

Professional Self-Regulation in Medicine

Will the Rise of Intelligent Tools Mean the End of Peer Review?

Anthony P. Weiss and Barak D. Richman

18.1 INTRODUCTION

Medicine has a longstanding compact with society to prioritize the needs of patients above other aims, a concept as old as the Hippocratic Oath. As an extension of this professional commitment, physicians have been granted the discretion and authority to police their own – to establish professional standards of conduct and to enforce those standards in the practice of medicine against fellow professionals.

Peer review is the culmination of this social compact. Peer review is a decentralized process in which the formal medical staff structures across the country assess adverse medical events, determine whether a colleague's conduct fell short of a professional standard of care, and, when necessary, discipline errant physicians. It is led by physicians, and though the objective is to learn from all errors and improve care throughout the delivery system, its primary focus is on physician conduct. It is so central to the modern practice of medicine that it has shaped the organization of hospitals and embodies the central values of being a medical professional.

However, the contours of peer review are products of – and are thus sensitive to – medical technology. Physician self-regulation and its prevailing current structures are largely driven by the realities of knowledge. Historically, it was held that only physicians had the expertise to distinguish proper from improper care, and thus should be the sole authorities to assess the competence of any individual member of the profession.¹ But modern medicine is now guided by sources of knowledge that no longer lie in the sole possession of physicians. Shifts toward team-based care models rely on the expertise of nonphysicians and skills that lie beyond physician capabilities. Electronic medical records are now the primary repositories of patient medical information and are becoming more capable than physicians at detecting and correcting errors or deviations in practice. And digital technologies increasingly have the capacity to synthesize data to generate diagnoses and medical recommendations, which compare favorably to human experts. These new technologies

¹ Edward H. Livingston & John D. Harwell, *Peer Review*, 182 *Am. J. Surgery* 103 (2001).

and systems are challenging the supremacy of physician expertise in medicine, and consequently are eroding the underlying justifications for peer review.

And true disruption might be at hand. Artificial intelligence (AI) and machine learning (ML) algorithms are slowly enabling computer-driven medicine to perform the core functions of peer review: identifying medical errors, learning from adverse outcomes, and instituting reforms.

Even though it is widely known that new capabilities and skills from nonphysicians are playing an increasingly significant role in the practice of medicine, there has been little thought to whether these technological disruptions will necessitate broader changes in the organization and delivery of medicine. This chapter begins that inquiry by exploring the implications that AI and digital technologies have for peer review. We suggest that these technologies do much more than supplement the medical staff's ability to evaluate and improve medical practice. They have the capacity to disrupt traditional centers of authority that underlie peer review and may thereby force a reorganization of medicine and a recalibration of health care regulation.

18.2 PEER REVIEW EXPLAINED

The governance relationship between so-called learned professions and the rest of society has been likened to a social contract. In exchange for professionals investing in valuable expertise, inculcating a commitment to, and offering guidance to state officials, the state defers to professional expertise in both substance and practice.² Professionals thus are tasked by the state to define standards of conduct and to discipline themselves accordingly. Economists typically describe this arrangement as a product of information asymmetries – that lay people, even elected officials, have inadequate knowledge to scrutinize the conduct of scientific experts³ – and sociologists observe that this social arrangement reserves for professionals a privileged status and fierce autonomy that few other laborers enjoy.

Physicians historically have been the archetype of the learned professions, enjoying greater autonomy and self-governance privileges than that afforded to other professional societies, and the expanse of physician self-governance is reflected across numerous public and private mechanisms. States authorize medical societies to establish state licensure regimes and malpractice standards, thereby allowing the profession to define qualifications and minimal standards. States also provide their plenary powers to both licensure boards (thereby prohibiting nonphysicians from engaging in “the practice of medicine”) and courts (thereby disciplining those who fail to meet medical board standards). The incorporation of professional standards

² William J. Goode, *Community Within a Community: The Professions*, 22 *Am. Sociological Rev.* 194 (1957).

