To: Congress (Senate HELP; House Energy & Commerce; House Ways & Means); U.S. Department of Labor (Employee Benefits Security Administration); Centers for Medicare & Medicaid Services; Federal Trade Commission
From: Anna Patrick and Dwija Adamala
Date: December 15, 2025
Subject: Establishing a National Framework for AI-Influenced Health Insurance Coverage Decisions

Decision Requested

Adopt a national regulatory framework requiring (1) disclosure, (2) clinician accountability, and (3) independent auditing for any AI-influenced coverage determination, through an ERISA amendment complemented by coordinated federal and state implementation.

Executive Summary

- Artificial intelligence is increasingly shaping health insurance decisions that determine when patients receive care, how long they stay, and whether services are covered.
- Without clear guardrails, predictive tools can override clinical judgment and produce opaque denials that erode trust.
- This memo proposes a care-centered national framework for AI used in utilization management and prior authorization.
- The framework rests on three pillars: (1) transparent disclosure to patients whenever AI influences a determination, (2) accountability by a licensed clinician for any adverse decision, and (3) independent audit and validation of the tools and workflows that affect coverage.
- California's SB 1120 offers a workable model. Federal and state regulators can extend similar protections through aligned implementation, while Congress sets a uniform federal floor for ERISA plans.
- To encode the framework, Congress should amend ERISA to require disclosure of AI involvement, documented clinician responsibility for adverse determinations, and audit-ready logging and independent validation (including disparity monitoring).

Insurers are increasingly deploying predictive tools to estimate recovery time, length of stay, and medical necessity. These outputs often inform or drive utilization review and prior authorization, yet the criteria, logic, and limitations of the tools are invisible to patients and clinicians. Existing oversight is fragmented across agencies and plan types. Patients experience the system as a black box, and clinicians face denials keyed to averages rather than individual conditions. For clarity, we use the following definitions throughout:

- Predictive tool: Any algorithm, artificial intelligence system, or software that generates a recommendation, score, or classification used to inform or make coverage, authorization, or length-of-stay decisions, including rules-based engines and machine-learning models.
- Utilization management (UM) and utilization review (UR): Processes plans use to determine whether a requested service, setting, or duration is covered and medically necessary, including prior authorization, concurrent review, and retrospective review.

- AI-influenced determination: Any coverage, authorization, or length-of-stay decision where a predictive tool was used as an input, constraint, or default recommendation, regardless of whether the final decision was made by a human or automated workflow.
- Adverse decision: A denial, reduction, delay, or premature termination of coverage relative to the treating clinician's request, including shortened post-acute stays, limits on rehabilitation intensity, or step-therapy requirements that materially change the care plan.

A national, care-centered framework should preserve the benefits of responsible prediction while ensuring that decisions remain individualized, reviewable, and fair.

Scope of Problem

- The scale and impact of AI-influenced prior authorization decisions are substantial:
- Volume: Medicare Advantage insurers processed an estimated 49.8 million prior authorization determinations in 2023. Of the total, more than 3.2 million (6.4 percent) were unfavorable, with the denial rate being even higher at 7.4 percent in 2022 (KFF, 2025).
- Error and reversals: Federal oversight has found that about 13 percent of sampled prior authorization denials in Medicare Advantage met Medicare coverage rules and would likely have been approved under traditional Medicare (Office of Inspector General, 2022). Plans overturned roughly three quarters of appealed denials in some periods, yet only a small fraction of beneficiaries or providers file appeals.
- Timelines for Appeals: Internal and external appeal timelines can be from 15 to 30 days or longer, while post-acute episodes of care and stays for rehabilitation usually last days to several weeks. When appeals outlast the window of clinical concern, delayed care functions as a de facto denial (Medicare.gov, 2024; HealthCare.gov, 2025).
- Concentrated harm: Post-acute care, rehabilitation, and other services following a hospital stay are disproportionately represented among inappropriate denials and are particularly sensitive to lost days of coverage, as recent congressional investigations into Medicare Advantage post-acute care denials have documented (Senate Permanent Subcommittee on Investigations, 2024).
- Administrative burden: Prior authorization and associated UM activities are estimated to represent tens of billions of dollars in annual administrative spending. Hospitals and health systems report devoting significant staff time to appealing improper denials, which costs billions of dollars annually (American Hospital Association, 2024). Physicians report spending more than half a day each week on tasks related to prior authorization, often diverting clinical time to UM processes (Surescripts, 2025; AMA survey cited therein).

