

REVENUE CYCLE

Automation in the revenue cycle: coding, denials, and prior authorization

ABSTRACT

Revenue-cycle automation is sold on statistics that frequently cannot be traced to a source. This paper audits the evidence behind coding accuracy, claim denials, prior authorization burden and payer-side algorithms, separating peer-reviewed measurement from vendor extrapolation, trade-association restatement and plaintiff allegation. It finds a small set of durable figures: the instability of the human coding standard, a measured Medicare Advantage denial rate, and a federal prior-authorization rule with fixed compliance dates. It also finds a large set of numbers that need a denominator before they mean anything.

Type Evidence and regulatory review **References** 28 **Reading time** 15 min

Last reviewed September 2026

1 Auditing a literature before reviewing it

Three statements about the medical revenue cycle circulate constantly in vendor materials and trade coverage: that most medical bills contain errors, that the industry standard for coding accuracy is 95 percent, and that automation is worth a specific number of billions of dollars. None of the three can be traced to a study. The failure mode is not fabrication but circulation: a survey is published, a trade association restates it, a trade publication cites both, and three citations begin to look like three sources.

Premier's 2024 trend alert is real measurement: 516 hospitals across 36 states, for calendar year 2022, finding nearly 15% of private-payer claims initially denied, 54.3% of those denials ultimately overturned and paid, and rework costing an average of \$43.84 per claim.¹ The American Hospital Association's 2025 *Cost of Caring* report states that hospitals spent \$26 billion in 2023 managing insurance claims, up 23% year over year, and that 70% of denied claims were eventually paid.² Those AHA figures restate the Premier survey. Citing both is the same 516 hospitals counted twice.

This paper puts one question to each major claim in revenue-cycle automation: what would a reader have to accept in order to believe it? The answer sorts the field into four categories that are routinely blurred together: measured, modeled, self-reported and alleged.

15.2%

complete agreement among three auditors coding the same 1,069 charts, 1995 guidelines

42%

of sampled 2010 Medicare E/M claims incorrectly coded, on federal review

17.7%

of Medicare Advantage claims denied at initial submission

13 hours

weekly physician *and staff* time on prior authorization, self-reported

WIDELY REPEATED FIGURES THAT DO NOT SURVIVE A CITATION CHECK

“80% of medical bills contain errors.” Traces only to a for-profit patient-billing advocacy firm, with no published methodology, sample, error definition or date. Consumer-finance media have recycled it for over a decade and it now reappears in marketing as though it were a study.

“The industry standard for coding accuracy is 95%.” There is no empirical basis for it. Professional-association reporting traces the number to the 1979 implementation of ICD-9-CM and to a loose analogy with the 95% *confidence interval* used in federal improper-payment sampling, which is a confidence level, not an accuracy target. Surveyed organizations in fact set thresholds at 95% (28% of them), 90% (19%) and 80–89% (16%), and practicing auditors recommended 80% as a baseline.³

“The algorithm had a 90% error rate.” This appears at paragraph 1 of a complaint filed against UnitedHealth Group.⁴ It is an unproven allegation in a pleading; no court has found it and no audit has produced it.

2 The human baseline for coding is unstable

Any claim that a system codes accurately is a claim about agreement with a reference standard. In evaluation and management coding, that standard has never been stable.

Three auditors, a faculty physician, a resident and a professional coder, independently coded 1,069 established-patient charts under the 1995 and 1998 HCFA documentation guidelines. Complete three-way agreement occurred in 15.2% of cases under the 1995 guidelines and 29.2% under the 1998

guidelines; overall concurrence among all three raters was 31.0% and 44.3% respectively.⁵

Six hundred randomly selected Illinois family physicians coded six hypothetical office-visit notes against a reference standard set by five expert coders. They agreed with it on 52% of established-patient notes and 17% of new-patient notes, with undercoding predominant for established patients and overcoding for new ones.⁶

Certified professional coders reviewing a sample of 2010 Medicare E/M claims for the HHS Office of Inspector General found 42% incorrectly coded, in both directions, and 19% insufficiently documented; 21% of Medicare E/M payments, \$6.7 billion, were improper (report OEI-04-10-00181, issued 28 May 2014).⁷ The report gives no inter-reviewer agreement statistic for its own reviewers.

