Payments AI

Valuing Fraud Reduction, Authorisation Uplift and False-Decline Control

Authored by Chennakeshav Adya, Independent Researcher

September 2026

Abstract

Artificial intelligence is increasingly used in payment authorisation, fraud scoring, authentication, routing, tokenisation, dispute management and merchant risk. The commercial proposition appears simple: stop more fraud, approve more legitimate payments and reduce manual work. The economics are harder. A lower fraud rate can be produced by declining more good customers; a higher approval rate can admit loss-making transactions; chargebacks can arrive after revenue is recognised; issuer, acquirer, network, merchant and customer decisions interact; and model benefits can be confounded by product mix, geography, authentication, tokenisation, seasonality or policy changes. This paper develops an evidence-led framework for valuing payment AI through realised transaction and cash outcomes. It constructs a payment decision funnel, fraud-loss bridge, false-decline cohort analysis, net-value waterfall and valuation scenarios. Forty modules cover the economic perimeter, event lineage, authorisation denominators, legitimate approvals, fraud and scam taxonomies, false declines, authentication, tokenisation, routing, counterfactual testing, model drift, customer outcomes, disputes, operating cost, funding, regulation, independent validation and transaction readiness. Five figures, five tables, eight frequently asked questions and twenty-six authoritative references support fact-specific diligence. Payment-network case studies and management forecasts require independent validation before use. The framework does not substitute for legal, accounting, tax, regulatory, cybersecurity, valuation or investment advice.

JEL Classification: G21, G23, G24, G32, O31

Keywords: artificial intelligence, payments, fraud, authorisation, false declines, chargebacks, payment orchestration, model governance, equity valuation, fintech

1. Define the equity decision

The work should begin with the decision facing the board. A primary equity raise, strategic investment, acquisition or secondary sale requires an explicit amount, use of proceeds, runway, valuation range, operating milestone and decision date. Payment-AI analysis should serve that decision and identify the evidence required before claimed benefits enter the transaction model.

The central question is whether the company can increase risk-adjusted payment contribution at larger scale while controlling fraud, customer harm, disputes, liquidity, compliance and operational dependence. Deployed and independently validated capabilities should be distinguished from experiments and product roadmaps. The financing model should include implementation, data, network, cloud, review, remediation and working-capital requirements before assigning value to future uplift.

2. Establish the legal and economic perimeter

The perimeter should identify the merchant, payment facilitator, gateway, processor, acquirer, issuer, scheme, wallet, fraud vendor and settlement bank for each flow and geography. It should show who makes the authorisation decision, who bears fraud and chargeback liability, who holds customer funds, who controls the model and which entity earns each fee.

Economic exposure can return through contractual routes even when the platform describes itself as asset-light. Reserves, indemnities, rolling holds, scheme assessments, fraud guarantees, minimum commitments, merchant credits and dispute obligations can transfer loss or liquidity risk. Contracts, regulatory permissions, accounting and observed cash conduct should be reconciled before authorisation or fraud improvements are included in equity value.

3. Build a source-controlled payment ledger

The evidence base should retain a stable identifier from checkout or payment initiation through authentication, authorisation, capture, clearing, settlement, refund, dispute, chargeback and recovery. Each event should preserve timestamp, amount, currency, merchant, channel, geography, credential type, issuer, acquirer, model version, score, rule result, decision, response code and liability position.

Table 1. Payment evidence ledger

Stage	Primary evidence	Control question
attempt	checkout or initiation event	what customer intended to pay?
decision	authentication, score, rule and response	can the outcome be reproduced?
completion	capture, clearing and settlement	did authorised value become settled cash?
loss	refund, fraud, dispute and chargeback	who bore the economic loss and when?
customer outcome	abandonment, retry, complaint and retention	what happened after friction or decline?
economics	revenue, scheme fee, processing cost and reserve	what contribution reached cash?

Fields should be adapted to the product, data rights, regulatory perimeter and accounting model.

Derived fields should retain calculation logic, source timestamp, owner and quality status. Manual corrections, reversals and missing network messages should remain visible. The ledger should reconcile both ways: from the general ledger to transaction events and from event populations to revenue, receivables, reserves and cash. Dashboard summaries without event lineage cannot establish a transaction-grade counterfactual.

Reconciliation should retain message multiplicity and timing. One purchase intent can create authentication, authorisation, reversal, advice, capture and clearing messages across several systems, with local and UTC timestamps. The evidence model should identify the authoritative event for each purpose and document late-arriving or out-of-order messages. Currency conversion should preserve transaction, settlement and reporting currencies with the applicable rates. Partial capture, split shipment and incremental authorisation require explicit linkage because a one-to-one assumption can distort both approval and loss. Sampling should trace selected

purchase intents through every party's record and explain gaps. Population controls should then compare daily counts and values across gateway, processor, acquirer, settlement and accounting sources. This foundation allows technical, financial and customer teams to analyse one controlled universe.

4. Define a stable event taxonomy

Payments systems often use different meanings for attempt, authorisation, approval, capture, completion, fraud and chargeback. The metric dictionary should define every event, unit, currency conversion, cut-off and inclusion rule. It should distinguish customer-initiated and merchant-initiated payments, card-present and card-not-present channels, recurring credentials, wallets, account-to-account transfers and alternative methods.

A stable taxonomy prevents apparent uplift from a denominator change. Removing duplicate attempts can raise approval rate; excluding soft declines can change retry performance; counting an authorised transaction before capture can overstate revenue; and treating all disputes as fraud can misclassify service failures. Historical reporting should be rebuilt under the approved definitions before periods, models or counterparties are compared.

Metric ownership should follow the economic event. Product can own checkout conversion, risk can own fraud classification, finance can own revenue and cash, and operations can own disputes; a designated data owner should reconcile their definitions. Every board metric should state numerator, denominator, exclusions, source, refresh, maturation and accountable owner. Changes should create a new version with a restated history or a visible break. This discipline is particularly important during acquisitions and platform migrations, when legacy systems may use different response codes and cut-offs. The diligence team should produce a crosswalk and quantify unresolved differences. A transaction narrative should use only the reconciled series, with legacy management indicators retained as supplementary evidence.

5. Map the payment decision funnel

The decision funnel should connect attempted value to authenticated, authorised, captured, settled and retained legitimate value. It should also show policy decline, suspected fraud, issuer decline, technical failure, customer abandonment, cancellation, refund, dispute and confirmed fraud. Counts and values should reconcile at each transition.

