THE COLUMBIA
SCIENCE & TECHNOLOGY
LAW REVIEW

VOLUME XXV STLR.ORG SPRING 2024

A PRODUCTS LIABILITY FRAMEWORK FOR AI

Catherine M. Sharkey[*]

A products liability framework, drawing inspiration from the regulation of FDA-approved medical products—which includes federal regulation as well as products liability—holds great promise for tackling many of the challenges AI poses. Notwithstanding the new challenges that sophisticated AI technologies pose, products liability provides a conceptual framework capable of responding to the learning and iterative aspects of these technologies. Moreover, this framework provides a robust model of the feedback loop between tort liability and regulation.

The regulation of medical products provides an instructive point of departure. The FDA has recognized the need to revise its traditional paradigm for medical device regulation to fit adaptive AI/ML technologies, which enable continuous improvements and modifications to devices based on information gathered during use. AI/ML technologies should hasten an even more significant regulatory paradigm shift at the FDA away from a model that puts most of its emphasis on (and resources into) ex ante premarket approval to one that highlights ongoing postmarket surveillance. As such a model takes form, products liability should continue to play a significant information-production and deterrence role, especially during the transition period before a new ex post regulatory framework is established.

[*] Segal Family Professor of Regulatory Law and Policy, New York University School of Law. Callen DiGiovanni (NYU 2024) and Nicholas Tilmes (NYU 2025) provided excellent research assistance. I am grateful to Thomas Merrill for insights he provided as a commentator at the symposium and to Ryan Abbott, Bryan Choi, Chinmayi Sharma, and Benjamin Zipursky for feedback. I also appreciate input from my fabulous Artificial Intelligence seminar students, who inspired this Essay.

I. INTRODUCTION

As a transformative technology, artificial intelligence ("AI") promises revolutionary advances while simultaneously posing new risks to society.[1] This Essay proposes a products liability conceptual framework that would tackle many of the regulatory challenges posed by AI and machine learning ("ML").[2] It draws inspiration from the regulatory scheme governing medical products approved by the Food and Drug Administration (FDA), which includes both ex ante federal regulation and ex post products liability. FDA regulation of AI-enabled medical devices provides a particularly apt point of departure for consideration of a products liability framework for AI. The agency has been at the vanguard of revising its traditional paradigm for medical device regulation to fit adaptive AI/ML technologies, which continuously improve and modify devices based on data collected during use. However, this Essay further suggests that AI/ML technologies should inspire an even more significant regulatory paradigm shift away from a model that focuses on (and invests in) premarket approval to one that emphasizes ongoing postmarket surveillance.[3] As such a model emerges, products liability should continue to play a significant role in information-forcing and deterrence, especially during the transition period before a new ex post regulatory framework for adaptive AI/ML technologies is in place.

The turn to a products liability framework for AI is notable for (at least) two reasons. First, to date, legal commentators' focus has been elsewhere, as the

[1] President Biden's Executive Order on Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence repeatedly sounds the theme of AI's "extraordinary potential for both promise and peril." Exec. Order No. 14,110, 88 Fed. Reg. 75191, 75191 (Oct. 30, 2023).

[2] AI is "a branch of computer science, statistics, and engineering that uses algorithms or models to perform tasks and exhibit behaviors such as learning, making decisions, and making predictions." Shawn Forrest, *Artificial Intelligence/ Machine Learning (AI/ML)-Enabled Medical Devices: Tailoring a Regulatory Framework to Encourage Responsible Innovation in AI/ML*, FOOD & DRUG ADMIN. (2023), https://www.fda.gov/media/160125/download [https://perma.cc/L8BH-5GHB] (adapted from *Machine Learning-Enabled Medical Devices: Key Terms and Definitions*, INT'L MED. DEVICE REGUL. F. (Mar. 9, 2022), https://www.imdrf.org/documents/machine-learning-enabled-medical-devices-key-terms-and-definitions [https://perma.cc/PKR8-VF99]). ML is "a subset of AI that allows ML models to be developed by ML training algorithms through analysis of data, without models being explicitly programmed." *Id.*

[3] *See* Catherine M. Sharkey & Kevin M.K. Fodouop, *AI and the Regulatory Paradigm Shift at the FDA*, 72 DUKE L. J. ONLINE 86, 97–103 (2022).

proliferation of AI/ML has raised high-profile defamation and copyright legal issues, with products liability newly emergent on the scene.[4] Second, commentators have suggested that the "learning" aspect of AI/ML demands an entirely new regulatory approach, without serious consideration of the extent to which products liability can adapt to face such a challenge.

The *Gonzalez v. Google*[5] case before the U.S. Supreme Court last Term illustrates a subtle turn in focus to products liability. The Court granted certiorari to consider whether a web service forfeits the protection of Section 230 of the Communications Decency Act,[6] which immunizes web platforms from civil liability claims arising from third-party content, when it uses algorithmic recommendations (like the content that autoplays after YouTube videos) to organize and promote certain user-generated content.[7] While the Court declined to decide that specific issue, it drew attention to looming defamation and copyright issues on the horizon. For example, AI-powered search engines (like Google's Bard and Microsoft's Bing) and large language models ("LLMs") (like OpenAI's ChatGPT)[8] are known to "hallucinate" and generate misleading or false information;[9] if such mistakes cross the line into defamation, search providers could be at serious risk of lawsuits. Moreover, AI-generated content may expose developers to copyright infringement claims—which fall outside of Section 230's purview—as generative AI models are trained on vast swathes of copyrighted material that they sometimes reproduce without alteration or attribution.[10]

However, the most telling moment during the oral argument in *Gonzalez*—at least to me—has received comparatively less attention than these looming

[4] Moreover, the canonical types of algorithmic harms that received the most attention at the symposium—reputational, representational, etc.—are not classic products liability harms, which typically involve discrete physical injuries. *See* Symposium, *Accountability & Liability in Generative AI: Challenges & Perspectives*, 25 COLUM. SCI. & TECH. L. REV. 190 (2023).

[5] Gonzalez v. Google LLC, 598 U.S. 617 (2023). The suit stemmed from a 2015 Islamic State shooting in Paris that killed student Nohemi Gonzalez. *Id.* at 619–20. Her family argued that YouTube had recommended videos by terrorists and therefore violated laws against aiding and abetting foreign terrorist groups. *See* Gonzalez v. Google LLC, 2 F.4th 871, 881 (9th Cir. 2021), *vacated and remanded*, 598 U.S. 617 (2023), and *rev'd sub nom.* Twitter, Inc. v. Taamneh, 598 U.S. 471 (2023). While Section 230 of the Communications Decency Act typically protects sites from liability over user-generated content, the complaint argued that YouTube created its own speech with its recommendation algorithm. *Id.*

[6] 47 U.S.C. § 230.

[7] *See Gonzalez*, 598 U.S. at 622.

[8] *See* Tu Vu, *FreshLLMs: Refreshing Large Language Models with Search Engine Augmentation*, ARXIV (Nov. 22, 2023) (preprint), https://arxiv.org/pdf/2310.03214.pdf [https://perma.cc/K5WX-YFU2] (comparing LLMs (two generative pre-trained transformer ("GPT") models and Perplexity) with Google search and finding that, while the models struggled to answer questions that contained false information, they were as, or in some cases more, accurate than Google).

[9] *See* Yue Zhang et al., *Siren's Song in the AI Ocean: A Survey on Hallucination in Large Language Models*, ARXIV (Sept. 24, 2023) (preprint), https://arxiv.org/pdf/2309.01219.pdf [https://perma.cc/Y894-3D35].

