

Product Liability Suits for FDA-Regulated AI/ML Software

Barbara J. Evans and Frank Pasquale

The 21st Century Cures Act confirmed the FDA's authority to regulate certain categories of software that, increasingly, incorporate artificial intelligence/machine-learning (AI/ML) techniques. The agency's September 27, 2019 draft guidance on Clinical Decision Support Software proposed an approach for regulating CDS software and sheds light on plans for regulating genomic bioinformatics software (whether or not it constitutes CDS software). No matter how the FDA's regulatory approach ultimately evolves, the agency's involvement in this sphere has an important – and underexamined – implication: FDA-regulated software seemingly has the status of a medical product (as opposed to an informational service), which opens the door to product liability for defects causing patient injury. When a diagnostic or treatment decision relies on FDA-regulated CDS software, will mistakes invite strict liability, as opposed to being judged by the professional or general negligence standards of care that traditionally governed diagnostic and therapeutic errors? This chapter explores the policy rationales for product liability suits and asks whether such suits may have a helpful role to play as an adjunct to FDA oversight in promoting safety, effectiveness, and transparency of CDS software as it moves into wider use in clinical health care settings.

2.1 INTRODUCTION

The term “clinical decision support” (CDS) software includes various tools for enhancing clinical decision making and patient care. Examples include systems that provide alerts and reminders to health care providers and patients, or algorithms that offer recommendations about the best diagnosis or treatment for a patient.¹ The US Food and Drug Administration (FDA) conceives CDS software as data processing systems that combine patient-specific information (such as a patient's test results or clinical history) with generally applicable medical knowledge (such as clinical practice guidelines, information from drug labeling, or insights gleaned from outcomes observed in

¹ See, e.g., Clinical Decision Support, HealthIT.gov (Apr. 10, 2018), www.healthit.gov/policy-researchers-implementers/clinical-decision-support-cds [<https://perma.cc/JWV8-YUGQ>].

other patients) to provide a health care professional with patient-specific recommendations about how to diagnose, treat, or prevent disease in clinical health care settings.²

Congress defines an FDA-regulable medical device as an “instrument, apparatus, implement, machine, contrivance . . . ” or “any component, part, or accessory” thereof which is “intended for use in the diagnosis of disease or other conditions, or in the cure, mitigation, treatment, or prevention of disease.”³ Despite its physical intangibility, CDS software arguably meets this definition. For many years, the FDA has regulated “software in a medical device”⁴ – software embedded in traditional hardware devices like x-ray machines, where the software affects the safety and effectiveness of the device as a whole.⁵ In 2013, as CDS software was growing more common in clinical health care, the FDA worked with medical product regulators in other countries to develop the concept of “software as a medical device” (SaMD): standalone medical software, designed to run on diverse platforms such as smartphones, laptops, or in the cloud, that constitutes a medical device in its own right.⁶ The notion was that when software is intended for use in diagnosing, treating, or preventing disease, then the software is itself a medical device, and its status as a device does not hinge on being incorporated into specific hardware.

Concerned that the FDA might be contemplating broad regulation of standalone medical software, the software industry pressed Congress for clarification. In December 2016, Congress responded in Section 3060 of the 21st Century Cures Act (the “Cures Act”).⁷ Section 3060 includes some (but not all) CDS software in the definition of a device that the FDA can regulate and provides a jurisdictional rule distinguishing which software is – and which is not – a medical device.⁸ In two subsequent draft guidance documents,⁹ the FDA has attempted to clarify this distinction, but key uncertainties remain unresolved for CDS software that incorporates AI/ML techniques.

Whether a piece of software is subject to FDA oversight has important legal impacts apart from the immediate burden and delay of having to comply with the FDA’s regulations. This chapter explains why the FDA’s regulation of medical

² See US Food & Drug Admin., *Clinical Decision Support Software: Draft Guidance for Industry and Food and Drug Administration Staff* (Sept. 2019), www.fda.gov/regulatory-information/search-fda-guidance-documents/clinical-decision-support-software; see also US Food & Drug Admin., *Clinical and Patient Decision Support Software: Draft Guidance for Industry and Food and Drug Administration Staff* (Dec. 2017) (providing earlier draft guidance replaced in Sept. 2019).

³ 21 U.S.C. § 321(h).

⁴ See Int’l Medical Device Regulator’s Forum, *Software as a Medical Device (SaMD): Key Definitions* (Dec. 9, 2013), www.imdrf.org/docs/imdrf/final/technical/imdrf-tech-131209-samd-key-definitions-140901.pdf.

⁵ *Id.*; see also, US Food & Drug Admin., *What Are Examples of Software as a Medical Device?* (updated Dec. 6, 2017), www.fda.gov/medicaldevices/digitalhealth/softwareasamedicaldevice/ucm587924.htm.

⁶ See Int’l Medical Device Regulator’s Forum, *supra* note 4.

