Institute for Youth in Policy

5214F Diamond Heights Blvd #3055
San Francisco, CA 94109
+1 (715) 469-6884
fellowship@yipinstitute.org
www.yipinstitute.org/fellowship/

Fellowship Capstone | Policy Brief

AI Assisted Utilization in U.S. Health Insurance: Defining Human Review

Audrey Huang

I. Executive summary

In 2022, Cigna physicians denied more than 300,000 payment requests using algorithmic processing, spending an average of 1.2 seconds per case, according to a ProPublica investigation published March 25, 2023. One physician, Dr. Cheryl Dopke, denied roughly 60,000 claims in a single month. Patients received form denial letters with no individualized rationale, no clear decision-maker, and no efficient path to appeal. When automated triage shapes coverage outcomes but no enforceable standard defines what "meaningful human review" requires, accountability dissolves. The efficiency logic that makes AI attractive to insurers becomes dangerous when applied to coverage decisions without clear standards for oversight. This brief asks: What should "meaningful human review" legally require before an insurer can use AI to deny, delay, or reduce coverage? It proposes an enforceable operational standard across five domains: disclosure, reviewer obligations, contestability and timelines, auditing and reporting, and enforcement.

II. Overview

This brief examines AI-assisted utilization review in U.S. health insurance, with a defined focus on prior authorization determinations and claims denials. It proceeds in three steps: first, describing the current utilization review landscape and the incentives driving high-throughput decision-making; second, identifying how accountability gaps emerge when automated triage pairs with under-specified clinician oversight; and third, proposing an operational legal standard across five domains: disclosure, reviewer obligations, contestability and timelines, auditing and reporting, and enforcement.

A. Relevance

Coverage determinations in utilization review function as time-sensitive clinical decisions. Delays and denials can interrupt diagnostic pathways, postpone treatment, and impose financial pressure that leads patients to forgo care entirely. The introduction of AI-driven triage heightens these risks in four ways.

1. Automation expands throughput, increasing the volume of adverse determinations patients must contest through slow, resource-intensive appeals processes.
2. Algorithmic recommendations create deference effects, where reviewers default to the system's output, making "human review" nominal rather than substantive.
3. Proprietary tools and opaque workflows undermine procedural fairness: patients receive adverse decisions without a clear explanation tied to their record or an accountable decision-maker.

4. These systems optimize for what is easiest to quantify. A Kaiser Family Foundation report from October 5, 2023 found that insurers' algorithmic assessments frequently rely on standardized inputs like age or blood pressure while failing to incorporate individualized context such as injury severity or a physician's stated rationale for care.

If insurers may deploy AI to influence high-impact determinations without defined duties of review, transparent justification, and enforceable accountability, the system permits health-affecting decisions at industrial speed while externalizing error costs onto patients.

III. History

A. Current Stances

Utilization review expanded as insurers sought to manage rising expenditures and promote consistency in coverage decisions. As determination volumes increased, insurers had strong incentives to standardize and accelerate review. Algorithmic triage promised administrative efficiency, greater uniformity, and scalable management of high caseloads.

These same features, however, can weaken the link between the individual patient record and the adverse decision. When review is structured around rapid confirmation of system outputs, documentation may perform the appearance of reasoning without supplying record-based justification.

Three groups currently hold distinct positions on this issue. Insurers and vendors commonly frame AI-assisted review as modernization of an overburdened administrative process: improving timeliness, consistency, and cost containment. Physicians and patient advocates emphasize that utilization review already produces clinically meaningful delays and denials, and that algorithmic integration risks amplifying these harms by scaling adverse decisions and increasing the difficulty of meaningful challenge. Regulators have increasingly prioritized transparency, timeliness, and process standardization in utilization review, but these efforts leave a central gap unaddressed: absent a clear definition of meaningful human review, automation may still drive outcomes while accountability remains diffuse.

IV. Policy Problem

A. Stakeholders

Patients and families are the most directly affected. Coverage determinations function like time-sensitive clinical gates, determining whether care happens now, later, or not at all. These decisions also determine financial exposure: whether a patient absorbs a charge, delays treatment, or abandons care. In California, my home state, Medi-Cal beneficiaries, who are disproportionately low-income and Latino, face additional barriers when algorithmic denials arrive in English without plain-language explanations.