Figure 1. Payment decision funnel

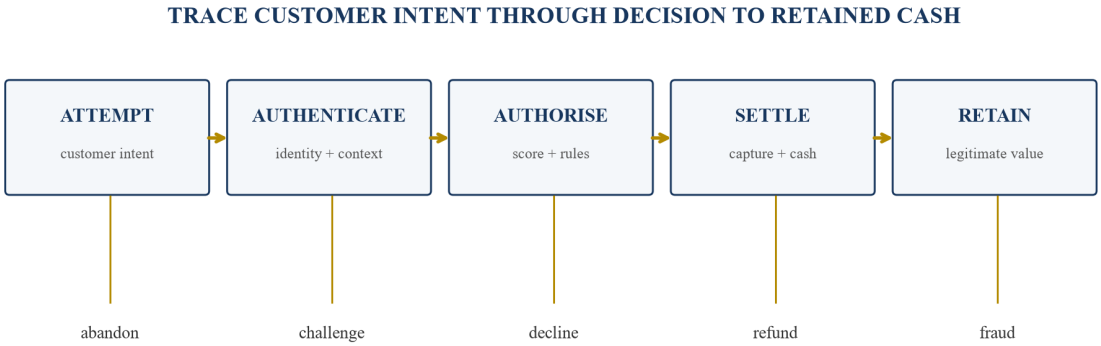

Populate each stage with transaction counts, value, conversion, ownership and observed loss.

The funnel should be segmented by product, geography, channel, credential, issuer, acquirer, merchant and customer cohort. It should preserve retry chains so one purchase intent is not counted as several independent opportunities. A company should state whether approval uplift relates to attempts, unique intents, authorised value or settled value and should connect it to retained cash after later losses.

6. Define authorisation rate correctly

Authorisation rate usually divides approved authorisation requests by eligible requests, by count or value. The definition should address reversals, partial approvals, retries, zero-value checks, balance enquiries, duplicate messages and technical failures. Unique customer intent provides a useful complementary denominator where repeated attempts are material.

Count and value rates should be shown together. A model can improve transaction-count approval while declining large purchases or can approve more value by admitting a small number of high-ticket risks. Approval should also be separated from capture and settlement. A transaction that is authorised but never fulfilled or collected does not create the same economic value as retained cash.

Benchmarking requires comparable eligibility. Issuer and network approval statistics may exclude different message types, regions or failure categories. Merchant performance can vary with basket, customer tenure, recurring status, card funding, local acquiring and cross-border share. The company should therefore benchmark internally across stable cohorts before using external figures. For each rate, a waterfall should explain policy decline, issuer decline, authentication failure, technical failure and abandonment. Volume-weighted and merchant-weighted results can answer different questions. A large merchant can dominate the portfolio rate while smaller merchants experience deteriorating acceptance. Distribution statistics and concentration should accompany averages, especially when the investment case assumes that a best-performing route or segment can be scaled across the portfolio.

7. Identify legitimate approvals

The desired outcome is legitimate completed payment rather than approval alone. A legitimate approval should satisfy the approved fraud definition, complete settlement, survive the relevant dispute window and remain economically positive after directly attributable costs. Where fraud labels mature slowly, the analysis should state the observed percentage and uncertainty.

The company should construct a matrix of approved legitimate, approved fraudulent, declined legitimate and declined fraudulent transactions. This confusion matrix should be measured by count, value and contribution. The cost of a false negative includes fraud, chargeback and operational consequences. The cost of a false positive includes lost margin, customer friction and potential lifetime-value effects. Thresholds should reflect both.

Outcome confirmation should be conservative. A transaction with no reported chargeback after a short window is not necessarily proven legitimate, and a customer complaint does not always establish fraud. The company should define a hierarchy of evidence and a finalisation rule. Where labels remain uncertain, probability-weighted value can be shown as sensitivity rather than converted into a definitive class. High-value transactions can be reviewed individually while large low-value populations use statistical estimation. The valuation model should reflect the same classification and maturation logic as model validation. Otherwise technical teams may report one legitimate-approval figure while finance forecasts another. A controlled bridge should identify every difference caused by cut-off, currency, liability, recovery or label maturity.

8. Build the approval-quality bridge

The approval-quality bridge begins with attempted value, removes ineligible or duplicate messages, adds incremental legitimate approvals, deducts incremental fraud and disputes, adjusts for capture and fulfilment, and arrives at retained contribution. It should be calculated for the operating company and for counterparties whose pricing, reserves or willingness to route volume affects enterprise value.

Attribution requires a controlled comparison. A higher rate after launch may reflect issuer mix, seasonality, new credentials, tokenisation, authentication, routing, merchant category or macro conditions. The bridge should state which variables changed, identify the counterfactual and report confidence intervals or outcome ranges. Benefits that cannot be separated from co-occurring changes should remain a combined programme result rather than a model-only claim.

Management should also test persistence. An initial approval improvement can decay as issuers, fraudsters, customer mix or traffic acquisition change. Monthly cohorts should track eligible volume, incremental decisions, legitimate settlement, later loss and net contribution for long enough to observe the relevant dispute cycle. Benefit concentration in one issuer or merchant should be stated because it may reflect a counterpart-specific fix rather than general model advantage. If teams use a synthetic control or matched cohort, they should report matching variables, pre-period fit and sensitivity to alternative matches. If no credible counterfactual exists, the result can support operational learning and should remain outside the validated valuation bridge until a better test is completed.

9. Construct the fraud taxonomy

Fraud should be divided into relevant modes: stolen credential, account takeover, application or identity fraud, merchant fraud, first-party misuse, friendly fraud, authorised push-payment scam, mule activity and internal abuse. Network and regulatory reporting definitions may differ from economic loss definitions, so crosswalks should be documented.

Each category has different signals, intervention rights, liability and time to confirmation. A credential model may perform well against stolen-card fraud and poorly against manipulation of legitimate customers. The 2025 EBA and ECB report found strong customer authentication remained effective for fraud types it was designed to mitigate while payer manipulation increased. The model inventory and valuation case should therefore avoid treating fraud as one stationary target.

10. Measure gross and net fraud loss

Gross fraud loss should include the value of confirmed fraudulent transactions under the approved definition. Net loss should then reflect recoveries, liability shifts, insurance, merchant or issuer allocation and direct dispute costs. The analysis should show who ultimately bears loss rather than assuming the platform's initial debit is final.

Fraud rate should be calculated by count and value using consistent transaction populations. Low rates can still be economically significant at high scale, and percentage changes can appear large when the base is small. Absolute loss, basis points, loss per active customer and risk-adjusted contribution should be presented together. Confirmation lag and unsettled disputes require an explicit reserve methodology.

