

UK Specialty Insurance Finance Using AI Claims Triage to Test Debt Capacity

*A Decision Framework for Claims Operations Loss Control Cash Conversion
Model Governance and Covenants*

Authored by Chennakeshav Adya

Independent Researcher

September 2026

Abstract

Background. UK specialty insurers and their holding companies can use debt for acquisition finance, refinancing, technology investment and working liquidity. Claims operations influence cash timing, reserve confidence, customer outcomes and management credibility, yet an artificial-intelligence initiative does not create debt capacity merely because it processes files faster. *Objective.* This paper develops a lender-oriented framework for testing whether AI-supported claims triage improves cycle time and loss control sufficiently to support debt sizing and covenants. *Approach.* It reconstructs a policy-to-claim event ledger, segments claims into decision cohorts, separates routing speed from settlement quality, reconciles claims movement with IFRS 17, Solvency UK and cash, and subjects the AI workflow to governance, conduct, third-party and operational-resilience tests. A hypothetical specialty-insurance group illustrates the translation from operating evidence to free cash, debt-service coverage, liquidity and covenant design. *Findings.* Triage creates credit value only when stable cohorts show that earlier classification leads to controlled decisions, reduced rework or leakage, reliable reserving and cash that is legally and prudentially transferable. False positives, model drift, outsourced-service concentration and customer remediation can reverse the benefit. *Implications.* Lenders, boards, insurers and advisers can use one evidence architecture to price implementation risk, distinguish run-rate benefit from temporary working-capital movement, and design financing protections without embedding an unvalidated model forecast. All worked figures are hypothetical analytical assumptions and require borrower-specific verification.

JEL Classification: G22, G21, G32, G17, C53

Keywords: UK specialty insurance, debt capacity, claims triage, artificial intelligence, loss control, cash conversion, covenants, Solvency UK, operational resilience, model governance

1 Define the credit decision before measuring automation

The financing question is whether a specialty-insurance group can meet debt service through cash that remains available after claims, expenses, tax, reinsurance, regulatory capital needs and policyholder obligations. AI claims triage is relevant only when it changes one or more of those cash drivers in a way that can be measured, governed and sustained. A lender should therefore begin with the proposed facility, borrower, guarantors, structural ranking, use of proceeds, maturity, amortisation and expected repayment source. The same automation evidence can have different value for a regulated carrier, a managing general agent, a service company and a holding company.

Claims triage classifies a newly reported claim, tests completeness, identifies urgency or complexity, recommends a route and creates a work queue. It does not determine coverage, liability, reserve adequacy or settlement authority by itself. Faster triage can reduce idle time and direct specialists toward high-risk files. It can also accelerate an incorrect decision, create unfair treatment, hide accumulating exceptions or

move cash payments forward without changing ultimate loss. The lender needs a causal chain from the operating intervention to economic benefit.

The PRA's 2026 insurance priorities identify a softening underwriting cycle in several general-insurance lines and expect boards to maintain pricing and reserving discipline.[1] The same letter notes scrutiny of optimistic underwriting assumptions in internal models. Lloyd's reported a 2025 market combined ratio of 87.6 per cent and a 5.8 per cent major-claims ratio, while emphasising that market conditions were evolving.[9] These system-level measures provide context and do not establish the debt capacity of an individual borrower.

The first deliverable is a credit-question register. Each question states the relevant legal entity, evidence source, observation period, threshold, owner and financing consequence. It prevents a technology case from becoming a substitute for a credit case.

Table 1. Credit questions for AI-supported claims triage

Credit question	Minimum evidence	Failure signal	Financing response
Is repayment cash legally available	entity cash flows, capital policy, dividend history and approvals	profit sits in a regulated entity with limited transfer capacity	size debt to upstreamable cash and add distribution conditions
Does triage improve operating performance	stable claim cohorts, timestamps, outcomes and control groups	speed improves while rework, complaints or severity deteriorate	exclude benefit from base case and require remediation
Are reserves reliable	case movement, actuarial views, back-tests and governance	favourable reserve movement depends on a model override	use conservative cash and capital assumptions
Can the workflow withstand disruption	service map, fallback, recovery tests and vendor dependencies	no tested manual route or data reconstruction	cap reliance and add resilience undertaking
Is the benefit durable	drift, mix, staffing, vendor cost and adoption evidence	performance decays after pilot or depends on one team	haircut run-rate value and require periodic evidence

Evidence should be defined before management selects the performance narrative.

