EC PULMONOLOGY AND RESPIRATORY MEDICINE
Review Article

AI-Driven Pricing in Pulmonary Care: Machine Learning Approaches for Optimizing Provider Revenue and Enhancing Patient Cost Transparency

Rahul Renaaud[1] and Alok S Shah[2]*

[1]Hamilton Southeastern High School, Fishers, Indiana, USA

[2]Pulmonary Department, University of Chicago, USA

***Corresponding Author:** Alok S Shah, Pulmonary Department, University of Chicago, USA.

Received: August 20, 2025; **Published:** September 08, 2026

Abstract

Pulmonary disease imposes one of the largest and fastest-growing economic burdens in modern healthcare, yet the pricing of respiratory care remains largely opaque to the patients who bear its costs. This review examines how machine learning (ML) can simultaneously optimize provider revenue and reduce patient financial burden in pulmonary medicine, advancing the non-zero-sum framing of AI-driven pricing as a hypothesis to be tested rather than a demonstrated result. We synthesize evidence across four domains: cost prediction and risk stratification, revenue cycle management, patient-facing transparency tools, and the surrounding ethical and regulatory landscape. Ensemble methods such as XGBoost and random forest consistently achieve areas under the receiver operating characteristic curve (AUC) of 0.80–0.88 for identifying high-cost respiratory patients from routinely collected claims data, and severe chronic obstructive pulmonary disease (COPD) exacerbation prediction now reaches an AUC of 0.866 on retrospective data, identifying the high-risk patients who could be targeted for interventions against hospitalizations costing $18,000–$32,000 per episode, though no prospective study has yet confirmed that such targeting produces net savings. Natural language processing is progressively automating clinical coding, and 80% of health systems now report exploring or implementing generative AI in revenue cycle operations. In contrast, patient-facing transparency applications remain underdeveloped, and cost-based prediction targets risk encoding racial and socioeconomic bias into pricing decisions. We conclude that the strongest current evidence supports COPD cost prediction, exacerbation prevention, and coding automation, and we outline research priorities including federated learning, prospective validation, and routine equity auditing of pricing algorithms.

Keywords: *Machine Learning; Artificial Intelligence; Pulmonary Medicine; Healthcare Costs; Price Transparency; Revenue Cycle Management; COPD; Financial Toxicity; Algorithmic Bias*

Introduction

The economic burden of pulmonary disease

Respiratory disease occupies a singular position in global health economics. COPD was the third leading cause of death worldwide in 2019, accounting for 3.23 million deaths, and its direct costs are dominated by hospitalization and acute care services [1]. The disease affects more than 250 million people globally, and both healthcare resource utilization and costs rise substantially with disease severity and exacerbation frequency [2]. Systematic analyses across Europe, Asia, and the Americas confirm that exacerbation frequency and disease severity are the primary cost drivers regardless of the healthcare financing system in which patients are treated [3].

The trajectory of this burden is steep. A recent open cohort Markov model projected cumulative global direct costs of $24.35 trillion and indirect costs of $15.43 trillion for COPD between 2025 and 2050, with 15.60 billion exacerbations expected over that period [4]. Prevalence modeling tells a similar story: global COPD cases among adults aged 25 years and older are projected to increase by 23% from 2020 to 2050, approaching 600 million patients, with growth concentrated among women (a 47.1% increase, compared with 9.4% for men) and in low- and middle-income countries [5]. These projections indicate that pulmonary care will become one of the dominant chronic disease cost challenges of the coming decades.

Within this aggregate burden, a small set of events and services drives most spending. Severe COPD exacerbations requiring emergency department visits or inpatient stays account for 90.3% of total COPD-related medical costs, estimated at $32.1 billion annually in the United States [6]. Critical care is a second concentration point: patients requiring mechanical ventilation incur mean intensive care unit (ICU) costs of $31,574, compared with $12,931 for non-ventilated patients [7]. At the individual level, the consequences are increasingly described as financial toxicity; among patients undergoing lung cancer resection at a major cancer center, 10.0% of surveyed patients experienced major financial toxicity after surgery [8]. The concentration of costs in predictable, high-acuity events is precisely the structure that machine learning is well suited to exploit.

The pricing transparency crisis in pulmonary medicine

While costs rise, the prices underlying them remain difficult for patients to see or compare. The Centers for Medicare and Medicaid Services (CMS) Hospital Price Transparency Rule, effective January 1, 2021, requires hospitals to post standard charges for all items and services, yet only approximately one-third of facilities had complied within the first 10 months of implementation [9]. Enforcement has since intensified: the U.S. Government Accountability Office (GAO) documented 1,287 CMS enforcement actions from 2021 through 2023 and over $4 million in civil monetary penalties, while also noting that rising hospital service prices contributed to a nearly 50% increase in private health plan spending from 2012 through 2022 [10].

A parallel regulatory response addresses surprise billing, a problem with particular relevance to respiratory medicine because emergency care for respiratory distress is a common setting in which patients receive out-of-network treatment without meaningful consent [11]. The No Surprises Act (NSA), effective January 1, 2022, prohibits surprise out-of-network billing for emergency services and limits patient cost-sharing to in-network rates [11]. Before the NSA, patient responsibility averaged $628 for emergency department surprise bills and $2,040 for inpatient admissions [12]. Yet even where pricing data now exist, they are used primarily by health systems and insurers for contract negotiation rather than by patients, who face data buried in complex spreadsheets that require deep billing code knowledge to interpret [9]. This gap between data availability and data usability defines the transparency crisis that artificial intelligence (AI) tools are increasingly asked to close.

