

2022

## Distributed Governance of Medical AI

W. Nicholson Price II  
*University of Michigan School of Law*

---

### Recommended Citation

W. Nicholson Price, *Distributed Governance of Medical AI*, 25 SMU SCI. & TECH. L. REV. 3 (2022)  
<https://scholar.smu.edu/scitech/vol25/iss1/2>

This Article is brought to you for free and open access by the Law Journals at SMU Scholar. It has been accepted for inclusion in SMU Science and Technology Law Review by an authorized administrator of SMU Scholar. For more information, please visit <http://digitalrepository.smu.edu>.

# DISTRIBUTED GOVERNANCE OF MEDICAL AI

*W. Nicholson Price II\**

## ABSTRACT

Artificial intelligence (AI) has the potential to democratize expertise in medicine, bring expertise previously limited to specialists to a variety of health-care settings. But AI can easily falter, and making sure that AI works well across that variety of settings is a challenging task. Centralized governance, such as review by the Food and Drug Administration, can only do so much, since system performance will depend on the particular health-care setting and how the AI system is integrated into setting-specific clinical workflows. This Essay presents the need for distributed governance, where some oversight tasks are undertaken in localized settings. It points out the resource-based challenges of such governance and offers policy suggestions to ease the burden.

## I. INTRODUCTION

Artificial intelligence (AI) promises to bring substantial benefits to medicine. In addition to pushing the frontiers of what is humanly possible, like predicting kidney failure or sepsis before any human can notice, AI can democratize expertise beyond the circle of highly specialized practitioners, like letting generalists diagnose diabetic degeneration of the retina. But AI doesn't always work; it doesn't always work for everyone, and it doesn't always work in every context. AI is likely to behave differently in well-resourced hospitals where it is developed than in poorly resourced frontline health environments, where it might well make the biggest difference for patient care. To make the situation even more complicated, AI is unlikely to go through the centralized review and validation process that other medical technologies undergo, like drugs and most medical devices. Even if it did go through those centralized processes, ensuring high-quality performance across a wide variety of settings, including poorly resourced settings, is especially challenging for such centralized mechanisms. What are policymakers to do? This Essay argues that the diffusion of medical AI, with its many potential benefits, will require policy support for a process of distributed governance, where quality evaluation and oversight take place in the settings

---

DOI: <https://doi.org/10.25172/smustr.25.1.2>

\* Professor of Law, University of Michigan Law School. JD, Columbia Law School, 2011. PhD (Biology), Columbia University, 2010. This work was presented as part of the SMU Science & Technology Law Review's 2022 Symposium on AI & Medicine: The Emerging Legal and Ethical Frameworks for Artificial Intelligence in Medicine. I thank Mark Sendak for helpful discussions and comments on an earlier draft; Karandeep Singh and Yindalon Aphinyanaphongs for helpful discussions; and Phoebe Roque for exemplary research assistance.

---

of application—but with policy assistance in developing capacities and making that oversight more straightforward to undertake. Getting governance right will not be easy—it never is—but ignoring the issue is likely to leave benefits on the table and patients at risk.

## II. THE NEED FOR GOVERNANCE

Why does medical AI require governance? Put plainly, quality, safety, and efficacy are quite difficult to assess. Many health-care technologies are credence goods,<sup>1</sup> requiring either blind faith in quality (not a great plan) or some form of rigorous, systemic evaluation (a better plan!), because one-off evaluations in the moment don't cut it. AI is very much a credence good.<sup>2</sup> It is novel; it deals in probabilities rather than certainties; and it relies on algorithms that are typically quite opaque, whether because of secrecy, inherent technological limits, or merely their underlying complexity.<sup>3</sup> Like any technology involved in medical care, whether AI works is hard knowledge to come by.

Quality concerns are well founded; researchers have demonstrated deep flaws with AI systems, including some already in use.<sup>4</sup> Sometimes the systems just aren't useful, and sometimes they are actively harmful.<sup>5</sup> Health-care vendor Epic's AI-powered system, designed to predict the risk of sepsis and distributed and used in hospitals around the country, has turned out to be a very poor predictor of risk—perhaps because the algorithm used as a predictive variable whether a physician had *already* ordered antibiotics, a typical *response* to sepsis.<sup>6</sup> COVID-19 prediction algorithms, developed rapidly and heralded as a success story for quick AI innovation in a global pandemic,

- 
1. Uwe Dulleck, Rudolf Kerschbamer & Matthias Sutter, *The Economics of Credence Goods: An Experiment on the Role of Liability, Verifiability, Reputation, and Competition*, 101 AM. ECON. REV. 526, 526 (2011).
  2. W. Nicholson Price II, *Regulating Black-Box Medicine*, 116 MICH. L. REV. 421, 432 (2017).
  3. *See id.* at 432–37.
  4. *Id.* at 436.
  5. *Id.*
  6. Andrew Wong, Erkin Otles, John P. Donnelly, Andrew Krumm, Jeffrey McCullough, Olivia DeTroyer-Cooley, Justin Pestrue, Marie Phillips, Judy Konye, Carleen Penozza, Muhammad Ghous & Karandeep Singh, *External Validation of a Widely Implemented Proprietary Sepsis Prediction Model in Hospitalized Patients*, 181 J. AM. MED. INTER. MED. 1065, 1065 (2021); Casey Ross, *Epic's Sepsis Algorithm Is Going Off the Rails in the Real World. The Use of These Variables May Explain Why*, STAT + (Sept. 27, 2021), <https://www.statnews.com/2021/09/27/epic-sepsis-algorithm-antibiotics-model/> [<https://perma.cc/R92J-JZJW>].

---

have not performed very well at all and have made little difference.<sup>7</sup> A tool for analyzing X-rays for pneumonia flopped when it was tested in another hospital (it was relying on X-ray procedural clues rather than actual patient image traits).<sup>8</sup> And a tool used by an insurer to allocate care coordinators was shown to be strongly biased against Black patients because of a careless proxy decision made in development.<sup>9</sup> AI has a lot of promise—but adopters are right to be cautious, and real governance is required to make sure that the systems being considered actually work in general, work where they are used, and work for the patients and providers in that particular setting.