[10] *See, e.g.*, Complaint at 2, New York Times Co. v. Microsoft Corp., No. 1:23-cv-11195 (S.D.N.Y. Dec. 27, 2023).

copyright and defamation issues. At the tail end of her argument, Lisa Blatt, Google's attorney, admitted that she was far less concerned about losing Section 230 immunity for defamation claims than for products liability claims.[11] Indeed, a host of pending cases allege that social media platforms are defectively designed, causing addiction and mental health issues among youths,[12] and some courts have already rejected motions to dismiss such claims on the grounds that Section 230 prevents these platforms from being held liable for content created by their users.[13]

Even as the emergence of social media/AI cases pushes the boundaries of products liability, all too often critics dismiss the potential of a products liability framework for AI. They insist that society has never before faced such a seismic technological shift and that we thus need an entirely distinct approach to regulate harms generated by AI.[14] Critics suggest that regulating AI/ML demands a unique regulatory approach because, as AI/ML technologies are sent out into the world and encounter new situations, they learn and change in real time.[15] The implication is that this adaptive "learning" aspect of ML and the interactive aspect of generative AI pose challenges different in kind from those raised by any prior technology regulated by products liability. But such critics have failed to appreciate how robust

[11] Blatt responded to a question about defamation liability, saying:

> No, I'm not worried about the defamation claim. *I'm worried for a products liability claim* or what the government kept saying, your design choices. Those could just be a product liability claim or a negligence claim. You negligently went alphabetical or you negligently featured whatever you featured that made my, you know, kid addicted to whatever it was. And that—those kind of claims happen because they're publishing. And the whole point of getting this statute was to protect against publishing. So whatever is publishing, inherent to publishing, yeah, has to be covered.

Transcript of Oral Argument at 161, *Gonzalez*, 598 U.S. 617 (emphasis added).

[12] The number of social media-related products liability complaints has dramatically increased. "Of the 186 federal complaints involving major social media platforms since 2016, 179 of them were filed in the last 12 months. Notably, 100 of those complaints have been filed since October 3, when the Supreme Court granted certiorari in *Gonzalez*." Peter Karalis & Golriz Chrostowski, *ANALYSIS: Product Claims Spike as SCOTUS Ponders Section 230 Fix*, BLOOMBERG L. (Mar. 2, 2023, 1:01 PM), https://news.bloomberglaw.com/bloomberg-law-analysis/analysis-product-claims-spike-as-scotus-ponders-section-230-fix [https://perma.cc/6MAV-7ER5].

[13] A California federal district court in an ongoing multidistrict litigation rejected social media companies' motions to dismiss products liability claims. *See* In re Soc. Media Adolescent Addiction/Pers. Inj. Prod. Liab. Litig., No. 4:22-MD-3047, 2023 WL 7524912 (N.D. Cal. Mar. 10, 2023). The court crafted a test to determine when platforms are subject to products liability. *Id.* at *24. Evaluating whether specific "functionalities" of platforms are analogous to physical products, Judge Rogers allowed design defect claims for platforms' omissions of parental and screen-time controls, among others, to survive the motion to dismiss. *Id.* at *29–31.

[14] *See, e.g.*, Andrew D. Selbst, *Negligence and AI's Human Users*, 100 B.U. L. REV. 1315, 1375 (2019) ("Where society decides that AI is too beneficial to set aside, we will likely need a new regulatory paradigm to compensate the victims of AI's use . . . .").

[15] *See, e.g.*, Jeannie Baumann, *ChatGPT Poses New Regulatory Questions for FDA, Medical Industry*, BLOOMBERG L. (June 21, 2023), https://www.bloomberglaw.com/bloomberglawnews/health-law-and-business/X9ISGPPG000000 [https://perma.cc/Q5GL-XEL8] ("Models that update and learn, potentially could veer off from their original tasks and learn in the wrong direction if we're not keeping track. . . . [S]oftware needs to be monitored to make sure it's continuing to perform like expected.").

a products liability framework (especially coupled with a regulatory framework) can be, and how it can embrace the learning and iterative aspects of AI/ML and generative AI technologies.

II. A PRODUCTS LIABILITY CONCEPTUAL FRAMEWORK

Society faces new and uncertain risks from AI.[16] Crafting a regulatory framework for AI using a products liability lens is the most promising approach to mitigate these risks. Products liability, as I have argued before,[17] is a microcosm of how the common law evolves over time to respond to new societal risks—historically, those posed by the automobile, mass-produced goods, digital e-commerce, and now, emerging technologies like AI. At each juncture, common law judges explicitly relied on prevention and mitigation of harm, or "cheapest cost avoider" deterrence rationales, to expand products liability to address new risks and prevent them from materializing into harms and, in so doing, they recognized new forms of harms.[18]

Now is the right time to act to regulate AI. We can draw lessons from historical examples where society faced new and uncertain risks to demonstrate that, even when risks are uncertain or not entirely understood, tort liability can serve an information-production function during a "transitional period" before an ex ante regulatory scheme is in place. This approach illustrates the essential feedback loop between tort liability and regulation, highlighting the ways in which tort liability can surface essential information on which regulators can act. Moreover, there are distinct dangers of holding off on tort liability and regulation, thereby creating a regulatory void into which private actors will race.

A. The Feedback Loop Between Tort Liability and Regulation

Few would gainsay that AI can pose risks to safety, but there is a great deal of uncertainty regarding the precise nature and scale of its potential risks and harms. The Office of Management and Budget's proposed guidance on President Biden's Executive Order No. 14,110, "Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence," outlines uses of AI that are "presumed to be safety-impacting," which include "movements of a robotic appendage or body," "movements of vehicles," and "[t]he design, construction, or testing of industrial

[16] President Biden's Executive Order on Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence recognizes that "[h]arnessing AI for good and realizing its myriad benefits requires mitigating its substantial risks." Exec. Order No. 14,110, *supra* note 1, at 75191. It proceeds to lay out "a coordinated, Federal Government-wide approach" to "governing the development and use of AI safely and responsibly." *Id.*

[17] Catherine M. Sharkey, *Products Liability in the Digital Age: Online Platforms as "Cheapest Cost Avoiders"*, 73 HASTINGS L.J. 1327, 1333–34 (2022) ("[P]roducts liability is a microcosm of how the common law evolves over times, specifically, here, to respond to new societal risks—posed by the automobile, mass-produced goods, and now, digital e-commerce.").

[18] *Id.*

equipment," among many other categories.[19] These categories cover only the federal government's internal uses of AI. Regulating societal uses of AI would implicate AI-enabled medical (and other) products, which likewise risk causing potentially widespread physical and emotional harms, and generative AI models, which risk causing privacy and reputational harms by spreading misinformation.

FDA Commissioner Robert Califf recently stated that LLMs are "going to be transformative and [they've] got to be regulated . . . . We've met with the officials of the companies that are in the lead, and they want to be regulated. But even they don't really have good suggestions now on how to do it."[20] This comment beautifully encapsulates the need for the common law of tort to play a "transitional" role in addressing emerging risks when there is great uncertainty regarding how to regulate such risks. Emerging AI/ML technologies have "got to be regulated" in some fashion and yet the most sophisticated market players "don't really have good suggestions on how to do it."[21] As I have argued before,[22] particularly in areas that pose emerging and incompletely understood health and safety risks, common law tort liability holds out the potential for a dynamic regulatory response, one that incentivizes gathering additional information about potential risks and means to mitigate or adapt to these risks. Tort law, in other words, could step in not only to fill the regulatory void emerging around AI/ML technologies but, even more so, to serve an essential information-production role, which is a necessary prerequisite for regulators to design an optimal policy strategy.

Hydraulic fracturing, or fracking, provides a historical illustration of the regulatory challenge posed by a new, controversial practice with highly uncertain risks.[23] When this technology emerged, there were many unknowns about the risks it posed to the environment, to groundwater, and to personal health.[24] Crafting an ideal ex ante regulatory framework was nearly impossible given these uncertainties;

[19] OFF. OF MGMT. & BUDGET, EXEC. OFF. OF THE PRESIDENT, ADVANCING GOVERNANCE, INNOVATION, AND RISK MANAGEMENT FOR AGENCY USE OF ARTIFICIAL INTELLIGENCE (Oct. 30, 2023), https://www.whitehouse.gov/wp-content/uploads/2023/11/AI-in-Government-Memo-draft-for-public-review.pdf [https://perma.cc/9L4Z-5M7U].

[20] Baumann, *supra* note 15.

[21] *Id*.

