

Artificial Intelligence in Health Care: Spotlight on Health Insurers

Artificial intelligence (AI) is rapidly becoming embedded in the American health insurance system, influencing how insurers evaluate medical necessity, process prior authorizations, and determine whether patients can access care. Yet despite the potentially life-altering consequences of these technologies, patients, providers, and policymakers have remarkably little information about how they operate or which companies are involved. And the information we do have about the consequences of using AI to make care decisions - formerly made by medical professionals - reveals patient harm and increased costs.

This report reveals previously undisclosed claims of integration of Palantir Technologies' products by major health insurers, including Aetna, Blue Shield of California, Cigna, and Humana. Documents obtained through a [Freedom of Information Act request by the Electronic Frontier Foundation](#) claim the use of Palantir technology by these insurers for purposes including "decision support," raising questions about what patient information these systems access and whether they influence coverage decisions. This report also examines the involvement of Palantir and other technology companies in Medicare's new AI-supported prior authorization program, WISeR, where early operational data reveal substantial delays in coverage decisions, low approval rates, and backlogs that threaten timely patient care and provider reimbursement. These findings come as a national survey shows that 71% of responding individual major medical insurers are using or exploring AI for utilization management, while 77% of insurers responding to a disclosure question do not inform providers how or when AI is used.

Beyond current practices, this report examines patent filings from major health insurers that reveal plans to automate increasingly complex aspects of coverage determination and appeals. UnitedHealth Group subsidiaries have pursued technologies capable of generating prior authorization requests, approving or denying coverage, automatically appealing denials, and reviewing appeals, potentially creating a system in which AI operates on multiple sides of the same coverage dispute. Other patents describe using historical appeal behavior to predict which denials will be overturned and using AI to discourage emergency department visits. While these patents do not establish that the technologies have been deployed, they reveal the direction insurers are exploring and the potential for AI to reinforce existing barriers to care.

Taken together, the findings expose an expanding network of health insurers, technology corporations, and utilization management vendors operating with little public transparency or independent accountability. Greater disclosure and oversight are urgently needed, but transparency alone cannot address the underlying problem: AI is being introduced into a profit-driven health insurance system that already creates unnecessary administrative complexity and barriers to medically necessary care. Rather than using increasingly sophisticated technologies to manage these obstacles, policymakers must address the financial incentives that created them in the first place.

1. Review of previously reported use of AI by health insurers

AI and machine learning (ML) are increasingly being incorporated into health insurers' utilization management and coverage determination processes, raising questions about how these technologies influence access to medically necessary care. According to the National Association of Insurance Commissioners' (NAIC) Health AI/ML Survey, 71% of responding companies selling individual major medical insurance reported using or exploring AI/ML for utilization management practices, while 12% reported using AI/ML to deny prior authorization requests. Despite the growing role of these technologies in coverage decisions, transparency remains limited. Among companies responding to a question about disclosure, 77% reported that they do not disclose to providers or physicians how and when AI/ML is used by the insurer or a third-party vendor, compared with just 23% that reported doing so.¹

These findings demonstrate not only the growing use of AI in health insurance but also a substantial information imbalance between insurers and the patients and providers affected by their decisions. Even when a physician receives a coverage determination, they may have no way of knowing whether an AI system contributed to that decision, what information it considered, or whether the system was evaluated for accuracy or bias. This information is key when appealing an insurer denial.

After a prior authorization or claim is denied, patients and providers have the ability to appeal the decision for a re-review by the insurer. However, a patient's likelihood of appealing a denial is not necessarily a reflection of whether the denied care was medically necessary. Patients may face barriers to appealing, including limited health literacy, language barriers, financial constraints, or a lack of assistance navigating the appeals process. In fact, in the Medicare Advantage program in which insurers deny over 4 million prior authorizations per year, only 11.5% of those denials are ever appealed, but over 80% of those appeals were successful.²

Congressional investigations have revealed that insurers are exploring applications of AI that assess the probability of appeals being filed and succeeding. In its October 2024 investigation of Medicare Advantage prior authorization and post-acute care, the Senate Permanent Subcommittee on Investigations (PSI) reported that a UnitedHealthcare working group met in December 2022 to explore using AI and ML to predict which post-acute care denials were likely to be appealed and which appeals were likely to be overturned.³ Although the stated objective included improving the accuracy of coverage decisions and avoiding unnecessary appeals, the distinction between predicting whether a denial is medically justified and predicting whether a patient will challenge that denial is consequential. Because appeal behavior is affected by factors like socioeconomic status, an algorithm that relies on historical appeal behavior could reproduce inequities already present in the health insurance system.

