

Developing a Computable Payer-Provider Contract

Lauren Bilbo, Elman Amador Medina, and Klara Klarowicz, Graduate School of Business, Stanford University; Dashiell Miner and David Scheinker, School of Medicine, Stanford University; Stefanos Zenios, Graduate School of Business, Stanford University; and Kevin Schulman, School of Medicine and Graduate School of Business, Stanford University

Contact: kevin.schulman@stanford.edu

Abstract

What is the message? Modular, computable payer-provider contracts could provide a pathway toward true digital transaction standards in healthcare, with significant financial benefits from improvements in transaction efficiency. Such standards could also encourage novel insurance solutions in the healthcare market. In this study, the authors define the technical, data, security, and governance requirements for implementing such payer-provider contracts and an accompanying digital transaction / adjudication platform.

What is the evidence? Using the computable contract PATIENTS framework as a reference, along with contract language from a large healthcare system, the study assessed the data standards available to develop a computable contract framework. To do so, the authors tested whether existing data standards can accommodate commercial payer-provider contract language used in the market, and addressed incentive structures that have impeded adoption of standardized transaction models.

Timeline: Submitted: February 5, 2026; accepted after review February 17, 2026.

Cite as: Lauren Bilbo, Elman Amador Medina, Dashiell Miner, David Scheinker, Stefanos Zenios, Kevin Schulman. 2026. Developing a Computable Payer-Provider Contract. *Health Management, Policy and Innovation* (www.HMPI.org), Volume 11, Issue 1.

Support: The Ludy Family Foundation, The Hirsch Family Foundation, The Mindshare Institute, and the Government, Business and Society of the Graduate School of Business, Stanford University.

Introduction

There is an administrative-cost crisis in U.S. healthcare. Administrative costs are estimated to be at about 34.2% of total U.S. healthcare spending [1]. The United States spends about 10 times more on average on healthcare administrative expenses than any other OECD country (2). The result is that the U.S. spends more to manage the healthcare system than France, Germany and Italy combined spend on healthcare services. (3) In our multi-payer system, an estimated two-thirds of administrative spending is likely related to transactions or billing costs [4,5]. Administrative costs have been estimated to include \$374 billion in waste in 2025 spending, although the data are not accurately tracked. (6,7,8)

The primary challenge in the market is that numerous components of an analog transaction system have been digitized separately, leaving the U.S. health sector without a uniform digital transaction platform. Each health insurance carrier must develop a contract with each in-network provider to describe the terms of how transactions will be documented and paid for by each patient covered through their plan.

With over 1,000 health insurance carriers offering 318,000 distinct health plan products in the U.S. market, this process is fraught with duplication. Each carrier develops its own contract terms such as payment schedules, documentation standards, and prior authorization processes. They can market multiple types of products, HMOs, PPOs, high-deductible health plans, and each plan can have their own set of in-network providers. The complexity of this process contributes to the high billing-related costs in the U.S. market (3,9-11). A large fraction of the transaction spend is due to transaction friction: interpreting heterogeneous payer-provider contracts, manual prior authorization (PA), serial claim denials/resubmissions, and reconciliation of remittance advice idiosyncrasies (8,9). Beyond direct transaction costs, both payers and providers must maintain substantial IT infrastructure and financial staff to support bespoke

claims processing software, with ongoing overhead for tracking and implementing frequent changes to contract terms and coverage policies (12).

One pathway to reduce these administrative costs is to develop a true digital transaction standard for the market. (4,12,13) Standardizing the adjudication process is one important step towards reducing administrative costs in the healthcare industry. This would start with the development of digital contracts which could provide computable transactions for payers and providers. Such an effort would greatly simplify the transaction process. For payers and providers who seek novel transaction models, computable contracts would allow for such complexity but would use an algorithmic model for these transactions. By converting contract terms into executable logic, the computable framework inherently enables validation and systematic updating of coverage rules through structured testing against real-world claims data, reducing ancillary support costs. Financial modeling based on empirical data and mathematical simulation suggests that standardization has the potential to save over 60% of administrative costs – over \$1 trillion dollars in 2025 and more than the potential savings of a single-payer system such as Medicare For All. (12,13)

