

University of Michigan Law School

University of Michigan Law School Scholarship Repository

Book Chapters

Faculty Scholarship

2018

Medical Malpractice and Black-box Medicine

W. Nicholson Price II

University of Michigan Law School, wnp@umich.edu

Available at: https://repository.law.umich.edu/book_chapters/169

Follow this and additional works at: https://repository.law.umich.edu/book_chapters



Part of the [Health Law and Policy Commons](#), and the [Medical Jurisprudence Commons](#)

Publication Information & Recommended Citation

Price, W. Nicholson, II. "Medical Malpractice and Black-Box Medicine." In *Big Data, Health Law, and Bioethics*, edited by I. Glenn Cohen, Holly F. Lynch, Effy Vayena, and Urs Gasser, 295-306. Cambridge: Cambridge University Press, 2018.

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

Medical Malpractice and Black-Box Medicine

*W. Nicholson Price II**

20.1 INTRODUCTION

The explosive proliferation of health data has combined with the rapid development of machine-learning algorithms to enable a new form of medicine: “black-box medicine.”¹ In this phenomenon, algorithms troll through tremendous databases of health data to find patterns that can be used to guide care, whether by predicting unknown patient risks, selecting the right drug, suggesting a new use of an old drug, or triaging patients to preserve health resources. These decisions differ from previous data-based decisions because black-box medicine is, by its nature, opaque; that is, the bases for black-box decisions are unknown and unknowable.

Black-box medicine raises a number of legal questions, ranging from how to shape incentives for its development to how to regulate its growth and quality.² One key question is how black-box medicine will influence the medical malpractice liability of healthcare providers. How should tort liability apply to providers who cannot know the mechanistic underpinnings of the treatment they recommend? Must they learn as much as they can about the way algorithms are developed and verified? Or can they rely on the assurances of the developer without more knowledge?

This chapter explores the medical malpractice implications of black-box medicine. It briefly introduces the phenomenon and then considers how the tort system does, can, and should regulate the behavior of providers and healthcare facilities using black-box medical techniques. It concludes that while providers and facilities are ill suited to evaluate the substantive accuracy of black-box medical algorithms,

* Many thanks to Nicholas Bagley, Ana Bracic, Sherman Clark, Rebecca Eisenberg, Roger Ford, Jessica Litman, Margo Schlanger, Kayte Spector-Bagdady, and Effy Vayena for helpful conversations and advice. All errors are my own.

¹ For a detailed description of black-box medicine, see W. Nicholson Price II, *Black-Box Medicine*, 28 *Harv. J.L. Technol.* 419 (2015).

² See W. Nicholson Price II, *Big Data, Patents, and the Future of Medicine*, 37 *Cardozo L. Rev.* 1401 (2016).

they could and perhaps should be required to exercise due care to evaluate procedural quality – the expertise of the developer and the availability of independent external validation – when implementing black-box algorithms in a healthcare facility or using them to care for patients.

20.2 BLACK-BOX MEDICINE

Black-box medicine is a response to the immense complexity of biological relationships. Although we are constantly developing tools to plumb that complexity, our explicit understanding of those biological relationships is necessarily slow to develop. Putting our understanding of complex relationships into medical practice is similarly slow and cumbersome because clinical trials and traditional drug development take many years and cost hundreds of millions of dollars.

Black-box medicine seeks to exploit the tremendous amount of data being generated in healthcare to find and use these underlying relationships even without understanding them and without undergoing the expense and delay of clinical trials. Health data are proliferating rapidly and becoming newly accessible. Clinical records, long kept on paper in doctors' offices, are shifting to electronic form, as are pharmacy records and medical test results.³ Newer systemic analyses, especially genomic sequencing, are creating large volumes of patient-specific data. These data can be combined into individual, group, or population-wide collections, whether by insurers, provider networks, or government entities like the United States' newly formed \$215 million Precision Medicine Initiative.⁴

These collections of data can be used to test and develop explicit hypotheses about the biology of the body and medical implications. But the data also contain patterns that are too complex or concealed for such explicit hypothesis testing. Machine-learning algorithms, using approaches such as deep learning and neural networks, can find that sort of complex underlying pattern in the data – but cannot explain or even state what those patterns are.⁵ Just as Facebook's DeepFace algorithm can match faces in a set of digital images without explicitly classifying the features it uses for recognition,⁶ medical algorithms can (for instance) predict tumor

³ Julia Adler-Milstein et al., Electronic Health Record Adoption in US Hospitals: Progress Continues, But Challenges Persist, 34 *Health Affairs* (published online November 2015), available at <https://www.healthaffairs.org/doi/abs/10.1377/hlthaff.2015.0992> [<https://perma.cc/6RQZ-LTW4>].