2 Map the insurer group and cash-transfer perimeter

An insurance group can contain regulated carriers, Lloyd's managing agents, corporate members, intermediaries, claims entities, service companies and a holding company. Premium, claim, commission, investment and capital cash flows follow different legal paths. A consolidated income statement can conceal restrictions at the entity that actually earns the cash. Debt capacity begins with a legal-entity map rather than a group earnings multiple.


The lender should locate external debt and intra-group funding, identify guarantees and security, and trace permitted distributions. It should understand the role of regulatory capital, liquidity buffers, reinsurance collateral, trust accounts, funds at Lloyd's and any ring-fenced assets. Cash that supports policyholders cannot automatically be treated as holding-company debt service. Management actions that require regulatory or board approval need realistic timing assumptions.

A specialty portfolio also produces uneven cash. Premium may arrive before claims, while catastrophe, casualty or professional-liability claims can settle much later. Brokerage, acquisition costs, delegated-authority settlements and reinsurance collections affect the path. An apparent cash release may be a timing difference that reverses in a later period. The model should therefore project monthly sources and uses by legal entity, currency and scenario.

Solvency UK reporting applies to relevant reporting dates from 31 December 2024 and provides the prudential frame within which firms report capital and risk.[19] IFRS 17 provides the accounting principles for insurance contracts, including fulfilment cash flows and risk adjustment.[14] The two views answer different questions. A lender needs bridges between accounting profit, regulatory capital, statutory distribution capacity and available cash.

The perimeter memo should list every dependency on claims performance. These include paid-loss timing, case-reserve development, incurred-but-not-reported estimates, reinsurance recoveries, loss-adjustment expense, customer remediation, outsourced adjuster cost and capital effects. Only then can the team determine which claims-triage metrics belong in a financing model.

Figure 1. Legal-entity cash-transfer map for a specialty-insurance group



Proposed map separates regulated cash generation from holding-company debt service.

3 Build the policy-to-claim event ledger

A lender cannot evaluate a claims intervention from a dashboard whose populations and timestamps are unclear. The evidence base should connect policy, exposure, first notification of loss, document receipt, coverage review, reserve movement, assignment, payment, recovery, closure and reopening. Each event needs an effective time, processing time, source system, legal entity, product, delegated authority, currency and accountable user or service.

The ledger should preserve raw events and version every transformation. A late system entry should not be mistaken for a late customer notification. A reopened claim should retain its earlier history. Migrated identifiers should be linked through an auditable crosswalk. Bulk updates, manual overrides and robotic actions should be identifiable. Claims notes may contain useful evidence, but text extraction should point back to the underlying passage and document.

Reconciliation is a prerequisite. Paid loss should connect to bank and general-ledger records. Case reserves, expenses and reinsurance balances should connect to actuarial and financial reporting. Counts should reconcile across claims, policy and workflow systems. Unresolved differences should stay visible by value, volume and age. A model trained on an unreconciled population can produce a precise score that has no defensible economic denominator.

The ledger also supports privacy and access controls. Claim files may contain medical, legal, financial, identity and location information. The ICO's AI and data-protection toolkit is designed to help organisations reduce risks to individual rights and freedoms from AI systems.[17] The project should apply data minimisation, purpose limitation, role-based access, retention, audit logging and current legal

review. Lenders generally need aggregated evidence and controlled samples rather than unrestricted personal data.

Data quality should be measured at the decision point. A field completed after settlement can make a historical model appear better than the live process. The evaluator should freeze snapshots at notification, first review and final outcome, then test which evidence was genuinely available. Missingness can itself be informative, but the team should establish whether it reflects the claimant, broker, product, channel or internal process before using it. A model that treats missing evidence as a low-risk signal can systematically misroute incomplete claims.