Computable contracts would have specific contract features reduced to computable terms. Making contract terms computable would require identifying standards for moving the contract language into an algorithm. For example, SMART on FHIR combines a standards-based data layer with the FHIR API. (14) This architecture is based on open standards including HL7's FHIR, OAuth2, and OpenID Connect. (14)

One computable contract framework has been publicly reported, the open-source PATIENTS framework. (15) This effort has developed standard modular contracts, an open payment system, a computable contract library, universal clinical coverage, standard plan mapping, and open transaction rails. (15)

We sought to extend this work to real-world payer-provider contracts. The PATIENTS framework proposes an idealized computable contract structure for optimal market conditions. Our work differs by testing whether existing data standards can accommodate actual commercial payer-provider contract language used in the market, and by addressing the incentive structures that have impeded adoption of standardized transaction models. Using the PATIENTS framework as a reference and contract language from a large healthcare system, we assessed the data

standards available to develop a computable contract examining each term in the contract. This work stands as an assessment of the feasibility of this approach.

Semantic Ambiguity in Source Artifacts

Three partially overlapping artifacts govern most commercial medical claims today:

1. The executed agreement / contract (legal and financial terms)
2. The provider manual (procedural, operational, and utilization policy guidance, often narrative and line-of-business specific)
3. The X12 companion guide (transaction formatting deviations, situational usage rules)

Critical adjudication logic (e.g., PA triggers, frequency limits, “lesser of” pricing precedence, quality modifier eligibility) frequently resides only as prose scattered across these documents. Each payer’s unique blend of a legal contract, provider manual, and X12 companion guide creates a bespoke semantic layer that revenue-cycle staff, clearinghouses, and claims systems must continually re-interpret. Furthermore, frequent changes in the governing artifacts necessitate frequent, nuanced reinterpretation of coverage rules.

The absence of a machine-readable, authoritative layer means every actor in the market rebuilds the same logic (clearinghouse edits, internal rules engines, third-party analytic overlays), propagating duplicated effort and inconsistency.

Divergent Incentives Lead to Divergent Technical Paradigms

Financial Incentives

At the point of service, *providers* generally maximize margin and operational efficiency when payment outcomes are predictable: given a set of clinical facts and billing codes, they want to know *ex ante* whether a claim will be paid, for what amount, and on what timeline. Determinism lets providers accurately staff clinical services, reduces working capital tied up in AR, and avoids costly billing rework cycles.

Payers, by contrast, benefit from ambiguity, vague coverage language, and discretionary prior-authorization criteria. Multi-stage denial rationales prolong decision windows, enable selective enforcement of payment rules, and increase the probability that some claims are abandoned or down-coded. Opacity acts as a frictional filter on total paid claims expense, making it possible for payers to utilize administrative complexity as a deliberate tool to support financial management of the benefit pool, both in terms of amount and timing of payments. In this environment, procedural friction functions as implicit policy, determining financial outcomes through attrition and exhaustion rather than clinical or contractual merit. (16)

Resulting Preference Split

These underlying incentives manifest as two preferred technical archetypes:

- **Provider-preferred architecture:** *Transparent, rules-based determinism.* Explicit, human-readable clauses compiled into executable logic (e.g., CQL/FHIRPath/DSL). Each decision produces a traceable proof: inputs → rule evaluations → allowance formula → outcome. Variance is eliminated; disagreements become focused on the adequacy of data, not on hidden logic.
- **Payer-preferred (status quo) architecture:** *Opaque, probabilistic discretion.* Black-box AI / heuristic scoring layers that classify claims or PA requests into “pay,” “pend,” or “deny” cohorts based on patterns in historical data rather than strictly enumerated contract rules. The internal model weights are not disclosed; justification can be post-hoc, preserving flexibility to modulate denial intensity.

Technical Implications

The divergence in incentives between providers and payers results in diverging technical goals, preferences, and approaches, as shown in the table below.

