International Journal of
INTELLIGENT SYSTEMS AND APPLICATIONS IN ENGINEERING
ISSN:2147-6799 www.ijisae.org Original Research Paper

AI-Driven Dynamic Pricing, Fee Optimization, and Incentive Intelligence Across the Transaction Lifecycle

Satheesh Kumar Kumara Chinnaian

Submitted:23/02/2026 **Revised**: 02/04/2026 **Accepted:** 10/04/2026

Abstract: Payment processors have historically relied on static billing models and broad merchant segmentation, creating structural inefficiencies in an increasingly dynamic digital commerce environment. Transaction-level costs, risks, and strategic value vary materially with context—channel, geography, funding source, payout timing, merchant behavior, and dispute outcomes—yet legacy pricing systems treat these dimensions as uniform. This article presents a modern pricing architecture that transforms the pricing engine into a real-time economic decision layer, combining transaction-level cost and loss forecasting, competitive and elasticity-aware optimization, continuous post-settlement learning, and an integrated incentive layer for promotions and merchant-funded campaigns. The platform employs machine learning for predictive components and large language models for unstructured signal extraction, enabling a pricing system that remains auditable, adaptive, and aligned with long-term network health. Implementation through governed architectural layers, deterministic fee construction with explainable components, event-driven lifecycle data contracts, and closed-loop learning mechanisms demonstrate how economic precision and transparency can coexist. Evaluation methods combining controlled experimentation, causality validation, and lifecycle measurement ensure that pricing decisions improve both processor profitability and merchant experience without sacrificing either regulatory compliance or competitive positioning.

Keywords: *Dynamic Pricing Optimization, Real-Time Transaction Economics, AI-Driven Fee Intelligence, Merchant Incentive Alignment, Causal Inference in Payment Networks*

1. Introduction

In payment processing, the pricing engine plays a critical role in determining network profitability and merchant experience. Unlike many software pricing domains, payment pricing is constrained by external network fees, regulatory and scheme rules, and risk-driven losses that can be realized days or weeks after the transaction occurs. A fee that appears profitable at authorization can become unprofitable after clearing, settlement, and dispute events are accounted for. This fundamental tension between immediate authorization economics and downstream financial outcomes requires that pricing be treated as a life-cycle-aware financial process rather than a single-point computation [1].

The introduction of digital payment technologies has changed the payment environment significantly. The modern transactions are fast-moving towards the electronic mode of payment as opposed to cash-based transactions due to technological innovation and evolving consumer behavior. Global non-cash transactions reached over 708 billion transactions in 2019, representing more than 14% growth and the highest volume in decades, demonstrating the scale at which digital payments now operate [1]. Digital payment technologies have been swiftly adapted as a means of how consumers and merchants conduct commerce across channels ever since the COVID-19 pandemic. Owing to this, payment processors in the contemporary era must optimize pricing across an increasingly complex ecosystem of transaction types, funding sources, and merchant segments.

The modern payment environment amplifies pricing complexity significantly. Some of the drivers are: 1) Increasing omnichannel and cross-border transactions; 2) faster payout and settlement expectations; and 3) a diverse range of merchant business models, ranging from subscription platforms and marketplaces to microtransaction digital services. In emerging markets like Kenya, the adoption of digital payment systems has directly strengthened the competitive positioning of financial institutions, with digital payments now recognized as both a functional service and a

---

Independent Researcher, USA

strategic asset that drives customer retention and transaction volume growth [2]. Simultaneously, competition has intensified globally, making pricing sensitivity and merchant churn practical constraints that cannot be ignored.

A modern pricing platform must therefore optimize along multiple dimensions: cost recovery, risk coverage, margin goals, adoption strategy, and competitive positioning while remaining explainable and controllable for finance, risk, and compliance stakeholders. The interplay between financial technology and internet-enabled transaction processing has triggered the emergence of multiple digital payment technologies simultaneously, each with distinct cost structures, risk profiles, and regulatory implications [1]. Thus, a novel pricing approach—recognizing context-dependent variations in cost, risk, and strategic value, rather than applying static rate cards uniformly across various merchant segments—is the need of the hour.

The combination of enhanced digital payment infrastructure and increased competition pressure is both a challenge and an opportunity for payment processors. Using AI-based, context-sensitive pricing methods, organizations can maintain merchant competitiveness while reclaiming the economic value of transactions. Certain impending dangers that businesses based on an inertial pricing framework face are 1) margin leakage, 2) the misalignment of risk coverage, and 3) market share loss to rivals with extra fine-grained and clever fee frameworks. The combination of these factors prompts the realization of the modern pricing architecture in which the pricing engine is constantly informed of the settlement results and responsive to network dynamics and yet deterministically auditable to facilitate regulatory and compliance requirements: a real-time economic decision layer.

Metric	Value
Global Non-Cash Transactions (2019)	708 billion transactions
Annual Growth Rate	14%

Table 1: Digital Payment Scale and Growth Metrics [1, 2]

2. Problem Statement and Key Drivers of Pricing Complexity

A fundamental assumption in legacy pricing systems is that transaction cost and risk are stable within a merchant segment and can be managed through static rates. In practice, pricing requires continuous reconciliation between predicted and realized economics, and various aspects make static pricing structurally insufficient [3].

First, the processor's cost base is multi-layered. External costs include interchange, assessments, cross-border and currency conversion fees, and scheme-driven pricing categories that can change based on transaction attributes. Internal costs include infrastructure, authentication, risk operations, customer support, compliance checks, and dispute servicing. Some costs are visible immediately; others are only known after settlement and downstream events. Investigations on payment card pricing reveal that consumers experience regressive cost structures where pricing mechanisms create dissimilar impacts across income categories and consumer segments. The financial outcomes of payment card pricing and merchant cost pass-through mechanisms vary considerably based on consumer characteristics and transaction patterns; thus, static pricing methods are not effective for heterogeneous impacts across diverse populations [3].

Second, risk and loss are highly context-dependent. Fraud probability and chargeback exposure vary by device confidence, customer history, merchant behavior, goods category, and geographic corridor. Loss is not only a probability estimate; it depends on the exposure window, representation success, and recovery likelihood. Payment processors must account for how merchants' behavior influences cost pass-through and consumer outcomes. Most merchants do not differentiate prices based on payment method, instead passing through their costs of accepting payments uniformly across all transactions. This uniform pricing approach creates cross-subsidies and regressive effects where certain consumer groups bear disproportionate costs relative to their transaction patterns [3].

Third, market pressure creates a "pricing ceiling." A processor cannot simply add margin to cover volatility without considering merchant alternatives and routing behavior. Competitive pricing, merchant elasticity, and platform switching costs must be included as part of the pricing strategy. Modern e-commerce environments demonstrate the critical importance of dynamic pricing approaches that respond to inventory levels, competitive activity, and demand fluctuations. Data-intensive forecasting systems employing artificial intelligence models enable processors and merchants to optimize pricing strategies by analyzing complex transaction datasets and predicting demand patterns across cross-border commerce channels. These intelligent systems can identify optimal pricing by processing heterogeneous data streams and detecting patterns that static approaches miss [4].

Finally, the processor also operates as a financial network. Not all merchants contribute equal strategic value. A large marketplace platform may drive high ecosystem growth, while a smaller, high-risk merchant may create an outsized operational burden. Pricing decisions must incorporate merchant value and network strategy, not only immediate transaction economics. Because pricing depends on large volumes of heterogeneous financial data, the system must support constant learning through a feedback loop that incorporates settlement outcomes, disputes, and behavioral responses to pricing changes. The complexity of modern payment systems and cross-border commerce operations requires that pricing move beyond static models toward adaptive, learning-based architectures that continuously refine cost recovery, risk assessment, and competitive positioning as new settlement and outcome data become available [3] [4].