³ Kenneth J. Arrow, *Uncertainty and the Welfare Economics of Medical Care*, 53 *Am. Econ. Rev.* 141 (Dec. 1963).

into the law, and the use of the state's police power to enforce those standards, allow medical professionals to maintain a robust and self-sustaining system of self-regulation.

But it might be said that the cornerstone of physician self-governance – and the pinnacle of collaboration between medical professionals and the state – lies in peer review. Both medical boards and state authority defer to institutional peer reviews, and numerous laws cloak the peer review process both to secure its sanctity and reinforce its authority.

This legal arrangement has been described as a unique reflection of successive “devolutions” of medical authority, from federal government to state government, from state government to physician-led state licensing boards, and from these boards to the local medical staff structures within individual hospitals.⁴ In many respects, it is a historical compromise reflecting the need to protect the public from deviant medical practice, while recognizing that medicine is an imperfect science with significant regional variation. It is also a means to achieve the end of public safety while avoiding nonphysician (government) intrusion in the doctor-patient relationship, a politically challenging issue.

While voluntary and informal review of physician practice by peers may occur in a number of settings, formal peer review as considered in this chapter is largely a hospital-based function, conducted to meet accreditation requirements specified by the Joint Commission (TJC). Despite some guidance put forward by TJC for medical staff practice, peer review remains highly variable in its structure and process, with relatively little standardization across the country. Peer review is typically conducted by a multi-disciplinary committee of physicians, and supported by hospital personnel, including nurses, safety experts, and attorneys. The peer review committee reports to a physician-led medical executive committee (MEC), which has local responsibility for the oversight of medical practice at a hospital. The committee may serve a number of roles, including initial review of physician qualifications for inclusion on the medical staff. But the committee primarily reviews adverse patient outcomes for evidence of physician negligence or incompetence, typically comparing the actions of the responsible physician to a community standard, looking for gross deviation from usual practice. In some cases, the committee will undertake an exercise, known as root cause analysis, meant to uncover the specific factors which resulted in the adverse patient outcome. The work product of the committee may range from recommendations to the hospital's governing body (via the MEC) to rescind physician privileges, to dismissal of any concern related to the adverse incident.

Peer review is much more than a political compromise or a social compact of convenience. There are benefits that justify empowering local inquiry, by and for

⁴ Theodore W. Ruger, *Plural Constitutionalism and the Pathologies of American Healthcare*, 120 *Yale L. J.* 347 (2011).

professionals, in assessing medical error. First, because of the complexity of medicine, inquiries into errors are best done by physicians most familiar with the context in which the error took place. In addition to the natural complexity of medicine, which justifies deferring to the expertise of practitioners over the judgment of regulators, there is further reason to defer to physicians who are best acquainted with the surrounding environment, facilities, and personnel in which scrutinized care took place. For this reason, many applications of medical malpractice law recognize regional variations in the practice of medicine and apply locally determined standards of care.

Second, and more significant, local peer review is designed to enhance the benefits of scrutinizing adverse outcomes. Since the ultimate objective is to learn from mistakes and improve the quality of care, the priority of any review process is to acquire accurate information, which plausibly is best done between colleagues within a cloak of trust, reciprocity, and collective learning. The tacit nature of information, the sensitivity of disclosing information related to potential errors, and the formal and informal support structures that embed the disclosure of this information all counsel towards providing discretion and authority to local medical boards.

These purported benefits of peer review are also reflected in the law, as many doctrines explicitly protect the peer review process. For example, any materials generated during the peer review process, including admissions of error, are shielded from discovery in any subsequent malpractice suit. Moreover, if internal documents or materials related to a medical error are not part of a peer review, they then do become subject to discovery. Perhaps most important, the association of peer review with high-quality medicine is enshrined in state licensure law and accreditation standards. A hospital needs to institute a peer review process, administered by a physician-led medical board, to be permitted to care for patients and receive Medicare funds. Peer review is not just thought to be important for maintaining high-quality medical care, it is deemed to be an essential feature of quality assurance for medical facilities.

