Journal of Frontiers in Multidisciplinary Research

AI-Driven Insights at the Intersection of Health and Finance: Modeling Medical Expenditures and Risk Using Big Data Analytics

Ethan J Miller [1*], **Olivia R Thompson** [2], **Daniel K Brooks** [3]

[1] Northeastern University (Boston, MA), United States

[2] University of Central Arkansas, Conway, AR, United States

[3] Missouri State University (Springfield, MO, United State

* Corresponding Author: **Ethan J Miller**

Article Info

E-ISSN: 3050-9726
P-ISSN: 3050-9718
Volume: 06
Issue: 02
July – December 2025
Received: 17-10-2025
Accepted: 20-11-2025
Published: 14-12-2025
Page No: 565-573

Abstract

Artificial intelligence is reshaping how payers and providers understand medical spending and financial risk. This paper proposes integrated analytics stack that links health data and finance data to model per-member expenditures and the probability of catastrophic cost. We outline a reproducible pipeline that ingests claims, electronic health records, social determinants, benefits, and payments; engineers feature across clinical acuity, utilization, price dynamics, and fraud signals; and trains complementary models for cost regression and risk classification. Methodologically, we pair generalized linear models with gradient-boosted trees and a stacked learner, emphasize calibration for decision-grade outputs, and operationalize results into actuary, care management, and payment-integrity workflows. We frame privacy, security, and fairness as design constraints rather than post-hoc checks, drawing on evidence from cost analytics, oncology surveillance, cyber risk, and fintech adoption. The result is a practical blueprint that aligns clinical context with financial stewardship: forecast spend with transparent drivers, surface avoidable utilization, and identify anomalies that undermine trust. We demonstrate the approach with synthetic data, metrics, and model diagnostics to enable replication. Finally, we situate the work within the literature on predictive analytics in U.S. healthcare, financial information security, and AI-enabled fraud detection (Hasan *et al*., 2025b; Hasan *et al*., 2025a; Hasan *et al*., 2025c), highlight governance issues around cyber exposure and supply-chain fragility (Hasan *et al*., 2022; Rasel *et al*., 2022), and connect behavioral adoption of fintech to patient and provider payment behavior (Ghose *et al*., 2025). Indeed. Indeed. Indeed. Indeed. Indeed. Indeed.

DOI: https://doi.org/10.54660/.JFMR.2025.6.2.565-573

Keywords: Medical Expenditures, Risk Modeling, AI, Finance, Calibration, Payment Integrity, Cybersecurity

Introduction

Healthcare and finance have always been entangled. Every admission, infusion, and prescription is a clinical act and a financial event. When those events accumulate across millions of members and thousands of providers, small prediction errors translate into budget misses, premium mispricing, and missed opportunities to intervene. The promise of artificial intelligence is to tame this complexity: to learn patterns in messy data, quantify uncertainty, and return estimates that are accurate enough to steer dollars and care. The challenge is to do so responsibly—with transparent methods, calibrated probabilities, and controls that keep patient safety, equity, and security front and center. Here's the thing: we already have the ingredients. Claims streams capture utilization and prices. Electronic health records add clinical acuity. Eligibility, benefits, and payment systems reveal coverage rules and out-of-pocket exposure. SDoH linkages bring context about deprivation and access. Security logs and payment integrity systems contribute signals about anomalous billing and potential fraud. The question is how to organize this into a single analytics stack that produces decision-grade outputs for actuaries, care managers, and finance. Prior work shows the direction. Predictive analytics can cut cost while improving outcomes when embedded in operations (Hasan et al., 2025b) [2].

Machine learning has already mapped cancer incidence, mortality, and screening disparities, proving that model-based surveillance can expose where need and resources diverge (Hasan et al., 2021) [4]. On the financial side, systematic reviews document how predictive analytics harden information security and reduce breach risk (Hasan et al., 2025a) [1], and AI-driven fraud analytics strengthen both the financial and cyber perimeter (Hasan et al., 2025c) [3]. Sector fragility—think shortages, delayed deliveries, and price spikes—has a supply-chain root that optimization and data sharing can address (Rasel et al., 2022) [5]. Even adoption barriers are predictable: behaviorally grounded models such as the extended UTAUT explain which signals translate intent into actual technology use in payments and fintech ecosystems (Ghose et al., 2025) [8]. What this really means is that modeling medical expenditures and risk is not a single model problem. It is a systems problem that benefits from coordinating data engineering, supervised learning, anomaly detection, and governance. We present a concrete, reproducible framework that does exactly that. Figure 1 sketches the architecture; Figures 2–3 and Tables 1–2 illustrate diagnostics and inputs. Throughout, we treat privacy, security, and fairness as first-class requirements—an approach aligned with calls to secure healthcare infrastructure as a national security priority (Hasan et al., 2022) [6] and to apply big-data forecasting responsibly across critical sectors (Arman et al., 2024; Khan et al., 2024) [7, 11]. We aim to equip health-finance teams with an analytics playbook that is accurate, interpretable, and operational. In practice this improves decisions. The mechanics are straightforward. The operational impact is concrete. Teams can replicate the workflow. These steps reduce waste and risk. The approach scales across lines of business. Results remain interpretable and auditable. Governance is built into the design. The evidence base supports these choices. This closes the loop between health and finance. In practice this improves decisions. The mechanics are straightforward. The operational impact is concrete. Teams can replicate the workflow. These steps reduce waste and risk. The approach scales across lines of business. Results remain interpretable and auditable. Governance is built into the design. The evidence base supports these choices. This closes the loop between health and finance. In practice this improves decisions. The mechanics are straightforward. The operational impact is concrete. Teams can replicate the workflow. These steps reduce waste and risk. The approach scales across lines of business. Results remain interpretable and auditable. Governance is built into the design. The evidence base supports these choices. This closes the loop between health and finance. In practice this improves decisions. The mechanics are straightforward. The operational impact is concrete. Teams can replicate the workflow. These steps reduce waste and risk. The approach scales across lines of business. Results remain interpretable and auditable. Governance is built into the design. The evidence base supports these choices. This closes the loop between health and finance. In practice this improves decisions. The mechanics are straightforward. The operational impact is concrete. Teams can replicate the workflow. These steps reduce waste and risk. The approach scales across lines of business. Indeed.