3. Proposed Architecture: Pricing as a Real-Time Economic Decision Layer

This paper proposes a pricing platform design where the pricing engine is a real-time decision service operating during authorization and refined through post-settlement learning. The key principle is to use AI models to predict economic components while keeping fee construction deterministic and auditable. The foundation for such a system requires canonical data models that unify disparate merchant, terminal, transaction, and chargeback data streams into a coherent architecture capable of supporting real-time decisions and post-transaction reconciliation [5].

At authorization time, the pricing engine consumes structured signals such as merchant segment, MCC/category, funding source, transaction amount, channel, geography, authentication results, device confidence, and payout requirements. It produces a contextual fee composed of separable components: base processing fee, expected external cost recovery, risk premium, liquidity premium (where applicable), and a strategic adjustment factor for merchant value or growth priorities. The critical challenge in building such systems is that merchant data originates from multiple isolated sources—merchant master records residing in onboarding portals and CRMs, terminals monitored in dedicated asset systems, authorizations received in ISO 8583 format, clearing and settlement in scheme files with proprietary extensions, and fees calculated through ad hoc collection of rules engines [5]. These fragmented data architectures create flaky joins, non-uniform keys, and high levels of reconciliation breaks, particularly when merchants operate across different legal entities, countries, or schemes.

Following clearing and settlement, the platform consumes realized cost and outcome messages: actual interchange and assessments, settlement time, reversals, dispute and chargeback messages, and operational servicing costs. These outcomes update model parameters and pricing sensitivities, forming a closed learning loop. A canonical payment data model using event sourcing with conformed dimensions and effective-dated reference data through disciplined Slowly Changing Dimension (SCD) patterns ensures that all acquirers and schemes maintain point-in-time truth and eliminate reconciliation breaks [5].

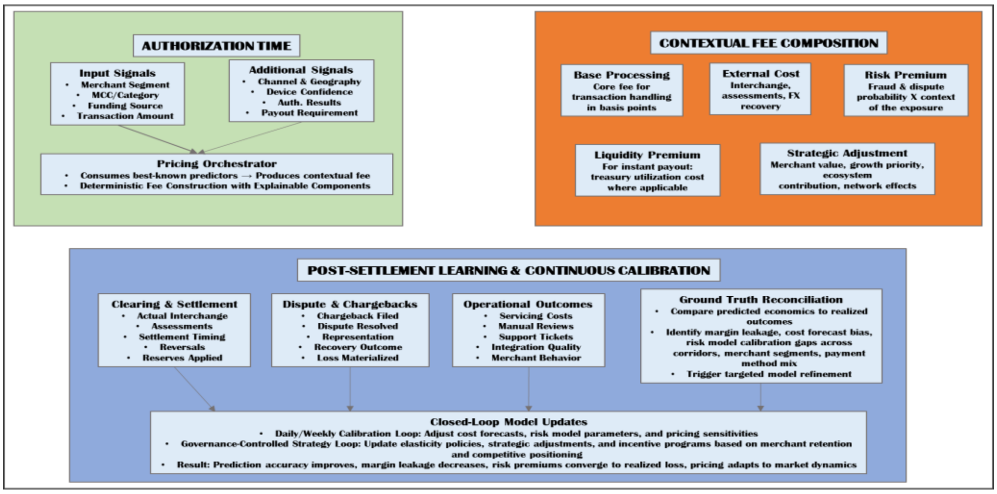

Figure 1: Authorization-Time Pricing with Post-Settlement Learning [5, 6]

Figure 1 illustrates the three-phase temporal architecture of this platform. The ‘Authorization time’ phase shows authorization-time signal consumption and pricing orchestration—where merchant context, transaction attributes, device signals, and authentication outcomes flow into the pricing engine to produce a contextual fee. The ‘Contextual fee composition’ phase decomposes this fee into five auditable components—base processing, external cost, risk premium, liquidity premium, and strategic adjustment—each driven by a specialized AI model yet composed through deterministic rules that auditors and merchants can reconstruct and validate. The ‘Post settlement learning and continuous calibration’ phase represents the post-settlement learning loop, where the outcomes flow in a loop to update and recalibrate the model. This separation of prediction from construction, combined with continuous post-settlement learning, achieves both precision and explainability—the pricing system becomes increasingly accurate as settlement data accumulates while remaining accountable to compliance and risk stakeholders at every decision point.

The adoption of dynamic pricing strategies in payment ecosystems requires understanding how design options and policy decisions influence real-world outcomes. Agent-based modeling simulations of retail payment ecosystem adoption demonstrate this principle empirically. When introducing a new payment instrument without attractive design features or stimulus policies, adoption remains low and slow—approximately 1.29% of transactions after one year in a baseline scenario [6]. When implementing reverse waterfall functionality combined with positive remuneration spreads, adoption reaches 8.88% of transactions—this shows how architecture choices directly influence economic outcomes [6].

A main advantage of this architecture is that full end-to-end fee prediction from a black-box model is not required; the platform predicts expected cost and loss for the drivers and outputs the fee in a human-understandable, auditable, and policy-compliant rule-based format. Deterministic decisions provide merchant accountability and regulatory compliance—they enable post-authorization learning and continuous improvement by integrating authorization-time pricing and post-transaction learning. These combine the advantages of leveraging the best-known predictors at authorization time with the ability to correct shortcomings through ground-truth labels and adapt the underlying models with the latest data. The canonical data model that combines merchant master records, terminal lifecycle management, transactional events, and fee rating into a single explainable schema provides the payment processor with an infrastructure that supports both real-time pricing decisions and post-transaction re-rating of pricing performance [5].

Architectural Element	Performance (%)
Agent-Based transaction adoption after 1 year (Baseline)	1.29%
Agent-Based Adoption (With Design Features)	8.88%

Table 2: Real-Time Decision Layer with Adoption Performance [5, 6]

4. AI/ML Models for Pricing Intelligence

A modern pricing platform benefits from multiple cooperating models, each specialized for a component of the transaction's economic profile [7].

4.1 Expected Loss Modeling

The platform applies supervised learning to estimate fraud and dispute probability, using signals such as customer novelty, device trust, historical merchant dispute ratios, geography corridor risk, and authentication outcomes. The predicted probability is multiplied by exposure to compute expected loss. Importantly, exposure is not always equal to the transaction amount; it may depend on refund windows. delivery timing, and representative success rates. Machine learning approaches for fraud detection in high-volume financial transaction datasets demonstrate the complexity of this task. A study analyzing over one million financial transaction records with advanced supervised learning techniques—including CatBoost, XGBoost, and ensemble methods combining anomaly detection—achieved a ROC-AUC score of 0.996, indicating near-perfect discrimination between fraudulent and legitimate transactions [7]. This level of predictive accuracy enables pricing engines to differentiate risk precisely across transaction contexts rather than applying uniform merchant rates.

A realistic illustration can be seen when two transactions have identical amounts and merchants but different contexts. A returning user on a familiar device may have a very low predicted dispute probability, while that of a first-time cross-border buyer with weak device confidence could be considerably higher. A static merchant rate prices both equally. The proposed system prices the second transaction with a higher risk premium while keeping low-risk transactions competitively priced.