Scope and objectives

This review examines the intersection of AI, pricing, and pulmonary care, asking whether ML can simultaneously optimize provider revenue and reduce patient financial burden. Our central proposition—advanced here as a hypothesis to be tested, not a conclusion the current evidence establishes—is that these goals are not inherently in conflict: the same predictive infrastructure that prevents revenue leakage for providers could also prevent financial surprise for patients, making ML-driven pricing potentially non-zero-sum. As Section 7 makes explicit, this dual-benefit claim has not been prospectively validated; the review's purpose is to assess how close the evidence comes and what would be required to close the remaining gap. The scope encompasses pulmonary-specific applications across COPD, lung cancer, asthma, critical care, and sleep medicine. The review is organized across four dimensions: cost prediction and risk stratification (Section 3), revenue cycle management (Section 4), patient-facing transparency tools (Section 5), and the challenges, ethical tensions, and research gaps that condition all three (Sections 6 and 7). Section 2 first establishes the theoretical framework linking the healthcare revenue cycle to ML methodology.

This synthesis was conducted as a narrative rather than systematic review. Relevant literature was identified through structured searches of PubMed, Scopus, Web of Science, and Google Scholar covering 2005 through 2026, combining terms for pulmonary disease (COPD, lung cancer, asthma, mechanical ventilation, respiratory failure) with terms for machine learning, artificial intelligence, healthcare cost prediction, revenue cycle management, price transparency, and algorithmic bias. Priority was given to peer-reviewed primary studies and systematic reviews. A small number of government reports and industry surveys were included where they supplied current, quantitative data not yet available in the peer-reviewed literature, and these sources are identified as such at the point of citation. Because the field is young and pulmonary-specific pricing studies remain sparse, no formal inclusion-exclusion protocol or quantitative pooling was applied; the objective is interpretive breadth across the pricing pipeline rather than exhaustive enumeration, and the strength of the underlying evidence is appraised critically throughout.

Theoretical framework: The pulmonary pricing pipeline

The healthcare revenue cycle in pulmonary medicine

The healthcare revenue cycle spans patient registration, charge capture, clinical coding, claim submission, payment posting, and collections, and revenue can leak at every stage. An industry analysis of hospital revenue operations identifies leakage points across front-end processes (registration errors and eligibility verification failures), mid-cycle processes (charge capture gaps and coding inaccuracies), and back-end processes (underpayments and failed appeals) [13]. Pulmonary medicine is exposed to these failure modes with unusual intensity because its billing spans multiple settings and rule sets: multi-specialty ICU billing, bundling rules for pulmonary function testing, and time-based critical care documentation all coexist within a single service line, from routine spirometry (Current Procedural Terminology [CPT] code 94010) to bronchoscopy (CPT 31622) and ventilator management.

Clinical coding is the cycle's structural bottleneck. The ICD-10-CM system contains approximately 68,000 diagnosis codes, and manual coding accuracy ranges from 50% to 98% with a median of 80%, depending on coder expertise, diagnosis complexity, and patient acuity [14]. Human coders require months of training, process roughly 60 cases per day, and still leave backlogs that often extend for months [14]. For a specialty in which the difference between coding an acute exacerbation versus stable chronic COPD materially changes reimbursement, this combination of volume, complexity, and error tolerance creates both the financial motivation and the technical opportunity for automation.

Machine learning taxonomy for healthcare pricing

The ML techniques applied to healthcare pricing fall into a small number of families. Supervised learning trains models on labeled examples to predict outcomes such as future cost or denial risk; unsupervised learning discovers hidden structure without labels; and reinforcement learning optimizes sequential decisions through trial and error [15]. Within supervised learning, ensemble tree-based methods—random forest, gradient boosting, and XGBoost—have become the dominant approaches for healthcare cost modeling, a pattern that recurs across nearly every study reviewed in Section 3. For text, automated coding systems have evolved through three generations: rule-based systems, conventional ML classifiers such as support vector machines and random forests, and deep learning architectures [16]. Recent work increasingly employs convolutional and recurrent neural networks and transformer-based models, although performance on rare codes and long documents remains challenging [17].

A final methodological strand concerns interpretability. Because pricing predictions influence clinical and financial decisions, explainable AI methods—including automatically generated rule-based explanations that identify the specific factors driving an individual patient's risk—are increasingly treated as prerequisites for deployment rather than optional refinements [18]. Regulatory pressure reinforces this direction; the European General Data Protection Regulation (GDPR) has been credited with driving development of more transparent AI models in healthcare [19].

Conceptual model

These elements can be integrated into a single conceptual pipeline. On the provider side, ML operates on claims, electronic health record (EHR), and remittance data to predict costs, assign codes, prevent denials, and detect underpayments. On the patient side, the same predictive infrastructure can generate personalized cost estimates, flag financial toxicity risk, and translate machine-readable pricing files into actionable information. A recent conceptual framework makes this connection explicit, proposing to link ML pricing analytics directly to the regulatory mandates of the CMS transparency rule and the NSA and suggesting how algorithms might assess pricing elasticity to determine billing thresholds that balance revenue generation with patient accessibility [20]; this proposal remains conceptual and has not been empirically validated. The stakeholders in this pipeline—providers, patients, payers, regulators, and data scientists—hold distinct and sometimes competing interests, a tension examined in Section 6. The following sections review the empirical evidence for each stage of the pipeline.