Clinicians and provider organizations bear the administrative burden of contesting adverse decisions and supplying documentation for services already rendered. AI-assisted triage may scale the volume and speed of denials that

providers must appeal, degrading clinical time and patient continuity.

Insurers benefit from scalable standardization and cost control, but they also carry reputational and legal risk when automation drives outcomes without a defensible review process, especially if "human review" is reduced to rubber-stamping.

Employers and plan sponsors face a cost shift rather than cost savings: when delayed care becomes emergency care, workforce absenteeism rises and downstream costs increase. Regulators and policymakers balance consumer protection, cost containment, and enforceability.

B. Risks of Indifference

If policymakers treat this as a standard technology efficiency issue, the likely outcome is industrialized adverse determinations with diffuse responsibility. No one is clearly accountable for the final decision. Patients receive adverse outcomes without an explanation they can act on, without timely appeal pathways that match clinical reality, and without a clearly accountable reviewer who can be questioned, corrected, or audited.

Indifference also risks a cost shift rather than cost control. When inappropriate denials delay care, patients may deteriorate and re-enter the system at higher cost through urgent care visits, emergency department visits, and avoidable complications. California's Department of Managed Health Care has flagged this pattern repeatedly in its annual report on prior authorization denials.

C. Nonpartisan Reasoning

This issue should be nonpartisan because it is fundamentally about governance, due process, and medical integrity. Three areas demonstrate this.

1. Rule of law and accountability: If an insurer can deny care at scale while the rationale is opaque and responsibility is difficult to locate, the system violates basic expectations of reason-giving for high-impact decisions. This is a procedural fairness concern, not a partisan one.
2. Market function and consumer trust: Transparent standards for meaningful review reduce arbitrary outcomes and stabilize trust in insurance as a product. People pay premiums expecting predictable rules. When AI makes those rules unpredictable, the market breaks down.
3. Administrative efficiency without patient harm: Clear standards can preserve efficiency while preventing "automation theater," where the process looks compliant but functionally bypasses individualized review. Both fiscal conservatives and patient advocates gain from eliminating wasteful downstream costs.

V. Tried Policy

Current policy efforts have largely targeted process standardization, including timelines, disclosure basics, and appeals infrastructure, but they fail to define the central term this brief focuses on: what "meaningful human review" must actually require. The current landscape relies

on three existing mechanisms, each with significant gaps.

1. Appeals rights and external review frameworks give patients a route to challenge adverse benefit determinations, but often only after delays have already caused harm. The harm is irreversible by the time the appeal succeeds.
2. Prior authorization reforms at the state and federal levels improve timelines and require some justification, but do not reliably address whether the clinician reviewer evaluated the individual record versus confirming an algorithmic output. California's AB 2348 (2022), for example, improved prior authorization timelines but left reviewer obligations under-specified.
3. Transparency and reporting initiatives increase visibility into denial rates or process metrics, but still leave "human review" undefined and therefore unenforceable.

The result is a persistent gap: insurers may comply with procedural requirements while still operating review systems where accountability is diffuse, explanations are generic, and the human role is nominal, especially when algorithmic triage is used to reduce human decision volume.

VI. Policy Options

Option 1: Baseline Standard

Define "Meaningful Human Review" as an enforceable duty across five domains.

Disclosure: Require insurers to disclose when AI or algorithmic tools materially influenced a denial, delay, or reduction. Require patient-facing explanation that identifies (a) the clinical guideline or basis used, (b) what key record facts were considered, and (c) how the decision connects to the patient's specific case.

Reviewer obligations: Require that a licensed clinician reviewer attest that they reviewed the relevant portions of the patient record and could state, in plain language, why the request does not meet medical necessity or coverage criteria. Ban "signature-only" review for adverse determinations driven by automated triage.

Contestability and timelines: Require expedited review pathways when delay risks harm. Require clear timelines that match clinical reality, distinguished between urgent and routine cases, plus confirmation that appeals are reviewed by a qualified clinician not bound to deference to the original output.