11. Build the fraud-loss bridge

Figure 2. Fraud-loss bridge

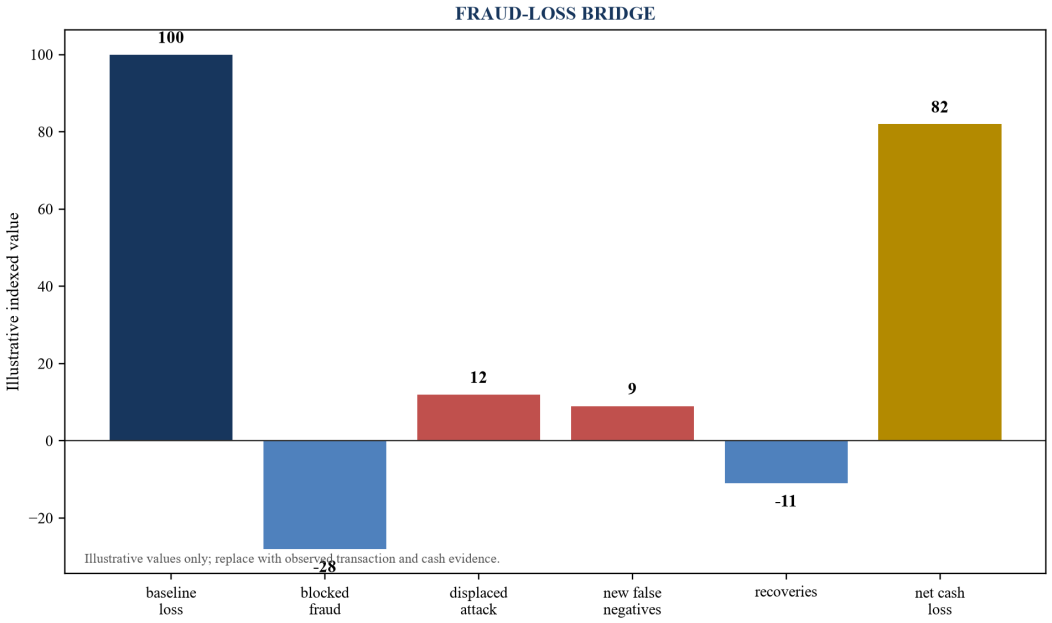

The bridge separates prevention, displacement, liability and cash timing using company-specific evidence.

The bridge should begin with a comparable baseline and separate model-detected fraud, rule-detected fraud, authentication effects, manual review and external network controls. It should

identify displacement across merchants, channels, credentials and fraud types. Blocking one attack can redirect adversaries rather than remove system-wide loss.

The final step should reconcile confirmed economic loss to cash. Chargeback timing, provisional credits, representment, recoveries and scheme assessments can span reporting periods. The bridge should therefore contain development triangles or vintage views for immature loss, with adjustments clearly separated from observed outcomes.

12. Define the false-decline problem

A false decline is a rejected transaction that would have been legitimate and economically acceptable if approved. The outcome is often unobserved because the declined payment never completes. A later approval on another credential or route can provide partial evidence, while customer surveys and sampled review can supplement without becoming ground truth.

The company should state its labelling method, coverage and error. Merchant-side systems may observe customer retry and abandonment but not whether an issuer's fraud suspicion was justified. Issuers may observe account history but not merchant fulfilment. Shared evidence can improve understanding, subject to law and contract. False-decline claims should retain these information boundaries.

Estimation can combine several evidence tiers. The strongest observable tier includes a declined purchase intent that succeeds shortly afterwards with the same merchant and materially similar amount after a controlled change. A second tier includes alternative credential or route success, which demonstrates customer intent but does not prove the original risk decision was wrong. Sampled manual review can estimate likely legitimacy, subject to reviewer error and incomplete information. Customer contact can reveal experience and intention, subject to response bias. Statistical reject-inference methods can provide ranges and require transparent assumptions. The company should report results by tier and avoid merging them into one precise headline. Investment analysis should use a conservative range and connect only observable recovery to realised contribution.

13. Build false-decline cohorts

Declined transactions should be grouped by reason, model version, issuer, acquirer, merchant, channel, customer tenure, credential, ticket size and time. Outcomes can include same-card retry, alternative-card success, alternative-payment success, purchase abandonment, customer support, complaint and future retention. Cohorts should use a defined observation window.

Table 2. False-decline cohort analysis

Cohort	Evidence	Economic question
soft-decline retry	response code and later approval	was timing or message quality causal?
suspected-fraud decline	score, rule and review	would controlled approval have been legitimate?
technical decline	gateway and network logs	was value lost through availability or configuration?
new customer	acquisition source and later activity	did friction destroy first-order conversion?
recurring payment	mandate, credential and updater status	could token or credential maintenance recover cash?
high-value purchase	amount, identity and fulfilment	was threshold policy economically calibrated?

Labels and windows require management validation and data-rights review.

The analysis should avoid counting every retry as a fresh lost sale. Unique purchase intent and customer outcome should be the principal units, supplemented by transaction messages for operations. A cohort with high later approval may indicate recoverable friction; a cohort with no retry remains uncertain rather than automatically false.

14. Quantify customer friction

Friction includes authentication challenge, extra data request, review delay, decline, retry, call-centre contact and account restriction. Some friction prevents loss; some merely moves cost to legitimate customers. The model should measure challenge rate, completion, latency, abandonment, complaint and subsequent activity by cohort.

Lifetime-value effects require caution. A declined customer may return, switch credential, use another merchant or leave entirely. The company should connect observed identity-level behaviour where lawful and should keep survey-based estimates separate. Transaction contribution, customer acquisition cost and expected retention can define a value range without presenting an unobserved counterfactual as fact.

15. Optimise fraud and approval jointly

Fraud and approval are linked by the decision threshold. Lowering a threshold can block more fraud and more legitimate payments; raising it can approve more of both. The appropriate objective is net economic value subject to risk appetite, customer-protection, compliance and operational constraints.

The optimisation function can include retained legitimate margin, fraud and dispute loss, manual review, authentication cost, customer-service burden, expected retention, scheme assessments and liquidity. Constraints can set maximum loss, prohibited customer outcomes, fairness limits and resilience requirements. Management should approve the function and review how changing weights alters decisions.