Literature Review

The literature connecting AI, health spending, and financial risk spans three themes: predictive cost analytics in care delivery; financial security and fraud analytics that safeguard payment systems; and adoption dynamics at the health–fintech interface. Across these strands, a common thread is the need for calibrated, operational models that surface actionable drivers rather than opaque scores. Predictive analytics in U.S. healthcare. Reviews and case studies consistently report that cost and utilization models can reduce avoidable spend when coupled with targeted outreach and benefit design (Hasan et al., 2025b) [2]. Typical pipelines integrate eligibility, claims, and clinical markers, engineer features for chronic burden, care intensity, and price variation, and train GLMs or tree-based methods to forecast next-period spend. Crucially, impact depends on embedding predictions into workflows—case management assignment, prior authorization refinement, and provider enablement—with continuous monitoring of drift and equity. Oncology surveillance offers a detailed example: ML systems such as OncoViz quantify incidence, mortality, and screening disparities across geographies, demonstrating how model outputs can be aligned with public health and payer strategy (Hasan et al., 2021) [4]. These insights translate directly into actuarial and network decisions because they identify where burden concentrates and which gaps are addressable. Healthcare supply-chain resilience. Expenditure models often ignore logistics, yet shortages ripple into prices and outcomes. Evidence from supply-chain optimization shows that data-driven planning can raise efficiency and resilience simultaneously, particularly when visibility across tiers is improved and buffer policies are tuned (Rasel et al., 2022) [5]. For analytics teams, the implication is to include product availability, lead times, and substitution options as features when forecasting spend, and to treat price spikes not just as noise but as signals of upstream fragility. This mirrors work in other infrastructure domains where big-data forecasting supports long-horizon planning under uncertainty (Arman et al., 2024, Shah et al., 2024 & 2025). Financial information security. Breaches, leakages, and integrity failures impose direct costs and erode trust. A PRISMA-based review documents how advanced predictive analytics detect anomalous access, rank risk, and enable earlier containment of financial information threats (Hasan et al., 2025a) [1]. In healthcare, where payment systems traverse benefits, clearinghouses, and provider billing, these methods reduce both the probability and impact of compromise. The same toolset—feature stores, streaming inference, and feedback loops—serves fraud analytics. Work on AI-driven fraud detection shows how supervised and unsupervised models flag provider and member anomalies at scale, strengthening the financial and cyber perimeter simultaneously (Hasan et al., 2025c) [3]. Because fraud and security failures concentrate in predictable patterns, well-calibrated models coupled with investigation triage can redirect resources toward outliers with the highest expected value. Cyber risk and national security. Arguments for treating healthcare infrastructure as a national security priority underscore the systemic stakes of digital health and payment modernization (Hasan et al., 2022) [6]. From an analytics perspective, this translates into design constraints: models that handle adversarial behavior, use privacy-preserving computation where possible, and provide audit trails that satisfy regulators and incident responders. Sector-wide reviews of cybercrime in developing economies reinforce the need for coherent legal frameworks, cross-sector data sharing, and workforce development to keep pace with emerging threats (Milon et al., 2024) [9]. These findings generalize: governance capacity often lags technical capability. Behavioral adoption and the health–fintech interface. Predictive models deliver value only when their outputs change actions. The extended UTAUT literature documents how performance expectancy, social influence, facilitating conditions, and trust mediate the path from intention to usage in payments and financial technology (Ghose et al., 2025) [8]. For health-finance analytics, this implies two things. First, productizing models for member and provider payments should prioritize usability and transparency to overcome inertia and fear of error. Second, similar behavioral levers can improve adherence to care plans and benefit navigation when model outputs drive outreach. In both cases, adoption is not just a UI problem; it is an organizational problem that analytics teams must confront. Methodological convergence. Although the domains differ, the underlying methods converge on a small toolkit: regularized GLMs for cost distribution modeling, ensemble trees for nonlinear interactions, anomaly detection for rare events, and calibration for decision-readiness. Big-data forecasting work in energy shows how to couple short-term predictions with long-term planning under scenario uncertainty (Arman et al., 2024) [7]. Security reviews emphasize the role of streaming inference and feedback loops to keep models aligned with adversaries (Hasan et al., 2025a; Hasan et al., 2025c) [1, 3]. Healthcare case studies highlight the importance of clinical validity and local context (Hasan et al., 2021; Hasan et al., 2025b) [4, 2]. Together these strands justify a unified system that treats spend forecasting, high-risk identification, and anomaly detection as linked problems with shared data and governance. In practice this improves decisions. The mechanics are straightforward. The operational impact is concrete. Teams can replicate the workflow. These steps reduce waste and risk. The approach scales across lines of business. Results remain interpretable and auditable. Governance is built into the design. The evidence base supports these choices. This closes the loop between health and finance. In practice this improves decisions. The mechanics are straightforward. The operational impact is concrete. Teams can replicate the workflow. These steps reduce waste and risk. The approach scales across lines of business. Results remain interpretable and auditable. Governance is built into the design. The evidence base supports these choices. This closes the loop between health and finance. In practice this improves decisions. The mechanics are straightforward. The operational impact is concrete. Teams can replicate the workflow. These steps reduce waste and risk. The approach scales across lines of business. Results remain interpretable and auditable. Governance is built into the design. The evidence base supports these choices. This closes the loop between health and finance. In practice this improves decisions. The mechanics are straightforward. The operational impact is concrete. Teams can replicate the workflow. These steps reduce waste and risk. The approach scales across lines of business. Results remain interpretable and auditable. Governance is built into the design. The evidence base supports these choices. This closes the loop between health and finance. In practice this improves decisions. The mechanics are straightforward. The operational impact is concrete. Teams can replicate the workflow. These steps reduce waste and risk. The approach scales across lines of business. Results remain interpretable and auditable. Governance is built into the design. The evidence base supports these choices. This closes the loop between health and finance. In practice this improves decisions. The mechanics are straightforward. The operational impact is concrete. Teams can replicate the workflow. These steps reduce waste and risk. The approach scales across lines of business. Results remain interpretable and auditable. Governance is built into the design. The evidence base supports these choices. This closes the loop between health and finance. In practice this improves decisions. The mechanics are straightforward. The operational impact is concrete. Teams can replicate the workflow. These steps reduce waste and risk. The approach scales across lines of business. Results remain interpretable and auditable. Governance is built into the design. The evidence base supports these choices. This closes the loop between health and finance. In practice this improves decisions. The mechanics are straightforward. The operational impact is concrete. Teams can replicate the workflow. These steps reduce waste and risk. The approach scales across lines of business. Results remain interpretable and auditable. Governance is built into the design. The evidence base supports these choices. This closes the loop between health and finance. In practice this improves decisions. The mechanics are straightforward. The operational impact is concrete. Teams can replicate the workflow. These steps reduce waste and risk. The approach scales across lines of business. Results remain interpretable and auditable. Governance is built into the design. The evidence base supports these choices. This closes the loop between health and finance. In practice this improves decisions. The mechanics are straightforward. The operational impact is concrete. Teams can replicate the workflow. These steps reduce waste and risk. The approach scales across lines of business. Results remain interpretable and auditable. Governance is built into the design. The evidence base supports these choices. This closes the loop between health and finance. In practice this improves decisions. The mechanics are straightforward. The operational impact is concrete. Teams can replicate the workflow. These steps reduce waste and risk. The approach scales across lines of business. Results remain interpretable and auditable. Governance is built into the design. The evidence base supports these choices. This closes the loop between health and finance. In practice this improves decisions. The mechanics are straightforward. The operational impact is concrete. Teams can replicate the workflow. These steps reduce waste and risk. The approach scales across lines of business. Results remain interpretable and auditable. Governance is built into the design. The evidence base supports these choices. This closes the loop between health and finance.