Auditing and reporting: Require reporting of denial rates, overturn rates on appeal, average time-to-decision, and stratification by service category. Require documentation retention sufficient for audit, including inputs considered, tool involvement, and reviewer identity.

Enforcement: Establish penalties for patterns of noncompliance, including repeated inadequate explanation, repeated rubber-stamp signatures, and abnormal denial-to-overturn ratios. Give regulators authority to require corrective action plans.

Option 2: Stronger Standard

In addition to Option 1, add a record-citation requirement and an anti-deference rule. Any adverse determination must cite specific record elements, such as diagnosis, clinical findings, prior treatments, and physician rationale, that support the decision. Generic boilerplate is noncompliant. Require insurers to implement reviewer

workflows that reduce automation deference: reviewers must document at least one individualized factor not contained in standardized inputs.

Option 3: Comprehensive Oversight Model

In addition to Options 1 and 2, require periodic external audits of AI-assisted utilization review programs. These audits would include error pattern detection for services with unusually high denial and overturn rates, bias and fairness checks, documentation compliance checks, and validation that meaningful human review is occurring through workflow evidence and record-citation sampling. Expand patient contestability by requiring a clear pathway for patients to request the decision basis in plain language, with faster external review triggers when internal appeals are slow and the service is time-sensitive.

VII. Conclusions

AI-assisted utilization review is a governance problem that currently operates without sufficient accountability. Coverage decisions are made at industrial speed, responsibility is difficult to locate, and justification is difficult to contest. The core gap is definitional: if "meaningful human review" is not operationally defined, "human review" becomes a checkbox rather than a safeguard.

This brief recommends adopting an enforceable standard for meaningful human review across five domains: disclosure, reviewer obligations, contestability and timelines, auditing and reporting, and enforcement. In practical terms, the policy aim is straightforward: if an insurer is going to deny care, it must be clear who decided, what they reviewed, why they decided it, and how a patient can challenge it in time for the decision to still matter.

This is a call to prevent a system where automation shapes outcomes while accountability dissolves. For patients in California and across the country, the difference between a compliant denial and a reviewed denial can be the difference between receiving care and going without it.

Acknowledgment

The Institute for Youth in Policy sincerely thanks Patrick Pickren, Tiffany Li, Nikki Wu, Asher Cohen, Taylor Beljon-Regen, Paul Kramer, and other contributors for developing and maintaining the Fellowship Program within the Institute.

References

[1] Marshall Allen, "How Cigna Saves Millions by Having Its Doctors Reject Claims Without Reading Them," ProPublica, March 25, 2023. https://www.propublica.org/article/cigna-pxdx-medical-health-insurance-rejection-claims

[2] Kaiser Family Foundation, "Prior Authorization and Utilization Management," October 5, 2023. https://www.kff.org/tag/prior-authorization/

[3] California Department of Managed Health Care, Annual Report, 2022. https://www.dmhc.ca.gov/Portals/0/Docs/DO/2022AR.pdf

[4] California AB 2348 (2022), Prior Authorization Reform Act. https://leginfo.legislature.ca.gov/faces/billNavClient.xhtml?bill_id=202120220AB2348

[5] Obermeyer, Z. et al., "Dissecting Racial Bias

in an Algorithm Used to Manage the Health of Populations," Science, 2019. https://www.science.org/doi/10.1126/science.aax2342

[6] HHS Office of Inspector General, "Medicare Advantage: Insurers Used Strategies That May Have Delayed or Denied Beneficiaries' Access to Care," 2022. https://oig.hhs.gov/oei/reports/OEI-09-18-00260.asp

[7] CMS Final Rule on Prior Authorization, CMS-0057-F, January 2024. https://www.cms.gov/priorities/key-initiatives/burden-reduction/interoperability/policies-and-regulations/cms-0057-f

[8] American Medical Association, "2022 Prior Authorization Physician Survey," 2022. https://www.ama-assn.org/practice-management/prior-authorization/2022-ama-prior-authorization-survey

[9] Pasquale, Frank. The Black Box Society. Harvard University Press, 2015.