18.3 CRITICISMS AND SHORTCOMINGS OF PEER REVIEW

Peer review is not without its detractors, however, both within and outside the field of medicine. There are three broad categories of criticism: 1) It promotes a singular societal aim (safety) over other health care ends of importance to the public (like innovation, cost, and access); 2) Like all human processes, it is liable to bias, self-preservation, and abuse; 3) It is ineffective in achieving its primary aim of promoting safe care.

Perhaps owing to the ancient dictum, *primum non nocere*, peer review almost solely focuses on safety of health care to the exclusion of other valuable aims. While nominally, peer review purports to drive learning and improvement, reviews of its

success in this area have been disappointing.⁵ This is, in part, inherent to the nature of the process – the widespread dissemination of lessons learned (needed for innovation) stands in opposition to the privacy needed to maintain a trusting arrangement for local review. Furthermore, peer review has limited impact on other aspects of the “iron triangle” of health care, including cost containment and access.⁶ Given the increasing importance of these parameters in promoting both affordability and public health, the local self-regulatory nature of peer review may be insufficient.

The process of peer review itself has been criticized as being biased,⁷ with concerns of both underreporting and overreporting of poor physician practice. Peer review requires physician colleagues to assess each other’s work product (patient care) in a reciprocal manner. Despite best intentions, policing one’s peers can prove difficult, with a variety of conflicts of interest and social connections serving as barriers to even the most well-meaning and thoughtful peer review committees. Underreporting is an unsurprising result of a process in which unpaid physician committee members are asked to make potential career-ending calls on classmates, friends, and patient-referral sources.⁸ On the other hand, peer review has also been used in an overaggressive manner, as one group of physicians attempts to drive out a competing member of another group. Indeed, this anticompetitive practice (sometimes known as “sham” peer review) was the basis for a \$2.2 million settlement for the plaintiff in the 1986 antitrust lawsuit brought by Dr. Timothy Patrick against Dr. William Burget and the Astoria Clinic. This settlement sent a strong message to curtail the practice of sham peer review, but further dampened physician interest in participating in peer review, for fear of legal downsides. This led to the passage of the Health Care Quality Improvement Act (HCQIA), which attempted to rectify this by providing legal immunity to those physicians who participate in peer review in good faith.

The efficacy of peer review in evaluating the adequacy of medical care has also come under criticism. Despite the push toward evidence-based medicine, large swaths of care remain outside the evidence-based map, leaving ample room for variation in practice. In addition, many medical errors are the result of cognitive biases that are challenging to elucidate and difficult to mitigate.⁹ Indeed, this point highlights the real challenge of the peer review process. To work well, peer review requires: self-reflection and insight, accurate recollection, capacity for thoughtful

⁵ Mohammad Farhad Peerally et al., *The Problem with Root Cause Analysis*, 26 *BMJ Quality & Safety* 417 (2017); Albert W. Wu et al., *Effectiveness and Efficiency of Root Cause Analysis in Medicine*, 299 *JAMA* 685 (2008).

⁶ William L. Kissick, *Medicine’s Dilemmas: Infinite Needs versus Finite Resources* (Yale Univ. Press ed., 1994).

⁷ Dinesh Vyas & Ahmed E. Hozain, *Clinical Peer Review in the United States: History, Legal Development and Subsequent Abuse*, 20 *World J. Gastroenterology* 6357 (2014).

⁸ George E. Newton, *Maintaining the Balance: Reconciling the Social and Judicial Costs of Medical Peer Review Protection*, 723 *Ala. L. Rev.* 723 (2001).

⁹ Jerome P. Kassirer, *Diagnostic Reasoning*, 110 *Annals of Internal Med.* 893 (1989); Geoffrey R. Norman & Kevin W. Eva, *Diagnostic Error and Clinical Reasoning*, 44 *Med. Educ.* 94 (2010).

interrogation, cooperativity, and clear communication within a trustworthy circle of colleagues. Deficiencies in any aspect of this set of conditions may limit the adequacy of the evaluation and diminish its impact of provision of safe care.

