

Supportive Structures for Healthcare AI Innovation and Adoption

Michelle M. Mello, JD, PhD¹

Stanford University

*Testimony presented to the U.S. House Committee on Energy and Commerce,
Subcommittee on Health, September 3, 2025.*

¹ Professor of Health Policy, Department of Health Policy, Stanford University School of Medicine; Professor of Law, Stanford Law School; Faculty Affiliate, Stanford Institute for Human-Centered Artificial Intelligence; Institute Faculty, Freeman Spogli Institute for International Studies, Stanford University.

Executive Summary

Although poorly designed regulation can hinder innovation, the government has a critical role to play in ensuring the conditions for innovation to translate into actual adoption of healthcare AI. In the U.S. market today, the key challenge isn't stunted innovation—it's low *uptake* of AI innovation. A major reason why adoption lags behind innovation and interest in AI is a foundational *trust deficit*. There are four areas where experts agree policy changes could promote AI adoption by building confidence in its performance:

1. **Ensure that the entities that develop and use AI adequately assess, disclose, and mitigate the risks of these tools.** Healthcare organizations and health insurers should be required to show that they have an AI governance process in place that meets certain standards. AI developers should be required to document and disclose key pieces of information about their products' design and performance.
2. **Support independent research on how AI tools perform in practice.** Such research can help healthcare organizations and insurers answer important questions about where investments in AI solutions can generate the most benefit and to minimize risks. It also ensures that this knowledge is disseminated broadly.
3. **Modify healthcare reimbursement policies to better support adoption and monitoring of effective AI tools.** Many AI tools will not save healthcare organizations money, and monitoring them properly can be costly.
4. **Address shortcomings in the Food and Drug Administration's statutory framework to make the agency a more constructive partner in AI development and adoption.** In some areas, the agency's authority doesn't go far enough; in others, it burdens developers with an antiquated regulatory framework that did not anticipate the AI revolution.

Statement

Chairman Griffith, Ranking Member DeGette, and Members of the Subcommittee, thank you for the opportunity to speak with you.

I have the privilege of being part of a group of data scientists, physicians, and ethicists at Stanford University—one of the earliest birthplaces of AI innovation—that helps govern how healthcare AI tools are used in patient care. I co-direct a lab at Stanford’s School of Medicine that evaluates risk and other ethical issues pertaining to proposed uses of AI tools in Stanford Health Care facilities, which care for over one million patients per year. I bring to that work more than two decades of research on patient safety and how law and policy shape healthcare delivery. I should note that I am here in my personal capacity and the views I’ll share reflect my own professional experience.

AI is already influencing life-and-death decisions in U.S. hospitals, often without many safeguards. That creates both a historic opportunity and a serious risk. My work has led me to understand that **although poorly designed regulation can hinder innovation, the government has a critical role to play in ensuring the conditions for innovation to translate into greater adoption of healthcare AI. The key problem isn’t that there isn’t a lot of innovation; it’s that uptake of new innovations is low. A major reason why adoption lags behind innovation and interest in AI is what experts have called “a foundational trust deficit.”**² I would like to share four important areas where experts agree policy changes could promote effective integration of AI tools into healthcare by strengthening trust.

² Elemento O, Sternberg CN, Khozin S. [Why Silicon Valley should demand clinical trials for its medical AI](#). STAT News, Aug. 28, 2025.

First, policymakers can help ensure that the entities that develop and use AI adequately assess, disclose, and mitigate the risks of these tools. You may be surprised to hear that most healthcare organizations and health insurers do little vetting of AI tools before they put them into use, and often no meaningful monitoring of their impact afterwards. There are good examples of AI governance processes and organizations willing to share what they've learned about how to do it. But governance takes work and resources, and the law doesn't require it, so most healthcare organizations don't do it.

Furthermore, developers are not required to make any particular disclosures when they pitch their tools to healthcare organizations and health insurers. The law also permits developers to disclaim liability and warranties for their products when they license them to healthcare organizations and insurers. As a result, developers currently have little incentive to reveal weaknesses of their AI tools and face little consequence when things go wrong.

So what does this rule-free space mean? It means that when your mother, spouse, or child seeks care at a U.S. hospital today, AI tools may strongly influence how they're prioritized for attention and services, which tests they receive, what records are kept about their care, and many other important decisions—sometimes solely on the basis of a sales pitch and monitored simply by checking that the software is on and working.