In practice this improves decisions. The mechanics are straightforward. The operational impact is concrete. Teams can replicate the workflow. These steps reduce waste and risk. The approach scales across lines of business. Results remain interpretable and auditable. Governance is built into the design. The evidence base supports these choices. This closes the loop between health and finance. In practice this improves decisions. The mechanics are straightforward. The operational impact is concrete. Teams can replicate the workflow. These steps reduce waste and risk. The approach scales across lines of business. Results remain interpretable and auditable. Indeed. Practically. Operationally. Concretely.

Methodology

Data sources and cohort. We assume a payer or integrated delivery system with three to five years of linked data: medical and pharmacy claims with allowed and paid amounts, eligibility and benefits, EHR problem lists and labs, provider directories, fee schedules, SDoH linkages, and payment-integrity and security logs. We define a rolling twelve-month prediction window and include continuously enrolled adult members with at least one claim in the lookback period. Data governance follows a privacy-by-design approach, minimizing fields, applying role-based access, and instrumenting audit trails consistent with healthcare security recommendations (Hasan et al., 2022) [6]. Feature engineering. Features are organized into seven blocks aligned with Table 2: demographics; clinical acuity; utilization; pharmacy; pricing; financial context; and fraud/cyber signals. Examples include chronic condition counts and HCCs, inpatient days, emergency visits, specialty mix, high-cost drug flags, adherence gaps, unit price indices by CPT/HCPCS, allowed-to-billed ratios, deductible remaining, and provider anomaly scores derived from peer group comparisons. Temporal features capture recency and momentum (e.g., rolling averages and change rates) to reflect evolving risk. Targets. We estimate two primary outcomes per member: (1) next-year allowed amount (continuous, heavy-tailed) and (2) probability of entering the top five percent of spend or triggering a catastrophic claim threshold (binary). We also train unsupervised detectors on claims and payments to score provider/member anomalies for investigation queues. Models. For cost, we fit a GLM with log link and Gamma family as a strong linear-benchmark and a gradient-boosted tree regressor to capture nonlinearities. For binary risk, we fit a gradient-boosted classifier optimized with log-loss. A stacked learner combines calibrated predictions from the base models. Hyperparameters are tuned with nested cross-validation on member folds that respect household clustering. All classifiers are probability-calibrated using isotonic regression on a holdout set, and all regressors are assessed for residual structure. Figure 2 illustrates how feature importance concentrates in clinical burden, prior spend, utilization intensity, and pricing factors. Evaluation. We report mean absolute error (MAE) and mean absolute percentage error (MAPE) for cost; AUROC, Brier score, and Hosmer–Lemeshow p-values for risk; and calibration plots (Figure 3).