[22] *See* Catherine M. Sharkey, *Common Law Tort as a Transitional Regulatory Regime: A New Perspective on Climate Change Litigation*, *in* CLIMATE LIBERALISM: PERSPECTIVES ON LIBERTY, PROPERTY AND POLLUTION 103 (Jonathan H. Adler ed., 2023); *see also* Catherine M. Sharkey, *Tort as Backstop to Regulation in the Face of Uncertainty*, JOTWELL (Nov. 26, 2013) (reviewing Thomas Merrill & David Schizer, *The Shale Oil and Gas Revolution, Hydraulic Fracturing, and Water Contamination: A Regulatory Strategy*, 98 MINN. L. REV. 145 (2013)) ("At the core is the need for a dynamic regulatory response, one that generates additional information about potential risks and stimulates innovation to reduce these risks.").

[23] *See* Sharkey, *Tort as Backstop to Regulation in the Face of Uncertainty*, *supra* note 22 ("Hydraulic fracturing is a controversial process whereby energy companies pump fluid into shale formations at high pressure to crack the rock and release the gas and oil trapped inside.").

[24] *Id.* (summarizing risks, including "increased air pollution, traffic and congestion (all risks associated with conventional oil and gas drilling) and, most significantly, potential contamination of groundwater (a unique risk associated with fracturing)").

moreover, overly hasty regulation might have hampered innovation or entrenched flawed incentives.[25]

Rather than settle for a potential regulatory void, tort liability can be relied on during a "transitional" period, where harmed individuals come forward to sue and thereby serve an information-production function with regard to identifying and mitigating risks. Common law tort liability thereby allows for experimentation with various risk-minimization methods and remedial strategies until optimal approaches emerge, which can then be enshrined in more uniform regulations. Thus, it plays an essential role in the transition to an ex ante regulatory scheme.

B. Products Liability as a Robust Conceptual Framework

The affirmative case for products liability as a framework to regulate AI remains to be stated; here, I take a first stab at suggesting that such a framework is robust enough to handle both the learning and iterative aspects of AI/ML. Creative scholars have put forth competing frameworks, and, in the process, some have disparaged products liability. I thus respond to some of the naysayers.

1. Building the Affirmative Case for Products Liability

a. Learning Aspect

In developing a regulatory framework that reflects the learning aspect of AI/ML, we might draw on the lessons of regulating pharmaceutical drugs. To be sure, regulators have a considerable body of information regarding the benefits and risks of drugs before releasing them onto the market. The FDA requires manufacturers to conduct three phases of clinical trials and to submit evidence of the trials to the FDA, which engages in a stringent ex ante regulatory review for safety and efficacy.[26] This premarket approval process is information-producing in

[25] Professor Thomas Merrill, commenting at the symposium, contrasted the U.S. approach to new technology of "try first, apologize later" with the European approach of "regulate first" in order to understand the consequences of the technology before releasing it. Symposium, *Accountability & Liability in Generative AI: Challenges & Perspectives*, *supra* note 4. According to Professor Merrill, worries stopped Europe from pursuing fracking, whereas the U.S. pursued fracking despite its risks—many of which never materialized. *Id.* The U.S. approach—which also encourages innovation—thus proved superior to the European one. *See id.*; *see also* Thomas Merrill & David Schizer, *The Shale Oil and Gas Revolution, Hydraulic Fracturing, and Water Contamination: A Regulatory Strategy*, 98 MINN. L. REV. 145, 215 (2013) ("When technology is new, we can predict some harm that it could cause, but not all of them, and not always with confidence about their magnitude and severity. Also, it is especially difficult to devise solutions for these harms. Effective predictions and solutions—and, thus, effective *ex ante* regulation—require experience. Without experience, we generally will be better off with some form of *ex post* regulation.").

[26] *See* FOOD & DRUG ADMIN., HOW TO STUDY AND MARKET YOUR DEVICE (Oct. 12, 2022), https://www.fda.gov/medical-devices/device-advice-comprehensive-regulatory-assistance/how-study-and-market-your-device [https://perma.cc/TP3L-HQ5H].

the sense that it mandates that manufacturers conduct clinical testing and hand over data-driven studies to the FDA for review.[27]

But, notwithstanding this rigorous ex ante premarket review, there is still relatively little known about the properties of drugs, in particular given their interactions with the much larger and more diverse populations who will consume those drugs.[28] Thus, once an FDA-approved drug is released onto the market, and is prescribed by doctors to a much larger population of individuals, we learn quite a lot about that drug's risks and benefits. A drug's interactions with multitudes of patients will produce new information regarding its properties and its safety and efficacy across different patient demographics.

Given this "learning," it is imperative to design a feedback loop whereby this new risk evidence, once unearthed, tweaks the regulatory framework. A manufacturer either has to go back to the FDA with this new evidence—whether it be about additional risks or a new use—and let the FDA weigh in on, for example, the addition of new warnings or the approval of new intended uses; or, if the manufacturer does not do so, then tort lawsuits should be available as a private mechanism of information-forcing and deterrence.[29]

In sum, a products liability framework can address the learning aspect of AI/ML. The regulatory framework for pharmaceutical drugs provides inspiration, as it demonstrates that tort liability can inform regulations by taking into account new risk evidence that becomes available only after a given product is released onto

[27] *See* Rebecca S. Eisenberg, *The Role of the FDA in Innovation Policy*, 13 MICH. TELECOMM. TECH. L. REV. 345, 347 (2007) (claiming, contrary to the common criticism that FDA regulation of drugs stifles innovation, that it "promot[es] a valuable form of pharmaceutical innovation—the development of credible information about the effects of drugs"); *see also* Catherine M. Sharkey & Daniel J. Kenny, *FDA Leads, States Must Follow*, 102 WASH. U. L. REV. (forthcoming 2024), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4753170 [https://perma.cc/KMU8-VQ2R] (highlighting the extent to which the FDA approval process generates high-quality data about medical products).

[28] One aspect of this is racial diversity. Despite the FDA's 2015 action plan to increase diversity in clinical drug trials, researchers found that it failed to improve representation of Black patients. *See* Angela K. Green et al., *Despite the FDA's Five-Year Plan, Black Patients Remain Inadequately Represented in Clinical Trials for Drugs*, 41 HEALTH AFFAIRS 368, 368 (Mar. 2022). The FDA has made improving clinical trial diversity a priority. *See* FOOD & DRUG ADMIN., DIVERSITY PLANS TO IMPROVE ENROLLMENT OF PARTICIPANTS FROM UNDERREPRESENTED RACIAL AND ETHNIC POPULATIONS IN CLINICAL TRIALS; DRAFT GUIDANCE FOR INDUSTRY; AVAILABILITY (Apr. 13, 2022), https://www.fda.gov/regulatory-information/search-fda-guidance-documents/diversity-plans-improve-enrollment-participants-underrepresented-racial-and-ethnic-populations [https://perma.cc/Y48V-RBYW].

[29] *See* Catherine M. Sharkey, *Products Liability Preemption: An Institutional Approach*, 76 GEO. WASH. L. REV. 449, 519 (2008) (setting forth the "agency reference" model for preemption that "takes into account the dynamic nature of FDA regulation of drugs by providing manufacturers with incentives to go to the FDA upon discovery of new risks"); Catherine M. Sharkey, *Federalism Accountability: "Agency-Forcing" Measures*, 58 DUKE L.J., 2125, 2129 (2009) ("The 'agency reference model' that I have proposed would be 'information-forcing' in the sense that it would require product manufacturers to come forward to the FDA with new safety risk information, including clinical studies, adverse event reports and the like, as a precondition for court determinations of preemption.").

the market. For this reason, notwithstanding the FDA's stringent ex ante premarket regulatory scrutiny, tort law still has an information-forcing and deterrence role, identifying new risks in the postmarket period.

b. Interactive Aspect

Generative AI models, such as ChatGPT, are distinct from other technologies due to their interactive nature. For example, they allow users to contribute by writing prompts. While this interactivity raises some vexing challenges, they differ more in degree than in kind from other complicated issues that products liability has tackled, such as liability for component parts manufacturers, or for manufacturers of a product to which another product is added down the line.

