


The AI medical services act: what it gets right, where it falls short, and what it matters for the next decade of digital health

MAR 01, 2026



Abstract

- Core argument of the bill: prohibition fails, unregulated consumer tools already in the vacuum, access crisis is present not future
- Key smart elements: tiered licensure model, supervised deployment, regulatory sandbox, bias monitoring requirements
- Key weaknesses: underspecified clinical validation standards, reimbursement framework ignores ERISA preemption and CPT mechanics, thin liability allocation, interstate delivery blind spots

Table of Contents

The Setup: Why This Bill Exists

What the Bill Actually Does

The Smart Stuff

Where the Argument Gets Thin

Reimbursement: The Glaring Gap

The Liability Question

The Setup: Why This Bill Exists

Start with the honest framing: healthcare's access problem is not hypothetical. There are roughly 100 million Americans living in primary care shortage areas according to HRSA data. The AAMC projects a physician shortage of 40,000 to 124,000 by 2030. Rural hospitals have been closing at approximately 20 per year for the past decade. Specialists in behavioral health, nephrology, and geriatrics are particularly scarce outside major metros. These are not theoretical future risks. They are the current operating environment.

Into this environment, AI tools are already proliferating at scale, and not always in careful, clinically integrated ways that anyone serious in this space would prefer. Companies are releasing consumer-facing diagnostic apps, mental health chatbots, and chronic disease management tools that operate entirely outside clinical oversight structures. These products are not being reviewed for safety or efficacy in any meaningful way. They are not billable under Medicare or Medicaid. They tend to attract cash-pay users, which by definition skews toward higher-income patients away from the populations that most need care capacity expansion.

The AMSA is responding to a real dynamic. The choice is not between regulating AI in healthcare and keeping AI out of healthcare. The choice is between regulating AI and watching it proliferate in the least accountable form possible. That framing is the right spin. It is actually correct. The consumer health app market is enormous and largely ungoverned, and the people building those products are not going to stop because the state legislature did nothing. The bill is trying to move the action inside the tent rather than outside it, and that instinct is the right one.

This framing aligns directly with what Sebastian Caliri, Adam Meier, and Joe Lonsdale are explicitly asking for feedback on in the thread that has been circulating around this draft. Their ask is whether the bill lets builders harness AI for maximum

impact on the US healthcare system. The honest answer is: it creates the right conditions, but the execution details will determine whether this becomes a real market or just a new compliance category.

What the Bill Actually Does

The structural mechanics of the AMSA center on creating a new licensure category called the AI Medical Services Provider, or AIMSP. This is a legal entity designation that allows an AI system, operated by a licensed entity, to deliver defined clinical services to patients. The framework is tiered based on risk, which is a sensible design choice borrowed from how the FDA approaches device classification. Lower-risk services like health assessments, triage guidance, and chronic disease monitoring carry lighter oversight requirements than higher-risk applications like diagnostic imaging interpretation or medication management.

The bill requires that any AIMSP operate under the supervision of a licensed physician or advanced practice provider. That supervision requirement has actual teeth, not nominal ones. The supervising clinician bears accountability for the AI's clinical outputs, creating real skin in the game. There are requirements for regular performance auditing, bias monitoring, and adverse event reporting. AI systems seeking licensure must demonstrate clinical validation through outcomes data, not just bench testing or theoretical performance metrics. The framework includes a regulatory sandbox that allows innovators to apply for limited, monitored deployment before full licensure, which is a practical acknowledgment that you cannot validate these systems in a vacuum before any patient ever sees them.

On reimbursement, the bill asserts that AIMSP-delivered services should qualify for state Medicaid reimbursement and mandates private insurer coverage for licensure services. Payment rates are to be established by a newly created AI Medical Services Board, composed of clinicians, technologists, patient advocates, and ethicists. Liability flows through the supervising physician and through the AIMSP entity, with specific provisions for software developers under a framework that distinguishes between design defects, training data deficiencies, and deployment errors.