16. Construct the lift-versus-loss matrix

Figure 3. Authorisation uplift versus fraud-loss matrix

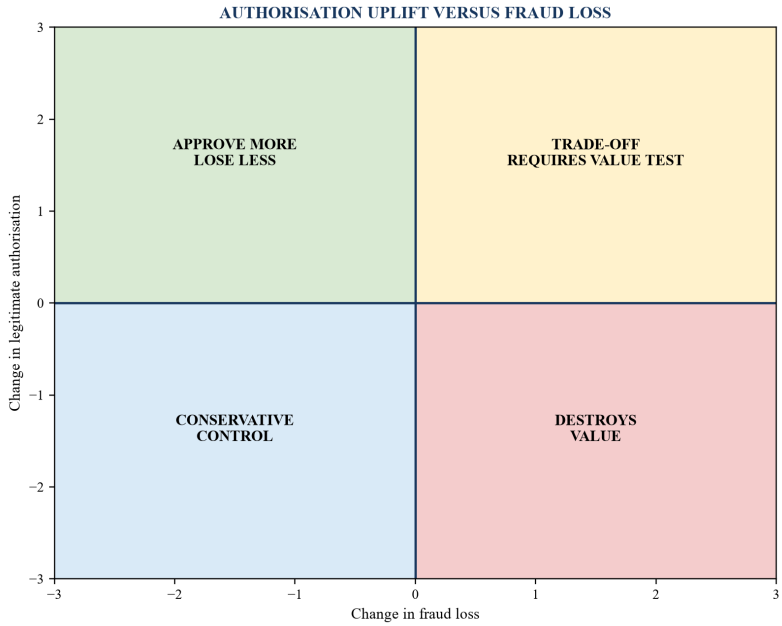

Quadrants require observed and seasoned results; movement between quadrants should be tracked by cohort.

The preferred quadrant increases legitimate approval and reduces fraud loss. Other quadrants may still be rational when contribution, customer protection or regulatory obligations justify the trade. The matrix should use the same baseline, maturity window and economic perimeter for both axes.

Threshold curves can add operational detail. For a range of candidate thresholds, the team should plot legitimate approval, confirmed fraud, manual-review volume, customer challenges and net contribution. The selected point should be compared with the current production threshold and approved limits. Segment-specific thresholds may improve value when risk and economics differ, although they increase governance complexity and potential customer-outcome concerns. The team should test whether segmentation is lawful, stable and supported by adequate observations. Threshold changes should be versioned and linked to the resulting cohorts. During a rapid attack, temporary tightening may be appropriate; its authority, expiry and customer effect should be documented. A permanent valuation assumption should use steady-state controlled performance rather than emergency conditions.

Results should be shown by value and count, with confidence and observation lag. A model may occupy the preferred quadrant overall while harming a high-value segment or geography. Limits and escalation should therefore operate below portfolio average where exposure or customer consequences are material.

17. Design the counterfactual

Randomised controlled testing can provide strong attribution where lawful and operationally safe. Champion-challenger deployment, shadow scoring, stepped rollout, matched cohorts and difference-in-differences can support other settings. The design should prevent leakage between treatment and control and should preserve model, policy and network versions.

The counterfactual must reflect the actual decision being changed. Comparing a new AI model with every historical decline can overstate opportunity if policy, funding or regulation would still reject transactions. Incremental value should include only decisions the intervention could alter,

using matured fraud and customer outcomes. External changes should be measured or included in sensitivity.

Experiment power should reflect fraud rarity, transaction dependence and the value distribution. Millions of low-value payments do not automatically provide adequate evidence for a rare high-value loss. Customer, merchant and account clustering should be addressed in design and analysis. Holdout duration should cover relevant fraud and dispute maturation, while sequential safety monitoring should use pre-agreed boundaries. Teams should report treatment exposure, cross-over, missing outcomes and implementation failures. A statistically detectable approval change may remain economically immaterial after costs, while a small basis-point loss change can be material at scale. The analysis should therefore pre-specify both statistical and economic decision criteria. Where experimentation would expose customers or firms to unacceptable risk, shadow testing and conservative exposure limits can build evidence before any broader controlled rollout.

18. Govern champion-challenger testing

The champion is the approved production decision system; challengers should run under documented objectives, traffic allocation, safety limits and rollback triggers. Shadow tests avoid direct customer impact but cannot fully observe acceptance or loss from changed decisions. Live tests provide stronger causal evidence and require tighter governance.

Each experiment should register hypothesis, population, duration, minimum sample, primary metric, guardrails and analysis method before launch. Repeated monitoring can create selection bias if teams stop when results look favourable. Results should include all planned metrics and deviations. Production promotion should require independent review and accountable approval.

Experiment governance should include operational readiness. Customer support, fraud operations, disputes, finance and incident teams should understand what treatment can change and how to identify affected transactions. Rollback should restore the complete decision stack rather than only a model endpoint. Data pipelines should verify treatment assignment and outcome capture in near real time. An experiment with missing assignment logs cannot support causal value even if portfolio metrics improve. Where several challengers run simultaneously, multiplicity and shared controls should be addressed. Post-test analysis should reconcile experimental outcomes to production reporting and explain any difference. The approved result, limitation and decision should be retained in the model inventory and board evidence pack.

19. Measure authentication economics

Authentication can reduce unauthorised transactions and shift liability, while challenge can add latency and abandonment. The analysis should separate frictionless and challenged flows, exemption use, success, technical failure, authorisation, fraud, chargeback and customer outcome. Geography and issuer behaviour can materially change results.

The 2025 EBA and ECB report provides evidence that strong customer authentication remained effective against fraud types it was designed to mitigate, while manipulation of payers was rising. This supports a segmented control design rather than an assumption that one authentication step addresses every fraud mode. Liability shift should be modelled separately from system-wide fraud prevention.

Authentication optimisation should distinguish risk-based selection from protocol performance. A low challenge rate can reflect effective frictionless authentication, broad exemptions, incomplete issuer participation or missing messages. A high success rate among challenged customers may coexist with substantial pre-challenge abandonment. The funnel should therefore begin before authentication and retain every eligible intent. The company should compare frictionless, challenged, exempt and failed populations on approval, confirmed fraud, disputes and customer return. Regulatory and scheme requirements set constraints; economic optimisation operates inside them. Any change to exemptions, data sharing or challenge routing should receive legal, security and customer-outcome review before launch.

20. Analyse tokenisation and credential quality

EMVCo defines payment tokenisation as replacing the primary account number with a constrained alternative value, supporting security and interoperability. Token status, lifecycle events, domain restrictions and payment account references can affect approval, fraud screening and credential continuity. The company should distinguish network token, device token, merchant token and internal vault reference.

Payment-network case studies may report approval or fraud improvements from token use. These are useful product evidence and require independent validation for the company's merchant, issuer, geography and mix. The transaction ledger should compare eligible token and PAN populations, adjust for selection and show credential-update, approval, fraud and retained contribution outcomes.

21. Test routing and orchestration

Payment orchestration may choose acquirer, route, authentication path, message fields or retry timing. Apparent AI uplift can arise from counterparty differences, local acquiring, currency, credential support and operating configuration. The routing model should record eligible routes, quoted and realised cost, latency, approval, fraud, dispute and settlement reliability.