A system reported as 90% accurate is therefore 90% concordant with whatever human or panel produced the labels. When independent expert reviewers agree completely on the same chart in fewer than a third of cases, concordance with one of them measures similarity to a particular reader, not correctness.

What changed in 2021, and what did not

The office and outpatient E/M rules were rewritten effective 1 January 2021. History and exam were removed from code selection and retained only “as medically appropriate”; the level is now chosen by medical decision making or by total time on the date of the encounter; code 99201 was deleted and prolonged-services code 99417 added; and medical decision making requires two of three elements covering problems, data and risk.⁸

An observational study of 303,547 providers across 389 organizations on a single EHR platform compared coding before and after the change. Level 3 visits fell while levels 4 and 5 rose, including a 22.6% relative increase in level 5 visits, and note length and EHR time showed no meaningful change.⁹ Simplifying the rule for choosing a code changed what was billed within months. It did not change what clinicians wrote.

3 Automated coding: what the research actually measures

The published literature on automated diagnosis coding is overwhelmingly a literature about one dataset: the MIMIC critical-care database of ICU discharge summaries. Three problems follow, and each alone breaks the inference from a benchmark score to ambulatory performance.

The labels are incomplete

MIMIC-III’s ICD assignments never underwent secondary validation. An experimental evaluation that built a silver-standard comparison found the most frequently assigned codes undercoded by up to 35%.¹⁰ Models scored on those labels are fitted to incomplete human coding, and a model that assigns a code the original coder omitted is penalized for it.

The metrics are not comparable across papers

A replication study that reproduced the major MIMIC-III and MIMIC-IV coding models found published results cannot be placed on a common scale. Macro-F1 had been computed suboptimally in prior work, and correcting it roughly doubles the score, while training splits and decision thresholds were handled inconsistently between papers.¹¹ A reported state-of-the-art F1 is not comparable to the previous one in the same table.

The population is wrong for the question

A systematic review of 73 automated ICD coding studies published between 2014 and 2024 found F1-micro improved by 6.4% on the full MIMIC-III task and by 10.2% on the top-50 code subset since 2019. Its central limitation is structural: MIMIC contains ICU patients only, with a correspondingly biased disease distribution, and the field suffers from what the authors call insufficient validation by clinical coders.¹²

Direct large-language-model performance has been benchmarked against human coders on inpatient notes in a preprint that has not been peer reviewed. Exact-match accuracy for ICD-10-CM extraction ranged from 1.4% to 15.2% across six models, with GPT-4 highest at 15.2% and 26.4% at category level; the authors concluded that current models perform poorly compared against a human coder.¹³

Nothing here supports unsupervised code assignment, and a MIMIC benchmark result is not evidence about ambulatory coding. An ICU discharge summary and a fifteen-minute follow-up visit differ in length, code distribution, documentation convention and payer rule. Citing MIMIC performance as evidence of ambulatory accuracy is a category error, not a conservative approximation.

4 Denials: the traceable numbers and their denominators

Denial statistics are the most quoted and least comparable figures in the field. Three are traceable to a described method, and they measure different populations, in different units, over different years.