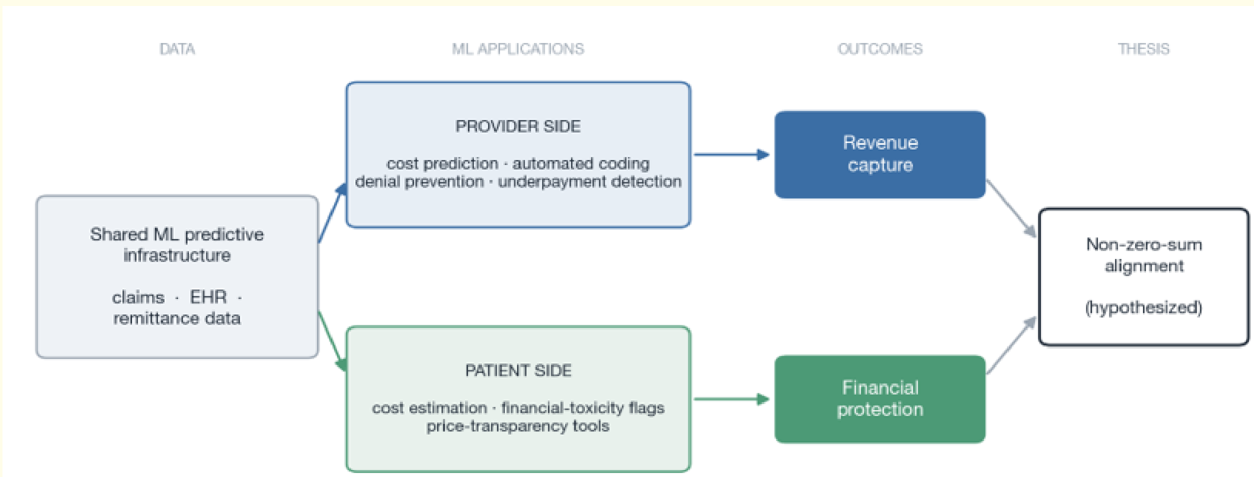

Figure 1

ML for pulmonary cost prediction and risk stratification

Predicting high-cost COPD patients

Claims-based cost prediction models

The most direct evidence for ML-driven pulmonary pricing comes from models that identify high-cost COPD patients in routine administrative data. Luo., *et al*. [21] applied logistic regression, random forest, and XGBoost to medical insurance data from a large city in western China, using 50 attributes from COPD admission records collected between 2011 and 2013. XGBoost achieved the highest discrimination (AUC 0.801), compared with 0.792 for random forest and 0.787 for logistic regression, and a reduced model using only the top 30 ranked variables retained an AUC of 0.798, suggesting that parsimonious, deployable feature sets are feasible [21]. Comparable performance has been demonstrated in European claims data: Langenberger., *et al*. [22] compared four algorithms in German sickness fund data and found that random forest performed best (AUC 0.883), with prior-year healthcare costs, number of hospitalizations, and chronic disease burden as the strongest predictors. In Korea's single-payer system, an XGBoost model trained on National Health Insurance Service check-up and claims data predicted next-year high-cost status with 77.1% accuracy, and respiratory disease was among the disease groups most strongly associated with high future expenditure [23].

These results warrant tempered enthusiasm. A systematic review and meta-analysis of ML and deep learning prognostic models in COPD found pooled C-statistics of only 0.624–0.769 for regression-based models and identified incorrect handling of missing data, unreported model uncertainty, and datasets too small for the number of predictors as pervasive sources of bias [24]. Notably, the review found limited evidence that ML models outperform existing disease severity scores, attributing this partly to methodological weaknesses rather than inherent model failure, and called for stricter adherence to TRIPOD and PROBAST reporting standards [24].

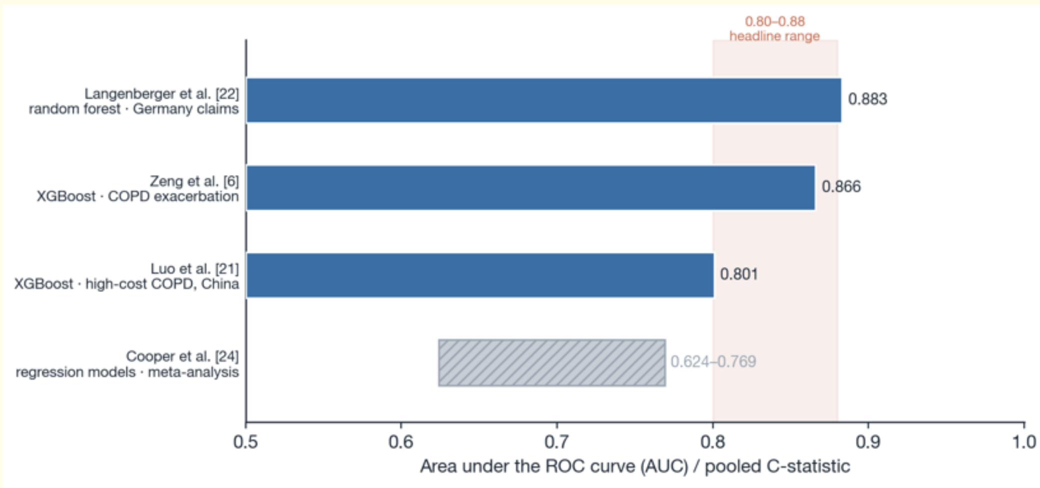

Figure 2

Exacerbation prediction as financial prevention

Because severe exacerbations account for 90.3% of COPD-related medical costs [6], predicting them is financially equivalent to predicting cost itself. Zeng., *et al.* [6] developed an XGBoost-based model using 43,576 data instances from University of Washington Medicine, screening 278 candidate features and 39 classification algorithms; the final model achieved an AUC of 0.866 (95% CI 0.838–0.892) for severe exacerbation within 12 months, exceeding all previously published models. Using the top 10% of predicted risk as a threshold, the model reached a sensitivity of 56.6%, specificity of 91.17%, and negative predictive value of 98.83% [6]. The economic logic is direct, but it remains a projection rather than a demonstrated outcome: many severe exacerbations are preventable with appropriate outpatient management, and each avoided hospitalization would save an estimated $18,000–$32,000 [6]. Two inferential steps separate the reported metric from that saving, however—an AUC of 0.866 measures how well the model ranks patients by retrospective risk, not how many admissions a prevention program actually averts, and no prospective study has yet closed that gap (Section 7.1).

