

Navigating the Frontiers of MedTech

Benjamin Sundholm*

The medical community is increasingly focused on the rise of adaptive and opaque artificial intelligence tools. These systems improve over time but produce results through complex calculations that are difficult for humans to fully understand. Although promising, these features challenge existing legal doctrines. To date, efforts to overcome these challenges have been too fragmented and limited in scope. As a result, the full promise of adaptive and opaque artificial intelligence systems remains untapped.

This Article proposes a framework for unleashing the potential of these systems through a blend of forward- and backward-looking reforms. The ex ante feature of my framework calls on the U.S. Food and Drug Administration revise its existing regulatory approach, which is too rigid and retrospective. Alternatively, the agency should adopt a more flexible and forward-looking model. The ex post feature of my proposal recognizes that, pursuant to Supreme Court precedent, medical tools approved through the sort of regulatory model I recommend are not guaranteed protection from civil liability. The specter of liability is a problem because existing tort doctrines are ill-equipped to handle harms caused by opaque artificial intelligence systems. To address these challenges, I propose leveraging common enterprise liability. Doing so would complement the ex ante regulatory reforms I suggest by ensuring the availability of an adequate ex post framework for responding to harms resulting from the use of these promising artificial intelligence systems.

* Assistant Professor, St. John's School of Law. This Article has benefitted from feedback received at the Seventh Annual Junior Faculty Forum for Law and STEM at Penn Carey Law School, the 2024 Health Law Works-in-Progress Retreat at Seton Hall Law, the 2023 Health Law Professors Conference, faculty workshops at Tulane Law, the 2023 Conference on the Governance of Emerging Technologies and Science at ASU Law, the 2023 Legal Scholars Roundtable on Artificial Intelligence at Emory Law, and the 2022 Weill Cornell Medicine-Baylor Medicine Medical Ethics Consortium. I also thank Charlotte Alexander, Yonathan Arbel, Jane Bambauer, Oren Bracha, Bryan Choi, Carl Coleman, Rebecca Crootof, Greg Day, Deven Desai, Anjali Deshmukh, Adam Feibelman, Brian Frye, Sharona Hoffman, Gus Hurwitz, Mark Lemley, Anat Lior, Ann Lipton, Gennie Mansi, Lisa Larrimore Ouellette, W. Nicholson Price II, Christopher Robertson, Ana Santos Rutschman, Matt Sag, Peter Salib, Jeremy Sheff, David Simon, Charlotte Tschider, Robin West, and Christopher S. Yoo.

INTRODUCTION.....	1405
I. TRADITIONAL U.S. FOOD AND DRUG ADMINISTRATION	
REGULATION.....	1410
A. <i>Shortcomings in the Traditional Regulatory Framework</i>	1414
B. <i>Outlining a Forward-Looking Regulatory Framework</i>	1422
1. People.....	1423
2. Processes	1425
3. Production	1429
C. <i>Statutory Limitations</i>	1433
II. THE CIVIL LIABILITY PROBLEM	1437
A. <i>Agency and Corporate Liability</i>	1441
B. <i>Products Liability</i>	1445
C. <i>Medical Malpractice</i>	1448
III. THE PROMISE OF ENTERPRISE LIABILITY	1451
A. <i>Traditional Versions of Enterprise Liability</i>	1451
B. <i>Common Enterprise Liability</i>	1459
C. <i>Common Enterprise Liability for Adaptive & Opaque Systems</i>	1461
IV. CONCLUSION.....	1468

INTRODUCTION

In general terms, artificial intelligence can be defined as a constellation of capabilities and technologies enabling a computer system to accomplish tasks that ordinarily require the application of human intelligence.¹ This technology has been used in medical care for decades.² Although early clinical and research applications for artificial intelligence were relatively narrow, its use has greatly expanded over the years as these systems have become more sophisticated.³ For example, powerful artificial intelligence tools are now being used to quickly review medical images to identify disease,⁴ detect irregular heart rhythms,⁵ analyze speech patterns for indications of neurological conditions,⁶ operate robotics to perform surgical tasks,⁷ analyze large datasets to streamline medical research,⁸ and more.

In recent years, an interdisciplinary community of healthcare scholars and practitioners have been buzzing about the emergence of a new and particularly promising sort of artificial intelligence.⁹ Unlike artificial

1. See, e.g., Ryan Calo, *Artificial Intelligence Policy: A Primer and Roadmap*, 51 U.C. DAVIS L. REV. 399, 404 (2017).

2. See, e.g., Aliza Becker, *Artificial Intelligence in Medicine: What Is It Doing for Us Today?*, 8 HEALTH POL'Y & TECH. 198, 198 (2019).

3. See, e.g., Marly van Assen et al., *Artificial Intelligence: A Century-Old Story*, in ARTIFICIAL INTELLIGENCE IN CARDIOTHORACIC IMAGING 3, 9–11 (Carlo N. De Cecco et al. eds., 2022).

4. See, e.g., Varun Gulshan et al., *Development and Validation of a Deep Learning Algorithm for Detection of Diabetic Retinopathy in Retinal Fundus Photographs*, 316 JAMA 2402, 2407 (2016); Andre Esteva et al., *Dermatologist-Level Classification of Skin Cancer with Deep Neural Networks*, 542 NATURE 115, 118 (2017).

5. See, e.g., Venkat D. Nagarajan et al., *Artificial Intelligence in the Diagnosis and Management of Arrhythmias*, 38 EUR. HEART J. 3904, 3914 (2021); Zachi I. Attia et al., *An Artificial Intelligence-Enabled ECG Algorithm for the Identification of Patients with Atrial Fibrillation During Sinus Rhythm: A Retrospective Analysis of Outcome Prediction*, 394 LANCET 861, 861 (2019).

6. See, e.g., Björn Herrmann, *The Perception of Artificial-Intelligence (AI) Based Synthesized Speech in Younger and Older Adults*, 26 INT'L J. SPEECH TECH. 395, 412 (2023); Georgia Zellou et al., *Age- and Gender-Related Differences in Speech Alignment Toward Humans and Voice-AI*, 5 FRONTIERS COMMUN., 600361, at 9 (Jan. 20, 2021), <https://www.frontiersin.org/journals/communication/articles/10.3389/fcomm.2020.600361/full> [https://perma.cc/5PZG-SDW9].

7. See, e.g., Sandip Panesar et al., *Artificial Intelligence and the Future of Surgical Robotics*, 270 ANNALS SURGERY 223, 223 (2019); Chi Zhang et al., *The Integration of Artificial Intelligence in Robotic Surgery: A Narrative Review*, 176 SURGERY 552, 556 (2024).

8. See, e.g., Pranav Rajpurkar et al., *AI in Health and Medicine*, 28 NATURE MED. 31, 36 (2022).

9. See, e.g., Sara Gerke et al., *The Need for a System View to Regulate Artificial Intelligence/Machine Learning-Based Software as Medical Device*, 3 NPJ DIGIT. MED., 53 (Apr. 7,

intelligence tools currently used in medicine, the latest and most powerful systems are adaptive and opaque.¹⁰ Adaptive systems leverage a type of artificial intelligence known as deep learning, which enables these tools to teach themselves to continually improve their ability to perform certain tasks.¹¹ These systems are opaque because the calculations they use to improve their capabilities are too complex and enigmatic to be fully understood by humans.¹²

Adaptive and opaque artificial intelligence systems have enormous potential to improve healthcare in a variety of ways, from optimizing workflows, reducing spending, and enhancing treatment quality to increasing diagnostic accuracy, expanding access to treatment, and more.¹³ However, at present, they are not being leveraged in the care of patients.¹⁴ This is in large

2020), <https://www.nature.com/articles/s41746-020-0262-2> [<https://perma.cc/5KF6-X6LW>]; Boris Babic et al., *Algorithms on Regulatory Lockdown in Medicine*, 366 SCIENCE 1202 (2019); Jessa Boubker, *When Medical Devices Have a Mind of Their Own: The Challenges of Regulating Artificial Intelligence*, 47 AM. J.L. & MED. 427, 428 (2021); Sam Surette, *How Should the FDA Regulate Adaptive AI, Software that Designs Itself?*, STAT (Oct. 2, 2020), <https://www.statnews.com/2020/10/02/how-should-fda-regulate-adaptive-ai> [<https://perma.cc/7QHR-KTV5>]; U.S. 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 3 n.7 (2019), <https://www.fda.gov/media/122535/download> [<https://perma.cc/BES6-BME3>]; W. Nicholson Price II, *Black-Box Medicine*, 28 HARV. J.L. & TECH. 420, 421 (2015).

10. See Daniele Ravi et al., *Deep Learning for Health Informatics*, 21 J. BIOMED. & HEALTH INFORMATICS 4, 4 (2017); Vivek Kaul et al., *History of Artificial Intelligence in Medicine*, 92 GASTROINTESTINAL ENDOSCOPY 807, 807 (2020); Harry Surden & Mary-Anne Williams, *Technological Opacity, Predictability, and Self-Driving Cars*, 38 CARDOZO L. REV. 121, 147–48 (2016); Harry Surden, *Machine Learning and Law*, 89 WASH. L. REV. 87, 88 (2014); George Maliha et al., *Artificial Intelligence and Liability in Medicine: Balancing Safety and Innovation*, 99 MILBANK Q. 629, 629–30 (2021).

11. See, e.g., Geetha Mahadevaiah et al., *Artificial Intelligence-Based Clinical Decision Support in Modern Medical Physics: Selection, Acceptance, Commissioning, and Quality Assurance*, 47 MED. PHYSICS e228, e228–29 (2020); PRAC. L. HEALTH CARE, KEY AI CONTRACTING ISSUES FOR HEALTH CARE PROVIDERS, Westlaw W-026-1072 (database updated 2024); Ravi et al., *supra* note 10, at 4; Kaul et al., *supra* note 10, at 807; Surden & Williams, *supra* note 10, at 147–48; Surden, *supra* note 10, at 88; Maliha et al., *supra* note 10, at 630.

12. See, e.g., Juan Manuel Durán & Karin Rolanda Jongsma, *Who Is Afraid of Black Box Algorithms? On the Epistemological and Ethical Basis of Trust in Medical AI*, 47 J. MED. ETHICS 329, 329 (2021); Michael Lang et al., *Artificial Intelligence in Cardiovascular Imaging: “Unexplainable” Legal and Ethical Challenges?*, 38 CAN. J. CARDIOLOGY 225, 228 (2022).

13. Gerke et al., *supra* note 9, at 1; see also Nathan Cortez, *The Mobile Health Revolution?*, 47 U.C. DAVIS L. REV. 1173, 1190–99 (2014).

14. See, e.g., Gerke et al., *supra* note 9, at 1; Boubker, *supra* note 9, at 428–30; Surette, *supra* note 9; Anish Bhardwaj, *Promise and Provisos of Artificial Intelligence and Machine Learning in Healthcare*, 14 J. HEALTHCARE LEADERSHIP 113, 116 (2022); Daniel A. Hashimoto

part due to the ill fit between these systems and existing regulatory and liability frameworks.¹⁵

On the regulatory front, the U.S. Food and Drug Administration (“FDA”) is the agency with primary authority to ensure the safety and effectiveness of medical technologies, including computer systems leveraging artificial intelligence.¹⁶ Traditionally, before approving medical tools for use in clinical workflows, the FDA requires developers of medical products—including medical software—subject to regulation to provide the agency with a backward-looking description of the medical tool’s past performance.¹⁷ The agency uses these descriptions as evidence of the tool’s ability to perform safely and effectively going forward.¹⁸ This regulatory approach works well for traditional medical products, but it is ill-suited for adaptive and opaque tools.¹⁹ This is so because pursuant to the FDA’s traditional regulatory framework, each change undergone by adaptive systems triggers an additional round of review.²⁰ These additional review requirements are cumbersome both for the agency and developers of adaptive systems.²¹ As a result of the burdens imposed by the current regulatory environment, many artificial intelligence developers are likely to opt for creating locked systems rather than adaptive ones, depriving clinicians, patients, and the broader public of the potential benefits these systems could yield.²²

et al., *Artificial Intelligence in Surgery: Promise and Perils*, 268 ANNALS SURGERY 70, 71 (2018); A. Michael Froomkin et al., *When AIs Outperform Doctors: Confronting the Challenges of a Tort-Induced Over-Reliance on Machine Learning*, 61 ARIZ. L. REV. 33, 50, 64, 66–67 (2019).

15. See, e.g., Walid Ben Ali et al., *Implementing Machine Learning in Interventional Cardiology: The Benefits Are Worth the Trouble*, 8 FRONTIERS CARDIOVASCULAR MED., 711401, at 12 (Dec. 8, 2024), <https://www.frontiersin.org/journals/cardiovascular-medicine/articles/10.3389/fcvm.2021.711401/full> [<https://perma.cc/BCC4-AJSJ>]; Gerke et al., *supra* note 9, at 1; Babic et al., *supra* note 9, at 1202–04; Fei Wang et al., *Deep Learning in Medicine—Promise, Progress, and Challenges*, 179 JAMA INTERNAL MED. 293, 294 (2019).

16. See Jeffrey M. Senger & Patrick O’Leary, *Big Data and Human Medical Judgment: Regulating Next-Generation Clinical Decision Support*, in BIG DATA, HEALTH LAW, AND BIOETHICS 283, 285 (I. Glenn Cohen et al. eds., 2018); W. Nicholson Price II, *Regulating Black-Box Medicine*, 116 MICH. L. REV. 421, 437 & n.77 (2017); Cortez, *supra* note 13, at 1200–05.

17. See, e.g., Price, *supra* note 16, at 437–40; Cortez, *supra* note 13, at 1200–11.

18. See Price, *supra* note 16, at 437–38; Cortez, *supra* note 13, at 1200–11.

19. See Senger & O’Leary, *supra* note 16, at 292.

20. Gerke et al., *supra* note 9, at 1.

21. See *id.*

22. See *id.*

As Part I will discuss in greater detail, the FDA has recognized this problem and begun developing a response. In April 2023, the agency released draft guidance detailing its plan to approve adaptive artificial intelligence tools if developers provide a sufficiently detailed “predetermined change control plan.” U.S. FOOD & DRUG ADMIN., *MARKETING SUBMISSION*

In response to these regulatory challenges, I suggest that the FDA abandon its traditional backward-looking regulatory scheme. Alternatively, I draw on proposals laid out by Sara Gerke, Boris Babic, Glenn Cohen, Nicholson Price, and others in recommending that the FDA adopt a forward-looking framework.²³ Specifically, I suggest that the agency condition its approval of adaptive artificial intelligence systems on receipt of a forward-looking description of the processes interested parties will follow when using adaptive artificial intelligence tools in the care of patients.²⁴ In so doing, the agency can unleash the enormous potential of these systems in a way that is safe and effective.

If the FDA adopts my forward-looking framework, reforms to existing tort doctrines will be needed. This is so because, pursuant to Supreme Court precedent, protection from civil liability is not guaranteed for medical tools approved by the sort of regulatory model I recommend.²⁵ The specter of civil liability is a problem because traditional tort doctrines are ill-suited to address harms resulting from using opaque artificial intelligence systems in the care of patients. Ordinarily, if a medical tool causes patient harm, it is possible to determine whether the injury is the result of an engineering defect, clinician negligence, or the conduct of some other party.²⁶ But as Nicholson Price and others have explained, the opacity of certain adaptive artificial intelligence systems makes it difficult—if not impossible—to identify precisely what went wrong and who is responsible.²⁷ This indeterminacy strains the ability

RECOMMENDATIONS FOR A PREDETERMINED CHANGE CONTROL PLAN FOR ARTIFICIAL INTELLIGENCE/MACHINE LEARNING (AI/ML)-ENABLED DEVICE SOFTWARE FUNCTIONS: DRAFT GUIDANCE FOR INDUSTRY AND FOOD AND DRUG ADMINISTRATION STAFF 1–29 (2023). According to the FDA, these plans will illustrate the sort of changes developers expect the system to undergo and detail the steps that will be taken to ensure safety and effectiveness. *Id.*

Although this proposal from the FDA is a step in the right direction, the opacity of the most powerful adaptive systems poses a serious problem. The opacity of these systems is due to the fact that they produce their outputs by making calculations that are too complex to be understood by humans. *See* Price, *supra* note 16, at 430; W. Nicholson Price II & Arti K. Rai, *Clearing Opacity Through Machine Learning*, 106 IOWA L. REV. 775, 784–85 (2021). Given the opacity of these calculations, any attempt by developers to provide the sort of sufficiently detailed predetermined change control plan pursuant to the agency’s 2023 guidance will either be unhelpfully vague or unreliable. Babic et al., *supra* note 9, at 1203.

23. *See* Maliha et al., *supra* note 10, at 629–30; Gerke et al., *supra* note 9, at 1; Babic et al., *supra* note 9, at 1203; W. Nicholson Price II & I. Glenn Cohen, *Locating Liability for Medical AI*, 73 DEPAUL L. REV. 339, 341–42 (2024).

24. Gerke et al., *supra* note 9, at 1.

25. *See infra* Part II.

26. *See infra* Part II.

27. *See infra* Part II; Price, *supra* note 16, at 433–34.

of traditional tort doctrines to assign liability for injuries caused by black-box systems.²⁸

In response to these challenges, I suggest applying common enterprise liability to the use of adaptive and opaque artificial intelligence systems in the care of patients.²⁹ Pursuant to this doctrine, joint and several liability is imposed on all those participating in a common aim that results in harm.³⁰ Because physicians, health systems, and software developers are engaged in the common objective of leveraging artificial intelligence technology to treat patients, it seems appropriate to impose joint liability on them. Common enterprise liability can overcome the inability of traditional tort doctrines to assign liability for injuries caused by black-box systems. In so doing, the doctrine can provide a backstop to address any injuries that aren't prevented through the forward-looking regulatory scheme that I recommend the FDA adopt.³¹

Of course, I am not the first to identify and respond to the potential benefits and risks of adaptive and opaque artificial intelligence systems.³² The primary contribution of this Article is its recognition that the safe and effective use of adaptive and opaque artificial intelligence tools in medicine requires a blend of forward- and backward-looking reforms to doctrines in both public and

28. See, e.g., Maliha et al., *supra* note 10, at 629–30; Price, *supra* note 16, at 434 & n.63; W. Nicholson Price II, *Medical Malpractice and Black-Box Medicine*, in *BIG DATA, HEALTH LAW, AND BIOETHICS* 295, 298–99 (I. Glenn Cohen et al. eds., 2018). Medical devices that are subject to, and have completed, the FDA's most stringent backward-looking review processes are exempt from civil liability. Charlotte A. Tschider, *Medical Device Artificial Intelligence: The New Tort Frontier*, 46 *BYU L. REV.* 1551, 1573 (2021); Barbara J. Evans, *The Streetlight Effect: Regulating Genomics Where the Light Is*, 48 *J.L. MED. & ETHICS* 105, 112 (2020). But the same protection is not afforded to medical software approved through a forward-looking regulatory regime. Tschider, *supra*, at 1575; Evans, *supra*, at 112. So, any patient harm resulting from the use of adaptive and opaque artificial intelligence systems could trigger tort liability for developers, clinicians, or health systems. See Evans, *supra*, at 112; Barbara J. Evans & Frank Pasquale, *Product Liability Suits for FDA-Regulated AI/ML Software*, in *THE FUTURE OF MEDICAL DEVICE REGULATION: INNOVATION AND REGULATION* 22, 30 (I. Glenn Cohen et al. eds., 2022).

29. See generally Maliha et al., *supra* note 10, at 629–30 (examining how AI/ML intertwines liability of physicians and other actors); Benny Chan, *Applying a Common Enterprise Theory of Liability to Clinical AI Systems*, 47 *AM. J.L. & MED.* 351 (2021) (discussing how common enterprise theory functions when applied to AI); Scott J. Schweikart, *Who Will Be Liable for Medical Malpractice in the Future? How the Use of Artificial Intelligence in Medicine Will Shape Medical Tort Law*, *MINN. J.L. SCI. & TECH.*, Mar. 2021, at 1 (assessing different tort paradigms for liability of AI).

30. See David C. Vladeck, *Machines Without Principals: Liability Rules and Artificial Intelligence*, 89 *WASH. L. REV.* 117, 129 n.39 (2014).

31. See Tschider, *supra* note 28, at 1570–73.

32. See, e.g., Maliha et al., *supra* note 10, at 629–30; Tschider, *supra* note 28, at 1551; Evans, *supra* note 28, at 105; Gerke et al., *supra* note 9, at 1; Babic et al., *supra* note 9, at 1202–04.

private law. These respective domains are siblings in a family of legal institutions capable of responding to thorny questions raised by the use of new technologies in the medical context and beyond.³³ Sometimes, the difficulties associated with new technologies are discreet enough to be adequately addressed through targeted reforms to doctrines in either public or private law. Other times, though, the challenges these systems pose are systemic and pervasive enough to warrant a more comprehensive set of reforms to multiple legal institutions, including those that are sometimes thought of as wholly separate and autonomous domains. I suggest that the latter approach is appropriate given the uniquely thorny questions raised by using adaptive and opaque artificial intelligence systems in the care of patients.

This Article proceeds in three parts. Part I illustrates the FDA's traditional backward-looking regulatory framework and explains why this conventional approach does not map on well to adaptive artificial intelligence systems. That Part will also outline the forward-looking regulatory framework that I suggest should replace the agency's existing approach. Next, Part II explains why traditional tort doctrines are poorly equipped to address harms resulting from the use of opaque artificial intelligence tools. Finally, Part III illustrates why common enterprise liability is uniquely capable of complementing a forward-looking regulatory framework.

I. TRADITIONAL U.S. FOOD AND DRUG ADMINISTRATION REGULATION

The FDA's regulatory scope applies to medical devices or products, broadly defined to include any "instrument, apparatus, implement, machine, [or] contrivance" used to treat, diagnose, or prevent various conditions.³⁴ Pursuant to the Medical Device Amendments of 1976, the agency classifies devices based on risk.³⁵ Those posing the lowest risk fall in Class I and are only subject to what are known as general controls, such as adverse-event reporting, registration, and listing requirements.³⁶ Class II concerns moderate-risk devices, which are subject to both general and "special controls."³⁷ The latter sort of controls are usually device-specific and require

33. See Gregory C. Keating, *Is Tort Law "Private"?*, in *CIVIL WRONGS AND JUSTICE IN PRIVATE LAW* 361–62 (Paul B. Miller & John Oberdiek eds., 2020).

34. 21 U.S.C. § 321(h).

35. Medical Device Amendments of 1976, Pub. L. No. 94-295, § 2, 90 Stat. 539, 540 (codified as amended at 21 U.S.C. § 360c).