18.4 THE PROMISE AND CHALLENGES OF ARTIFICIAL INTELLIGENCE (AI)

Against this backdrop, a new era of intelligent tools, including advanced decision support with artificial intelligence, pose an attractive alternative to the peer review process. Artificial intelligence (AI) and machine learning (ML) are predictive modeling approaches which combine data in unique ways, via algorithms, to identify optimal solutions.¹⁰ With improvements in programming, massive increases in computing power, and digitization of nearly all aspects of health care delivery, AI/ML has achieved a series of impressive results, now surpassing human capabilities in many domains, including image recognition. The latest iteration of this type of technology, known as deep learning, shows the capacity for perpetual enhancement in accuracy, self-modifying the algorithms it uses based on the relation of outputs to inputted data.

There is significant enthusiasm for the use of AI/ML within health care. The myriad of potential applications currently fall largely into two domains: image analysis and clinical decision support. The former application has multiple use-cases within health care, from visual analysis of radiographic images,¹¹ to pathologic tissue diagnosis,¹² to interpretation of retinal scans.¹³ These solutions are focused primarily on the diagnostic aspect of health care – distinguishing normal from abnormal and applying the taxonomy of human pathology to abnormal findings.

AI is also being explored to assist with the cognitive decision making so critical to the work of many physicians. Taking disparate bits of information from multiple sources (e.g., patient history, physical exam, diagnostic tests) and determining a diagnosis, prognosis, and therapeutic plan. The range of potential decision-support applications is as broad as the expanse of all of medical practice, and includes recent examples, including early warning prediction of intraoperative hypotension,¹⁴ mortality after heart failure,¹⁵ and suicidal behavior after hospital

¹⁰ Michael D. Howell & Jennifer P. Stevens, Chapter 17: Predictive Modelling 3.0: Machine Learning, in *Understanding Healthcare Delivery Science* 341 (McGraw Hill eds., 2020).

¹¹ Liu et al., (2018).

¹² Yahui Jiang et al., Emerging Role of Deep Learning-Based Artificial Intelligence in Tumor Pathology, 40 *Cancer Comm.* 154 (2020).

¹³ Ting et al., (2017).

¹⁴ Marije Wijnberge et al., Effect of a Machine Learning-Derived Early Warning System for Intraoperative Hypotension vs. Standard Care on Depth and Duration of Intraoperative Hypotension During Elective Noncardiac Surgery, 323 *JAMA* 1052 (2020).

¹⁵ Joon-myoungh Kwon et al., Artificial Intelligence Algorithm for Predicting Mortality of Patients with Acute Heart Failure, *PLOS One* (2019), <https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0219302>.

discharge.¹⁶ Increasingly, applications are moving beyond diagnosis or prediction, to autonomously enacting therapeutic decisions. For example, closed-loop neurostimulatory devices are being designed to both identify epileptic signals in the brain and immediately treat them via electrical stimulation.¹⁷ Chatbots powered by AI are also now being deployed, to field patient chief complaints and triage them, in some cases recommending basic treatments.¹⁸

Despite this potential, AI and other digital diagnostics are largely kept out of the peer review process. One reason for this is the traditional structure of peer review. Peer review is focused on physician conduct, more than systems-centered care, and is by design reactive to significant adverse outcomes. Of the many errant and potentially harmful actions that take place regularly in a large and complex hospital, very few attract the attention of peer review scrutiny. Those that do are examined by fellow physicians through traditional means of physician judgment, without the deep computational support that is more commonplace in advanced analytics.

Another reason is that digital monitoring, the foundation on which AI technology is built, has developed around a quality assurance infrastructure that is largely parallel to physician-centric peer review. Modern hospitals institute software that monitors the administration of patient care that contain safeguards against predictable errors. For example, when medication orders are entered into most hospital computer systems, software reviews these orders for obvious deviations or harmful interactions. This electronic monitoring occurs downstream of most physician conduct and has the capacity to scrutinize daily conduct that usually is outside the domain of peer review.