The data dictionary should reconcile commercial language with system implementation. Terms such as open, closed, settled, repudiated, recovered and reopened can differ across platforms or delegated handlers. The dictionary should state business definition, source field, transformation, owner and effective date. When a definition changes, historical series should be restated or visibly segmented. This control matters because an apparent performance improvement can result from a reporting rule rather than an economic change.

Table 2. Minimum fields in the claims-credit evidence ledger

Evidence layer	Representative fields	Credit use	Primary control
Policy and exposure	insured, limit, deductible, wording, geography, inception	define risk and cohort denominator	policy-system reconciliation
Claim notification	occurrence, notification, channel, documents, claimant	measure lag and completeness	event-time validation
Triage decision	route, score, reason, reviewer, override	test automation performance	decision log and authority
Reserve and payment	case reserve, expense, payment, currency, approval	connect workflow to cash and liability	finance and actuarial bridge
Recovery	treaty, facultative, notice, collection, dispute	estimate net cash timing	counterparty and ageing review
Outcome and conduct	closure, reopening, complaint, remediation	test quality and customer effect	independent outcome review

The field set connects operating performance with cash and control evidence.

4 Define claims cohorts before comparing performance

The evaluation population should be defined before the result is known. Specialty claims differ by line, wording, jurisdiction, severity, complexity, catastrophe involvement, litigation, coverage uncertainty, delegated handler and maturity. A single average can improve merely because the mix shifted toward simpler claims. The analysis should compare like with like and retain an explicit unknown category.

Useful cohorts include straight-through candidates, low-complexity claims requiring human confirmation, specialist claims, large losses, coverage disputes, litigation, catastrophe claims, vulnerable-customer cases and sanctions or fraud exceptions. The routing model can recommend a queue, while authority remains with designated people. The team should record the original recommendation, final route, reason for override and subsequent outcome.

Historical comparison needs a common observation window. A recent cohort can look favourable because difficult claims remain open. The evaluator should use maturity-adjusted measures and survival-style views where appropriate. Calendar effects also matter. A catastrophe, court decision, repair-cost shock or

staffing change can affect both treatment and comparison groups. A credible design may combine phased rollout, matched cohorts and time-based holdouts.

The baseline should cover enough time to observe reopening, late payment, recovery and complaint outcomes. An eight-week pilot may support a workflow decision but rarely supports a long-tail credit conclusion. Benefits should be separated into observed, annualised and management-estimated categories. Only observed cash and durable cohort evidence should receive full weight in the lender base case.

Volume capacity should accompany accuracy. A route can be statistically correct while sending more work to a specialist team than it can process. Queue age, handler capacity, service level and handoff delay should therefore be measured by cohort. The model threshold should be tested against peak volumes and staff absence, including catastrophe conditions. If a specialist queue saturates, the resulting delay can reverse the claimed benefit and create a new operational concentration.

The evaluation should report confidence intervals or an equivalent uncertainty range where sample size permits. Rare high-severity cohorts may never support a narrow statistical estimate. In those populations, expert review, conservative assumptions and file-level evidence carry greater weight. The credit memorandum should state which conclusions are statistically supported, operationally observed or dependent on judgement.

Figure 2. Claims funnel from notification to controlled route



Proposed funnel keeps exception, specialist and human-review paths visible.

5 Separate routing speed from settlement quality

Cycle time has several components: notification to first contact, evidence completion, coverage decision, reserve establishment, expert appointment, liability decision, settlement agreement, payment and closure. AI triage mainly affects early routing. A faster first assignment does not establish faster or better settlement. The evaluation should retain each timestamp so that management can identify where time moved.

Quality measures include correct routing, first-time completeness, rework, reassignment, reserve change, reopening, complaint, customer contact, expert cost, litigation escalation and recovery preservation. A model can improve median speed while worsening the tail. The 90th and 95th percentile are relevant because long delays can concentrate customer harm and cash uncertainty. Measures should also be segmented by claim type and protected or vulnerability-related characteristics where lawful and appropriate.