⁴ Francis S. Collins & Harold Varmus, A New Initiative on Precision Medicine, 372 *N. Engl. J. Med.* 793 (2015).

⁵ See Joseph A. Cruz & David S. Wishart, Applications of Machine Learning in Cancer Prediction and Prognosis, 2 *Cancer Inform.* 59 (2007).

⁶ DeepFace: Closing the Gap to Human-Level Performance in Face Verification, Research at Facebook (June 24, 2014), available at <https://research.facebook.com/publications/480567225376225/deepface-closing-the-gap-to-human-level-performance-in-face-verification/> [<https://perma.cc/TNT3-JSUR>].

response to a particular drug based on allelic patterns among thousands of genes or predict lung cancer prognosis by analyzing microscopic images⁷ – all without understanding or identifying why or how those patterns matter.⁸ This opacity is not deliberate, though some secrecy by developers could compound it.⁹ Instead, the opacity is unavoidable. Sometimes patterns are opaque because they are too complicated; that is, even if the computer could state the set of, for example, thousands of genes and interacting patient-history factors, we could not understand it. Other times the opacity is a result of the machine-learning techniques used to find patterns; a trained neural network (one such technique) typically cannot output the artificial neurons’ “connections” in any meaningful sense and thus does not demonstrate how it reached its result. To be clear, this opacity is also not desirable; it would be preferable to know and understand the relationships being used. And there are some machine-learning methods that are more transparent. But opaque black-box methods open up for use a broad swath of nuanced biological patterns currently too complex or hidden for explicit understanding and are the focus of this chapter.

Black-box medicine accordingly has tremendous potential benefits. Most important for the context of this chapter, it can direct care, predicting a patient’s risk profile, helping choose between a selection of known interventions, or suggesting an off-label use of an approved intervention. A doctor might feed the genetic sequence of a patient’s tumor into a black-box algorithm, for instance, and receive a recommendation as to what drug is most likely to treat the tumor effectively.¹⁰ Alternatively, an opaque algorithm could continuously evaluate a trauma patient’s electronic vital signs and sound an alarm at the earliest sign of trouble, perhaps even before trained providers could observe the need.¹¹ Black-box medicine can also be used to allocate scarce healthcare resources by suggesting which patient might benefit most from an organ transplant, a hospital bed, or the attention of the first available healthcare provider.¹² In addition, black-box medicine could potentially generate hypotheses for traditional biomedical research that might eventually uncover and understand the underlying mechanisms; that is, the box need not stay black forever. While all

⁷ Kun-Hsing Yu et al., Predicting Non-Small Cell Lung Cancer Prognosis by Fully Automated Microscopic Pathology Image Features, 7 *Nat. Commun.* 12474 (2016).

⁸ See, e.g., Hoiin Moon et al., Ensemble Methods for Classification of Patients for Personalized Medicine with High-Dimensional Data, 41 *Artificial Intelligence Med.* 197 (2007).

⁹ For a description of deliberately secret algorithms, see Frank Pasquale, *The Black Box Society: The Secret Algorithms That Control Money and Information* (2015).

¹⁰ See Moon et al., Moon et al., Ensemble Methods for Classification of Patients for Personalized Medicine; Michael P. Menden et al., Machine Learning Prediction of Cancer Cell Sensitivity to Drugs Based on Genomic and Chemical Properties, 8 *PLoS ONE* e61318 (2013).

¹¹ See Nehemiah T. Liu et al., Development and Validation of a Machine Learning Algorithm and Hybrid System to Predict the Need for Life-Saving Interventions in Trauma Patients, 52 *Med. Biol. Eng. Comput.* 193 (2014).