36. See 21 U.S.C. § 360c(a)(1)(A).

37. *Id.* § 360c(a)(1)(B).

the provision of information (e.g., performance standards or post-market surveillance strategies) sufficient to demonstrate the device's safety and effectiveness.³⁸ High-risk devices fall in Class III, and most must be approved through a full pre-market pathway, which is the most rigorous submission type.³⁹

The years leading up to and following enactment of the 1976 amendments were marked by a rapidly growing interest in computerized medical devices.⁴⁰

38. *Id.*

39. *Id.* § 360c(a)(1)(C). Full premarket pathway approval requires a determination by the FDA that there is valid scientific evidence that the tool is safe and effective. *Id.* According to the agency, valid scientific evidence is defined as resulting

from well-controlled investigations, partially controlled studies, studies and objective trials without matched controls, well-documented case histories conducted by qualified experts, and reports of significant human experience with a marketed device, from which it can fairly and responsibly be concluded by qualified experts that there is reasonable assurance of the safety and effectiveness of a device under its conditions of use.

21 C.F.R. § 860.7(c)(2) (2024). In addition to analyzing data supporting claims of safety and effectiveness, the FDA also examines manufacturing facilities to ensure compliance with relevant process requirements. 21 U.S.C. § 360e(d)(2).

There are three exceptions to the requirement that Class III devices complete the full premarket approval process. The first is for devices that were on the market prior to the enactment of the Medical Device Amendments of 1976. *Id.* §§ 360e, 360c(f). When the 1976 amendments were enacted, “Congress realized that existing medical devices could not be withdrawn from the market” and, therefore, included a grandfathering provision exempting such devices from full premarket requirements. *Medtronic, Inc. v. Lohr*, 518 U.S. 470, 477–78 (1996). Second, premarket approval is not required for devices that have been granted an exemption for the purpose of conducting investigations concerning that device. 21 U.S.C. § 360j(g). Finally, premarket approval is not required for devices that are determined to be substantially equivalent to a device already on the market. *Id.* §§ 360c(f)(1), 360e(b). The majority of devices are cleared through 510(k). AMANDA K. SARATA, CONG. RSCH. SERV., RL47374, FDA REGULATION OF MEDICAL DEVICES 8 (2023); Nathan Cortez, *Digital Health and Regulatory Experimentation at the FDA*, 21 YALE J.L. & TECH. (SPECIAL ISSUE) 4, 19 (2019) (citing INST. OF MED., MEDICAL DEVICES AND THE PUBLIC'S HEALTH: FDA 510(K) CLEARANCE PROCESS AT 35 YEARS 15, 85 (2011)).

40. See Cortez, *supra* note 39, at 7. Limitations in early artificial intelligence systems prevented them from being widely applied in medicine, but by the 1970s, medical software systems had developed to a point where they were being used to support clinical decision-making. See Kaul et al., *supra* note 10, at 807–08. When artificial intelligence tools were first used in the clinical context, they were standalone systems that analyzed information input to the systems by clinicians. Mahadevaiah et al., *supra* note 11, at e228–29; see also I. Glenn Cohen et al., *The Legal and Ethical Concerns That Arise from Using Complex Predictive Analytics in Health Care*, 33 HEALTH AFFS. 1139 (2014) (explaining that probability analytics can enhance clinician judgment by suggesting treatment options or how to best allocate resources). Some tools enhanced networking and collaboration among physicians to improve clinical and research capabilities.

For example, in the 1960s, the National Institutes of Health funded programs exploring the use of computer systems to assist humans in making clinical decisions.⁴¹ By the 1970s, the FDA had begun granting premarket approval to computerized medical tools, such as cardiac pacemaker programs, magnetic resonance imaging (“MRI”) machines, and patient monitors.⁴² The FDA’s interest in this technology became more systematic in the early 1980s, when it created a few task forces and committees devoted to studying medical software.⁴³

By 1987, the agency had published its first draft guidance, which identified devices and products that would be subject to regulation.⁴⁴ That draft policy was updated in 1989,⁴⁵ though it was never finalized.⁴⁶ Pursuant to the agency’s 1989 draft guidance, which came to be known as the “Draft Software Policy,” the FDA aimed to apply “the least degree of [regulatory] control necessary to provide reasonable assurance of safety and

Kaul et al., *supra* note 10, at 808–09. Other artificial intelligence systems applied sets of rules to patient information that physicians would input into the system to produce a list of potential diagnoses and recommend treatment options. *Id.* Over the next several decades, artificial intelligence tools continued improving to provide better assistance to clinicians in a variety of ways. *Id.* Eventually, such systems evolved to make predictions and alert clinicians of issues or recommend actions. *See* Mahadevaiah et al., *supra* note 11, at e228–29. Today, clinicians are supported by an array of artificial intelligence tools performing a variety of functions, from making predictions and developing treatment plans, to image analysis and recommending preventative care. *Id.*

41. Nathan Cortez, *Analog Agency in a Digital World*, in *FDA IN THE TWENTY-FIRST CENTURY* 438, 439–40 (Holly Fernandez Lynch & I. Glenn Cohen eds., 2015) (citing *Computers in Health Care: Hearings Before the Subcomm. on Dom. & Int’l Sci. Plan., Analysis & Coop. of the H. Comm. on Sci. & Tech.*, 99th Cong. 76–77 (1986) (statement of Arnold W. Pratt, Director, Division of Computer Research and Technology, National Institutes of Health)). These programs laid the groundwork for future developments of artificial intelligence systems to be used in medical care. *Id.*; Casimir A. Kulikowski, *Beginnings of Artificial Intelligence in Medicine (AIM): Computational Artifice Assisting Scientific Inquiry and Clinical Art—with Reflections on Present AIM Challenges*, in *IMIA YEARBOOK OF MEDICAL INFORMATICS* 249, 251 (2019); GREGORY FREIHERR, U.S. DEP’T OF HEALTH, EDUC. & WELFARE, NIH PUB. NO. 80-2071, *THE SEEDS OF ARTIFICIAL INTELLIGENCE: SUMEX-AIM* (1980).

42. Cortez, *supra* note 39, at 7 (citing *Information Technologies in the Health Care System: Hearing Before the Subcomm. on Investigations & Oversight of the H. Comm. on Sci. & Tech.*, 99th Cong. 199 (1986)).

43. *Id.* at 7–8; Cortez, *supra* note 41, at 444.

44. FDA Draft Policy Guidance for Regulation of Computer Products, 52 Fed. Reg. 36104 (Sept. 25, 1987).

45. Draft FDA Policy for the Regulation of Computer Products (Nov. 13, 1989), 1989 WL 1178702.

46. *See* Nathan Cortez, *Regulating Disruptive Innovation*, 29 *BERKELEY TECH. L.J.* 175, 192 (2014).

effectiveness.”⁴⁷ For years, this meant that the agency would apply a light-touch approach to software systems that required human intervention before having any effect on patients.⁴⁸ In the 1990s, the agency indicated that it might publish comprehensive rules for medical software, but it never did so.⁴⁹ The FDA withdrew its 1989 draft policy in 2005 without comment.⁵⁰

As the foregoing illustrates, although the FDA has long had an interest in regulating computerized medical technologies, the agency’s attempts to act on this interest throughout the years have been halting and sporadic.⁵¹ The lack of a clear framework for regulating medical software can be explained in part by the ambiguity in the agency’s statutory authority. For approximately forty years, the FDA has regulated medical software based on the broad definition of “device” provided by the 1976 Medical Device Amendments, which did not specify how the agency should review and approve advancements in medical computing.⁵² As a result of this statutory ambiguity, “the FDA has been both blessed and cursed with significant discretion in how to adapt the 1976 statutory framework to computer hardware and software products.”⁵³

There are a few reasons contributing to the agency’s reluctance to make bold use of its regulatory discretion. One of the FDA’s enduring concerns is that creating more robust regulations for technology that develops at such a fast pace would strain the agency’s case-by-case review processes.⁵⁴ Indeed, in 2011, the FDA cited the speed and complexity of software developments as a factor contributing to its reluctance to implement overarching policies.⁵⁵ Additionally, interested parties have claimed that imposing stringent governance standards on medical software would stifle innovation.⁵⁶ Another hesitation is rooted in uncertainties over whether software is sufficiently different than other medical devices, such as bedpans, to warrant the former

47. Draft FDA Policy for the Regulation of Computer Products (Nov. 13, 1989), 1989 WL 1178702, at *1; E. Stewart Crumpler & Harvey Rudolph, *FDA Software Policy and Regulation of Medical Device Software*, 52 FOOD & DRUG L.J. 511, 513 (1997).

48. See Cortez, *supra* note 46, at 192.

49. Cortez, *supra* note 41, at 443–44.

50. Cortez, *supra* note 46, at 192; see also Draft Guidance for Industry and Food and Drug Administration Staff; Mobile Medical Applications; Availability, 76 Fed. Reg. 43689 (July 21, 2011) (announcing the FDA’s proposed guidance for mobile medical application regulations).

51. See Cortez, *supra* note 41, at 439, 442–43.

52. *Id.* at 443; Cortez, *supra* note 39, at 7.

53. Cortez, *supra* note 39, at 7.

54. See *id.* at 12–14; Cortez, *supra* note 41, at 447–48. Some wonder whether the FDA has the requisite technical expertise to regulate software. *Id.* at 449.

55. Cortez, *supra* note 41, at 448.

56. *Id.* at 447–48.

being subject to its own set of specific regulations.⁵⁷ Still another concern is that proposals to regulate medical software come uncomfortably close to authorizing the agency to regulate the practice of medicine, which would contravene one of the FDA's most fundamental limitations.⁵⁸

A. Shortcomings in the Traditional Regulatory Framework

Regardless of the merits of the various factors contributing to the FDA's regulatory reluctance, recent advancements in artificial intelligence are pushing the agency to develop more predictable and structured oversight of medical software.⁵⁹ Currently, artificial intelligence is good at analyzing data to recognize patterns, weigh probabilities, and more.⁶⁰ As these capabilities continue to improve, artificial intelligence will become better at image compression and identification.⁶¹ These possibilities make artificial intelligence tools particularly promising candidates for detecting pathologies in "image-based sources, such as radiographs, electrocardiograms, or biopsies."⁶² In addition to identifying pathologies, we might be approaching the day when algorithms, rather than clinicians, are leading decisions regarding the most appropriate combination of treatments for a given patient.⁶³

To some degree, that day has already arrived. For example, LumineticsCore—or IDx-DR as it was known when first developed—is an artificial intelligence system that received FDA approval in 2018.⁶⁴ LumineticsCore can diagnose diabetic retinopathy and diabetic macular edema, which are diseases caused by high levels of blood sugar damaging the

57. *Id.* at 448–49.

58. *Id.* at 450–51.

59. *See id.* at 448; Cortez, *supra* note 39, at 9–14.

60. JAMES E. BAKER, THE CENTAUR'S DILEMMA: NATIONAL SECURITY LAW FOR THE COMING AI REVOLUTION 3 (2021).

61. *Id.*

62. Maliha et al., *supra* note 10, at 630.

63. Senger & O'Leary, *supra* note 16, at 283, 291; Froomkin et al., *supra* note 14, at 39. Of course, it is not a foregone conclusion that artificial intelligence will develop such that it supplants many—if not all—of the tasks ordinarily performed by humans. *See* I. Glenn Cohen, *Informed Consent and Medical Artificial Intelligence: What to Tell the Patient?*, 108 GEO. L.J. 1425, 1462–63 (2020); Roger Allan Ford & W. Nicholson Price II, *Privacy and Accountability in Black-Box Medicine*, 23 MICH. TELECOMM. & TECH. L. REV. 1, 5–7 (2016). Nevertheless, it is a possibility that is likely enough to warrant our thinking hard about the challenges the technology presents.

64. Louis Pilla, *AI Comes to Diagnostics*, OPTHALMOLOGY MGMT., Jan.–Feb. 2024, at 14, 14.

retina's blood vessels.⁶⁵ These conditions can lead to light sensitivity and possible blindness if not detected and treated soon enough.⁶⁶ LumineticsCore works by using an artificial intelligence algorithm to analyze images of eyes taken by a retinal camera.⁶⁷ "The images are uploaded to a cloud server, and the software provides two results: [1] 'more than mild diabetic retinopathy, refer to an eye care professional' or [2] 'negative for more than mild diabetic retinopathy; rescreen in 12 months.'"⁶⁸ Normally, trained physicians need to examine and diagnose patients. But LumineticsCore enables non-physician clinicians to screen for disease without the need to see a physician unless follow-up is warranted. This means more patients can be seen and screened in various locations, including at the office of their primary care doctor. LumineticsCore also frees time for physicians to perform a variety of additional tasks.

This example represents only a fraction of the growing number of increasingly sophisticated artificial intelligence products currently being used by clinicians across the United States.⁶⁹ Although there are many important

65. Press Release, U.S. Food & Drug Admin., FDA Permits Marketing of Artificial Intelligence-Based Device to Detect Certain Diabetes-Related Eye Problems (Apr. 11, 2018), <https://www.fda.gov/news-events/press-announcements/fda-permits-marketing-artificial-intelligence-based-device-detect-certain-diabetes-related-eye> [<https://perma.cc/MWK3-6DWZ>]; Andrzej Grzybowski & Piotr Brona, *Analysis and Comparison of Two Artificial Intelligence Diabetic Retinopathy Screening Algorithms in a Pilot Study: IDx-DR and Retalyze*, 10 J. CLINICAL MED. 2352, 2353 (2021).

66. See Press Release, U.S. Food & Drug Admin., *supra* note 65.

67. Melissa M. Chen et al., *Who Will Pay for AI?*, 3 RADIOLOGY: A.I., e210030, at 3 (Mar. 3, 2021), <https://pubs.rsna.org/doi/10.1148/ryai.2021210030> [<https://perma.cc/WKH8-NNEU>].

68. *Id.* (quoting Pam Kassing & Christina D. Berry, *Hospital Outpatient Prospective Payment System: A Maturing Prospective Payment System*, 17 J. AM. COLL. RADIOLOGY 534 (2020)). The creator of LumineticsCore claims that the system is "autonomous" in the sense that it can diagnose patients without the need for physician involvement. See Gerke et al., *supra* note 9, at 1. However, the system is not truly autonomous because "[i]t is one part of a larger system involving various kinds of human involvement—from health care teams inputting the data to physicians reacting to the [artificial intelligence] recommendation to insurers deciding whether to reimburse only for certain courses of action." *Id.*

69. See Sara Gerke, *Health AI for Good Rather Than Evil? The Need for a New Regulatory Framework for AI-Based Medical Devices*, 20 YALE J. HEALTH POL'Y L. & ETHICS 433, 506 (2021); Press Release, U.S. Food & Drug Admin., FDA Permits Marketing of Clinical Decision Support Software for Alerting Providers of a Potential Stroke in Patients (Feb. 13, 2018), <https://www.fda.gov/news-events/press-announcements/fda-permits-marketing-clinical-decision-support-software-alerting-providers-potential-stroke> [<https://perma.cc/G2HE-ZUW6>]. Caption Guidance is another example of artificial intelligence software approved by the FDA. See Gerke, *supra*, at 506. Caption Guidance's software captures ultrasound images of patients' hearts. *Id.* Though Caption Guidance resembles other clinician-support tools in many respects, a notable "peculiarity of the software is that it can be used by non-experts"; for example, nurses can be

differences between the various systems now in use, they are alike in that they are “locked.”⁷⁰ Locked systems do not evolve over time as they analyze new data to modify how they perform the tasks they have been assigned.⁷¹ Instead, they provide “the same result each time the same input is applied.”⁷²

It is important to note that LumineticsCore is not locked because it is incapable of learning over time.⁷³ Rather, LumineticsCore is locked because the FDA’s existing framework for reviewing and approving medical devices is not well equipped to handle adaptive systems.

A large part of what makes adaptive systems so powerful and promising is their use of a type of artificial intelligence that is distinct from conventional software.⁷⁴ The development of conventional software involves humans using principles of formal logic to write code needed to perform a specific task.⁷⁵ For example, a typical approach to designing a tool to assist physicians in diagnostic tasks might begin by identifying all symptoms and diagnostic ranges of a disease.⁷⁶ Then, developers would manually write lines of code mapping all relationships between these symptoms and diagnoses.⁷⁷ In this way, conventional software involves lines of code that are designed to follow rules and pathways created by humans.⁷⁸

trained to operate Caption Guidance with only a few days of instruction. *Id.* In this way, Caption Guidance is like IDx-DR insofar as both can be operated by non-experts. *Id.*

Another example is the Viz.AI Contact application. Like LumineticsCore, Viz.AI was approved by the FDA in 2018. *See* Press Release, U.S. Food & Drug Admin., *supra*. Viz.AI’s artificial intelligence software analyzes computed tomography images of the brain to identify indicators associated with a stroke. *Id.* If Viz.AI identifies an indicator, the application notifies a neurovascular specialist. *Id.*

Of course, these examples are only a few of many. *See, e.g.*, Laura M. Holdsworth et al., *Predicting and Responding to Clinical Deterioration in Hospitalized Patients by Using Artificial Intelligence: Protocol for a Mixed Methods, Stepped Wedge Study*, 10 JMIR RSCH. PROTOCOLS, e27532 (July 7, 2021), <https://www.researchprotocols.org/2021/7/e27532> [<https://perma.cc/L7DB-92PP>]; Samer Ellahham & Nour Ellahham, *Use of Artificial Intelligence for Improving Patient Flow and Healthcare Delivery*, 12 J. COMPUT. SCI. & SYS. BIOLOGY 1 (2019).

70. *See* Boubker, *supra* note 9, at 428; Surette, *supra* note 9.

71. Gerke et al., *supra* note 9, at 1.

72. U.S. FOOD & DRUG ADMIN., *supra* note 9, at 3 n.7.

73. *See* Surette, *supra* note 9.

74. *See* Ravi et al., *supra* note 10, at 4; Kaul et al., *supra* note 10, at 807; Surden & Williams, *supra* note 10, at 147–48; Surden, *supra* note 10, at 88; Maliha et al., *supra* note 10, at 630.

75. *See* Marta Garnelo & Murray Shanahan, *Reconciling Deep Learning with Symbolic Artificial Intelligence: Representing Objects and Relations*, 29 CURRENT OP. BEHAV. SCIS. 17, 17 (2019).

76. *See* STUART RUSSELL & PETER NORVIG, ARTIFICIAL INTELLIGENCE: A MODERN APPROACH 23 (4th ed. 2020).

77. *Id.*

78. Garnelo & Shanahan, *supra* note 75, at 17.

Techniques for developing modern artificial intelligence tools are distinct from this conventional approach. Much of the excitement surrounding modern artificial intelligence flows from machine learning, which is a particular type of artificial intelligence.⁷⁹ Machine learning refers to the capacity of a computer system to, over time, improve its ability to perform programmed tasks by building statistical associations between various data inputs.⁸⁰ The most cutting-edge practitioners today are exploring deep learning, which is a set of artificial intelligence techniques falling under the umbrella of machine learning.⁸¹ Deep learning occurs when layers of a computer system's artificial neural networks make cascades of calculations to adjust the way data is received and processed at various stages based on what it recognizes as "correct" outputs.⁸² Originally modeled on the neurons in human brains, artificial neural networks are based on mathematical models enabling layers of interconnected neurons to communicate with each other.⁸³ Different "weights" are assigned to the point at which the nodes of different neurons connect, and this weighting guides input data as it travels through each layer to produce outputs.⁸⁴ Stacking multiple layers on top of each other enables the system as a whole to identify complex concepts and patterns.⁸⁵ For example, whereas one layer might be concerned with identifying color pixels in an image, the next might be tasked with identifying shapes and so

79. BAKER, *supra* note 60, at 15; Calo, *supra* note 1, at 424.

80. Ryan Marshall Felder, *Coming to Terms with the Black Box Problem: How to Justify AI Systems in Health Care*, HASTINGS CTR. REP., July–Aug. 2021, at 38, 38; BAKER, *supra* note 60, at 14; Calo, *supra* note 1, at 405; Surden, *supra* note 10, at 88. Machine learning can also be defined as a "type of AI that uses algorithms whose performance improves as they are exposed to more data over time." Maliha et al., *supra* note 10, at 630.

There are three main ways in which machine learning can occur. BAKER, *supra* note 60, at 14–15. The first is supervised learning, which involves feeding a computer mathematically weighted data so the computer can be trained to better recognize new data and make forward-looking predictions. *Id.* at 15; Surden, *supra* note 10, at 93. The second is unsupervised learning, which occurs when a computer leverages algorithmic models to identify patterns without first being trained to do so. BAKER, *supra* note 60, at 15. Finally, reinforced learning occurs when a computer is prompted, via algorithmic programming, to optimize its performance by learning from its experiences. *Id.* at 15.

81. Calo, *supra* note 1, at 405.

82. See CHRISTOPHER BISHOP, *PATTERN RECOGNITION AND MACHINE LEARNING* 32 (Michael Jordan et al. eds., 2006); BAKER, *supra* note 60, at 16–17, 79–80; Price, *supra* note 16, at 426; Surden & Williams, *supra* note 10, at 147–48.

83. See Yavar Bathaee, *The Artificial Intelligence Black Box and the Failure of Intent and Causation*, 31 HARV. J.L. & TECH. 889, 901–03 (2018).

84. Bryan H. Choi, *AI Malpractice*, 73 DEPAUL L. REV. 301, 313 (2024).

85. *Id.*

on until the system as a whole can determine whether an image contains indicators of a particular disease, such as diabetic retinopathy.⁸⁶

Perhaps the most important feature of deep learning systems distinguishing them from conventional software systems, and even other types of artificial intelligence tools, is that they are autodidactic—which is to say, humans do not design the neural networks of these systems.⁸⁷ Although a developer might be responsible for the initial way the system weights input data, over time, these multi-layered networks learn from experience to improve their ability to perform tasks by reinforcing or decaying individual connections between neurons.⁸⁸ The data itself, rather than human design, determines the number of layers in these neural networks.⁸⁹

To make the concept of adaptive artificial intelligence systems leveraging deep learning a little more concrete, consider the following example. Imagine a modified version of the LumineticsCore tool; let's call this tool DLeye. Unlike LumineticsCore, DLeye is adaptive because it evolves over time as it makes cascades of calculations to weight and re-weight data as it improves its ability to perform certain tasks. Suppose that DLeye consistently outperforms human clinicians in identifying disease in radiology images. Due to this record of consistent out-performance, it becomes standard practice for human clinicians to rely on the system to review images in the first instance; humans only analyze these images themselves if a concern is flagged by the tool. As designed, DLeye continues learning over time as non-physician clinicians use the tool to review images of patients' eyes for indications of disease.⁹⁰ As the system continues learning over time, it evolves in a way that causes it to fail to identify disease in an image of a patient's eye. Because no disease was detected by DLeye, no follow-up appointment was made, and the patient's health irreversibly declined as a result of not being referred.