FCA ICOBS 8 requires insurers to handle claims promptly and fairly, provide reasonable guidance, avoid unreasonable rejection and settle promptly once terms are agreed.[6] The FCA reported in 2025 that motor insurers expected to pay £200 million of compensation to approximately 270,000 customers after historic underpayments and process changes.[8] The example concerns motor insurance, but the control lesson is broader: an automated deduction or settlement logic can create remediation exposure when it produces unfair outcomes.

Lloyd's Claims Management Principle expects a strategy aligned with the business, appropriate expertise, proactive management, accurate and timely case reserving, oversight of delegated services and executive governance.[22] A credit case should therefore include service quality and control evidence. Cost reduction achieved by suppressing review or customer communication is not a durable benefit.

Table 3. Operating measures and their credit interpretation

Measure	What it can show	What it cannot show alone	Credit interpretation
notification-to-route time	early workflow speed	settlement quality or ultimate loss	useful leading indicator
first-time correct route	decision quality at entry	reserve adequacy	reduces rework if stable
reassignment rate	process friction	reason for complexity	cost and delay signal
paid severity	cash outcome by cohort	adequacy for open claims	combine with maturity and mix
reopening rate	closure quality	long-tail emergence not yet observed	adverse-quality indicator
complaints and remediation	conduct outcomes	full population fairness	downside and control signal
recoveries preserved	net-loss protection	collection timing and counterparty quality	cash benefit after collection evidence

Each measure requires a cohort, denominator, observation window and owner.

6 Test loss control and leakage mechanisms

Loss control means reducing avoidable economic cost while honouring policy obligations. Potential mechanisms include earlier specialist assignment, quicker evidence preservation, fraud referral, subrogation, salvage, medical or legal intervention, repair-network routing and duplicate detection. Each mechanism should identify the affected claims, action, counterfactual, cost and observed outcome. A generic statement that AI reduces leakage is insufficient.

The evaluator should separate gross paid loss, case reserve, allocated loss-adjustment expense, unallocated expense and reinsurance recovery. A cheaper repair can lower gross loss. Faster documentation can preserve recovery. Earlier settlement can reduce legal expense while bringing cash payment forward. An aggressive referral rule can increase investigation cost and customer delay. These paths affect free cash differently.

Counterfactual construction is difficult because claims are heterogeneous. Randomised allocation may be possible for workflow prompts within safe boundaries. More often, a phased rollout, matched cohort or difference-in-differences design is needed. The team should test pre-existing trends and document selection logic. Large claims should be reviewed individually because a few files can dominate the aggregate result.


Savings should be net of model development, data engineering, vendor fees, human review, monitoring, appeals, remediation and resilience costs. Staffing reductions require evidence that capacity has actually been removed or redeployed to productive work. Avoided future recruitment remains a management estimate until the workload and budget change are demonstrated.

Leakage measurement should avoid circular labels. A claim later settled above its early estimate is not automatically leakage; the early estimate may have been incomplete. A negotiated saving against a supplier quote is not automatically economic value; the alternative price may never have been payable. The team should define avoidable cost through documented policy entitlement, benchmark, counterfactual and approval. Recoveries and fraud prevention should be recognised only when the action, amount and cash outcome can be traced.

The lender should examine adverse selection created by process design. If simple claims settle faster while difficult files accumulate, average closure time may improve even as uncertainty rises. An ageing waterfall should show the value and count moving into longer-duration buckets. Large files, disputed coverage and litigation should be reviewed separately so that aggregate efficiency does not conceal tail concentration.

Claims inflation must remain visible. The PRA's review of general-insurance reserving and capital modelling emphasised up-to-date, credible information and realistic assumptions, including claims inflation in technical provisions and capital.[20] Automation performance should therefore be evaluated in real and nominal terms and separated from severity changes caused by price, wage, legal or social factors.