¹² See, e.g., I. Glenn Cohen et al., The Legal and Ethical Concerns that Arise from Using Complex Predictive Analytics in Health Care, 33 *Health Affairs* 1139 (2014).

these possibilities are significant, this chapter focuses on the first: using black-box medicine to direct patient care.

As black-box medicine becomes increasingly capable of predicting patient outcomes and suggesting interventions, it will – and should – become an important part of patient care. Among many implementation issues, a key question for healthcare providers and facilities is the legal obligation of a provider or facility under the standards of medical malpractice. What must providers and facilities legally know and do when practicing black-box medicine?

20.3 LIABILITY FOR THE USE OF OPAQUE MEDICAL ALGORITHMS

Medical liability law, like tort law in general, typically serves at least two purposes: first, to compensate injured parties for their injuries and, second, to deter unreasonably dangerous behavior. To accomplish these goals, injured patients may recover from providers who provided substandard care. Patients may also sometimes recover from the healthcare enterprises involved in the provision of care, including hospitals and clinics, either vicariously based on the actions of the relevant professional or directly based on the enterprise's own duties to the patient.

Negligence actions are not the only possibility for patient recovery. In other contexts, patients can recover under a strict liability theory for injuries arising from products that are defective due to manufacturing defects, design defects, or failure to warn of risks. However, neither healthcare providers nor healthcare facilities are typically held strictly liable for defects in the products they provide, sell, or use.¹³ Such cases might be brought against black-box medicine developers.¹⁴ Patients may

¹³ See *Hollander v. Sandoz Pharm. Corp.*, 289 F.3d 1193, 1217, n. 22 (10th Cir. 2002) (“[A]n overwhelming majority of jurisdictions have refused to apply strict liability principles to claims against hospitals and physicians involving the distribution of allegedly dangerous drugs or medical devices”); see also Randolph A. Miller & Sarah M. Miller, Legal and Regulatory Issues Related to the Use of Clinical Software in Health Care Delivery, in *Clinical Decision Support* 423, 426 (Robert A. Greenes, ed., 2007) (arguing against the application of strict products liability for clinical decision-support software to providers and hospitals).

¹⁴ Potential liability for black-box medical algorithm developers could be conceived of under a product liability framework, typically based on strict liability – that is, liability for injury without any determination of fault. However, such liability is complicated by several doctrines, including the learned intermediary doctrine (limiting recovery against manufacturers where doctors prescribe drugs or devices to patients; see Timothy Hall, Reimagining the Learned Intermediary Rule for the New Pharmaceutical Marketplace, 35 *Seton Hall L. Rev.* [2005]); for the long-time immunity of software to product liability suits, see 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.* 745 (2004); to see whether software is properly classified as a good or a service, see Michael D. Scott, Tort Liability for Vendors of Insecure Software: Has the Time Finally Come, 67 *Md. L. Rev.* 425, 436–42 (2007); for the difficulty of proving causation and the possibility of preemption by regulatory regimes, see *Riegel v. Medtronic*, 552 U.S. 312, 330 (2008) (finding preemption of common-law tort claims against makers of medical devices requiring FDA preapproval).

also recover from providers who fail to obtain informed consent before undertaking a treatment that results in the patient's injury.¹⁵ Both of these complex situations are outside the scope of this chapter, which focuses on malpractice liability against providers and against healthcare enterprises.

20.3.1 Medical Malpractice by Providers

Jurisdictions vary, but in general, providers must treat patients with due expertise and care. Typically, this means that the provider must provide the level of care that would be expected of relevant members of the profession. In most jurisdictions, courts insist on a national reference group for this comparison; in a minority, a provider is held, instead, to the level of care offered by those practicing in the same locality.¹⁶ The care required may be modified based on the provider's specialized credentials or on the facilities available to the provider; a physician practicing in a small rural hospital will not be required to use the same specialized equipment available to the most well-resourced urban medical centers. Hornbook law is that adhering to customary practice will typically shield a provider from liability. The reality is more complicated, however; one observer has noted that "judicial deference to physician customs is eroding. Gradually, quietly and relentlessly, state courts are withdrawing this legal privilege."¹⁷ Moreover, what tort law requires can be explicitly modified, for instance, to encompass a certain type of care or to immunize provision of some type of care from liability regardless of common medical