4.2 Real-Time Cost Forecasting

External cost forecasting uses transaction attributes to estimate the likely interchange and assessment category and expected cross-border/FX cost where applicable. Because final costs can depend on clearing details, the model forecasts using best-known authorization-time signals, then learns from realized post-clearing costs to correct systematic bias. This reduces margin leakage caused by underestimating network fees in particular corridors or categories. The structure of transaction costs directly influences pricing strategy and merchant behavior. Research on liquidity premiums and transaction cost dynamics reveals that the shape of transaction costs—whether linear, quadratic, or linear-quadratic—creates substantially different economic outcomes for market participants. Specifically, liquidity providers who facilitate transactions by absorbing transitory price impacts charge risk premiums that can reach several hundred percent of the risk premium observed in frictionless markets [8], demonstrating how transaction cost structures fundamentally reshape the pricing landscape that processors must navigate.

4.3 Liquidity and Payout Cost Modeling

When instant payout or accelerated settlement is offered, the processor incurs liquidity costs due to treasury utilization and market conditions. A time series for the marginal cost of funding allows the pricing engine to only apply a liquidity premium with context-aware and need-based pricing that reflects actual cost, as opposed to conventional "instant payout fee" structures. Understanding how equilibrium responses differ under these cost structures is useful for dynamic pricing. While in this class of models, including many heterogeneous agents with different risk, dislike, and frequency of trading, transaction cost structures can have large implications for market demand, prices tend to be more stable across various cost structure specifications due to the diversity of agents [8].

4.4 Merchant Elasticity and Competitive Optimization

Even with accurate cost and loss estimation, pricing must consider merchant response. An elasticity model estimates the likelihood that a merchant shifts volume away after fee changes by using facets like volume drops that follow any rate adjustment, increased drops in some payment methods, and routing pattern changes. This supports proactive pricing improvements for high-value merchants with strong retention potential and allows the processor to optimize pricing without triggering churn.

4.5 Continuous Learning and Outcome Calibration

The platform employs post-settlement feedback to recalibrate cost forecasts and risk models. When chargebacks like region, device type, or merchant sub-segment arise for a particular pattern, the model updates quickly. The pricing engine adjusts the risk premium specifically for that context rather than penalizing the entire merchant portfolio, thereby reducing false penalties and maintaining fairness.

5. LLM Layer for Unstructured Signals and Merchant Guidance

Payment ecosystems generate large amounts of unstructured text that contain valuable risk and operational signals, while traditional ML models operate on structured transaction features. Dispute narratives, support tickets, merchant onboarding descriptions, and integration logs often explain why disputes occur or why operational cost rises. Recent surveys have identified multiple domains where LLMs (large language models) enhance decision-making and operational efficiency, thereby emphasizing them as emergent, transformative tools for financial applications. Investigations on LLM applications in the financial industry indicate that these models have shown incredible contextual understanding, large-scale data processing, and content creation preferred by humans, making them applicable to all linguistic tasks, sentiment analysis, financial time series, and financial reasoning [9].

Large language models (LLMs) can transform unstructured content into structured features used by the pricing and recommendation system. For instance, dispute narratives can be grouped into emerging clusters that indicate new abuse patterns. Merchant support tickets and logs can reveal integration gaps—missing device signals, improper capture practices, and incomplete shipping evidence—that increase disputes and chargeback losses. The integration of LLMs into financial analysis addresses challenges such as lookahead bias in backtesting, legal concerns surrounding AI-generated content, data pollution, signal decay, inference speed, cost, uncertainty estimation, and interpretability [9].

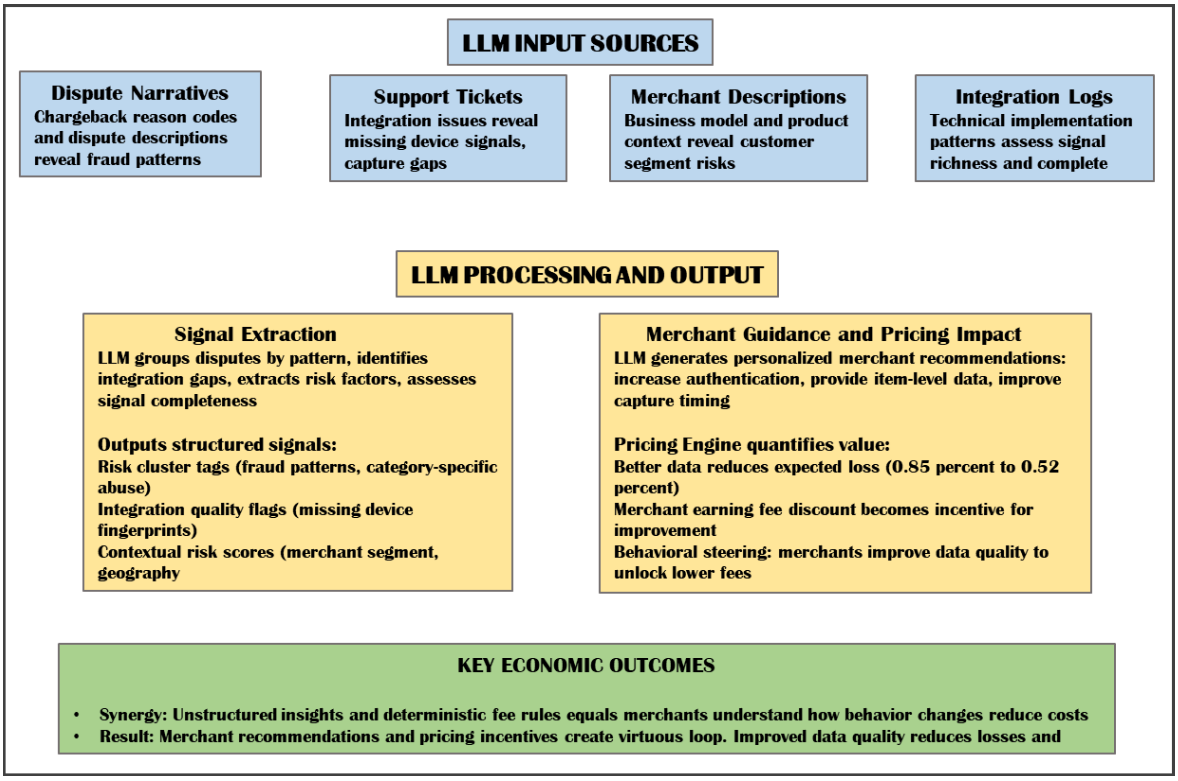

Figure 2: Unstructured Signal Extraction to Merchant Guidance Pipeline [9, 10]

Figure 2 visualizes how unstructured payment data becomes structured risk intelligence through LLM processing. The top row shows four LLM input sources: 1) dispute narratives reveal emerging fraud patterns and category-specific abuse clusters; 2) support tickets expose integration gaps and missing device signals that increase operational burden; 3) merchant descriptions contextualize customer segment risks and geographic exposures; and 4) integration logs assess the completeness and richness of transaction signal transmission. The 'LLM processing and output' section shows how LLMs perform specialized extraction: grouping disputes into actionable risk clusters, identifying integration quality gaps, extracting contextual risk factors, and scoring signal completeness. The output is in 2 parts: structured signals feed back into the ML models (improving fraud and loss predictions).

Personalized merchant recommendations—backed by narrative evidence from disputes and tickets—become economic levers. When merchants implement recommended improvements (better authentication, item-level data, optimized capture timing), the pricing engine quantifies the resulting loss reduction and translates it into fee incentives, creating a virtuous cycle where merchant behavior improvements directly unlock lower pricing.