The strongest is peer reviewed. Using a commercial claims database covering roughly 30% of Medicare Advantage enrollment, with 2019 services tracked through adjudication to mid-2022, 17.7% of claims were denied at initial submission. Sixty percent of denied claims were resubmitted and 56.6% of denied dollars were ultimately overturned, leaving providers 7.2% short of the dollars they initially billed.¹⁴

The most widely quoted is an analysis of federal transparency data. Of approximately 496 million claims received by HealthCare.gov issuers in 2024, 85 million of 451 million in-network claims were denied: a 19% in-network rate and 37% out-of-network. Fewer than 1% of denied claims were appealed, and 66% of the 262,982 internal appeals upheld the denial.¹⁵

Two corrections attach to that figure. It is routinely generalized to “insurers deny 20% of claims,” when it describes marketplace issuers on HealthCare.gov and nothing else. And it cannot carry an argument about clinical judgment, because the reasons are largely unrecorded: 36% of denials are categorized as “other” and 25% as administrative, while medical necessity accounts for 5%. The plurality reason for denial in the best public dataset is a residual category.

TABLE 1

The three traceable denial datasets and the populations they describe. These are not alternative estimates of one quantity.

Source and type	Population and period	Denial measure	Downstream outcome
Health Affairs, peer-reviewed ¹⁴	~30% of Medicare Advantage enrollment; 2019 services, adjudicated to mid-2022	17.7% denied at initial submission (16.6% dollar-weighted)	60% resubmitted; 56.6% of denied dollars overturned; net 7.2% of billed dollars lost
KFF, federal transparency data ¹⁵	HealthCare.gov issuers; 451 million in-network claims, 2024	19% in-network; 37% out-of-network; 3%–36% across 157 insurers	Under 1% appealed; 66% of internal appeals upheld the denial
Premier, industry survey ¹	516 hospitals, 36 states; private-payer claims, CY2022	Nearly 15% initially denied	54.3% overturned and paid; \$43.84 average rework cost

WHAT THE DENIAL FIGURES DO NOT SUPPORT

These three rows cannot be pooled, averaged or ranked. They cover different populations, count claims in one case and dollars in another, span 2019, 2022 and 2024, and mean different things by “overturned”. No single national denial rate exists in this evidence base. Nor is there corroboration among them, since the American Hospital Association’s \$26 billion and 70%-eventually-paid figures restate the Premier survey.^{1,2}

5 Prior authorization: burden, regulation and pledges

Prior authorization is the one part of the revenue cycle with a firm regulatory timetable, which makes the line between measured and asserted unusually easy to draw.

The burden figures, read correctly

The AMA’s 2025 prior authorization physician survey, fielded in December 2025 and released in May 2026 with 1,000 practicing physicians, reports about 40 prior authorizations per physician per week consuming 13 hours of time; 40% of practices employing staff who work exclusively on prior authorization; 95% reporting that it delays care; 26% reporting that it has led to a serious adverse event; and 74% reporting that denials have risen over five years.¹⁶ This is self-reported experience, not a time-and-motion measurement.

The 13-hour figure requires care, because the instrument measures *physician and staff time combined*. It is regularly restated as physician-only time, including in a peer-reviewed *npj Digital Medicine* analysis that describes physicians as spending an average of 13 hours per week submitting prior-authorization requests.¹⁷ The measured physician-only burden is much smaller: about 3.0 hours per week interacting with health plans, at \$68,274 per physician annually and \$31 billion nationally.¹⁸ The 13 hours describes a practice, not a doctor.

The regulation

CMS-0057-F, published at 89 FR 8758–8988 on 8 February 2024 under RIN 0938-AU87 and effective 8 April 2024, binds Medicare Advantage organizations, state Medicaid and CHIP fee-for-service programs, Medicaid and CHIP managed care plans, and issuers of qualified health plans on the federally-facilitated exchanges.¹⁹

TABLE 2

What CMS-0057-F requires and when. The 2026 obligations are process rules; the 2027 obligations are interfaces.