Translating such predictions into action requires interpretability and capacity planning. Luo., *et al.* [18] extended this model with an automatic rule-based explanation method that identifies the specific factors driving each patient's predicted risk and suggests tailored interventions such as correct inhaler use, peak flow monitoring, and pulmonary rehabilitation enrollment. Their deployment scenario is instructive for pricing applications: the care management program could accommodate only 3% of patients, so the model's role was to surface the top 10% at highest risk, from which care managers selected enrollees after chart review [18]. ML risk scores, in other words, do not allocate resources autonomously; they concentrate scarce clinical and financial attention where it prevents the most cost.

Predicting in-hospital charges for respiratory admissions

A complementary line of work predicts the magnitude of charges once admission occurs. Arnold., *et al.* [25] built six ML models—linear regression, ridge regression, support vector machine, random forest, gradient boosting, and XGBoost—to predict total in-hospital charges for COPD exacerbations, congestive heart failure exacerbations, and diabetic ketoacidosis using 26,190 national discharge records from 2016 to 2019. In the COPD cohort (n=9,552), where the average exacerbation hospitalization cost $16,794 (±$22,003), R-squared values ranged from 0.680 to 0.710 in training and 0.694 to 0.724 in testing, with remarkably consistent performance across all six algorithms [25]. Variable importance analyses identified mechanical ventilation duration, length of stay, and comorbidity index as the dominant cost drivers [25]. The replication of one pipeline across three high-burden conditions offers pulmonary departments a directly transferable template for charge forecasting.

Financial toxicity prediction in lung cancer

Cost prediction can also be oriented toward the patient's finances rather than the provider's. Deboever., *et al*. [8] conducted the first study applying ML to predict financial toxicity in thoracic surgery, analyzing 1,477 patients who underwent lung cancer resection at MD Anderson Cancer Center between 2016 and 2021, of whom 462 (31.3%) completed a financial toxicity survey and 46 (10.0%) reported major financial toxicity. The investigators ensembled decision tree, random forest, gradient boosting, and XGBoost predictions and identified the most informative preoperative features, including age, race and ethnicity, smoking status, household income, credit score, marital and employment status, tumor histology, extent of resection, and preoperative lung function [8]. The clinical implication is that financial vulnerability can be identified before surgery, creating a window for financial counseling, insurance navigation, and payment plan initiation before costs accumulate [8].

Longer-horizon modeling remains less developed. Predicting out-of-pocket cost trajectories across the full course of advanced lung cancer—integrating treatment modality, insurance design, and disease progression—would effectively pair a financial prognosis with the clinical one, but published pulmonary-specific models of this kind are scarce. This gap is revisited in Section 7.

ICU and mechanical ventilation cost prediction

Critical care provides the granular cost structure on which ML forecasting depends. In the landmark U.S. analysis of 51,009 ICU patients across 253 hospitals, Dasta., *et al*. [7] found that daily costs were highest on ICU day 1 ($10,794 for ventilated versus $6,667 for non-ventilated patients), decreased sharply on day 2, and stabilized after day 3 at $3,968 and $3,184, respectively; after adjustment for patient and hospital characteristics, mechanical ventilation added a mean incremental cost of $1,522 per day. European microcosting corroborates the pattern: in a German university hospital cohort of 10,637 ICU patients, daily costs were EUR 999 for non-ventilated and EUR 1,590 for ventilated care, a 59% increase, with the differential attributable primarily to ventilation itself rather than underlying disease severity [26].

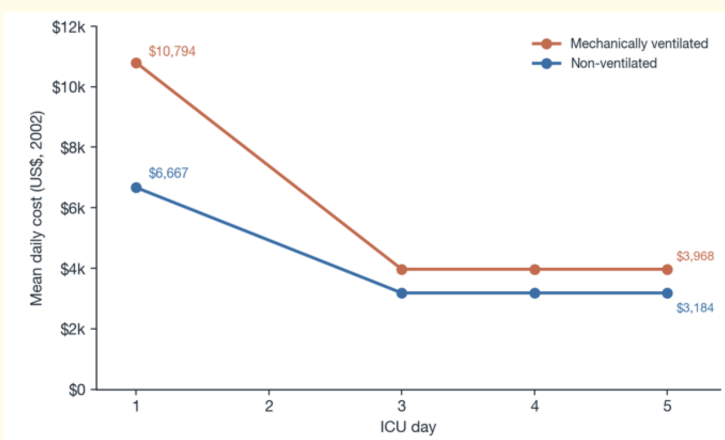

Figure 3

These day-by-day trajectories make ventilation duration and length of stay the natural prediction targets for ML-based financial planning, and the charge prediction literature confirms that ventilation duration is the single most important cost driver in respiratory admissions [25]. EHR-trained models have demonstrated the feasibility of predicting prolonged length of stay and 30-day readmission [15], and applying such models to respiratory failure populations connects clinical prediction directly to resource allocation, capacity planning, and cost avoidance in pulmonary critical care. Having established what ML can predict, the review now turns to how those predictions are being embedded in the billing infrastructure itself.

AI in pulmonary revenue cycle management

Automated clinical coding for respiratory services

NLP for ICD-10 and CPT code assignment

Automated clinical coding is the most mature AI application in the revenue cycle. Systematic reviews document a clear technical evolution: an analysis of 38 studies from 2010 to 2020 catalogued the progression from rule-based systems through conventional classifiers to deep learning, with an AdaBoost classifier achieving an F-score of 0.9141 for ICD-10-AM classification in an Australian hospital dataset [16]. A subsequent review of 42 studies confirmed the shift toward transformer-based architectures capable of capturing complex relationships between clinical text and ICD codes, while noting persistent difficulty with rare codes, long documents, and heavy reliance on the English-language MIMIC-III benchmark [17]. The most recent synthesis identified six recurring challenges—extensive label space, imbalanced label distribution, lengthy clinical documents, interpretability requirements, ethical concerns, and lack of transparency—and observed that automated coding systems have not yet been widely adopted in practice, largely because real-time performance remains unproven [27].