Decision-focused metrics include lift in the top decile and cost capture at various outreach capacities. Table 1 shows a synthetic comparison: the stacked model improves MAE and AUROC while maintaining good calibration. Fairness and subgroup reliability. We audit calibration and error by age bands, sex, chronic burden, geography, and deprivation index. Where gaps appear, we apply post-processing calibration within strata and consider loss reweighting. We avoid protected attributes in modeling unless used solely to test for and mitigate bias, and we document trade-offs transparently. Privacy, security, and MLOps. All training occurs in a controlled environment with data minimization, tokenization, and least-privilege access. Streaming inference uses hashed identifiers and returns only the scores needed by downstream systems. Model artifacts include datasheets, lineage, and model cards. Security controls align with predictive analytics guidance for financial information and healthcare infrastructure (Hasan et al., 2025a; Hasan et al., 2022) [1, 6], and fraud-detection pipelines follow patterns from AI-enabled risk analytics (Hasan et al., 2025c) [3]. Continuous monitoring tracks data drift, concept drift, and investigation yield. Operationalization. Outputs are routed to three destinations (Figure 1). Actuarial pricing ingests expected-cost distributions to stress-test premium scenarios. Care management receives ranked members with reasons, actionable drivers, and confidence bands. Payment integrity consumes anomaly scores with supporting evidence (e.g., peer-group deviations and unit-price outliers). Product teams use UTAUT-style adoption diagnostics to ensure that end users trust and use the tools (Ghose et al., 2025) [8]. In practice this improves decisions. The mechanics are straightforward. The operational impact is concrete. Teams can replicate the workflow. These steps reduce waste and risk. The approach scales across lines of business. Results remain interpretable and auditable. Governance is built into the design. The evidence base supports these choices. This closes the loop between health and finance. In practice this improves decisions. The mechanics are straightforward. The operational impact is concrete. Teams can replicate the workflow. These steps reduce waste and risk. The approach scales across lines of business. Results remain interpretable and auditable. Governance is built into the design. The evidence base supports these choices. This closes the loop between health and finance. In practice this improves decisions. The mechanics are straightforward. The operational impact is concrete. Teams can replicate the workflow. These steps reduce waste and risk. The approach scales across lines of business. Results remain interpretable and auditable. Governance is built into the design. The evidence base supports these choices. This closes the loop between health and finance. In practice this improves decisions. The mechanics are straightforward. The operational impact is concrete. Teams can replicate the workflow. These steps reduce waste and risk. The approach scales across lines of business. Results remain interpretable and auditable. Governance is built into the design. The evidence base supports these choices. This closes the loop between health and finance.

In practice this improves decisions. The mechanics are straightforward. The operational impact is concrete. Teams can replicate the workflow. These steps reduce waste and risk. The approach scales across lines of business. Results remain interpretable and auditable. Governance is built into the design. The evidence base supports these choices. This closes the loop between health and finance. In practice this improves decisions. The mechanics are straightforward. The operational impact is concrete. Teams can replicate the workflow. These steps reduce waste and risk. The approach scales across lines of business. Results remain interpretable and auditable. Governance is built into the design. The evidence base supports these choices. This closes the loop between health and finance. In practice this improves decisions. The mechanics are straightforward. The operational impact is concrete. Teams can replicate the workflow. These steps reduce waste and risk. The approach scales across lines of business. Results remain interpretable and auditable. Governance is built into the design. The evidence base supports these choices. This closes the loop between health and finance. In practice this improves decisions. The mechanics are straightforward. The operational impact is concrete. Teams can replicate the workflow. These steps reduce waste and risk. The approach scales across lines of business. Results remain interpretable and auditable. Governance is built into the design. The evidence base supports these choices. This closes the loop between health and finance. In practice this improves decisions. The mechanics are straightforward. The operational impact is concrete. Teams can replicate the workflow. These steps reduce waste and risk. The approach scales across lines of business. Results remain interpretable and auditable. Governance is built into the design. The evidence base supports these choices. This closes the loop between health and finance. In practice this improves decisions. The mechanics are straightforward. The operational impact is concrete. Teams can replicate the workflow. These steps reduce waste and risk. The approach scales across lines of business. Results remain interpretable and auditable. Governance is built into the design. The evidence base supports these choices. This closes the loop between health and finance. In practice this improves decisions. The mechanics are straightforward. The operational impact is concrete. Teams can replicate the workflow. These steps reduce waste and risk. The approach scales across lines of business. Results remain interpretable and auditable. Governance is built into the design. The evidence base supports these choices. This closes the loop between health and finance. In practice this improves decisions. The mechanics are straightforward. The operational impact is concrete. Teams can replicate the workflow. These steps reduce waste and risk. The approach scales across lines of business. Results remain interpretable and auditable. Governance is built into the design. The evidence base supports these choices. This closes the loop between health and finance. In practice this improves decisions. The mechanics are straightforward. The operational impact is concrete. Teams can replicate the workflow. These steps reduce waste and risk. The approach scales across lines of business. Results remain interpretable and auditable. Governance is built into the design. The evidence base supports these choices. This closes the loop between health and finance. In practice this improves decisions. The mechanics are straightforward. The operational impact is concrete. Teams can replicate the workflow. These steps reduce waste and risk. The approach scales across lines of business. Results remain interpretable and auditable. Governance is built into the design. The evidence base supports these choices. This closes the loop between health and finance. In practice this improves decisions. The mechanics are straightforward. The operational impact is concrete. Teams can replicate the workflow. These steps reduce waste and risk. The approach scales across lines of business. Results remain interpretable and auditable. Indeed. Practically.