Dimension	Rules-Based, Deterministic (Provider-preferred)	AI / ML, Probabilistic (Payer-preferred)
Primary Objective	Predictable cash flow & minimized rework	Reduced aggregate payout & increased denial leverage
Representation	Declarative rule sets mapped to standard terminologies; versioned & diff-able	Proprietary model artifacts (weights, feature sets for algorithms), periodic validation and retraining cycles
Explanation	Intrinsic: rule execution trace is the explanation	Extrinsic: generated rationale templates; limited transparency
Governance Load	Up-front modeling & normalization effort; low runtime ambiguity	Lower up-front specification; ongoing tuning to maintain denial yield
Error Surface	Missing or stale rule → identifiable gap	Drift, bias, adversarial gaming; harder to contest individual outputs, such as clinical/administrative decisions
Provider Appeal Path	Challenge specific rule or input datum	Demand model transparency; since model explainability is often infeasible, fall back on manual override channels

Comparison of Technical Approaches

Neither the provider approach nor the payer approach is inherently “correct”. Rather, each technical approach offers unique advantages and disadvantages, as described in the table below.

Dimension	Rules-Based, Deterministic	AI / ML, Probabilistic
Primary Use	Coverage, PA, pricing, denials (explicit logic)	Pattern anomaly detection (e.g. detection of potentially inappropriate care), document parsing, data extraction; used to inform clinical decision makers when denials are initially recommended
Strengths	Transparency, auditability, low variance	Adaptable to noisy inputs, learns from historical denials/fraud patterns
Weaknesses	Brittle if inputs missing; manual rule authoring	Opacity, bias risk, harder to make legally binding
Failure Mode	Silent gaps, unhandled edge cases	False positives/negatives; adversarial gaming
Governance Fit	High (directly traceable to contract)	Low for binding decisions (should be advisory); suitable only as a decision-support tool for clinicians and administrators

Systemic Consequence

The coexistence of these paradigms produces a structural stalemate: providers escalate investment in contract “translation” layers to re-impose determinism locally; payers escalate ML-driven utilization management to retain flexibility. *Neither approach reduces inherent administrative complexity.*

Regulators have begun to intervene in this stalemate. In January 2024, the Centers for Medicare & Medicaid Services (CMS) issued a final rule for Medicare Advantage requiring that medical necessity determinations be “based on the circumstances of the specific individual...as opposed to using an algorithm or software that doesn’t account for an individual’s circumstances.” (17,18) The rule mandates physician review of determinations and public disclosure of evidence supporting algorithmic criteria. These requirements effectively limit the black-box AI approaches favored in the payer-preferred architecture, while remaining compatible with the rule-based

determinism of computable contracts. This regulatory shift creates both pressure and a pathway for industry transformation.

A computable modular contract explicitly realigns both sides on a deterministic core for enforceable payment logic while confining probabilistic methods to *assistive* roles (e.g., anomaly detection, document extraction) that do not unilaterally affect payment without a corresponding explicit rule. The computable contract core serves as the foundation for more effective AI/ML deployment and agentic AI workflows: rather than training models to reverse-engineer opaque contract language, machine learning can focus on higher-value tasks like fraud pattern detection, care pathway optimization, and predictive risk stratification. These are solutions that can be built from the structured data that result from computable transactions (some of these elements can be built into the transaction such as fraud detection, some will use the data to build new applications to support and evaluate clinical care). AI becomes dramatically more accurate, useful and economically efficient when operating on clean, structured, deterministic rule sets (a novel optimized process) rather than parsing inconsistent natural language across thousands of contract variations (an inefficient “legacy” process). (17)

This approach aligns with recent CMS requirements for Medicare Advantage that mandate physician review of coverage determinations and prohibit purely algorithmic denials (18,19). Further, it creates a pathway to reconcile incentive divergence: payers gain fraud and leakage controls plus faster legitimate adjudication; providers gain predictability, shifting competition toward efficiency and quality rather than exploitation of ambiguity. These results are included as an [appendix](#) to this paper.

Conclusion

Overall, we found that data standards exist for almost all contract terms in payer-provider contracts. We were able to develop a consistent structure and typography while preserving the original substance of the contract templates we identified. This effort consolidates requirements across data standards, rules computability, and operational workflows.