The combination of opacity, speed, and scale argues not just for general principles, but for specific, enforceable standards for transparency, clinical accountability, and independent validation.

Current Approaches

California's SB 1120 will take effect in 2025 and requires plans to disclose when algorithms are used in utilization review, makes sure that a qualified human makes and owns the determination, and preserves auditability of how a tool influenced an outcome for a patient (California Legislature, SB 1120, 2024). California regulators also reminded carriers that AI-

assisted decisions remain subject to market-conduct review and anti-discrimination law, including new guidance implementing SB 1120's requirements for non-discrimination, human medical-necessity review, and auditability of decision support tools (California Department of Insurance, 2025). Federally, existing authorities touch pieces of the problem. CMS guides Medicare Advantage utilization management, the Department of Labor oversees ERISA plans, state insurance departments regulate fully insured markets, and the FTC enforces against unfair and deceptive practices, including in use of automated decision systems. What is missing is nationwide and uniform practices that patients be told when AI was used and why, that a licensed clinician is accountable for denials, and that plans maintain auditable logs and submit tools to independent validation with disparity monitoring. A national framework should harmonize these expectations across plan types, minimize administrative burden through standard templates, and phase in requirements to protect access without chilling responsible innovation.

Proposal: A Care-Centered Framework for National Regulation

To address the limitations of current approaches to regulating the healthcare AI space, we propose a national AI for Health Insurance Framework specifically for the domain of predictive, automated decision making with the ability to deny patients care. The **AI in Healthcare Bill of Rights** should contain three pillars: mandating transparency of the use of AI in medical decision making, requiring a licensed clinician to be held accountable, and making these tools open to independent audit and validation.

I. Transparency when Care is Influenced by AI

When patients are not aware of the ways artificial intelligence shapes the decisions made about their care by their insurance companies, their ability to make their own informed decisions is weakened. A core tenet of respecting patients' agency and dignity is allowing them to make their own decisions about their care (Sedig, 2016); while patients may not have full control over their insurance, they should at least have the information necessary to petition or challenge an insurance claim if in their best interest. The policy would require the following:

- Mandate insurance companies to create an accessible AI policy, including information about patient privacy protections and data collection
- Disclose to patients the use of AI anytime predictive technology is utilized to make decisions about their care, via an established line of communication created at the onset of the insurance package

II. Accountability of Licensed Human Physicians

The current state of predictive AI in insurance often contributes to a disconnect between the needs of a patient and the care the company licenses them to receive. For example, consider the case where a patient receives a procedure, and the technology predicts a four day recovery based on similar cases and the patient's history. While this prediction could be accurate, it's also highly possible that nuances in the patient's procedure or medical records that a human physician would notice could suggest the need for additional care or hospital time. In these situations, it's crucial that the output of a predictive technology does not deny coverage to a patient in need. For these reasons, the framework for regulating AI in health insurance decisions must include these provisions:

- National regulations should mandate insurance companies to get physician sign off before algorithmically denying a patient care

- Physicians must be allowed to petition to insurance companies on behalf of their patients when their expertise contradicts AI decisions about patient care
- Insurance companies must designate a team responsible for reviewing physician petitions in a timely manner