### III. THE LIMITS OF CENTRALIZED REGULATION

The default turn for regulation of new medical technologies is to the Food and Drug Administration (FDA). FDA regulates drugs and medical devices, and software—including AI software—is explicitly within the bounds of regulated “medical devices”<sup>10</sup> (or, at least, may be<sup>11</sup>), whether embedded in other devices (Software *in* a Medical Device or SiMD) or on its own (Software *as* a Medical Device or SaMD).<sup>12</sup> FDA regulates, or has the au-

- 
7. E.g., Will Douglas Heaven, *Hundreds of AI Tools Have Been Built to Catch Covid. None of Them Helped*, MIT TECH. REV. (July 30, 2021) <https://www.technologyreview.com/2021/07/30/1030329/machine-learning-ai-failed-covid-hospital-diagnosis-pandemic/> [<https://perma.cc/3Z8G-HD3M>]; Jordana Cepelewicz, *The Hard Lessons of Modeling the Coronavirus Pandemic*, QUANTA MAG. (Jan. 28, 2021) <https://www.quantamagazine.org/the-hard-lessons-of-modeling-the-coronavirus-pandemic-20210128/> [<https://perma.cc/9PAW-YXX7>].
  8. W. Nicholson Price II, Rachel E. Sachs & Rebecca S. Eisenberg, *New Innovation Models in Medical AI*, 99 WASH U. L. REV. 1121, 1171 (2022).
  9. Ziad Obermeyer, Brian Powers, Christine Vogeli & Sendhil Mullainathan, *Dissecting Racial Bias in an Algorithm Used to Manage the Health of Populations*, 366 SCI. 447 (2019); see also Jenna Wiens, W. Nicholson Price II & Michael W. Sjoding, *Diagnosing Bias in Data-Driven Algorithms for Healthcare*, 26 NATURE MED. 22, 25 (2020) (responding to Obermeyer et al.).
  10. Barbara Evans & Frank A. Pasquale, *Product Liability Suits for FDA-Regulated AI/ML Software*, in THE FUTURE OF MEDICAL DEVICE REGULATION: INNOVATION AND PROTECTION 22, 23 (I. Glenn Cohen, Timo Minssen, W. Nicholson Price II, Christopher Robertson & Carmel Shachar eds., 2022).
  11. *Id.*; see also Nathan Cortez, *Substantiating Big Data in Health Care*, 14 I/S: A.J. L. & POL’Y INFO. SOC’Y, 61, 72–81 (2017).
  12. See *Software as a Medical Device (SaMD): Key Definitions*, INT’L MED. DEVICE REG. FORUM (“IMDRF”) (Dec. 19, 2013), <http://www.imdrf.org/docs/imdrf/final/technical/imdrf-tech-131209-samd-key-definitions-140901.pdf> [<https://perma.cc/F57Y-PGJN>]; “*Software as a Medical Device*”: *Possible Framework for Risk Categorization and Corresponding Considerations*, IMDRF (Sept. 18, 2014), <http://www.imdrf.org/docs/imdrf/final/technical/imdrf-tech-140918-samd-framework-risk-categorization-141013.pdf> [<https://>

thority to regulate, software via the familiar Class I-II-III risk-based categorization of medical devices, with scrutiny titrated to risk.<sup>13</sup> This process is meant to ensure that medical devices are safe and effective before entering into commerce and then into practice.<sup>14</sup> This process, however, has severe limitations on how well it can govern the quality of the broad sweep of AI products developed and used today.<sup>15</sup>

#### A. What Does FDA See?

The first problem is one of coverage: many AI systems in development or already in use have not gone through FDA review of any kind, and many likely never will.<sup>16</sup> As I have written with Rachel Sachs and Rebecca Eisenberg, there is substantial user innovation in the space of medical AI, where health systems, hospitals, and insurers are developing AI systems for use within their own walls.<sup>17</sup> AI innovation is well within the capacity of many well-resourced actors in this space, in a way that the development of novel drugs or complex physical medical devices, for instance, may not be.<sup>18</sup> Academic medical centers have developed in-house predictors for the likelihood of sepsis,<sup>19</sup> hospital systems have developed AI systems to model patient flow (and, accordingly, resource allocation),<sup>20</sup> and insurers have developed

---

perma.cc/F6YH-SWGA]; *Global Approach to Software as a Medical Device*, FDA (Dec. 6, 2017) <https://www.fda.gov/medical-devices/software-medical-device-samd/global-approach-software-medical-device> [<https://perma.cc/Z2SN-HR66>] (adopting IMDFR's regulatory framework).

13. *Classify Your Medical Device*, FDA (Feb. 7, 2020), <https://www.fda.gov/medical-devices/overview-device-regulation/classify-your-medical-device> [<https://perma.cc/X8A7-CUMY>].
14. *Id.*
15. *Id.*
16. Price et. al., *supra* note 8, at 1140.
17. *Id.*
18. *Id.* at 1139–40 (Notably, health systems have begun developing the capacity to manufacture drugs, though so far their efforts are entirely focused on generic manufacturing and not developing new products); *see, e.g.*, CIVICARX, <https://civicarx.org/> [<https://perma.cc/D758-HG6B>] (describing a collaborative generic drug manufacturing initiative undertaken by health systems).
19. *See, e.g.*, Mark Sendak et al., *Real-World Integration of a Sepsis Deep Learning Technology into Routine Clinical Care: Implementation Study*, 8 JMIR MED. INFORM. 1, 11 (2020) (discussing Duke Health's Sepsis Watch program).
20. *See, e.g.*, Michael Thompson, *New Ways to Improve Hospital Flow with Predictive Analytics* (March 18, 2020), <https://www.slideshare.net/healthcatalyst1/new-ways-to-improve-hospital-flow-with-predictive-analytics> [<https://perma.cc/62RS-ZFFR>] (outlining Cedar-Sinai Medical Center's AI tool that aims to reduce capacity strain by predicting patient census).

---

models to allocate care coordinators to reduce costs.<sup>21</sup> These user innovations are very unlikely to see any sort of formal FDA review for multiple reasons, including FDA’s long-time exercise of enforcement discretion over laboratory-developed tests,<sup>22</sup> the exclusion of many clinical decision support software products (CDS) from the definition of medical devices under the 21st Century Cures Act,<sup>23</sup> and, potentially, other jurisdictional issues.<sup>24</sup> These systems are already being used<sup>25</sup> and they undoubtedly have quality problems—indeed, multiple examples of quality failures described above were in-house systems,<sup>26</sup> but FDA typically does not review them.<sup>27</sup>

On a broader scale, AI products are being embedded within electronic health record (EHR) systems that are then distributed to health care providers, sometimes quite broadly.<sup>28</sup> Epic, the market leader for EHR systems, has