The real-time digital quality assurance alerts and the retrospective human-led peer review represent largely parallel solutions to improving patient safety, both required for Joint Commission accreditation. Even now, digital monitoring exhibits capabilities that reach much deeper and more objectively into the practice of medicine. Whereas peer review examines perhaps a few dozen adverse outcomes per month, order alert software monitors thousands of inputs daily. Moreover, these software systems can collate and analyze these events to identify common sources of error.

The superimposition of AI offers enormous opportunity for deeper analysis and more sophisticated monitoring that could improve quality assurance. First, AI and data analytics could do more than provide simple alerts to known errors. Deep learning algorithms could examine population health data and tailor recommendations to an individual patient's needs, identify improvements to accepted medical protocols, or anticipate systemic sources of provider error. Second, AI algorithms

¹⁶ Trehani M. Fonseca et al., *The Utility of Artificial Intelligence in Suicide Risk Prediction and the Management of Suicidal Behaviors*, 53 *Austl. & N.Z. J. Psychiatry* 954 (2019).

¹⁷ Urvish Patel et al., *Artificial Intelligence as an Emerging Technology in the Current Care of Neurological Disorders*, *J. of Neurology* (2019), <https://www.researchgate.net/publication/335406868>.

¹⁸ Mary Bates, *Health Care Chatbots are Here to Help*, 10 *IEEE Pulse* 12 (2019).

could institute real-time guidance to providers at the point of care, both anticipating moments of likely errors and interjecting with medical treatments that data analysis determines is superior to a human's judgment. Perhaps ultimately, such deep learning algorithms could remove the role of human judgment and human implementation altogether. AI could, on its own, identify and implement a tailored treatment, assess preliminary results, change course to alternatives if necessary, and integrate data from a patient's progress to a broader body of knowledge. The role of humans would be to peripherally monitor the AI's implementation, perhaps play some supervisory role, and ensure that the algorithms have the data and resources they need to operate effectively.

18.5 PEER REVIEW AND ARTIFICIAL INTELLIGENCE

The very aspects of AI that make it an exciting adjunct or alternative within health care are those that pose the greatest challenges to peer review. As much as software and digital capabilities can enhance and improve the quality assurance mechanisms that culminate in peer review, those same technologies might also prove to undermine peer review.

The self-learning, black-box nature of AI makes it difficult to interrogate to human knowledge, even if one has expertise in computer technology. As AI puts together information in novel and unique ways, the resulting algorithms may in fact be more accurate. But this reliance on alternative models makes it impenetrable to the physician review process which is framed by longstanding medical Western Medical tradition of the factors to be incorporated in diagnosis, prognosis, and treatment. Physicians simply do not have an understanding of the algorithms used to generate these decisions; decisions that may be swayed by information a physician may not have considered *prima facie* relevant. Whereas peer review is designed to be by physicians and for physicians, the skillset required to monitor AI-guided medicine would more likely involve computer scientists and software engineers.

In addition, since AI may identify patterns in existing data that will guide treatments where there is current clinical equipoise, the best course of treatment may no longer be determined solely on the basis of published literature. A peer review committee determining if an error occurred will have less information than the machine that directed treatment. Since AI will constantly be at the forefront of medical practice, any human oversight of AI and deep learning algorithms will be, by definition, deficient.

Therefore, however helpful AI might be, it is critical to recognize it contains qualities that are starkly different from human physicians (to state the obvious, machines are different from humans). In some ways, AI is like an uncooperative physician, one whose logic is impenetrable, and speaks another language entirely. To the degree that this new member of the medical staff is guiding care which may

in some cases be harmful, it will fall upon local peer review processes to mitigate their impact on other patients. It is uncertain how the traditional approach to reciprocal peer review can effectively manage this task.