To take one illustrative example—high-profile, if atypical given that it involves the U.S. Supreme Court confronting state tort law—in 2019, in *Air & Liquid Systems Corp. v. DeVries*, the Court considered whether the manufacturer of a "bare-metal" product such as a turbine, blower, or pump has a duty to warn of dangers arising from the later incorporation of asbestos-laden parts into the product.[30] To answer this question, the Justices turned to first principles from tort theory. In a 6–3 decision, Justice Brett Kavanaugh, drawing heavily from Judge Guido Calabresi's "cheapest cost avoider" theory, held for the majority that the bare-metal product manufacturer did have a duty to warn, reasoning that "the product manufacturer will often be in a better position than the parts manufacturer to warn of the danger from the integrated product."[31] Justice Neil Gorsuch's dissent likewise hinged on Judge Calabresi's cheapest cost avoider theory but reasoned that the subsequent parts manufacturer is the one that "is in the best position to understand and warn users about its risks; in the language of law and economics, those who make products are generally the least-cost avoiders of their risks."[32] While the majority and the dissent disagreed as to which party—the bare-metal product manufacturer or the subsequent parts manufacturer—was in fact the cheapest cost avoider, they were unanimous in using the lens of law-and-economics, incentive-driven tort theory.[33] So, we have a well-established framework for decision-making, even though courts may disagree about how to apply it.

A similar approach could be applied to interactive AI/ML technologies if users are understood as contributing to the product down the line. Moreover, the cheapest cost avoider framework bypasses attribution issues created by the AI "black box" because it aims to reduce the societal cost of accidents. Instead of attempting to attribute each AI output to a single party, courts would focus on whether the

[30] *See* Air & Liquid Sys. Corp. v. DeVries, 586 U.S. 446, 451 (2019).

[31] *Id.* at 455.

[32] *Id.* at 460 (Gorsuch, J., dissenting).

[33] *See* Catherine M. Sharkey, *Modern Tort Theory: Preventing Harms, Not Recognizing Wrongs*, 134 HARV. L. REV. 1423, 1423–24 (2021) (describing *DeVries* and arguing that "[t]he law and economics-inspired view of tort law is ascendant, not only in the legal academy but also in the decisions of influential state and federal courts, including the U.S. Supreme Court").

interactive user or the AI developer is in the best position to mitigate or prevent harms.[34]

2. Responding to the Naysayers

Two scholars have taken up the gauntlet of resisting a products liability framework for AI. An early mover in the AI space, Professor Ryan Abbott, argues that acts of "autonomous computer tortfeasors" should be assessed under a negligence framework where the computer is playing the role of a reasonable person and automation will increase safety.[35] Also advocating for a negligence framework, but moving away from the nebulous "reasonable person" standard, Professor Bryan Choi argues that we should draw on professional malpractice law to create a framework for regulating AI developers.[36]

Abbott rejects a products liability framework for AI on two grounds. First, he argues that we should not regulate computers as products subject to strict liability (instead of negligence) because doing so would hamper socially valuable AI innovation, as "[i]t is easier to establish strict liability than negligence" since the former does not require showing intent.[37] While Abbott acknowledges that the promise and peril of automation is up for debate,[38] his advocacy for negligence liability relies on it incentivizing the adoption of iteratively safer technologies, like self-driving cars that are ten- or one hundred-times safer than human drivers.[39] Second, Abbott argues that "[c]omputers are no longer just inert tools directed by individuals. Rather, in at least some instances, computers are taking over activities once performed by people and . . . stepping into the shoes of a reasonable person."[40]

However, Abbott fails to appreciate how products liability fosters incentives to take care. For instance, it would enable liability if consumers came to expect self-driving cars to be inhumanly safe or if one could slightly reduce risk at scale

[34] *Cf.* Presidential Strategy, National Cybersecurity Strategy, 21 (Mar. 2023), https://www.whitehouse.gov/wp-content/uploads/2023/03/National-Cybersecurity-Strategy-2023.pdf [https://perma.cc/B59W-4PUH] ("Responsibility must be placed on the stakeholders most capable of taking action to prevent bad outcomes, not on the end-users that often bear the consequences of insecure software nor on the open-source developer of a component that is integrated into a commercial product.").

[35] *See* Ryan Abbott, *The Reasonable Computer: Disrupting the Paradigm of Tort Liability*, 86 GEO. WASH. L. REV. 1, 4 (2018).

[36] *See* Bryan H. Choi, *AI Malpractice*, 73 DEPAUL L. REV. 301, 301 (2024). Building on his work about "software as a profession," Choi argues that AI design should be governed by professional malpractice law because: (a) it necessarily involves subjective judgments; (b) those subjective judgments risk harmful outcomes; and (c) AI design is socially productive. *Id.* at 306.

[37] *See* Abbott, *supra* note 35, at 20–22.

[38] *See id.* at 42–43.

[39] *See id.* at 22. Other advocates of a negligence model for AI liability have expanded on this argument. *See, e.g.*, Karni A. Chagal-Feferkorn, *How Can I Tell If My Algorithm Was Reasonable?*, 27 MICH. TECH. L. REV. 213, 251 (2021) ("In the context of rapidly evolving technology, the neutrality of the reasonableness assessment has great value. When law is said to chase technological improvements, the neutral reasonableness assessment may be counted on without the expensive—and often futile—need to constantly reshape the legal framework when the technology advances.").

[40] Abbott, *supra* note 35, at 23.

through a reasonable alternative design.[41] Moreover, the dominant test for design defects is a negligence-inflected risk-utility balancing test (whether or not the more demanding version of risk-utility, namely the reasonable alternative design formulation favored by the *Restatement (Third) of Torts: Products Liability*, is adopted). And, as Professors Chinmayi Sharma and Benjamin Zipursky have highlighted in arguing for federal legislation to impose products liability rather than negligence for software security, "[b]y focusing on the fault in the product—the potential defect—rather than the faultiness of the defendant's practices, products liability chooses forward-looking proactivity rather than backward-looking conventionalism."[42]

Choi criticizes the products liability framework for "evaluating the AI system as a discrete entity or object, while minimizing the human processes behind the development of that AI system."[43] But, a products liability framework in no way precludes consideration of "human processes." For example, manufacturing defects often relate to human errors made in production. Similarly, design decisions always occur in the shadow of the law, namely, potential liability in the form of the consumer expectations and risk-utility tests, which hinge on the attitudes and product use of reasonable humans, respectively. Moreover, choices AI developers make, which range from identifying hyperparameters to cleaning training data to testing models,[44] could be cited in support of a products liability claim.

C. Why Involve a Regulator Now?

There are perils in waiting too long to regulate risks posed by emerging technologies. The saga of direct-to-consumer genetic testing provides a cautionary tale. The FDA initially held off regulating these tests—considered by some to be information that falls within the category of "the practice of medicine" not regulated by the FDA—and belatedly decided to regulate them as medical products,

[41] Abbott does nonetheless suggest that manufacturers should weigh the cost of improvements to reduce accident risk. *See id.* at 23.

[42] Chinmayi Sharma & Benjamin Zipursky, *Who's Afraid of Products Liability? Cybersecurity and the Defect Model*, LAWFARE (Oct. 19, 2023), https://www.lawfaremedia.org/article/who-s-afraid-of-products-liability-cybersecurity-and-the-defect-model [https://perma.cc/A2SG-HY75]. However, Jim Dempsey has criticized the products liability model in this context:

> If developers of software are to be held responsible for the harm caused by defects in their products, we cannot risk the impact on innovation that would result from a lack of clarity as to the standard of care. Nor, given the urgency of the cybersecurity threat, can we afford to proceed at the pace of common law, with its incremental and often inconsistent articulation by judges across many cases over many years.