This example illustrates that, on the one hand, the adaptiveness of the system enables it to continually improve its capabilities over time, and this is a desirable feature. We want to take advantage of the ways in which these systems could lead to many exciting and dramatic improvements in medical care, from optimizing workflows, reducing spending, and enhancing treatment quality to increasing diagnostic accuracy, expanding access to care, and more.⁹¹ On the other hand, the dynamic nature of DLeye makes it difficult

86. *See id.*

87. Eric J. Topol, *High-Performance Medicine: The Convergence of Human and Artificial Intelligence*, 25 *NATURE MED.* 44, 45 (2019).

88. *See Bathae, supra* note 83, at 901–03.

89. Topol, *supra* note 87, at 45.

90. *See Mahadevaiah et al., supra* note 11, at e228–29.

91. *See, e.g., Gerke et al., supra* note 9, at 1; Cortez, *supra* note 13, at 1190–99.

to predict whether it will evolve in a problematic way at some point in the future. Given this latter risk, it is understandable that the FDA would want to ensure the continued safety and effectiveness of medical tools that are sensitive to rapid technological advancements and new data encountered in clinical workflows.⁹² The FDA's traditional method for doing so is by requiring technologies subject to regulation—including medical software—to pass through a set of trials and tests to ensure their safety and effectiveness before approving them to be used in clinical workflows.⁹³ But this conventional regulatory approach is backward-looking insofar as it requires any adaptive artificial intelligence tools to be re-reviewed and approved by the FDA each time they learn to improve their capabilities.⁹⁴ Because of the significant up-front cost of shepherding tools through the agency's traditional processes, it will be cumbersome for the agency and developers to do so on a repeat basis each time these systems evolve.⁹⁵ Given these burdensome re-submission requirements, developers are likely to opt for creating "locked" systems rather than adaptive ones.⁹⁶ For these reasons, treating medical software like static medical devices (e.g., bedpans or scalpels) could deprive clinicians, patients, and the broader public of the benefits made possible by deep learning systems.⁹⁷

The FDA has recognized both the promise of these systems and that the agency's existing regulatory framework "is not well suited for the faster iterative design, development, and type of validation used for" deep learning systems.⁹⁸ In response to this dilemma, the FDA released a discussion paper in 2019 to solicit feedback on how to best address the regulatory challenges posed by adaptive medical software.⁹⁹ That 2019 discussion paper led to a 2021 action plan, which proposed the idea of using a "Predetermined Change Control Plan" to ensure the safety and effectiveness of adaptive medical

92. See Cortez, *supra* note 39, at 7–9.

93. See Senger & O'Leary, *supra* note 16, at 292.

94. Surette, *supra* note 9; Boubker, *supra* note 9, at 435 (citing U.S. FOOD & DRUG ADMIN., *supra* note 9, at 5).

95. Gerke et al., *supra* note 9, at 1.

96. See *id.*; Kaul et al., *supra* note 10, at 807; Surden & Williams, *supra* note 10, at 147–48; Surden, *supra* note 10, at 88; Maliha et al., *supra* note 10, at 630.

97. See Gerke et al., *supra* note 9, at 1; Kaul et al., *supra* note 10, at 807; Surden & Williams, *supra* note 10, at 147–48; Surden, *supra* note 10, at 88; Maliha et al., *supra* note 10, at 630.

98. U.S. FOOD & DRUG ADMIN., DIGITAL HEALTH INNOVATION ACTION PLAN 2 (2017), <https://www.fda.gov/downloads/MedicalDevices/DigitalHealth/UCM568735.pdf> [<https://perma.cc/7NE5-AS56>].

99. See U.S. FOOD & DRUG ADMIN., *supra* note 9.

software.¹⁰⁰ In April 2023, building on feedback it received in response to its 2019 discussion paper and 2021 action plan, the FDA released draft guidance further outlining the agency's conception of a predetermined change control plan.¹⁰¹

This plan has three main components. The first is a description of the anticipated modifications the system is expected to undergo in the real world.¹⁰² Because each anticipated modification must be included in the predetermined change control plan, the agency recommends including only a limited number of specific updates that can be verified and validated.¹⁰³ The description must also make clear whether all devices in the market—rather than a select few—will be modified and whether the changes will be implemented automatically (i.e., pursuant to instructions the adaptive tool gives itself without human oversight) or manually (i.e., via human-directed updates).¹⁰⁴

Second, predetermined change control plans must detail the developers' protocols for ensuring that the system will remain safe and effective as it evolves over time.¹⁰⁵ Among the primary components that should be included in such protocols are: how data will be collected, analyzed, used, and stored; strategies for the initial and ongoing training of data over time; protocols for performance evaluation; and procedures for updating practices concerning data management, re-training, and performance evaluation.¹⁰⁶

Finally, predetermined change control plans must provide an assessment of the benefits and risks—along with steps taken to reduce risks—of implementing proposed changes.¹⁰⁷ The FDA's draft guidance recommends that these assessments compare modified versions of the device to non-modified ones, discuss benefits and risks of each modification, detail efforts to ensure safety and effectiveness, evaluate how an update to one device might affect another, and assess the collective effect of all proposed modifications.¹⁰⁸

100. U.S. FOOD & DRUG ADMIN., ARTIFICIAL INTELLIGENCE/MACHINE LEARNING (AI/ML)-BASED SOFTWARE AS A MEDICAL DEVICE (SAMD) ACTION PLAN 1 (2021), <https://www.fda.gov/media/145022/download> [<https://perma.cc/55NH-MVEE>].

101. See U.S. FOOD & DRUG ADMIN., *supra* note 22, at 1–29; U.S. FOOD & DRUG ADMIN., *supra* note 100; U.S. FOOD & DRUG ADMIN., *supra* note 9.

102. See U.S. FOOD & DRUG ADMIN., *supra* note 22, at 16–18.

103. *Id.* at 16.

104. *Id.* at 16–17.

105. See *id.* at 18–23.

106. See *id.* at 20–24; Gerke et al., *supra* note 9, at 3.

107. See U.S. FOOD & DRUG ADMIN., *supra* note 22, at 24–25.

108. *Id.*

The predetermined change control plan laid out in the FDA’s 2023 draft guidance is a laudable example of the agency taking steps to address the challenges posed by adaptive artificial intelligence tools. However, the FDA’s proposed approach to regulating adaptive systems is unworkable. The main problem with the agency’s 2023 guidance is that, like the agency’s traditional regulatory model, it is fundamentally backward-looking. That is, both the 2023 guidance and the FDA’s conventional regulatory scheme look to evidence of how a medical tool has performed in the past when determining whether to approve it for use in the care of patients. This backward-looking regulatory approach works well for medical products that are locked and don’t evolve over time. But this framework is ill-suited for adaptive tools. Because such systems continually evolve, they cannot be guaranteed to function consistently over time and in different settings.¹⁰⁹ As a result, “[a]ny predetermined change control plan risks being either uninformative or impractical.”¹¹⁰ At one end of the spectrum, developers might attempt to describe DEye’s anticipated changes at a very high level given its expectation of updates as it encounters new data and its susceptibility to various developments in the people and processes using the system over time.¹¹¹ But such a description would be too vague and uninformative.¹¹² At the other extreme, it would be impractical to expect developers of a tool like DEye to describe precisely how the system will adapt over time as it encounters new data and other changes in the clinical environment.¹¹³ Hence, the adaptiveness of a system like DEye means that developers will struggle to comply with the FDA’s 2023 guidance by providing a reliable *ex ante* prediction of how the tool will learn and evolve over time.¹¹⁴

109. Gerke et al., *supra* note 9, at 2.

110. Babic et al., *supra* note 9, at 1203.

111. *Id.*

112. *Id.*

113. *Id.*

114. *Id.*; *see also* M. Jason Brooke, Brooke & Assocs., Comment on Marketing Submission Recommendations for a Predetermined Change Control Plan for Artificial Intelligence/Machine Learning (AI/ML)-Enabled Device Software Functions—Draft Guidance for Industry and FDA Staff (July 5, 2023), <https://www.regulations.gov/comment/FDA-2022-D-2628-0026> [<https://perma.cc/2YM6-BZXN>]; Inflammatrix Inc., Comment on Marketing Submission Recommendations for a Predetermined Change Control Plan for Artificial Intelligence/Machine Learning (AI/ML)-Enabled Device Software Functions—Draft Guidance for Industry and FDA Staff (June 25, 2023), <https://www.regulations.gov/comment/FDA-2022-D-2628-0009> [<https://perma.cc/X6E6-9UCB>]; Blair Anderson, Amazon Web Servs., Comment on Marketing Submission Recommendations for a Predetermined Change Control Plan for Artificial Intelligence/Machine Learning (AI/ML)-Enabled Device Software Functions—Draft Guidance for Industry and FDA Staff (July 5, 2023), <https://www.regulations.gov/comment/FDA-2022-D->

B. Outlining a Forward-Looking Regulatory Framework

In response to these challenges, I build on the work of Sara Gerke, Nicholson Price, Boris Babic, Glenn Cohen, and others to suggest a forward-looking alternative to the FDA's traditional backward-looking regulatory approach.¹¹⁵ The framework I have in mind draws on insights from sociotechnical systems theory.¹¹⁶

With roots tracing back to the years following World War II, sociotechnical systems theory studies how human and technical systems interact with each other when deployed in workplaces to achieve collective goals.¹¹⁷ Although several distinct schools of sociotechnical theory have

2628-0028 [<https://perma.cc/UQZ3-VSX6>]; Joseph Corrigan, Cambridge Consultants, Comment on Marketing Submission Recommendations for a Predetermined Change Control Plan for Artificial Intelligence/Machine Learning (AI/ML)-Enabled Device Software Functions—Draft Guidance for Industry and FDA Staff, (July 5, 2023), <https://www.regulations.gov/comment/FDA-2022-D-2628-0021> [<https://perma.cc/X9N5-QXQ4>].

115. Gerke et al., *supra* note 9, at 1; Price, *supra* note 9, at 424.

116. I thank Ryan Calo and Gennie Mansi for very illuminating conversations on sociotechnical systems theory.

117. Ken Eason, *Afterword: The Past, Present and Future of Sociotechnical Systems Theory*, 45 APPLIED ERGONOMICS 213, 213 (2014). The origins of sociotechnical systems theory can be traced to the years following World War II. Enid Mumford, *The Story of Socio-Technical Design: Reflections on Its Success, Failures, and Potential*, 16 INFO. SYS. J. 317, 319–320 (2006). At that time, coal was the primary source of power in Britain, but the industry was struggling to produce enough to keep pace with demand. ERIC TRIST, THE EVOLUTION OF SOCIO-TECHNICAL SYSTEMS: A CONCEPTUAL FRAMEWORK AND AN ACTION RESEARCH PROGRAM 7 (1981). In response, Britain's National Coal Board asked the Tavistock Institute of Human Relations to examine the British coal mining industry and explain low productivity in the sector. *Id.* at 7–8. The Institute's researchers observed that in many instances during this post-war mining era, engineers without hands-on mining experience were enlisted to create and implement new mining technologies and practices. William Pasmore et al., *Reflections: Sociotechnical Systems Design and Organization Change*, 19 J. CHANGE MGMT. 67, 68 (2019). Although these new practices and systems were intended to boost productivity and efficiency, problems arose almost as soon as they were implemented. *Id.* The complaints from miners were many and varied: the new technologies created safety issues; mechanized mining methods were too rigid to be compatible with unpredictable underground conditions; new work systems requiring miners to perform rote tasks in siloes depressed miners' camaraderie, job satisfaction, and productivity; individual compensation schemes pitted miners against each other; and more. *Id.*; TRIST, *supra*, at 7–10. Based on these findings, the Institute concluded that employee morale, workplace safety, and enterprise productivity suffer when engineers design and organize workplace technology without considering the effects those decisions will have on human capital. Mumford, *supra*, at 318; TRIST, *supra*, at 7–10; Mark Govers & Pierre van Amelsvoort, *A Theoretical Essay on Socio-Technical Systems Design Thinking in the Era of Digital Transformation*, 54 GRUPPE. INTERAKTION. ORGANISATION. 27, 28 (2023). To avoid these negative consequences and boost productivity in the mining industry, the Institute recommended that when assessing workplace

emerged over the years, a common theme bridges the various flavors of this concept: a focus on wrestling with challenges that emerge when new workplace technologies interact with the people, institutional structures, and other components of any given workplace.¹¹⁸

Drawing inspiration from this literature, I suggest that adaptive tools like DLeye should be regulated pursuant to a forward-looking framework that takes into account how the tools interact with the people and institutional structures in the environment where they operate. The sort of forward-looking framework I have in mind would closely resemble the predetermined change control plan that the FDA outlined in its 2023 draft guidance. However, the framework I propose would be different in a few important ways. In particular, it would require developers, health systems, and clinicians to maintain adequate processes for developing and using adaptive and opaque artificial intelligence tools in the care of patients.¹¹⁹ The description of processes to be submitted to the FDA would focus less on predicting precisely how DLeye itself would evolve over time. Instead, a forward-looking regulatory scheme would provide the agency with a detailed description of the collaborative efforts undertaken by various interested parties in the healthcare setting as they work alongside DLeye on an ongoing basis.¹²⁰ This detailed description could be divided into three components: (1) the people responsible for using the tool; (2) the processes governing the use of these use various instruments in the care of patients; and (3) the design and production of adaptive artificial intelligence tools.

1. People

Leveraging a tool like DLeye in the care of patients will surely implicate a multidisciplinary collection of professionals with expertise in a variety of

productivity, the technical and social aspects of organizations be jointly analyzed. Pasmore et al., *supra*, at 68–69; TRIST, *supra*, at 7–10.

Initially, sociotechnical theory focused on heavy manufacturing and industries, such as coal and textiles. Gordon Baxter & Ian Sommerville, *Socio-Technical Systems: From Design Methods to Systems Engineering*, 23 INTERACTING WITH COMPUTS. 4, 6 (2011). Eventually, these ideas made their way into other workspaces, including the professional services industry. Mumford, *supra*, at 330. Over the years, sociotechnical systems theory continued developing across various industries and established distinct approaches based on the particular challenges in a given context. Matthew C. Davis et al., *Advancing Socio-Technical Systems Thinking: A Call for Bravery*, 45 APPLIED ERGONOMICS 171, 172 (2014).

118. Davis et al., *supra* note 117, at 172.

119. *See id.* at 173.

120. *See* Baxter & Sommerville, *supra* note 117, at 13.

areas.¹²¹ In the name of economy, rather than require that all these individuals and entities be listed, I propose that the forward-looking plan submitted to the FDA focus at a high level on software developers, clinicians, and health administrators.¹²²

The team of software developers will include both those with primary responsibility for the business and technical aspects of developing adaptive and opaque artificial intelligence tools. Because a tool like DLeye will evolve over time, developers will be involved on an ongoing basis both in the system's continued development and in the re-training of clinicians to use the tool as it learns from new data.

The clinical team will include physicians needed to help train the artificial intelligence system to recognize disease on an initial and ongoing basis. Because a significant benefit of a system like DLeye is that it can be used by clinicians who do not have an M.D., non-physician clinicians will also be crucial members of the team. For these reasons, I use the term clinician broadly to include physicians, nurses, and health technicians.

Additionally, administrators will be important actors given their responsibility for operating the health system in which DLeye is deployed on a day-to-day basis. Administrators will also play a large role in acquiring the services of tools like DLeye and setting the credential standards for clinicians tasked with using such systems in the care of patients.

121. See, e.g., T. Martin-Noguerol et al., *Artificial Intelligence in Radiology: Relevance of Collaborative Work Between Radiologists and Engineers for Building a Multidisciplinary Team*, 76 *CLINICAL RADIOLOGY* 317, 317–22 (2021); Joseph C. Gambone & Michael S. Broder, *Embedding Quality Improvement and Patient Safety—The UCLA Value Analysis Experience*, 21 *BEST PRAC. & RSCH. CLINICAL OBSTETRICS & GYNAECOLOGY* 581, 583 (2007); Claire Rupert & Sylvia Kitzman, *The Clinical Voice in Product Selection and Implementation*, 8 *SSM* 24, 24 (2002); Quinn Grundy, “*Whether Something Cool Is Good Enough*”: *The Role of Evidence, Sales Representatives and Nurses’ Expertise in Hospital Purchasing Decisions*, 165 *SOC. SCI. & MED.* 82, 82–83 (2016).

122. Other actors that are likely to play a role in the use of this technology are interest groups representing patients, technology industry experts, ethicists, and lawyers. These actors and many other individuals are likely to be involved in a health system's value analysis committee, which have been created at many institutions in response to pressure to cut costs and remain competitive in the modern healthcare landscape. David M. Kalainov, *Value-Based Healthcare: Controlling Costs Through a Value Analysis Committee*, 482 *CLINICAL ORTHOPAEDICS & RELATED RSCH.* 769, 769 (2024); Grundy, *supra* note 121, at 82. The primary focus of such committees is to review the health system's procurement of medical products to maximize efficacy, cost-savings, safety, and quality. Daniel T. Engelman et al., *Addressing the Imperative to Evolve the Hospital New Product Value Analysis Process*, 155 *J. THORACIC & CARDIOVASCULAR SURGERY* 682, 682 (2018). A value analysis committee will likely be responsible for deliberating about, and ultimately acquiring the rights to use, a tool like DLeye in the care of a patient population.

2. Processes

Once these actors have been identified at the outset of the forward-looking plan submitted to the FDA, the document can proceed to detail the guidelines that will govern their conduct. These processes and procedures will cover at least three areas: (1) communication channels; (2) training and credentialing; and (3) ongoing monitoring of the tool and those using it.

The specific details of communication channels between developers, clinicians, and administrators will depend at least in part on the type of technology and the environment in which DEye is deployed. But in general, a mixed framework of bottom-up and top-down communication between front-line professionals and organizational leadership seems appropriate for a dynamic and opaque tool like DEye. Such a model can facilitate open, honest, and reliable lines of communications between developers, clinicians, and administrators on a variety of topics, from developing innovative ideas to the identification of risks.¹²³

Another important set of standards and procedures to be established are those concerning the relevant training and credentialing benchmarks for clinicians and developers. Many people have heard of artificial intelligence, and they interact with it daily.¹²⁴ However, even very sophisticated clinical parties often do not have a detailed understanding of how artificial intelligence systems operate.¹²⁵ For this reason, it will be prudent to develop a set of foundational education and training standards for clinicians who routinely use tools like DEye.¹²⁶

Of course, clinicians will need to know *something* about how these tools work. But precisely what level of technical expertise clinicians can and

123. Jee Young Kim et al., *Organizational Governance of Emerging Technologies: AI Adoption in Healthcare*, in PROCEEDINGS OF THE 2023 ACM CONFERENCE ON FAIRNESS, ACCOUNTABILITY, AND TRANSPARENCY 1395, 1398–1400 (2023).

124. See, e.g., Kanadpriya Basu et al., *Artificial Intelligence: How Is It Changing Medical Sciences and Its Future?*, 65 INDIAN J. DERMATOLOGY 365, 367 (2020); Kim et al., *supra* note 123, at 1402.

125. See, e.g., Santiago Romero-Brufau et al., *A Lesson in Implementation: A Pre-Post Study of Providers' Experience with Artificial Intelligence-Based Clinical Decision Support*, 137 INT'L J. MED. INFORMATICS, 104072, at 4–5 (May 2020), <https://www.sciencedirect.com/science/article/pii/S1386505619310123>; Kim et al., *supra* note 123, at 1402.

126. See, e.g., Mingyang Chen et al., *Acceptance of Clinical Artificial Intelligence Among Physicians and Medical Students: A Systematic Review with Cross-Sectional Survey*, 9 FRONTIERS MED., 990604, at 12–15 (Aug. 31, 2022), <https://www.frontiersin.org/journals/medicine/articles/10.3389/fmed.2022.990604/full> [<https://perma.cc/BW4L-NGMG>]; Kim et al., *supra* note 123, at 1402; Kim V. Garvey et al., *Considering Clinician Competencies for the Implementation of Artificial Intelligence-Based Tools in Health Care: Findings from a Scoping Review*, 10 JMIR MED. INFORMATICS, e37478, at 1 (Nov. 16, 2022), <https://medinform.jmir.org/2022/11/e37478> [<https://perma.cc/PT9U-Q4QG>].

should possess is a matter of debate. On the one hand, some suggest that clinicians should possess a familiarity with mathematical concepts, principles of data and computer science, artificial intelligence fundamentals, and corresponding legal and ethical considerations.¹²⁷ But others suggest that it is more important that clinicians know *that* such systems work rather than know precisely *how* they work.¹²⁸ After all, physicians routinely rely on all sorts of very sophisticated medical technologies like MRI scanners. Physicians might understand at a high-level what these tools are doing to produce their outputs. But few—if any—have sufficient training in physics to fully understand the effect of placing the human body in a strong magnetic field like an MRI scanner.¹²⁹ Whatever level of training is deemed sufficient, poorly trained operators could cause patient harm if they misuse DLeye.¹³⁰

Regardless of the final answer to this question concerning the appropriate level of technical skill to be required of clinicians, it is a matter that will likely be determined in large part by the leaders of health systems.¹³¹ Of course, health administrators will solicit and weigh input from clinicians and software developers when developing the sort of core competencies to be

127. See, e.g., Ketan Paranjape et al., *Introducing Artificial Intelligence Training in Medical Education*, 5 JMIR MED. EDUC., e16048, at 3–4 (Dec. 3, 2019), <https://mededu.jmir.org/2019/2/e16048> [<https://perma.cc/HH57-6P5S>].

128. See, e.g., Robin C. Feldman et al., *Artificial Intelligence in the Health Care Space: How We Can Trust What We Cannot Know*, 30 STAN. L. & POL'Y REV. 399, 412–14 (2019). I thank Gennie Mansi for encouraging me to note this point.

129. *Id.* at 414. There are many other examples indicating that it might not be important that physicians have a detailed understanding of precisely *how* the “guts” of an artificial intelligence system works. Rather, it seems most important that clinicians are able to understand *what* the system is supposed to do, and that relevant specialists (e.g., tech developers, ethicists, and more) have indicated *that* the tool will operate safely and effectively in the care of patients. For instance, “modern clinicians prescribed aspirin as an analgesic for nearly a century without understanding the mechanism through which it works. Lithium has been used as a mood stabilizer for half a century, yet why it works remains uncertain.” Alex John London, *Artificial Intelligence and Black-Box Medical Decisions: Accuracy Versus Explainability*, HASTINGS CTR. REP., Jan.–Feb. 2019, at 15, 17.