Compliance date	Requirement
1 January 2026	Standard prior authorization decisions within 7 calendar days and expedited decisions within 72 hours, with QHP issuers on the federally-facilitated exchanges excluded from the timeframe requirement
1 January 2026	A specific denial reason regardless of transmission method; annual public reporting of prior authorization metrics, first set due 31 March 2026
1 January 2027	FHIR-based Prior Authorization API; expanded Patient Access API carrying prior authorization data, excluding drugs; Provider Access API; Payer-to-Payer API with a five-year lookback

Medicare Advantage insurers made 52.8 million prior authorization determinations in 2024, about 1.7 per enrollee, of which 4.1 million were denied, a 7.7% denial rate against 6.4% in 2023. Only 11.5% of denials were appealed, and 80.7% of appeals were fully or partially overturned.²⁰

The first analysis of the metrics CMS-0057-F newly requires payers to publish reports standard-request denial rates of 12% in Medicare Advantage, 14% in Medicaid managed care and 18% in the ACA Marketplace, with appeal overturn rates of 67%, 47% and 43% respectively.²¹ The authors flag six gaps in the mandated data, one of which undermines everything built on it: insurers need not report request volumes, so the published denial counts have no denominator. It is also why this 12% and the 7.7% above cannot be reconciled.

The pledge

In June 2025, 48 health plans covering 257 million Americans committed through their trade association to six actions, among them a standardized FHIR-based electronic prior authorization framework operational by 1 January 2027, demonstrated reductions in authorization scope by 1 January 2026 defined plan by plan, and medical review of all clinically denied requests.²² The commitment sets no numeric target for scope reduction, provides for no independent audit, and carries no penalty.

A ten-month progress release in April 2026 reported that participating plans had eliminated 11% of prior authorizations across a range of medical services, 6.5 million fewer authorizations, with more than a 15% reduction in Medicare Advantage.²³ That figure is self-reported by the plans that made the pledge, with no response rate, no standardized denominator and no independent verification. The contemporaneous physician survey found 74% reporting that denials had increased, and a companion release reported that only one in three physicians expected the pledge to produce meaningful change.¹⁶

6 Algorithms on the payer side

The most consequential automation in the revenue cycle is the payer's. The evidence is one federal audit, one congressional staff report, one piece of sub-regulatory guidance and a set of unresolved lawsuits, which are not equivalent kinds of document.

The audit is firmest. The HHS Office of Inspector General drew a stratified random sample of 250 prior authorization denials and 250 payment denials from the fifteen largest Medicare Advantage organizations during one week in June 2019, and had coding experts and physicians review them. It found that 13% of denied prior authorization requests met Medicare coverage rules, meaning they would have been approved under fee-for-service, and that 18% of payment denials met both coverage rules and the organization's own billing rules.²⁴

That 13% is among the most frequently inverted numbers in this literature. It does not mean that 13% of prior authorization requests were denied, and it does not mean that 87% of denials were wrong. It is the share of *denied* requests that met Medicare coverage rules, and the sentence should be quoted rather than paraphrased.

A Senate Permanent Subcommittee on Investigations majority-staff report of October 2024 examined post-acute care denials. In 2022, post-acute denial rates ran three times the overall denial rate at UnitedHealthcare and CVS and sixteen times at Humana; UnitedHealthcare's post-acute denial rate rose from 10.9% to 22.7% between 2020 and 2022. The report attributes the increases to insurer investment in predictive and AI-based utilization-management tools.²⁵ It is a majority-staff document rather than an audit finding, and the attribution of cause is the staff's.

What is permitted is set out in CMS guidance accompanying the Medicare Advantage final rule CMS-4201-F, issued 6 February 2024 and anchored in 42 C.F.R. § 422.101(b)(6) and section 1557 of the Affordable Care Act. Algorithms *may assist* a coverage determination but may not drive it. An algorithm "that determines coverage based on a larger data set instead of the individual patient's medical history... would not be compliant."²⁶ This is sub-regulatory guidance interpreting a rule, not a rule of its own.

The litigation demands the most discipline. A putative class action in the District of Minnesota alleges that UnitedHealth used the nH Predict algorithm to terminate post-acute care in Medicare Advantage. The complaint alleges at paragraph 1 that the defendants knew the model had "a 90% error rate," at paragraph 38 that "over 90 percent of patient claim denials are reversed through either an internal appeal process or through federal Administrative Law Judge (ALJ) proceedings," and at paragraph 2 that roughly 0.2% of policyholders appeal.⁴ Each is an allegation in a pleading, yet the error rate circulates as a fact about an AI system's accuracy.