Jim Dempsey, *Standards for Software Liability: Focus on the Product for Liability, Focus on the Process for Safe Harbor*, LAWFARE (Jan. 23, 2024), https://www.lawfaremedia.org/article/standards-for-software-liability-focus-on-the-product-for-liability-focus-on-the-process-for-safe-harbor [https://perma.cc/AVQ8-DRLW]. Neither Dempsey nor Sharma and Zipursky contend with the information-forcing role of common law liability, which I highlight above.

[43] Choi, *supra* note 36, at 305.

[44] *See id.* at 314–15.

given the risks they posed.[45] But, meanwhile, the private company 23&Me had amassed a large private database of highly sensitive genetic information, giving it an edge over new entrants to the market and potentially hampering future competition. Similarly, when it comes to regulating data-hungry AI, it is prudent to consider the role of the FDA not only as a safety regulator but also as a medical information regulator.[46]

III. A POINT OF DEPARTURE: AI-ENABLED MEDICAL PRODUCTS

My overarching goal in this Essay is to put forward a products liability conceptual framework for the regulation of AI. The FDA's evolving regulatory landscape for AI is an auspicious starting point for thinking about the viability of a products liability framework for the regulation of AI.

In this Part, as a jumping off point to consider potential regulation of AI, I showcase how high-risk medical products, such as medical devices that incorporate AI/ML, are subject to a nuanced regime of ex ante federal regulation by the FDA and ex post products liability. Under a products liability framework, courts would pin liability on the "cheapest cost avoider" to mitigate AI harms stemming from design defects and failures to warn. Courts would also need to consider the interplay between federal regulation and common law products liability. Finally, I consider the extent to which liability insurance may play a role in shaping the contours of the development of products liability in this realm.

A. *The FDA's Regulation of AI-Enabled Medical Devices*

In recent years, the FDA has made AI an area of particular focus:[47]

[45] While 23&Me claims to have begun its dialogue with the FDA in 2008, the FDA took until 2010 to notify it that direct-to-consumer genetic tests were medical devices that required approval. *See An Update Regarding The FDA's Letter to 23andMe*, 23&ME (Nov. 26, 2013), https://blog.23andme.com/articles/update-on-fda-letter [https://perma.cc/D9TT-KJGT]; Andrew Pollack, *F.D.A. Faults Companies on Unapproved Genetic Tests*, N.Y. TIMES (June 11, 2010), https://www.nytimes.com/2010/06/12/health/12genome.html [https://perma.cc/8MAB-DVJH]. 23&Me was slow in communicating with the FDA, and, in 2013, the FDA ordered it to stop marketing its tests. *See* Andrew Pollack, *F.D.A. Orders Genetic Testing Firm to Stop Selling DNA Analysis Service*, N.Y. TIMES (Nov. 26, 2013), https://www.nytimes.com/2013/11/26/business/fda-demands-a-halt-to-a-dna-test-kits-marketing.html [https://perma.cc/665T-ZQWU]. 23&Me's tests secured FDA approval in 2017. *See* FOOD & DRUG ADMIN., FDA ALLOWS MARKETING OF FIRST DIRECT-TO-CONSUMER TESTS THAT PROVIDE GENETIC RISK INFORMATION FOR CERTAIN CONDITIONS (Apr. 6, 2017), https://www.fda.gov/news-events/press-announcements/fda-allows-marketing-first-direct-consumer-tests-provide-genetic-risk-information-certain-conditions [https://perma.cc/UB6F-6MLK]. *See generally* Catherine M. Sharkey et al., *Regulatory and Medical Aspects of DTC Genetic Testing*, *in* CONSUMER GENETIC TECHNOLOGIES: ETHICAL AND LEGAL CONSIDERATIONS 277 (I. Glenn Cohen et al. eds., 2021); Kenneth Offit et al., *Regulation of Laboratory-Developed Tests in Preventive Oncology: Emerging Needs and Opportunities*, 41 J. CLINICAL ONCOLOGY 11 (2023).

[46] *See* Catherine M. Sharkey, *Direct-to-Consumer Genetic Testing: The FDA's Dual Role as Safety and Health Information Regulator*, 68 DEPAUL L. REV. 343, 377–79 (2019).

[47] *See* FOOD & DRUG ADMIN., FOCUS AREA: ARTIFICIAL INTELLIGENCE (Sept. 6, 2022), https://www.fda.gov/science-research/focus-areas-regulatory-science-report/focus-area-artificial-

> As technology continues to advance every aspect of health care, software incorporating artificial intelligence (AI), and specifically the subset of AI known as machine learning (ML), has become an important part of an increasing number of medical devices. One of the greatest potential benefits of AI/ML resides in its ability to create new and important insights from the vast amount of data generated during the delivery of health care every day. Digital health technologies are playing an increasingly significant role in many facets of our health and daily lives, and AI/ML is powering important advancements in this field. Ensuring that these innovative devices are safe and effective, and that they can reach their full potential to help people, is central to the FDA's public health mission.[48]

Moreover, the FDA is by no means new to regulating AI. Going back nearly three decades, the FDA has granted marketing authorizations for roughly 700 AI/ML-enabled medical devices; the vast majority of devices are radiological or cardiovascular, and the AI/ML models underpinning them range widely in complexity.[49] The FDA has also been active in developing action plans, engaging industry, issuing guidance documents, and devoting an increasing number of resources to oversee this growing area. For instance, the FDA extended its device

intelligence. Whether AI/ML is embedded in Software as a Medical Device ("SaMD"), or whether it is used in the development, clinical investigation, postmarket data analysis, or quality control of medical products and their quality management systems, the proliferation of AI/ML in the healthcare industry shows no signs of slowing down.

[48] FOOD & DRUG ADMIN., ARTIFICIAL INTELLIGENCE AND MACHINE LEARNING (AI/ML)-ENABLED MEDICAL DEVICES (Oct. 19, 2023), https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-and-machine-learning-aiml-enabled-medical-devices.

[49] *See id.* The FDA maintains a database of AI/ML-enabled medical devices, which it periodically updates, to showcase the growing number of medical devices authorized by the agency; as of the most recent update, the FDA granted marketing authorization to 692 AI/ML-enabled medical devices. *Id.*

According to the FDA: "The year-over-year increase of AI/ML-enabled devices slowed in 2021 (15%) and 2022 (14%) after an increase of 39% in 2020 (compared to 2019). Based on projected volume in 2023, the increase of AI/ML-enabled devices (compared to 2022) is expected to reach 30+%." *Id.*

"Through the end of July 2023, 79% of devices authorized in 2023 are in Radiology (85), 9% in Cardiovascular (10), 5% in Neurology (5), 4% in Gastroenterology/Urology (4), 2% in Anesthesiology (2), and 1% each in Ear, Nose and Throat (1), and Ophthalmic (1)." *Id.* "87% of devices on this list authorized in calendar year 2022 are in Radiology (122), followed by 7% in Cardiovascular (10) and 1% each in Neurology (2), Hematology (1), Gastroenterology/Urology (1), Ophthalmic (2), Clinical Chemistry (1) and Ear, Nose and Throat (1)." *Id.*

regulations to govern software intended for medical purposes[50] and released an action plan governing related AI/ML applications.[51]

The FDA has recognized that novel AI/ML-enabled medical devices warrant a novel regulatory approach.[52] Perhaps most significantly, this new paradigm must be aimed at mitigating risks over the entire lifecycle of the ever-changing devices.[53] That said, a word of caution is in order: "[t]he landscape of FDA oversight of AI/ML-enabled medical devices is dynamic, and the only constant that industry should expect in the coming years is change."[54] AI/ML regulatory principles are seemingly in greater flux and progressing at a faster rate than any other area of FDA regulatory oversight.

1. AI/ML

AI/ML-enabled medical devices are products "intended to treat, diagnose, cure, mitigate, or prevent disease" that use algorithms to perform their intended medical purposes.[55] The FDA grants marketing authorization for medical devices that pass one of three premarket review pathways: (1) the rigorous premarket approval ("PMA") process; (2) the more streamlined premarket notification ("PMN") or 510(k) clearance (for devices substantially similar to ones already on the market); and (3) de novo classification (for novel devices with low to moderate risk).[56] Initially, the FDA only approved AI/ML-enabled devices with "locked algorithms," the code and outputs of which do not change with use, as such changes would likely

[50] *See* FOOD & DRUG ADMIN., SOFTWARE AS A MEDICAL DEVICE (SAMD) (Dec. 4, 2018), https://www.fda.gov/medical-devices/digital-health-center-excellence/software-medical-device-samd.