⁷ 21st Century Cures Act, Pub. L. No. 114–255, § 3060(a), 130 Stat. 1033 (2016).

⁸ 21 U.S.C. § 360j(o)(1)(E).

⁹ See US Food & Drug Admin., *supra* note 2.

software could increase the likelihood that state courts would view it as a product that is subject to strict product liability tort regimes. Software liability has long been a contested topic. Courts have shown reluctance to apply product liability to software, whether because its intangible nature seems at odds with the notion of a product, or because software seems better characterized as a service.¹⁰ If classified as a service, professional malpractice or ordinary negligence regimes would apply to software vendors. If classified as a product, they could face product liability (which encompasses both negligence and strict liability claims). The fact that product vendors face product liability does not prevent plaintiffs from also bringing malpractice suits against physicians and other health care professionals who ordered, prescribed, or used a defective product in the course of treating the patient. Product liability and malpractice coexist in the medical setting, and a single injury can generate both types of suit.

This chapter briefly explains the jurisdictional rule Congress set out in Section 3060 of the Cures Act and identifies key uncertainties after the FDA's two recent attempts at clarification. The chapter next summarizes some of the policy rationales for product liability and their applicability to CDS software. The chapter then explores two intriguing types of product liability suits that could emerge in connection with FDA-regulated AI/ML CDS software.

2.2 THE FDA'S AUTHORITY TO REGULATE CDS SOFTWARE

The very fact that the FDA regulates a piece of software militates in favor of its classification as a product, as opposed to an informational or professional service, potentially subjecting it to product liability suits. This proposition may strike readers as nonobvious, but it is an artifact of how the FDA's jurisdiction is defined under the Food, Drug, and Cosmetic Act (FDCA).

A key divide in health law is between FDA regulation of medical products versus state-level licensure directed at health care services such as the practice of medicine. "The scope of FDA's power is defined almost entirely by the list of product categories over which it has jurisdiction."¹¹ The major exception is that the FDA shares broad powers with the Centers for Disease Control and Prevention to manage the spread of communicable diseases, but those powers arise under a different statute.¹² Under the FDCA, the FDA's ability to regulate persons or entities rests on whether they are developing, manufacturing, shipping, storing, importing, or selling an item that fits within one of the product categories that Congress authorizes the FDA to regulate: drugs, devices, biological products, food,

¹⁰ See Joseph L. Reutiman, *Defective Information: Should Information Be a Product Subject to Products Liability Claims*, 22 *Cornell J. L. and Pub. Pol'y* 194–6 (2012) (discussing cases that have treated software as a service).

¹¹ Peter Barton Hutt et al., *Food and Drug Law* 77 (4th ed. 2014).

¹² See *id.* (discussing the FDA's jurisdiction under the Public Health Service Act).

animal drugs, et cetera.¹³ The FDA's regulatory authority under the FDCA extends to products rather than services.¹⁴ Medical devices, as FDA-regulated products, are routinely subject to product liability suits.¹⁵

When the FDA asserts that it has jurisdiction to regulate something, the agency is making a determination that that thing fits into one of these congressionally defined product categories, and therefore is not a service. Once the FDA determines that something is a product, it is conceivable that a state court hearing a tort lawsuit might disagree, but this is unlikely. Doing so would amount to a state court finding that the FDA regulated outside of its lawful jurisdiction. Suits challenging the FDA's jurisdiction pose federal questions to be heard in federal court, not state court. Moreover, the FDA is making scientific/technical determinations when it classifies something as a medical product, and courts (both state and federal) tend to give "super deference" to such decisions.¹⁶ If the FDA determines that software fits within Congress's definition of a medical device, and therefore is a product, state courts seem likely to defer.

The jurisdictional rule for CDS software under the Cures Act carefully respects the line between products and services, as has all FDA-related legislation dating back to the 1930s when the scope of the FDA's power to regulate medical practice was hotly debated before Congress passed the FDCA.¹⁷ Congress denied intent for FDA regulation of medical products to encompass regulation of health care services, a traditional province of the states.¹⁸ As a policy matter, the FDA seeks to avoid regulating physicians' activities, even though courts have never found constitutional limits on the FDA's power to do so.¹⁹ "There is little doubt under modern law that Congress has ample power to regulate the manufacture, distribution, and use of drugs and medical devices."²⁰ Regulating use is tantamount to regulating health care

¹³ See 21 U.S.C. § 321 (defining these and other product categories that trigger FDA jurisdiction).

¹⁴ US Food & Drug Admin., What Does FDA Regulate? (2018), www.fda.gov/about-fda/fda-basics/what-does-fda-regulate.