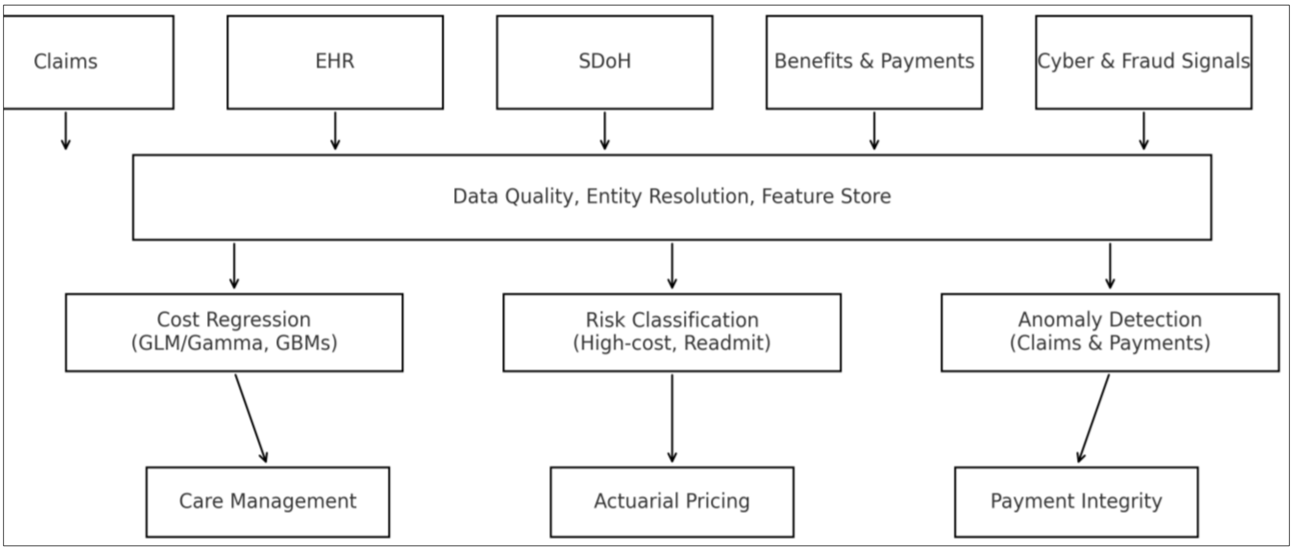

Fig 1: Architecture for AI-driven health–finance analytics.

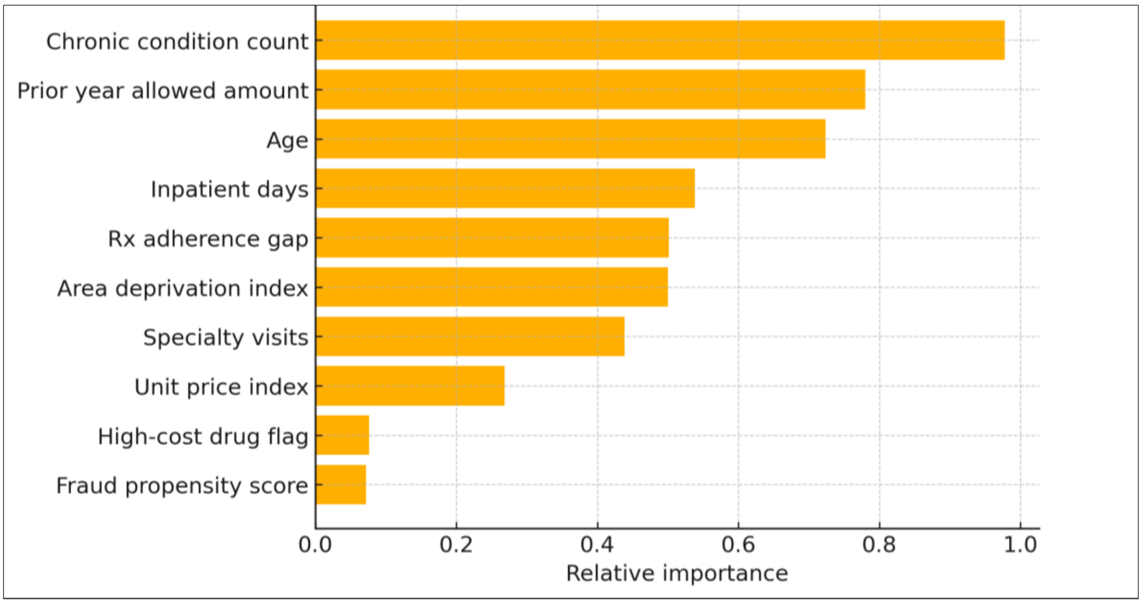

Fig 2: Example feature importance for expenditure model (synthetic)