¹⁵ Black-box medicine raises nuanced informed-consent issues. At an intuitive level, it is hard to imagine precisely what "informed" means in the context of a recommendation where no one knows exactly how it works. But informed consent aims to facilitate treatment decisions made in the context of a relationship between the trained doctor and the lay patient and recognizes that many things need not be disclosed. Providers must disclose to the patient information that a reasonable provider (in some jurisdictions) or a reasonable patient (in others) would find material and must give the patient a choice about accepting the treatment. See Jaime S. King & Benjamin Moulton, Rethinking Informed Consent: The Case for Shared Medical Decision-Making, 32 *Am. J.L. Med.* 429, 493–501 (2006) (appendix) (finding that about half of the United States follows a patient-based standard, about half a physician-based standard, and two a hybrid standard); see also *Canterbury v. Spence*, 464 F.2d 772 (D.C. Ct. App. 1972) (rejecting a physician-based standard and developing a patient-based standard). It is entirely possible that in most circumstances neither a reasonable provider nor a reasonable patient would find information about black-box medicine's development or opacity material to disclose, just as patients need not be informed about the strength of clinical trial evidence for most interventions recommended today.

¹⁶ See *Hall v. Hilbun*, 466 So.2d 856 (Miss. 1985) (discussing the locality-based standard of care and adopting a modified nationally based standard of care); Michelle H. Lewis et al., The Locality Rule and the Physician's Dilemma: Local Medical Practices vs. the National Standard of Care, 297 *JAMA* 2633, 2635 (2007) (tallying which states follow various versions of the locality rule).

¹⁷ Philip G. Peters, Jr., The Quiet Demise of Deference to Custom: Malpractice Law at the Millennium, 57 *Wash. Lee L. Rev.* 163, 164 (2000).

practice.¹⁸ In outlier cases, judges have held that failure to provide a specific type of care, even if uncommon, is tortious.¹⁹

Providers thus could be held liable for harmful use of black-box medical algorithms depending on the prevailing customary practice and the extent that custom is considered dispositive. As with medical innovation more generally, there is a risk of liability during this transition phase,²⁰ which presents an opportunity to consider how tort law might encourage the most beneficial medical practices.

The closest useful analogy is clinical decision-support software, designed to help providers diagnose and treat patients.²¹ This software can, for instance, provide relevant patient information, collate test results, and suggest diagnoses or treatments based on well-known explicit relationships. Clinical decision-support software has been analyzed as an aid to physicians exercising independent judgment when directing care; it merely “augments the physician’s existing knowledge by providing further information.”²² Under this logic, a trained provider should be subject to the exact same standard of negligence irrespective of whether clinical decision-support software is used because any treatment decisions are ultimately his or her own. Software provides information, but the knowledgeable provider intervenes to make the final choice.

This knowledgeable intervention, however, is precisely what is different about black-box medicine. Because neither providers nor developers know the relationships underlying the recommendations of black-box medicine, the physician cannot stand as merely the final step in a sequence of care. Once he or she has decided to use a particular black-box algorithm – itself a complex choice²³ – he or she cannot understand and thus verify the algorithm’s recommendation against his or her body of substantive expertise; the physician can only accept what the algorithm

¹⁸ See, e.g., 18 Vt. Stat. § 5281 (immunizing physicians for civil liability for prescribing life-ending medication in compliance with certain statutory provisions).

¹⁹ See *Helling v. Carey*, 83 Wn.2d 514 (Wash. 1974) (holding that ophthalmologists must test for glaucoma in patients under age forty as well as those over age forty, although under-forty testing was not common practice). On the impact of *Helling*, which was controversial and soon countermanded by state statute, see Jerry Wiley, *The Impact of Judicial Decisions on Professional Conduct: An Empirical Study*, 55 *S. Cal. L. Rev.* 345 (1982).

²⁰ See Anna B. Laakmann, *When Should Physicians Be Liable for Innovation?*, 36 *Cardozo L. Rev.* 913 (2015).

²¹ See Miller & Miller, *Legal and Regulatory Issues Related to the Use of Clinical Software in Health Care Delivery*, at 424.