¹⁵ Elizabeth O. Tomlinson, Proof of Defective Design of Medical Device in Products Liability Action, 149 Am. Jur. Proof of Facts 407 (2015); See also Barbara J. Evans & Ellen Wright Clayton, Deadly Delay: The FDA's Role in America's COVID-Testing Debacle, 130 Yale Law Journal Forum 78–100 (2020), www.yalelawjournal.org/forum/deadly-delay-the-fdas-role-in-americas-covid-testing-debacle (discussing the product/service distinction in FDA regulation of diagnostics).

¹⁶ Emily Hammond Meazell, Super Deference, the Science Obsession, and Judicial Review as Translation of Agency Science, 109 Mich. L. Rev. 733 (2010–11).

¹⁷ See Joel E. Hoffman, Administrative Procedures of the Food and Drug Administration, in 2 Fundamentals of Law and Regulation: An In-Depth Look at Therapeutic Products 13, 17–24 (David G. Adams et al. eds., 1999) [hereinafter *Fundamentals of Law and Regulation*].

¹⁸ See Legal Status of Approved Labeling of Prescription Drugs; Prescribing for Uses Unapproved by the Food and Drug Administration, 37 Fed. Reg. 16,503 (Aug. 15, 1972) (discussing Congress's legislative intent).

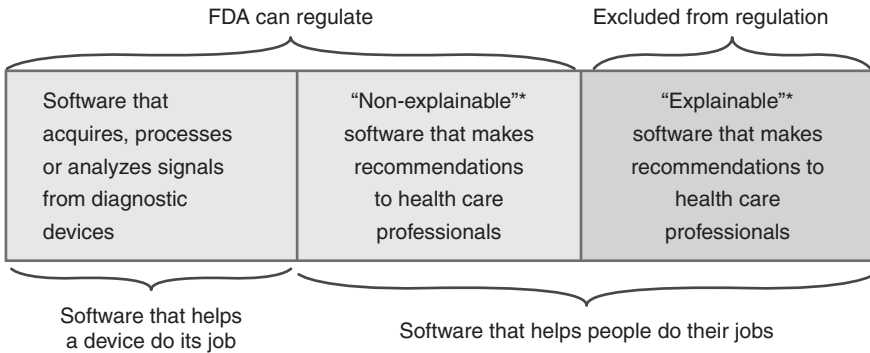
¹⁹ *Id.*; see also David G. Adams, The Food and Drug Administration's Regulation of Health Care Professionals, in 2 Fundamentals of Law and Regulation, *supra* note 17, at 423–5.

²⁰ Richard A. Epstein, Why the FDA Must Preempt Tort Litigation: A Critique of Chevron Deference and a Response to Richard Nagareda, 1 J. Tort L. 7 (2006), www.bepress.com/jtl/vol1/iss1/art5 (emphasis added).

services when, for many types of devices used in health care facilities, the provider rather than the patient is the user.²¹ Tension is seen in the 1976 Medical Device Amendments, which authorize the FDA to approve medical devices subject to restrictions on their use,²² but expressly forbid the FDA to interfere with physicians’ discretion to use those devices however they see fit in the context of practitioner–patient relationships.²³

The product/service distinction grows even more strained when the device is CDS software, which by its very design is intended to influence the practice of medicine. The Cures Act traces a line between CDS software that performs device-like functions (which the FDA appropriately can regulate) versus CDS software whose functions resemble medical practice (which the FDA should not regulate).²⁴ The baseline assumption is that CDS software performs a practice-related function and should be excluded from the FDA’s oversight. Congress recognizes two situations, however, where FDA oversight is appropriate. These are portrayed in [Figure 2.1](#).

At the far left, the FDA can regulate CDS software when its “function is intended to acquire, process, or analyze a medical image or a signal from an



§360j(o)(1)(E)(iii) exclusion criterion – Is the software intended to enable the health care professional to independently review the basis for its recommendations, so that it is not the intent for the health care professional to rely primarily on the software to make a clinical diagnosis or treatment decision for an individual patient? If No, it falls under FDA’s regulations as a device. If Yes, FDA cannot regulate it.

FIGURE 2.1: The FDA’s jurisdiction to regulate CDS software under the Cures Act

²¹ See Patricia J. Zettler, *Toward Coherent Federal Oversight of Medicine*, 52 *San Diego L. Rev.* 427 (2015) (exploring de facto FDA regulation of medical practice).

²² Medical Device Amendments of 1976, Pub. L. No. 94–295, § 2, 90 Stat. 539, 565 (adding Section 520(e) of the Food, Drug, and Cosmetic Act) (codified as amended at 21 U.S.C. § 360j(e) (2006)).

²³ 21 U.S.C. § 396.