For pulmonary medicine, the stakes of solving these problems are high. Respiratory coding requires distinctions that are semantically subtle but financially consequential: acute exacerbation versus chronic stable COPD, sleep study staging, and critical care time documentation all hinge on precise language in clinical notes. Given that manual coding accuracy has a median of only 80% [14], NLP systems that reliably parse such distinctions could simultaneously raise coding accuracy and release coder capacity for complex cases.

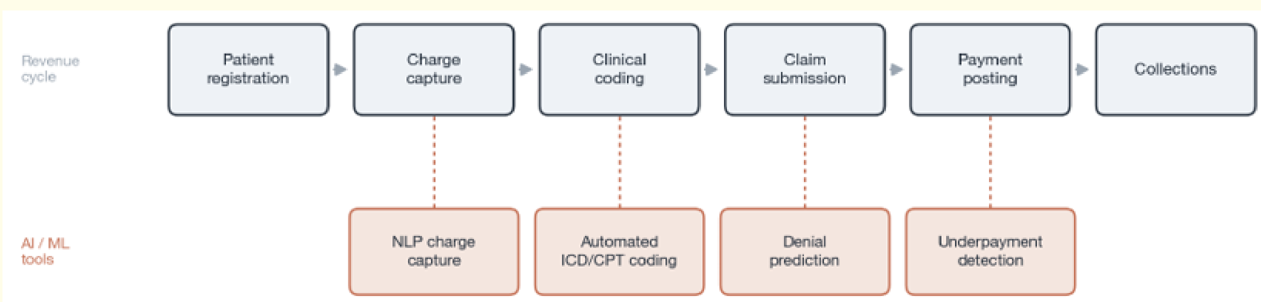

Figure 4

Reducing claim denials through pre-submission AI

A second application intervenes before claims leave the building. ML models can forecast the denial probability of a claim prior to submission, allowing staff to correct coding errors, missing modifiers, and documentation gaps pre-emptively; one hospital reported a 19% reduction in denial rates within six months of implementing denial prediction tools [13]. This capability is diffusing rapidly through an infrastructure that already exists: 65% of U.S. hospitals use predictive models, 79% of those obtain models through their EHR vendor, and billing optimization is among the reported use cases alongside clinical decision support and scheduling [28]. The same survey revealed substantial variation in whether hospitals locally evaluate their models for accuracy and bias [28], a governance gap with direct implications for the equity concerns discussed in Section 5.2.

Revenue recovery and contract compliance

Underpayment detection and charge capture

Beyond denials, AI addresses revenue that is earned but never collected. NLP analysis of physician notes and operative reports can identify services that were performed but not billed, while AI-enabled contract compliance tools analyze remittance data to detect systematic underpayments and payer deviations from negotiated rates [13]. Because leakage accumulates across front-end, mid-cycle, and back-end processes simultaneously [13], these tools are most valuable when deployed as a portfolio rather than in isolation—an argument for the integrated pipeline framing introduced in Section 2.3.

Generative AI versus conventional ML in the revenue cycle

The arrival of large language models (LLMs) has forced a build-versus-buy calculation. Burns., *et al*. [29] compared locally developed Clinical-BigBird transformer models against commercial GPT-4 for medical free-text classification in revenue cycle operations, training domain-specific models on 149,702 chronic kidney disease notes and 94,965 heart failure notes; the local, domain-specific models demonstrated superior accuracy at substantially lower cost. The economics are non-trivial in both directions: on-premise GPU servers can exceed $200,000, cloud GPU services run upwards of $40 per hour, and development team payroll exceeds $100,000 per person annually [29]. Adoption is nonetheless accelerating. A 2025 industry survey of 519 chief financial officers and revenue cycle leaders found that 80% of health systems were exploring, piloting, or implementing generative AI for revenue cycle management—up from 58% in 2023—with roughly 40% already piloting or implementing, and with coding optimization and revenue capture as the primary use cases; cost and budget constraints were the largest reported obstacles [30].

International applicability and the economics of implementation

The evidence base for ML-driven pricing is notably international, spanning fee-for-service claims in China [21], statutory sickness funds in Germany [22], single-payer data in Korea [23], and emerging adoption in Saudi Arabia, where reviews describe AI assessing pricing elasticity to set billing thresholds that balance revenue with patient accessibility [31]. This breadth suggests that the underlying methods transfer across reimbursement architectures, even as the prediction targets differ. Large aggregate savings projections are frequently invoked in this literature, but they are typically extrapolations, and rigorous economic evaluation lags far behind such claims. A systematic review of cost-effectiveness studies of AI in healthcare found that very few publications had thoroughly assessed economic impact, that most failed to consider whether equivalent savings could have been achieved by non-AI methods, and that nearly all measured resource use through top-down approaches likely to miss stakeholder-level cost drivers [32]. Model-based projections indicate that AI-assisted diagnosis can reduce costs through time savings and faster throughput [19], but for pulmonary departments weighing investment, the honest summary is that implementation costs are well characterized [29] while returns remain largely extrapolated.

ML-driven transparency tools for pulmonary patients

Making pricing data actionable

The provider-side sophistication documented above contrasts sharply with the tools available to patients. Hospital machine-readable pricing files exist because regulation requires them, but the data are buried in complex spreadsheets, demand billing code fluency, lack standardization across hospitals, and are consumed primarily by insurers and health systems for negotiation [9]. Data quality compounds the problem: health plans, employers, and researchers have reported unusually high or low posted prices that appear to be errors, prompting CMS in 2024 to mandate standardized file formats and hospital affirmations of accuracy [10]. AI-driven frameworks have been proposed to bridge this gap by aggregating and standardizing pricing files and generating patient-comprehensible cost analytics [20], and the NSA's good faith estimate requirement creates a natural insertion point for ML-generated, personalized cost predictions that incorporate insurance coverage and deductible status [11]. Whether such estimators actually change patient decision-making or shopping behavior, however, remains essentially untested—a gap noted in Section 7.