The Smart Stuff

The tiered risk framework is genuinely well-designed. Applying the same regulatory overhead to an AI triage tool as to an AI-assisted surgical guidance system would be absurd and would effectively prohibit lower-risk innovation while doing nothing to constrain higher-risk deployment. The AMSA explicitly calibrates requirements to risk level, which is how every mature regulatory framework in medicine actually operates. The idea that AI should be categorically different is a bias toward novelty, not a defensible policy position.

The supervised deployment requirement is also smart in ways that are easy to underestimate. One of the genuine unsolved problems in clinical AI is the accountability gap. When an AI system produces a bad outcome, who is responsible? The patient's physician? The hospital? The software company? The company that trained the underlying model? Right now the honest answer is that nobody has clear legal accountability in most jurisdictions, which creates perverse incentives across the board. Physicians defensively disclaim responsibility. Software companies hide behind learned intermediary doctrine and terms of service. The AMSA plants a flag: the supervising clinician is responsible, and the AIMSP entity is responsible. That is a complete answer to the liability question, but it is better than the current answer which is effectively nobody.

The regulatory sandbox is worth flagging specifically for founders because the current path to market for clinical AI is genuinely broken. FDA breakthrough device designation takes years and costs millions. Operating without FDA clearance means living in permanent regulatory ambiguity and being locked out of serious hospital procurement. Releasing as a consumer app means forfeiting any reimbursement pathway and getting written off by institutional buyers. The sandbox creates a viable fourth option: operate under active regulatory supervision with defined patient safety requirements, generate real-world evidence, and build toward full licensure. This is a workable business model. It also reflects how most medical innovation actually unfolds in practice, through iterative deployment and refinement rather than obtaining approval for a fully perfected product before deployment.

The bias monitoring and adverse event reporting requirements reflect genuine technical sophistication about how AI systems fail. These tools do not fail the way traditional software fails, through bugs that produce consistent incorrect output. They fail through distributional shift, through training data that underrepresents certain patient populations, through feedback loops that amplify existing clinical biases. Requiring ongoing monitoring of algorithmic performance across demographic subgroups is not regulatory theater. It is the correct technical response to how these systems actually behave in production. Any founder who has deployed a model in the real world knows that the performance profile visible in validation bears limited resemblance to what shows up six months into live deployment.

Where the Argument Gets Thin

The bill's weakest section is its treatment of clinical validation standards. The requirement that AI systems demonstrate clinical validity through outcomes data is stated as a principle but left almost entirely undefined mechanically. What outcomes? Measured over what time horizon? Against what comparator? Using what statistical threshold for sufficiency? These questions are not minor details. They are the essential substance of what it means to validate a clinical AI system, and the bill essentially delegates all of this to the AI Medical Services Board to work out later.

That is not necessarily wrong as a legislative drafting strategy. Legislatures are not the right bodies to define sensitivity thresholds for diagnostic AI tools. But it does mean that the actual regulatory substance of the bill will be determined by whoever ends up on that Board, under whatever political pressures apply at the time of their appointments, and against whatever industry lobbying is most effective during rulemaking. From a founder or investor perspective, this is not a reason to dismiss the bill, but it is a strong reason to stay deeply engaged with the Board composition rulemaking process if the framework gets enacted. The real game will be played

The bill also leans heavily on the claim that supervised deployment will naturally improve safety outcomes relative to the current unregulated consumer app environment. That argument is probably right directionally, but the bill does not

establish mechanisms that actually guarantee it. A supervised deployment requirement means a physician signs off on using the system. It does not mean the physician reviews the AI's outputs for each patient encounter, has the technical capacity to evaluate algorithmic reasoning, or will catch errors that a sophisticated system produces in ways that look superficially plausible. The research on human oversight of automated systems is genuinely discouraging on this point. Automation bias is real and well-documented. People monitoring AI systems tend to trust the machine and miss errors, particularly when the AI is usually right. The bill acknowledges this dynamic nowhere.