---

21. Obermeyer et al., *supra* note 9, at 447 (“Large health systems and payers rely on this algorithm to target patients for ‘high-risk care management’ programs. These programs seek to improve the care of patients with complex health needs by providing additional resources, including greater attention from trained providers, to help ensure that care is well coordinated. Most health systems use these programs as the cornerstone of population health management efforts, and they are widely considered effective at improving outcomes and satisfaction while reducing costs.”).
22. Price et. al., *supra* note 8, at 1162; *see also* U.S. Food & Drug Admin., Off. on In Vitro Diagnostics and Radiological Health, *Draft Guidance for Industry, Food and Drug Administration Staff, and Clinical Laboratories, Framework for Regulatory Oversight of Laboratory Developed Tests (LDTs)*, FDA at 5–7 (Oct. 3, 2014), <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/framework-regulatory-oversight-laboratory-developed-tests-ldts> [<https://perma.cc/Z5NQ-H67A>]; U.S. Food & Drug Admin., *Discussion Paper on Laboratory Developed Tests (LDTs)*, at 4–5 (Jan. 2017), <https://www.fda.gov/medical-devices/in-vitro-diagnostics/laboratory-developed-tests> [<https://perma.cc/7G9Q-5U9Y>].
23. 21st Century Cures Act, Pub. L. No. 114–255, §3060, 130 Stat. 1033, 1130–33 (2016) (codified at 21 U.S.C. § 360).
24. Price et. al., *supra* note 8, at 1143 (“The FDCA only applies to products that are introduced, delivered, or received in interstate commerce, an important limitation that may exclude many user innovations.”) (citing 21 U.S.C. § 331).
25. *See generally* Sendak et al., *supra* note 1919; Thompson, *supra* note 20.
26. *See, e.g.*, Monika K. Goyal, et.al., *Racial Disparities in Pain Management of Children With Appendicitis in Emergency Departments*, 169 JAMA PEDIATRICS 996 (2015).
27. This is not to say that FDA is totally uninvolved; at least some developers have had discussions with FDA officials about how to *avoid* FDA review by ensuring sufficient presence of a human in the algorithmic loop to stay within the CDS exclusion of the Cures Act. Price et. Al., *supra* note 8, at 1145.
28. *See generally* Sehj Kashyap, Keith E Morse, Birju Patel & Nigam H Shah, *A Survey of Extant Organizational and Computational Setups for Deploying Pre-*

---

developed several AI-based tools that are integrated into its EHR suites, including some addressing COVID-19 risk and deterioration probability for patients more generally.<sup>29</sup> It is also developing a marketplace for developers to interact with its EHR systems.<sup>30</sup> Other EHR vendors, such as Cerner, have also developed AI products, though Epic appears to have taken an early lead.<sup>31</sup> Although these products are already widely distributed, it appears that no EHR-vendor-developed AI systems have gone through FDA premarket review.<sup>32</sup> These systems have quality issues, too; a wide-ranging review

---

*dictive Models in Health Systems*, 28 J. AM. MED. INFORM. ASSOC. 2445 (2021).

29. Alicia Reale-Cooney, *COVID-19 Risk Model Developed by Cleveland Clinic Now Available to Health Systems Around the World Through Epic*, CLEVELAND CLINIC (Nov. 9, 2020), <https://newsroom.clevelandclinic.org/2020/11/09/covid-19-risk-model-developed-by-cleveland-clinic-now-available-to-health-systems-around-the-world-through-epic/> [<https://perma.cc/EBN9-FF6X>] (Epic integrated AI systems into its COVID-19 prediction models—models that are used to determine an individual’s likelihood of testing positive for COVID-19 as well as their likelihood of needing critical care after testing positive (e.g., the Deterioration Index)); *Epic AI Helps Clinicians Predict When COVID-19 Patients Might Need Intensive Care*, EPIC (May 18, 2020), <https://www.epic.com/epic/post/epic-ai-helps-clinicians-predict-covid-19-patients-might-need-intensive-care> [<https://perma.cc/7MJP-G4DS>] (prior to the pandemic, Epic produced SlicerDicer, an AI-based tool that “allows a provider to tap into patient data to investigate clinical conjectures or make new discoveries about patient populations.”); Christina DuVernay, *SlicerDicer Reveals Practice-Based Data*, JOHNS HOPKINS MED. (Aug. 31, 2017), <https://www.hopkinsmedicine.org/office-of-johns-hopkins-physicians/best-practice-news/slicerdicer-reveals-practice-based-data> [<https://perma.cc/PJZ9-ZYCB>]. (Epic is also currently developing EHR document assistant that would enable AI to transcribe patient-clinical interactions); Christopher Jason, *Epic in Process of Developing AI EHR Documentation Assistant*, EHR INTEL. (Feb. 26, 2021), <https://ehrintelligence.com/news/epic-in-process-of-developing-ai-ehr-documentation-assistant> [<https://perma.cc/TLC9-F8UT>].
30. *Epic App Orchard*, EPIC, <https://apporchard.epic.com> [<https://perma.cc/8KCR-QKYW>].
31. Christopher Jason, *Epic Systems, Cerner Lead EHR Vendors in AI Development*, EHR INTEL. (May 12, 2020), <https://ehrintelligence.com/news/epic-systems-erner-lead-ehr-vendors-in-ai-development> [<https://perma.cc/96P4-MGNJ>].
32. *Artificial Intelligence and Machine Learning (AI/ML)-Enabled Medical Devices*, FDA (Sept. 22, 2021), <https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-and-machine-learning-aiml-enabled-medical-devices> [<https://perma.cc/FM6H-Y8TL>] (In September 2021, FDA released a list of AI and machine-learning enabled medical devices that have been reviewed and approved for the U.S. market. The list does not include any products by Epic or Cerner); *Artificial Intelligence and Machine Learning (AI/ML)-Enabled Medical Devices*, FDA (Sept. 22, 2021), <https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-and-machine-learning-aiml-enabled-medical-devices> [<https://perma.cc/FM6H-Y8TL>].

---

found problems in many of Epic’s AI products,<sup>33</sup> and another documented problems with its model reporting guidelines.<sup>34</sup>

Overall, a substantial number of AI products already being used in care settings do not appear to pass through centralized, national-level review for safety and efficacy.<sup>35</sup> This lack of governance isn’t because the products are great, but the limited scope of FDA review is not the only problem of centralized governance.

## B. What Can FDA Do?

The second problem is more systematic: Even for those AI products that do go through FDA review, that review only addresses a subset of issues for which governance is necessary.<sup>36</sup>

A burgeoning literature addresses the limitations of FDA review for AI products writ large. For instance, the vast majority of AI products that have received some sort of FDA marketing authorization have gone through the Section 510(k) clearance process rather than a full approval.<sup>37</sup> Section 510(k) clearance requires demonstrating substantial equivalence to a product that is

---

[www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-and-machine-learning-aiml-enabled-medical-devices](https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-and-machine-learning-aiml-enabled-medical-devices) [https://perma.cc/FM6H-Y8TL]; see also Jodi K. Scott, Lina Kontos, Suzanne Levy Friedman, Kristin Zielinski Duggan, Kelliann Payne, *FDA Launches List of AI and Machine Learning-Enabled Medical Devices*, HOGAN LOVELLS (Sept. 23, 2021), <https://www.engage.hoganlovells.com/knowledgeservices/news/fda-launches-list-of-ai-and-machine-learning-enabled-medical-devices> [https://perma.cc/KQ9X-FTLV].