130. W. Nicholson Price II, *Medical AI and Contextual Bias*, 33 HARV. J.L. & TECH. 65, 77–79 (2019). Another consideration affecting the level of expertise to be required of clinicians is the possibility that they lack sufficient training to use the tool, leading to additional stress (i.e., “technostress”), or if the system is poorly integrated into the clinical environment and disrupts workflows rather than streamlining them. Kelly J. Thomas Craig et al., *The Burden of the Digital Environment: A Systematic Review on Organization-Directed Workplace Interventions to Mitigate Physician Burnout*, 28 J. AM. MED. INFORMATICS ASS'N 985, 992–95 (2021). As indicated in the sociotechnical systems literature, low morale and job satisfaction among clinicians could lead to decreased production across the healthcare enterprise. Pasmore et al., *supra* note 117, at 68; TRIST, *supra* note 117, at 7–10.

131. See Paranjape et al., *supra* note 127, at 3–4.

required of those using systems like DLeye.¹³² Clinicians, for example, will be best positioned to identify how a design decision made by software developers might affect the actual use of an artificial intelligence tool in the care of patients.¹³³ Complementing clinical expertise will be software engineers, who are best positioned to understand the limitations, or possibilities, of bringing artificial intelligence tools to bear on clinical issues.¹³⁴ But ultimately, health administrators will be responsible for finalizing, disseminating, and enforcing the educational systems, training requirements, and program structures across the healthcare enterprise to ensure the adequate training of clinicians.¹³⁵ Based upon a robust dialogue between administrators, clinicians, and developers, a set of baseline qualifications and training standards can be outlined. To formalize these baseline standards, health administrators might require completion of a fellowship program as a prerequisite for credentialing medical professionals responsible for using tools like DLeye.¹³⁶

Once initial training and credentialing standards have been established, both these standards and the system itself will need to be continually monitored and updated to keep up with adaptive systems as they continue developing over time.¹³⁷ Given the constantly evolving nature of the tools themselves, as well as the monitoring and training processes surrounding them, a robust set of standards and practices will be required to effectively oversee the evolution of tools like DLeye on an ongoing basis.¹³⁸ The full range of benchmarks and practices to be tracked is expansive and complex, but key examples include measuring outcomes across a variety of contexts, including patient health, financial considerations, clinician and patient satisfaction with the tool, and more.¹³⁹ Other metrics to be measured might include evolving protocols with regard to initial treatment, follow-up care,

132. *See id.*

133. *See, e.g.,* Yvonne W. Lui et al., *How to Implement AI in the Clinical Enterprise: Opportunities and Lessons Learned*, 17 J. AM. COLL. RADIOLOGY 1394, 1395–97 (2020); Kim et al., *supra* note 123, at 1397.

134. *See* Kim et al., *supra* note 123, at 1397; Gerke et al., *supra* note 9, at 3 (citing U.S. FOOD AND DRUG ADMIN., DEN180001, DE NOVO CLASSIFICATION REQUEST FOR IDX-DR (2018), https://www.accessdata.fda.gov/cdrh_docs/reviews/DEN180001.pdf [<https://perma.cc/YC4E-EEH7>]).

135. *See* Paranjape et al., *supra* note 127, at 3–4.

136. Gerke et al., *supra* note 9, at 2.

137. *See, e.g.,* Junaid Bajwa et al., *Artificial Intelligence in Healthcare: Transforming the Practice of Medicine*, 8 FUTURE HEALTHCARE J. e188, e190 (2021); Kim et al., *supra* note 123, at 1402–03.

138. Kim et al., *supra* note 123, at 1402–04.

139. *Id.* at 1403.

and changes in responsibilities with regard to clinical, technical, and administrative members of the team.¹⁴⁰ In addition to creating a standardized monitoring plan to be observed on a regular basis, a channel for raising any unexpected issues that might be identified by members of the team will also be appropriate.¹⁴¹

Once the monitoring and testing responsibilities of various actors are identified, a forward-looking plan will need to provide additional details about how it will be enforced. One option is for the FDA to follow its current premarket approval process and conduct periodic examinations of the developer's manufacturing facilities and the clinical settings in which DLeve systems will operate.¹⁴² For example, the FDA's Sentinel system could be modified to monitor the testing and quality assurance processes for adaptive black-box systems.¹⁴³ Launched in 2008 pursuant to congressional mandate, Sentinel collects and monitors data to ensure the safety of medical products.¹⁴⁴ Any time patients interact with the U.S. health care system, data is created (e.g., patient billing records or electronic health records documenting care).¹⁴⁵ Organizations participating in the Sentinel system collect and maintain data in a standard format, which enables the FDA and partnering organizations to analyze data across the distributed network to identify patterns in links between various products and adverse events.¹⁴⁶ The FDA's Sentinel system could be expanded to monitor the performance of tools like DLeve.¹⁴⁷

Alternatively, the FDA could enlist the assistance of an organization like the Joint Commission.¹⁴⁸ As a non-profit organization tasked with accrediting and certifying healthcare programs and institutions, the Joint Commission conducts periodic visits of healthcare facilities to evaluate the degree to which they satisfy safety and quality metrics.¹⁴⁹ The Joint Commission's

140. *Id.*

141. *See, e.g.,* Davis et al., *supra* note 117, at 173–75; Kim et al., *supra* note 123, at 1403.

142. 21 U.S.C. § 360e(d)(2).

143. *See* Babic et al., *supra* note 9, at 1204.

144. *FDA's Sentinel Initiative*, U.S. FOOD & DRUG ADMIN. (Mar. 8, 2024), <https://www.fda.gov/safety/fdas-sentinel-initiative> [<https://perma.cc/PL8H-Q7KM>].

145. *See About the Food and Drug Administration (FDA) Sentinel Initiative*, SENTINEL, <https://www.sentinelinitiative.org/about> [<https://perma.cc/899N-KT84>].

146. Steven Findlay, *Health Policy Brief: The FDA's Sentinel Initiative*, HEALTHAFFAIRS 3–4 (June 4, 2015), https://www.healthaffairs.org/doi/10.1377/hpb20150604.936915/full/healthpolicybrief_139-1534166665937.pdf [<https://perma.cc/6FCL-AZL2>].

147. Babic et al., *supra* note 9, at 1204.

148. *See* Cortez, *supra* note 39, at 18–19.

149. *See Facts About the Joint Commission*, JOINT COMM'N, <https://www.jointcommission.org/who-we-are/facts-about-the-joint-commission> [<https://perma.cc/4463-TRVY>].

mandate could be expanded to conduct periodic assessments to certify that developers and relevant clinical parties are maintaining and observing adequate safety and effectiveness procedures.¹⁵⁰

To be sure, steps must be taken to ensure the enforcement of sufficiently stringent safety and quality standards by the Joint Commission or a similar organization charged with evaluating institutions and entities using adaptive black-box tools in the care of patients.¹⁵¹ Both the Sentinel system and the Joint Commission have been criticized for taking a light-touch approach to enforcing relevant standards.¹⁵² However, if meaningful safety and quality standards can be adequately monitored and enforced, enlisting third parties could be a key component of a comprehensive forward-looking approach to regulating tools like DLeye.¹⁵³ If I am on the right track in proposing a forward-looking regulatory framework like the one outlined in this Section, a future piece can elaborate on whether a modified version of the Sentinel program, expanding the Joint Commission's scope, or some other enforcement scheme is most appropriate.

3. Production

Finally, a sufficiently detailed forward-looking plan will provide an explanation of the technology underling DLeye. At a high level, such a description will include illustrations of: (1) how tools like DLeye are designed; and (2) the safety and effectiveness specifications of such systems.

An illustration of a tool like DLeye's design specifications could closely resemble the disclosures required by the FDA's existing premarket approval requirements. Pursuant to those obligations, developers provide results from investigations, studies, and expert opinions demonstrating how the system will be developed to ensure its safety and effectiveness.¹⁵⁴ This description will likely include an evaluation of whether DLeye's capabilities and design features make it compatible with the workflows in which it is planned to be placed, metrics regarding its accuracy, the degree to which it satisfies patient

150. See Cortez, *supra* note 39, at 18–19, 22–23.

151. See *id.* at 22–23.

152. See *id.*; Findlay, *supra* note 146.

153. See Cortez, *supra* note 39, at 22–23; Feldman et al., *supra* note 128, at 414.

154. 21 U.S.C. § 360e(c)(1); 21 C.F.R. § 860.7 (2023).

and clinician expectations, any economic or operational efficiencies it will achieve, and more.¹⁵⁵

With respect to an assessment of the safety and effectiveness of a tool like DLeye, perhaps the most important factor is the quality of the data upon which the model was trained.¹⁵⁶ Tools like DLeye are typically “trained in high-resource settings,” such as state-of-the-art teaching hospitals or medical centers.¹⁵⁷ The high-quality data collection and analysis systems in these settings can go a long way in demonstrating that tools like DLeye were developed pursuant to well-controlled trials and studies.¹⁵⁸ However, training adaptive black-box systems in resource-rich settings does not guarantee that they will perform equally well in poorly resourced settings.¹⁵⁹ Granted, because retinas that DLeye examines look mostly the same across different patient populations, it is unlikely that small differences between retinal images will matter as much as, for example, variations in skin images, which can look quite “different depending on whether the skin is fair or not.”¹⁶⁰ Indeed, the lack of many medically significant distinctions between images of eyes is one of the reasons I have selected DLeye as a high-level example to illustrate its broader point. But other factors in a tool like DLeye’s training environment can matter a great deal.¹⁶¹

To better understand why this is the case, let’s assume that images of the same retina differ by only a few pixels depending on whether a high- or low-quality camera is used. Now imagine that DLeye is trained on images from high-quality cameras in a resource-rich setting like Massachusetts General Hospital. If DLeye is then deployed in a setting where only low-quality cameras are available, discrepancies in just a few pixels could cause the system to produce different outputs despite there being no medically significant distinction between the images taken by high- and low-quality cameras.¹⁶²

155. See, e.g., Trishan Panch et al., *Artificial Intelligence and Algorithmic Bias: Implications for Health Systems*, 9 J. GLOB. HEALTH, 020318, at 1–3 (Dec. 2019), <https://jogh.org/documents/issue201902/jogh-09-020318.pdf> [<https://perma.cc/72LD-DTKH>]; Kim et al., *supra* note 123, at 1400.

156. See Kim et al., *supra* note 123, at 1401; Price, *supra* note 130, at 66–67, 85–87.

157. Price, *supra* note 130, at 66–67.

158. *Id.* at 67, 85–87.

159. W. Nicholson Price II, *Distributed Governance of Medical AI*, 25 SMU SCI. & TECH. L. REV. 3, 3 (2022).

160. Price, *supra* note 130, at 94.

161. Robert Challen et al., *Artificial Intelligence, Bias and Clinical Safety*, 28 BMJ QUALITY & SAFETY 231, 232 (2019); Price, *supra* note 130, at 94; Babic et al., *supra* note 9, at 1203.

162. See Babic et al., *supra* note 9, at 1203–04.

Other factors that could contribute to bias will be important to consider and address in the forward-looking plan submitted to the FDA. For example, the various ways in which datasets might be biased could be revealed by validating the model through comparison to real-world data prior to integration in the clinical environment.¹⁶³ These validation efforts might include comparing DLeye's outputs to patient charts, conducting pilot trials using the tool on a subset of patients before implementing it in the care of a broader patient population, and more.¹⁶⁴ Throughout the validation process, developers should closely collaborate with clinicians to ensure the equity of models like DLeye by training them on large samples that include representation of underserved populations.¹⁶⁵ In pursuit of this goal, developers and clinicians might look to standards proposed by artificial intelligence researchers and the FDA's existing guidance on collecting electronic source data in clinical studies.¹⁶⁶ The Food, Drug, and Cosmetic Act also contains potentially helpful and informative detail regarding baseline standards for safety and quality in the pharmaceutical context.¹⁶⁷

163. Kim et al., *supra* note 123, at 1401.

164. *Id.*

165. *Id.*

166. Timnit Gebru et al., *Datasheets for Datasets*, COMM'NS ACM, Dec. 2021, at 86, 86–88; U.S. FOOD & DRUG ADMIN., GUIDANCE FOR INDUSTRY ELECTRONIC SOURCE DATA IN CLINICAL INVESTIGATIONS (2013), <https://www.fda.gov/downloads/drugs/guidances/ucm328691.pdf> [<https://perma.cc/N698-UD7G>]; *Facts About Current Good Manufacturing Practices (CGMP)*, U.S. FOOD & DRUG ADMIN. (Feb. 16, 2024), <https://www.fda.gov/drugs/pharmaceutical-quality-resources/facts-about-current-good-manufacturing-practice-cgmp> [<https://perma.cc/7RSH-A3JG>].

167. 21 U.S.C. §§ 301–399; *see also* 21 C.F.R. §§ 211.22, 211.68, 211.192, 211.194(a), 212.110(b) (2024). The FDA has also provided guidance to ensure that pharmaceutical companies take steps to ensure the integrity of data. Data Integrity and Compliance with Drug CGMP: Questions and Answers, 83 Fed. Reg. 64132 (Dec. 13, 2018).

Still other bias-related considerations could be the respective interests and goals of clinicians, developers, and health administrators. For instance, developers, clinicians, and administrators will all be interested in ensuring the safety, security, and effectiveness of DLeye by increasing the accuracy of the tool, reducing errors, and mitigating security risks. *See* Ian A Scott et al., *Exploring Stakeholder Attitudes Towards AI in Clinical Practice*, 28 BMJ HEALTH & CARE INFORMATICS, e100450, at 2 (Dec. 9, 2021), <https://informatics.bmj.com/content/bmjhci/28/1/e100450.full.pdf> [<https://perma.cc/V3MY-6N96>]. In addition to the goals shared by clinicians and administrators, developers will have the added aim of protecting any proprietary intellectual property involved in the software of systems like DLeye. *See, e.g.*, Price, *supra* note 9, at 435–37. Achieving efficiencies will also be a goal valued by clinicians, developers, and health administrators. But compared to clinicians, the focus of developers and health administrators will be slightly different. Developers, for example, will be interested in how efficiencies achieved by DLeye can lead to continued improvements in, and expanded use of, the tool. *See, e.g.*, FN Media Grp., *Global A.I. Healthcare Market Expected to Expand at CAGR of 36.4% From 2024 to 2030*,

In addition to anticipating and addressing the possibility of various biases, a forward-looking plan will need to account for cybersecurity risks.¹⁶⁸ To mitigate these risks, borrowing from cybersecurity practices like red teaming and adversarial attack testing, developers might illustrate in their forward-looking plan how they will collaborate with health systems and clinicians to conduct periodic stress testing of tools like DLeye.¹⁶⁹

At this point, it is worth noting that collaboration between developers and frontline clinicians will be key.¹⁷⁰ Relative to administrators, whose knowledge base will be most specialized with respect to the business considerations of utilizing artificial intelligence in the clinical context, clinicians and developers will have important and nuanced understandings of how tools like DLeye will be used in the care of patients and the vulnerabilities that could arise from *ex ante* or *ex post* design decisions.¹⁷¹

* * *

This illustration of a forward-looking plan for regulating a tool like DLeye is unavoidably “sketchy[] and sketchily argued.”¹⁷² Many nuances and details are omitted from this high-level overview.¹⁷³ Notwithstanding these and other important omissions, I am optimistic that this overview can improve our appreciation of what a forward-looking regulatory framework might look like. My modest aim in this Article is to provide a rough outline of the technical and social details to be included in forward-looking plans for using tools like DLeye in the clinical context. If these types of systems and processes are thoughtfully implemented, they could ensure that a realistic level of safety is observed while simultaneously obtaining the maximum benefit from medical software that evolves over time.

PR NEWSWIRE (Apr. 9, 2024, 8:45 AM), <https://www.prnewswire.com/news-releases/global-ai-healthcare-market-expected-to-expand-at-cagr-of-36-4-from-2024-to-2030--302111098.html> [<https://perma.cc/T9UK-QBZH>]. Similarly, health administrators will, in addition to valuing increased patient and clinician satisfaction, value the reputational advantages following from the operational efficiencies and cost savings brought about by using DLeye. Scott et al., *supra*, at 2–3.

168. See Gerke et al., *supra* note 9, at 3; Babic et al., *supra* note 9, at 1204.

169. Babic et al., *supra* note 9, at 1204.

170. Kim et al., *supra* note 123, at 1401.

171. *Id.*

172. Here, I borrow a turn of phrase from Christine Korsgaard. CHRISTINE M. KORSGAARD ET AL., *THE SOURCES OF NORMATIVITY* 91 (Onora O’Neill ed., 1996).

173. It does not, for example, comment on safety and effectiveness concerns related to privacy, compliance with regulatory requirements, and more. Kim et al., *supra* note 123, at 1402–03.

C. Statutory Limitations

Despite the promising benefits coinciding with a forward-looking approach like the one just outlined, efforts to implement such a plan will face a few hurdles.¹⁷⁴ Perhaps the most serious impediment is the conventional wisdom that the FDA serves as the gatekeeper for medical products but does not regulate the practice of medicine.¹⁷⁵ Indeed, that the FDA does not regulate the practice of medicine is a shibboleth that has been widely accepted for years by the American public, courts, medical practitioners, legislators, and even the FDA itself.¹⁷⁶ Many statutes and courts define the practice of medicine as the activity of diagnosing and treating disease.¹⁷⁷ Because implementing the sort of forward-looking approach just described would involve the FDA in the certification of medical professionals and authorization of clinical standards, it comes precariously close to regulating the practice of medicine. Hence, there is reason to think the FDA might need additional statutory authority to implement the forward-looking regulatory approach I recommend.¹⁷⁸

Despite the pervasive influence of this conventional wisdom, there are a few reasons to think it is misplaced. For one, there is no constitutional limitation on the FDA regulating the practice of medicine.¹⁷⁹ Granted, various statutory provisions include language indicating that the legislation is *not* intended to authorize the federal government to encroach on the autonomy of physicians.¹⁸⁰ For example, the Food, Drug, and Cosmetic Act stipulates that nothing in that law “shall be construed to limit or interfere with the authority of a health care practitioner to prescribe or administer any legally marketed device to a patient for any condition or disease.”¹⁸¹ Similar language is present in other statutes, such as the Drug Addiction Treatment Act of 2000 and the Food and Drug Administration Amendments Act of 2007.¹⁸² Notably, though, these provisions do not cite any constitutional limitation on the FDA’s

174. Gerke et al., *supra* note 9, at 3.

175. *Id.*; Patricia J. Zettler, *Toward Coherent Federal Oversight of Medicine*, 52 SAN DIEGO L. REV. 427, 430 (2015); *FDA’s Role in Regulating Medical Devices*, FDA (Aug. 31, 2018), <https://www.fda.gov/medical-devices/home-use-devices/fdas-role-regulating-medical-devices> [<https://perma.cc/3ADB-PY5R>].

176. Gerke et al., *supra* note 9, at 3; *see, e.g.*, Zettler, *supra* note 175, at 435–38.

177. Zettler, *supra* note 175, at 435–36.

178. *Id.* at 435–36; Gerke et al., *supra* note 9, at 3.

179. Lars Noah, *Ambivalent Commitments to Federalism in Controlling the Practice of Medicine*, 53 U. KAN. L. REV. 149, 168 (2004).

180. *Id.* at 165.

181. 21 U.S.C. § 396.

182. *Id.* § 823(g)(2)(H)(i); 42 U.S.C. § 247d-5a(d) (repealed 2016); Noah, *supra* note 179, at 166–67; *see also* Zettler, *supra* note 175, at 443.

authority to regulate the practice of medicine; indeed, it is doubtful whether any such prohibition exists.¹⁸³ Given the lack of a constitutional ban on the agency regulating the practice of medicine, these legislative provisions appear to reflect boilerplate language that was developed and reflexively reproduced in response to political pressure imposed on legislators by organized medicine, which has long had an interest in limiting the ability of Congress to infringe on physician autonomy.¹⁸⁴ Thus, in brief, despite language in statutes endorsing deference to professional autonomy, “[n]othing in the Constitution requires that [the FDA give] doctors . . . such a wide berth.”¹⁸⁵

In a similar vein, certain cases over the years have concluded that the federal government does not regulate the practice of medicine, but these holdings are often narrow and do not fully support the proposition that agencies like the FDA cannot oversee medical practitioners.¹⁸⁶ Additionally, some of the most important cases that could be construed as limiting federal authority to regulate the practice of medicine have since been rejected by courts.¹⁸⁷

Finally, contrary to the conventional wisdom that the federal government does not regulate the practice of medicine, there are multiple instances where it does exactly that—both directly and indirectly.¹⁸⁸ An example of direct regulation by the FDA is the Controlled Substances Act of 1970.¹⁸⁹ Pursuant to that Act, the federal government dictates which physicians may use controlled substances, and it restricts how such substances can be used in the practice of medicine.¹⁹⁰ Similarly, § 333(e) of the Federal Food, Drug, and Cosmetic Act prohibits practitioners from prescribing, dispensing, or administering Human Growth Hormone for any purpose other than those approved by the FDA.¹⁹¹ This restriction differs from the broad authority physicians have to prescribe most drugs approved by the FDA for off-label use.¹⁹²

183. See Noah, *supra* note 179, at 166–67.

184. See *id.* at 154–71; Zettler, *supra* note 175, at 441–46.

185. Noah, *supra* note 179, at 168.

186. See Zettler, *supra* note 175, at 438–40.

187. *Id.* at 440.

188. See *id.* at 454–66.

189. See Comprehensive Drug Abuse Prevention and Control Act of 1970, Pub. L. No. 91-513, §§ 100–709, 84 Stat. 1236, 1242–84 (codified as amended at 21 U.S.C. §§ 801–864, 871–904).

190. See 21 U.S.C. §§ 802(10), 812(b), 823.

191. *Id.* § 333(e).

192. Rebecca Dresser & Joel Frader, *Off-Label Prescribing: A Call for Heightened Professional and Government Oversight*, 37 J.L. MED. & ETHICS 476, 476 (2009).