The section on interstate AI systems is also underspecified in ways that will create real operational problems for anyone building at scale. A substantial portion of delivered clinical services will involve AI systems trained and operated in one state being deployed to patients in other states. The bill covers services delivered within an enacting state, but it does not address how an AIMSP licensed in one state interacts with regulatory frameworks in other states, how liability allocates when the supervising physician is in State A and the patient is in State B, or how to handle systems operated by national health systems that cannot maintain separate state-by-state compliance architectures. These are not edge cases. They describe the operating model of every serious national digital health company.

Reimbursement: The Glaring Gap

The reimbursement section reads like the part that was written last, after someone realized the framework needed an economic engine but did not have time to work through the mechanics. The assertion that private insurers shall cover AIMSP-delivered services is stated as a mandate without any of the actuarial or rate-setting substance that would make it real. Coverage at what rate? Using what CPT code? Subject to what prior authorization requirements? Under what medical necessity criteria? These questions are not answered, and the answers matter enormously because they determine whether this creates an actual market or just a theoretic one.

Health insurance reimbursement is a system built on CPT codes developed by the AMA, on relative value units assigned through a politically fraught committee process, and on coverage determination processes that routinely take years even for well-evidenced interventions. The AMSA asserts that a state board will solve all this without grappling with the reality that most private insurance reimbursement in the US is governed by ERISA-preempted employer plan documents that are not subject to state insurance mandates. A state law mandating insurer coverage of AIMSP services will have zero effect on the majority of commercially insured Americans, who are covered by self-funded employer plans that fall under federal jurisdiction. This is not an obscure technicality. It is the central structural fact of commercial insurance regulation in the US, and the bill does not acknowledge it.

The Medicaid pathway is more credible because states actually have authority over their own Medicaid programs. But Medicaid reimbursement rates are notoriously low, reimbursement processes are notoriously slow, and Medicaid managed care organizations have their own coverage determination processes that operate semi-independently of state fee schedules. Founders who have tried to build sustainable businesses on Medicaid reimbursement alone know that the theoretical availability of a payment pathway and the practical economics of getting paid are very different things.

What would a more rigorous reimbursement framework look like? At a minimum, it would identify the specific service categories eligible for reimbursement with proposed CPT code crosswalks, establish a rate-setting methodology that accounts for both AI service delivery costs and physician supervision costs, address the ERISA preemption problem for commercial insurance honestly, specify what evidence of clinical efficacy is required to trigger the coverage mandate, and establish a timeline for coverage determinations. None of that is in the bill. What is in the bill is a mandate that a Board will figure it out, which is effectively equivalent to having no reimbursement provision at all until rulemaking plays out years from now. The authors acknowledge this themselves in the thread, noting that reimbursement is a topic for future discussion. That is an honest concession, but it is also the thin

determines whether any of the rest of the framework produces a real business ecosystem.

The Liability Question

The liability framework in the AMSA is the section that will generate the most litigation if this bill is enacted, and likely the most investor anxiety in the meantime. The bill creates layered liability: supervising clinicians bear professional liability through existing malpractice frameworks, AIMSP entities bear organizational liability through a new statutory cause of action, and AI developers bear product liability for design defects and training data failures. That structure is defensible in theory. In practice there are several problems worth working through.

The distinction between a design defect, a training data failure, and a deployment error sounds clean but is extremely difficult to establish in the context of modern machine learning systems. When an AI tool produces harmful clinical advice, tracing that failure to its root cause requires extensive forensic analysis of model architecture, training data composition, validation methodology, deployment configuration, and specific inputs provided at the time of failure. The legal system is not equipped for this analysis, which means liability will effectively be determined by whoever retains the most convincing expert witness rather than by any principled fault allocation. Attorneys who work in medical device product liability will find this environment familiar and lucrative. The innovation community will find it difficult to predict the price.