²⁴ 21 U.S.C. § 360j(o)(1)(E).

in vitro diagnostic device or a pattern or signal from a signal acquisition system.”²⁵ The FDA has, for many years, regulated this type of software which includes, for example, software that enhances mammogram images to highlight areas suspicious for disease.²⁶ In one sense, this is CDS software because it helps a human actor – the radiologist – make a diagnosis. Still, another way to view it is that the software is helping a device (the imaging machine) do its job better by transforming outputs into a user-friendly format. By leaving such software under the FDA’s oversight, the Cures Act treats it as mainly enhancing the performance of the device rather than the human using the device. The software is, in effect, a device accessory, and an accessory to a device is itself a device that the FDA can regulate.²⁷

The FDA can regulate some, but not all, of the remaining CDS software which more directly aims to bolster human performance. The Cures Act allows the FDA to regulate CDS software if it is not intended to enable the “health care professional to independently review the basis for such recommendations that such software presents” so that there is an intent that the “health care professional rely primarily on any of such recommendations to make a clinical diagnosis or treatment decision regarding an individual patient.”²⁸ This murky language expresses a rather simple concept: the FDA can regulate CDS software if the developer intends for it to operate as a “black box,” to use Nicholson Price’s phrase.²⁹ To use engineering parlance, the FDA can regulate AI/ML CDS software if it is not intended to function as explainable artificial intelligence (XAI).³⁰ CDS software that makes recommendations falls under the FDA’s regulatory jurisdiction if those recommendations are not intended to be transparent to the health care professionals using the software. On the other hand, if CDS software is transparent enough that a health care professional would be able to understand its recommendations and challenge them – that is, when it is not a black box – then Congress excludes it from FDA regulatory oversight.

The Cures Act parses the product/practice regulatory distinction as follows: Congress sees it as a medical practice issue (instead of a product regulatory issue) to make sure health care professionals safely apply CDS software recommendations that are amenable to independent professional review. In that situation, safe and effective use of CDS software is best left to clinicians and to their state practice regulators, institutional policies, and the medical

²⁵ *Id.*

²⁶ Bradley Merrill Thompson, Learning from Experience: FDA’s Treatment of Machine Learning, *Mobile Health News* (Aug. 23, 2017), www.mobihealthnews.com/content/learning-experience-fda%E2%80%99s-treatment-machine-learning; [<https://perma.cc/Q95C-9R22>].

²⁷ 21 U.S.C. § 321(h).

²⁸ 21 U.S.C. § 360j(o)(1)(E)(iii).

²⁹ W. Nicholson Price II, *Regulating Black-Box Medicine*, 116 *Mich. L. Rev.* 421–74 (2017).

³⁰ Enrico Tjoa & Cuntai Guan, *A Survey on Explainable Artificial Intelligence (XAI): Towards Medical XAI*, <https://arxiv.org/pdf/1907.07374.pdf>.

profession. When CDS software is not intended to be independently reviewable by the health care provider at the point of care, there is no way for these bodies to police appropriate clinical use of the software. In that situation, the Cures Act tasks the FDA with overseeing its safety and effectiveness. Doing so has the side effect of exposing CDS software developers to a risk of product liability suits. Product liability regimes may provide a useful legal framing for some of the problems CDS software presents.

2.3 WHY THERE MAY BE A ROLE FOR PRODUCT LIABILITY

An emerging literature on the limits of AI and big data analytics has raised serious concerns about the safety of these technologies, including in their CDS software applications. Lack of reproducibility may occur because of nonrepresentative datasets, or because vendors and developers refuse to permit others to scrutinize their wares. As Rebecca Robbins reported in 2020, “Some of these AI models are fraught with bias, and even those that have been demonstrated to be accurate largely haven’t yet been shown to improve patient outcomes. Some hospitals don’t share data on how well the systems work.”³¹ Narrow validity undermines some models’ applicability in certain health care settings, but overblown claims of accuracy or assistance can lead physicians not to mention that the software is in use, much less to seek patients’ informed consent to it.³² Data opacity also creates situations where even those who might be concerned about CDS software cannot adequately complete due diligence or otherwise explore its limits. Dr. Eric Topol summarized many examples of these problems (lack of reproducibility, narrow validity, overblown claims, and nontransparent or hidden data).³³

There is also concern that the data involved may not merely lack representativeness generally but may be biased in particularly troubling ways. Datasets may inadequately reflect all groups in society,³⁴ or may underinclude women³⁵ and

³¹ Rebecca Robbins, *An Invisible Hand: Patients Aren’t Being Told about the AI Systems Advising Their Care*, *Stat News* (July 15, 2020), www.statnews.com/2020/07/15/artificial-intelligence-patient-consent-hospitals/.

³² See W. Nicholson Price II, *Medical AI and Contextual Bias*, 33 *Harv. J. L. & Tech.* 66 (2019) (discussing narrow validity of AI systems developed in resource-rich contexts when implemented in lower-resource settings).

³³ Eric Topol, *Deep Medicine: How Artificial Intelligence Can Make Healthcare Human Again* (Basic Books ed., 2019); See also Matthew Zook et al., 10 *Simple Rules for Responsible Big Data Research*, 13 *PLoS Computational Biology* (2017) (identifying similar limits); Danah Boyd & Kate Crawford, *Critical Questions for Big Data*, 15 *J. Info., Comm’n and Soc’y* 662–79 (2012).