LLMs can also generate structured information for merchant action, such as: 1) More authentication coverage for first-time buyers above a certain transaction amount, 2) Item-level data on high-dispute categories, 3) Adjusted capture timing in order to become less likely to be reversed. These recommendations connect directly to economic outcomes: better data and practices reduce expected loss, and the pricing engine can reward those improvements with lower fees. Research on merchant recommendation systems demonstrates how personalized approaches substantially improve performance when domain-specific information is incorporated. A neural collaborative filtering model enhanced with merchant and customer information was implemented in a study. It demonstrated a 3% performance improvement based on HR@10 and a 5% improvement based on NDCG@10 compared to baseline models—thus making the value of incorporating contextual data into recommendation systems eminent [10].

Furthermore, merchant guidance systems powered by LLMs can leverage collaborative filtering methodologies to identify patterns across merchant segments. By analyzing payment data from millions of customer interactions, LLMs can detect which merchants provide the greatest value to specific customer segments and recommend improvements that reduce friction in payment processing. The merchant recommender system, when augmented with customer demographic information and merchant characteristics such as industry category and regional location, shows substantially improved accuracy in predicting future merchant usage patterns [10]. This capability enables pricing engines to dynamically adjust fees based not only on transaction-level risk but also on the likelihood of merchant adoption and customer satisfaction.

The synergy between LLM-powered signal extraction with deterministic pricing rules generates a system in which unstructured insights are directly used to shape fee composition. When merchants are provided with actionable advice produced by LLMs, supported by evidence in the form of narratives of disputes, patterns of customer feedback, and integration histories, they are able to make changes that can be shown to lower the costs of operations. These enhancements are then identified by the pricing engine by use of reduced forecasted loss and risk indicators, which convert improvement of merchant behavior into reduced charges. This virtuous loop aligns the incentives of merchants with the profitability of processors and creates transparency in the process of setting fees and in how merchants can gain a better pricing tier.

LLM Component	Output and Impact
Dispute Clustering	Risk pattern tags; fraud scheme detection; geographic abuse surges
Integration Gap Detection	Quality flags; missing device signals; capture and evidence gaps
Risk Factor Extraction	Contextual risk scores by merchant segment, product, geography
Signal Completeness Assessment	Quality scores for device signals, customer IDs, item-level data
Merchant Recommendations	Evidence-backed guidance: authentication, item-level data, capture timing, 3D Secure

Table 3: LLM Signal Extraction and Merchant Guidance [9, 10]

6. Promotions and Incentives: Pricing as Behavioral Steering

Beyond pricing transactions, modern processors require a promotion and incentive layer that supports adoption growth without unmanaged margin erosion. Conditional promotions that align with measurable improvements in cost and risk are the most effective. The manner in which organizations approach customer engagement and incentive design has undergone a sea change with the integration of machine learning and artificial intelligence into financial services. Contemporary research demonstrates that analyzing spending patterns through advanced computational methods enables marketers to craft targeted promotional strategies that respond to actual consumer behavior rather than demographic assumptions alone [11].

A potential approach involves providing fee discounts in exchange for merchants offering additional context in the transaction. This can aid in preventing chargebacks (e.g., device fingerprint signals, customer-identifying details, or item-level data) and thereby reduce the estimated chargebacks and fees towards the resolution of the dispute. This allows the pricing engine to evaluate the economic value of the signals to expected loss and to the cost of servicing them, leading to economically driven rather than solely marketing-driven incentives. Machine learning of spending and transaction behavior can look at the merchants and customer segments that are most influenced by each type of incentive and can therefore better allocate promotion dollars [11].

The incentive structure should also be somewhat adaptive so as to increase the baseline pricing based on a given transaction's context. If a merchant signals better and fewer disputes arise in subsequent settlement cycles, the algorithm must gradually adjust and increase its confidence in sensing the merchant as low-risk. Thus, instead of one-off discounts, promotions are on-ramps into better pricing through improving behavior. The success of promotional mechanics depends on how monetary incentives could influence consumer and merchant behavior across different segments. Modern incentive design thinking takes behavioral factors, beyond the customary economic factors, into account to attract and retain users in payment systems [12].

Financial literacy and informed decision-making are crucial aspects affecting merchants' and consumers' responses to promotional offers. The economic rationale behind incentive structures is how better data quality translates to lower risk premiums: when merchants understand this, they become more engaged partners in improving transaction quality. Similarly, consumers who receive personalized incentive offers based on their spending patterns and preferences demonstrate higher engagement and loyalty. Machine learning systems analyzing transaction data can identify optimal incentive levels for different customer cohorts. This enables incremental adoption of promotional spending drives, as opposed to a de facto subsidizing behavior [12].

The processor can be thought of as a two-sided market approach, with the value perceived by both merchants and consumers triggering their participation. It is feasible to keep both merchants and consumers aligned with value generation by evaluating historical spend patterns and financial behavior in the promotion design and offering incentives accordingly. Merchants receive discounted rates for higher-quality data and lower fraud, and consumers get cashback and discounts when spending at merchants with higher-quality data. As such, the processor sees lower risk, higher transactions per consumer, and sustainable growth, creating returns on promotions rather than permanent margin erosion, as is the case with retailer promotions.

Promotion design also needs to be adjusted based on realized outcomes. If merchants make changes to signal to consumers, or if consumers respond to the incentives offered, the pricing engine compares realized economics to the economics implied by the incentive structure and adjusts the latter. Processors that run promotions as experiments and measure the results are able to continuously increase the promotional yield (value per dollar of promotional spend) as merchants change their behavior [11] [12].

7. Merchant-Funded Campaigns and Consumer Adoption Optimization

Merchants commonly run consumer campaigns using predetermined funds deposited into a campaign budget. The consumer receives cashback or discounted purchase pricing until the exhaustion of campaign funds pertaining to their account—often to increase store adoption, drive repeat purchases, or promote specific in-house products. A modern pricing platform integrates campaign execution with transaction processing so that campaign funding, consumer incentives, and processor fees are computed consistently at authorization and reconciled after settlement. The design and implementation of merchant-funded campaigns represent a critical intersection between consumer incentives, merchant economics, and payment processor profitability. Research on instant payments and merchant economics demonstrates that payment timing and cost structures fundamentally influence how merchants allocate marketing budgets and structure consumer promotions [13].

AI improves campaign efficiency by optimizing targeting and pacing. Instead of spending the campaign budget uniformly, the system estimates conversion lift by segment, time, corridor, and customer cohort, allocating incentives where incremental adoption is highest. Merchants benefit from campaigns where a data-driven approach is utilized for pricing incentive allocation, rather than ad hoc allocation, to maximize promotional return on investment. The pricing platform captures the responsiveness of customer segments to campaign offers, the time frame of customers' maximum purchasing intensity, and the geographies with the most incremental volume. Merchants can use this intelligence to focus promotional expenditures on moments and groups of customers with high opportunities, instead of allocating budgets uniformly in low-value segments [13].