In addition to these examples of the federal government directly regulating the practice of medicine, it has also done so indirectly.¹⁹³ One such instance is the FDA's authority to approve medical "drugs, devices, and biologics []before they enter the market."¹⁹⁴ "By determining what medical products may be sold" and used in the United States, the federal government indirectly regulates the options available to medical professionals in caring for patients.¹⁹⁵ Another example is the FDA's authority to impose heightened requirements on manufacturers to ensure the safe and effective use of certain drugs.¹⁹⁶ If the agency determines that the risks from a particular drug are severe enough, in addition to requiring manufacturers to provide warnings and use information, the FDA can mandate that manufacturers ensure medical professionals have adequate training to use the drug, restrict the areas in which the substance may be prescribed, and require that specific tests be conducted before the drug is distributed.¹⁹⁷

These and other examples indicate that, contrary to conventional wisdom, the federal government routinely conducts both direct and indirect regulation of the practice of medicine.¹⁹⁸ As Patricia Zettler suggests, the justification for doing so turns on a comparison of the potential benefits and drawbacks. Sometimes, patients' interests are best served by deferring to the judgment of medical professionals, who might be best positioned to weigh the potential risks and benefits of a particularly specialized medical practice or discreet clinical issue. Other times, though, the potential severity and scale of medical risks are large enough that federal actors are "better positioned than individual practitioners, professional organizations, or state governments to detect" and regulate the potential problems that might arise from certain medical practices.¹⁹⁹

I suggest that the use of adaptive tools in the care of patients constitutes a medical practice that warrants a forward-looking regulatory approach from the FDA. As the example concerning DLEYE illustrates, medical professionals are not solely responsible for the opaque outputs produced by dynamic artificial intelligence tools. Rather, physicians and other clinicians are members of a joint enterprise comprised of social actors—such as health systems, developers, and various other human medical professionals—interacting with technical systems comprised of medical software and

193. Zettler, *supra* note 175, at 460.

194. *Id.*

195. *Id.* at 461.

196. 21 U.S.C. § 355-1(a).

197. *See id.* § 355-1(f)(3)(A), (C), (D).

198. *See* Zettler, *supra* note 175, at 454–66.

199. *Id.* at 488.

hardware. Importantly, no single individual or entity exercises complete control over other members of these sociotechnical systems. As a result of the knowledge and power dispersion in these systems, we can't solely rely on the judgment of medical professionals to identify and respond to risks created by using adaptive systems. For these reasons, a forward-looking regulatory approach enforced by the FDA represents a promising means of overseeing the complex web of actors and interests comprising the sociotechnical systems responsible for leveraging adaptive tools in the clinical context.

Although these considerations indicate that the FDA has the authority to regulate the practice of medicine by adopting a forward-looking approach, the sheer inertia of conventional wisdom represents a significant barrier to doing so. Indeed, in September 2022, citing the agency's long-standing tradition of not regulating the practice of medicine, the FDA concluded that it lacked the statutory authority to implement a program that resembles the forward-looking approach I recommend.²⁰⁰ In addition to hesitation within the agency, the current political climate and makeup of the courts suggest that a critical mass of the general public, lawmakers, and judges will take a dim view of agencies wielding authority that is not explicitly granted by Congress.²⁰¹ For these reasons, as Gerke and others have noted, the FDA is unlikely to move in the near term toward a full-scale implementation of the forward-looking regulatory framework laid out in this Article.²⁰²

Nevertheless, the framework outlined in this Article can inform incremental progress toward a more forward-looking regulatory scheme if the political and judicial climate becomes more hospitable to doing so.

200. U.S. FOOD & DRUG ADMIN, THE SOFTWARE PRECERTIFICATION (PRE-CERT) PILOT PROGRAM: TAILORED TOTAL PRODUCT LIFECYCLE APPROACHES AND KEY FINDINGS 2–4 (2022), <https://www.fda.gov/media/161815/download> [<https://perma.cc/4K4Q-F5FH>]. A group of U.S. senators questioned the FDA's statutory authority to rely on the De Novo pathway for the Pre-Cert Program. *Senators Warren, Murray, and Smith Raise Further Questions about the FDA's Oversight of Digital Health Services*, ELIZABETH WARREN (Oct. 30, 2019), <https://www.warren.senate.gov/oversight/letters/-senators-warren-murray-and-smith-raise-further-questions-about-the-fdas-oversight-of-digital-health-devices> [<https://perma.cc/P3EA-NMP5>].

201. *See, e.g.*, *West Virginia v. EPA*, 597 U.S. 697 (2022); Daniel G. Aaron & Christopher L. Peterson, *Consumer Protection, Agencies, and the Supreme Court*, 5 JAMA HEALTH F., e240254, at 1–2 (Apr. 5, 2024), <https://jamanetwork.com/journals/jama-health-forum/fullarticle/2817282> [<https://perma.cc/HBP9-CRF3>].

202. Gerke et al., *supra* note 9, at 1.

II. THE CIVIL LIABILITY PROBLEM

At this point, it is worth noting that even if the FDA determines that it has the requisite authority to adopt my proposal, the agency will need to consider how such an approach would interact with existing tort doctrines. Civil liability applies to devices that have received FDA approval via abbreviated pathways.²⁰³ LumineticsCore, for example, was approved by the FDA as a Class II device pursuant to the agency's streamlined De Novo pathway.²⁰⁴ The benefit of bringing a tool like DLeye to market via a more abbreviated review process would be avoiding the rigorous and costly Class III pathway.²⁰⁵ The drawback to doing so is the possibility of facing state tort lawsuits even after seeking and receiving FDA approval. Thus, if DLeye's developers seek to bring the system to market via an abbreviated review process, they could face civil liability if the system results in patient harm.

However, for devices that are required to complete the FDA's full Class III premarket pathway because they have been assigned a higher risk classification, the benefit of navigating this more rigorous process is receiving protection from civil liability.²⁰⁶ The justification for this protection is rooted in the Supremacy Clause of the U.S. Constitution, which establishes that federal law in the United States "shall be the supreme Law of the Land."²⁰⁷ Pursuant to this clause, the 1976 Medical Devices Amendments preclude civil liability for injuries resulting from failures to observe requirements that are "different from" or "in addition to" those mandated by the federal government.²⁰⁸ The preemption protection afforded to such

203. Tschider, *supra* note 28, at 1576.

204. W. Nicholson Price II et al., *New Innovation Models in Medical AI*, 99 WASH. U. L. REV. 1121, 1127 (2022). The De Novo process is an expedited approval pathway available to developers who determine that there is no substantially equivalent device already on the market. Letter from Angela C. Krueger, Acting Deputy Dir., Eng'g & Sci. Rev., Off. of Device Evaluation, to Janice Hagan, Regul. Couns., Hogan Lovells US LLP 1–2 (Apr. 11, 2018), https://www.accessdata.fda.gov/cdrh_docs/pdf18/DEN180001.pdf [<https://perma.cc/KP4P-YX2M>]. The De Novo pathway allows devices for which there is no substantially equivalent device to avoid first submitting a 510(K) and instead ask the FDA to make a risk-based classification of the device under 513(a)(1) of the Food, Drug and Cosmetic Act. *Id.*; see also Aaron Gin & Bryan Helwig, *FDA Signals Fast-Track Approval for AI-Based Medical Devices*, BLOOMBERG L. (May 9, 2018), <https://news.bloomberglaw.com/tech-and-telecom-law/fda-signals-fast-track-approval-for-ai-based-medical-devices-1> [<https://perma.cc/3SF3-6NXF>].

205. Price, *supra* note 130, at 112.

206. *Id.*; Tschider, *supra* note 28, at 1574–76.

207. U.S. CONST. art. VI, § 2.

208. 21 U.S.C. § 360c; see also Tschider, *supra* note 28, at 1574–75 ("The preemption provision . . . made sense from the perspective of a fully regulating central agency, if Congress believed that the FDA would have the appropriate resources to fully regulate safety and efficacy

devices reflects a balance between incentivizing developers to expend resources to create higher-risk devices while maintaining plaintiffs' access to recovery for products that have undergone less comprehensive premarket processes.²⁰⁹

Given these provisions in the 1976 amendments, it is reasonable to think that if a tool like DLeye is assigned a higher-risk classification by the FDA and is therefore subject to the agency's most rigorous premarket approval standards, it would receive preemption protection.²¹⁰ However, Supreme Court precedent limiting the scope of preemption in this space suggests that DLeye could face civil liability despite being subject to a forward-looking review framework like the one I propose.

In *Medtronic, Inc. v. Lohr*, the Court indicated that civil liability imposed by state statute would be preempted for devices that had received a detailed review through the FDA's rigorous premarket approval process.²¹¹ Of note, the preemption protection established by *Lohr* does not extend to devices reviewed through more abbreviated processes, such as 510(k).²¹² In other words, *Lohr* grants heightened protection against civil liability for devices that have undergone more rigorous and comprehensive FDA review while allowing plaintiffs to seek recovery for injuries sustained by devices that have received more abbreviated scrutiny.²¹³

Subsequently, in *Riegel v. Medtronic, Inc.*, the Court extended the preemption protection established in *Lohr*.²¹⁴ Whereas the preemption reasoning in *Lohr* applied to "state statutes and regulations" that conflicted with federal requirements, it did not apply to "general duties enforced by common-law actions."²¹⁵ In *Riegel*, the Court expanded preemption protection to apply to common law tort claims that impose obligations that are different from or in addition to those required by the FDA.²¹⁶

The collective effect of these cases is that preemption protection from both state statutes and common law tort actions is afforded to devices that have undergone a sufficiently rigorous, device-specific premarket approval

for the medical device sector. . . . This preemption clause seemed, at its passage, to embrace the role of the FDA as primary regulator for the medical device industry by barring state laws that establish different or additional state requirements beyond federal requirements.").

209. Tschider, *supra* note 28, at 1572–76.

210. *Id.* at 1574–76; Price, *supra* note 130, at 112.

211. *See* 518 U.S. 470, 477–79, 494 (1996).

212. *Id.* at 492–94.

213. Tschider, *supra* note 28, at 1575–76.

214. 552 U.S. 312, 321–25 (2008).

215. *Lohr*, 518 U.S. at 489.

216. *Riegel*, 552 U.S. at 323–30.

process.²¹⁷ This means that medical devices that have undergone full product-based review by the FDA are exempt from civil liability, but medical software approved through a forward-looking approach of the sort I proposed in the previous Part lacks the same protection.²¹⁸

Unfortunately, the specter of civil liability raises a few thorny questions because artificial intelligence systems like DLeye are not only adaptive but also opaque. A tool like DLeye could be opaque for several reasons, which have been extensively discussed elsewhere.²¹⁹ Hence, in the interest of economy, I will focus on two that are most relevant to the DLeye hypothetical.

The opacity of DLeye could be the result of the fact that even though the algorithm relies on explicit rules, its calculations are too complex for even the most technically sophisticated humans to fully comprehend.²²⁰ The complexity of these systems could be the result of several factors.²²¹ But perhaps the most important factors concern variability between inputs and outputs.²²² For example, humans will have an easier time understanding how an algorithm works if the relationship between variables is monotonic, which is to say that for every increase or decrease in input, there is a corresponding and consistent increase or decrease in output.²²³ But our ability to understand

217. See *id.* at 322–23; see also *Lohr*, 518 U.S. at 477–79. Of note, neither *Lohr* nor *Riegel* support preemption for parallel claims, which concerns failure to comply with FDA requirements. See *Riegel*, 552 U.S. at 330; *Lohr*, 518 U.S. at 495.

218. See *Riegel*, 552 U.S. at 321–23; see also Tschider, *supra* note 28, at 1573; Evans, *supra* note 28, at 112.

219. See, e.g., Price & Rai, *supra* note 22, at 784–88; Andrew D. Selbst & Solon Barocas, *The Intuitive Appeal of Explainable Machines*, 87 *FORDHAM L. REV.* 1085, 1089–94 (2018). Included among the factors not discussed in this Article is that developers might wish to maintain a competitive advantage by keeping secret the algorithm’s source code or associated parameters, training processes or data, or some combination of these aspects. Price & Rai, *supra* note 22, at 788. Somewhat relatedly, the opacity of certain artificial intelligence systems could be due to a lack of specialized training. See, e.g., Jenna Burrell, *How the Machine ‘Thinks’: Understanding Opacity in Machine Learning Algorithms*, *BIG DATA & SOC’Y*, Jan.–June 2016, at 4, <https://journals.sagepub.com/doi/full/10.1177/2053951715622512> [<https://perma.cc/U4SH-TC7A>]; Selbst & Barocas, *supra*, at 1093–94. If the opacity is only due to the technical sophistication of a reviewer, though, presumably the opacity issue could be resolved by consulting a technical expert. Price & Rai, *supra* note 22, at 784–88.

220. See Price, *supra* note 16, at 430; see also Price & Rai, *supra* note 22, at 784–85; Lang et al., *supra* note 12, at 227.

221. Complexity theory can quickly take us into deep waters. See, e.g., MELANIE MITCHELL, *COMPLEXITY: A GUIDED TOUR* (2009). Fortunately, though, our present purposes don’t require that we delve into the details of this field.

222. Price & Rai, *supra* note 22, at 785.

223. Selbst & Barocas, *supra* note 219, at 1095 (citing *Monotonicity Function*, *CONCISE OXFORD DICTIONARY OF MATHEMATICS* (3d ed. 2014)).

and interpret the workings of a model dwindles when, for instance, there is a haphazard relationship between a steady upward input trend and outputs, the latter of which might go up or down without a predictable correlation to input changes.²²⁴ Even if an upward trend in inputs reliably produces output increases, humans can have difficulty understanding a model if it is non-linear or discontinuous.²²⁵ If, for instance, an input increase of two units in a model could produce output increases of more or less than two units, the model would be more difficult for humans to understand and predict.²²⁶ The number of variables considered by a model could also increase its opacity.²²⁷ The sheer scope of variables and dimensions analyzed by these models makes them difficult for humans to understand.²²⁸

[Machine learning (“ML”)] models consider hundreds or thousands of different variables and many ML methodologies, including most [deep learning] techniques, and ascribe weights to each identified variable. These weights reflect the degree to which changes in one variable affect the model’s output prediction. Consequently, developing a complete descriptive account of an ML algorithm’s functioning . . . is not achievable for complicated ML models such as deep neural networks.²²⁹

Yet another factor contributing to the opacity of systems like DLeye is the fact that they make decisions based on machine, rather than human, intuition.²³⁰ As layers of DLeye’s adaptive artificial neural networks make cascades of calculations to adjust how data is received and processed at various stages, each layer or cluster of neurons weights or encodes certain aspects of data (e.g., a wheel or mirror in an image) in unique ways relative to other neurons.²³¹ The result is neurons that learn intuitively and heterogeneously.²³² We cannot define or understand how a large, multilayered network of neurons learns by pointing to what any single neuron or group of neurons determined to be important.²³³ This is because even if we

224. *Id.*

225. *Id.*; Price & Rai, *supra* note 22, at 785.

226. Selbst & Barocas, *supra* note 219, at 1095.

227. *Id.* at 1095–96.

228. See Duran & Jongsma, *supra* note 12, at 329; see also Price & Rai, *supra* note 22, at 786–87; David Lehr & Paul Ohm, *Playing with the Data: What Legal Scholars Should Learn About Machine Learning*, 51 U.C. DAVIS L. REV. 653, 700 (2017).

229. Lang et al., *supra* note 12, at 227.

230. Selbst & Barocas, *supra* note 219, at 1096–98; Price & Rai, *supra* note 22, at 785–87.

231. Bathaee, *supra* note 83, at 901–03.

232. *Id.* at 902–03.

233. *Id.*

could develop a complete descriptive account of what various neurons are doing, our ability to understand why they produce certain outputs is limited.²³⁴ Learning to ride a bike, for example, requires repeated experience that can't be substituted with an explanation of how to balance on two wheels.²³⁵ In the same way, understanding the intuitive learning occurring inside systems like DLeye cannot be acquired via a step-by-step description of how and why certain neurons are weighting various data.²³⁶

To be sure, the black-box nature of the calculations systems like DLeye use to produce their outputs are a big part of what makes them so powerful and promising. At the same time, though, the opacity of these tools strains traditional tort doctrines.

A. Agency and Corporate Liability

If outputs from DLeye result in a patient's injury, hospitals or health systems—as principals—could be held liable for the actions of clinicians—as agents—using adaptive black-box tools in the care of patients.²³⁷ Ordinarily, employers are considered principals and their employees are

234. The interrelationship between features that deep learning models use to make predictions or take actions reflect connections and patterns that the algorithm has identified and determined to be meaningful, and these relationships between features “do not necessarily reflect phenomenon that are human-comprehensible.” Lang et al., *supra* note 12, at 227.

235. Bathaee, *supra* note 83, at 902.

236. *See id.* at 901–03; *see also* GILBERT RYLE, *THE CONCEPT OF MIND* 14–47 (Routledge 2009) (1949).

237. *See* Price, *supra* note 28, at 295, 303; Efthimios Parasidis, *Clinical Decision Support: Elements of a Sensible Legal Framework*, 20 J. HEALTH CARE L. & POL'Y 183, 213–14 (2018). Some have suggested that it might be appropriate to regard artificial intelligence systems as agents of corporations. *See* Anat Lior, *AI Entities as AI Agents: Artificial Intelligence Liability and the AI Respondeat Superior Analogy*, 46 MITCHELL HAMLINE L. REV. 1043, 1071 (2020); Matthew U. Scherer, *Of Wild Beasts and Digital Analogues: The Legal Status of Autonomous Systems*, 19 NEV. L.J. 259, 262–63 (2018). Without wading too far into the dialectic concerning this issue, I suggest that artificial intelligence currently lacks the degree of autonomy required to be regarded as a legal person in a way that would be practically and theoretically needed for the purposes discussed in this Article. *See* SAMIR CHOPRA & LAURENCE F. WHITE, *A LEGAL THEORY FOR AUTONOMOUS ARTIFICIAL AGENTS* 159–160 (2011); *see also* Lior, *supra*, at 1071; Scherer, *supra*, at 262–63; Lawrence B. Solum, *Legal Personhood for Artificial Intelligence*, 70 N.C. L. REV. 1231, 1239–40 (1992). Perhaps AI could manage a fund in the market to make compensation available for injured patients. But we would still need the people behind the system to represent the tool in court and maintain/continue developing the tool. So even if we sued the artificial intelligence system, we would effectively be suing the same sort of enterprise I subsequently outline. *See infra* Part III. Further, the opacity and adaptiveness of the artificial intelligence systems discussed in this Article makes them somewhat different than ships, corporations, or other entities that have traditionally been regarded as legal persons.

regarded as agents because the former task the latter with acting on their behalf.²³⁸ In the healthcare context, however, agency liability is complicated by the fact that, although the majority of physicians are technically designated as “staff” of the hospitals and health systems with which they are affiliated, physicians have traditionally been independent contractors rather than employees of such enterprises.²³⁹ Physicians admit patients to hospitals and treat patients therein.²⁴⁰ Historically, though, physicians have often been self-employed or members of physician groups contracting with hospitals and, therefore, are not agents of hospitals in the traditional sense.²⁴¹ But today, an increasing number of physicians (indeed, the majority by some counts) are employed by hospitals.²⁴² This makes holding hospitals liable for the conduct of their agents a bit more straightforward in some circumstances.

One approach is to hold health systems or hospitals vicariously liable for the negligent actions of their agents.²⁴³ However, the opacity of adaptive black-box systems complicates the application of this doctrine. Perhaps the most important unanswered question is which entity should be considered an agent’s principal for liability purposes. According to the *Restatement (Third) of Agency*, a common law “principal” is defined as a party that “has authorized another to act on [their] account and subject to [their] control.”²⁴⁴ Adaptive tools like DLeve evolve with each new patient encounter and datum added; and because these tools are black boxes, human clinicians cannot fully understand the opaque calculations used to produce their outputs.²⁴⁵ This makes it challenging to determine whether the hospital or health system that

238. See RESTATEMENT (THIRD) OF AGENCY § 1.01 cmt. g (AM. L. INST. 2006).

239. ARTHUR F. SOUTHWICK, *THE LAW OF HOSPITAL AND HEALTH CARE ADMINISTRATION* 213–14 (1978); BARRY R. FURROW ET AL., *HEALTH LAW: CASES, MATERIALS AND PROBLEMS* 441, 445 (6th ed. 2008); Kenneth S. Abraham & Paul C. Weiler, *Enterprise Medical Liability and the Evolution of the American Health Care System*, 108 HARV. L. REV. 381, 387 (1994).

240. FURROW, *supra* note 239, at 441; see Abraham & Weiler, *supra* note 239, at 387.

241. SOUTHWICK, *supra* note 239, at 198, 213; Abraham & Weiler, *supra* note 239, at 387.

242. See Press Release, Am. Med. Ass’n, *AMA Analysis Shows Most Physicians Work Outside of Private Practice* (May 5, 2021), <https://www.ama-assn.org/press-center/press-releases/ama-analysis-shows-most-physicians-work-outside-private-practice> [https://perma.cc/LDU4-PGD2].

243. Sharon Hoffman & Andy Podgurski, *E-Health Hazards: Provider Liability and Electronic Health Record Systems*, 24 BERKELEY TECH. L.J. 1523, 1535–36 (2009) (“The doctrine of ‘respondeat superior,’ which literally means ‘let the superior answer,’ establishes that employers are responsible for the acts of their employees in the course of their employment.”).

244. RESTATEMENT (THIRD) OF AGENCY § 1 cmt. c (AM. L. INST. 2006) (emphasis added).

245. See Senger and O’Leary, *supra* note 16, at 292; see also Mark A. Chinen, *The Co-Evolution of Autonomous Machines and Legal Responsibility*, 20 VA. J.L. & TECH. 338, 360 (2016) (describing the difficulty of predicting the processes and outputs of computer driven devices).

hired clinicians responsible for using DLeye, clinicians themselves, software developers, or some other party is in control of the adaptive black-box system resulting in injury.

In addition to the question of control, establishing negligence for the use of tools like DLeye will be challenging. In the DLeye example discussed earlier, the standard of care was for clinicians to rely on the system's output and only conduct their own review if disease is detected. They did not act negligently by failing to review the image that DLeye did not flag as indicating disease. Because the question of negligence is at the heart of vicarious liability, this doctrine is ill-suited to address the sort of harms arising from the use of systems like DLeye.²⁴⁶

Alternatively, health systems could be held liable through the doctrine of apparent authority.²⁴⁷ The rationale for apparent authority is perhaps clearest in the emergency room, where patients only know that hospitals have arranged for all the attending professionals to be there.²⁴⁸ As a result of such surroundings, courts have held that it is reasonable for patients to assume that hospitals are responsible for the negligence of physicians, whom hospitals have made available to care for patients.²⁴⁹ Thus, in effect, actual agents and those acting with apparent authority create the same degree of liability for principals.²⁵⁰

With these considerations in mind, it seems reasonable to hold hospitals or health systems liable pursuant to the doctrine of apparent authority. After all, clinicians caring for patients in a hospital setting appear to be acting as actual or apparent agents of health systems or hospitals, and injury caused by their use of tools like DLeye seems to be the sort of tortious act that would trigger liability.²⁵¹

But I suggest that we have a few reasons to reject the application of this doctrine to patient injuries resulting from the use of tools like DLeye. As a normative matter, health systems are not on an island in the complex web of actors and institutions comprising the sociotechnical systems responsible for

246. See Maliha et al., *supra* note 10, at 633.

247. SOUTHWICK, *supra* note 239, at 199–200, 214, 379; Abraham & Weiler, *supra* note 239; Arthur F. Southwick, *Hospital Liability: Two Theories Have Been Merged*, 4 J. LEGAL MED. 1, 9–13 (1983).