³⁴ Ziad Obermeyer et al., *Dissecting Racial Bias in an Algorithm Used to Manage the Health of Populations*, 366 *Science* 447–53 (2019); Ruha Benjamin, *Assessing Risk, Automating Racism*, 366 *Science* 421–2 (2019). But see Frank Pasquale & Danielle Keats Citron, *Promoting Innovation While Preventing Discrimination: Policy Goals for the Scored Society*, 89 *Wash. L. Rev.* 1413–24 (2014) (discussing ways to address biases).

³⁵ Carolyn Criado Perez, *Invisible Women: Data Bias in a World Designed for Men* (Abrams ed., 2019).

overrepresent persons of European ancestry,³⁶ causing the software to provide unreliable or unsafe recommendations for the underrepresented groups.³⁷

Some observers hope that the FDA or National Institute of Standards and Technology will gradually nudge CDS software vendors toward better practices. However, the current path of development of medical software casts doubt on whether FDA oversight can fulfil this role. The agency's Digital Innovation Action Plan³⁸ and its Digital Health Software Precertification (Pre-Cert) Program³⁹ acknowledge these concerns:

FDA's traditional approach to moderate and higher risk hardware-based medical devices is not well suited for the faster iterative design, development, and type of validation used for software-based medical technologies. Traditional implementation of the premarket requirements may impede or delay patient access to critical evolutions of software technology.⁴⁰

In response, the FDA is "reimagining its approach to digital health medical devices,"⁴¹ but the agency's policies are still a work in progress. Roiled by long-term trends toward underfunding, politically motivated attacks on its expertise, and flagging public confidence in the wake of the US COVID-19 debacle, the FDA faces a difficult path ahead and may be particularly challenged when it comes to regulating the safety and effectiveness of AI/ML CDS software.⁴² The agency's priorities may justifiably be elsewhere, and its ability to recruit experts at government salary scales is suspect when AI/ML experts command significantly more than current public compensation levels.

When diagnostic AI ignores problems with inclusivity and bias yet still manages to deliver better results than unaided human observation for many or most patients, the patients who do suffer an injury may not have a tort remedy under a negligence standard – particularly if the standard of care is unaided human observation. Even if standard-setting bodies enunciate standards for database inclusion, many states continue to base negligence liability on customary standards of care.⁴³ The *next*

³⁶ Alice B. Popejoy et al., *The Clinical Imperative for Inclusivity: Race, Ethnicity, and Ancestry (REA) in Genomics*, 39 *Human Mutation* 1713–20 (2018).

³⁷ Adewole S. Adamson & Avery Smith, *Machine Learning and Health Care Disparities in Dermatology*, 154 *JAMA Dermatology* 1247 (2018).

³⁸ US Food & Drug Admin., *Digital Health Innovation Action Plan* (2017), www.fda.gov/downloads/MedicalDevices/DigitalHealth/UCM568735.pdf.

³⁹ US Food & Drug Admin., *Digital Health Software Precertification (Pre-Cert) Program*, www.fda.gov/medical-devices/digital-health-center-excellence/digital-health-software-precertification-pre-cert-program.

⁴⁰ US Food & Drug Admin., *supra* note 38, at 2.

⁴¹ *Id.* at 5.

⁴² On the history of challenges to the FDA, see Frank Pasquale, *Grand Bargains for Big Data: The Emerging Law of Health Information*, 72 *Md. L. Rev.* 682 (2013); Efthimios Parasidis, *Patients over Politics: Addressing Legislative Failure in the Regulation of Medical Products*, 5 *Wis. L. Rev.* 929 (2011).

⁴³ Frank Pasquale, *Data Informed Duties*, 119 *Colum. L. Rev.* 1917–39 (2019).

section explores whether failures to use more representative databases might be deemed actionable in strict product liability.

The FDA's announced approaches for regulating Software as a Medical Device (SaMD) seemingly would not preempt state product liability suits under doctrines announced in prior medical device cases like *Medtronic v. Lohr*⁴⁴ and *Riegel v. Medtronic*.⁴⁵ This opens the door for product liability suits to help fill the regulatory gaps and help incentivize quality improvement, accountability, and responsibility that an overburdened FDA may be incapable of ensuring.⁴⁶

Other factors suggesting a need for product liability include medical software contract practices that blunt the impact of negligence suits against software developers. It is common for developers to shield themselves from negligence through license terms that shift liability to (or require indemnification from) health care providers that use their software.⁴⁷ Such terms are seen, for example, in vendor contracts for electronic health record (EHR) systems, which may also include alternative dispute resolution procedures and gag clauses that stifle public disclosure of safety problems.⁴⁸ Patients hurt by defective medical software might attempt to sue their health care provider but would face challenges establishing negligence of the software developer. The provider, who might possess facts bearing on the developer's negligence, cannot pursue claims under terms of the licensing agreement. The result is to channel negligence claims toward providers while the software developer goes unscathed. In contrast, product liability widens opportunities for patients to sue any party in the chain of commerce that resulted in their injuries. Private contracts between software developers and health care providers can foreclose suits between those two signatories but cannot waive the rights of patients to sue developers whose defective software causes medical injuries.