III. Enforcement of Independent Auditing

The third pillar of the framework involves guardrails for auditing insurance companies to validate that the regulations proposed by the policy are adhered to. The framework hinges on insurance companies' capability to adhere to regulations, specifically in relation to processing petitions and maintaining fair AI policies. Thus, we propose that:

- The FTC should designate (or create) a specialized enforcement and technical review function for AI-enabled coverage determinations, coordinating with DOL/EBSA and CMS.
- Verified third party auditors be granted permissions to investigate insurance companies' AI policies and use cases
- Permissions for auditors to access necessary materials from insurance companies and model developers granted by ERISA

IV. Implementation

In addition to the implementation details included above, we propose that Congress pass an amendment to the Employee Retirement Income Security Act (ERISA), which sets standards for health plans. Many amendments have already been passed to ERISA to improve equity and care in the healthcare insurance industry (US Department of Labor). Additionally, insurance reform has high bipartisan support, making a congressional bill realistic for legislators (Kirzinger et al., 2025). In the case that Congress was unable to pass an amendment, the Department of Labor would need to work within the regulatory power already granted by ERISA to set standards for health insurance companies to follow.

Then, the standards would be enforced by a committee within the FTC to validate that insurance companies were adhering to necessary practices. To roll out the standards, we suggest the new committee should commit to in the next six months producing a public set of benchmarks and metrics for auditing of predictive AI models, including with relation to disparate impacts across groups. Additionally, the Department of Health and Human Services could monitor and research the strength of the policy in decreasing the number of incorrect denials of care over time. Together, these agencies would be able to implement the change necessary to address the current gap in regulating insurance companies' use of predictive AI.

V. Limitations

One possible limitation of the proposed policy framework is that the regulations place extensive responsibility on physicians to advocate for their patients and participate in the insurance pipeline. This could be argued as an undue pressure for medical care workers, who may already be overworked and unable to take on this additional component. However, we find it necessary to include physicians in our policy, as they are the ones most closely qualified to comment on the needs of their patients. By allowing for patients and their care providers to together petition against unjust denials, the policy aims to build trust in the healthcare system and improve patient outcomes.

Others may argue that the use of predictive artificial intelligence should be outright prohibited. Our policy takes a more agnostic stance toward the nature of predictive AI. If well-regulated as to control for the number of unjustified denials, use of AI could support insurance companies in freeing up resources to shift to client support and communication. Furthermore, it

is safer to regulate predictive AI in insurance—requiring companies to disclose their use such that further knowledge can be gained on the efficacy of these programs—than it is to have companies use them behind closed doors. Thus, this policy works at the balance of creating safeguards for patients and the reality of AI usage by major insurance companies.

Major Constituencies

- **Patients:** Patients benefit from the framework, as they obtain more information and transparency from their insurance companies and can better advocate for themselves. There may be additional stress in their processes of petitioning for the care they need, but this arises from an opportunity that they previously wouldn't have had.
- **Physicians:** Physicians take on an additional responsibility through this policy, as they may be asked to contribute to petitioning on behalf of their patients. However, this is reasonably in line with physicians' existing responsibility to create a plan of care for their patients and ultimately allows them to follow through on treatments while decreasing the likelihood of losing patients to insurance denials.
- **Insurance Companies:** Insurance companies take on the greatest share of accountability from this policy. By creating stricter regulations on insurance companies' AI policies and review teams, the framework increases the resources that companies will have to put into safeguards if they choose to use predictive analytics. Ultimately, however, this policy will serve to increase trust between insurance companies and their clients and protect companies from unintended adverse effects on the populations they serve.

Conclusion

The health insurance space is a dangerous context in which to experiment with the potential of predictive AI. While using artificial intelligence to make predictions about patient outcomes could save insurance companies time and resources, the cost for incorrect predictions is the health of patients across the country. We can mitigate these risks by improving national regulations for health insurance companies, building on the work of California's SB 1120 to promote transparency and human in the loop decision making. This framework can be enacted by Congress as an amendment to ERISA with established quality metrics managed by the FTC, such that citizens on federal healthcare plans can receive the same protections as citizens on state plans. By creating safer guardrails for predictive technology, especially relating to the ability to petition against denials, we can move toward a more equitable healthcare system in a world with artificial intelligence.