Figure 3. Hypothetical bridge from triage activity to annual cash benefit



Illustrative values show why gross operating claims require deductions before entering the credit case.

7 Reconcile claims movement with reserving and IFRS 17

Claims triage changes operational data before it changes the ultimate liability. Faster first contact may lead to earlier case estimates. A specialist route may strengthen reserves before it reduces ultimate loss. A settlement initiative may accelerate payments and reduce open counts. The credit model should distinguish operational movement from economic improvement.

The reconciliation begins with opening case reserves, incurred-but-not-reported estimates, loss-adjustment expense and reinsurance assets. It then identifies new claims, reserve changes, payments, closures, reopenings, recoveries, discounting, risk adjustment and foreign exchange. The claims system, actuarial analysis and IFRS 17 accounts should bridge to controlled totals. Any difference attributed to AI needs an evidenced cohort and a documented mechanism.

IFRS 17 measurement includes estimates of future cash flows, discounting and a risk adjustment for non-financial risk.[14] A triage model can affect the timing or information set used in those estimates, but it does not remove actuarial judgement. Changes in fulfilment cash flows may affect insurance service

results and balances differently depending on the contract group and service relationship. The lender should avoid translating an accounting movement directly into cash.

The review should test whether management overlays changed around implementation. A favourable reserve release can obscure higher current-period severity. A more complete early case estimate can make incurred results look worse before ultimate outcomes improve. The bridge should show observed cash, expected future cash, modelled ultimate loss and accounting presentation side by side.

Table 4. Claims and finance reconciliation

Layer	Key movement	Evidence	Lending treatment
Operational	routing, documents, assignment, closure	timestamped event ledger	leading evidence only
Case reserve	handler estimate and movement	file history and authority	test consistency and bias
Actuarial	ultimate loss and expense	triangles, models and judgement	central and stress cash view
IFRS 17	fulfilment cash flows and risk adjustment	accounting roll-forward	reconcile; do not equate to cash
Reinsurance	recognised and collected recovery	treaty, notice, ageing and dispute	haircut for timing and credit risk
Cash	paid claims, expenses and collections	bank and ledger	primary debt-service input

The bridge prevents timing, accounting and economic effects from being combined.

8 Establish a champion-challenger triage design

The incumbent process should remain measurable. A rules-based route, human triage team or existing model provides the champion. Candidate machine-learning or language models act as challengers. The purpose is to demonstrate incremental value against a controlled benchmark, not to compare a sophisticated model with no process at all.

Training and validation splits should follow time and claim boundaries. Records from one claim must not appear in both training and test populations. Text created after the routing decision should not leak into a model intended for first notification. Evaluation should cover calibration, accuracy, false-positive and false-negative cost, segment stability, override, missing data and operational capacity. A rank score can help prioritise work without granting settlement authority.

The chosen threshold is an economic and conduct decision. A low threshold may send too many files to scarce specialists. A high threshold may miss severe or vulnerable cases. The team should state the cost matrix and test it under volume and mix stress. Performance should be reported for material cohorts, including rare but high-consequence outcomes.

Bank of England and FCA discussion material identifies claims management as an insurance AI use case and notes that concept drift and lack of explainability can affect the ability to claim and policyholder protection.[3] The IAIS application paper states that insurers remain responsible for understanding and managing AI systems and their outcomes.[15] The practical response is accountable ownership, traceable evidence and authority that remains with the insurer.

Figure 4. Champion-challenger sequence for claims triage



Drift, override, complaints, rework, cash and service capacity are monitored after approval.

Proposed control path combines model comparison with accountable human authority.

9 Govern explainability, authority and model change

Every material decision should have an owner. The claims function owns the service and authority framework. Data and technology teams own controlled implementation. Risk and compliance test conduct and policy obligations. Actuarial and finance functions reconcile reserve and cash effects. Internal audit provides independent assurance according to its mandate. The board or delegated committee approves risk appetite and material changes.