The second allegation is checkable against measured data, and it does not hold. The overturn rate for appealed Medicare Advantage prior authorization denials is 80.7%,²⁰ and 56.6% of denied dollars were ultimately overturned in the peer-reviewed claims analysis.¹⁴ Those are different quantities from each other and from the pleading's, but neither is "over 90 percent."

7 Administrative cost: a modeled opportunity is not a measured saving

The economic case for revenue-cycle automation rests on the size of administrative spending, which is genuinely large. The difficulty is that the biggest numbers estimate an opportunity rather than observe a result, and the two are quoted interchangeably.

The standard waste estimate puts total US healthcare waste at \$760 billion to \$935 billion annually, roughly 25% of total spending, with realistic potential savings of \$191 billion to \$286 billion. Administrative complexity is the largest single domain at \$265.6 billion, and it is the one domain for which the authors identified no studies of interventions that reduce it.²⁷

The projection most often quoted for AI estimates that adoption using existing technology could yield \$200 billion to \$360 billion annually in 2019 dollars, 5% to 10% of US healthcare spending, within five years, with administrative savings of roughly \$65 billion to \$135 billion.²⁸ It is an unrefereed working

paper modeling an opportunity under an assumption of adoption, and it measures savings realized nowhere. Nearly every claim that AI will save US healthcare hundreds of billions descends from it, usually with the working-paper status and the conditional dropped.

The measured administrative figures are smaller and more specific: \$68,274 per physician per year and \$31 billion nationally for practice interaction with health plans.¹⁸ That is cost incurred, not savings achievable. No study reviewed here measures a practice's administrative cost before and after deploying an automation tool, which is the whole difference between a modeled ceiling and an observed reduction.

8 What survives

A practice evaluating revenue-cycle automation can rely on a short list. Human E/M coding is not reproducible: three independent auditors agreed completely on 15.2% of charts under one guideline set and 29.2% under another,⁵ and federal reviewers found 42% of sampled Medicare E/M claims incorrectly coded.⁷ Medicare Advantage denies 17.7% of claims at initial submission and leaves providers 7.2% short of what they billed.¹⁴ CMS-0057-F imposes decision timeframes, specific denial reasons and public metrics reporting from 1 January 2026, and four FHIR APIs from 1 January 2027.¹⁹ The regulation is the only item on the list with a deadline attached.

The rest needs a denominator or a source. A prior authorization denial rate published under the new federal metrics has no request-volume denominator.²¹ A coding accuracy figure benchmarked on ICU discharge summaries says nothing about an ambulatory visit.^{11,12} A self-reported 11% reduction in authorization volume, from the plans that pledged it, is a claim rather than a finding.²³ A projected national saving from AI adoption estimates what might be available, not what anyone has recovered.²⁸

The useful question to ask of any revenue-cycle claim is not whether automation helps. It is what population the number was measured on, against what reference standard, and by whom. In this field that question eliminates most of the evidence, and what remains is worth more for it.

REFERENCES

Entries 5, 6, 9–12, 14, 17, 18 and 27 are peer reviewed; 13 is a preprint and 28 an unrefereed working paper, and both are flagged as such where they are used. Entries 7, 24 and 25 are federal or congressional documents, 8 and 19 are regulations, and 26 is sub-regulatory guidance. Entries 15, 20 and 21 are think-tank analyses of federal data. Entries 1, 3, 16, 22 and 23 are industry surveys or trade-association statements and carry the limitations of self-report; entry 2 restates entry 1. Entry 4 is a court filing, and the figures quoted from it are allegations rather than findings.