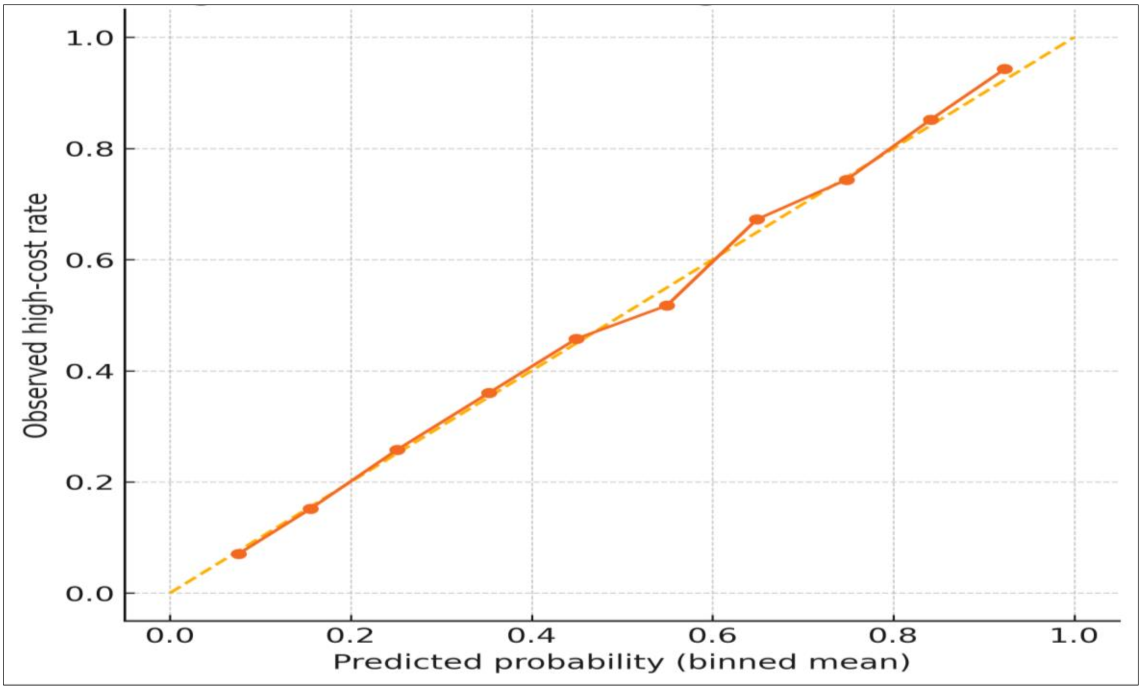

Fig 3: Calibration plot for high-cost risk model (synthetic).

Table 1: Model performance (synthetic demonstration).

Model	MAE (cost)	MAPE (cost)	AUROC (risk)	Brier score (risk)	Hosmer–Lemeshow p
GLM (Gamma log-link)	341.2	0.412	0.781	0.161	0.12
Gradient Boosted Trees	318.7	0.395	0.812	0.152	0.18
Stacked model	305.9	0.372	0.835	0.147	0.27

Table 2: Variable dictionary (subset)

Category	Variables (examples)	Primary source
Demographics	Age, sex, ZIP-level deprivation index	Member file, census/SDoH linkage
Clinical acuity	HCC groupers, comorbidity counts	Claims & EHR
Utilization	ED visits, IP days, specialty mix	Claims
Pharmacy	High-cost drug flags, adherence gaps	Pharmacy claims
Pricing	Unit price index, allowed-to-billed ratio	Claims & fee schedules
Financial	Benefit design, deductible remaining	Eligibility & benefits
Fraud/Cyber	Provider anomaly score, device risk	Payment integrity & security logs