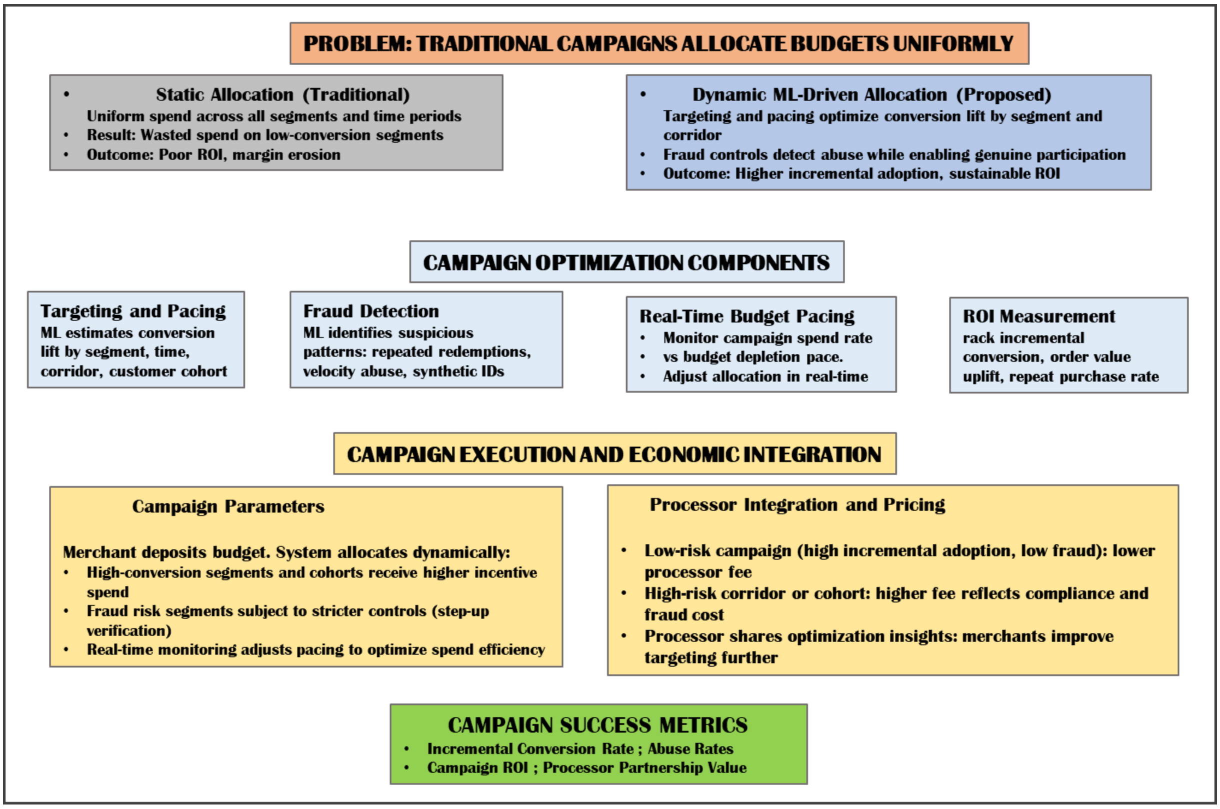

Figure 3: Dynamic Campaign Execution: From Budget Allocation to Incremental ROI [13, 14]

Figure 3 contrasts static (traditional) versus dynamic (proposed) campaign execution and illustrates how AI optimization addresses three simultaneous challenges. The top part shows the problem: traditional uniform budget allocation wastes spend on low-conversion segments; and the proposed solution: ML-driven targeting concentrates incentives where conversion lift is highest by segment, time corridor, and customer cohort. The 'Campaign optimization components' show the execution mechanics: 1) targeting and pacing concentrate spend on incremental adoption; 2) Fraud detection identifies suspicious patterns (repeated redemptions, velocity abuse, and synthetic identities) and applies contextual controls (throttling, step-up verification, and tighter eligibility); 3) real-time budget pacing monitors spend depletion and adjusts allocation dynamically; and 4) ROI measurement distinguishes incremental conversion from baseline volume. The 'Campaign execution and economic integration' section shows economic integration: merchants deposit campaign budgets and the system allocates dynamically, while processor fees vary by campaign risk profile—low-risk, high-adoption campaigns earn lower fees, positioning the processor as an active partner in merchant success rather than a passive fee taker. The result is sustainable campaign ROI measured by incremental conversion rate, abuse rate mitigation, and repeatable merchant profitability across cohorts. Campaign systems must also manage abuse risk. Promotions attract fraud attempts, such as repeated redemptions, synthetic identities, or velocity abuse. Risk models can detect abnormal redemption patterns and apply controls like throttling, step-up verification, or tighter eligibility rules, protecting the merchant's campaign funds while maintaining genuine consumer incentives. The balance between enabling legitimate consumer participation and preventing fraud exploitation is delicate and requires sophisticated detection systems. Fraud detection methodologies that operate within promotional payment systems employ machine learning to identify suspicious patterns while minimizing false positives that would unnecessarily restrict genuine customers from participating in campaigns [14].

Finally, the platform must measure the true return on investment (ROI), not just spend. It needs to track aspects like incremental conversion, average order value uplift, and repeat purchase rates and needs to recommend campaign strategies and budget allocations. Merchants running a campaign need to clearly visualize the effect of a campaign's spend on consumer behavior. Thus, an ideal analytics platform would determine the volume increment caused by the promotion, allowing the impact of a campaign to be accurately monitored without interference from the volume baseline (the sales figure that would happen anyway without the promotion), giving merchants the opportunity to assess a campaign's success and plan future expenditure accordingly [13].

The processor's role in merchant-funded campaigns extends beyond simple transaction processing to encompass campaign design consultation, fraud prevention, outcome measurement, and optimization recommendations. By providing merchants with insights into which promotional strategies drive the highest incremental adoption, processors position themselves as valuable partners in merchant success rather than merely charging fees. Campaign data aggregated across thousands of merchants reveals patterns that individual merchants cannot observe in isolation—seasonal variations in effectiveness, cohort-specific response rates, optimal incentive levels for different customer segments, and emerging fraud patterns specific to promotional environments [14]. Integration of campaign management with the pricing engine enables dynamic fee adjustments based on campaign characteristics. Campaigns driving high incremental volume with low fraud risk might warrant lower processor fees, aligning processor incentives with merchant campaign success. Conversely, campaigns in high-risk corridors or with fraud indicators might carry higher fees, reflecting increased processing and compliance costs. This approach transforms the processor from a passive fee taker into an active partner in campaign optimization, where both parties benefit from campaigns that drive genuine incremental adoption while controlling fraud and operational risk [13] [14].

Campaign Component	Optimization Approach
Targeting and Pacing	Conversion lift estimated by segment and corridor
Fraud Detection in Campaigns	Machine learning identifies suspicious patterns
ROI Measurement	Tracks incremental conversion and order value uplift
Processor Partnership Role	Provides campaign design and optimization consultation

Table 4: AI-Driven Campaign Execution and Fraud Prevention [13, 14]