33. Ross, *supra* note 6.
34. Jonathan H. Lu, Alison Callahan, Birju S. Patel, Keith E. Morse, Dev Dash & Nigam H. Shah, *Low Adherence to Existing Model Reporting Guidelines by Commonly Used Clinical Prediction Models*, MEDRXIV (Jul. 23, 2021) <https://www.medrxiv.org/content/10.1101/2021.07.21.21260282v1> [https://perma.cc/QK3L-QY5Q].
35. FDA review is not the only form of such review; insurers also sometimes play a quality-review function—but appear not to be playing such a role here. Price et. al., *supra* note 6, at 1161.
36. See generally Stan Benjamens, Pranavsingh Dhunoo & Bertalan Meskó., *The State of Artificial Intelligence-based FDA-approved Medical Devices and Algorithms: An Online Database*, NPJ (Sept. 11, 2020) <https://www.nature.com/articles/s41746-020-00324-0> [https://perma.cc/3S9L-PU6E].
37. Charlotte Tschider, *Medical Device Artificial Intelligence: The New Tort Frontier*, 46 BYU L. REV. 1551, 1597 (2021). The minority that have not gone through the 510(k) process have been de novo classified as Class I and II devices rather than undergoing the full Class III premarket approval process. Medical Device De Novo Classification Process, 86 Fed. Reg. 54,826 (Oct. 5, 2021) (to be codified at 21 C.F.R. pt. 860); see also *The De Novo Pathway: What Has Changed in 10 Years?*, GUIDED SOLUTIONS (Aug. 19, 2019), <https://>



---

already marketed,<sup>38</sup> and FDA has made clear that an AI product can be demonstrated to have substantial equivalence to an already approved non-AI product.<sup>39</sup> But the “substantial equivalence” standard has long been criticized for insufficient rigor.<sup>40</sup> Such arguments apply in the context of AI systems as well.<sup>41</sup> Critics also note FDA’s potential deficit in terms of AI-expert personnel<sup>42</sup> (a deficit FDA is trying to remedy<sup>43</sup>) and raise the challenges of regulating products that can and potentially should be updated relatively frequently as new data and performance metrics become available.<sup>44</sup> Some note the challenge of regulating systems incorporating AI as whole systems, rather

---

[www.guidedsolutions.co.uk/blog/the-de-novo-pathway/](http://www.guidedsolutions.co.uk/blog/the-de-novo-pathway/) [https://perma.cc/J9M3-7H9K].

38. *Premarket Notification 510(k)*, FDA (March 13, 2020), <https://www.fda.gov/medical-devices/premarket-submissions/premarket-notification-510k#se> [https://perma.cc/NS9C-YT9T].
39. *See How FDA Regulates Artificial Intelligence in Medical Products*, PEW TRUSTS (Aug. 15, 2021) <https://www.pewtrusts.org/en/research-and-analysis/issue-briefs/2021/08/how-fda-regulates-artificial-intelligence-in-medical-products> [https://perma.cc/YJA3-795Q].
40. Jeffrey K. Shapiro, *Substantial Equivalence Premarket Review: The Right Approach for Most Medical Devices*, 69 FOOD & DRUG L.J. 365, 365 (defending the substantial equivalence standard but acknowledging that it “is frequently compared unfavorably to [pre-market] approval as a means of establishing the safety and effectiveness of new devices, as an affront to the original intent of the Medical Device Amendments of 1976 (MDA), and as a “regulatory loophole” that should be scrapped or, if that is not practical, at least limited to the extent possible.”) (internal citations omitted). Indeed, the Supreme Court has recognized the difference in rigor between approval and clearance, finding that approval preempts state tort litigation for certain types of defect, while clearance does not. *Riegel v. Medtronic, Inc.*, 552 U.S. 312, 340–42 (2008); *see also* Tschider, *supra* note 37, at 1597.
41. *See* Soleil Shah & Abdul El-Sayed, *Medical Algorithms Need Better Regulation*, SCI. AM. (Oct. 7, 2021) <https://www.scientificamerican.com/article/the-fda-should-better-regulate-medical-algorithms/> [https://perma.cc/RMJ4-ZSY7].
42. Tschider, *supra* note 37, at 1586 (“To what degree FDA personnel or panel members actually provide expert direction in this [pre-market approval] review process is unknown, though facially it seems unlikely that personnel and panel members are equipped to review software design and anticipate real patient risks for new software technology like AI from a position of deep expertise.”).
43. Dave Muoio, *Commissioner Hahn: FDA Hiring More Data Experts to Help Healthcare ‘Unleash the Power of Data’*, MOBI HEALTH NEWS (Oct. 12, 2020, 3:09 PM), <https://www.mobihealthnews.com/news/commissioner-hahn-fda-hiring-more-data-experts-help-healthcare-unleash-power-data> [https://perma.cc/KUE2-2U79].
44. *Id.*

---

than trying to focus on the algorithm itself.<sup>45</sup> And finally, the set of entities with the resources and capabilities of taking an AI system through the FDA evaluation process is itself limited—especially given uncertainty about payment and reimbursement mechanisms<sup>46</sup>—impacting who is able to innovate effectively if the default model runs through FDA.<sup>47</sup>

Even setting aside the questions about how to get centralized governance functioning well in the first place, centralized governance can only do so much for AI systems that frequently need to be adapted and responsive to local environments.<sup>48</sup> Some products might reasonably expect to work essentially the same irrespective of context; whether a retina shows signs of diabetic retinopathy when examined through a uniform camera system *hopefully* does not change depending on whether that system is used in an academic medical center in Boston or a clinic in Alabama.<sup>49</sup> Other systems that similarly rely on measurements expected to be universally applicable might fit well into a centralized, national (or international) regulatory paradigm,<sup>50</sup> whether those systems calculate heart volumes using machine learning,<sup>51</sup> identify cerebral hemorrhage on CT scans,<sup>52</sup> or perform other presumably generalizable tasks.<sup>53</sup>

---

45. See Sara Gerke, Boris Babic, Theodoros Evgeniou & I. Glenn Cohen, *The Need for a System View to Regulate Artificial Intelligence/Machine Learning-based Software as Medical Device*, NPJ DIGIT. MED., April 7, 2020, at 1.
46. See Price et. al., *supra* note 8, at 1157–62.
47. See Price, *supra* note 2, at 452–53.
48. See Mark Sendak, Michael Gao, Marshall Nichols, Anthony Lin & Suresh Balu, *Machine Learning in Health Care: A Critical Appraisal of Challenges and Opportunities*, 7 eGEMS at 1, 2, (Jan. 24, 2019) <http://doi.org/10.5334/egems.287> [<https://perma.cc/M4E8-2RZM>] (“Personalized medicine will require mass customization of models that are trained and re-calibrated at the hospital and cohort level. Modern machine learning techniques focus on generalization beyond a training dataset, not on generalization to different sites.”).
49. See Michael D. Abràmoff, Philip T. Lavin, Michele Birch, Nilay Shah & James C. Folk, *Pivotal Trial of an Autonomous AI-based Diagnostic System for Detection of Diabetic Retinopathy in Primary Care Offices*, NATURE PARTNER J. DIGIT. MED., at 1 (Aug. 28, 2018).
50. I say “might” because differences in how data are recorded and processed or underlying populations could still limit widespread applicability.
51. See Akhil Narang et. al, *Machine Learning Based Automated Dynamic Quantification of Left Heart Chamber Volumes*, 20 EUR. HEART J. 541 (2018).
52. See Mohammad R. Arbabshirani et. al, *Advanced Machine Learning in Action: Identification of Intracranial Hemorrhage on Computed Tomography Scans of the Head with Clinical Workflow Integration*, 1 NPJ DIGIT. MED., Art. 9, at 1 (2018).
53. The commercially distributed IDx-DR diabetic retinopathy detection system requires a specific camera system for precisely this reason. See *generally* DIGIT.