A successful authorisation on a second route should be linked to the original purchase intent. The analysis should deduct extra processing attempts, customer delay and potential duplicate risk. Route choice should also respect scheme rules, regulatory permissions, merchant agreements and data-transfer restrictions. A routing algorithm that improves headline approval while concentrating dependency or weakening settlement resilience can reduce enterprise value.

22. Govern decline codes and retries

Decline messages should be mapped into a controlled taxonomy covering hard decline, soft decline, suspected fraud, insufficient funds, invalid credential, limit, technical failure and issuer unavailable. Network, processor and internal codes can differ; mappings should be versioned and tested. Ambiguous codes should remain visible rather than forced into a favourable category.

Retry policy should define which transactions may be retried, when, how often and with what updated information. Uncontrolled retry can create fees, duplicate activity, customer frustration and scheme risk. The company should measure recovery at unique-intent level, deduct incremental cost and test whether timing or message enrichment changed the outcome.

23. Measure model discrimination and calibration

Discrimination measures whether higher-risk transactions receive higher scores. Precision, recall, receiver-operating characteristics and lift can support technical assessment, subject to class imbalance and cost. Calibration measures whether predicted probabilities correspond to observed fraud frequency. A well-ranked but poorly calibrated model can lead to unstable thresholds and inconsistent economics.

Metrics should be reported on out-of-time production populations and by material segment. Model performance on confirmed labels should be accompanied by label maturity and representativeness. Because fraudsters respond to controls, validation should distinguish stationary historical patterns from emerging attacks. Threshold decisions should use economic cost and risk constraints rather than a technical statistic alone.

24. Address delayed and disputed labels

Fraud labels can arrive through customer reports, issuer notifications, disputes, chargebacks, merchant review and investigations. These sources have different timing and error. Friendly fraud, service disputes and confirmed unauthorised use should not be merged without a documented purpose. Label reversals should preserve history.

Training and evaluation datasets should apply a consistent maturity window. Recent transactions with incomplete labels can make performance look better than seasoned cohorts. Development triangles should show how reported fraud evolves after transaction date. Forecast loss should include an allowance for unreported and unresolved cases, reconciled to accounting and cash treatment.

25. Monitor drift and adversarial adaptation

Input drift can arise from geography, merchant mix, device, channel, credential, ticket size, issuer or customer behaviour. Concept drift occurs when the relationship between signals and fraud changes. Adversaries can probe thresholds, imitate legitimate patterns and move to weaker routes. Monitoring should therefore cover population, score, decision, outcome and attack indicators.

Alerts need defined owners and actions. A statistical shift may require investigation, tighter limits, rule changes, retraining or rollback. Monitoring thresholds should reflect exposure and expected variation. Teams should record every intervention so subsequent results are not wrongly attributed to the model. Backtests should reproduce the exact production version and decision path.

26. Govern rules, models and human review together

The decision system may combine vendor scores, internal models, deterministic rules, authentication, policy and manual review. The inventory should show precedence, aggregation and exception handling. A rule can dominate a model, making reported model performance irrelevant to actual outcomes. Manual review can create value and hidden cost.

Override records should state reason, authority, direction, amount and outcome. Concentrated overrides can reveal model gaps, policy conflict or operational pressure. Reviewer queues should be assessed for selection, response time, productivity and consistency. Human decisions require

quality assurance and escalation; they should not become an undocumented substitute for model control.

27. Build the governance map

Figure 4. Payments-AI governance map

ACCOUNTABILITY FROM DECISION TO CUSTOMER AND CASH

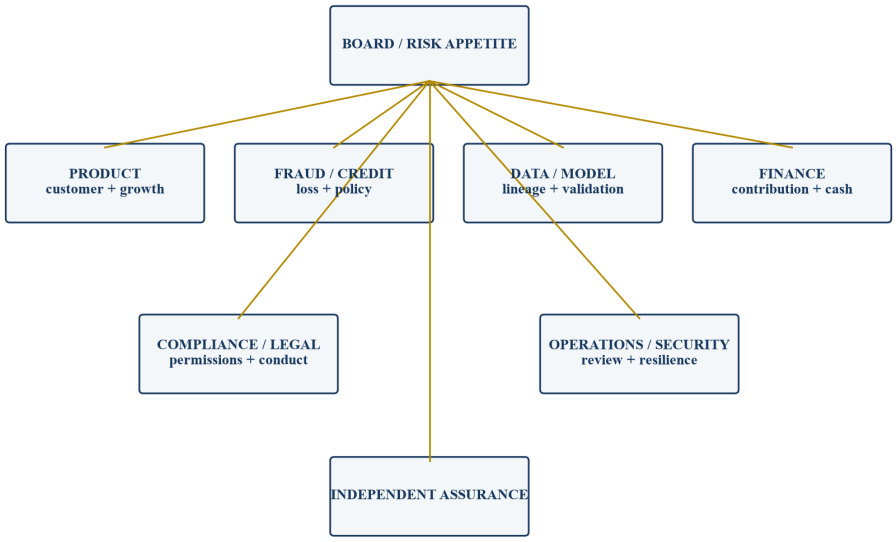

The map should identify named accountable owners, independent review and emergency authority.

The board should approve risk appetite and the economic objective; management should allocate decision rights across product, fraud, data, compliance, operations, finance and security. Model owners should be distinct from independent validators. Emergency action should have narrow authority, logged rationale, expiry and retrospective review.

Governance should connect technical thresholds to financial and customer consequences. A model committee may approve statistical performance while a risk committee sets loss limits and a product team owns conversion. The governance map should create one escalation path when these metrics conflict. Material changes should trigger validation, financial reforecast and customer-outcome review.

28. Perform independent validation

Independent validation should test purpose, data lineage, target, sample construction, model method, implementation, performance, calibration, thresholds, stability, security, explainability and limitations. It should reproduce a sample of production decisions from source events. Findings should have owners, severity, dates and closure evidence.

Table 3. Independent validation programme

Test	Evidence	Transaction relevance
data and labels	point-in-time lineage and maturity	supports reproducible results
implementation	code, version and decision replay	identifies production divergence
performance	out-of-time fraud and approval	tests realised benefit
economics	contribution and cash reconciliation	connects model to value
customer outcomes	decline, challenge and complaint cohorts	identifies harmful friction
resilience	fallback, latency and vendor failure	tests continuity under stress

Scope and frequency should be proportionate to exposure, change and regulatory obligations.

Validation should examine the complete decision system, including rules, vendors and human review. A technically strong component can be neutralised by implementation or policy. Material limitations should affect allowed use, exposure limits and valuation sensitivities. Closure should require evidence rather than management assertion.