[51] *See* FOOD & DRUG ADMIN., ARTIFICIAL INTELLIGENCE AND MACHINE LEARNING IN SOFTWARE AS A MEDICAL DEVICE (Sept. 22, 2021), https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-and-machine-learning-software-medical-device. The FDA's Center for Devices and Radiological Health (CDRH) Digital Health Center of Excellence is considering adopting lifecycle monitoring for adaptive AI/ML SaMD. *Id.*

[52] *See* FOOD & DRUG ADMIN., PROPOSED REGULATORY FRAMEWORK FOR MODIFICATIONS TO ARTIFICIAL INTELLIGENCE/MACHINE LEARNING (AI/ML)-BASED SOFTWARE AS A MEDICAL DEVICE (SAMD) - DISCUSSION PAPER AND REQUEST FOR FEEDBACK, FDA-2019-N-1185-0001 (Apr. 2, 2019), https://www.fda.gov/media/122535/download?attachment (recognizing that its "traditional paradigm of medical device regulation was not designed for adaptive AI/ML").

[53] *See* FOOD & DRUG ADMIN., MARKETING SUBMISSION RECOMMENDATIONS FOR A PREDETERMINED CHANGE CONTROL PLAN FOR ARTIFICIAL INTELLIGENCE/MACHINE LEARNING (AI/ML)-ENABLED DEVICE SOFTWARE FUNCTIONS, FDA-2022-D-2628, 1–2 (Apr. 3, 2023), https://www.fda.gov/regulatory-information/search-fda-guidance-documents/marketing-submission-recommendations-predetermined-change-control-plan-artificial.

[54] *The Evolving FDA Regulatory Landscape of Artificial Intelligence*, JD SUPRA (Mar. 29, 2023), https://www.jdsupra.com/legalnews/the-evolving-fda-regulatory-landscape-2823241/ [https://perma.cc/228R-5A5K].

[55] *See* FOOD & DRUG ADMIN., PROPOSED REGULATORY FRAMEWORK FOR MODIFICATIONS TO ARTIFICIAL INTELLIGENCE/MACHINE LEARNING (AI/ML)-BASED SOFTWARE AS A MEDICAL DEVICE (SAMD) - DISCUSSION PAPER AND REQUEST FOR FEEDBACK, *supra* note 52, at 4.

[56] *See* Sharkey & Fodouop, *supra* note 3, at 90–91, n.21.

require further regulatory approval.[57] For instance, the FDA first granted PMA approval to an AI/ML-enabled medical device in 1995 when it approved the PAPNET Testing System, a semi-automated test that assists in rescreening negative cervical smears and detecting cervical epithelial abnormalities in order to support clinician decision-making.[58] More recently, the FDA has begun approving increasingly complex and autonomous AI/ML-enabled systems that "learn" and adapt based on data inputs and analyses; for example, in 2018, it granted de novo classification to the first autonomous AI/ML-powered diagnostic device, the IDx-DR software, which uses AI for the early detection of diabetic retinopathy.[59]

The FDA has answered the question—long beguiling products liability law—whether software is a product or a service: it regulates software as a medical device (i.e., a product).[60] In April 2019, the FDA proposed regulating AI/ML-enabled software as a medical device ("SaMD") and sought feedback about issues such as premarket review of AI/ML-enabled medical devices, governance frameworks for adaptive SaMD systems, and the role of lifecycle monitoring.[61] After receiving a significant amount of feedback, the FDA published its "Artificial Intelligence/Machine Learning (AI/ML)-Based Software as a Medical Device (SaMD) Action Plan" in January 2021.[62] The five-part plan sought to create a new regulatory paradigm for AI/ML-enabled SaMD that facilitates certain postmarket changes; harmonizes principles for good AI/ML design; increases patient transparency about AI/ML data sources, training methods, risks, etc.; supports technical methods that reduce AI bias; and incorporates real-world performance

[57] *See* FOOD & DRUG ADMIN., PROPOSED REGULATORY FRAMEWORK FOR MODIFICATIONS TO ARTIFICIAL INTELLIGENCE/MACHINE LEARNING (AI/ML)-BASED SOFTWARE AS A MEDICAL DEVICE (SAMD) - DISCUSSION PAPER AND REQUEST FOR FEEDBACK, *supra* note 52, at 3.

[58] *See The Evolving FDA Regulatory Landscape of Artificial Intelligence*, *supra* note 54.

[59] *See* FOOD & DRUG ADMIN., FDA PERMITS MARKETING OF ARTIFICIAL INTELLIGENCE-BASED DEVICE TO DETECT CERTAIN DIABETES-RELATED EYE PROBLEMS (Apr. 12, 2018), https://www.fda.gov/news-events/press-announcements/fda-permits-marketing-artificial-intelligence-based-device-detect-certain-diabetes-related-eye.

[60] While there has been a debate—especially active in recent years—software has traditionally been classified as a service; moreover, the *Restatement (Third) of Torts: Products Liability* implicitly supports that view by defining products as "tangible personal property." RESTATEMENT (THIRD) OF TORTS: PRODUCTS LIABILITY § 19. But modern cases like *Lemmon v. Snap* and *In re Social Media Adolescent Addiction* suggest that software can at least sometimes be treated as a product. Lemmon v. Snap, Inc., 995 F.3d 1085, 1093 (9th Cir. 2021) (finding that Snap had a duty to design its in-app software to be a reasonably safe product); In re Soc. Media Adolescent Addiction/Pers. Inj. Prod. Liab. Litig., No. 4:22-MD-03047-YGR, 2023 WL 7524912 (N.D. Cal. Nov. 14, 2023) (holding that whether a functionality of an intangible platform is a product turns on whether that functionality is analogous to tangible personal property, and holding that functionalities such as parental controls, screen time limits, barriers to deletion, and failure to label filtered content are products).

[61] *See* FOOD & DRUG ADMIN., PROPOSED REGULATORY FRAMEWORK FOR MODIFICATIONS TO ARTIFICIAL INTELLIGENCE/MACHINE LEARNING (AI/ML)-BASED SOFTWARE AS A MEDICAL DEVICE (SAMD) - DISCUSSION PAPER AND REQUEST FOR FEEDBACK, *supra* note 52.

[62] *See* FOOD & DRUG ADMIN., ARTIFICIAL INTELLIGENCE/MACHINE LEARNING (AI/ML)-BASED SOFTWARE AS A MEDICAL DEVICE (SAMD) ACTION PLAN (Jan. 2021), https://www.fda.gov/media/145022/download.

metrics into FDA monitoring.[63] And, in September 2022, the FDA completed its Digital Health Software Precertification (Pre-Cert) Pilot Program, which aimed to inform the development of adaptive regulatory approaches, many of which are applicable to AI/ML-enabled devices, and signaled openness to further paradigm shifts.[64]

Recognizing the need to update its traditional framework for premarket review of changes to medical devices, in April 2023, the FDA outlined a "Predetermined Change Control Plan for Artificial Intelligence/Machine Learning Devices" ("PCCP"), intended to reduce regulatory hurdles for adaptive devices that can improve performance through iterative modifications.[65] The FDA authorizes manufacturers to include a PCCP in marketing submissions to "pre-specify and seek premarket authorization for intended modifications (and their method of implementation) to an [AI/ML-enabled device] without necessitating additional marketing submissions for each modification delineated and implemented in accordance with the PCCP."[66] Central to PCCPs is the requirement that developers conduct an ex ante impact assessment: "[a]n Impact Assessment, in the context of a PCCP, is the documentation of the assessment of the benefits and risks of implementing a PCCP for an [AI/ML-enabled device], as well as the mitigations of those risks."[67]

The FDA's PCCP guidance strikes a balance between allowing AI/ML-enabled medical devices to be released and to learn in the market while mitigating harm that could result from the use of unregulated AI. The PCCP framework does so by requiring a certain level of premarket review while loosening restrictions so as to allow for certain, limited adjustments to be made without requiring developers to submit new proposals too frequently.