On the other hand, many systems are inherently quite hard to generalize. Most obviously, systems that predict or recommend based on site workflow are closely tied to that site—a patient or staff volume predictor may only be applicable to the venue where it was developed.<sup>54</sup> Predictors of patient deterioration may not generalize well, whether because of differing patient populations,<sup>55</sup> differing care infrastructure and treatment patterns,<sup>56</sup> or simply differing data infrastructures so that different data are recorded and available for systems to use.<sup>57</sup> Predictors of sepsis, for instance, have been quite difficult to generalize across contexts, as have predictions of infection by *Clostridium difficile*.<sup>58</sup>

It is not always straightforward to figure out which generalizability bucket a particular product will fit into. Looking at skin lesions to identify potential skin cancer seems like it might be generalizable—except that it turns out patient skin color makes a big difference, and dermatological image databases from different places are, you guessed it, very different in patient demographics as well.<sup>59</sup> Questions of generalization arise not only across

---

DIAGNOSTICS, *IDx-DR*, <https://www.digitaldiagnostics.com/products/eye-disease/idx-dr/> [<https://perma.cc/HK6V-9TD4>].

54. See, e.g., Thompson, *supra* note 20.
55. W. Nicholson Price II, *Contextual Bias and Medical AI*, 33 HARV. J.L. & TECH. 66, 91–96 (2019).
56. See *id.* at 96–97; see also M. Scottie Eliassen, Ashleigh King, Christopher Leggett, Sukdith Punjasthitkul & Jonathan Skinner, *The Dartmouth Atlas of Health Care: 2018 Data Update*, DARTMOUTH ATLAS PROJECT (2021) (providing the latest update in a series dedicated to reporting the variation in care and medical services throughout the United States).
57. Price, *supra* note 55, at 100–101.
58. Wong et al., *supra* note 6 (finding that Epic’s sepsis prediction model correctly identified patients’ risk of sepsis only 63% of the time due to hospitals’ varying definitions of and billing codes for sepsis); Jeeheh Oh et al., *A Generalizable, Data-Driven Approach to Predict Daily Risk of Clostridium difficile Infection at Two Large Academic Health Centers*, 39 INFECTION CONTROL & HOSP. EPIDEMIOLOGY 425, 425 (2018) (finding that institution-specific models for estimating risk for *Clostridium difficile* provide “earlier and more accurate identification of high-risk patients and better targeting of infection prevention strategies” than a one-size-fits-all approach).
59. Veronica Rotemberg et al., *A Patient-Centric Dataset of Images and Metadata for Identifying Melanomas Using Clinical Context*, 8 SCI. DATA 34, 41 (2021); see also Nicole Westman, *Data Used to Build Algorithms Detecting Skin Disease is Too White*, THE VERGE (Nov. 9, 2021), <https://www.theverge.com/2021/11/9/22770852/data-dermatology-algorithms-skin-tone-ethnicity> [<https://perma.cc/LQY5-5J8Q>]; but see Angela Lashbrook, *AI-Driven Dermatology Could Leave Dark-Skinned Patients Behind*, THE ATLANTIC (Aug. 16, 2018), <https://www.theatlantic.com/health/archive/2018/08/machine-learning-dermatology-skin-color/567619/> [<https://perma.cc/NSB6-8VPF>] (citing Haenssle et

high-resource contexts in the U.S. but especially in contexts with different levels of resources (and presumably different care patterns),<sup>60</sup> and of course are likely to be even more prominent in international contexts,<sup>61</sup> though international implications are largely beyond the scope of this Essay.

All of this is not to say that AI systems developed specifically for a particular context cannot usefully be applied to other contexts—just that that application can take substantial work and is often not straightforward; more importantly for present purposes, centralized, national-level governance is a poor fit for localized application.

#### IV. DISTRIBUTED GOVERNANCE

Distributed, localized governance will be an essential complement to national regulators in providing robust oversight for medical AI. That recognition in itself is important; the normal approach to validating biomedical credence goods simply will not cover the gamut of useful medical AI, at least any time in the near future.<sup>62</sup> With the need for local governance taken as given for the near term, then, what could such local governance look like?

##### A. Examples of Local Governance

One starting point for the normative is the positive; here, some well-resourced health environments are already engaged in organized local governance efforts to validate AI before, during, and after deployment.<sup>63</sup> I focus here on academic medical centers which have the capacity both for in-house

---

al.'s study and arguing that, despite its clear racial disparities, machine-learning software could aid marginalized communities that don't have access to a dermatologist and can be improved over time).

60. Price, *supra* note 55, at 95–97.

61. See Daniel E. Weissglass, *Contextual Bias, The Democratization of Healthcare, and Medical Artificial Intelligence in Low- and Middle-Income Countries*, *BIOETHICS* 201 (Aug. 2021), <https://pubmed.ncbi.nlm.nih.gov/34460977/> [<https://perma.cc/XJV2-WLDW>].

62. Could we imagine AI engines so powerful, stable, and flexible that they can be validated at a national level and then deployed locally to integrate into various data ecosystems, collecting and refining data and self-validating through reports to centralized regulators like FDA? Of course! But that health-algorithmic utopia is a long way off.

It's also worth noting that localized governance is the norm for much of the *rest* of medical practice, such as physician oversight or hospital safety. Biomedical technologies have typically been treated differently, as described above.

63. Mark Sendak, *How Clinical AI is Being Deployed at Duke*, HIMSSTV (Nov. 3, 2021), <https://himsstv.brightcovegallery.com/detail/videos/latest-videos/video/6095739252001/how-clinical-ai-is-being-deployed-at-duke?autoStart=true> [<https://perma.cc/RF8E-YTTF>].

---

development and outsourcing models, with the recognition that governance capabilities and procedures may differ substantially in different contexts.