8. Evaluation and Operational Metrics

A modern pricing intelligence platform should be measured using both financial and network health metrics. Financially, the key outcomes include reduced margin leakage (the difference between expected and realized margins), improved loss ratio, and higher net revenue per transaction adjusted for churn. On the network side, outcomes include reduced dispute rate, improved merchant retention, and increased adoption of risk-reducing signals and preferred integration patterns. The measurement of pricing system effectiveness requires comprehensive frameworks that capture multiple dimensions of performance across the payment ecosystem. Research on intelligent pricing strategies in e-commerce demonstrates that effective pricing systems must adapt continuously to market conditions while maintaining measurable alignment between predicted and realized economics [15]. Margin leakage represents one of the most critical financial metrics for pricing system evaluation. When a pricing engine predicts that a transaction will generate a certain expected margin based on authorization-time signals, but actual post-settlement economics differ materially, that variance indicates systematic prediction error. Reducing margin leakage—the gap between expected and realized margin—demonstrates that the pricing system has improved its understanding of true transaction costs and risks. A processor with tight margin leakage variance can confidently commit to pricing commitments and margin targets, whereas high leakage indicates that either the pricing model is systematically biased or that post-settlement events (disputes, reversals, or operational costs) are unpredictable [15]. Loss ratio improvement measures how effectively the pricing system has incorporated risk assessment into fee structures. When transaction-level fraud or chargeback risk is accurately priced, the realized loss ratio should approximate the expected loss ratio. Deviations indicate either that risk models are inaccurate or that pricing does not sufficiently reflect identified risks. Network-side metrics provide equally important signals about system health. Dispute rate reduction indicates that either transaction quality has improved or merchants are implementing better practices to prevent disputes. Merchant retention improvement suggests that pricing changes have not driven unacceptable churn while still recovering appropriate costs and margins. Integration quality adoption reflects whether merchants perceive sufficient value in implementing single improvements to justify implementation effort [15]. A realistic operational target is not simply "maximize fee." It is to improve expected margin accuracy while maintaining competitive positioning. Even small improvements—such as reducing underpriced high-loss flows and rewarding low-risk behavior—compound significantly at scale. The challenge in evaluating pricing systems lies in distinguishing between correlation and causation. When merchant retention improves after implementing a new pricing system, it may reflect the new pricing structure itself, or it may reflect broader market improvements, seasonal variations, or competitive dynamics. Rigorous evaluation frameworks must account for these confounding factors through appropriate control group comparisons or time-series analysis [16]. Strategic data integration across payment processing infrastructure enables detection of margin leakage sources and supports continuous improvement in pricing accuracy. By linking authorization-time pricing decisions to post-settlement outcomes through consistent transaction identifiers and time stamps, processors can trace each pricing decision back to its realized economic outcome. This traceability reveals which merchant segments, transaction types, or geographic corridors exhibit the highest margin leakage, directing attention to areas where pricing models require refinement [16]. Furthermore, evaluation frameworks should incorporate leading indicators that predict future financial performance. Merchant adoption of signal improvements, for example, typically precedes dispute reduction by one to three settlement cycles. By monitoring signal adoption rates, processors can predict future loss ratio improvements before they materialize. Similarly, changes in merchant routing patterns or payment method preferences may signal satisfaction or dissatisfaction with pricing, allowing processors to adjust strategies before churn accelerates [15] [16].

9. Implementation Approach: Building an Auditable, AI-Driven Pricing Platform

9.1 Target Architecture and Data Flow

Modernizing pricing in a payment processor requires separating real-time authorization decisioning from post-transaction learning. At authorization time, the pricing engine consumes best-known predictors—merchant segment, MCC, geography, device confidence, authentication outcomes, and payout requirements—and produces a contextual fee composed of transparent, separable components: base processing fee; external cost recovery; risk premium; liquidity premium, where applicable; and strategic adjustment factors. Later events—clearing, settlement, and chargebacks—provide ground truth that updates model parameters and identifies systematic bias. This temporal architecture prevents the common failure mode where authorization-time economics appear sound but downstream events reveal persistent underpricing or overpricing in specific corridors or risk segments [17]. The cleanest implementation structures around a few non-negotiable principles: deterministic fee construction, model-predicted components, event-driven lifecycle data contracts, and closed-loop learning from clearing, settlement, and disputes.

9.2 Core Services and Responsibilities

Production-grade implementation typically breaks into specialized services with clear ownership boundaries. A Pricing Orchestrator manages the real-time authorization API with stringent latency and availability requirements, calling enrichment and scoring services while applying deterministic fee logic that can be reconstructed from audit logs. A Signal Enrichment Service builds transaction context using merchant metadata, historical behavior, risk indicators, and payout preferences, relying on caches and precomputed aggregates to minimize runtime joins. A cost intelligence service predicts external costs—interchange category, assessments, and FX fees—and later reconciles realized costs from clearing files. A risk & loss prediction service scores fraud & dispute probability, returning not only a numerical score but also top reason drivers for explainability. A liquidity pricing service computes liquidity premiums for instant payout using treasury utilization signals. Policy and Compliance Guardrails enforce contract commitments, jurisdiction constraints, and non-discrimination rules. A Promotion and Campaign Service manages conditional discounts and merchant-funded campaign wallets with fraud controls. Finally, a recommendation service generates merchant-specific actions that reduce disputes and unlock fee incentives [17].

9.3 Deterministic Fee Construction with ML Components

To preserve auditability aligned with SOX and compliance expectations, the platform avoids black-box fee prediction in favor of ML-estimated drivers composed through transparent rules. This structure produces an explainability receipt per transaction documenting the decision path: merchant segment, device confidence, predicted fraud probability, predicted loss exposure, predicted interchange category, external cost recovery percentage, base fee in basis points, risk premium, liquidity premium, strategic adjustment, and final fee. This granularity enables auditors to validate logic, identify systematic issues, and defend pricing decisions to regulators and disputing merchants. Explainability remains the core differentiator between responsible AI implementation and speculative algorithmic scoring [17].

9.4 Event-Driven Data Contracts Across the Lifecycle

Lifecycle truth requires explicit event contracts spanning authorization through dispute resolution. Core immutable events include PaymentAuthorized, PaymentCaptured, ClearingReceived (with interchange category and assessments); SettlementCompleted (net settlement, payout timing, and reserves); DisputeOpened, ChargebackFiled, DisputeResolved, and ReversalPosted. Each event carries stable identifiers—transaction ID, merchant ID, instrument hash, and network reference—enabling deterministic joins back to original pricing decisions. Without these contracts, institutions lose the ability to correlate predicted economics with realized outcomes, perpetuating blind spots in margin analysis. A systematic review of AI-enabled financial decision systems demonstrated that governance maturity mediates the relationship between AI integration and financial outcomes, with strong structural pathways confirmed ($\beta = 0.76$ from AI integration to governance, $\beta = 0.73$ from governance to outcomes) across 1155 peer-reviewed studies [17].

9.5 Continuous Learning Loop

Closed-loop learning operates in two stages with different governance cadences. A calibration loop—running daily or weekly—adjusts model calibration and cost forecasts as interchange, assessment, and dispute outcomes accrue, enabling rapid tactical response to emerging patterns. A strategy loop—governed and rolled out gradually—updates elasticity policies, strategic adjustments, and incentive programs based on merchant cohort retention, volume sensitivity, and competitive positioning. This separation prevents hasty policy changes while enabling rapid tactical calibration. Model convergence and validation are confirmed through robust fit indices, with comparative fit indices (CFI = 0.947) exceeding thresholds that demonstrate structural soundness and negligible residual variance [17].

9.6 Promotions, Campaign Wallets, and Fraud Controls

Promotions and campaign wallets are implemented as policy-driven fee modifiers rather than ad-hoc overrides. Conditional processor-funded incentives reward merchants who provide enriched data—device signals, customer identifiers, and item-level detail—that measurably reduces expected loss. The pricing engine quantifies this value by measuring impact on expected loss and operational cost, grounding incentives in economics. Merchant-funded consumer campaigns allocate incentives dynamically using uplift models that estimate conversion lift by segment and corridor, concentrating spend where incremental adoption is highest. Real-time pacing and abuse detection protect campaign budgets while maintaining genuine consumer incentives [18].