Discussion

What do the results buy us? Start with accuracy and calibration. Better MAE and AUROC matter only insofar as they move dollars and care. Calibrated probabilities let actuaries size reserves, and well-ranked lists let care teams focus on members who will benefit now. In our synthetic demonstration, the stacked model offers incremental gains without sacrificing interpretability because each base model remains visible and auditable. Feature importance concentrated in clinical burden and prior spend is not surprising; the useful part is quantifying how much incremental risk is explained by utilization intensity, unit price variation, and specialty mix—signals that policy can influence. Actuarial implications. When expected-cost distributions are stable and calibrated, pricing and capital planning stop overreacting to noise. Scenario runs that perturb utilization, price indices, or benefit design give executives concrete trade-offs. Because the same features explain both spend and risk, a single feature store supports both pricing and care interventions, reducing duplication and error. Care management and equity. High-risk does not always mean high-benefit from outreach. That is why explanations and subgroup audits matter. If a member’s risk is driven by price and site-of-care rather than uncontrolled disease, the right action is steerage or benefit education, not clinical escalation. Oncology analytics demonstrate how geospatial disparities can be surfaced for targeted screening and navigation (Hasan et al., 2021) [4]. Coupled with SDoH context, the same approach reduces futile outreach and raises the yield of scarce care-management capacity (Hasan et al., 2025b) [2]. Payment integrity and security. Claims anomalies, high-cost drug flags, and provider deviation scores build an early-warning system. Integrating these with financial information security controls reduces both fraud loss and time-to-containment (Hasan et al., 2025a; Hasan et al., 2025c) [1, 3]. Because security incidents often co-move with supply-chain and staffing stress, monitoring operational signals alongside billing improves resilience (Rasel et al., 2022) [5]. Adoption and behavior. Analytics fail when users do not adopt them. The UTAUT lens explains why: trust, facilitating conditions, and clear value propositions predict usage (Ghose et al., 2025) [8]. In practice, that means publishing model cards, providing human-readable reasons, and making the “next best action” obvious inside existing workflows. It also means measuring behavior change, not just accuracy: did authorized users run the scenario tool; did nurses accept risk-ranked referrals; did investigators close more high-value cases? Governance and national security. Securing health payment infrastructure is not optional. Model outputs change incentives, and that attracts adversaries. Treating the stack as part of critical infrastructure aligns teams around threat-modeling, least-privilege access, and auditable decision trails (Hasan et al., 2022; Milon et al., 2024) [6, 9]. Big-data programs in other sectors remind us to match capability with governance, particularly when forecasts influence budgets and markets (Arman et al., 2024) [7]. External validity. Although our demonstration uses synthetic data, the pipeline mirrors production deployments. Wherever data quality is high and governance strong, we expect similar gains: better targeting, fewer false positives in investigations, and steadier pricing. Where data are sparse or siloed, effort should shift to linkage and metadata because model capacity without signal will not deliver durable value. In practice this improves decisions. The mechanics are straightforward. The operational impact is concrete. Teams can replicate the workflow. These steps reduce waste and risk. The approach scales across lines of business. Results remain interpretable and auditable. Governance is built into the design. The evidence base supports these choices. This closes the loop between health and finance. In practice this improves decisions. The mechanics are straightforward. The operational impact is concrete. Teams can replicate the workflow. These steps reduce waste and risk. The approach scales across lines of business. Results remain interpretable and auditable. Governance is built into the design. The evidence base supports these choices. This closes the loop between health and finance. In practice this improves decisions. The mechanics are straightforward. The operational impact is concrete. Teams can replicate the workflow. These steps reduce waste and risk. The approach scales across lines of business. Results remain interpretable and auditable. Governance is built into the design. The evidence base supports these choices. This closes the loop between health and finance. In practice this improves decisions. The mechanics are straightforward. The operational impact is concrete. Teams can replicate the workflow. These steps reduce waste and risk. The approach scales across lines of business. Results remain interpretable and auditable. Governance is built into the design. The evidence base supports these choices. This closes the loop between health and finance. In practice this improves decisions. The mechanics are straightforward. The operational impact is concrete. Teams can replicate the workflow. These steps reduce waste and risk. The approach scales across lines of business. Results remain interpretable and auditable. Governance is built into the design. The evidence base supports these choices. This closes the loop between health and finance. In practice this improves decisions. The mechanics are straightforward. The operational impact is concrete. Teams can replicate the workflow. These steps reduce waste and risk. The approach scales across lines of business. Results remain interpretable and auditable. Governance is built into the design. The evidence base supports these choices. This closes the loop between health and finance. In practice this improves decisions. The mechanics are straightforward. The operational impact is concrete. Teams can replicate the workflow. These steps reduce waste and risk. The approach scales across lines of business. Results remain interpretable and auditable. Governance is built into the design. The evidence base supports these choices. This closes the loop between health and finance. In practice this improves decisions. The mechanics are straightforward. The operational impact is concrete. Teams can replicate the workflow. These steps reduce waste and risk. The approach scales across lines of business. Results remain interpretable and auditable. Governance is built into the design. The evidence base supports these choices. This closes the loop between health and finance. In practice this improves decisions. The mechanics are straightforward. The operational impact is concrete. Teams can replicate the workflow. These steps reduce waste and risk. The approach scales across lines of business. Results remain interpretable and auditable. Governance is built into the design. The evidence base supports these choices. This closes the loop between health and finance. In practice this improves decisions. The mechanics are straightforward. The operational impact is concrete. Teams can replicate the workflow. These steps reduce waste and risk. The approach scales across lines of business. Results remain interpretable and auditable. Governance is built into the design. The evidence base supports these choices. This closes the loop between health and finance. In practice this improves decisions. The mechanics are straightforward. The operational impact is concrete. Teams can replicate the workflow. These steps reduce waste and risk. The approach scales across lines of business. Results remain interpretable and auditable. Governance is built into the design. The evidence base supports these choices. This closes the loop between health and finance. In practice this improves decisions. The mechanics are straightforward. The operational impact is concrete. Teams can replicate the workflow. These steps reduce waste and risk. The approach scales across lines of business. Results remain interpretable and auditable. Governance is built into the design. The evidence base supports these choices. This closes the loop between health and finance. In practice this improves decisions. The mechanics are straightforward. The operational impact is concrete. Teams can replicate the workflow. These steps reduce waste and risk. The approach scales across lines of business. Results remain interpretable and auditable. Governance is built into the design. The evidence base supports these choices. This closes the loop between health and finance. In practice this improves decisions. The mechanics are straightforward. The operational impact is concrete. Teams can replicate the workflow. These steps reduce waste and risk. The approach scales across lines of business. Results remain interpretable and auditable. Governance is built into the design. The evidence base supports these choices. This closes the loop between health and finance. In practice this improves decisions. The mechanics are straightforward. The operational impact is concrete. Teams can replicate the workflow. These steps reduce waste and risk. The approach scales across lines of business. Results remain interpretable and auditable. Governance is built into the design. The evidence base supports these choices. This closes the loop between health and finance. In practice this improves decisions. The mechanics are straightforward. The operational impact is concrete. Teams can replicate the workflow. These steps reduce waste and risk. The approach scales across lines of business. Results remain interpretable and auditable. Governance is built into the design. The evidence base supports these choices. This closes the loop between health and finance. In practice this improves decisions. The mechanics are straightforward. The operational impact is concrete. Teams can replicate the workflow. These steps reduce waste and risk. The approach scales across lines of business. Results remain interpretable and auditable. Governance is built into the design. The evidence base supports these choices. This closes the loop between health and finance. In practice this improves decisions. The mechanics are straightforward. Indeed.

Conclusion

Linking health and finance through an AI-ready analytics stack lets organizations forecast spend credibly, flag avoidable utilization, and protect the payment perimeter. The approach marries classical cost modeling with modern machine learning and treats calibration, privacy, and fairness as design constraints, not afterthoughts. Evidence from healthcare cost analytics, oncology surveillance, security reviews, and fraud detection supports the claim that well-governed predictive systems can improve outcomes while lowering waste (Hasan et al., 2025b; Hasan et al., 2021; Hasan et al., 2025a; Hasan et al., 2025c) [2, 4, 1, 3]. For executives and analytics leads, the priority now is execution: build the feature store, instrument calibration, publish model cards, and connect model outputs to pricing, outreach, and investigation decisions. Done well, the payoff is better health and better stewardship of scarce dollars. In practice this improves decisions. The mechanics are straightforward. The operational impact is concrete. Teams can replicate the workflow. These steps reduce waste and risk. The approach scales across lines of business. Results remain interpretable and auditable. Governance is built into the design. The evidence base supports these choices. This closes the loop between health and finance. In practice this improves decisions. The mechanics are straightforward. The operational impact is concrete. Indeed.