I am enthusiastic about healthcare AI. There are so many problems we've been trying to solve for decades that it may help with. I'm especially optimistic about the prospects for addressing the tragedy of missed and delayed diagnoses and the menace of physician burnout. But in healthcare, before we do things to patients, we test. We study. We gain physicians' and nurses' confidence that a new treatment or technology generates more benefit than harm. Yet

most healthcare organizations don't do that for AI and most developers don't provide much support to those that do.

Things are moving fast in the AI space and our traditional ways of testing new treatments, such as clinical trials, often aren't a good fit for the pace of innovation. But we still need some way of assessing and managing risk; we owe that to patients. I interview patients about their perceptions of healthcare AI tools in my work at Stanford, and they consistently express hope about AI but send us a stern message: *We expect you to keep us safe*. Research suggests that patients don't think we've risen to the challenge: 60% of US adults say they would be uncomfortable with their physician relying on AI and only one-third trust healthcare organizations to use AI responsibly.^{3,4,5}

Physicians are concerned, too. They are trained to be cautious, and savvy enough to realize they're usually the ones on the hook when patients get hurt. This is a huge problem for the AI market because uncertainty and fear about the risks of healthcare AI is currently chilling adoption. Again, the trust deficit is a major reason why AI adoption isn't higher—and why the largest area of AI implementation in healthcare isn't the exciting tools that promise great leaps forward in saving lives, it's tools that perform administrative tasks like taking notes during clinic visits.⁶

So how can we help? We need a supportive infrastructure to build trust in AI. For starters, **healthcare organizations and health insurers should be required to show that they have an AI governance process in place that meets certain standards.** The federal

³ Tyson A, Pasquini G, Spencer A, Funk C. [60% of Americans would be uncomfortable with provider relying on AI in their own health care](#). Pew Research Center. Feb. 22, 2023.

⁴ Platt J, Nong P, Carmona G, Kardia S. [Public attitudes toward notification of use of artificial intelligence in health care](#). JAMA Netw Open. 2024;7(12):e2450102.

⁵ Mello MM, Char D, Xu SH. [Ethical obligations to inform patients about use of AI tools](#). JAMA. Published online ahead of print, July 21, 2025.

⁶ Bessemer Venture Partners. [The Healthcare AI Adoption Index](#). Apr. 15, 2025.

government already does this for the human research it funds: it [requires](#) research organizations to show they review and monitor human subjects research using institutional review boards to protect participants. Analogously, through their operation of and certification processes for Medicare, Medicaid, the Veterans Affairs Health System, and other health programs, Congress and federal agencies can require that participating hospitals and clinics have a process for vetting any AI tool that affects patient care before deployment and a plan for monitoring it afterwards.^{7,8,9} The Joint Commission, an independent organization that inspects healthcare facilities for purposes of certifying their compliance with the Medicare Conditions of Participation, has already developed a voluntary certification standard for the Responsible Use of Health Data that focuses on how patient data will be used to develop algorithms and is exploring a similar certification for facilities' use of AI tools.

There are many good models for institutional governance processes. At Stanford, a C-suite-level committee decides which AI tools the hospital will deploy. To inform these decisions, the hospital provides support for data scientists to evaluate the tool for clinical utility and performance in subgroups, and for my team of ethicists to interview clinicians, AI developers, and patients to learn what matters to them and what they're worried about.¹⁰ We find that with a small investment of effort, we can spot potential risks, mismatched expectations, and questionable assumptions. In most cases, our recommendations don't halt deployment but rather

⁷ Mello MM. [Facilitating responsible governance of healthcare AI tools](#). Testimony Before the U.S. Senate Committee on Finance, Feb. 8, 2024.

⁸ Pencina MJ, Silcox C, Economou-Zavlanos N, McClellan M. [Bridging the gap between developers and implementers in AI](#). JAMA Health Forum. 2025;6(6):e251692.

⁹ Fleisher LA, Economou-Zavlanos NJ. [Artificial intelligence can be regulated using current patient safety procedures and infrastructure in hospitals](#). JAMA Health Forum. 2024;5(6):e241369.

¹⁰ Callahan A, McElfresh D, Banda JM, Bunney G, Char D, Chen J, et al. [Standing on FURM ground: a framework for deploying fair, useful, and reliable AI models in health care systems](#). NEJM Catalyst. 2024;5(10).

strengthen plans for monitoring the tool and training users. We designed this process to be scalable and exportable to other organizations.