Explainability should match the decision. A route recommendation may need reason codes, source evidence and confidence. A specialist should see the features or documents that caused escalation. A customer-facing decision needs an understandable explanation and an accessible review path. A lender does not need model source code in every case, but it does need evidence that the borrower can explain, challenge, reproduce and suspend the workflow.

The model inventory should identify version, owner, purpose, training period, inputs, vendor dependencies, validation, materiality and fallback. Change control should distinguish retraining, threshold change, new feature, prompt change, vendor-model update and infrastructure change. Each can alter performance. Silent vendor updates are a material control risk when the borrower cannot reproduce prior decisions.

Monitoring should combine technical, operational, financial and conduct signals. Technical drift without outcome deterioration may prompt investigation. Stable technical metrics with rising complaints can indicate a missing dimension. A model may remain accurate while workload shifts overwhelm the specialist queue. Thresholds should trigger investigation, restriction, rollback or manual operation, with named decision rights.

Validation should test implementation as well as mathematics. Production features, preprocessing, thresholds, user interfaces and permissions should match the approved design. A model can pass offline testing and fail in operation because documents arrive differently, staff misunderstand a reason code or downstream systems truncate information. Release evidence should include code and configuration hashes, user acceptance, access testing, sample decision reproduction and reconciliation of the first production cohorts.

Management information should show denominator changes. If the automated route expands from one product to five, an apparently stable error rate may conceal deterioration in a new segment. The dashboard should preserve product, jurisdiction, handler and version views. The governance committee should approve material population expansion and record whether finance may recognise additional benefit.

The PRA's 2025/26 annual report records continued supervision and support for safe AI adoption, alongside operational-resilience work.[2] Governance therefore belongs in the credit case because weak model control can create financial, regulatory and operational loss.

10 Protect customer outcomes and claims fairness

Specialty insurance often serves sophisticated commercial policyholders, yet claims can also affect individuals, employees, small businesses and vulnerable parties. The workflow should identify the applicable conduct perimeter and policy obligations by product and jurisdiction. Automation must preserve access, communication, review and escalation.

The testing plan should look beyond protected-characteristic averages. Proxy variables, missing documentation, language, digital access, geography and claims complexity can create uneven outcomes. False-negative analysis is important when the model fails to route a severe, urgent or vulnerable case. Human override data should be reviewed for systematic disagreement, workload pressure and inconsistent use.

The FCA's Consumer Duty guidance requires firms within scope to act to deliver good outcomes for retail customers.[7] ICOBS claims rules require prompt and fair handling.[6] Where a specialty group handles both commercial and retail-linked claims, the operating design should avoid allowing a less protective workflow to spill into a regulated population. The borrower should maintain product-level rules and evidence.

Complaints, ombudsman outcomes, remediation, ex gratia payments and litigation should feed the credit model. A small apparent saving can be overwhelmed by a systemic remediation programme. The lender should receive timely notice of material control failures, regulatory enquiries or remediation estimates where relevant to the financing documents.

Customer outcome testing should remain independent of the team rewarded for efficiency. The evidence pack should include samples of correctly routed, incorrectly routed, overridden and appealed cases. Findings should lead to threshold, process, training or data changes, with closure verified.

11 Test third-party and operational resilience

Claims triage may depend on cloud infrastructure, model providers, data vendors, loss adjusters, document processors and communications platforms. A failure can interrupt customer contact, routing, reserve updates and payments. The service map should identify critical components, data location, subcontractors, concentration and recovery dependencies.

The PRA's operational-resilience framework requires firms within scope to identify important business services, set impact tolerances and remain within them during severe but plausible disruption.[4] The PRA's 2026 policy on incident and third-party reporting responds to increasing reliance on external services and introduces reporting expectations for certain incidents and material arrangements.[5] Existing outsourcing expectations also require disciplined due diligence, contracting, access, audit, security, continuity and exit.[21]

A lender should examine tested fallback. A manual route needs current procedures, trained staff, capacity and access to the necessary data. A promise that staff can revert to spreadsheets is weak when claim volume, permissions or document location make that impossible. Recovery tests should demonstrate queue reconstruction, decision integrity and reconciliation after restoration.