The next section explores two possible product liability causes of action that offer promise for this gap-filling role. The first is manufacturing defect suits for lack of explainability, when software fails to live up to developers' claims that the algorithm is transparent to physicians tasked with using it. The second is design defect suits when software uses training or operational datasets that are too small, inaccurate, biased, or otherwise inappropriate for the actual patients for whom the software renders recommendations.

⁴⁴ 518 U.S.C. § 470 (1996).

⁴⁵ 552 U.S.C. § 312 (2008); See also Barbara J. Evans, *The Streetlight Effect: Regulating Genomics Where the Light Is*, 48 J. L., Med. Ethics Supp: LawSeq 105 (2020) (discussing why the FDA's proposed approaches do not appear to preempt failure-to-warn suits).

⁴⁶ There is not complete harmony between tort and regulation here; some preemption issues may arise, for example, if high-risk software were regulated as a Class III medical device. See, e.g., Charlotte Tschider, *Medical Device Artificial Intelligence: The New Tort Frontier*, 46 *BYU L. Rev.* 1551, 1573–86 (2021), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3443987.

⁴⁷ Liis Vihul, *The Liability of Software Manufacturers for Defective Products*, Tallinn Paper No. 2 (Cooperative Cyber Defense Center of Excellence ed., 2014).

⁴⁸ See Jim Hawkins et al., *Nontransparency in Electronic Health Record Systems*, in *Transparency in Health and Health Care in the United States* 273–85 (Holly F. Lynch et al. eds., 2019).

2.4 CAN MANUFACTURING DEFECT SUITS PROMOTE EXPLAINABILITY OF AI/ML CDS SOFTWARE?

In its 2017 and 2019 draft guidance documents on CDS software,⁴⁹ the FDA failed to clarify the central jurisdictional enigma in the Cures Act: How will the agency determine whether AI/ML CDS software is intended to enable the health care professional to independently review the basis for the software's recommendations? This section describes the problem and explores whether product liability suits might help.

The FDA has a clear process for assessing device manufacturers' intent,⁵⁰ but needs to explain how this process applies to developers of AI/ML software: How, exactly, will the FDA assess whether a software developer intends for its software to be explainable? The agency could, for example, describe algorithmic features that support an inference that CDS software is medical XAI. Alternatively, the agency could prescribe a clinical testing process (such as having physicians use the software and surveying whether they understand the basis for its decisions). The FDA has done neither.

The draft guidance documents both view simple, rule-based CDS systems – those that merely apply existing knowledge that is “publicly available (e.g., clinical practice guidelines, published literature)”⁵¹ – as meeting the § 360j(o)(1)(E)(iii) “explainability” standard, thus escaping FDA regulation. The 2017 draft guidance did not directly discuss AI/ML systems that derive and apply new insights from real-world evidence. It seemed to presume that all such systems would be subject to FDA regulation – an overly expansive view of the FDA's authority that ignored the jurisdictional rule in the Cures Act. The 2019 draft guidance acknowledges that the explainability of AI/ML software is a key jurisdictional issue but failed to provide any standards or processes for judging whether software is intended to be explainable.

This default has serious consequences. If the FDA deems all but the simplest CDS systems to be unexplainable, this could have detrimental impacts on innovation and on patients. What incentive will software developers have to invest in making AI/ML medical software more explainable to physicians, if the FDA deems all such software to be unexplainable no matter what they do? Simple, rule-based CDS software would escape FDA oversight. Promising AI/ML software to enable a learning health care system informed by real clinical evidence might face long regulatory delays.

The FDA's failure to set standards – or at least a process – for assessing software explainability leaves the agency with no basis to rebut developers' claims that their software is intended to be explainable and to allow independent review by

⁴⁹ See US Food & Drug Admin., *supra* note 2.

⁵⁰ See 21 C.F.R. § 801.4.

⁵¹ See US Food & Drug Admin 2017 Draft Guidance, *supra* note 2, at 8.

physicians. 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 healthcare professional” and “not intended to serve as the primary basis for making a clinical diagnosis or treatment decision regarding an individual patient.”⁵² Developers have strong incentives to pursue this strategy. They might escape FDA regulation under the Cures Act’s jurisdictional rule. This in turn would let them argue that their software is a service, rather than an FDA-regulated product subject to product liability. Any physician that relies on the software’s recommendations as the main basis for decision making would be using it off-label, and negligence liability for off-label use rests with the physician, rather than the software developer. Why would a rational software developer not try this strategy?