18.6 IMPLEMENTING AI INTO HOSPITAL QUALITY ASSURANCE

In short, while the hype and promise surrounding AI is clear, it is also clear that health care AI will not readily fit into existing hospital governance. The oversight of AI will depend heavily on the nature of its use, whether AI is a tool used by providers or whether it becomes a provider in and of itself.¹⁹

The introduction of AI processes could be maintained and modified to become safely incorporated into the current delivery system. For example, local medical staff policies could be implemented to limit the use of AI, by precluding its use in making immediate patient-facing decisions or requiring physician signoff on all diagnostic or therapeutic decisions. In so doing, human physicians would retain their current authority over medical care and buffer patients from harm related to aberrant AI decisions. The medical staff could also monitor the novel use of AI for a period of time until deemed safe, akin to the Joint Commission-mandated initial focused professional peer evaluations (FPPE) used to evaluate a new member of the medical staff. Furthermore, the medical staff could invite assistance from outside technical experts, such as computer programmers, to assist in interrogating the algorithm used in association with a medical error. This would be akin to the common practice of asking a biomedical engineer to inspect a machine (like an intravenous pump) after a serious event, to help distinguish operator error from device failure. These are not mutually exclusive – some combination of these solutions (and more) could be deployed while still maintaining local regulatory control through the classical peer review process.

A scenario that we think is more likely is that AI applications will chip away at traditional physician-centered medical care and even physician self-governance. The entry of AI, eventually, is likely to surpass the capabilities of local-level provider control, and the physician's primacy in the delivery of health care will wane. One might say that this trend has already begun with greater corporatization in health care and the expansion of nonphysician health care providers. Uptake of AI adds a degree of technical complexity that likely lies outside the physician's expertise, and the logic underlying physician self-governance will collapse. More immediately, AI would spell the end for peer review as a practice. The complexities of an autonomous intelligent health care machine may simply be too challenging for a quaint process from the late-nineteenth century.

¹⁹ Anthony Weiss et al., *How AI Will Change the Regulation and Organization of Medicine*, Health Affairs Blog (May 3, 2020), <https://www.healthaffairs.org/doi/10.1377/hblog20200430.833066/full/>.

A related challenge will confront the traditional powers that now govern peer review. What is now a physician-dominated process – often exclusively so – must also be able to handle AI-based approaches to health care. A hospital's most critical personnel decisions, including the allocation of admission privileges, remains (by tradition, law, and otherwise) under the control of the physician-led hospital medical staff. If AI and the peer review process become at odds, the continued use of either one will be determined ultimately by the medical staff. The medical staff would have to prove it was capable of managing the growing significance of AI, or it would have to narrow its authority over hospital operations.

This in turn begs the broader question of whether AI will force a change in the organization of the hospital. Peer review is central to a hospital's governance structure precisely because it is a window into the hospital's authority structure. If the locus of control changes in reviewing patient care and scrutinizing provider conduct, the locus of authority will also change in health care delivery. It certainly is hard to imagine that the superiority of AI (if, indeed it earns that superiority) over human-governed quality assurance is compatible with an American hospital's traditional governance structure. If AI changes peer review, it stands to reason that it will require changes to the hospital as well. The power to define standards and determine when those standards are not met – and by whom – is the power that controls the care delivered within the hospital.

A reallocation and reorganization of power in the hospital would certainly signal broader changes in health care delivery. A hospital system would have little trouble leveraging its AI capabilities to reach a broader scale of patients, and hospital administrators untethered to physician limits might reconfigure delivery. Although removing the human element from medicine will introduce meaningful drawbacks – and presumably, AI and digital services will never duplicate human wisdom of physicians – one could imagine how an operations-centric, AI-centric delivery system could advance the aims of population health. An AI-focused health system might monitor and sustain the health of a large population better than one that services patients with physician visits and hospital-based procedures.

18.7 LEGAL AND POLICY IMPLICATIONS

Just as AI might change the operation and governance of the hospital, it will also likely demand changes to health care regulation. The nation's current angst over increased health insecurity, political demands for distributive justice, and the fiscal strain of health care expenditures already fuel public demands for delivery reform. There may be an acute need to develop policies that thoughtfully usher in AI medicine while assuring a worrying public.