Vendor exit should cover model artefacts, data, prompts, configurations, decision logs, intellectual property, transition support and deletion. Concentration matters when many insurers depend on the same provider or model. A shared failure can occur during a catastrophe, when claim volume is already elevated.

Operational-resilience cost belongs in the business case. Redundant services, monitoring, recovery testing and manual capacity reduce headline savings while making the benefit more financeable. The lender should prefer a smaller evidenced benefit with resilient delivery over a larger estimate that disappears during stress.

Table 5. AI claims-control matrix

Risk	Preventive control	Detective control	Recovery action
incomplete or corrupt input	schema, reconciliation and mandatory fields	exception volume and source drift	quarantine cohort and reconstruct data
wrong route	bounded authority and threshold design	sample review, override and outcome testing	return affected queue to human triage
model drift	time-based validation and change control	calibration and cohort monitoring	restrict, retrain or roll back
unfair outcome	feature review and product rules	complaint, appeal and segment testing	suspend rule and remediate cases
vendor outage	resilient architecture and exit plan	service and queue monitoring	tested fallback and reconciliation
unauthorised change	role controls and version approval	configuration audit	restore approved version and investigate

Controls should have evidence, owners, thresholds and tested actions.

12 Convert operating evidence into cash conversion

The credit model should begin with observed cash and then add supported adjustments. Claims paid, loss-adjustment expense, recoveries collected, vendor cost, payroll, remediation and implementation expenditure should be mapped monthly. Reserve movement informs future cash, but it is not itself cash. Faster settlement can reduce expense and uncertainty while temporarily increasing cash outflow.

Benefit categories should be explicit. Permanent cost removal can enter the run-rate case after implementation evidence. Reduced leakage can enter after mature matched cohorts and controls. Faster reinsurance collection can improve liquidity when notices and recoveries are documented. Earlier claim payment can improve customer outcomes and reduce expense, but its near-term working-capital effect may be negative. Avoided catastrophe losses or assumed lower ultimate severity require especially strong evidence.


The model should distinguish regulated-entity free cash from group free cash. It should deduct required capital growth, tax, investment, policyholder obligations and permitted-distribution constraints. Management fees or service charges should be tested for substance, transfer pricing, contractual basis and regulatory acceptability. A historical dividend provides evidence but does not guarantee future approval.

Sensitivity should cover claim volume, severity, inflation, model precision, rework, vendor cost, outage, recovery delay, capital requirement and distribution timing. The downside should assume that management cannot instantly remove staff or investment when performance deteriorates. Cure actions should have owners, lead times and operational consequences.

Working-capital interpretation requires particular care. Earlier settlement can reduce uncertainty, expense and complaints while increasing near-term cash outflow. Faster recovery can offset part of that movement, subject to reinsurer notice and collection. The monthly model should therefore separate timing benefit from ultimate benefit and show the maximum cumulative cash requirement. A lender may support a temporary liquidity need through a revolving tranche without treating it as permanent earnings.

Capital effects should remain scenario-specific. Better data and control can improve confidence, yet regulatory capital does not fall automatically when an operating KPI improves. Model, parameter or internal-model changes can require governance and supervisory engagement. The base case should retain current capital assumptions until the borrower has an approved and evidenced basis for change.

Figure 5. Claims KPI to debt-service cash waterfall



Proposed bridge keeps regulated capital and transferability between operating benefit and debt service.

13 Size debt through scenarios and structural subordination

Debt sizing should use recurring cash under central and downside cases, together with minimum liquidity and capital constraints. A lender may calculate interest coverage, fixed-charge coverage, debt-service coverage and leverage, but the denominator must reflect cash actually available to the borrower. Consolidated EBITDA can overstate capacity when value is trapped in regulated entities or supported by reserve releases.

The base case should exclude benefits that remain in pilot, depend on unapproved staffing changes or lack mature outcome evidence. A partial benefit can enter a lender case with a haircut and milestone. Downside scenarios should include higher severity, slower reinsurance, automation rollback, regulatory remediation, vendor failure and constrained distributions. These stresses can interact.