1. Premier Inc. *Trend Alert: Private Payers Retain Profits by Refusing or Delaying Legitimate Medical Claims*. Premier Inc. March 21, 2024. [premierinc.com](https://www.premierinc.com) INDUSTRY SURVEY

2. American Hospital Association. *The Cost of Caring: Challenges Facing America's Hospitals in 2025*. AHA. April 2025. [aha.org](https://www.aha.org) TRADE-ASSOCIATION REPORT
3. Scott S. *The Art of Setting E/M Pass Rate Thresholds*. Published by AAPC; no publication date is given on the record consulted. [aapc.com](https://www.aapc.com) POSITION PAPER
4. *Estate of Gene B. Lokken, et al. v. UnitedHealth Group, Inc., et al.* U.S. District Court, District of Minnesota, No. 0:23-cv-03514. Complaint filed November 14, 2023; motion to dismiss granted in part February 13, 2025. litigationtracker.law.georgetown.edu LITIGATION
5. Zuber TJ, Rhody CE, Muday TA, et al. *Variability in code selection using the 1995 and 1998 HCFA documentation guidelines for office services*. The Journal of Family Practice. 2000;49(7):642–645. PMID 10923576 CROSS-SECTIONAL
6. King MS, Sharp L, Lipsky MS. *Accuracy of CPT Evaluation and Management Coding by Family Physicians*. Journal of the American Board of Family Practice. 2001;14(3):184–192. pdxscholar.library.pdx.edu CROSS-SECTIONAL
7. U.S. Department of Health and Human Services, Office of Inspector General. *Improper Payments for Evaluation and Management Services Cost Medicare Billions in 2010*. Report OEI-04-10-00181. May 28, 2014. oig.hhs.gov GOVERNMENT REPORT
8. American Medical Association; Centers for Medicare & Medicaid Services. *Evaluation and Management (E/M) Office Visits—2021; and Medicare Program; CY 2021 Payment Policies Under the Physician Fee Schedule (CMS-1734-F)*. Federal Register. 2020;85(248):84472–85377. Effective January 1, 2021. [govinfo.gov](https://www.govinfo.gov) REGULATION
9. Apathy NC, Hare AJ, Fendrich S, et al. *Early Changes in Billing and Notes After Evaluation and Management Guideline Change*. Annals of Internal Medicine. 2022;175(4):499–504. doi:10.7326/M21-4402 OBSERVATIONAL
10. Searle T, Ibrahim Z, Dobson R. *Experimental Evaluation and Development of a Silver-Standard for the MIMIC-III Clinical Coding Dataset*. Proceedings of the 19th SIGBioMed Workshop on Biomedical Language Processing (BioNLP 2020). 2020:76–85. doi:10.18653/v1/2020.bionlp-1.8 BENCHMARK
11. Edin J, Junge A, Havtorn JD, et al. *Automated Medical Coding on MIMIC-III and MIMIC-IV: A Critical Review and Replicability Study*. Proceedings of SIGIR '23. 2023. doi:10.1145/3539618.3591918 REPLICATION STUDY
12. Zhang Y, Lyu C, Chang L, et al. *A systematic review of automated International Classification of Diseases coding models using the Medical Information Mart for Intensive Care dataset*. Digital Health. 2025;11. doi:10.1177/20552076251404518 SYSTEMATIC REVIEW