Footnotes

1. <https://content.naic.org/sites/default/files/inline-files/Health%20Survey%20Report%20-%20FINAL%205.9.25.pdf>
2. <https://www.kff.org/medicare/medicare-advantage-insurers-made-nearly-53-million-prior-authorization-determinations-in-2024/#6e420acb-2fc1-4707-8689-ac19594e493a>
3. <https://www.hsgac.senate.gov/wp-content/uploads/2024.10.17-PSI-Majority-Staff-Report-on-Medicare-Advantage.pdf>



The working group's discussions raise the possibility that efforts to correct inappropriate denials could be concentrated on cases most likely to generate appeals, rather than on all cases in which coverage was improperly denied. The report does not establish that UnitedHealthcare ultimately deployed such a model or used appeal propensity to deny care, but it documents the company's consideration of patient and provider appeal behavior as a predictive input in its coverage decision-making processes.

Investigative reporting by ProPublica and The Capitol Forum has similarly raised questions about the use of automated decision-making by eviCore, a major utilization management company owned by The Cigna Group.⁴ In response to scrutiny of its prior authorization practices, eviCore emphasized that physicians can appeal denied requests and that the company routinely monitors appeal outcomes as part of its continuous quality-improvement efforts to ensure accurate and timely medical-necessity determinations. That acknowledgment is significant because appeal outcomes provide a potentially valuable dataset for evaluating and training systems that predict which initial denials are likely to be reversed. Monitoring overturned denials can help identify errors and improve patient access. However, historical appeal data can also reveal which patients and providers are most likely to challenge an adverse determination. Whether insurers or their vendors use that information to improve the clinical accuracy of decisions, predict appeal behavior, or influence the handling of future requests remains largely unknown.

Taken together, the NAIC survey, Senate investigation, and investigative reporting establish that AI/ML is already part of insurers' utilization management activities, that transparency about its use remains limited, and that insurers and their vendors collect information about appeals and coverage decisions that could support increasingly sophisticated predictive systems. These findings expose a playbook in which insurers are developing and using technologies that influence access to medical care while their bottom line is increasing profits for shareholders, not patient care. Further, they are doing so without routinely disclosing how those technologies function or providing sufficient information for the public to determine whether the AI is safe.

2. Problematic Partnerships: Big Insurance and Big Tech

As health insurers utilize AI more frequently, they are partnering with some of the nation's largest technology corporations. Yet many of these relationships, including the technologies being used and the purposes they serve, remain largely unknown to the public.

For example, documents obtained through a Freedom of Information Act request by the Electronic Frontier Foundation (EFF) identify the use of Palantir Technologies' systems for "AI-driven decision-making solutions" by Aetna, Humana, Cigna, and Blue Shield of California.⁵

Footnotes

4. <https://www.propublica.org/article/evicore-health-insurance-denials-cigna-unitedhealthcare-aetna-prior-authorizations>

5. <https://www.eff.org/document/wiser-foia-response-rfis-caps-and-applications> (pages 123, 264)



The documents, which were part of the records surrounding the Wasteful and Inappropriate Service Reduction (WISeR) model launched by the Centers for Medicare and Medicaid Services (CMS) in January 2026 also include an application for Zyter, one of the vendors ultimately chosen for program participation. In its application, Zyter wrote that they are applying as a consortium of companies including Palantir and that Palantir infrastructure “securely ingests, processes, and audits structured and unstructured data, including OCR records, chart notes, ABNs, and decision history.” The company additionally touts that Palantir “provides mission-critical software infrastructure for CMS, HHS, the VA, and multiple national payers.” The application notes that Palantir provides “secure data integration and decision support” for Humana, Aetna, and others, raising questions about whether the technology companies are directly involved in or profiting from denial of care.⁶

Given that these partnerships already exist, health insurers should immediately release to the public all vendor agreements involving the use of patient data. Without these documents we cannot answer:

- Which technology companies provide AI systems or related infrastructure to health insurers;
- What functions those technologies perform, and if they involved in prior authorization, claims review, or coverage determinations;
- How protected health information shared from health insurers is used;
- When and how are patient records combined with information from other insurers or data sources;
- What measures exist to prevent algorithms from reproducing health disparities or generating inaccurate coverage recommendations; and
- Who is responsible for evaluating the accuracy of these systems and correcting errors that affect patient care.