²² *Ibid.*, at 433; see also Randolph A. Miller, Kenneth F. Schaffner, & Alan Meisel, *Ethical and Legal Issues Related to the Use of Computer Programs in Clinical Medicine*, 127 *Ann. Intern. Med.* 842 (1985).

²³ A physician might choose an algorithm based on the recommendations of a professional society, the scholarly literature, or the choices of his or her hospital, as discussed below. The mechanics of approving, validating, ranking, and selecting black-box algorithms are outside the scope of this chapter. For an exploration of these issues, see Roger A. Ford & W. Nicholson Price II, *Privacy and Accountability in Black-Box Medicine*, 22 *Mich. Telecomm. Tech. L. Rev.* 1 (2016); W. Nicholson Price II, *Regulating Black-Box Medicine*, 116 *Mich. L. Rev.* 421 (2017).

recommends or not. In some instances, this is unlikely to make a difference: where black-box algorithms suggest an increased risk to be managed with closer monitoring, the provider's actions take over from there. But where black-box algorithms suggest taking an unrelated drug based on previously unknown secondary effects or changing a drug's dosage or schedule without conforming to existing medical knowledge, the resulting care is fundamentally different from care not directed by black-box algorithms. Imposing the same standard of negligence would make little sense.

So what should medical malpractice law require of providers? Ordinarily, medical malpractice law responds to medical practices developed by providers (although the law may indirectly shape those practices). As long as a set of reasonable providers follows a certain practice, that practice will not typically lead to liability.

But when a practice is too innovative to have many adherents, it runs significant liability risk. Such is the situation of black-box medicine. So right now, in this period of early development, legislative action and practice guidelines set by professional organizations could be particularly influential.²⁴ These types of interventions could incorporate different levels of skepticism associated with the severity of different interventions.²⁵ For minimal-risk interventions, such as otherwise unindicated testing, increased monitoring, or taking widely used low-side-effect drugs such as aspirin, the standard of care might require no particular inquiry of the recommendations of a black-box algorithm. For riskier interventions, such as taking higher doses of a powerful drug or avoiding such a course when otherwise suggested, providers might require some validation before relying on a black-box algorithm. The form of validation would differ from traditional evidence because it would likely need to be based on procedural checks or independent computation by third parties, not clinical trials.²⁶ Professional societies, the FDA, and other intermediaries might serve a role in that validation to make it more feasible for an individual physician to reasonably check an algorithm's quality.²⁷ Such a process suggests the exercise of

²⁴ See Michelle M. Mello, *Of Swords and Shields: The Role of Clinical Practice Guidelines in Medical Malpractice Litigation*, 149 *U. Penn. L. Rev.* 645 (2001).

²⁵ This structure parallels the risk-based framework the FDA applies to preapproval for medical devices; low-risk devices are subject only to agency notification, high-risk devices to rigorous clinical trials, and intermediate-risk devices to an intermediate standard of evidence. See, e.g., Aaron V. Kaplan et al., *Medical Device Development from Prototype to Regulatory Approval*, 109 *Circulation* 3068 (2004). Ex ante classification raises its own challenges, including who would classify the devices. See, e.g., "Software as a Medical Device": Possible Framework for Risk Categorization and Corresponding Considerations, *International Medical Device Regulators Forum* 13–15 (September 18, 2014), available at www.imdrf.org/docs/imdrf/final/technical/imdrf-tech-140918-samd-framework-risk-categorization-141013.pdf [<https://perma.cc/GK54-UGVB>] (describing potential risk classification of software). Ex post classification, whether tiered or not, raises risks of chilling physician adoption.

²⁶ See Price, *Black-Box Medicine*, at 440–42 (discussing computational validation of black-box medicine).

²⁷ See Price, *Regulating Black-Box Medicine* (discussing potential roles for the FDA and collaborating healthcare actors in validating black-box algorithms).

intervening procedural judgment, even if substantive judgment is impossible. In somewhat analogous situations, courts have been willing to treat the use of an outdated information source as evidence of substandard care.²⁸ For the riskiest and most counterintuitive interventions – for instance, prescribing high doses of thalidomide to a pregnant woman – it is possible that under current standards, no black-box verification could be strong enough to overcome the presumption of harm under a reasonable standard of care. Such validation and reliability concerns could also be incorporated into potential liability for choosing a poor-quality algorithm against, for instance, the recommendations of a professional society.