Even as AI grows to replace human roles, it will be treated as a device or product and thus will receive the same legal treatment as other technologies currently in use. For example, AI tools are likely to be subject to product liability law, unlike

malpractice law that governs physicians. Moreover, even when physicians maintain responsibility for the care provided by AI tools, the law will have difficulty navigating between the two tort regimes.²⁰ This might also mean a greater regulatory focus on the corporations providing health care, rather than on individual providers or specific tools, and thus might stimulate demand for enterprise liability regimes that will feature employed, rather than independent, physicians.

The shift to product liability or enterprise liability, and an erosion of individual professional liability, could be further fueled by a need to scrutinize the core of malpractice law. If malpractice liability law evolved under a professionalism paradigm – one that nurtured and protected peer review and deferred to professional sources of authority – then the supremacy of AI and the inadequacies of human knowledge will require a rethinking of medical malpractice law. Specifically, if expertise lies more within the domain of the computer engineer than the physician, and if general standards of practice succumb to deep learning algorithms, then the malpractice law's deference to professional standards would be inapposite. Changes would include altering Daubert rules that define expertise, discovery rules that determine what evidence is authoritative and what is not, and substantive rules of tortious negligence and the sources of knowledge from which they are derived.

There would also be a shift towards federal medical device regulation, governed under FDA and intellectual property laws, and away from local regimes that govern the practice of medicine. And because digital products will naturally disseminate in a national (and international) market, there will be diminished tolerance for localities with their own rules and quality standards. We should expect to see a continued loss of local control over medicine, both in law and in practice, and a corresponding shift from local to state and from state to federal oversight of health care delivery.

Licensure regimes might change as well. Instead of licensing boards scrutinizing which humans warrant credentials, the FDA would approve machines and algorithms, and AI certifications would emerge in a to-be-developed federal product approval process. This could both dilute the effect of state medical boards, perhaps the longest surviving institution in modern medicine, and usher in the rise of health care corporations with national reach. Health care systems, insurance companies, and technology companies are the most likely to deploy AI tools, and these providers would be responsible for the outcomes, including safety, cost, and access, much as Ford is responsible for the cars they sell nationwide.

And even if peer review continues, the specific laws surrounding it will be ripe for reform. For example, discovery rules include immunities that protect disclosure during the peer review process, under the logic that peer review requires that immunity. But if quality review relies on digital analytics, which in turn requires the exchange of data across hospital systems, then the logic of limiting discovery is

²⁰ W. Nicholson Price et al., *Potential Liability for Physicians Using Artificial Intelligence*, 322 JAMA 1765 (2019).

undermined. In fact, even without the presence of malpractice suits, AI would function best with vigorous reporting requirements, subject to appropriate privacy rules. The laws originally designed to nurture peer review will be redesigned to nurture alternative mechanisms to assure medical quality.

This might suggest a broader need to rethink health care regulation more generally. Regulators should be receptive to the possibility that health care regulation should accommodate the needs of AI, rather than the reverse, and deem the emergence of AI as an occasion for a broader reframing of medical regulation.²¹ If AI genuinely represents a potential to improve the quality of medical care, provide the digital products that can achieve population health, and avoid the shortcomings of physician self-governance, then policymakers and medical leaders should plot out the legal rules that would support the growth, improvement, and accessibility of new digital products.

At the same time, there will be significant hesitation to move both the delivery of health care and the regulation of health care away from a professional paradigm. Despite its imperfections, locally controlled, physician-led peer review is a process that has served an important role in ensuring patient safety for more than a century, adapting to innumerable changes in health care practice over that time.

In the end, these questions might be answered by political intuitions and popular perceptions. Peer review has remained a pillar of medical practice because it has succeeded in maintaining the trust of the public, and a move to algorithms could undermine that trust. But the integrity of peer review has been showing cracks of its own and may not continue to win the confidence of a digitally connected public. Perhaps the incorporation of AI tools will require a very different set of rules to maintain public trust.

²¹ Barak Richman, *Health Regulation for the Digital Age*, 379 *N. Eng. J. Med.* 1694 (Nov. 1, 2018).