248. See SOUTHWICK, *supra* note 239, at 196, 199–200; Abraham & Weiler, *supra* note 239, at 388.

249. SOUTHWICK, *supra* note 239, at 199–200; Abraham & Weiler, *supra* note 239, at 388.

250. See 3 AM. JUR. 2D *Agency* §§ 14–16, 65–71, 73–75, Westlaw (database updated Aug. 2024).

251. See RESTATEMENT (THIRD) OF AGENCY §§ 7.04, .06, .08 (AM. L. INST. 2006).

leveraging adaptive black-box tools in the clinical context.²⁵² Holding health systems solely liable would not reflect the moral ledger of the clinicians and developers contributing to the use of DEye in the care of patients. Of course, if health systems are held solely liable, they could seek indemnity and contribution from these other actors. However, for reasons to be discussed in Section III, it will be practically challenging for health systems to succeed in doing so.²⁵³ Also discussed in Section III are reasons why clinicians and developers might be resistant to health systems assuming sole liability for patient harms.²⁵⁴ Hence, a different approach seems warranted.

An alternative to vicarious liability and apparent authority is the long-recognized doctrine of corporate liability.²⁵⁵ Whereas agency liability holds hospitals or health systems liable for the negligence of their agents even when the former are not at fault,²⁵⁶ corporate liability concerns negligence on the part of health systems or hospitals.²⁵⁷ Building on the leading 1965 case *Darling v. Charleston Community Memorial Hospital*, courts have held hospitals responsible for failing to maintain a clinical environment that is conducive to providing safe and effective care to patients.²⁵⁸ Discharging this duty requires hospitals to, for example, establish reasonable accreditation standards, comply with state licensing and safety standards, and investigate the credentials and qualifications of the physicians with whom they have contracted to treat patients.²⁵⁹ Today, the Joint Commission on the Accreditation of Health Care Organizations conducts periodic assessments to ensure healthcare facilities observe these quality and safety standards.²⁶⁰

If the use of DEye results in patient injury because of improper physician operation of the tool, health systems or hospitals might be liable for failing to

252. See *infra* Section II.C.

253. See *infra* Section III.A.

254. See *infra* Section III.A. For example, developers and clinicians might worry about the reputational damage that could result if hospitals are inclined to settle rather than fight claims. What's more, the enterprise to which liability is assigned might seek to reduce injuries through micromanagement or pass the cost of harms onto developers and clinicians in the form of reduced compensation.

255. Abraham & Weiler, *supra* note 239, at 389.

256. Maliha et al., *supra* note 10, at 633; Abraham & Weiler, *supra* note 239, at 391.

257. SOUTHWICK, *supra* note 239, at 212, 349–50; Abraham & Weiler, *supra* note 239, at 391.

258. Abraham & Weiler, *supra* note 239, at 390 (citing *Darling v. Charleston Cmty. Mem'l Hosp.*, 211 N.E.2d 253 (Ill. 1965)); I. Trotter Hardy, Jr., *When Doctrines Collide: Corporate Negligence and Respondeat Superior When Hospital Employees Fail to Speak Up*, 61 TUL. L. REV. 85, 90–91 (1986).

259. Hardy, Jr., *supra* note 258, at 90–91; Abraham & Weiler, *supra* note 239, at 389.

260. Abraham & Weiler, *supra* note 239, at 390.

verify that their agents possessed the requisite skill to properly use the tool.²⁶¹ Again, though, the problem of negligence arises.²⁶² In the example concerning DLeve, the patient injury might have occurred despite the hospital in which the system was used observing all the relevant quality and safety standards, including verifying the credentials of clinicians tasked with using DLeve. The opaque and adaptive nature of black-box technology makes it uniquely challenging to determine “the relative responsibilities of hospitals and developers in training physicians and developing or enforcing protocols for the” general or specific use of a tool like DLeve.²⁶³ For these reasons, vicarious liability, apparent authority, and corporate liability are not well-suited to resolve thorny questions concerning what constitutes fault—let alone which party has acted negligently—when outputs from adaptive black-box systems harm patients.

B. Products Liability

Products liability is an alternative to agency and corporate liability.²⁶⁴ Plaintiffs are entitled to recovery when they are injured by devices or products that are defective in either manufacturing, design, or warning.²⁶⁵

A product: (a) contains a manufacturing defect when the product departs from its intended design even though all possible care was exercised in the preparation and marketing of the product; (b) is defective in design when the foreseeable risks of harm posed by the product could have been reduced or avoided by the adoption of a reasonable alternative design by the seller or other distributor, or a predecessor in the commercial chain of distribution, and the omission of the alternative design renders the product not reasonably safe; (c) is defective because of inadequate instructions or warnings when the foreseeable risks of harm posed by the product could have been reduced or avoided by the provision of reasonable instructions or warnings by the seller or other distributor, or a predecessor in the commercial chain of distribution, and the

261. See RESTATEMENT (THIRD) OF AGENCY § 7.05 cmt. b (AM. L. INST. 2006).

262. See Abraham & Weiler, *supra* note 239, at 391.

263. Nicolas Terry, *Of Regulating Healthcare AI and Robots*, 21 YALE J.L. & TECH (SPECIAL ISSUE) 133, 162 (2019).

264. See Evans & Pasquale, *supra* note 28, at 30.

265. See Maliha et al., *supra* note 10, at 633–34; A. Mitchell Polinsky & Steven Shavell, *The Uneasy Case for Product Liability*, 123 HARV. L. REV. 1438, 1453–54 nn.59–62 (2010); Alexander B. Lemann, *Autonomous Vehicles, Technological Progress, and the Scope Problem in Products Liability*, 12 J. TORT L. 157, 181 (2019); Parasidis, *supra* note 237, at 216, 218–19.

omission of the instructions or warnings renders the product not reasonably safe.²⁶⁶

Determining whether a device contains a manufacturing defect does not pose a unique challenge for the transparent components of any artificial intelligence system. Developers of a tool like DLeye could be liable for a manufacturing defect if they fail to engineer the system according to its intended specifications.²⁶⁷ Assigning liability for dynamically inscrutable systems becomes more challenging, though, when the transparent aspects of such systems are manufactured without any defects but the system nevertheless produces an unexpected output that results in patient injury.²⁶⁸ In such instances, plaintiffs seeking to establish a claim for products liability will need to prove the existence of a deficiency in either design or warning.

With respect to design defects, one way that DLeye's developers could be held liable is if the harm could have been avoided by adopting additional precautions or an alternative design, and the cost of undertaking such alternatives is foreseeably lower than the expected safety benefits.²⁶⁹ However, it will be challenging for plaintiffs to satisfy this foreseeability requirement. Whereas locked systems like LumineticsCore perform tasks pursuant to a detailed set of instructions, tools like DLeye are adaptive. Coinciding with the enormous potential benefits of tools like DLeye is the possibility that they will produce unexpected outputs.²⁷⁰ If it is difficult to foresee the potential risks and benefits resulting from these systems, it will be hard for plaintiffs to succeed in arguing that the cost of an alternative design would be less than the expected upshots from adopting such a design.²⁷¹

Plaintiffs could avoid these challenges if they are successful in pursuing an alternative strategy for establishing the existence of a design defect: demonstrating that the danger of the product exceeds consumer expectations.²⁷² Again, though, the problem of foreseeability arises. Unexpected results are an inescapable byproduct of dynamically inscrutable

266. RESTATEMENT (THIRD) OF TORTS: PRODS. LIAB. § 2 (AM. L. INST. 1998).

267. *Id.* § 2(a).

268. Omri Rachum-Twaig, *Whose Robot Is it Anyway?: Liability for Artificial-Intelligence-Based Robots*, 2020 U. ILL. L. REV. 1141, 1155; Marguerite E. Gerstner, Comment, *Liability Issues with Artificial Intelligence Software*, 33 SANTA CLARA L. REV. 239, 245–48 (1993).

269. RESTATEMENT (THIRD) OF TORTS: PRODS. LIAB. § 2(b) cmt. a (AM. L. INST. 1998); David G. Owen, *Design Defects*, 73 MO. L. REV. 291, 310–11 (2008).

270. Rachum-Twaig, *supra* note 268, at 1156.

271. See Owen, *supra* note 269, at 310–15.

272. RESTATEMENT (SECOND) OF TORTS § 402A cmt. g (AM. L. INST. 1965); Owen, *supra* note 269, at 300–01.

systems.²⁷³ For this reason, it seems contradictory to engineer these systems to learn and innovate in ways that are opaque to humans while simultaneously assigning liability for the lack of foreseeability resulting from such design choices.²⁷⁴

In response to these challenges, Frank Pasquale and Barbara Evans propose revising some of the principles underlying liability standards for design defects in artificial intelligence systems.²⁷⁵ They suggest recognizing a design defect for software that is trained on datasets that are “too small, inappropriate, inaccurate, or biased and nonrepresentative of patients the software will analyze.”²⁷⁶

Although this is an interesting proposal that seems to provide a promising framework for many contexts, it does not address the unique challenges posed by using adaptive black-box tools. The evolving and opaque nature of such systems makes it challenging to know whether the patient’s injury was caused by some problematic feature in the dataset on which the software was trained; absent a defect in the dataset, the harm could have resulted from some aspect within the cascade of calculations responsible for producing the system’s outputs. Of course, it is possible to avoid this concern by holding developers strictly liable for any shortcoming in the software once it leaves their possession.²⁷⁷ However, as Section II.C and Part III will explain, given the complex ecosystem of multiple parties responsible for designing, training, and using adaptive black-box tools, it is inappropriate to single-out developers as being solely accountable for harm resulting from the use of such tools.²⁷⁸ Even if a defect could be identified, it will not be easy to

273. Steven J. Frank, *Tort Adjudication and the Emergence of Artificial Intelligence Software*, 21 SUFFOLK U. L. REV. 623, 638–39 (1987); Surden & Williams, *supra* note 10, at 150–63; Andrew D. Selbst, *Negligence and AI’s Human Users*, 100 B.U. L. REV. 1315, 1323–25 (2020).

274. Rachum-Twaig, *supra* note 268, at 1156.

275. Evans & Pasquale, *supra* note 28, at 32–33; *see also* Frank Pasquale, *Data-Informed Duties in AI Development*, 119 COLUM. L. REV. 1917, 1921, 1923–37 (2019).

276. Evans & Pasquale, *supra* note 28, at 32–33.

277. *Id.*

278. Of note, the shortcomings Evans and Pasquale have in mind concern software that purports, but fails, to be explainable. Evans & Pasquale, *supra* note 28, at 32 (“Developers seemingly could escape FDA regulation by simply asserting that they intend for software to be explainable (whether or not it actually is) and by labeling the software as ‘not intended for use without independent review by a health care professional’ and ‘not intended to serve as the primary basis for making a clinical diagnosis or treatment decision regarding an individual patient.’ . . . This in turn would let them argue that their software is a service, rather than an FDA-regulated product subject to product liability.”). But the concern with systems like DEye is that they are not designed to be explainable because of the significant benefits that could result from

determine which person or entity is responsible for a design flaw when multiple parties are involved in the design and manufacturing of a black-box tool.²⁷⁹

Finally, for reasons related to the challenges regarding claims for a design defect, establishing a warning deficiency will also be difficult. Like design defects, the *Restatement (Third) of Torts* indicates that products contain a warning defect “only when the risks are foreseeable.”²⁸⁰ Because adaptive black-box tools like DLEYe continually evolve, it will be difficult to predict the sort of risks that might result from its outputs as the system continues evolving over time.²⁸¹ Thus, it is doubtful whether products liability is a good candidate for responding to adaptive black-box tools.

C. Medical Malpractice

An alternative to the foregoing liability systems is medical malpractice,²⁸² which requires plaintiffs to establish: (1) duty; (2) breach; (3) causation; and (4) damages.²⁸³ For harm caused by tools like DLEYe, breach is the most fundamental element.²⁸⁴

their powerful, albeit opaque, capabilities. Hence, for reasons indicated in Section II.C, it seems inappropriate to single out developers as being solely liable for harms that result from the coordinated activities of multiple parties.

279. See Gerstner, *supra* note 268, at 247–49; see also Kyle Graham, *Strict Products Liability at 50: Four Histories*, 98 MARQ. L. REV. 555, 600–01 (2014).

280. RESTATEMENT (THIRD) OF TORTS: PRODS. LIAB. § 2 cmt. a (AM. L. INST. 1998).

281. Rachum-Twaig, *supra* note 268, at 1156.

282. See Price, *supra* note 9, at 298–303; VICTOR SCHWARTZ ET AL., PROSSER, WADE, AND SCHWARTZ’S TORTS: CASES AND MATERIALS 188, 190, 201–02 (13th ed. 2015); Maliha et al., *supra* note 10 at 632; Froomkin et al., *supra* note 14 at 51, 53.

283. See David A. Hyman & Charles Silver, *Medical Malpractice and Compensation in Global Perspective: How Does the U.S. Do It?*, 87 CHI.-KENT L. REV. 163, 168 (2012).

284. See *id.* at 168. For instance, because a clinician–patient relationship is usually considered sufficient to establish the existence of a duty, we don’t need to spend too much time on the first element. Maliha et al., *supra* note 10, at 632 (“First, courts have allowed malpractice claims to proceed against health professionals in cases where there were mistakes in medical literature given to patients or when a practitioner relies on an intake form that does not ask for a complete history. Second, courts have generally been reluctant to allow physicians to offset their liability when a pharmaceutical company does not adequately warn of a therapy’s adverse effect. Third, most courts have been unwilling to use clinical guidelines as definitive proof of the standard of care and require a more individualized determination in a particular case. Fourth, many courts are disinclined to excuse malpractice based on errors by system technicians or manufacturers.”).

We also don’t need to take up the question of damages, which can be hotly contested and implicates thorny questions about punitive versus compensatory damages, whether compensatory

Although the appropriate standard for using new adaptive black-box tools in the care of patients will initially be an open question, the example concerning DEye presumed that the use of that tool satisfied the standard of care.²⁸⁵ Even if a new technology becomes part of the standard of care at one point in time, though, precisely what such standards require will need to be intermittently revised as the tool evolves.²⁸⁶ Periodically updating such standards will be practically and doctrinally difficult in part because U.S. case law has not developed a unified set of principles for determining how to incorporate new medical technologies into the relevant standard of care.²⁸⁷

Notwithstanding the fact that judicial precedent “on physician use of black-box technology is not yet well developed, several lines of cases suggest that physicians bear the burden of errors that result from black-box outputs.”²⁸⁸ But even if case law indicates that clinicians will be held liable,

damages are economic or non-economic, and more. See SCHWARTZ ET AL., *supra* note 282, at 169–70; see also David A. Hyman & Charles Silver, *Medical Malpractice, Litigation and Tort Reform: It’s the Incentives, Stupid*, 59 VAND. L. REV. 1085, 1117 (2006). For our purposes, we don’t need to wade into the damages thicket because this element of a malpractice claim follows once breach and injury have been established. In this Article, we are primarily concerned with the question of whether clinicians have injured patients by breaching their duty of care.

Even if all the other elements of a malpractice claim are satisfied, questions remain regarding causation, which can be a challenging element to satisfy. See Bathaee, *supra* note 83, at 923; Selbst, *supra* note 273, at 1361–62. At its core, this element requires plaintiffs to demonstrate that a reasonable person would foresee that the conduct in question could result in the plaintiff’s injury. See H.L.A. HART & TONY HONORE, *CAUSATION IN THE LAW* 255 (2d ed. 1985). The thought animating this doctrine is that defendants should not be held liable for events that they could not have prevented through some reasonably foreseeable exercise of their agency. *Id.* Holding clinicians liable for risks of harm that they reasonably should have foreseen will incentivize them to either take appropriate precautions or not participate in risky behavior altogether. See Mark F. Grady, *Proximate Cause Decoded*, 50 UCLA. L. REV. 293, 294 (2002). The opaque nature of black-box systems makes it challenging for the technology’s creator—let alone clinicians using the tool—to foresee risks resulting from the system’s decisions or outputs. See Selbst, *supra* note 273, at 1331–33. For example, black-box tools might “make decisions based on higher-dimensional relationships between variables that no human can visualize.” See Bathaee, *supra* note 83, at 924. Thus, it can be difficult for any user or creator of such technologies to foresee the cascades of calculations informing a decision or recommendation that led to a patient’s injury. See SCHWARTZ ET AL., *supra* note 282, at 193–94, 198; Bathaee, *supra* note 83, at 924.

285. See Michael D. Greenberg, *Medical Malpractice and New Devices: Defining an Elusive Standard of Care*, 19 HEALTH MATRIX 423, 428, 434 (2009); Froomkin et al., *supra* note 14, at 51; Hyman & Silver, *supra* note 283, at 168; Parasidis, *supra* note 237, at 212–13.

286. See Froomkin et al., *supra* note 14, at 51.

287. See Greenberg, *supra* note 285, at 434.

288. Maliha et al., *supra* note 10, at 632. Even if the breach element is satisfied, questions remain regarding causation, which can be a challenging element to satisfy. See Bathaee, *supra* note 83, at 923; Selbst, *supra* note 273, at 1361–62. At its core, this element requires plaintiffs to

it seems unreasonable to ask clinicians to bear the burden of liability for injuries resulting from a technology like DLeye. As previously illustrated, the use of an adaptive and opaque system like DLeye is made possible by the cooperation of larger ecosystem of health systems, their agents, and various vendors providing a range of services.²⁸⁹ This means that as a practical matter, clinicians' liability is "inextricably linked to the liability of these other actors" through a complicated web of contractual and business relationships.²⁹⁰

In addition to these practical considerations, the complex mosaic of interested parties illustrated in Part I also indicates that, as a normative matter, the moral ledger of clinicians does not reflect sole liability for patient harms resulting from tools like DLeye. The operation of tools like DLeye depends upon the cooperation of clinicians, administrators, developers, and others. Each of these parties have unique skillsets and roles they bring to bear in the initial and ongoing operation of systems like DLeye. As a result of this collaborative effort, each bears at least some responsibility when the use of such systems results in patient harm. Individual clinicians, software engineers, and entire insurance and health systems are not on an island with regard to their responsibility for the development and use of adaptive black-box systems. For these reasons, holding any one of these actors solely liable does not do a good job reflecting the accountability of each party in the larger ecosystem providing a range of services needed to use dynamic and opaque artificial intelligence systems in the care of patients.

demonstrate that a reasonable person would foresee that the conduct in question could result in the plaintiff's injury. See HART & HONORE, *supra* note 284, at 255. The thought animating this doctrine is that defendants should not be held liable for events that they could not have prevented through some reasonably foreseeable exercise of their agency. *Id.* Holding clinicians liable for risks of harm that they reasonably should have foreseen will incentivize them to either take appropriate precautions or not participate in risky behavior altogether. See Mark F. Grady, *Proximate Cause Decoded*, 50 UCLA L. REV. 293, 294 (2002). The opaque nature of black-box systems makes it challenging for the technology's creator—let alone clinicians using the tool—to foresee risks resulting from the system's decisions or outputs. See Selbst, *supra* note 273, at 1331–33. For example, black-box tools might "make decisions based on higher-dimensional relationships between variables that no human can visualize." See Bathaee, *supra* note 83, at 924 (citations omitted). Thus, it can be difficult for any user or creator of such technologies to foresee the cascades of calculations informing a decision or recommendation that led to a patient's injury. See SCHWARTZ ET AL., *supra* note 282, at 193–94, 198; Bathaee, *supra* note 83, at 924.

289. See *supra* Part I; see also Maliha et al., *supra* note 10, at 630.

290. Maliha et al., *supra* note 10, at 630.

III. THE PROMISE OF ENTERPRISE LIABILITY

In response to the doctrinal challenges discussed earlier,²⁹¹ some have suggested leveraging a version of enterprise liability to address patient injuries resulting from the use of tools like DLeve.²⁹² Because enterprise liability has a long and contested history in medicine, I will first provide a bit of background regarding the doctrine's proposed application to medical injuries and illustrate some of the usual objections to doing so. I will then explain how a "first cousin" of enterprise liability—known as common enterprise liability—can avoid these traditional concerns and supplement the forward-looking regulatory approach I propose in Section I.B.

A. Traditional Versions of Enterprise Liability

Around the middle of the twentieth century, the doctrine of enterprise liability emerged as a break from the traditional allocation of liability for injuries in medicine.²⁹³ Prior to the 1940s, hospitals were not held responsible for the negligent acts of physicians treating patients at their facilities.²⁹⁴ Instead, physicians generally bore sole responsibility for medical injuries caused by their negligence.²⁹⁵ There were a few rationales for these traditional limitations on liability.

The first can be traced to the doctrine of charitable immunity. Before the middle of the twentieth century, courts regarded hospitals as charitable institutions immune from liability for the negligence of their salaried physicians.²⁹⁶ Although the main justification for this doctrine was that patients waived their right to sue for malpractice by receiving care free of charge, immunity applied to care provided both to paying and non-paying patients.²⁹⁷ The immunity for charitable services began eroding with the D.C. Court of Appeals' 1942 decision in *Georgetown College v. Hughes*.²⁹⁸ In that case, the court took a dim view of the notion that hospitals could fund a

291. See *infra* Part II.

292. See Chan, *supra* note 29, at 32; Schweikart, *supra* note 29.

293. See Abraham & Weiler, *supra* note 239, at 385.

294. *Id.* at 385.

295. Randall R. Bovbjerg & Robert Berenson, *Enterprise Liability in the Twenty-First Century*, in *MEDICAL MALPRACTICE IN THE U.S. HEALTH CARE SYSTEM* 227, 230 (2006).

296. David H. Rutchik, *The Emerging Trend of Corporate Liability: Courts' Uneven Treatment of Hospital Standards Leaves Hospitals Uncertain and Exposed*, 47 *VAND. L. REV.* 535, 551 (1994).

