, educational background, or work experience. All participants answered all questions in the survey, including for all of the responses discussed above,¹⁵⁰ and the demographic questions summarized below.

Participants answered survey questions 1–9 as seven-point Likert scale items, ranging from “strongly agree” to “strongly disagree”. The positive responses reported in Table 4 and Table 5¹⁵¹ correspond to the sum of respondents who selected “strongly agree,” “agree,” or “somewhat agree,” out of the total number of participants. Survey question 10 (Table 4) allowed respondents to select multiple items from a list of nine reasons (including “not applicable,” with a narrative option to provide a reason) for not communicating their ethical concerns. Table 5¹⁵² reflects the number of times each reason was reported by one of the survey participants, out of the total number of participants who reported concerns.

The authors calculated descriptive statistics for demographic variables collected from the sample. Participants average 3.68 years of professional experience (standard deviation of 1.02 years), and 45.12% of them were female. All respondents supervised other employees, as summarized in Table 7.

TABLE 7. NUMBER OF EMPLOYEES SUPERVISED BY PARTICIPANTS

Number of Employees Supervised	Percentage
1-9	13.4%
10-99	54.9%
100+	31.7%

149. Original data from the study remains on file with the authors.

150. See discussion *supra* Part IV.

151. See *supra* Tables 4–5.

152. See *supra* Table 5.

Industry sectors included leisure and hospitality, professional and business services, and health services (see Table 8 for details).

TABLE 8. SECTORS REPRESENTED AMONG THE TECHNOLOGY MANAGERS AND EXECUTIVES SURVEYED

Sector	Number of Respondents Surveyed
Leisure and Hospitality	29
Professional and Business Services	14
Health Services	10
Trade, Transportation, and Utilities	9
Natural Resources and Mining	7
Construction and Manufacturing	5
Education	3
Other	3
Financial Activities	1
Information Technology	1

The most common company size among those surveyed was 100 to 499 employees, as set forth in Table 9.

TABLE 9. SIZES OF THE COMPANIES OF THE TECHNOLOGY MANAGERS
AND EXECUTIVES SURVEYED

Percentage	Company Size
1 to 9	2.44%
10 to 99	20.73%
100 to 499	32.93%
500 to 999	20.73%
1000+	23.17%