Duke, for instance, has devoted substantial resources to developing and deploying AI in medical practice.<sup>64</sup> It uses a stage-gate process, where a new AI model needs to move through various procedural steps before implementation.<sup>65</sup> Outcome measures are defined with the input of clinical leaders (that is, what is the algorithm supposed to do in practice), and then performance and improvement targets are set relative to the baseline of standard practice.<sup>66</sup> An algorithm is validated on retrospective data, then on current data while development continues.<sup>67</sup> Once the algorithm is trained, if adequate performance metrics are met, it is integrated into Duke's electronic health records system—but in the background, and with a smaller set of beta testers, who may see delayed results rather than real-time data and who do not use the algorithm to alter patient care.<sup>68</sup> This period is used to evaluate performance and determine what effective integration into the real clinical workflow would require.<sup>69</sup> If a model passes this background beta testing, it is evaluated in a prospective study, and only after passing that, can it be considered for integration into routine care.<sup>70</sup> At each stage, the process involves IT, individuals trained in technical evaluation and data analysis (such as the Duke Institute for Health Innovation<sup>71</sup>), and a set of clinical stakeholders across clinical departments: End users who will need to buy into the new technology, clinical leadership (including both nursing and physician staff), and operational leadership.<sup>72</sup>

The University of Michigan has a similar governance process in substantive terms, but processes initiatives through the Clinical Informatics Committee (The Committee).<sup>73</sup> The Committee includes academics, care

---

64. *Id.*; Video Interview with Mark Sendak, Population Health & Data Science Lead, Duke Institute for Health Innovation (Nov. 3, 2021); *see also Algorithm-Based Clinical Decision Support (ABCDS) Oversight*, DUKE AI HEALTH, <https://aihealth.duke.edu/algorithm-based-clinical-decision-support-abcds/> [https://perma.cc/4KJM-LNJL].

65. Video Interview with Mark Sendak, *supra* note 64; Sendak, *supra* note 63.

66. Video Interview with Mark Sendak, *supra* note 64; Sendak, *supra* note 63.

67. Video Interview with Mark Sendak, *supra* note 64; Sendak, *supra* note 63.

68. Video Interview with Mark Sendak, *supra* note 64; Sendak, *supra* note 63.

69. Video Interview with Mark Sendak, *supra* note 64; Sendak, *supra* note 63.

70. Video Interview with Mark Sendak, *supra* note 64; Sendak, *supra* note 63.

71. Video Interview with Mark Sendak, *supra* note 64; Duke Institute for Health Innovation, DUKE UNIV. SCH. OF MED., <https://dihl.org> [https://perma.cc/JHP9-3YG6].

72. Video Interview with Mark Sendak, *supra* note 64; Sendak, *supra* note 63.

73. Video Interview with Mark Sendak, *supra* note 64; AMIA Clinical Informatics Fellows, *Making Machine Learning Models Clinically Useful with Dr.*

---

providers, and other professionals.<sup>74</sup> Its membership includes informatics representatives from physicians, pathology, nursing, and research; it also includes several representatives from Health IT and Learning Health Systems and adds guests and ad hoc members as needed.<sup>75</sup> The Committee evaluates requests and suggestions for models to incorporate, and similarly runs them as background models for at least six months, after which decisions are made about potential integration into clinical workflow.<sup>76</sup>

New York University's Langone Medical Center takes a more siloed approach: it has an in-house team, conceived of as part of the IT department but including researchers and clinical staff, that develops and deploys models, including the process of evaluating model performance and clinical workflow integration.<sup>77</sup>

What do these various models tell us? Most basically, evaluating, implementing, and governing a model is a challenging, complex task. It involves substantial time, effort, and expertise across a significant timeframe and across both technical and clinical domains. Some entities treat this as a relatively enclosed process, largely entrusting it to one unit within the organization (e.g., NYU);<sup>78</sup> others treat it more as a committee service requirement developed on top of other responsibilities (e.g., Michigan);<sup>79</sup> and still others as a process involving both built-for-purpose subunits and broader clinical involvement (e.g., Duke).<sup>80</sup> This is a complex, multimodal, demanding process of governance.

Governance does not stop at implementation. After the process of deciding to develop or import a model, evaluating it, and then integrating it, effective local governance also requires ongoing monitoring of performance and undertaking maintenance and updating efforts as required. Among other

---

*Karandeep Singh*, AMERICAN MEDIA INFORMATICS ASSOCIATION (Nov. 29, 2021), <https://podcasts.apple.com/us/podcast/s3e8-virtual-care-information-security-and-data/id1211273977?i=1000552016091> [<https://perma.cc/JG6E-68SY>].

74. Video Interview with Mark Sendak, *supra* note 64; AMIA Clinical Informatics Fellows, *supra* note 73.

75. Video Interview with Mark Sendak, *supra* note 64; AMIA Clinical Informatics Fellows, *supra* note 73.

76. Video Interview with Mark Sendak, *supra* note 64; AMIA Clinical Informatics Fellows, *supra* note 73.

77. *Predictive Analytics and Machine Learning*, NYU LANGONE HEALTH CTR. FOR HEALTHCARE INNOVATION AND DELIVERY SCI., <https://med.nyu.edu/centers-programs/healthcare-innovation-delivery-science/predictive-analytics-unit> [<https://perma.cc/JB3L-R92U>].

78. *Id.*

79. AMIA Clinical Informatics Fellows, *supra* note 73.

80. Sendak, *supra* note 63.

---

things, data drift within a particular environment means that an AI system will tend to lose performance over time, absent regular updating.<sup>81</sup>

## B. Problematic Capacity Variations

Relying on the variety of sites that can use medical AI to create their own governance structures is, frankly, a recipe for failure. Resources vary wildly across contexts. While the sort of well-resourced academic medical center that has the capacity to develop AI may well have the capacity to evaluate, deploy, and otherwise govern that AI (with a bit of emphasis on “may,” since those capacities are distinct), contexts with fewer resources generally are less likely to have the governance infrastructure for evaluation and ongoing monitoring.<sup>82</sup>

At a basic level, health entities differ substantially with respect to the information technology resources necessary for deploying and evaluating models.<sup>83</sup> Running models silently in the background of systems for evaluative purposes demands more technical expertise and intervention than simply turning on a vendor-provided model and letting it run (though even that task is often fraught). And accurately collecting performance metrics brings similar challenges when deciding whether to go forward and integrate a model into the clinical workflow.

Perhaps more significantly, the human resources necessary for model evaluation are also highly disparate between settings.<sup>84</sup> Michigan, for in-

---

81. Adarsh Subbaswamy & Suchi Saria, *From Development to Deployment: Dataset Shift, Causality, and Shift-Stable Models in Health AI*, 21 *BIostatistics* 345 (arguing that datasets need to be monitored and maintained overtime to account for data shifts and their accompanying performance decay).