297. Abraham & Weiler, *supra* note 239, at 385.

298. See Southwick, *supra* note 247, at 1; Abraham & Weiler, *supra* note 239, at 385 (citing *Georgetown Coll. v. Hughes*, 130 F.2d 810 (D.C. Cir. 1942)).

variety of operating costs without needing to insure against liability for negligently injured patients.²⁹⁹ In the years following that decision, courts continued the trend of eroding the charitable immunity doctrine.³⁰⁰

Although removal of charitable immunity for hospitals was well underway following the *Hughes* decision, additional layers of protection for hospitals remained in the form of the conventional distinction between administrative and medical errors.³⁰¹ Pursuant to the doctrine of respondeat superior, employers are generally liable for their employees' acts because the former exercise control over the latter's activities.³⁰² However, courts did not traditionally hold hospitals liable for the negligent acts of physicians treating patients at their facilities.³⁰³ Instead, physicians generally bore sole responsibility for medical injuries caused by their negligence.³⁰⁴ This exception to the traditional application of respondeat superior in the medical context can be explained by the fact that, prior to the middle of the twentieth century, hospital administrators were only responsible for facilitating *administrative* services, whereas physicians were responsible for providing *medical* services to patients.³⁰⁵ "For example, giving blood transfusions to the wrong patient was labeled an 'administrative error,' but giving the wrong blood to the right patient was a 'medical error.'"³⁰⁶ Because hospital administrators lacked the medical expertise needed to control physicians' discretionary acts, physicians were regarded as ship captains who were responsible for their own actions and those of hospital staff, such as nurses, operating under their direction.³⁰⁷ Hospitals were widely considered to be the

299. Abraham & Weiler, *supra* note 239, at 385–86.

300. William H. Payne, *Recent Developments Affecting a Hospital's Liability for Negligence of Physicians*, 18 S. TEX. L.J. 389, 390 (1976); Abraham & Weiler, *supra* note 239, at 386.

301. Abraham & Weiler, *supra* note 239, at 386; Arthur F. Southwick, *The Hospital as an Institution: Expanding Responsibilities Change Its Relationship with the Staff Physician*, 9 CAL. W. L. REV. 429, 440 (1973).

302. Rory Van Loo, *The Revival of Respondeat Superior and Evolution of Gatekeeper Liability*, 109 GEO. L.J. 141, 143, 148 (2020); William M. Sage, *Enterprise Liability and the Emerging Managed Health Care System*, 60 LAW & CONTEMP. PROBS. 159, 173 (1997).

303. Abraham & Weiler, *supra* note 239, at 385–86; *see also* Southwick, *supra* note 301, at 430–31 (observing that hospitals could be held legally responsible for non-treatment tasks involved in the delivery of care, even though hospitals were not liable for negligent treatment).

304. Bovbjerg & Berenson, *supra* note 295, at 21.

305. Southwick, *supra* note 301, at 431; Clark C. Havighurst, *Doctors and Hospitals: An Antitrust Perspective on Traditional Relationships*, 33 DUKE L.J. 1071, 1075 (1984).

306. Abraham & Weiler, *supra* note 239, at 386 n.22 (citing *Bing v. Thunig*, 143 N.E.2d 3, 4–5 (N.Y. 1957)).

307. Rutchik, *supra* note 296, at 552; Bovbjerg & Berenson, *supra* note 295, at 222.

mere “workshops” of highly skilled physicians.³⁰⁸ For these reasons, courts have traditionally declined to hold hospitals liable for the acts of their salaried physicians.³⁰⁹

But gradually, the difficulty in determining whether the hospitals’ activities were administrative or medical in nature led courts to erode the line between administrative and medical services.³¹⁰ Part of the reason for abandoning the distinction was the difficulty in drawing lines between medical and administrative services in hard cases.³¹¹ These difficulties in many hard cases convinced courts to abandon the distinction as a guiding principle in deciding whether to impose liability on hospitals or their salaried physicians.³¹²

The breakdown of these barriers paved the way for imposing liability on hospitals for the negligence of physicians, including through the theories of agency and corporate liability discussed in Section II.A. The expanding legal responsibilities of hospitals caught the attention of legal theorists who proposed assigning liability at the enterprise level (e.g., requiring health systems, hospitals, or health plans to be liable for all patient injuries).³¹³ The rationale underlying this philosophy is that large institutions are in the best position to manage risk by preventing injuries.³¹⁴ Requiring health plans, hospitals, or health systems to bear liability for all injuries occurring in their facilities positions them to incorporate the price of accidents into the cost of operating the enterprise as a whole and spread liability expenses among those providing and receiving medical care.³¹⁵

308. Bovbjerg & Berenson, *supra* note 295, at 222; *see* Southwick, *supra* note 301, at 431, 434.

309. Southwick, *supra* note 301, at 440. In addition to liability exemptions for hospitals, insurers were exempt from responsibility for negligently caused medical injuries based on the rationale that they were “third parties” too far removed from clinical care to be considered liable. Bovbjerg & Berenson, *supra* note 295, at 222.

310. Abraham & Weiler, *supra* note 239, at 386.

311. Hardy, *supra* note 258, at 124.

312. *See id.* *See generally* GEORGE GLEASON BOGERT ET AL., THE LAW OF TRUSTS AND TRUSTEES § 402 (2020) (noting that “more than three quarters of the states had abolished or limited the doctrine of charitable immunity” by the end of the twentieth century).

313. Bovbjerg & Berenson, *supra* note 295, at 224; Jeffrey O’Connell, *Expanding No-Fault Beyond Auto Insurance: Some Proposals*, 59 VA. L. REV. 749, 827 (1973); Clark C. Havighurst & Laurence R. Tancredi, “Medical Adversity Insurance”—*A No-Fault Approach to Medical Malpractice and Quality Assurance*, 51 MILBANK MEM’L FUND Q. HEALTH & SOC’Y 125, 125–26 (1973); Robert E. Keeton, *Compensation for Medical Accidents*, 121 U. PA. L. REV. 590, 616–17 (1973).

314. Bovbjerg & Berenson, *supra* note 295, at 224.

315. *Id.*

The work of these scholars informed the 1991 Reporters' Study for the American Law Institute, which concluded that physician liability could not adequately respond to the increasingly institutional nature of medical care in the United States.³¹⁶ The 1991 report determined that imposing liability on those in the best position to identify and mitigate the risk of patient injury would incentivize them to create and comply with organization-wide safety standards to reduce the frequency of patient injuries and enhance the overall quality of care.³¹⁷ Imposing liability at the enterprise level would also enable physicians to stop worrying about liability and focus more squarely on providing the best treatment for patients.³¹⁸

One of the 1991's study's authors, Paul Weiler, was also part of a famous team of researchers that produced the Harvard Study of Medical Practice in New York. Like the 1991 report, the Harvard Study recommended implementing a hybrid system of enterprise liability and no-fault liability.³¹⁹ The 1991 report and the Harvard Study influenced the Clinton Administration's 1993 Task Force on National Health Reform, which included a version of enterprise liability.³²⁰

By the middle of the 1990s, several scholars and practitioners—including some of the most influential ones—had become convinced that implementing a form of institutional liability was needed to address modern medical accidents.³²¹ To be sure, proponents of enterprise liability did not agree on everything, including whether to impose liability at the level of the hospital, insurer, or managed care organization.³²² Nevertheless, they agreed that institutional liability would improve the quality, and reduce the cost, of medical care.³²³

Despite the appeal of enterprise liability, several interested parties strongly opposed the doctrine's foray into medicine.³²⁴ Of the concerns raised over

316. *Id.* at 227–28; 2 AM. L. INST., REPORTER'S STUDY ON ENTERPRISE RESPONSIBILITY FOR PERSONAL INJURY: APPROACHES TO LEGAL AND INSTITUTIONAL CHANGE 125–26 (1991).

317. Bovbjerg & Berenson, *supra* note 295, at 227–28; 2 AM. L. INST., *supra* note 316, at 123.

318. Philip G. Peters, Jr., *Resuscitating Hospital Enterprise Liability*, 73 MO. L. REV. 369, 374 (2008).

319. *Id.* at 373–74; 2 AM. L. INST., *supra* note 316, at 123–26; HARV. MED. PRAC. STUDY, PATIENTS, DOCTORS, AND LAWYERS: MEDICAL INJURY, MALPRACTICE LITIGATION, AND PATIENT COMPENSATION IN NEW YORK 11-9 (1990).

320. Sage, *supra* note 302, at 162–63; Abraham & Weiler, *supra* note 239, at 382–83; Bovbjerg & Berenson, *supra* note 295, at 228.

321. Peters, *supra* note 318, at 374–75.

322. *Id.* at 375; Sage, *supra* note 302, at 162–63.

323. Peters, *supra* note 318, at 375.

324. William M. Sage & James M. Jorling, *A World That Won't Stand Still: Enterprise Liability by Private Contract*, 43 DEPAUL L. REV. 1007, 1010 (1994); Abraham & Weiler, *supra* note 239, at 383; Bovbjerg & Berenson, *supra* note 295, at 230.

enterprise liability, perhaps the most notable ones were those raised by organized medicine. Physicians viewed enterprise liability as a threat to their autonomy.³²⁵ They claimed that enterprise liability would incentivize health plans to terminate high-risk physicians.³²⁶ Similarly, organized medicine anticipated that hospitals would be motivated to micromanage clinical care and restrict privileges to uncooperative or “underperforming” physicians.³²⁷ Additionally, if hospitals or other entities were empowered to settle cases without consulting physicians, the latter could suffer reputational damage.³²⁸ For these and other reasons, physicians expressed strong opposition to hospital enterprise liability in its various forms. At least one representative for physician groups went so far as to describe enterprise liability as an effort to take away what they viewed as their “constitutional right to be sued.”³²⁹ Although physicians are not the only party in the complex mosaic of parties responsible for caring for patients, they have enjoyed substantial clout since the turn of the twentieth century.³³⁰ Today, even though the power and solidarity of organized medicine is not as significant as it once was, the profession retains significant political influence.³³¹

Regardless of the merits of these and other concerns, they proved to be a significant blow to momentum for enterprise liability, which never gained the widespread support its advocates had hoped for.³³² Although the doctrine is not as popular as it once was, support for enterprise liability in its various forms remains.³³³ In fact, two of the most prominent contemporary health law scholars have recently proposed applying a version of enterprise liability to the use of certain types of artificial intelligence tools in medicine.³³⁴

Notwithstanding the enduring appeal of enterprise liability, I suggest that we have two good reasons to reject it. The first has to do with the inability of the doctrine to distribute liability among health systems, clinicians, and developers. The rationale for assigning liability solely to health systems, hospitals, or health plans is that large institutions are in the best position to

325. Sage, *supra* note 302, at 170; Sage & Jorling, *supra* note 324, at 1010; Bovbjerg & Berenson, *supra* note 295, at 230.

326. Sage, *supra* note 302, at 170.

327. Bovbjerg & Berenson, *supra* note 295, at 227–28.

328. *Id.* at 230.

329. Sage, *supra* note 302, at 170 n.46; Bovbjerg & Berenson, *supra* note 295, at 230.

330. Zettler, *supra* note 175, at 441.

331. *Id.*

332. Peters, *supra* note 318, at 376; Sage & Jorling, *supra* note 324, at 1008; Abraham & Weiler, *supra* note 239, at 383; Bovbjerg & Berenson, *supra* note 295, at 230.

333. See, e.g., Gregory C. Keating, *The Theory of Enterprise Liability and Common Law Strict Liability*, 54 VAND. L. REV. 1285 (2001).

334. Price & Cohen, *supra* note 23, at 342.

predict the risk of accidents, take steps to mitigate the likelihood of harm, and spread the cost of injuries.³³⁵ For example, relative to physicians and developers, health systems are likely to have more information about the institutions and processes in which adaptive black-box systems will be deployed, details regarding patient populations, the credentials of the clinicians using such tools, and more.³³⁶ But as illustrated in Parts I and II, health systems have an insufficiency of information problem.³³⁷ Health systems lack key insights, such as detail regarding the datasets on which systems are trained or parameters used to validate adaptive black-box tools.³³⁸ Even if health systems had access to this sort of information, they likely lack the requisite capacity and expertise to analyze it.³³⁹ Indeed, the work of training deep neural networks is sometimes described as an “art” rather than a “science” because it involves intuition acquired through repeated trial and error.³⁴⁰ Of course, health systems are not on an island in this regard. Physicians lack the time and expertise needed to identify and address all problems in artificial intelligence systems, and the developers lack knowledge of significant institutional and clinical details that will affect how adaptive black-box tools will operate in real life.³⁴¹ But for the same reasons illustrated in Section II.C, it does not seem appropriate to make health systems solely liable for patient harms resulting from the use of tools like DLeye.³⁴²

To be sure, health systems could seek indemnity or contribution from clinicians and developers if using DLeye results in patient harm.³⁴³ But the opacity of tools like DLeye complicates indemnification efforts.³⁴⁴ Because DLeye is a black box, it is very challenging to determine whether harms caused by DLeye are due to health systems failing to implement appropriate protocols and credentialing standards, inadequate training of the system by developers, negligent clinician operation of the system, or some combination thereof.³⁴⁵ Courts are inclined to limit liability when it can’t be established that a party’s conduct contributed to the cause of an accident.³⁴⁶ Hence,

335. Bovbjerg and Berenson, *supra* note 295, at 224.

336. Price & Cohen, *supra* note 23, at 341–42; *see also* Price, *supra* note 16, at 462.

337. Price & Cohen, *supra* note 23, at 341.

338. *Id.* at 341–42.

339. *See* Choi, *supra* note 84, at 310.

340. *Id.* at 314.

341. Price & Cohen, *supra* note 23, at 341.

342. *See* Vladeck, *supra* note 30, at 148.

343. *See id.* at 128.

344. *See id.*

345. *See* Terry, *supra* note 263, at 162.

346. *See, e.g.,* Vladeck, *supra* note 30, at 128 n.38.

without the ability to identify the cause of a patient's injury, health systems won't have a basis upon which to leverage the law in their attempt to seek indemnification.³⁴⁷

The second reason for rejecting traditional forms of enterprise liability has to do with the enduring concerns that have long surrounded proposals to apply the doctrine in the medical context. Holding health systems, hospitals, or health plans strictly liable for patient injuries would relieve clinicians and developers of liability concerns. But doing so would implicate the autonomy-based objections mentioned just a few paragraphs ago.³⁴⁸ For example, developers and clinicians might worry about the reputational damage that could result if hospitals are inclined to settle rather than fight claims.³⁴⁹ What's more, the enterprise to which liability is assigned might seek to reduce injuries through micromanagement, or it could pass the cost of harms onto developers and clinicians in the form of reduced compensation.³⁵⁰

Of course, clinicians and developers could attempt to exercise their autonomy by influencing the process of negotiating the contract for using a tool like DLeye. Value analysis committees have been created at many hospitals and health systems to control costs and maximize quality when contracting for the services of medical tools like DLeye.³⁵¹ As members of such committees deliberate over the contractual terms for acquiring the services of adaptive and opaque artificial intelligence tools, clinicians and developers could attempt to influence the drafting of liability and indemnification provisions to better reflect the fair value created by tools like DLeye for all parties.³⁵² In the unfortunate circumstances when patients are injured, these provisions could fairly distribute the cost of these harms among health systems, developers, and physicians, and the broader community of patients benefitting from the use of these systems.³⁵³

347. *See id.* at 128.

348. *See supra* notes 324–30 and accompanying text.

349. *See, e.g.,* Bovbjerg & Berenson, *supra* note 295, at 230.

350. *See id.* at 227–28, 230.

351. *See* Kalainov, *supra* note 122, at 769; Grundy, *supra* note 121, at 82.

352. *See* Bovbjerg & Berenson, *supra* note 295, at 224; *see also* Sage, *supra* note 302, at 166–69; Price & Cohen, *supra* note 23, at 342.

353. *See* Bovbjerg & Berenson, *supra* note 295, at 224–26; *see also, e.g.,* Sage, *supra* note 302, at 166–69. Using tools like DLeye could help many patients receive more timely access to diagnosis treatment than they otherwise receive if such systems were not in operation. Unfortunately, there will likely be rare instances in which these tools fail to operate as intended and, as a result, contribute to patient harm. In such cases, the broader patient population served by this tool, developers, clinicians, and health systems could coordinate to make compensation available to injured patients. This compensation scheme could be funded through a combination of a surcharge paid by patients, a portion of the cost savings, and increased revenue from the use of tools like DLeye by developers, clinicians, and health systems.

But given the lack of equal bargaining power among health systems, clinicians, and developers, this approach to alleviating autonomy concerns would likely be an empty gesture.³⁵⁴ For decades, the delivery of healthcare has become increasingly bound up with the marketplace.³⁵⁵ In recent years, that trend has continued playing out with private equity firms purchasing more and more established healthcare practices.³⁵⁶ Once an acquisition has been completed, private equity firms seek to achieve cost savings through a variety of methods, including by cutting staffing and wages, attempting to capture referral services, resisting unionization, and more.³⁵⁷ As private equity investors continue gaining control of more practices, one of the many consequences is reduced autonomy for clinicians.³⁵⁸ Developers might be able to exercise more influence if they have a medical tool that is in high demand. But as already large institutional investors continue acquiring established healthcare practices, the bargaining power of smaller developers and clinician groups will struggle to influence decisions on a variety of topics, including liability distribution.³⁵⁹ Of course, if use of adaptive black-box technology becomes sufficiently widespread in the clinical context, standard indemnification clauses and liability-spreading provisions could develop to ensure an appropriate distribution of risk sharing across all relevant actors.³⁶⁰ But the pervasiveness of existing power asymmetries suggests that the development of such standard provisions is far from guaranteed.³⁶¹ Hence, the lack of influence available to developers and clinicians implicates the

354. See Maliha et al., *supra* note 10, at 637.

355. See Abraham & Weiler, *supra* note 239, at 398.

356. See, e.g., Field et al., *Private Equity in Health Care: Barbarians at the Gate?*, 15 DREXEL L. REV. 821 824–25 (2023); MARYANN P. FELDMAN & MARTIN F. KENNEY, PRIVATE EQUITY AND THE DEMISE OF THE LOCAL 36–39 (Arie Y. Lewin & Till Talaulicar eds., 2024).

357. See, e.g., Field et al., *supra* note 356, at 832–33; Suhas Gondi & Zirui Song, *Potential Implications of Private Equity Investments in Health Care Delivery*, 321 JAMA 1047 (2019); FELDMAN & KENNEY, *supra* note 356, at 37–38.

358. See Lola Butcher, *The Future of Private Equity in HealthCare*, 7 PHYSICIAN LEADERSHIP J. 57, 58 (2020); see also Erin C. Fuse Brown & Mark A. Hall, *Private Equity and the Corporatization of Health Care*, 76 STAN. L. REV. 527 (2024).

359. See, e.g., Abraham & Weiler, *supra* note 239, at 398.

360. See Maliha et al., *supra* note 10, at 637.

361. See, e.g., Claire E. O’Hanlon, *Impacts of Health Care Industry Consolidation in Pittsburgh, Pennsylvania: A Qualitative Study*, INQUIRY, Jan.–Dec. 2020, <https://journals.sagepub.com/doi/full/10.1177/0046958020976246> [<https://perma.cc/AC6P-XQQZ>]; Stephanie M. Topp et al., *Power Analysis in Health Policy and Systems Research: A Guide to Research Conceptualization*, 6 BMJ GLOB. HEALTH, e007268, at 1 (Nov. 5, 2021), <https://gh.bmj.com/content/6/11/e007268> [<https://perma.cc/4BBC-NQPN>]; Sage, *supra* note 302, at 169–73.

autonomy concerns that have long accompanied traditional versions of the doctrine.³⁶²

B. Common Enterprise Liability

These considerations provide reason to consider an alternative to enterprise liability in its traditional form. The alternative I have in mind is a close cousin of enterprise liability known as common enterprise liability.³⁶³ Common enterprise liability imposes joint and several liability on entities coordinating their shared pursuit of a common aim.³⁶⁴ To better understand the doctrine and how it could be applied to adaptive black-box tools like DEye, consider the leading case of *FTC v. Tax Club, Inc.*³⁶⁵

In *Tax Club*, the U.S. District Court for the Southern District of New York heard a case involving multiple defendants.³⁶⁶ Over the course of several years, defendants coordinated their efforts to operate an interrelated set of companies, which contacted prospective customers to sell products and services for the purpose of creating small businesses.³⁶⁷ The operation of defendants' companies resulted in several customer allegations of deceptive marketing practices and, ultimately, the complaint before the court.³⁶⁸ Included among the arguments defendants provided for dismissing the complaint was the rule against group pleadings.³⁶⁹

Generally, courts take a dim view of grouping defendants together without specifying the misconduct of each.³⁷⁰ However, the court in *Tax Club* noted an exception to this rule: when defendants are engaged in a "common enterprise."³⁷¹ Pursuant to the theory of common enterprise liability, each defendant in a group can be held jointly and severally liable for the actions of other defendants in that group.³⁷²

To determine whether a common enterprise exists among a group of defendants, courts consider whether they maintain a common set of

362. See Sage, *supra* note 302, at 170 n.46; Bovbjerg & Berenson, *supra* note 295, at 230.

363. See Vladeck, *supra* note 30, at 129 n.39.

364. *Id.*

365. 994 F. Supp. 2d 461 (S.D.N.Y. 2014).

366. *Id.* at 465.

367. *Id.* at 466.

368. *Id.* at 467.

369. *Id.* at 469.

370. See *id.*

371. *Id.* The Federal Trade Commission frequently invokes the doctrine of "common enterprise" to impose joint and several liability on companies that worked together to commit fraud. *Id.*; see also *FTC v. Network Servs. Depot, Inc.*, 617 F.3d 1127, 1142–43 (9th Cir. 2010).

372. 994 F. Supp. 2d at 469.

employees and officers, share common control over the enterprise in question, maintain common offices, share workspaces, conduct the enterprise through a web of interrelated companies, comingle funds, conduct joint advertising, and lack features distinguishing themselves from other parties to the enterprise.³⁷³ It is immaterial, the court noted, whether each individual defendant engaged in the wrongdoing in question.³⁷⁴ “The very nature of this theory is that corporate entities that are a part of the common enterprise are liable for the conduct of other entities in the enterprise, regardless of whether the particular entity engaged in the behavior at issue.”³⁷⁵ In other words, the theory of common enterprise liability is concerned with wrongdoing by the enterprise as a whole.³⁷⁶

Extending this theory to autonomous vehicles driving on highways, David Vladeck proposed holding manufacturers and designers of autonomous vehicles liable for injuries caused by such cars.³⁷⁷ The rationale for doing so is that these parties are jointly involved in the common aim of designing and manufacturing self-driving cars.³⁷⁸ When these vehicles result in harm, courts face doctrinal uncertainty.³⁷⁹ For example, if an autonomous car fails to stop for a pedestrian in a walkway, it will be challenging to determine whether this failure was the result of a manufacturing or design defect for reasons similar to those discussed in Section II.B. These challenges make it difficult to trace the failure resulting in an injury “to, or reasonably imputed to, the activity of an identifiable person or legal entity.”³⁸⁰

In response to such situations, Vladeck proposes that a version of common enterprise liability can be applied so that courts need not wrestle with thorny questions concerning fault or causation.³⁸¹ When it is impossible to assign fault to any individual party, he suggests joint and several liability for the multiple entities (e.g., a vehicle’s designer, programmer, and manufacturer) collectively engaged in the enterprise responsible for the existence of autonomous cars on roads and highways.³⁸² According to Vladeck’s version of common enterprise liability, we not need concern ourselves with whether

373. *See id.* at 470.