Algorithmic bias and financial equity

Any ML system trained on cost data inherits the inequities embedded in historical spending. The canonical demonstration is Obermeyer., *et al*. [33], who showed that a widely used commercial algorithm predicting healthcare costs as a proxy for health need systematically disadvantaged Black patients: because less money had historically been spent on Black patients at equivalent illness levels, Black patients at a given risk score were considerably sicker, and correcting the bias would have raised the share of Black patients flagged for additional care management from 17.7% to 46.5%. Substituting cost proxies with direct measures of health need substantially reduced the bias [33]. For pulmonary pricing models—whose prediction targets are, by definition, costs—this finding is not peripheral but foundational: target choice is an equity decision.

The bias problem extends beyond race. A case study using the housing-based HOUSES index found that asthma exacerbation prediction models performed worse for children of lower socioeconomic status, with disparities exceeding those associated with age, sex, or race and ethnicity, driven mechanistically by less complete and less accurate EHR data among disadvantaged patients [34]. Conceptual frameworks now catalogue the sources of such bias—training data reflecting historical inequities, proxy prediction targets, and differential subgroup performance—and recommend routine bias audits and transparent subgroup reporting [35]. Insurers, whose data ecosystems feed many pricing models, face parallel calls for "algorithmoveillance," and accreditation bodies such as the National Committee for Quality Assurance have begun incorporating bias evaluation into standards [36]. Encouragingly, mitigation is practical: in the largest U.S. public health system, threshold adjustment reduced subgroup disparities in deployed asthma and readmission models to below 5 percentage points of equal opportunity difference while sacrificing less than 10% accuracy, and the authors released a replicable playbook for low-resource settings [37].

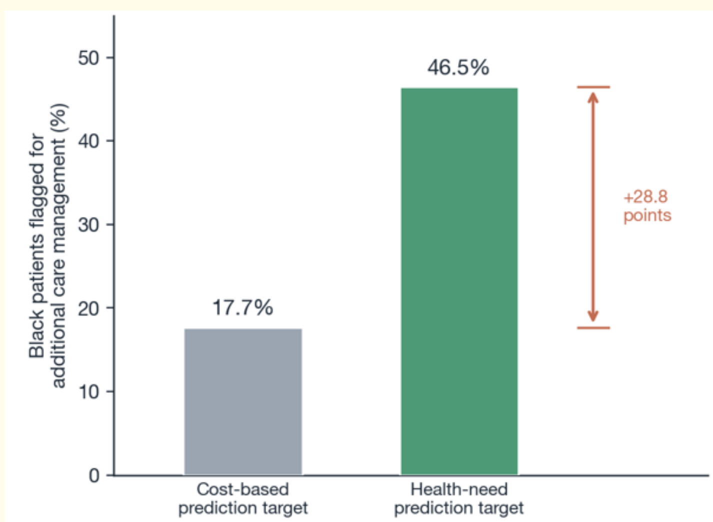

Figure 5

The No surprises act and regulatory impact

Regulation defines the boundary conditions for patient-facing pricing AI. The NSA's protections apply directly to emergency pulmonary care and ICU services, where out-of-network exposure was historically concentrated [11]. Federal evaluation shows that out-of-network professional claims were already declining before implementation, from 6.0% to 4.7% between 2012 and 2020, and that more than 200 million privately insured Americans stand to benefit from the law's provisions [12]. Early results are nonetheless equivocal: independent dispute resolution decisions have often resembled previous out-of-network payment amounts, raising doubt about net cost reduction [12], and a quasi-experimental difference-in-differences analysis found that patient healthcare spending may have remained stagnant, potentially because providers raised in-network charges to offset lost out-of-network revenue [38]. For AI developers, these findings pose a design question with ethical weight: pricing tools can be built to complement transparency mandates—or to help providers optimize around them.

Challenges and ethical considerations

Data and technical barriers

The most fundamental constraint on ML-driven pricing is that the relevant data—clinical, billing, and insurance—sit in silos, and privacy concerns restrict the access needed to train robust models [39]. Hospital pricing files remain inconsistently formatted and of uncertain accuracy despite CMS schema guidance [10]. Within the modeling literature itself, small samples relative to feature counts, imbalanced datasets, and mishandled missing data are the dominant sources of bias in pulmonary prognostic models [24], while automated coding contends with a label space of tens of thousands of codes distributed with extreme imbalance [27]. Generalizability compounds these problems: model performance degrades when applied to populations that differ from the training data, which is precisely the situation whenever a pricing model trained at one institution is deployed at another [15].

Privacy and security

Pricing models are unusual in requiring the joint processing of clinical, financial, and demographic data, which raises de-identification and re-identification challenges beyond those of purely clinical AI. Deployment architecture is the pivotal privacy decision: cloud-based LLM processing routes sensitive billing and clinical text through external infrastructure, whereas on-premise models keep data within the institution at higher fixed cost [29]. Regulatory frameworks push in the same direction from different jurisdictions—HIPAA governing U.S. patient billing data, and the GDPR shaping model transparency expectations in Europe [19]. Privacy-preserving training paradigms, discussed in Section 7, offer a technical route through this impasse [39].