82. See, e.g., Tiankai Wang, Yangmei Wang & Alexander McLeod, *Do Health Information Technology Investments Impact Hospital Financial Performance and Productivity?* 28 *INT’L J. ACCT. SYS.* 1 (2018) (finding that health information technology investments, like electronic health records adoption, lead to positive financial performance and productivity); AM. HOSPITAL ASS’N, *RURAL REPORT: Challenges Facing Rural Communities and the Roadmap to Ensure Local Access to High-Quality, AFFORDABLE CARE* 14 (Feb. 2019) (“Rural hospitals are committed to improved care through use of HIT in order to meet past and current regulatory requirements . . . Rural hospitals must meet the same regulatory requirements [in this area] as other hospitals, yet often do not need the additional technology functionality contained in required, expensive system upgrades; nor do they have the available infrastructure such as adequate broadband to support them.”); Jordan Rau & Emmarie Huetteman, *Some Urban Hospitals Face Closure or Cutbacks as the Pandemic Adds to Fiscal Woes*, *NPR* (Sept. 15, 2020) <https://www.npr.org/sections/health-shots/2020/09/15/912866179/some-urban-hospitals-face-closure-or-cutbacks-as-the-pandemic-adds-to-fiscal-woe> [<https://perma.cc/VL3M-X6XQ>].

83. See Wang et al., *supra* note 82

84. See Rau & Huetteman, *supra* note 82.

---

stance, relies on committee work that typically sits atop other responsibilities.<sup>85</sup> Settings stretched for personnel simply may not be able to call on those human resources (and may not have staff with the relevant technical expertise in any case). Standalone units, like NYU Langone’s in-house development team,<sup>86</sup> are similarly out of reach for all but the most well-resourced medical environments.

Existing, intense, location-specific modes of governance are simply infeasible for many medical environments—including those that might benefit most from the ability of medical AI to democratize expertise, expand capacity, and improve care.

## V. POLICY INTERVENTIONS

Given the need for local, ongoing governance of medical AI and the currently unevenly distributed resources to conduct that governance, how can policymakers help enable that governance in the future? My aim here is not to solve the problem, but to sketch two directions for potential solutions: offloading and building capacity.

### A. Offloading

One approach is to provide resources for low-resource institutions to offload whatever governance tasks can be reasonably offloaded. At a federal level, FDA review of AI tools could presumably perform a partial oversight role (for those tools that go through FDA<sup>87</sup>)—demonstrating that a tool works in principle, for some set of assumptions—and then leaving the last-mile task of localized validation to the local entity. But FDA is not the only option. OCHIN, for instance, vets AI products and provides interfaces between those products and health system IT infrastructures to more easily adopt the products of trusted partners.<sup>88</sup> Other organizations, including for-

---

85. See W. Nicholson Price II, *Risk and Resilience in Health Data Infrastructure*, 16 COLO. TECH. L.J. 65, 76 (2017).

86. See Price, *supra* note 55, at 79.

87. See *supra* Section II.A.

88. See *OCHIN’s Preferred Technology Partners*, OCHIN, <https://ochin.org/technology-partners> [<https://perma.cc/XF6F-Q6KK>] (OCHIN collaborates with a range of technology partners) (“As leaders in the EHR space, we take great pride in being able to provide add-on functionality, services and upgrades to our members from our vast array of technological partnerships . . . OCHIN coordinates the interface and build required to use these products so our members don’t have to, and we are able to make it available for a fraction of market cost.”); see also Christopher Jason, *eHealth Exchange Taps Electronic Case Reporting for Interoperability*, EHR INTEL. (Aug. 17, 2020), <https://ehrintelligence.com/news/ehealth-exchange-taps-electronic-case-reporting-for-interoperability> [<https://perma.cc/7QM2-EKQD>]; Christopher Jason, *Epic Systems, OCHIN Launch COVID-19 Preparedness Screening App*, EHR INTELLIGENCE



profit entities that help health systems make procurement decisions, could develop similar capabilities.<sup>89</sup>

## B. Building Capacity

Offloading is unlikely to get us all the way there; policymakers should also consider investments in the infrastructure necessary to enable distributed governance, whether that infrastructure is technical, data-based, or procedural.

Technical infrastructure refers to the IT tools necessary to process data, run AI tools, integrate those tools into the care workflow, and—crucially—monitor and evaluate the outputs to measure performance.<sup>90</sup> Technical infrastructure, in addition to facilitating in-house distributed governance, can also facilitate the flow of monitoring data *out* of the low-resource context, to enable the sort of outsourced monitoring described above.<sup>91</sup> Technical infrastructure also includes developing programs or technical tools to monitor performance, to integrate AI systems developed elsewhere into local clinical data ecosystems and workflow, and to facilitate the training of local care providers on new tools.<sup>92</sup>

---

(March 30, 2020) <https://ehrintelligence.com/news/epic-systems-ochin-launch-covid-19-preparedness-screening-app> [<https://perma.cc/8XLR-B6YC>]; OCHIN, *OCHIN Joins NIH Funded AI/ML Consortium to Advance Health Equity and Researcher Diversity*, PR NEWSWIRE (Oct. 7, 2021), <https://www.prnewswire.com/news-releases/ochin-joins-nih-funded-aiml-consortium-to-advance-health-equity-and-researcher-diversity-301395495.html> [<https://perma.cc/U9W4-AS6A>].

89. See, e.g., Vizient, *Clinical Cost Management*, <https://www.vizientinc.com/our-solutions/clinical-solutions/clinical-cost-management> [<https://perma.cc/M4T4-XAX2>].
90. See, e.g., IBM, *What is IT Infrastructure?*, <https://www.ibm.com/topics/infrastructure> [<https://perma.cc/V27L-W7YJ>]; cf. Price, *supra* note 85, at 65, 67, 77–79 (“[H]ealth data infrastructure would be infrastructure *for* health data—that is, infrastructure on which health data can be stored and transmitted (such as computer systems, shared data standards, and the like). But it should also be infrastructure *of* health data—that is, a platform of shared data on which to base further efforts to increase the efficiency or quality of care. In an infrastructure *of* data, the data themselves are a resource to enable productive downstream activity that can improve the health care system.”).
91. Cf. Kai Hu et al., *Federated Learning: A Distributed Shared Machine Learning Method*, COMPLEXITY, Aug. 2021, at 1, 1 (analyzing federated learning: a machine learning (ML) framework where “multiple clients collaborate to solve traditional distributed ML problems under the coordination of the central server without sharing their local private data with others.”).
92. For instance, technical infrastructure can facilitate the reporting of adverse events, itself typically left to widely distributed individuals. See *FDA Adverse Event Reporting System (FAERS) Electronic Submissions*, FDA <https://>