Though partnerships and data sharing between Big Insurance and Big Tech have already been initiated, and therefore, it is important to use transparency to assess the damage done by those partnerships thus far, it is equally important for policymakers to prevent these partnerships in the future that put patients’ health and data privacy at risk.

3. A CMS Model Provides a Window into AI’s Human Consequences in Health Care

Health insurers operate behind a veil of secrecy, rarely sharing operational data such as prior authorization outcomes, claims denials, and more. This makes evaluating the effects of the use of AI by health insurers nearly impossible. The WISeR model provides a unique window into the impacts of AI on patient care and provider solvency since the model runs in the traditional Medicare program, a federal program with data reporting requirements. The WISeR model added a prior authorization requirement for 17 new services in traditional Medicare and contracted with private AI companies to process those prior authorizations. The hypothesis behind the model is that AI will speed up the prior authorization process and streamline things for patients, providers, and CMS.

Footnotes

6. <https://www.eff.org/document/wiser-foia-response-rfis-caps-and-applications> (pages 123-132)



Close examination of the WISeR model's performance shows failure of AI-based prior authorization to meet the requirements for decisions set by CMS. As reported in Health Care uncovered, as of March 30, 2026, Virtix, the WISeR contractor operating in Washington, had issued 6,096 prior authorization decisions, of which only 2,863 (47%) were affirmed or approved, leaving 3,233 decisions (53%) not approved or affirmed.⁷

Separately, the program's prior authorization and prepayment review process was experiencing substantial backlogs. In January of 2026, 30% of prior authorization requests sent to Virtix and Zyter, two WISeR AI vendors, took more than five days to be approved or denied. This is in spite of CMS setting a requirement for decisions to be rendered within 72 hours of the request and the point of the model being that AI can speed up prior authorization review.⁸ Another AI vendor, Genzeon had more than 700 cases older than four days waiting for decisions in February 2026.⁹ Physicians surveyed as part of the model expressed extreme frustration at the delays with one saying, "We have patients calling our offices crying in pain because their procedures are being delayed while awaiting approvals or guidance tied to this model. A 3-4 day delay for necessary pain procedures is already difficult for vulnerable patients, but when providers cannot obtain answers for weeks, the situation becomes unacceptable."¹⁰

As of March 26, Noridian, the Medicare Administrative Contractor overseeing Arizona and Washington, had referred 4,373 claims to Zyter, Arizona's WISeR contractor, for prepayment review but had received only one decision in return. These delays were particularly concerning because providers who chose not to seek prior authorization were still subject to prepayment review, potentially delaying reimbursement after care had already been delivered. Provider surveys and complaints documented malfunctioning portals, delayed procedures, and patients waiting in pain for medically recommended treatment.¹¹

Further, it is difficult to know who is behind these AI companies now determining life-impacting medical care for patients in this model. For example, as described above the WISeR vendor Zyter described in its application that it was applying as a "consortium [that] unites industry leaders, including an AI-enabled utilization management and care management platform (Zyter), a secure federal-grade software infrastructure provider (Palantir), and nationally trusted clinical review organizations (AllMed and MRIOA), to transform how prior authorization and pre-payment reviews are delivered for Medicare beneficiaries."¹² In the same way Palantir has never been publicly disclosed as a partner of large health insurers, until now, it was not disclosed to patients, providers, or policymakers that Palantir, AllMed, and MRIOA are involved in the WISeR model. Virtix Health is also cloaked in a veil of secrecy with little known about them and a disclosure only made public through EFF's freedom of information request that they are in fact the same company as private-equity-owned CorroHealth.¹³