13. Simmons A, Takkavatakarn K, McDougal M, et al. *Benchmarking Large Language Models for Extraction of International Classification of Diseases Codes from Clinical Documentation*. medRxiv. 2024. [doi:10.1101/2024.04.29.24306573](https://doi.org/10.1101/2024.04.29.24306573) PREPRINT
14. Vabson B, Hicks AL, Chernew ME. *Medicare Advantage Denies 17 Percent Of Initial Claims; Most Denials Are Reversed, But Provider Payouts Dip 7 Percent*. Health Affairs. 2025;44(6):702–706. [doi:10.1377/hlthaff.2024.01485](https://doi.org/10.1377/hlthaff.2024.01485) COHORT
15. Long M, Lo J, Pestaina K. *Claims Denials and Appeals in ACA Marketplace Plans in 2024*. KFF. March 24, 2026. kff.org POLICY ANALYSIS
16. American Medical Association. *2025 AMA Prior Authorization Physician Survey*. AMA. Fielded December 2025; released May 2026; n = 1,000 practicing physicians. ama-assn.org. Companion release: *AMA survey: Prior authorization reform pledge falls short with physicians*, May 13, 2026. ama-assn.org INDUSTRY SURVEY
17. Raza S, Gerke S, Silcox C, et al. *Medicare advantage becoming a disadvantage with use of artificial intelligence in prior authorization review*. npj Digital Medicine. 2026;9:208. [doi:10.1038/s41746-026-02387-x](https://doi.org/10.1038/s41746-026-02387-x) POSITION PAPER
18. Casalino LP, Nicholson S, Gans DN, et al. *What Does It Cost Physician Practices To Interact With Health Insurance Plans?* Health Affairs. 2009;28(4):w533–w543. commonwealthfund.org CROSS-SECTIONAL
19. Centers for Medicare & Medicaid Services. *Medicare and Medicaid Programs; Patient Protection and Affordable Care Act; Advancing Interoperability and Improving Prior Authorization Processes (CMS-0057-F)*. Federal Register. 2024;89(27):8758–8988. RIN 0938-AU87. Effective April 8, 2024. federalregister.gov REGULATION
20. Fuglesten Biniek J, Sroczynski N, Freed M, et al. *Medicare Advantage Insurers Made Nearly 53 Million Prior Authorization Determinations in 2024*. KFF. January 28, 2026. kff.org POLICY ANALYSIS
21. Long M, Fuglesten Biniek J, Lo J, et al. *Prior Authorization Metrics Provide New Insights into Insurer Practices, but Gaps Remain*. KFF. August 13, 2026. kff.org POLICY ANALYSIS
22. AHIP. *Health Plans Take Action to Simplify Prior Authorization*. AHIP press release. June 23, 2025. ahip.org INDUSTRY STATEMENT
23. AHIP; Blue Cross Blue Shield Association. *Health Plans Reduce Prior Authorization, Support Continuity of Care and Enhanced Consumer Communications*. AHIP. April 7, 2026. ahip.org INDUSTRY STATEMENT

24. U.S. Department of Health and Human Services, Office of Inspector General. *Some Medicare Advantage Organization Denials of Prior Authorization Requests Raise Concerns About Beneficiary Access to Medically Necessary Care*. Report OEI-09-18-00260. April 27, 2022. oig.hhs.gov
GOVERNMENT REPORT
25. U.S. Senate Permanent Subcommittee on Investigations (Majority Staff). *Refusal of Recovery: How Medicare Advantage Insurers Have Denied Patients Access to Post-Acute Care*. October 17, 2024. hsgac.senate.gov
GOVERNMENT REPORT
26. Centers for Medicare & Medicaid Services. *FAQs Related to Coverage Criteria and Utilization Management Requirements in CMS Final Rule (CMS-4201-F)*. CMS. February 6, 2024. healthaipolicy.org
GUIDANCE
27. Shrank WH, Rogstad TL, Parekh N. *Waste in the US Health Care System: Estimated Costs and Potential for Savings*. JAMA. 2019;322(15):1501–1509. [doi:10.1001/jama.2019.13978](https://doi.org/10.1001/jama.2019.13978)
MODELING STUDY
28. Sahni NR, Stein G, Zimmel R, et al. *The Potential Impact of Artificial Intelligence on Healthcare Spending*. NBER Working Paper No. 30857. January 2023 (revised October 2023). nber.org
WORKING PAPER

RELATED PAPERS

Regulating AI in certified EHR technology

What the certification program and the FDA actually require of a vendor.

Ambient documentation and the clinician's day

The measured effect of ambient scribes, at the size the trials support.

A write-path architecture for an AI-native EMR

Why the claim should be a projection of the record rather than a reconstruction of it.