---

Data infrastructure refers to the development of data resources on a broad, representative level for the development of AI tools.<sup>93</sup> Why is this necessary for distributed *governance*, rather than just for the development of tools? Three reasons: First, representative datasets allow at least some types of variation to be built into the development of AI tools, whether developed for a national audience or in-house with cross-validation on large, infrastructural datasets.<sup>94</sup> Accordingly, local governance should be easier, because some problems will be weeded out earlier. Second, the variations in large datasets can help illuminate the quirks and complexities of application to varied subsets of data—which can correspondingly flag issues that local governance should take into account.<sup>95</sup> And third, infrastructural datasets can help establish performance baselines against which tools can be measured.<sup>96</sup>

Procedural infrastructure refers to development of processes for governance so that each entity doing tasks need not reinvent the wheel but can instead rely on best practices and guides prepared by experts.<sup>97</sup> Many standards have been developed for determining how best to conduct AI studies

---

[www.fda.gov/drugs/questions-and-answers-fdas-adverse-event-reporting-system-faers/fda-adverse-event-reporting-system-faers-electronic-submissions](http://www.fda.gov/drugs/questions-and-answers-fdas-adverse-event-reporting-system-faers/fda-adverse-event-reporting-system-faers-electronic-submissions) [<https://perma.cc/MX2U-3HY5>] (describing distributed reporting of adverse events).

93. See Price, *supra* note 85.

94. See, e.g., *Core Values*, NAT'L INST. HEALTH ALL OF US RSCH. PROGRAM, <https://allofus.nih.gov/about/core-values> [<https://perma.cc/E63P-2HAC>] (“The All of Us Research Program is guided by a set of core values . . . Participants reflect the rich diversity of the United States. To develop individualized plans for disease prevention and treatment, researchers need more data about the differences that make each of us unique. Having a diverse group of participants can lead to important breakthroughs. These discoveries may help make health care better for everyone.”); *Data and Research Core*, AIM-AHEAD, <https://aim-ahead.net/home/leadership/research> [<https://perma.cc/6Z86-YQ6A>] (describing the goal of linking and preparing multiple sources and types of research data to form an inclusive basis for AI / ML).

95. Jeffrey S. Brown, Michael Kahn & Sengwee Toh, *Data Quality Assessment for Comparative Effectiveness Research in Distributed Data Networks*, 51 MED. CARE S22, S28 (2013).

96. See *id.*

97. See, e.g., *Responsible AI Practices*, GOOGLE, <https://ai.google/responsibilities/responsible-ai-practices> [<https://perma.cc/FZQ8-C6HE>]; *Good Machine Learning Practice for Medical Device Development: Guiding Principles*, FDA, <https://www.fda.gov/medical-devices/software-medical-device-samd/good-machine-learning-practice-medical-device-development-guiding-principles> [<https://perma.cc/5MWR-YBRN>]; Reid Blackman, *A Practical Guide to Building Ethical AI*, HARV. BUS. REV. (Oct. 15, 2020), <https://hbr.org/2020/10/a-practical-guide-to-building-ethical-ai> [<https://perma.cc/F9ZY-JP9C>].

---

and report results,<sup>98</sup> including performance metrics.<sup>99</sup> Similar tools can ease the process of governance once tools are developed and reported. What steps need to be taken to evaluate an AI tool for deployment—and if only a subset is feasible, what steps are most crucial? Correspondingly, what steps are so essential that insufficient capacity to perform them means that model deployment is too risky to go forward? This question is a tricky one! It’s easy to answer with the assumption that everything is necessary, an assumption which is only realistic from the comfortable seat of a high-resource setting. Knowing which corners can be reasonably cut is a key question for democratizing expertise in low-resource settings, in AI implementation just as much as in care.<sup>100</sup> Accordingly, process-based infrastructure tools should include some way to evaluate the costs, not only of deploying an AI system, but of effectively governing that system going forward.

Both infrastructure development and offloading have the potential for significant scaling, whether in a domestic or an international context. For instance, tools to monitor performance, processes for implementing governance structures, and dataset-driven knowledge about places to seek performance glitches should all be at least potentially deployable broadly—ideally to

---

98. See, e.g., Gary S. Collins, Johannes B. Reitsma, Douglas G. Altman & Karel G.M. Moons, *Transparent Reporting of a Multivariable Prediction Model for Individual Prognosis or Diagnosis (TRIPOD): The TRIPOD Statement*, 162 ANNALS INTERN. MED. 55 (2015) (describing the development of the TRIPOD Statement, “a set of recommendations for the reporting of studies developing, validating, or updating a prediction model, whether for diagnostic or prognostic purposes”); Robert F. Wolff, Karel G.M. Moons, Richard D. Riley, Penny F. Whiting, Marie Westwood, Gary S. Collins, Johannes B. Reitsma, Jos Kleijnen, Sue Mallett & PROBAST Group, *PROBAST: A Tool to Assess the Risk of Bias and Applicability of Prediction Model Studies*, 170 ANNALS INTERN. MED. 51 (2019) (describing PROBAST, “a tool for assessing the risk of bias and applicability of diagnostic and prognostic prediction model studies”); Gary S. Collins et al., *Protocol for Development of a Reporting Guideline (TRIPOD-AI) and Risk of Bias Tool (PROBAST-AI) for Diagnostic and Prognostic Prediction Model Studies Based on Artificial Intelligence*, 11 BMJ OPEN e048008, at 1 (July 2021) (explaining how PROBAST and the TRIPOD statement will be extended to prediction models that utilize machine learning techniques).
99. See, e.g., Mark P. Sendak, Michael Gao, Nathan Brajer & Suresh Balu, *Presenting Machine Learning Model Information to Clinical End Users with Model Facts Labels*, 3 NPJ DIGIT. MED. 1 (2020).
100. Cf. Price, *supra* note 55, at 113–14 (“The most straightforward way for AI algorithms to address cost issues would be to add those issues to the AI’s optimization function: that is, when scoring outcomes as desirable or undesirable (for the purposes of care recommendations, at least), the cost of care could be included in the score, rather than just patient health measures. Algorithms would then prioritize not simply outcomes or duplicating the patterns prevalent in High-Resource Hospitals, but also cost-effectiveness.”).

---

environments globally (with the obvious need for tweaks). Other interventions, like simply purchasing computer systems or deploying roving integration-and-evaluation teams, will be less easy to scale.

## VI. CONCLUSION

None of these interventions will be a magic bullet. Low-resource environments will still have few resources, barring massive structural change across the health system (domestic or worldwide). The people charged with adopting new technology, AI or otherwise, will still be overworked, under-resourced, and generally charged with doing too much with too little. Adding the initial and ongoing governance of AI tools onto overfull plates seems unfair. Nevertheless, there is hope! To the extent that AI has the possibility of making that task easier in *other* domains—of allowing more to be done with less for patient care, or resource management, or whatever else—it should be distinctly worth it in the end. Making the responsible adoption and governance of AI tools as easy and straightforward as possible looks likely to pay considerable dividends down the road and should be a focus for policymakers going forward.