Strict product liability might provide an answer. Under the Third Restatement of Torts, a plaintiff establishes a manufacturing defect by showing that a product “departs from its intended design even though all possible care was exercised in the preparation and marketing of the product.”⁵³ The plaintiff merely needs to show that the product deviated from the intended design when it left the developer’s possession.⁵⁴ If an AI/ML CDS software developer states that it intended for its software to allow independent review by physicians, and perhaps even escaped FDA oversight by making that claim, then that proves that the software was intended to be explainable. If the software later lacks explainability, hindering independent physician review, then the software clearly departs from its intended design and has a manufacturing defect. Plaintiffs seemingly face low evidentiary hurdles to establish the defect: they could call their physician to the witness stand and ask the physician to explain to the jury how the AI/ML software reached its recommendations that affected patient care. If the physician cannot do so, the plaintiff would have proved the defect. In light of the FDA’s ongoing failure to enunciate standards for AI/ML software explainability, manufacturing defect suits are a promising tool to incentivize investment in improved explainability (and frank disclosures when explainability is lacking).

2.5 CAN DESIGN DEFECTS PROMOTE THE USE OF APPROPRIATE TRAINING DATASETS?

The FDA’s 2017 draft guidance on CDS software suggested that the Cures Act explainability standard cannot be met unless physician users have access to the data underlying a software product’s recommendations.⁵⁵ The agency backed away from this position in its 2019 draft guidance, possibly reflecting the reality that

⁵² See 21 U.S.C. § 360j(o)(1)(E)(iii).

⁵³ Restatement (Third) of Torts: Products Liability § 2(a).

⁵⁴ *Id.*

⁵⁵ 21 U.S.C. § 360j(o)(1)(E)(iii).

software developers are deeply opposed to sharing their proprietary training datasets with anyone – neither users nor regulators. At most, developers express willingness to share summary statistics, such as the kinds of health conditions, demographics, and number of patients included in the training dataset. The FDA’s oversight of AI/ML training datasets thus seems destined to be cursory.

There have been calls for software developers to have legal duties relating to the accuracy and appropriateness (representativeness) of training datasets, as well as the integrity of all data inputs and the transparency of outputs.⁵⁶ Prospective regulation by the FDA is proving an uncertain legal vehicle for establishing such duties. Can design defect suits address this deficiency?

“Strict” product liability/design defect suits allege that a product is unreasonably dangerous even though it may conform to its intended design. Complex products like CDS software are unsuitable for a consumer expectations test, which applies only if jurors would be able to understand a product’s risks without the aid of expert witnesses. Courts likely would apply a risk-utility test, which usually involves requiring the plaintiff to show that a reasonable alternative design (RAD) existed at the time of sale and distribution. This reliance on reasonability concepts causes strict liability suits for design defects to bear a considerable resemblance to negligence suits, which is why this paragraph put “strict” between quotation marks.

Selection of the training dataset is a central design decision when developing AI/ML software. If the training dataset is too small, inappropriate, inaccurate, or biased and nonrepresentative of patients the software later will analyze, then the software – by its design – cannot provide accurate recommendations for their care. An alternative design seemingly always exists: that is, train the software on a larger, more appropriate, more accurate, less biased dataset that better reflects the intended patient population. However, the “R” in RAD stands for “reasonable,” and it would be left for the trier of fact to decide whether it would have been reasonable for the software developer to have used that alternative, better dataset, in view of the cost, delay, availability, and accessibility of additional data.

Framing the problem as a design defect of the AI/ML software (which in most cases will be an FDA-regulated product) may avert some of the difficulties seen in prior product liability suits alleging defects in information itself. Because information is intangible, some courts struggle with treating it as a product and applying strict liability.⁵⁷ Suits for defective graphic presentation of information occasionally succeed, as in *Aetna Casualty and Surety v. Jeppeson & Co.*,⁵⁸ involving a deadly air crash after the pilot relied on a Jeppeson instrument approach chart – a product consisting almost entirely of the graphic presentation of information, which the

⁵⁶ Pasquale, *supra* note 43.

⁵⁷ Reutiman, *supra* note 10, at 183.

⁵⁸ 642 F.2d 339 (9th Cir. 1981); see also *Brocklesby v. United States*, 767 F.2d 1288, 1294–5 (9th Cir. 1985), and discussion in Oren Bracha & Frank Pasquale, *Federal Search Commission? Access, Fairness, and Accountability in the Law of Search*, 93 *Cornell L. Rev.* 1149, 1194 (2008).

district court found defective. That case is considered anomalous, however, and many courts hesitate to allow design defect suits over deadly information, whether on First Amendment grounds or reluctance to hinder free flows of information in our society.⁵⁹ Suits for defective information seem most likely to fail when the information in question resembles expressive content,⁶⁰ which might not be an issue for AI/ML training datasets, which are not expressive. Still, courts have a well-known reluctance to treat information as a “product” that was “defective.” The approach proposed here avoids this problem. The alleged defect is not in the information itself, but in the design of the software product that relied on the information. The information in an AI/ML training dataset is best conceived as a design feature of the software rather than a product in its own right.