Using a risk-based approach to evaluate the recommendations of black-box medicine brings with it challenges of implementation, overcaution, and undercompensation. First, a risk-based approach might be hard to implement efficiently, requiring expert testimony about the appropriate level of procedural assurance required for a particular intervention's risk. This real concern is, unfortunately, an unavoidable aspect of defining negligence based on a practitioner-based standard of care. Demonstrating adequate procedural care does not on its face appear to raise substantially greater evidentiary challenges than substantive standards for non-black-box care.

Second, providers might be too cautious, avoiding beneficial interventions out of concern for potential liability. This concern is true about new medical interventions in general, but to the extent that opaque algorithms seem to have a higher risk of causing harm, it may be more salient. Notably, however, some of the limited work done on diagnostic aids suggests that juries are less likely to punish doctors who act in accordance with such aids²⁹; similar patterns might occur for black-box medicine, easing adoption. Eventually, liability concerns might actually help drive adoption as black-box medicine graduates from its current status of untested innovation, becoming sufficiently prevalent to get the protection of the ordinary medical malpractice deference rules or even sufficiently dominant that *nonadherence* is deemed unreasonable. Like other medical advances, once black-box medicine becomes a more routine and better-accepted tool for medical care, providers not incorporating it could be liable for negligence if patients sustain reasonably avoidable injuries.³⁰

Third and finally, a risk-based standard of care might undercompensate patients if observing failures to follow the standard of care were particularly difficult. In medical malpractice in general, identifying errors and demonstrating causation raise

²⁸ See, e.g., *Boswer v. Craig Ranch Emergency Hosp.*, No. 05-14-00501-CV, 2015 WL 3946371, at *4–5 (Tex. App. Jun. 29, 2015) (finding that a hospital's use of an outdated textbook to establish its policy could imply a "failure to promulgate policies and procedures regarding the standard of care").

²⁹ Hal R. Arks et al., The Influence of a Physician's Use of a Diagnostic Decision Aid on the Malpractice Verdicts of Mock Jurors, 28 *Med. Decision Making* 201, 204–5 (2008).

³⁰ See Miller & Miller, Legal and Regulatory Issues Related to the Use of Clinical Software in Health Care Delivery, at 346 (making a similar argument in the context of clinical decision support software).

substantial challenges. Observing failures to evaluate the riskiness of a black-box algorithm's recommendation might be similarly challenging. However, limited observability and compensation are a long-term reality of medical malpractice in any context, not just black-box medicine.

In sum, liability for providers using black-box medicine is both familiar and quite novel, like black-box medicine itself. Providers will typically be held to the level of care of other comparable providers, which will develop over time as black-box medicine enters care. However, choices might best be shaped – judicially, legislatively, or professionally – to rely on risk-based procedural validation as the touchstone for practice and liability because the underlying physiologic cause of injury will be both unavoidably opaque and unevaluable by the expertise of the provider.

20.3.2 Liability of Healthcare Enterprises

In addition to providers, larger healthcare enterprises such as hospitals and clinics owe a duty of care to patients, especially in modern medical settings involving coordination of complex care. Hospitals may be vicariously liable for the negligence of healthcare providers who are actual or apparent agents of the enterprise but may also have duties directly to patients that are especially relevant for black-box medicine. In particular, hospitals have been held to have a duty to provide adequate facilities for patient care, including well-functioning equipment necessary for adequate care.³¹ In some jurisdictions, hospitals have other direct duties to patients, including a duty to coordinate care and sometimes a nondelegable duty to actually provide care for patients.³²

Under these theories, hospitals could be liable for negligently choosing, implementing, and using black-box medical systems. As with liability for providers, hospitals are typically held to the standard of a reasonable hospital; also as with providers, this means that a hospital's responsibilities can evolve as industry custom does. But policymakers could try to move hospitals' standard of care for implementing black-box algorithms toward one that would involve procedural tools to make sure that algorithms are well validated and competently developed before implementation.

Although hospitals are typically not liable for defects in the products they provide and/or sell, they may have a duty to nonnegligently evaluate the quality of those products and may be liable for failures of products that they fail to evaluate.³³

³¹ See *Washington v. Washington Hospital Center*, 579 A.2d 177 (D.C. Ct. App. 1990) (finding that a hospital could be directly liable for failure to provide carbon dioxide monitors in operating rooms).