References

American Hospital Association. *Skyrocketing Hospital Administrative Costs, Burdensome Commercial Insurer Policies Are Impacting Patient Care*. American Hospital Association, 10 Sept. 2024, www.aha.org/guidesreports/2024-09-10-skyrocketing-hospital-administrative-costs-burdensome-commercial-insurer-policies-are-impacting.

California Department of Insurance. "SB 1120-1 Guidance: Use of Artificial Intelligence, Algorithms and Other Software Tools in Utilization Management." *California Department of Insurance*, 2025, https://www.insurance.ca.gov/0250-insurers/0500-legal-info/0200-regulations/HealthGuidance/upload/SB-1120-1-Guidance-Use-of-Artificial-Intelligence-Algorithms-and-Other-Software-Tools-in-Utilization-Management.pdf.

California Legislature. *Senate Bill No. 1120, Health Care Coverage: Utilization Review (SB 1120, Becker).* Stats. 2024, ch. 879. LegiScan, https://legiscan.com/CA/bill/SB1120/2023.

Centers for Medicare & Medicaid Services. "Appeals in Medicare Health Plans." *Medicare.gov*, 2024, www.medicare.gov/providers-services/claims-appeals-complaints/appeals/medicare-health-plans/.

Centers for Medicare & Medicaid Services. "Internal Appeals." *HealthCare.gov*, 2025, www.healthcare.gov/appeal-insurance-company-decision/internal-appeals/.

KFF. "Medicare Advantage Insurers Made Nearly 50 Million Prior Authorization Determinations in 2023." *KFF*, 28 Jan. 2025, www.kff.org/medicare/issue-brief/medicare-advantage-insurers-made-nearly-50-million-prior-authorization-determinations-in-2023/.

Kirzinger, Ashley, Isabelle Valdes, Alex Montero, Liz Hamel, and Mollyann Brodie. "5 Charts About Public Opinion on the Affordable Care Act." *KFF*, 15 May 2024, updated 3 Oct. 2025, www.kff.org/affordable-care-act/5-charts-about-public-opinion-on-the-affordable-care-act/.

Office of Inspector General, U.S. Department of Health and Human Services. *Some Medicare Advantage Organization Denials of Prior Authorization Requests Raise Concerns About Beneficiary Access to Medically Necessary Care*. Apr. 2022, oig.hhs.gov/reports/all/2022/some-medicare-advantage-organization-denials-of-prior-authorization-requests-raise-concerns-about-beneficiary-access-to-medically-necessary-care/.

Sedig, Laura. "What's the Role of Autonomy in Patient- and Family-Centered Care When Patients and Family Members Don't Agree?" *AMA Journal of Ethics*, vol. 18, no. 1, Jan. 2016, pp. 12–17. *DOI.org*, https://doi.org/10.1001/journalofethics.2017.18.1.ecas2-1601.

Surescripts. *Healthcare Professionals Highlight Medication Prior Authorization Challenges & Solutions (Data Brief).* Surescripts, Aug. 2025, surescripts.com/prior-authorization-challenges-data.

United States Senate Permanent Subcommittee on Investigations. *Refusal of Recovery: How Medicare Advantage Insurers Have Denied Patients Access to Post-Acute Care for Vulnerable Seniors*. Majority Staff Report, 17 Oct. 2024, www.blumenthal.senate.gov/newsroom/press/release/senate-permanent-subcommittee-on-investigations-releases-majority-staff-report-exposing-medicare-advantage-insurers-refusal-of-care-for-vulnerable-seniors.

U.S. Department of Labor. "ERISA." *U.S. Department of Labor*, www.dol.gov/general/topic/health-plans/erisa.