This shows that governance is feasible and scalable—but only if developers provide basic information about how their tools work and where they fall short. That’s why transparency is essential. **Developers should be required to document and disclose key information about AI design and performance.** Some developers are forthcoming about this; others are less so. Getting important safety information shouldn’t depend on the customer being sophisticated enough about AI to know what questions to ask.

There is a great deal of careful work on so-called “model card” design to draw on here—short, standardized documents that summarize how an AI system was built, report on performance testing results, and describe known limitations.¹¹ Here is one example of how such information can help avoid harm to patients: Hospitals adopting ambient scribe tools usually tell clinicians to check the draft notes but don’t give specific guidance.¹² The reason hospitals adopt scribe tools is that clinical staff are overburdened and pressed for time, so the risk that clinicians may miss errors in the draft notes isn’t trivial. If the developer discloses that the scribe tool performs worse for certain visits—such as where more than two health problems are discussed, where the patient brings a third person into the conversation, or where the patient speaks with an Indian accent—the hospital can train clinicians to give special attention to notes from those visits and it can set up a study of whether clinicians successfully catch errors in those high-risk situations.

¹¹ See, for example, [work products from the Coalition for Health AI](#) (CHAI), a collaboration of nearly 3,000 technology companies, health systems, standards organizations, medical advocacy groups, and other organizations.

¹² Ambient scribe tools listen to a clinical conversation, use AI to generate a complete transcript of the conversation, and then use AI to draft a summary of important points in the conversation and place it in the patient’s medical record for the clinician to review and approve.

A second step the government can take to foster greater adoption of AI in healthcare is to support independent research on how AI tools perform in practice. The Agency for Healthcare Research and Quality, the National Cancer Institute, and other agencies have a long tradition of sponsoring research to ensure that care innovations, as actually implemented, produce benefit for patients. Independent research on AI can help healthcare organizations and insurers answer important questions like, *Where could investment in AI solutions generate the greatest benefit? What are the risks of AI tools and what works in mitigating them? What do we need to understand about how doctors and nurses interact with AI to design workflows to make best use of promising tools?* By funding and disseminating this kind of work, the government can make sure knowledge doesn't just stay in elite health systems but gets into the public domain so that smaller healthcare organizations have the information they need.

Third, healthcare reimbursement policies need to support adoption and monitoring of effective AI tools. Some healthcare AI tools enhance revenue for healthcare organizations—for instance, a tool that helps interpret imaging scans can help private-practice radiologists review (and bill for) more scans per day. But for many other tools, healthcare organizations can have no reasonable expectation that adoption will mean cost savings. Time savings are often modest because the humans-in-the-loop take time to review and evaluate the model output. Where there are efficiency gains, hospitals are reluctant to respond by asking physicians and nurses to see more patients because that could exacerbate burnout and breed resistance to the AI tools. Even where AI tools are inexpensive to buy, they can be costly to monitor. These realities are why market forces alone, with current reimbursement policies, may not support widespread adoption and monitoring of healthcare AI tools. Adapting Medicare and Medicaid

reimbursement to provide mechanisms to help cover costs would bring the benefits of AI tools to patients cared for by rural hospitals and children's hospitals, among other places.

Finally, the Food and Drug Administration (FDA) should be empowered to be a more constructive partner in AI development and adoption. There is only so much the agency can do with a [statutory framework](#) that dates to the Ford Administration and the year Apple Computer was founded. Most healthcare AI tools aren't cleared by the FDA, and *can't* be, under FDA's current scope of authority. This includes many tools that have important implications for patient safety and quality of care. Moreover, the standards and processes laid out in the law for clearing medical devices often aren't a good fit for evaluating AI tools, which may have hundreds of different clinical indications and which continue to learn and change after marketing clearance is issued. Addressing these problems requires action by Congress. Modernizing the agency's authority over software tools can help fill the trust deficit while also fixing areas where the current regulatory framework makes it unduly burdensome to bring AI products to market.

In summary, AI holds enormous promise for improving healthcare, but its adoption is being slowed by a fundamental trust deficit. By taking practical steps now, Congress can help close that gap, make healthcare providers and the public more excited to receive the products coming out of industry, and ensure that innovation truly reaches the bedside.

Thank you, and I welcome your questions.