29. Govern data privacy, security and resilience

Payment data can include account information, device signals, behavioural data and identity attributes. The company should document lawful basis, purpose, consent where applicable, retention, access, transfer, vendor use and deletion. Feature usefulness does not override legal or contractual restrictions.

PCI DSS provides baseline technical and operational requirements for entities that store, process or transmit cardholder data or can affect the cardholder-data environment. The security review should map tokenisation, encryption, privileged access, logging, software integrity and incident response. Model services also require availability targets, fallback rules, latency budgets and recovery testing. A fraud control that fails during peak volume can expose both loss and revenue.

30. Assess third-party and network dependency

The platform may depend on schemes, acquirers, cloud providers, device-intelligence vendors, identity services, consortium data and fraud APIs. Contracts should be reviewed for data rights, change control, service levels, audit, incident notification, liability, termination, portability and change of control. Minimum fees and volume commitments belong in the economics model.

Dependency concentration should be quantified by transaction and revenue exposure. A proprietary orchestration layer can still rely on external risk scores or network tokens that are unavailable on some routes. The valuation case should distinguish company-owned capability from access to counterparties. Exit and continuity plans should be tested against realistic migration time and regulatory approval.

Concentration can be multidimensional. A company may have several acquirers that rely on the same scheme, cloud region, device-intelligence provider or identity consortium. The dependency map should look through direct contracts to critical sub-processors where evidence is available. Stress tests should model score unavailability, delayed response, degraded data, price increase, contract termination and counterparty risk tightening. Fallback rules can preserve continuity while changing fraud and approval performance; their economics should be tested rather than assumed

equivalent. During diligence, vendors may restrict transfer of historical data or models after change of control. These clauses can affect integration, portability and valuation and should be raised before signing.

31. Map regulation and customer protection

Regulatory obligations depend on activity, entity and geography. In the UAE, the CBUAE Retail Payment Services and Card Schemes Regulation establishes requirements for licensed providers and card schemes, while its responsible-AI guidance addresses consumer protection in AI adoption. The CBUAE 2025 annual report describes work to establish a central anti-fraud operations centre.

The European framework includes payment-services rules, strong customer authentication and fraud reporting. The EBA and ECB publish aggregated fraud evidence under defined reporting methods. The analysis should map permissions, safeguarding, AML/CFT, sanctions, outsourcing, operational resilience, incident reporting, data protection and complaints. Legal conclusions require qualified counsel and should not be inferred from technical architecture.

Regulatory mapping should operate at product and flow level. The same platform can act as technology provider in one route and regulated payment service provider in another. Merchant-of-record, payment-facilitator, agent and outsourcing structures can alter safeguarding, complaint, reporting and liability obligations. Cross-border flows add data-transfer, currency, sanctions and local-acquiring considerations. The company should maintain a permissions matrix that connects entity, activity, jurisdiction, customer, method and responsible regulator. Model changes should be assessed against that matrix, particularly when new data, automated decisions or customer communications are introduced. Open findings, supervisory correspondence and remediation cost should enter transaction diligence and scenario analysis.

32. Measure dispute and chargeback economics

Disputes should be classified by fraud, service, processing error, duplicate, cancellation and other reason. The ledger should connect original transaction, evidence submission, provisional credit, representment, arbitration, liability and final cash. Win rate without dispute selection and direct cost can mislead.

The economics should include staff, vendor, evidence, scheme and customer-service cost. A model may reduce fraud and increase disputes if approved merchants or customers have poor fulfilment outcomes. Merchant-risk controls should therefore connect onboarding, monitoring, reserves and termination to transaction decisions. Chargeback development should be analysed by cohort and reason.

Representment productivity should not be optimised through win rate alone. Teams can select easy cases and improve win rate while leaving recoverable value unworked. The operating model should show eligible disputes, submitted cases, evidence completeness, time to resolution, recoveries and net contribution after labour and fees. Automation can assemble evidence and prioritise cases; source documents, deadlines and accountable approval should remain controlled. Recurring patterns may reveal merchant service problems, misleading descriptors, cancellation failures or account takeover. These findings should feed product and merchant governance rather than remain a back-office statistic.

33. Reconcile operating cost

Payment-AI cost includes data, model development, cloud inference, vendor fees, authentication, orchestration, manual review, disputes, compliance, security, monitoring and customer support. Fixed platform investment should be separated from activity-driven and step-up costs. Claimed automation savings require a cash bridge.

Released review time can create value through vacancies avoided, contractor reduction, increased capacity or improved outcomes. Theoretical hours are not cash savings until converted into an observed operating result. Costs should be allocated at product and segment level where decisions differ. Shared cost allocation should use stable drivers and retain reconciliations.

Inference economics deserve direct measurement. Cost per scored transaction should include feature retrieval, vendor calls, compute, storage, observability and fallback. Peak capacity and latency requirements can create committed cost above average use. A more complex model can improve discrimination while increasing response time, timeout or vendor dependency. The analysis should compare model versions on net value at production service levels. Development and retraining cost should be separated from recurring inference, with capitalisation and amortisation following the applicable accounting policy. Implementation effort for merchants, issuers and acquirers can constrain adoption and should be reflected in rollout timing. Transaction forecasts should include sales, integration, certification and support capacity required to convert technical capability into eligible volume.

34. Build risk-adjusted payment contribution

Risk-adjusted contribution should begin with transaction or subscription revenue and deduct scheme and processing fees, incentives, refunds, fraud, chargebacks, authentication, data, variable cloud, manual review, support and directly attributable compliance. It should reconcile to recognised revenue and cash under the company's accounting policy.

Table 4. Net payment contribution bridge

Bridge item	Evidence	Key risk
monetised volume and revenue	contracts, transactions and ledger	inconsistent take-rate denominator
legitimate approval uplift	controlled cohort and settlement	attribution or maturity gap
fraud and chargeback	seasoned loss and liability	delayed or misclassified labels
network and processing cost	invoices and route records	volume or route step-up
AI and review cost	vendor, cloud and workforce records	omitted implementation burden
retained contribution	cash and accounting reconciliation	timing or perimeter mismatch

Categories and accounting treatment require company-specific validation.

Contribution should be segmented by merchant, product, geography, method and risk cohort. Portfolio averages can hide a loss-making growth segment subsidised by mature merchants. The company should disclose how reserves, settlement timing and working capital affect cash even where accounting contribution appears positive.