³² See *Thompson v. Nason Hospital*, 427 Pa. 330 (Penn. 1991).

³³ See *Parker v. St. Vincent Hospital*, 122 N.M. 39, (N.M. Ct. App. 1996) (finding that a hospital is not strictly liable for defects in a physician-selected implant but acknowledging potential liability for negligently failing to examine the safety of the implant before provision).

Thus hospitals might reasonably be held liable for failing to ensure that the algorithms they make available to providers and patients are, as a whole, high quality and safe. Because substantive validation may be impossible in many cases – given the opaque nature of black-box medicine – procedural validation could be required instead. Parallels could be drawn to a more familiar responsibility of hospitals: their requirement to adequately credential the physicians who work in them to ensure that patients are seen by high-quality, well-trained doctors.³⁴ While a hospital cannot ensure that each decision of its doctors is correct, it can ensure that the doctors it brings through its doors are reasonably proficient. Applying a similar duty to black-box medicine would recognize the inherent opacity of the technology while leaving some responsibility on hospitals to take care in selection and implementation.³⁵

20.4 RECOMMENDATIONS AND CONCLUSIONS

The development and adoption of black-box medicine are complex processes that will naturally evolve over time. But black-box medicine is coming fast. Knowing the risks of legal liability under current law is important, but perhaps more important is the opportunity to shape that liability as black-box medicine becomes increasingly powerful and prevalent. Liability for the developers of black-box medicine is important and complicated and demands careful study to set incentives for development and validation. But providers and healthcare facilities will be the crucial frontline of black-box medicine, choosing and implementing algorithms and interfacing directly with patients. Correct liability rules for providers and facilities could help to ensure that they will exercise due care in that task; they could also direct compensation for patients injured by the absence of such care.

Setting the right standards of liability for providers and facilities may require more than just waiting for the professions to eventually evolve the proper standard of care. To reiterate, the baseline grounding of the standard of care in customary practice privileges hewing to tradition and may therefore slow the adoption of new technologies. Deliberate standard setting could provide a different path. This is not to suggest that deliberate standard setting is especially *likely* to occur; the shield of deference to customary practice has deep roots. But it remains worth considering whether medical malpractice could have a role in the responsible adoption of black-box medical algorithms and what that might look like.

Standards must tread a careful middle ground between two extremes. On the one hand, overly lax liability rules could result in potentially haphazard application of new algorithms that may be insufficiently validated. While this might speed

³⁴ See Torin A. Dorros & T. Howard Stone, Implications of Negligent Selection and Retention of Physicians in the Age of ERISA, 21 *Am. J.L. Med.* 383, 396–97 (1995).

³⁵ Hospitals could also potentially be involved in developing their own black-box algorithms, as custodians of substantial amounts of patient data. In such circumstances, they would potentially incur liability as algorithm developers, a topic outside the scope of this chapter.

adoption of black-box medicine, it could also insufficiently protect patients and would create fewer incentives for algorithm developers to demonstrate the quality of their products. On the other hand, overly stringent liability rules, such as strict liability for providers or facilities for injuries or requirements that they comprehend the inherently incomprehensible mechanisms of black-box algorithms, would stymie adoption and the benefits that black-box medicine could bring.

A potentially workable middle ground would require providers and facilities to exercise due care in procedurally evaluating and implementing black-box algorithms. Providers and facilities should evaluate black-box algorithms for hallmarks of careful development, including independent validation of algorithmic results and the qualifications of the developers. Facilities are best suited to evaluate algorithms at the point of implementation and should ensure that algorithms – as a whole – are high quality according to measurable characteristics. Providers are able to measure the risk associated with a particular intervention and should accordingly measure the level of validation and confidence against the risks entailed. Such a duty would help facilitate the development of independent private or regulatory mechanisms for developing and providing that type of validation.³⁶ Finally, these duties or individual determinations should not remain static; black-box medicine will develop and evolve rapidly, and the appropriate role for providers and healthcare facilities should evolve with it.

³⁶ See Price, *Black-Box Medicine*, at 457–62.