9.7 Supply Chain Pricing Integration

Within a closed-loop supply chain optimization framework, analogous principles apply to inventory and pricing decisions. An integrated model optimizing selling prices and cycle times demonstrates how synchronized pricing and operational strategies enhance profitability. Research implementing such models found that optimal selling prices for manufactured products at \$837.5750 and remanufactured products at \$421.6721, with cycle times of 1.5314 months, yielded total system profits of \$1,378,007.354 [18]. This case illustrates how pricing decisions interact with production-remanufacturing rates, inventory holding costs, and collection center efficiency, demonstrating that profit optimization requires simultaneous consideration of pricing, cycle planning, and operational configuration rather than isolated fee adjustment.

9.8 Governance, Monitoring, and Control Architecture

Governance infrastructure anchors the entire system. Model versioning with change logs; explainability logs stored per transaction for audit and support; kill switches and safe-mode fallbacks if models degrade; bias and fairness checks ensuring pricing alignment with policy and regulations; drift monitoring detecting feature, prediction, or outcome drift; A/B experimentation frameworks enabling incremental merchant-cohort rollouts and separation of duties ensuring pricing policy edits require approval before deployment—these controls transform AI from an experimental tool into production infrastructure. A mature implementation treats the pricing engine not as a black-box optimization artifact but as a governed financial decision layer accountable to merchants, regulators, and stakeholders, embedding transparency and accountability at every architectural layer [17].

Metric	Value
AI Integration to Data Governance Coefficient (β)	0.76
Data Governance to Financial Outcomes Coefficient (β)	0.73
Peer-reviewed Studies analyzed	1155
Comparative Fit Index (Model Validation)	CFI = 0.947
Optimal Selling Price (Manufactured Products)	$837.575
Optimal Selling Price (Remanufactured Products)	$421.6721

Table 5: Causal AI and Governance Integration Metrics [17, 18]

10. Evaluation Methodology: Proving Economic Lift and Network Health Improvements

10.1 Experimental Design Principles

A pricing modernization effort is only credible if it isolates causality, protects merchants from unexpected impacts, and satisfies finance and risk governance expectations. The evaluation framework must combine controlled experimentation, lifecycle-aware measurement spanning authorization through disputes, and robust financial reconciliation that distinguishes true cost from true loss. Pricing changes influence merchant routing behavior and volume decisions, making global rollout testing problematic. Instead, controlled rollouts with clear segmentation and strong guardrails enable causal validity assessment. Causal artificial intelligence models have been increasingly applied to supply chain and financial decision contexts to identify genuine treatment effects and optimize policy outcomes, as demonstrated in aerospace supply chain financing where researchers evaluated nine suppliers with two operational treatments applied through causal frameworks to determine effective mitigation strategies [19].

10.2 Progressive Rollout Phases

Implementation progresses through three phases designed to isolate impacts while minimizing merchant exposure to adverse effects. Phase 1 operates in shadow mode without affecting merchant economics: the AI pricing engine computes a shadow fee while merchants continue paying current rates, validating prediction accuracy and margin leakage without production impact. Phase 2 introduces limited cohort A/B testing with conservative caps and floors applied to small merchant or corridor sets. Randomized assignment is preferred; when contract constraints prevent randomization, matched cohorts using propensity score matching reduce selection bias. Phase 3 scales multi-dimensional optimization across incentive policies and campaign structures, testing these separately before joint implementation since incentives alter transaction mix and confound pricing measurements [19].

10.3 Cohort Construction and Segmentation

Testing respects merchant heterogeneity through structured cohort grouping. Stable cohorts are organized by merchant segment (SMB, mid-market, enterprise), merchant category (digital goods, marketplace, travel, subscriptions), geography corridor (domestic, cross-border, FX), payment method mix (card, wallet, bank transfer, BNPL), risk tier (historical dispute/chargeback rate bands), and payout type (standard versus instant). Treatment and control assignments operate at the merchant level when possible to avoid contamination where merchants observe different fees for similar transactions. For large merchants, routing partitions enable testing by geography or payment method, allowing safe assessment without contract-level disruption [19].

10.4 Financial and Risk Metrics

A comprehensive evaluation framework incorporates financial, risk, behavioral, and operational metrics measured across the full payment lifecycle. Net Revenue per Transaction (NRPT)—calculated as revenue minus external costs, internal costs, and realized losses after settlement and dispute resolution—serves as the primary financial metric. Margin leakage quantifies the gap between predicted and realized margin after clearing and settlement, measuring not only margin height but also predictability and tightness. Cost Recovery Accuracy tracks how closely predicted interchange and assessments match realized network fees by corridor and category. Liquidity premium accuracy evaluates whether priced liquidity costs align with realized funding costs under treasury conditions when instant payout is offered. Chargeback Rate and Loss Rate are measured separately to distinguish transaction volume from dollar impact after representment and recovery [19].

10.5 Merchant Behavior and Adoption Metrics

Volume Retention and Churn Signals track processing volume share, routing shifts, or transaction attempt changes, revealing whether pricing changes trigger unintended merchant behavior. Elasticity indicators quantify volume sensitivity to fee changes by cohort, which is essential for tuning strategy. Integration Quality Adoption measures the percentage of transactions with device signals, customer identifiers, and item-level data when incentives are offered for such enrichment, directly indicating whether incentive policies drive desired behavior improvements. Dispute Mix Shift identifies emerging patterns through reason code clustering, detecting new abuse modes early. Support tickets corresponding to transactions reveal whether pricing changes generate merchant confusion; explainability receipts reduce this friction [19].

10.6 Lifecycle Measurement Windows

Correct evaluation establishes multiple measurement points reflecting payment outcomes that mature at different horizons throughout the transaction lifecycle. Authorization-time metrics capture initial decisioning quality, clearing and settlement phases provide ground truth regarding actual costs and payout timing, dispute and chargeback periods surface realized loss exposure and recovery outcomes, and merchant retention and routing behavior reveal whether pricing preserves competitive positioning. Causal inference methods demonstrate that treatment effects vary substantially across supplier segments and time horizons, with some exhibiting valid causal relationships while others show no significant causality despite identical treatment application, highlighting why evaluation must account for longitudinal variation through structural equation modeling frameworks that capture indirect effects through multiple mediating pathways accumulating across time [19] [20]. This structured approach prevents common failure modes where authorization economics appear sound but downstream losses expose systematic underpricing.

10.7 Causal Methods and Quasi-Experimentation

When full randomization is constrained by contracts, credible quasi-experimental methods establish causality. Difference-in-Differences (DiD) compares treatment merchant changes relative to control merchants over the same period, controlling for seasonality. Matched Cohorts with Propensity Score Matching pairs merchants with similar pre-treatment volume, dispute history, corridor mix, and growth rate, then compares outcomes. Interrupted Time Series tracks time-series metrics and detects shifts after rollout while controlling for trend. A case study in aerospace supply chain financing using causal machine learning methods found that nine suppliers with different buyer relationships exhibited heterogeneous treatment effects, where some demonstrated valid causal relationships between operational treatments and payment outcomes while others showed no significant causality, demonstrating the necessity for company-specific analysis rather than population-level assumptions [19].

10.8 Guardrails, Controls, and Publishable Results

Pricing experiments embed strict risk controls protecting merchant experience. Maximum fee change caps per cohort lead to bounding of exposure. Stop-loss triggers activate rollback if loss rates exceed thresholds. Merchant opt-out controls preserve contract flexibility where monitoring dashboards track margin, loss, disputes, and volume shift continuously. Results reporting breaks down outcomes by cohort rather than global averages, because pricing improvements derive from precision.