The FDA's PCCP framework is still primarily oriented towards ex ante premarket regulatory review. I have argued that AI/ML technologies are accelerating a paradigm shift at the FDA from a model that devotes the majority of its emphasis (and resources) to premarket approval to one that focuses on ongoing

[63] *See id.* at 3–6.

[64] *See* FOOD & DRUG ADMIN., THE SOFTWARE PRECERTIFICATION (PRE-CERT) PILOT PROGRAM: TAILORED TOTAL PRODUCT LIFECYCLE APPROACHES AND KEY FINDINGS (Sept. 26, 2022), https://www.fda.gov/medical-devices/digital-health-center-excellence/digital-health-software-precertification-pre-cert-pilot-program [https://perma.cc/93KB-KKJA].

[65] *See* FOOD & DRUG ADMIN., MARKETING SUBMISSION RECOMMENDATIONS FOR A PREDETERMINED CHANGE CONTROL PLAN FOR ARTIFICIAL INTELLIGENCE/MACHINE LEARNING (AI/ML)-ENABLED DEVICE SOFTWARE FUNCTIONS, *supra* note 53. Note that there are additional elements, including cybersecurity guidelines, that are part of the FDA's overall plan but are not the focus here. *See, e.g.*, Jennifer Korn, *FDA Requires Medical Devices Be Secured Against Cyberattacks*, CNN BUSINESS (Mar. 29, 2023), https://www.cnn.com/2023/03/29/tech/fda-medical-devices-secured-cyberattacks/index.html [https://perma.cc/P3KG-JEAL].

[66] FOOD & DRUG ADMIN., MARKETING SUBMISSION RECOMMENDATIONS FOR A PREDETERMINED CHANGE CONTROL PLAN FOR ARTIFICIAL INTELLIGENCE/MACHINE LEARNING (AI/ML)-ENABLED DEVICE SOFTWARE FUNCTIONS, *supra* note 53, at 6.

[67] *Id.* at 24.

postmarket surveillance.[68] Given that AI models learn from the data they gather, it is even more important to buttress this shift to ongoing postmarket scrutiny.

2. LLMs and Generative AI

LLMs are AI models trained on sprawling text datasets, enabling them to recognize, summarize, translate, predict, and generate content; some prominent LLMs include ChatGPT, Llama, Claude, and PaLM. Generative AI is "the class of AI models that emulate the structure and characteristics of input data in order to generate derived synthetic content. This can include images, videos, audio, text, and other digital content."[69]

As of October 19, 2023, the FDA has not authorized any device that is powered by LLMs or that uses any form of generative AI.[70] But, to the extent that such models are used to treat or diagnose conditions or diseases, the FDA may soon find itself regulating generative AI. For example, a ChatGPT-like product specific to clinicians would be subject to FDA oversight if a company were selling it for the purpose of providing diagnostic, treatment, or clinical decision support.

Generative AI has the ability to increase the efficacy and safety of the medical devices and drugs put onto the market. For example, "[b]y training on data related to known drugs' chemical properties, [generative AI] can generate new candidate[] [drugs] with similar properties but different structures, potentially resulting in safer and more effective drugs."[71] It can also help doctors better understand how particular individuals will respond to a certain drug, making possible far more tailored drug design.[72] In addition, generative AI has been found to be more effective than traditional AI/ML models in certain domains, such as diagnosing the initial stages of certain heart conditions.[73] But, the downside risk is that individuals might rely on the information generated by generative AI models at the expense of ever consulting their physicians,[74] as can sometimes occur with direct-to-consumer genetic testing.[75]

[68] *See* Sharkey & Fodouop, *supra* note 3, at 87.

[69] Exec. Order No. 14,110, *supra* note 1, at 75195.

[70] *See* FOOD & DRUG ADMIN., ARTIFICIAL INTELLIGENCE AND MACHINE LEARNING (AI/ML)-ENABLED MEDICAL DEVICES, *supra* note 48.

[71] Anupriya Ramraj, *Transforming Healthcare with Generative AI*, FORBES (June 27, 2023), https://www.forbes.com/sites/forbestechcouncil/2023/06/27/transforming-healthcare-with-generative-ai/ [https://perma.cc/W455-2LRV].

[72] *See id.* ("[Generative AI] can also predict the efficacy and safety of new drug candidates by analyzing large data on drug-target interactions. Generative AI identifies patient subgroups more likely to respond to a drug by analyzing clinical data patterns, helping personalize drug therapy and improve patient outcomes."); *see also* Kevin B. Johnson et al., *Precision Medicine, AI, and the Future of Personalized Health Care*, 14 CLINICAL & TRANSLATIONAL SCI. 86 (2021).

[73] *See* Salah S. Al-Zaiti et al., *Machine Learning for ECG Diagnosis and Risk Stratification of Occlusion Myocardial Infarction*, 29 NATURE MED. 1804–13 (2023).

[74] *See* Matthew Huddle et al., *Generative AI Will Transform Health Care Sooner Than You Think*, BCG (June 22, 2023), https://www.bcg.com/publications/2023/how-generative-ai-is-transforming-health-care-sooner-than-expected [https://perma.cc/M9SC-5ZZW].

[75] *See* Sharkey et al., *supra* note 45, at 285–86; Sharkey, *supra* note 46, at 370–71.

Experts disagree about whether generative AI poses a challenge different in kind or merely in degree from other technologies. One view places perhaps undue emphasis on what it deems categorical differences: "[i]f you take, for example, the 100 or 150 or so software as a medical device approvals that have gone through the FDA, what's going on with many of them is a supervised learning algorithm," whereas, "[w]hen it comes to large language models, they're quite different."[76] "'These models are more flexible, maybe than some of those supervised learning classifiers, in that they can do a number of different things,' such as answering questions about how well certain therapies worked or summarizing [medical research] papers in conversational language."[77] The contrary view—which seems to better reflect the reality of the evolution of AI technologies to date—holds that the principles embedded in the FDA's SaMD framework would transfer readily to ChatGPT: "[a]s we become more familiar with the technology, there might be changes that are made specifically for this. But it strikes me that a lot of it is going to be following on the principles that already exist."[78]

B. Products Liability and Federal Preemption

Federal preemption is an affirmative defense to a state products liability action, whereby the defendant claims that a federal statute and regulatory scheme reigns supreme and ousts conflicting state law. The courts' preemption determinations—namely the extent to which federal regulation ousts conflicting state tort actions—are not only a matter of conventional statutory interpretation, but are also highly dependent upon the stringency of the ex ante review by the relevant regulator (in this case the FDA).

With regard to products liability actions involving prescription medical products, the courts' starting point is the Federal Food, Drug and Cosmetic Act ("FDCA"), enacted by Congress, which includes an express preemption provision governing medical devices, but not drugs. But, even when faced with preemption decisions pertaining to medical devices, courts have looked beyond the statutory language of the FDCA—and the preemption provision enacted as part of the Medical Device Amendment Act of 1976—to consider the stringency of the regulatory review conducted by the FDA. Thus, in *Riegel v. Medtronic*, the U.S. Supreme Court decided that the vast majority of products liability claims arising from high-risk medical devices that were approved by the FDA via the PMA pathway—which entails rigorous scrutiny of empirical evidence resulting in a determination of safety and efficacy—are preempted.[79] On the other hand, claims

[76] Baumann, *supra* note 15 (quoting Alan Karthikesalingam, research lead at Google Health).

[77] *Id.*

[78] *Id.* (quoting David Peloquin, health care attorney at Ropes & Gray LLP).