374. *Id.*

375. *Id.*

376. *Id.*

377. Vladeck, *supra* note 30, at 128–29.

378. *Id.*

379. *See id.* at 130–41 (showing through case illustration that over the last fifty years courts are still uncertain of the correct doctrine to apply in these cases).

380. *Id.* at 141.

381. *Id.* at 129 n.39.

382. *Id.*; *see also id.* at 149.

there is a comingling of funds, joint advertisement, or other factors traditionally indicating the existence of a common enterprise.³⁸³ Instead, his version of the doctrine would only require that the parties work together in pursuit of a common aim.³⁸⁴

C. Common Enterprise Liability for Adaptive & Opaque Systems

Although the common aim Vladeck had in mind was designing and manufacturing autonomous vehicles, his proposal can be extended to the joint enterprise of developing and using adaptive black-box tools in the clinical context. Specifically, pursuant to common enterprise liability, responsibility for a patient injury resulting from the use of an adaptive black-box tool would be apportioned among all those participating in the “common objective” of using that technology in the clinical context.³⁸⁵ The discussion in Section II.C illustrates the range of regulatory, technical, and business considerations connecting a complex web of individuals and entities—from health systems and hospitals to clinicians and software vendors—participating in the joint aim of bringing tools like DEye to bear when caring for patients. For this reason, it is unreasonable to ask any one of these individuals or entities to bear the sole burden of liability for injuries.³⁸⁶ Instead, it seems appropriate to impose common enterprise liability on all those participating in the joint aim of developing and using adaptive black-box tools in the clinical context.³⁸⁷

But because adaptive and opaque medical tools are a little different than self-driving cars, a few modifications to Vladeck’s version of common enterprise liability are in order. One adjustment worth making is in response to the fact that common enterprise liability implicates the same sort of autonomy concerns that accompany enterprise liability. That is, the power asymmetries inhibiting the fair development of liability and indemnity distribution pursuant to traditional forms of enterprise liability remain under a common enterprise liability regime. Let’s assume that the use of DEye results in patient harm, and in response, a court holds all parties to a suit jointly and severally liable. If an injured patient sues a health system, that

383. *See id.* at 149 & n.94.

384. *Id.*

385. *Id.* at 128–29.

386. *See* Maliha et al., *supra* note 10, at 630 (“[P]hysicians exist as part of an ecosystem that also includes health systems and AI/ML device manufacturers. Physician liability over use of AI/ML is inextricably linked to the liability of these other actors.”).

387. Vladeck, *supra* note 30, at 129 n.39.

party is free to then seek contribution from the other members of the common enterprise. But if, for instance, health systems have much more power relative to other interested parties, they could foist the bulk of the liability burden on developers, clinicians, or even patient populations in the form of higher prices for using tools like DLeve, micromanagement, or reduced compensation.³⁸⁸ Given this state of affairs, one could wonder whether there is a good reason to prefer common enterprise liability to more conventional versions of the doctrine.

In response, I suggest that Vladeck's version of common enterprise liability be modified such that when courts impose joint and several liability, they also require a proportional distribution of liability among all relevant parties. A framework outlining respective parties' liability burdens could be proposed by members of the common enterprise themselves. Such a proposal could result from their engaging in a sort of Coasian bargaining to distribute the cost of potential injuries among physicians, developers, and other relevant actors through, perhaps, indemnification clauses in the contracts drafted pursuant to a value analysis committee's procurement process.³⁸⁹ Although the Coase theorem is controversial and lacks a singular definition, in effect, it states that "where a market does not exist to internalize costs, the assignment of liability will have the effect of creating a set of market or market-like interconnections between classes of agents that will cause relevant costs to be reflected in the prices faced by all agents."³⁹⁰ The theorem presumes the absence of factors that complicate the time and effort relevant parties need to make decisions about allocating liability.³⁹¹ As a result of this frictionless world, the parties will bargain for an efficient resolution, regardless of which party is initially assigned with liability.³⁹²

But the real world is not frictionless.³⁹³ For example, developers, clinicians, health systems, and other relevant parties might evaluate risk differently for reputational or other reasons, and these differences could

388. See Bovbjerg & Berenson, *supra* note 295, at 227–28.

389. See, e.g., Kyle D. Logue & Joel Slemrod, *Of Coase, Calabresi, and Optimal Tax Liability*, 63 TAX L. REV. 797, 804–09 (2010); Steven G. Medema, *Debating Law's Irrelevance: Legal Scholarship and the Coase Theorem in the 1960s*, 2 TEX. A&M L. REV. 159, 163–67 (2014); Price & Cohen, *supra* note 23, at 366–67; Maliha et al., *supra* note 10, at 637.

390. Steven G. Medema, *Juris Prudence: Calabresi's Uneasy Relationship with the Coase Theorem*, 77 LAW & CONTEMP. PROBS. 65, 75 (2014) (citing Guido Calabresi, *The Decision for Accidents: An Approach to Nonfault Allocation of Costs*, 78 HARV. L. REV. 713, 729–30 (1965)).

391. See *id.* at 75–76.

392. *Id.*

393. See Guido Calabresi, *The Pointlessness of Pareto: Carrying Coase Further*, 100 YALE L.J. 1211, 1222–23 (1991).

translate into asymmetrical abilities to insure against risk.³⁹⁴ What's more, the power disparities mentioned earlier could result in disproportionate liability burdens being foisted on developers, clinicians, or other parties with less sway when contracting for the use of a tool like DLEYe.³⁹⁵

Given these considerations, we have reason to prefer an alternative to expecting the relevant parties to arrive at a proportional allocation of liability on their own. One option is for courts to consult a panel of experts comprised of representatives from interest groups. This panel might be comprised of experienced professionals representing the interests of patients, clinicians, health administrators, technology industry experts, ethicists, and lawyers to review the matter in question and determine how to apportion responsibility.³⁹⁶ Of course, doing so would inject an additional layer of uncertainty and cost into such proceedings. Additionally, asking a panel to apportion responsibility would be an incredibly difficult task for the same reasons that make it so hard to determine fault in the first place.

But these costs might only be temporary obstacles. Standard indemnification clauses and liability-spreading provisions could develop over time as the use of adaptive black-box technology becomes sufficiently widespread and judges become more comfortable making determinations about what constitutes a fair distribution of liability in these cases.³⁹⁷ What's more, the possibility of a judge rejecting an inequitable distribution could incentivize a fair allocation of liability between relevant parties when negotiating for the use of tools like DLEYe in the clinical workflow. Even if this decentralized scheme for an equitable distribution of accountability is adopted not in the near term, it need not prevent the eventual realization of a more standard, equitable, and widely applicable distribution of responsibility in the future.

A second modification worth making to Vladeck's version of common enterprise liability addresses line-drawing questions. Concerns over the boundaries of enterprise liability, in any form, predate the doctrine itself, which arose from the cannon of strict liability.³⁹⁸ One of the most influential and controversial strict liability cases, *Rylands v. Fletcher*, was decided in

394. See, e.g., Guido Calabresi, *Some Thoughts on Risk Distribution and the Law of Torts*, 70 YALE L.J. 499, 532 (1961).

395. See Bovbjerg & Berenson, *supra* note 295, at 227–28.

396. See Jessica S. Allain, *From Jeopardy! to Jaundice: The Medical Liability Implications of Dr. Watson and Other Artificial Intelligence Systems*, 73 LA. L. REV. 1049, 1075–77 (2013).

397. Maliha et al., *supra* note 10, at 639.

398. Keating, *supra* note 333, at 1285–87.

England during the 1860s.³⁹⁹ In the United States, although *Rylands* was not immediately and unanimously endorsed by courts, the precedent for strict liability established in that case eventually won support in a majority of jurisdictions.⁴⁰⁰ Throughout the doctrine's history in the United States, there has been considerable debate over whether and how to limit strict liability to certain kinds of activities, such as those that are "ultrahazardous" or "unusual."⁴⁰¹ The twists and turns of this debate over the years have largely been shaped by evolving attitudes among the public and scholars over laissez-faire capitalism.⁴⁰² Proponents of expanding strict liability claim that it is fair to require those who are the principal beneficiaries of advancements in technology and industry to bear the burden of the harm resulting from their activities.⁴⁰³ On the other hand, critics of the doctrine have expressed concern that imposing liability too broadly hinders economic progress in a way that harms both injurers and victims.⁴⁰⁴ In the case of the DLeve hypothetical, in

399. See WILLIAM LLOYD PROSSER, *The Principle of Rylands v. Fletcher*, in *SELECTED TOPICS ON THE LAW OF TORTS* 135, 135 (1953). *Rylands* concerned a reservoir used for supplying a textile mill with waterpower. *Fletcher v. Rylands* (1865) 159 Eng. Rep. 737; 3 H. & C. 774, *rev'd*, (1866) 1 LR Exch. 265, *aff'd*, (1868) 3 LRE & I App. 330 (HL). The reservoir burst and flowed into the shaft of an abandoned coal mine and, from there, flooded into the underground mine shafts of a neighboring property. *See id.* at 737–45. The cause of the burst reservoir was likely due to the negligence of independent contractors responsible for constructing the reservoir. *Id.* However, the doctrine of privity prevented the owners of the coal mine from bringing a cause of action against the contractors. *Id.* Additionally, because the mill owners were not aware of the mine shafts beneath the reservoir, they could not be held liable for negligence. *Id.* at 737–46. The facts also did not support a trespass claim because the flood damage was indirect and consequential rather than immediate and direct; nor did the facts constitute nuisance without a recurring event or one that was "hurtful or injurious to the senses." *Id.* at 737–45. Because neither the mill owners nor the independent contractors could be held liable, and the case did not fit into one of tort's existing pigeonholes, the English courts eventually held that liability could be applied to activities that did not constitute "natural" use of the land without proof of negligence. *See Rylands*, 3 LRE & I App. at 338–39. Subsequent cases in England and elsewhere wrestled with what constitutes the natural versus unnatural use of property and how liability should be apportioned based on such distinctions. *See* W.T.S. Stallybrass, *Dangerous Things and the Non-Natural User of Land*, 3 CAMBRIDGE L.J. 376 (1929) (discussing various courts and commentators interpreting the meaning of dangerous and non-natural use of property). For example, courts and commentators debated whether the rule in *Rylands* was limited to "extraordinary," "abnormal," or especially "dangerous" activities. *Id.* at 377, 390, 392.

400. Gregory C. Keating, *Recovering Rylands: An Essay for Robert Rabin*, 61 DEPAUL L. REV. 543, 582–84 (2012); *see* Gerald W. Boston, *Strict Liability for Abnormally Dangerous Activity: The Negligence Barrier*, 36 SAN DIEGO L. REV. 597 (1999) (tracking and discussing the evolution of the *Rylands* decision).

401. *See* Keating, *supra* note 400, at 565–84; Boston, *supra* note 400, at 605–28.

402. *See* Keating, *supra* note 400, at 558–82.

403. *See id.* at 576–79.

404. *See id.*

particular, some might worry that imposing liability too broadly could stifle innovation.

This discourse concerning the limits of strict liability carries over into the debates over enterprise liability.⁴⁰⁵ Theoretically, it is possible to assign liability in a way that makes entities responsible for the costs of accidents resulting from their activities.⁴⁰⁶ However, in practice, it can be difficult to determine how much control particular individuals and entities have over the accidents resulting from their activities.⁴⁰⁷ These difficulties have long been a source of concern for the doctrine's critics, who claim that without a threshold like fault to limit responsibility for accidents, enterprise liability is unworkable.⁴⁰⁸

Regardless of the merits of these objections, adopting a slightly modified version of common enterprise liability can blunt the force of these line-drawing concerns. The form of common enterprise liability developed by Vladeck imposes liability on all entities involved in a joint enterprise that results in harm when fault cannot be determined.⁴⁰⁹ Unlike Vladeck, I suggest that the imposition of liability for patient injuries caused by adaptive black-box systems be conditional on the enterprise as a whole failing to satisfy a relevant standard of care.

To better understand the line-drawing benefits of my proposal, consider a stylized example. Let's assume that based on the collective efforts of clinicians, health systems, and developers, DLEye correctly identifies disease 90% of the time. This is an improvement over the performance of reasonably competent human clinicians, who correctly identify disease 70% of the time. Because of this 20% delta, it becomes the standard of care for clinicians to

405. See Keating, *supra* note 333, at 1287–88; Gregory C. Keating, *The Heroic Enterprise of the Asbestos Cases*, 37 SW. U. L. REV. 623, 627 (2008). The concept of enterprise liability first originated outside of tort law in legislative schemes like workers compensation. See Keating, *supra* note 333, at 1287–88; Keating, *supra* note 405, at 627. But the rationale underlying enterprise liability eventually made its mark on tort law by, for example, shaping doctrinal developments in vicarious liability and influencing the rise of products liability. See Keating, *supra* note 333, at 1287–88; Keating, *supra* note 405, at 627. In fact, by the middle of the twentieth century, enterprise liability had gained almost “complete support within the academic community.” George L. Priest, *The Invention of Enterprise Liability: A Critical History of the Intellectual Foundations of Modern Tort Law*, 14 J. LEGAL STUD. 461, 463 (1985). At that time, it seemed possible—perhaps even likely—that some version of enterprise liability would replace the fault principle as the foundation of tort theory in the coming years. *Id.* at 463–64.

406. James A. Henderson, Jr., *The Boundary Problems of Enterprise Liability*, 41 MD. L. REV. 659, 662–63 (1982).

407. *Id.*

408. See, e.g., James A. Henderson, Jr. & Aaron D. Twerski, *The Unworkability of Court-Made Enterprise Liability: A Reply to Geistfeld*, 67 N.Y.U. L. REV. 1174, 1174–75 (1992).

409. Vladeck, *supra* note 30, at 129 n.39.

rely on DLeye to screen patients in the first instance. Non-physician clinicians take pictures of patients' eyes, and DLeye analyzes these images to determine whether disease is present. If no disease is detected, patients are not referred for a follow-up with a human physician.

Over time, developers, clinicians, and the health system in which DLeye is deployed collaborate to use the tool for the benefit of patients. As they do so, the tool undergoes four successful updates based on the efforts of all the relevant parties.⁴¹⁰ Developers monitor the system, train clinicians, and update the tool based on new data.

However, after the fifth update is made to DLeye, it fails to identify disease in an image that would have been recognized by previous versions (i.e., versions one through four) of the system. Notably, although previous versions of DLeye would have recognized disease in this particular image, a reasonably competent human clinician would not have done so. In other words, the fifth version of DLeye performs better than a reasonably competent human physician, but it performs less well than previous versions of itself.

This example illustrates what we might regard as the standard of care for adaptive black-box tools like DLeye. Contra Vladeck's version of common enterprise liability, I do not suggest imposing liability on all those involved in the joint enterprise of bringing DLeye to bear in the care of patients liable for *any* injury caused by its outputs, regardless of fault. Rather, I propose imposing liability only in instances where the tool fails to live up to its own

410. Here, it is worth asking: Why wouldn't developers stop updating the tool once it becomes 90% effective? This is a good question, so let's assume that DLeye arrives at its 90% effectiveness rate after continuing to improve from its baseline, which was only 80% effective. If the developers stopped improving the system once it achieves 90% effectiveness, there would likely be competitors that could take advantage of developing a system that out-performed DLeye. Indeed, the market for artificial intelligence tools is expected to expand at a compound annual growth rate of over forty percent in the coming decade, so there will be financial incentives for companies to explore the space. *See, e.g.*, MKTS. & MKTS., ARTIFICIAL INTELLIGENCE (AI) IN HEALTHCARE MARKET SIZE, SHARE & GROWTH (2024), <https://www.marketsandmarkets.com/Market-Reports/artificial-intelligence-healthcare-market-54679303.html> [<https://perma.cc/9PCQ-VBFZ>]. Let's assume, though, that there isn't a market for a tool that performs better than 90%; maybe clinicians and health systems wouldn't want to go through the hassle of acquiring a tool that offers only a marginal benefit over an already quite successful (i.e., 90% effective) system. In that case, even if the DLeye tool stopped updating once it arrived at a 90% effectiveness rate, we would still want it to continue improving from its 80% baseline until it arrived at that rate. So, we would still need a more forward-looking regulatory model from the FDA. We would also still need common enterprise liability to address patient harms that might result from DLeye encountering real world data. This is because the possibility remains that DLeye could fail to identify disease because it encounters some image in the real world that is sufficiently dissimilar to what it encountered in the training data. In that case, the system might fail to identify disease and, as a result, contribute to a patient harm.

standards, which might be established based on an aggregate of average performance.⁴¹¹

If common enterprise liability only imposes liability when the outputs from systems like DLeye fall below their own standard, we have reason to prefer this version of the doctrine to the one proposed by Vladeck. Because Vladeck's proposal would impose liability regardless of fault, it would trigger compensation in more instances than the form of the doctrine I propose. In the specific instances where adaptive black-box tools produce outputs that fall below the standard of performance that such systems usually observe, this version of common enterprise liability might assign responsibility to those involved in the joint activity of bringing that system to bear when caring for patients. In this way, the version of common enterprise liability I defend is less exposed to concerns over the inability to reliably limit liability without a threshold like fault.⁴¹² This sort of predictability with regard to liability risk could provide comfort to developers interested in exploring the market for

411. I thank Yonathan Arbel for pressing me to consider the following question: What if the newest version of DLeye fails to identify disease in a particular image that a previous version of the system would have caught, but the average performance of the newest version of DLeye performs as well or better than previous versions, including the one that would have identified disease in the image that the newest version failed to detect? Would liability attach in such a case? I lean toward thinking that liability should apply here. But my argument supporting that claim implicates deontological principles, which warrant more space than is available in this Article. Additionally, for these same deontological reasons, in an ideal world, I am inclined to prefer applying Vladeck's—stricter—version of common enterprise liability to all harms caused by DLeye. But if the aggregate-performance standard limitation on liability that I've proposed in this Article succeeds in opening the door to adopting *some* version of common enterprise liability to injuries caused by adaptive black-box tools, perhaps we could gradually move closer to applying Vladeck's version in the future.

412. *See, e.g.,* Henderson, Jr. & Twerski, *supra* note 408, at 1174–75. The reasons I've provided for preferring my version of common enterprise liability to Vladeck's also supply reasons to doubt the viability of using a no-fault system, perhaps modeled on workers compensation or vaccine funds, to address the use of adaptive black-box tools. *See, e.g.,* Price & Cohen, *supra* note 23, at 365. Pursuant to a no-fault liability system, patients injured by the conduct of clinicians would receive compensation even if the latter did not act negligently. David M. Studdert & Troyen A. Brennan, *Toward a Workable Model of "No-Fault" Compensation for Medical Injury in the United States*, 27 AM. J.L. & MED. 225, 227–29 (2001); Paul C. Weiler, *The Case for No-Fault Medical Liability*, 52 MD. L. REV. 908, 909–11 (1993). However, a no-fault system would raise the sort of line-drawing concerns that my version of common enterprise liability aims to weaken.

Additionally, despite these expected benefits of no-fault systems, it has become apparent that they have not operated as expected. Perhaps the main appeal of no-fault systems was their ability to drive down costs, but cost reductions have not materialized. *See, e.g.,* JAMES M. ANDERSON ET AL., THE U.S. EXPERIENCE WITH NO-FAULT AUTOMOBILE INSURANCE: A RETROSPECTIVE, at xiii–xiv (2010), <http://www.rand.org/pubs/monographs/MG860> [<https://perma.cc/2XZC-URX3>].

artificial intelligence tools, which is expected to expand at a compound annual growth rate of over forty percent in the coming decade.⁴¹³ Hence, there seems to be sufficient financial incentives for companies to explore the space even if they are subject to the version of common enterprise liability I recommend.

Of course, the aggregate-performance standard I've offered might not satisfy critics of enterprise liability in its various forms. But it is a response that positions us to continue a dialogue regarding how we might come to terms with the promise of adaptive black-box systems while implementing a liability system that provides a sufficient degree of safety and effectiveness.

IV. CONCLUSION

In this Article, I have claimed that a coming wave of adaptive and opaque artificial intelligence has many promising implications while simultaneously raising very thorny questions. I provide a two-part framework for coming to terms with the challenges posed by these systems and, in so doing, unleashing their promising potential.

The first part of my framework responds to the fact that traditionally, the FDA requires medical tools—including artificial intelligence systems—subject to regulation to pass through a gauntlet of rigorous tests and trials before they are approved to enter the clinical workflow. In this way, the agency's existing regulatory scheme is fundamentally backward-looking, which makes it ill-suited to regulate artificial intelligence systems that continually improve. In response, I recommend that the FDA adopt a forward-looking regulatory approach to adaptive artificial intelligence systems. Specifically, I propose that the agency make approval of such systems contingent upon receipt of a detailed list of processes and procedures interested parties will follow to ensure that the tool remains safe and effective as it continues evolves over time.

The second part of my framework recognizes that the forward-looking regulatory approach I recommend does not guarantee protection from civil liability. This is so because the specter of liability is a challenge for both plaintiffs and defendants: traditional tort doctrines like malpractice or products liability are ill-suited to address dynamically inscrutable artificial intelligence systems. I argue here in favor of a particular form of enterprise liability known as common enterprise liability. This doctrine can supplement a forward-looking regulatory approach for adaptive and opaque artificial

413. See, e.g., MKTS. & MKTS., *supra* note 410.

intelligence tools by both filling doctrinal gaps and addressing any injuries that aren't prevented through ex ante regulations.

Due to current regulatory and safety concerns, I acknowledge that the potential applications for truly adaptive artificial intelligence technologies are relatively narrow in the near term. Given the current political climate and makeup of the courts, it is unlikely that my two-part framework will soon be adopted to create a more hospitable environment for the development of adaptive and opaque artificial intelligence tools in medicine.

Notwithstanding these present limitations, this Article peeks around the corner with an eye toward the fast-approaching day when adaptive and opaque artificial intelligence tools will be used more widely in a variety of clinical contexts. The framework laid out in this Article can inform incremental progress toward a more forward-looking regulatory scheme if the political and judicial climate becomes more hospitable to doing so.

What's more, this Article recognizes that the safe and effective use of adaptive and opaque artificial intelligence tools in medicine requires a blend of forward- and backward-looking reforms to doctrines in both public and private law. These respective domains are siblings in a family of legal institutions capable of responding to thorny questions posed by new technologies being used in the medical context and beyond. Sometimes, the difficulties raised by new technologies are discreet enough to be adequately addressed through targeted reforms to doctrines in either public or private law. Other times, though, the challenges these systems pose are systemic and pervasive enough to warrant a more comprehensive set of reforms to multiple legal institutions, including those that are sometimes thought of as wholly separate and autonomous domains. I suggest that the latter approach is appropriate given the uniquely thorny questions raised by using adaptive and opaque artificial intelligence systems in the care of patients.