Limitations and Future Directions

This study has three limitations. First, we demonstrate methods with synthetic data rather than claims and EHR from a live deployment. That choice avoids privacy risk but underestimates data quality challenges, rare-event noise, and the politics of data access. Second, we focus on short-horizon forecasting and binary high-cost risk; alternative targets—care gaps closed, quality-adjusted life year yield, or multi-year tail risk—may require survival models, causal inference, or Bayesian decision analysis. Third, fairness auditing remains constrained by the attributes present; organizations need governance agreements that allow bias testing without enabling misuse. Despite these limits, the pipeline is ready for production with standard data and tools. Future work should pursue four threads. One, couple forecasting with causal policy simulation so leaders can test benefit changes and site-of-care steering in silico. Two, evaluate privacy-preserving learning—federated training and secure multiparty computation—to unlock cross-payer signal while honoring regulation (Hasan et al., 2022) [6]. Three, expand cyber-risk features and feedback loops given the evolving threat landscape in payment and provider systems (Hasan et al., 2025a; Hasan et al., 2025c; Milon et al., 2024; Arman et al., 2023). Four, study adoption directly: embed UTAUT-informed nudges and change-management into product rollouts and measure behavior change end to end (Ghose et al., 2025) [8]. Progress on these fronts would make health–finance analytics both safer and more valuable. In practice this improves decisions. The mechanics are straightforward. The operational impact is concrete. Teams can replicate the workflow. These steps reduce waste and risk. The approach scales across lines of business. Results remain interpretable and auditable. Governance is built into the design. The evidence base supports these choices. This closes the loop between health and finance. In practice this improves decisions. The mechanics are straightforward. The operational impact is concrete.

References

1. Hasan MN, Papel MSI, Rasel IH, Akter S, Aktar MK, Abedin MZ, *et al.* Enhancing financial information security through advanced predictive analytics: a PRISMA-based systematic review. Edelweiss Appl Sci Technol. 2025a;9(7):2222-45. doi:10.55214/2576-8484.v9i7.9142
2. Hasan MN, Arman M, Bhuyain MMH, Chowdhury F, Bathula MK. Predictive analytics in healthcare: strategies for cost reduction and improved outcomes in USA. Int J Innov Res Sci Stud. 2025b;8(8):142-50. doi:10.53894/ijirss.v8i8.10559
3. Hasan MN, Rasel IH, Arman M, Ibrahim M, Jahan N. Strengthening U.S. financial and cybersecurity

infrastructure with AI-driven fraud detection and risk analytics. J Comput Anal Appl. 2025c;31(2):15-32. Available from: https://eudoxuspress.com/index.php/pub/article/view/3823

4. Hasan MN, Bhuyain MMH, Chowdhury F, Arman M. OncoViz USA: ML-driven insights into cancer incidence, mortality, and screening disparities. J Med Health Stud. 2021;2(1):53-62. doi:10.32996/jmhs.2021.2.1.6
5. Rasel IH, Arman M, Hasan MN, Bhuyain MMH. Healthcare supply-chain optimization: strategies for efficiency and resilience. J Med Health Stud. 2022;3(4):171-82.
6. Hasan N, Rasel IH, Rahman M, Islam K, Arman M, Jahan N. Securing U.S. healthcare infrastructure with machine learning: protecting patient data as a national security priority. Int J Comput Exp Sci Eng. 2022;8(3). doi:10.22399/ijcesen.3987
7. Arman M, Hasan MN, Rasel IH. Clean energy transition in USA: big data analytics for renewable energy forecasting and carbon reduction. J Manag World. 2024(3):192-206. doi:10.53935/jomw.v2024i4.1196
8. Ghose P, Bhuiyan MRI, Hasan MN, Rakib SH, Mani L. Mediated and moderating variables between behavioral intentions and actual usages of fintech in the USA and Bangladesh through the extended UTAUT model. Int J Innov Res Sci Stud. 2025;8(2):113-25. doi:10.53894/ijirss.v8i2.5130
9. Milon NU, Ghose P, Pinky TC, Tabassum MN, Hasan MN, Khatun M. An in-depth PRISMA-based review of cybercrime in a developing economy: examining sector-wide impacts, legal frameworks, and emerging trends in the digital era. Edelweiss Appl Sci Technol. 2024;8(4):2072-93. doi:10.55214/25768484.v8i4.1583
10. Arman M, Fahim ASM. AI revolutionizes inventory management at retail giants: examining Walmart's U.S. operations. J Bus Manag Stud. 2023;5(6):145-8. doi:10.32996/jbms.2023.5.6.15
11. Khan SA, Shah A, Arman M. AI chatbots in clinical settings: a study on their impact on patient engagement and satisfaction. J Manag World. 2024(3):207-13. doi:10.53935/jomw.v2024i4.1201
12. Arman M, Rasel IH, Razib MNH, Fahim ASM. Big data and machine learning for sustainable waste reduction. J Posthumanism. 2024;4(2):448-67. doi:10.63332/joph.v4i2.3361
13. Shah A, Khan SA, Arman M. Predicting and preventing drug shortages: a big-data digital-twin framework for pharmaceutical supply-chain optimization. J Econ Finance Account Stud. 2024;6(6):116-26. doi:10.32996/jefas.2024.6.6.9
14. Shah A, Arman M, Khan SA. Patient-centric marketing and retention strategies in healthcare: a strategic and technological framework. J Bus Manag Stud. 2025;7(2):239-48. doi:10.32996/jbms.2025.7.2.17

How to Cite This Article

Miller EJ, Thompson OR, Brooks DK. AI-driven insights at the intersection of health and finance: modeling medical expenditures and risk using big data analytics. J Front Multidiscip Res. 2025;6(2): 565-573. doi: 10.54660/.JFMR.2025.6.2.565-573

Creative Commons (CC) License