[79] *See* Catherine M. Sharkey, *What* Riegel *Portends for FDA Preemption*, 103 NW. U. L. COLLOQUY 437, 445 (2009) ("*Riegel* is rife with details from the FDA's regulatory review process—though their precise legal effect, given the Court's insistence on governing statutory text, is rather opaque. The Court drilled down to the details of the FDA's review process, repeatedly stressing the 'rigorous' nature of its premarket approval (PMA) process for medical devices. This PMA process demands considerable resources and manpower hours, culminating in the FDA's determination of 'reasonable assurance' of the medical device's 'safety and effectiveness.'").

arising from medical devices that are approved via the more lax PMN pathway—whereby the FDA gives a more streamlined review, resulting in a decision that the medical device is "substantially similar" to one existing on the market—are not preempted.[80]

In prior work, I have explored at length the interplay between federal regulation and products liability for mitigating risks posed by pharmaceutical drugs and medical devices.[81] I have set forth a proposed "agency reference" model, whereby courts look to and scrutinize input from the FDA regarding the extent to which a products liability design defect or failure-to-warn action revisits the risk-benefit calculus the FDA made in either its approval decision or subsequent actions: if the action raises new risk evidence not yet considered by the FDA, the products liability action should proceed, and, if it does not, the action should be preempted.[82] In this way, state tort actions serve an information-forcing role, surfacing evidence of new risks that have come to light post-approval, after a drug or medical device has been used by patients.

The agency reference model can be applied to the realm of AI-enabled medical devices. The stringent ex ante PMA regulatory pathway, by which the FDA would approve high-risk (i.e., Class III) AI-enabled devices, would preempt products liability claims. But the evolution of SaMD and PCCP guidelines (discussed above)[83] suggests that the FDA is moving towards a minimally burdensome ex ante regulatory approach. Under the agency reference model, this relatively light-touch ex ante review should leave ample room for products liability claims to survive.

C. Liability Insurance

Finally, by invoking a products liability framework for AI, I mean also to invoke the potential role to be played by liability insurance. The availability of liability insurance has figured prominently in the development and expansion of products liability for three-quarters of a century and has been expressly stated to be a relevant consideration in numerous significant decisions.[84]

[80] *See id.* ("Premarket notification is a streamlined process, which is completed in an average of 20 hours (as compared to the PMA's 1,200-hour average). So, measured by average manpower hours, this type of regulatory review is sixty times more lax. Even more germane is the distinction the Court drew between the FDA's premarket notification 'equivalence' review, which essentially 'grandfathers' devices that are equivalent to those existing on the market at the time of the MDA's enactment, and the full-blown PMA 'safety' review.").

[81] *See id.*

[82] *See id.* at 441 ("The basic question at the core of implied conflict preemption inquiries is whether or not state common law actions are irreconcilable with, or would stand as an obstacle to, or frustrate, the command of federal regulatory directives and goals. To answer this question, courts need a fine-grained account of the precise regulatory review conducted by the agency and evidence as to its compatibility with state law tort claims. The agency reference model aims, as a general matter, to facilitate input from federal agencies on these issues.").

[83] *See supra* Part III.A.

[84] *See* Kenneth S. Abraham & Catherine M. Sharkey, *The Glaring Gap in Tort Theory*, 133 YALE L.J. (forthcoming 2024), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4585790 [https://perma.cc/HK5W-SHKE].

Liability insurance, moreover, can play a critical role in terms of fostering the development of standards that can be deployed to mitigate or prevent harms, including in this brave new world of AI harms. Liability insurers can aggregate risk-related information obtained about the expanding universe of policyholders as part of the process of underwriting and premium-setting.[85] As I have explored with Professor Kenneth Abraham in depth,[86] when liability insurers have more information than policyholders regarding certain methods of reducing risk, those insurers have an incentive to communicate that information to policyholders through "coaching" policyholders, either as part of their interactions over premiums or on an ongoing basis.[87]

As mentioned above, the IDx-DR software, which uses AI for the early detection of diabetic retinopathy, was the first autonomous AI/ML-powered diagnostic device to receive FDA approval.[88] Digital Diagnostic, the designer of IDx-DR, carries malpractice insurance and indemnifies physicians using its devices.[89] It is curious that the form of liability insurance is malpractice insurance, as distinct from products liability insurance.[90] But it does nonetheless suggest that AI-enabled software designers are anticipating potential tort liability.

It remains to be seen whether a third-party liability insurance market will emerge in response to the threat of products liability actions for AI/ML harms, in which case insurers would engage in further risk management and exert more potent regulatory control.[91] In this way, liability insurance joins in as part of the tort liability-regulatory feedback loop.

[85] *See* Peter Z. Grossman et al., *Uncertainty, Insurance and the Learned Hand Formula*, 5 L., PROBABILITY & RISK 1, 1 (2005) (arguing how insurance provides necessary, aggregated guidance to courts and litigants regarding the danger of certain activities, optimal precautions to take, and therefore—under the Hand Formula—the threshold below which a lack of caution becomes legally cognizable negligence); *see also* OMRI BEN-SHAHAR & ARIEL PORAT, PERSONALIZED LAW: DIFFERENT RULES FOR DIFFERENT PEOPLE 32 (2020) ("Insurance services are perhaps the pioneers in personalized treatments, which is not surprising given the wealth of personal data the industry has. Insurers 'rate' policyholders . . . . With the advent of digital data collection, the personalization has become more intensive, focusing on policyholders' conduct.").

[86] *See* Abraham & Sharkey, *supra* note 84.

[87] The greater the threat of tort liability, the more cost-effective premium refinement and individualized coaching will be. *Id.*

[88] *See* FOOD & DRUG ADMIN., FDA PERMITS MARKETING OF ARTIFICIAL INTELLIGENCE-BASED DEVICE TO DETECT CERTAIN DIABETES-RELATED EYE PROBLEMS, *supra* note 59 & accompanying text.

[89] *See* W. Nicholson Price & Glenn Cohen, *Locating Liability for Medical AI*, 73 DEPAUL L. REV. 339, 342 (2024); Michael D. Abramoff et al., *Lessons Learned About Autonomous AI: Finding a Safe, Efficacious, and Ethical Path Through the Development Process*, 214 AM. J. OPHTHALMOLOGY 134, 139 (June 2020) ("Just like a physician that would be held legally responsible for his or her diagnosis or other clinical decision, IDx, as creators of autonomous AI products, assume similar liability and have obtained medical malpractice insurance.").

[90] Commercial general liability ("CGL") policies cover liability for damages imposed because of bodily injury and property damage caused by anything that happens accidentally. *See* Abraham & Sharkey, *supra* note 84. Thus, when strict products liability was adopted in 1965 by the *Restatement (Second) of Torts* § 402A, CGL policies automatically covered it.

[91] Cyber liability is a good example:

IV. CONCLUSION

A products liability framework provides an auspicious model for regulating AI harms. Such a framework is robust enough to tackle the challenges posed by the adaptive learning aspect of ML as well as the interactive aspects of AI and generative AI technologies. Chief among the advantages of products liability is capturing the feedback loop between tort liability and regulation, which is well illustrated by the evolving regulatory framework for FDA-approved, AI-enabled medical devices. The feedback loop of an agency regulator and products liability allows for experimenting with standards before the agency regulator settles on an optimal one, which in turn will preempt certain tort actions under something like the agency reference model. Finally, liability insurance is a key ingredient of any form of tort liability, and, here too, it offers promise in terms of risk management and the development of safety standards.

> Initially, the demand for liability insurance was low because most defendants assumed they were protected from tort liability by the "no duty" economic loss rule. Some courts then began to impose duties to protect against security breaches. When this form of liability began to emerge, not all CGL policyholders wanted to purchase coverage against it. At first, an optional endorsement providing coverage was developed. But after a time, this approach became unsuitable and freestanding cyber insurance was developed and marketed. The process of developing such coverage is a major undertaking and requires that sufficient demand for the coverage exist or be anticipated in order to justify the investment, long before any premiums are earned. Cyber insurance premiums (typically for a combination of liability insurance and first-party coverage) grew 61% in 2021. But it remains to be seen whether such a thin sliver of new liability will satisfy this precondition.

Abraham & Sharkey, *supra* note 84.