Recommended reporting includes predicted versus realized margin distribution, margin leakage reduction by corridor and category, loss rate changes by risk band, and merchant retention impact by segment. A comprehensive evaluation framework examining financial innovation's impact on investment decisions across 575 respondents found that structural equation modeling revealed a total indirect effect of financial innovation on investment decisions of 0.362 through mediating pathways including bank performance (standardized coefficient 0.512 for BP→ID), customer satisfaction, and employee satisfaction, demonstrating that multi-layered causality analysis captures complex economic relationships in organizational systems [20]. Publishable outcomes demonstrate "improved economic accuracy and network health while maintaining competitive behavior" rather than only "increased fees," establishing credibility with both merchant and regulatory audiences.

Evaluation Component	Specification
Financial Innovation Study Sample	575 respondents
Indirect Effect of Financial Innovation on Investment Decisions	0.362
Bank Performance to Investment Decisions Coefficient	$\beta = 0.512$

Table 6: Causal Experimentation and Financial Impact Metrics [19, 20]

Conclusion

Modern payment processing requires moving beyond static rate cards toward dynamic, data-driven pricing architectures that reflect the true economic nature of each transaction while maintaining transparency, fairness, and auditability. The convergence of machine learning, large language models, and governed implementation frameworks enables processors to optimize across multiple dimensions simultaneously—cost recovery, risk mitigation, merchant retention, and ecosystem growth—without sacrificing the explainability and accountability that regulators and merchants demand. By separating cost and loss predictions from fee construction, processors achieve both precision and interpretability, allowing auditors and merchants to understand the economic logic behind every price. Merchant-funded campaigns and incentive structures grounded in measurable economic impact transform promotions from marketing expenses into strategic tools aligned with processor and merchant objectives. Continuous post-settlement learning closes the loop between predicted and realized economics, enabling systems to adapt to changing market conditions, emerging fraud patterns, and merchant behavior shifts. Implementation through specialized microservices with clear ownership, deterministic policies, and comprehensive monitoring ensures that complexity serves merchants and processors rather than creating opacity. Evaluation methods grounded in causal inference and lifecycle measurement provide credible evidence that pricing improvements translate into sustainable competitive advantage and network health rather than short-term margin extraction. As payment volumes accelerate and merchant sophistication increases, the ability to make economically sound, explainable, and adaptive pricing decisions will increasingly determine competitive positioning in payment processing markets globally.

References

[1] Khando Khando et al., "The Emerging Technologies of Digital Payments and Associated Challenges: A Systematic Literature Review," MDPI, 2022. [Online]. Available: https://www.mdpi.com/1999-5903/15/1/21

[2] Dianah Injendi Manyengo et al., "Effect of Digital Payment Systems on the Competitive Advantage of Commercial Banks in Nairobi," International Journal of Social Science and Humanities Research, Nov. 2025. [Online]. Available: https://www.researchpublish.com/upload/book/EFFECT%20OF%20DIGITAL%20PAYMENT%20SYSTEMS-14112025-4.pdf

[3] Marie-Hélène Felt et al., "Regressive effects of payment card pricing and merchant cost pass-through in the United States and Canada," ScienceDirect, 2023. [Online]. Available: https://www.sciencedirect.com/science/article/abs/pii/S0378426623001668

[4] Yuk Ming Tang et al., "Data-Intensive Inventory Forecasting with Artificial Intelligence Models for Cross-Border E-Commerce Service Automation," MDPI, 2023. [Online]. Available: https://www.mdpi.com/2076-3417/13/5/3051

[5] Ravi Kumar Vallemoni, "Canonical Payment Data Models for Merchant Acquiring: Merchants, Terminals, Transactions, Fees, and Chargebacks," SCSITR-IJCSE, 2022. [Online]. Available: https://iscsitr.in/index.php/ISCSITR-IJCSE/article/view/ISCSITR-IJCSE_03_01_006/ISCSITR-IJCSE_03_01_006

[6] Carlos León et al., "Simulating the Adoption of a Retail CBDC," De Gruyter Brill, 2024. [Online]. Available: https://www.degruyterbrill.com/document/doi/10.1515/jbnst-2024-0002/html?lang=en&srsltid=AfmBOooOdlkNN0JsIDOex1GASfD-6fd2QMHqoXPv7qb_iNoI_wIK4MDG

[7] Sourav Yallur et al., "Scalable Machine Learning Solutions for High-Volume Financial Transaction Fraud Detection," MDPI, Dec. 2025. [Online]. Available: https://www.mdpi.com/2813-0324/12/1/1

[8] Sergey Isaenko, "Liquidity premium and the shape of transaction costs," ScienceDirect, Nov. 2025. [Online]. Available: https://www.sciencedirect.com/science/article/abs/pii/S0264999325002640

[9] Yuqi Nie et al., "A Survey of Large Language Models for Financial Applications: Progress, Prospects, and Challenges," arXiv, 2024. [Online]. Available: https://arxiv.org/html/2406.11903v1

[10] Suyoun Yoo and Jaekwang Kim, "Merchant Recommender System Using Credit Card Payment Data," MDPI, 2023. [Online]. Available: https://www.mdpi.com/2079-9292/12/4/811

[11] Renuka Agrawal et al., "Analyzing and Rewarding Credit Card Spending Habits in India: A Machine Learning Approach," Springer Nature, Jul. 2025. [Online]. Available: https://link.springer.com/article/10.1007/s44196-025-00899-0

[12] RVS Praveen et al., "Enhancing Financial Literacy and Personal Investment Decisions Through AI and Machine Learning," Journal of Marketing & Social Research, Jun. 2025. [Online]. Available: https://jmsr-online.com/article/enhancing-financial-literacy-and-personal-investment-decisions-through-ai-and-machine-learning-229/

[13] William Cook et al., "Instant Payments and Merchants—Pricing Policy Considerations," CGAP, 2023. [Online]. Available: https://www.cgap.org/sites/default/files/publications/Instant%20Payments%20and%20Merchants-compressed.pdf

[14] Katherine Mayo et al., "Fraud Risk Mitigation in Real-Time Payments: A Strategic Agent-Based Analysis," IJCAI. [Online]. Available: https://www.ijcai.org/proceedings/2024/0018.pdf

[15] Mahesh Gavade et al., "Intelligent pricing strategies in E-Commerce: An AI-Based Framework For Market Trend Adaptation," International Journal of Applied Mathematics Financial Analytics and Performance Measurement, Dec. 2025. [Online]. Available: https://ijamjournal.org/ijam/publication/index.php/ijam/article/view/1499/1363

[16] Adegbola Oluwole Ogedengbe et al., "Strategic Data Integration for Revenue Leakage Detection: Lessons from the Nigerian Banking Sector," International Journal of Multidisciplinary Research and Growth Evaluation, 2022. [Online]. Available: https://www.allmultidisciplinaryjournal.com/uploads/archives/20250814114326_MGE-2025-4-277.1.pdf

[17] Phaktada Choowan and Hanvedes Daovisan, "Artificial Intelligence in Data Governance for Financial Decision-Making: A Systematic Review," MDPI, Dec 2025. [Online]. Available: https://www.mdpi.com/2504-2289/10/1/8

[18] Amrina Kausar et al., "Optimizing inventory and pricing decisions in a closed-loop supply chain: a sustainable approach towards manufacturing and remanufacturing," ScienceDirect, Sep. 2025. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2772390925000228

[19] Lingxuan Kong and Alexandra Brintrup, "A causal artificial intelligence model for payment delay optimization in supply chain financing," ScienceDirect, Sep. 2025. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S294986352500038X

[20] Jean Baptiste Bernard Pea-Assounga et al., "Effect of financial innovation and stakeholders' satisfaction on investment decisions: Does internet security matter?" ScienceDirect, 2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2405844024032730