The bridge should also reconcile economics among ecosystem participants. A merchant's incremental sale may create acquirer revenue, issuer interchange, network fees and platform take rate, while liability allocation determines who bears later fraud. A platform can improve its own

margin by shifting risk or cost to a merchant, yet that counterparty may respond through repricing, reserves, reduced routing or termination. Durable enterprise value therefore depends on a mutually sustainable result under actual contracts. The analysis should show gross benefit, contractual allocation and behavioural response separately. It should also test volume tiers, minimum fees, cross-border pricing, currency conversion, refund policies and reserve releases. These items can introduce step changes that a constant percentage assumption misses. Cash contribution should be shown by transaction cohort until settlement and dispute windows are substantially complete.

35. Construct the net-value waterfall

Figure 5. Payments-AI net-value waterfall

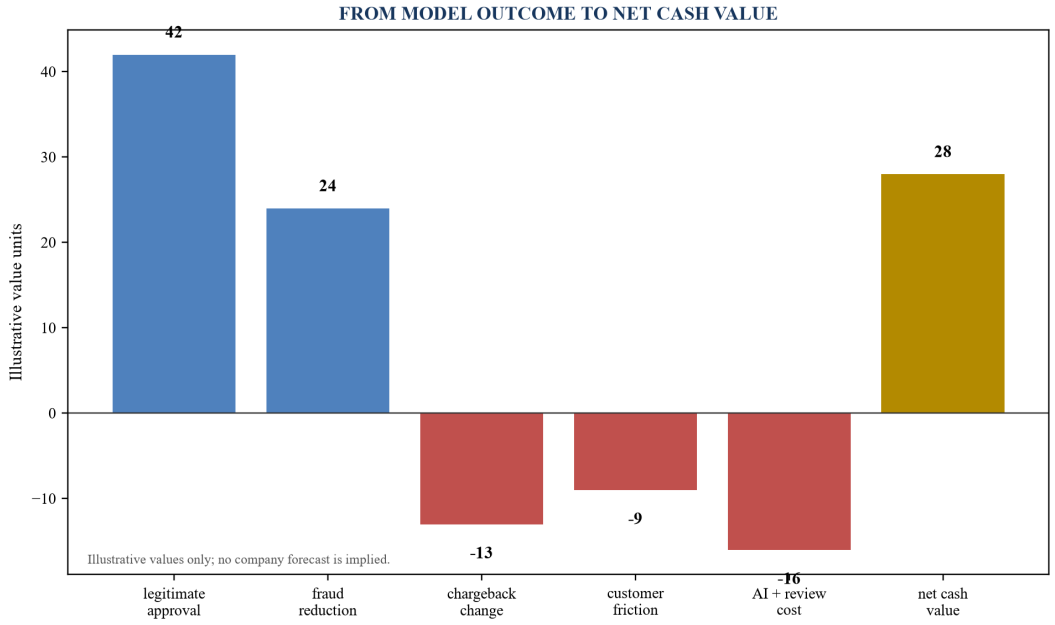

Illustrative structure only; replace values with controlled company evidence.

The waterfall should show incremental retained margin from legitimate approvals and avoided fraud, then deduct incremental fraud admitted, chargebacks, customer friction, authentication, data, cloud, review, remediation and implementation. Timing should be explicit because approved volume and fraud loss can mature in different periods.

Value should be reported as observed historical benefit, annualised run-rate and forecast, each separately. Annualisation requires stable seasonality and production exposure. Forecasts should include ramp, retraining, attack adaptation and capacity. The waterfall should reconcile to cash and should avoid counting the same uplift in revenue, margin and terminal value.

36. Build valuation scenarios

The valuation model should connect eligible payment volume, legitimate approval, fraud, chargeback, take rate, variable cost, operating expense, working capital and cash. Base, upside and downside scenarios should vary model effectiveness, transaction mix, adoption, counterparty pricing, loss emergence and remediation. Unsupported roadmap benefits should remain outside the base case.

Table 5. Payments-AI valuation scenarios

Scenario	Operating evidence	Valuation treatment
validated base	seasoned controlled benefit and stable costs	include in operating forecast with sensitivity
adoption case	validated benefit on limited volume	phase by eligible exposure and capacity
product expansion	limited transfer evidence	probability-weight and fund validation
attack stress	higher fraud, drift and review load	reduce contribution and increase liquidity need
dependency stress	vendor, route or scheme disruption	include migration, downtime and concentration cost

Scenario values are conceptual; no company value or investor return is forecast.

Enterprise value should arise from forecast free cash flow and risk, not from a model-performance multiple. Evidence can influence revenue durability, margin, growth capacity, capital needs and discount assumptions. The valuation bridge should identify which line moves, by how much, why and with what source. Market multiples remain cross-checks that require comparable perimeter, growth, margin and risk.

Scenario design should preserve interactions. Higher approval can increase volume, data and gross revenue while increasing fraud exposure, dispute operations, reserves and settlement concentration. Better fraud control can reduce loss and release risk capacity, although a tighter system may lower conversion. Vendor pricing can decline with scale or rise through premium features and minimum commitments. The model should use coherent combinations rather than independent upside assumptions on every line. Downside scenarios should include attack migration, degraded model performance, major merchant concentration, acquirer restriction, vendor failure, regulatory remediation and liquidity lock-up. Management should identify early indicators and funding responses for each scenario. Valuation conclusions should state the proportion supported by observed outcomes, the proportion dependent on scaling validated performance and the proportion dependent on unvalidated product development.

37. Score transaction readiness

Readiness can be scored across event lineage, metric definitions, counterfactual testing, label maturity, customer outcomes, governance, validation, economics, resilience and regulatory mapping. Each domain should receive a status, evidence reference, owner and remediation date. The score should reveal blockers rather than average them away.

A high-performing model with missing cash reconciliation is not transaction-ready. A complete data room with unresolved customer harm is not transaction-ready. The investment committee should see critical limitations, their financial sensitivity and the decision right to accept or remediate them. Readiness should be refreshed after material model, route, product or fraud changes.

38. Prepare the diligence data room

The data room should include the metric dictionary, legal-entity and payment-flow map, contracts, licences, event schema, reconciled populations, model inventory, development and validation, decision rules, override logs, counterfactual tests, fraud and false-decline cohorts, disputes, customer outcomes, vendor dependencies, incidents, forecasts and board approvals.

Files should have owners, dates, versions and cross-references. Management presentations should trace to source evidence. Personally identifiable and payment data require controlled access, minimisation and lawful sharing. A reproducible analysis pack can provide aggregate cohorts and scripts without exposing unnecessary customer information.

39. Run the ninety-day value-control sprint

Days one to thirty should lock definitions, perimeter and evidence. The team should reconcile attempts through retained cash, map fraud and decline taxonomies, identify model and rule versions, and establish label maturity. Critical data gaps and customer risks should be escalated immediately.