Footnotes

7. <https://healthcareuncovered.substack.com/p/the-secret-plan-to-hand-your-medicare>
8. <https://www.eff.org/document/wiser-foia-response-combined-records>
9. <https://www.eff.org/document/wiser-foia-response-rfis-caps-and-applications> (page 14)
10. <https://www.eff.org/document/wiser-foia-response-feedback-survey-1>
11. <https://www.eff.org/document/wiser-foia-response-combined-records>
12. <https://www.eff.org/document/wiser-foia-response-rfis-caps-and-applications> (pages 123-132)
13. <https://www.eff.org/document/wiser-foia-response-rfis-caps-and-applications> (page 92)



Although big insurance companies are not directly involved in the WISeR model, their influence is clear: in its Innovation Center Investment Plan (ICIP), CMS lists “most providers/suppliers, professional organizations, and consumer advocacy organizations” as opponents of prior authorization, while the expected supporters of the pilot were “groups representing payers (i.e., America’s Health Insurance Plans (AHIP)).”¹⁴ An examination of the 60 unsuccessful WISeR bids reveals that four bids came from subsidiaries of major insurers: Cigna, Elevance, UnitedHealth Group, and Blue Cross and Blue Shield of Alabama.¹⁵ And CMS’s internal plan for expanding the program cited Medicare Advantage, managed primarily by big insurers, as having used prior authorization as “an effective method for reducing spending in other contexts,” despite public reporting and data that Medicare Advantage’s prior authorization is often a disaster for patient care.¹⁶

The contracts between CMS and the WISeR vendors similarly show a pronounced lack of transparency or oversight. The participation agreements, which were standardized across all six vendors, allowed participants to “reuse original or derivative data without prior written authorization from CMS” for “improving health or reducing health care costs through the Prior Authorization Process,” which is alarmingly vague. In the documents released, there were no public limitations on the use of WISeR data for AI training, and while the contract requires vendors to “return or destroy” WISeR data when the contracts end, it allows vendors to keep data beyond that point if it’s “infeasible.”¹⁷

Perhaps the most concerning aspect of the WISeR model is the use of seniors and people with disabilities as test subjects for this new AI-driven care decision model. Using real patients on Medicare to test the use of AI in making medical decisions has caused documented harm across the country for traditional Medicare beneficiaries seeking these treatments that were previously automatically approved. Reports have shown patients seeking out-of-state care because of onerous and incorrect WISeR determinations, and patients left waiting in pain as these vendors test new algorithms on them. This model is not akin to other CMS models that test new payment methods. Instead, it puts real obstacles to care in place and causes painful collateral damage as the AI companies learn how to perform prior authorizations.

4. The Future of AI in Health Insurance is Worrisome

Aside from the publicly disclosed or unearthed use of AI by health insurers, there are hints found in patent applications from insurers that AI use may be more widespread and that insurers are planning to expand its use. For example, eviCore has an active patent on a technology that uses historical appeals data to preemptively approve prior authorizations that were originally going to be denied if they have a high likelihood of being overturned on appeal.¹⁸ On its face this may seem positive. However, research shows inequities in the likelihood that a patient appeals a denial, mirroring a lot of health inequities in our system.¹⁹ Using data tainted by health inequities to build predictive models only furthers those inequities.

Footnotes

14. <https://www.eff.org/document/wiser-foia-response-combined-records> (page 38)
15. <https://www.eff.org/document/wiser-foia-response-rfis-caps-and-applications> (page 250); <https://www.eff.org/document/wiser-foia-response-additional-applications-and-contracts>, (pages 212, 478, 314)
16. <https://www.eff.org/document/wiser-foia-response-combined-records> (page 30)
17. <https://www.eff.org/document/wiser-foia-response-participation-agreements> (pages 51, 126, 201, 276, 351, 426)
18. US11848097B2
19. <https://www.macpac.gov/wp-content/uploads/2024/03/Chapter-2-Denials-and-Appeals-in-Medicaid-Managed-Care.pdf>.accreport.pdf



Aetna pursued a patent, before abandoning it, that would use an AI algorithm to predict if a patient was likely to visit an emergency room in the coming weeks, predict whether that emergency room visit was necessary, and if the algorithm determined it was not necessary, send a notification to the patient's phone telling them they did not need to go to the emergency room.²⁰ While it is positive that it was abandoned, the basis of the patent and language in the application claiming the purpose of this technology was to decrease health care costs provided a worrying look into how insurers want to use AI and the care obstacles they could put in a patient's way.