Ethical tensions

Beneath the technical challenges lies a structural tension: the same models that optimize provider revenue could, deployed differently, simply extract more from patients. The transparency literature already documents a version of this paradox, with pricing data serving insurer and health system negotiation more than patient empowerment [9]. Cost-based prediction targets can silently encode discrimination [33], and the training data feeding these models over-represent a narrow slice of the world—more than half of published clinical AI models draw on datasets from the United States or China, and underrepresented groups suffer both degraded performance and algorithmic underestimation [40]. The populations bearing the heaviest respiratory burden are often those least represented: 92% of the world's population lives in regions exceeding World Health Organization air pollution guidelines, linking environmental injustice directly to the patients whom biased pricing algorithms would most affect [41]. Finally, automation raises workforce questions for the clinical coders and billing staff whose work these systems absorb, a transition that reviews of coding automation flag as an ethical concern in its own right [14]. None of these tensions is an argument against ML-driven pricing; they are arguments for governing it deliberately.

Research gaps and future directions

Knowledge gaps

Three gaps dominate the current literature. First, pulmonary-specific pricing studies are scarce; most cost prediction work is disease-general or system-wide, leaving the field to extrapolate from adjacent evidence. Second, virtually no AI pricing intervention has undergone prospective validation, and rigorous economic evaluation—including comparison against non-AI alternatives—is largely absent [32], with the methodological weaknesses of existing pulmonary models well documented [24]. Third, patient-facing applications are underexplored relative to provider-side tools: automated coding systems remain unadopted in routine practice [27], and no substantial evidence yet shows whether cost estimators alter patient decision-making. Standardized outcome metrics for evaluating pricing tools—capturing both provider margin and patient financial wellbeing—do not yet exist.

Technological frontiers

Federated learning offers the most direct path around the data fragmentation that currently caps model quality, enabling multi-institutional training without centralizing sensitive billing data, with topologies ranging from hub-and-spoke aggregation to peer-to-peer exchange and mechanisms such as differential privacy guarding against memorization of individual records [39]. This is particularly promising for pulmonary pricing, where no single institution observes enough rare, high-cost trajectories—prolonged ventilation, lung transplantation, advanced lung cancer—to model them robustly [39]. Multimodal integration represents a second frontier: combining imaging, genomics, wearable, and billing data could yield cost predictions grounded in physiology rather than utilization history alone, consistent with the broader convergence of human and artificial intelligence anticipated across medicine [42]. Domain-specific language models trained on respiratory billing corpora, which already outperform generalist LLMs at lower cost in revenue cycle tasks [29], are a natural near-term investment.

Policy and system-level priorities

The demand curve makes these questions urgent. Air pollution and smoking trends are projected to drive trillions of dollars in COPD costs through 2050 [4], with growth concentrated in low- and middle-income countries where pricing infrastructure is weakest [5]—settings that will require deliberately adapted, resource-appropriate AI tools rather than transplanted ones, built with the open science practices that mitigate dataset bias [41]. Longitudinal studies should track whether AI pricing tools improve provider margins and patient financial outcomes simultaneously, since the dual-benefit thesis is ultimately an empirical claim. Realizing it will require a genuinely interdisciplinary agenda spanning pulmonologists, data scientists, health economists, ethicists, and regulators.

Conclusions

This review examined how machine learning is reshaping the economics of pulmonary care on both sides of the ledger. The strongest current evidence supports three applications: claims-based identification of high-cost COPD patients, where ensemble methods consistently achieve AUCs of 0.80–0.88 [21,22]; severe exacerbation prediction as financial prevention, where an AUC of 0.866 identifies the patients who could be targeted to avert $18,000–$32,000 hospitalizations, pending prospective confirmation that such targeting yields net savings [6]; and NLP-based coding and denial automation, which is already diffusing through health systems at scale [28,30]. Provider-side applications are markedly more mature than patient-facing ones—an asymmetry that, left uncorrected, would make AI a tool for revenue optimization alone.

The central argument of this review is that such an outcome is neither inevitable nor economically necessary. The same predictive infrastructure that prevents revenue leakage can generate honest cost estimates, flag financial toxicity before surgery, and route vulnerable patients to assistance programs; provider and patient financial interests can be aligned rather than traded off. Achieving that alignment requires deliberate choices: prediction targets that measure need rather than historical spending [33], routine equity audits with practical mitigation [37], privacy-preserving architectures for multi-institutional learning [39], and prospective evaluation that treats patient financial wellbeing as a primary outcome. As the global burden of respiratory disease accelerates toward its projected mid-century peak [4,5], the question is not whether AI will shape pulmonary pricing, but whether the field will shape it toward both solvent providers and solvent patients.

Bibliography

1. HQ Pham., *et al*. "Economic burden of chronic obstructive pulmonary disease: A systematic review". *International Journal of Chronic Obstructive Pulmonary Disease* 87.3 (2024): 234-251.
2. I Iheanacho., *et al*. "Economic burden of chronic obstructive pulmonary disease (COPD): A systematic literature review". *International Journal of Chronic Obstructive Pulmonary Dis* 15 (2020): 439-460.
3. C Gutiérrez Villegas., *et al*. "Cost analysis of chronic obstructive pulmonary disease (COPD): A systematic review". *Health Economics Review* 11.1 (2021): article 31.
4. A N Amin., *et al*. "Forecasting the global economic and health burden of COPD from 2025 through 2050". *CHEST* (2025).
5. E Boers., *et al*. "Global burden of chronic obstructive pulmonary disease through 2050". J*AMA Network Open* 6.12 (2023): article e2346598.
6. S Zeng., *et al*. "Developing a machine learning model to predict severe chronic obstructive pulmonary disease exacerbations: Retrospective cohort study". *Journal of Medical Internet Research* 24.1 (2022): article e28953.
7. JL Dasta., *et al*. "Daily cost of an intensive care unit day: The contribution of mechanical ventilation". *Critical Care Medicine* 33.6 (2005): 1266-1271.
8. N Deboever., *et al*. "Machine learning prediction of financial toxicity in patients with resected lung cancer". *Journal of the American College of Surgeons* 241.2 (2025): 107-116.
9. J X Jiang., *et al*. "Price transparency in hospitals — Current research and future directions". *JAMA Network Open* 6.1 (2023): article e2249588.
10. U.S. Government Accountability Office. "Health care transparency: CMS needs more information on hospital pricing data completeness and accuracy". GAO-25-106995, Washington, DC (2024).
11. W Haque., *et al*. "Rapid review of 'No Surprise' medical billing in the United States: Stakeholder perceptions and challenges". *International Journal of Environmental Research and Public Health* 20.5 (2023): article 4459.
12. Office of the Assistant Secretary for Planning and Evaluation (ASPE). "No Surprises Act: Impact on surprise billing". No Surprises Act Reports to Congress Series, Washington, DC (2023).
13. "AI is a promising tool for eliminating hospitals' revenue leakage". *HFMA* (2026).
14. KP Venkatesh., *et al*. "Automating the overburdened clinical coding system: Challenges and next steps". *npj Digital Medicine* 6 (2023): article 16.
15. A Rajkomar., *et al*. "Machine learning in medicine". *The New England Journal of Medicine* 380.14 (2019): 1347-1358.