Days thirty-one to sixty should rebuild fraud, authorisation, false-decline and contribution cohorts; perform independent validation; and construct the net-value waterfall. The team should test threshold and routing scenarios, assess vendor and resilience exposure, and connect findings to the operating forecast.

Days sixty-one to ninety should close material findings, approve production limits, complete valuation sensitivities and prepare transaction evidence. The board should decide which benefits are validated, which remain experiments and which require funded milestones. The sprint should produce accountable actions and a controlled data room rather than a promotional AI claim.

40. Conclusion

Payments AI can create equity value when it increases retained legitimate contribution while controlling fraud, disputes, customer harm, operating cost, liquidity and dependency. Approval rate, fraud rate and model lift are incomplete measures on their own. A common event population and counterfactual are required to connect decisions with outcomes.

The proposed framework combines the payment funnel, fraud-loss bridge, false-decline cohorts, joint optimisation, independent validation, net-value waterfall and valuation scenarios. It enables a board or investor to distinguish reproduced operating evidence from roadmap aspiration. Every conclusion remains specific to the company's contracts, permissions, transaction mix, data, controls and observed cash.

References

[1] Central Bank of the UAE. Retail Payment Services and Card Schemes Regulation. https://rulebook.centralbank.ae/en/rulebook/312-retail-payment-services-and-card-schemes-regulation

[2] Central Bank of the UAE. Guidance Note on Consumer Protection and the Responsible Adoption and Use of Artificial Intelligence. https://rulebook.centralbank.ae/en/rulebook/guidance-note-consumer-protection-and-responsible-adoption-and-use-artificial-intelligence

[3] Central Bank of the UAE. Annual Report 2025. https://www.centralbank.ae/media/4qbn11cl/annual-report-2025-en.pdf

[4] European Central Bank and European Banking Authority. 2025 Report on Payment Fraud. https://www.ecb.europa.eu/press/intro/publications/pdf/ecb.ebaecb202512.en.pdf

[5] European Central Bank. Joint EBA-ECB report on payment fraud, 15 December 2025. https://www.ecb.europa.eu/press/pr/date/2025/html/ecb.pr251215~e133d9d683.en.html

[6] European Banking Authority. Payment services and electronic money policy resources. https://www.eba.europa.eu/regulation-and-policy/payment-services-and-electronic-money

[7] European Banking Authority. Guidelines on fraud reporting under PSD2. https://www.eba.europa.eu/sites/default/files/2024-09/767ddecf-75a1-4210-bf48-fc7d0950a7b4/Guidelines%20on%20fraud%20reporting%20%28EBA%20GL-2018-05%29_EN.pdf

[8] Bank for International Settlements, Committee on Payments and Market Infrastructures. Enhancing cross-border payments: addressing fraud. https://www.bis.org/cpmi/pietf/fraud_report_2026.pdf

[9] EMVCo. EMV Payment Tokenisation. https://www.emvco.com/emv-technologies/payment-tokenisation/

[10] EMVCo. EMV Specifications and Associated Bulletins. https://www.emvco.com/specifications/

[11] EMVCo. How EMV Specifications Support Online Commerce. https://www.emvco.com/knowledge-hub/quick-resource-how-emv-specifications-support-online-commerce/

[12] PCI Security Standards Council. PCI Data Security Standard. https://www.pcisecuritystandards.org/standards/pci-dss/

[13] National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework 1.0. https://www.nist.gov/publications/artificial-intelligence-risk-management-framework-ai-rmf-10

[14] National Institute of Standards and Technology. AI RMF Playbook. https://www.nist.gov/itl/ai-risk-management-framework/nist-ai-rmf-playbook

[15] International Organization for Standardization. ISO 20022. https://www.iso20022.org/

[16] International Financial Reporting Standards Foundation. IFRS 15 Revenue from Contracts with Customers. https://www.ifrs.org/issued-standards/list-of-standards/ifrs-15-revenue-from-contracts-with-customers/

[17] International Valuation Standards Council. International Valuation Standards. https://www.ivsc.org/standards/

[18] Financial Stability Board. The Financial Stability Implications of Artificial Intelligence. https://www.fsb.org/2024/11/the-financial-stability-implications-of-artificial-intelligence/

[19] OECD. OECD AI Principles. https://oecd.ai/en/ai-principles

[20] European Union. Regulation (EU) 2024/1689, Artificial Intelligence Act. https://eur-lex.europa.eu/eli/reg/2024/1689/oj

[21] European Union. Directive (EU) 2015/2366 on payment services in the internal market. https://eur-lex.europa.eu/eli/dir/2015/2366/oj

[22] European Union. Commission Delegated Regulation (EU) 2018/389 on strong customer authentication and secure communication. https://eur-lex.europa.eu/eli/reg_del/2018/389/oj

[23] Visa. Helping to Maximize Merchant Success through Authorization and Fraud Prevention. https://corporate.visa.com/content/dam/VCOM/corporate/services/documents/vca-helping-maximize-merchant-success.pdf

[24] Mastercard. Early Detect Insights. https://www.mastercard.com/gb/en/business/insights-intelligence/operations-intelligence/solutions/mastercard-early-detect-insights.html

[25] UK Information Commissioner's Office. Guidance on AI and data protection. https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/artificial-intelligence/

[26] ISO and IEC. ISO/IEC 42001 Artificial intelligence management systems. https://www.iso.org/standard/81230.html

About the Author

Authored by Chennakeshav Adya, Independent Researcher

Chennakeshav (CK) is a corporate finance and investment banking executive with 25+ years of global experience in deal origination, structuring and execution across M&A, growth capital and corporate strategy. He has led value-creation mandates for founders, corporates and funds — bridging the boardroom view to hands-on execution and close.

His career spans Morgan Stanley, HSBC, Lloyds Banking Group, EWEC, ADQ portfolio companies and Emirates Growth Fund, across TMT, real estate, fintech, deeptech, cleantech, infrastructure and energy. He has partnered with C-suite leaders, private equity and venture funds, sovereign wealth funds and family offices to finance complex fund raises and scale-up ventures, and has led M&A due diligence, post-merger integration and business-transformation initiatives to create value.

At Matchpoint Partners he is Managing Partner, leading the firm's corporate finance, M&A and capital-raising practice. He holds an MBA from London Business School, an engineering degree from VTU and a Master of Laws (LLM, in progress) from UCL London.

An active start-up mentor, CK mentors at Techstars, DIFC FinTech Hive, Startup Grind, Founder Institute and IN5, serves as Entrepreneur Mentor in Residence (EMiR) at London Business School, and judges the Entrepreneurship World Cup.