For investors, the liability question matters because it directly affects the insurability of AIMSP businesses and their exposure to catastrophic loss events. The bill does not specify what insurance an AIMSP must maintain, what capital requirements apply, or what indemnification structures are permissible between AIMSP entities and the developers whose systems they deploy. A startup that licenses a foundation model from a large AI company and deploys it as a clinical service under the AIMSP framework needs to understand exactly what it is assuming liability for and what

can contractually shift back to the model provider. The bill's treatment of this is thin to provide founders or their investors any real comfort.

What This Means for Founders and Investors

The practical implications depend on where you are in the capital stack and development cycle. For early-stage founders building clinical AI tools, the AMSA framework, if enacted, creates a cleaner path to market than the current environment provides. The regulatory sandbox is genuine upside. Being able to operate under active state regulatory supervision, generate real-world evidence, and build toward licensure beats the current options of waiting years for FDA breakthrough device designation or flying under the radar and being locked out of serious institution customers. The supervised deployment requirement adds operational complexity but also forces the physician partnership structures that serious clinical AI products need to develop anyway.

For growth-stage companies with existing revenue from hospital and health system customers, the bill's impact depends heavily on how validation and audit requirements get operationalized in rulemaking. If the AI Medical Services Board establishes rigorous, technically sophisticated validation standards that map to how these systems actually perform in production, compliance costs will be real but manageable for companies with mature MLOps infrastructure. If the Board creates checkbox compliance requirements that are easy to satisfy but do not assess what actually matters, the framework will not improve safety and will create a false sense of accountability. That second outcome is probably more likely than the first given the typical composition of state regulatory bodies, which is a reason for the technical community to engage aggressively during rulemaking rather than leaving it to the default stakeholders.

For angels and syndicate investors, the most important signal in the AMSA is not the specific provisions but the direction of travel. State-level AI medical licensure frameworks are coming regardless of whether this specific bill serves as the template.

The question is whether the frameworks that emerge get designed with input from people who understand the technology, the clinical workflows, and the business models, or whether they get designed primarily by people protecting existing professional monopolies and minimizing political risk. The AMSA reads like a genuine attempt at the former, which is relatively rare in health tech policy.

On portfolio construction, companies that have been building for a regulated environment have a different and more defensible competitive position than companies whose business models depend on the current regulatory vacuum continuing. Clinical validation data, established supervising physician networks, auditable MLOps infrastructure, and experience navigating regulatory sandboxes are durable competitive advantages in a licensed market. They are largely irrelevant in the consumer app market. If the AMSA or something like it gets enacted in even a handful of states, the business models that work will shift substantially, and portfolios that anticipated that shift will look very different from those that did not.

Closing Take

The AI Medical Services Act is better policy thinking than most of what gets produced in health tech legislation. The core diagnosis is correct. The access crisis is real, AI deployment is inevitable, and the choice is between accountable regulated deployment and unaccountable unregulated proliferation. The tiered risk framework is smart. The regulatory sandbox is genuinely useful for founders. The supervised deployment requirement is the right instinct even if the oversight mechanisms need strengthening. The bias monitoring and adverse event reporting requirements show technical sophistication that is rare in legislation.

The gaps are real too. Clinical validation standards need to be far more specific and meaningful. The reimbursement framework needs to engage with CPT codes, EHR integration, and rate-setting methodology rather than delegating everything to a state Board. The liability framework needs cleaner rules around developer responsibility and AIMSP insurance requirements. The interstate service delivery problem needs a serious answer before any national-scale company can rely on this framework.

None of these gaps are fatal, and some are probably better addressed in rulemaking than statute. The bill is a framework, not a complete regulatory regime, and it is honest about that limitation. For founders and investors who have been operating in the current environment of regulatory ambiguity, the AMSA represents a bet that clarity, even imperfect clarity, is better than the status quo. That bet is almost certainly right. The window to shape what that clarity looks like is open right now during the drafting and rulemaking phases, and the people who engage seriously during that window will have disproportionate influence on the framework that ultimately governs this market. Given that this draft is being circulated publicly with an explicit request for builder feedback, the opportunity to actually shape this bill is real and right now.



2 Likes • 1 Restack

← Previous

Next →

Discussion about this post

Comments

Restacks



Write a comment...