2.6 CONCLUSION

Some commentators express concern that applying product liability to software could have adverse impacts on innovation and might delay diffusion of software.⁶¹ We agree that these are valid concerns that courts will need to weigh carefully when considering claims by patients injured during the use of AI/ML CDS software. At the same time, however, a vast and growing literature on algorithmic accountability and critical algorithm studies has painstakingly documented that AI/ML software, even if it provides useful results for most people, can harm members of groups that were underrepresented in the datasets on which the software relies.⁶² Such injuries are predictable and need remedies when they do occur. Product liability should not be ruled out. Slowing the diffusion of software might well be justified if the software injures people or entrenches historical disparities in access to high-quality health care.

Other commentators question applying product liability to AI/ML continuous-learning software which can evolve independently of the manufacturer.⁶³ To date, the FDA has only cleared or approved software that is “locked” – that is, stops evolving – prior to the FDA’s review, which removes this concern. As continuously learning software does reach the market, the possibility of software evolution

⁵⁹ See Reutiman, *supra* note 10, at 188–9.

⁶⁰ See, e.g., *Winter v. G.P. Putnam’s Sons*, 938 F.2d 1033 (9th Cir. 1991) (no liability for dangerous misinformation in a book); *Wilson v. Midway Games, Inc.*, 198 F. Supp 2d 167 (D. Conn. 2002) (rejecting claim that a video game was dangerously defective for stimulating violent behavior in users, noting similarities to expressive media like movies and television).

⁶¹ See, e.g., Bryan H. Choi, *Crashworthy Code*, 94 Wash. L. Rev. 39 (Mar. 2019); Jamil Ammar, *Defective Computer-Aided Design Software Liability in 3d Bioprinted Human Organ Equivalents*, 35 Santa Clara High Tech. L. J. 58 (2019); Karni A. Chagal-Feferkorn, *Am I an Algorithm or a Product? When Products Liability Should Apply to Algorithmic Decision-Makers*, 30 Stan. L. & Pol’y Rev. 82 (2019).

⁶² Frank Pasquale, *The Black Box Society* (Harvard University Press ed., 2015); Frank Pasquale, *New Laws of Robotics* (Harvard University Press ed., 2020).

⁶³ Jacob Turner, *Robot Rules: Regulating Artificial Intelligence 94–100* (Palgrave Macmillan ed., 2018).

underscores the need to program in restraints and checks against problematic forms of evolution.⁶⁴ The FDA has not explained how it will (or whether it can) ensure such restraints. Product liability has long served alongside the FDA's oversight to promote patient safety.

Some commentators endorse product liability in the CDS context, pointing to the need for courts to recognize and counteract automation bias, which can arise when often overworked professionals seek tools to ease their workload.⁶⁵ More stringent product liability standards are a way of promoting a lower risk level in the health care industry and can ease the difficulties injured patients will face establishing negligence, given the extraordinarily complex and even trade-secret protected methods used to develop CDS software.⁶⁶

The time may be right to reconsider product liability for medical software. The FDA's foray into this regulatory sphere bestows "product" status on medical software that courts often have tended to view as information services. By doing so, the agency's regulation of CDS software opens the door to product liability suits. This chapter has suggested two examples that merit further study. Such suits could help nudge software developers to improve the explainability of their software and ensure appropriate training datasets and could promote greater industry transparency about CDS software on which patients' lives may depend.

⁶⁴ Stuart J. Russell, *Human Compatible: Artificial Intelligence and the Problem of Control* (New York: Penguin Random House ed., 2019).

⁶⁵ Efthimios Parasidis, *Clinical Decision Support: Elements of a Sensible Legal Framework*, 20 *J. Healthcare L. & Pol'y* (2018); see also Nicholas Carr, *The Glass Cage: How Our Computers Are Changing Us* (W.W. Norton ed., 2015); see also, Kevin R. Pinkney, *Putting Blame where Blame Is Due: Software Manufacturer and Customer Liability for Security-Related Software Failure*, 13 *Albany L. J. Sci. & Tech.* (2002) (focusing on security defects); Michael D. Scott, *Tort Liability for Vendors of Insecure Software: Has the Time Finally Come?*, 67 *Md. L. Rev.* 469–70 (2017) (same).

⁶⁶ Frances E. Zollers et al., *No More Soft Landings for Software: Liability for Defects in an Industry That Has Come of Age*, 21 *Santa Clara Computer & High Tech. L. J.* 777 (2005).