These patents and the use of AI in health care is also driving up system costs and leading to a system in which AI is leading actions on all sides of the transaction. For example, UnitedHealth Group through its Optum division has applied for separate patents using AI that 1) generates a prior authorization request by mining data from provider notes in a patient's chart, 2) approve or deny the prior authorization and transmit the decision, 3) automatically appeal a prior authorization denial, and 4) review and recommend decisions for appealed prior authorization or claims denials.²¹ This creates a potential process in which AI machinery ping-pongs requests, decisions and denials back and forth, adding complexity and cost to the health care system.

Rather than eliminate or decrease the barriers they created, these patents show that insurance companies are primarily using AI to accelerate timelines in the existing system - speeding up prior authorization requests without addressing the unnecessary burdens the system puts on patients and providers. AI cannot fix the problems caused by excessive prior authorization requirements from insurers; it only serves to accelerate the problem. These examples only scratch the surface of the AI capabilities that health insurers are envisioning or developing, all of which will likely make care more costly and complex.

5. The Need for Reform

The evidence reviewed in this paper reveals a significant gap between the expanding capabilities of AI in health insurance and the guardrails in place to protect patients from potential harms of these technologies.

Insurers are already using AI in utilization management. Technology corporations are becoming involved in the infrastructure supporting insurance administration. Early operational data from WISer raise concerns about the performance of AI-supported coverage review. Patent applications also show that insurers are pursuing technologies that can automate increasingly complex aspects of coverage determination and appeals. Yet critical information remains unavailable, and the use of AI in care decisions is largely unrestrained.

Footnotes

20. US20210407667A1

21. US20240331835A1; US20260087075A1; US20260134480A1; US20250148542A1



Protecting patients' well-being during the rise of AI should include the following:

- **Preventing AI from determining patient care:** Insurers should not be able to use AI in any care-altering decisions, including prior authorization, claims review, and other tasks.
 - Insurers will likely claim AI allows them to review the large amounts of prior authorization requests they get each day. The true solution to that problem is for insurers to stop requiring prior authorization for thousands of services and create a system that puts patient needs over profits.
- **End the Wasteful and Inappropriate Service Reduction pilot:** Medicare patients should not be denied care by AI, and the model's structure incentivizes delays and denials for AI companies to profit from. This pilot should not run until 2031, and CMS should immediately release a plan to wind it down.
- **Requiring proof of safety:** AI companies and insurers who wish to use AI in their plans for tasks other than care decisions must prove that the AI is safe, accurate, and does not replicate health disparities or overtake judgment that should come from a medical professional.
- **Disclosure of AI use:** Insurers should identify where AI or ML is used in administrative tasks for their plans.
- **Vendor transparency:** Insurers should disclose the companies providing AI technologies and related infrastructure, the functions those companies perform and the nature of their access to patient information.
- **Operational reporting:** Regulators should have access to meaningful information about approval rates, denial rates, appeals, overturns, review times, and system performance.
- **Accountability for decisions:** The use of outside technology vendors should not diminish an insurer's responsibility to make coverage determinations consistent with applicable clinical and legal requirements.
- **Stop profiteering in health care:** The use of AI to attempt to streamline the complexity and lower the costs of our health care system should not be seen as a solution to the byzantine system created by profit-driven entities such as health insurers and private equity firms. The true solution is removing the profit motive from health care.



Conclusion

The evidence reviewed in this paper raises serious concerns about patient harm from the rapid expansion of AI in health insurance. These technologies are becoming embedded in an already complex and opaque system in which patients and providers frequently struggle to understand, challenge, and overturn coverage decisions.

Previously reported insurer practices demonstrate that AI is already being incorporated into utilization management. Newly uncovered technology relationships raise questions about the growing involvement of outside corporations in health insurance administration. Early WISeR operational data illustrate the potential consequences of delayed or inefficient coverage review. And insurer patent applications reveal ambitions to automate additional stages of the coverage and appeals process. They also reveal how little information is available to independently determine whether these technologies are improving access to care or reinforcing the financial and administrative incentives that already shape insurer decisions.

Patients should not have to wonder whether an undisclosed algorithm influenced their coverage determination, whether their medical information was shared with an outside technology company, or whether an automated system was evaluated for accuracy before affecting their care. As AI becomes increasingly integrated into health insurance, it should not be allowed to make care-altering decisions, especially when it was designed by a company driven by profit rather than patient health.