Computable contracts could provide a pathway to the development of true digital transaction standards in healthcare and agentic AI workflows, with significant financial benefits from improving transaction efficiency. Such a set of transaction standards could also open a pathway

for entry of novel insurance solutions into the healthcare market. (20,21)

Appendix

Technical Requirements for a Computable Payer-Provider Contract Digital Adjudication Platform

References

1. David U. Himmelstein, Terry Campbell, Steffie Woolhandler. [Health Care Administrative Costs in the United States and Canada, 2017](#). Ann Intern Med.2020;172:134-142.
2. Turner A, Miller G, Lowry E. *High US health care spending: where is it all going? Published October 4, 2023. Accessed December 20, 2025.*
<https://www.commonwealthfund.org/publications/issuebriefs/2023/oct/high-us-health-care-spendingwhere-is-it-all-going>
3. Schulman KA, Nielsen PK Jr, Patel K. AI Alone Will Not Reduce the Administrative Burden of Health Care. JAMA. 2023 Dec 12;330(22):2159-2160.
4. Istvan B, Schulman KA, Zenios S. Addressing Health Care's Administrative Cost Crisis. JAMA. 2025 Mar 4;333(9):749-750.
5. Health Affairs. The Role Of Administrative Waste In Excess US Health Spending. October 6, 2022.
<https://www.healthaffairs.org/content/briefs/role-administrative-waste-excess-us-health-spending>
6. Sahni NR, Carrus B, Cutler DM. Administrative simplification and the potential for saving a quarter-trillion dollars in health care. JAMA. 2021;326(17):1677-1678.
7. 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 (waste as a percent was updated to 2025 National Health Expenditures).
8. Sahni NR, Istvan B, Bello Thornhill H, Joynt-Maddox KE, Cutler D, Emanuel EJ. Availability of consistent, reliable, and actionable public data on US hospital administrative expenses.

Health Aff Sch. 2025 May 22;3(5):qxaf069.

9. Tseng, P., Kaplan, R.S., Richman, B.D., Shah, M.A. and Schulman, K.A., 2018. Administrative costs associated with physician billing and insurance-related activities at an academic health care system. *Jama*, 319(7), pp.691-697.
10. National Association of Insurance Commissioners. U.S. Health Insurance Industry: 2022 Annual Results. Published 2023. Accessed February 15, 2026.
<https://content.naic.org/sites/default/files/industry-analysis-report-2022-health.pdf>
11. Sahni NR, Istvan B, Stafford C, Cutler D. Perceptions of prior authorization burden and solutions. *Health Aff Sch*. 2024 Aug 6;2(9):qxae096.
12. David Scheinker, Kevin Schulman, and Stefanos Zenios. It's Time To Bring Health Care Systems Into the Digital Age | Opinion. *Newsweek* April 09, 2025.
<https://www.newsweek.com/its-time-bring-health-care-systems-digital-age-opinion-2056429>
13. Scheinker, D., Richman, B.D., Milstein, A. and Schulman, K.A., 2021. Reducing administrative costs in US health care: Assessing single payer and its alternatives. *Health Services Management Research*, 34(3), pp.153-158
14. <https://smarthealthit.org/smart-on-fhir-api/>
15. <https://open.turquoise.health/docs/getting-started/introduction/>
16. Herd, P., & Moynihan, D. P. (2018). *Administrative burden: Policymaking by other means*. Russell Sage Foundation.
17. Oliver M., Faris R. A Blueprint for Enterprise-Wide Agentic AI Transformation. *HBR*. February 12, 2026.
<https://hbr.org/sponsored/2026/02/a-blueprint-for-enterprise-wide-agentic-ai-transformation>. Accessed Feb 17, 2026.
18. Mello MM, Rose S. Denial—Artificial Intelligence Tools and Health Insurance Coverage Decisions. *JAMA Health Forum*. 2024;5(3):e240622.
19. CMS. CMS Interoperability and Prior Authorization Final Rule (CMS-0057-F). January 17, 2024.
<https://www.cms.gov/cms-interoperability-and-prior-authorization-final-rule-cms-0057-f>
20. Brooke Istvan, Perry Nielsen Jr, Megan Eluhu, Bryan Kozin, Walt Winslow, David Scheinker, Kavita Patel, Kenneth Favaro, Stefanos Zenios, Kevin Schulman. 2024. Applying Precedents Thinking to the Intractable Problem of Transaction Costs in Healthcare. *Health*

Management, Policy and Innovation (www.HMPI.org). Volume 9, Issue 3.

21. Cutler DM. Reducing administrative costs in US health care. The Hamilton Project.

Accessed July 14,

2024. <https://www.hamiltonproject.org/publication/policy-proposal/reducing-administrative-costs-in-u-s-health-care/>