16. R Kaur and JA Ginige. "A systematic literature review of automated ICD coding and classification systems using discharge summaries". *arXiv preprint arXiv:2107.10652* (2021).

17. F Almaghrabi., *et al*. "AI-based ICD coding and classification approaches using discharge summaries: A systematic literature review". *Expert Systems with Applications* 213 (2023): article 119023.

18. G Luo., *et al*. "Automatically explaining machine learning predictions on severe chronic obstructive pulmonary disease exacerbations: Retrospective cohort study". *JMIR Medical Informatics* 10.2 (2022): article e33043.

19. NN Khanna., *et al*. "Economics of artificial intelligence in healthcare: Diagnosis vs. treatment". *Healthcare* 10.12 (2022): article 2493.

20. SK Pendyala. "Enhancing healthcare pricing transparency: A machine learning and AI-driven approach to pricing strategies and analytics". *International Journal of Scientific Research in Computer Science, Engineering and Information Technology* 10.6 (2024): 2334-2344.

21. L Luo., *et al*. "Using machine learning approaches to predict high-cost chronic obstructive pulmonary disease patients in China". *Health Informatics Journal* 26.3 (2020): 1577-1598.

22. B Langenberger., *et al*. "The application of machine learning to predict high-cost patients: A performance-comparison of different models using healthcare claims data". *PLOS ONE* 18.1 (2023): article e0279540.

23. K Yun., *et al*. "Development and evaluation of machine learning-based high-cost prediction model using health check-up data by the National Health Insurance Service of Korea". *Journal of Personalized Medicine* 12.10 (2022): article 1606.

24. LA Smith., *et al*. "Machine learning and deep learning predictive models for long-term prognosis in patients with chronic obstructive pulmonary disease: A systematic review and meta-analysis". *Lancet Digital Health* 6.2 (2024): e140-e157.

25. M Arnold., *et al*. "Development, evaluation and comparison of machine learning algorithms for predicting in-hospital patient charges for congestive heart failure exacerbations, chronic obstructive pulmonary disease exacerbations and diabetic ketoacidosis". *BioData Mining* 17 (2024): article 35.

26. K Kaier., *et al*. "Mechanical ventilation and the daily cost of ICU care". *BMC Health Services Research* 20 (2020): article 267.

27. M Alkhatib., *et al*. "Artificial intelligence-based automated International Classification of Diseases coding: A systematic review". *Journal of Health Informatics in Developing Countries* (2025).

28. "Current use and evaluation of artificial intelligence and predictive models in US hospitals". *Health Affairs* (2024).

29. ML Burns., *et al*. "Generative AI costs in large healthcare systems, an example in revenue cycle". *npj Digital Medicine* 8 (2025): article 579.

30. "Adoption of AI for hospital RCM surges, but cost and operational constraints slow progress". *Fierce Healthcare* (2025).

31. M Alshammari., *et al*. "The evolution of automated medical billing with artificial intelligence: A review with a global and Saudi perspective". *Cureus* (2025).

32. J Wolff., *et al*. "The economic impact of artificial intelligence in health care: Systematic review". *Journal of Medical Internet Research* 22.2 (2020): article e16866.

33. Z Obermeyer., *et al*. "Dissecting racial bias in an algorithm used to manage the health of populations". *Science* 366.6464 (2019): 447-453.

34. "Assessing socioeconomic bias in machine learning algorithms in health care: A case study of the HOUSES index". *Journal of the American Medical Informatics Association* (2022).

35. J K Paulus and DM Kent. "Algorithm bias and racial and ethnic disparities in health and health care". *JAMA Network Open* 6.12 (2023): article e2346378.

36. MA Gianfrancesco., *et al*. "The potential for bias in machine learning and opportunities for health insurers to address it". *Health Affairs* 41.2 (2022): 212-218.

37. S Mackin., *et al*. "Identifying and mitigating algorithmic bias in the safety net". *npj Digital Medicine* 8 (2025): article 335.

38. "Patient healthcare spending after the No Surprises Act: Quasi-experimental difference-in-differences study". *BMJ* (2025).

39. N Rieke., *et al*. "The future of digital health with federated learning". *npj Digital Medicine* 3 (2020): article 119.

40. "Bias in medical AI: Implications for clinical decision-making". *International Journal of Medical Informatics* (2024).

41. N Norori., *et al*. "Addressing bias in big data and AI for health care: A call for open science". *Patterns* 2.10 (2021): article 100347.

42. E J Topol. "High-performance medicine: The convergence of human and artificial intelligence". *Nature Medicine* 25.1 (2019): 44-56.

Volume 15 Issue 9 September 2026