Structural subordination matters when holding-company creditors depend on distributions from an operating insurer whose policyholders and regulatory requirements rank economically ahead. Security over holding-company assets may have limited value if the core cash and licences remain within regulated subsidiaries. Intercreditor arrangements, double leverage and intra-group instruments require review.

The facility purpose also shapes risk. Funding an acquisition introduces integration and reserve uncertainty. Refinancing can reduce maturity risk. Funding technology investment creates execution risk before benefit. A revolving facility can absorb volatility, but repeated drawings may reveal structural cash weakness. Each use needs a repayment path and stop-funding conditions.

Debt capacity should be presented as a range, not a single output. The range reflects cash volatility, distribution timing, capital headroom, evidence quality and covenant design. A lower amount with accessible liquidity may create more resilience than maximum opening leverage.

Refinancing risk should be tested independently of initial affordability. A five-year facility can mature during a different underwriting cycle, after a catastrophe or while the technology platform is being replaced. The model should show residual principal, expected cash accumulation, market-access

assumptions and contingency actions. A bullet repayment that depends on a higher future valuation deserves a separate downside and a funded preparation timetable.

Currency and rate risk also matter. Specialty groups can write risk, hold investments, pay claims and receive recoveries in several currencies while debt service may be sterling. Natural offsets should be demonstrated by entity and timing. Hedging cost, collateral, breakage and accounting treatment should enter cash forecasts. The lender should avoid counting the same liquidity as both claims support and derivative collateral.

14 Work through a hypothetical UK specialty-insurance case

Consider a hypothetical UK specialty-insurance group with a regulated carrier and a holding company. The example is an analytical construct and does not describe an identified firm. Management reports £900 million of net insurance revenue, £540 million of claims and related expenses paid, a 92 per cent combined ratio, £1.05 billion of gross claim liabilities, £735 million of eligible own funds and a £430 million solvency capital requirement. The holding company seeks a £180 million five-year term facility to refinance existing debt and fund claims-platform investment.

Management estimates that AI triage can create £18 million of annual gross benefit through lower handling cost, reduced leakage and faster recoveries. Diligence accepts £7 million in the central run-rate case after deducting vendor and compute cost, retained human review, rework, implementation support and timing-only movements. The lender excludes the remaining £11 million until cohort evidence matures. The central case assumes that only £32 million of annual cash can be distributed after capital, liquidity, investment and board constraints.

Annual cash interest and scheduled amortisation total £24 million in the first full year. Central debt-service coverage is therefore 1.33 times before discretionary management action. An adverse case adds £28 million of net claims cash, delays £12 million of reinsurance recovery, removes the £7 million triage benefit and restricts distribution to £22 million. Coverage falls below one time. A severe case adds an operational incident and remediation cost, requiring liquidity support and an amortisation cure.


The credit conclusion is conditional. The facility can be sized below management's maximum request, include a committed liquidity reserve, defer part of amortisation until the operating evidence matures and require capital and distribution headroom. The lender should not capitalise the full AI case at closing.

Table 6. Hypothetical debt-capacity case

Measure	Management plan	Lender central	Adverse	Severe
annual gross AI benefit	18	18	8	0
accepted net recurring benefit	14	7	0	0
transferable cash before debt service	43	32	22	12
annual debt service	24	24	24	24
debt-service coverage	1.79x	1.33x	0.92x	0.50x
minimum liquidity after debt service	38	27	13	4
model availability	full	full with fallback	restricted	unavailable

All figures are illustrative assumptions in £ million unless stated otherwise.

Figure 6. Hypothetical debt-service coverage sensitivity



Illustrative sensitivity combines accepted triage benefit with claims stress.

15 Design covenants around evidence and control

A covenant package should protect the repayment path without turning the lender into the claims manager. Financial covenants can include leverage, interest or debt-service coverage, minimum liquidity and limits linked to capital or distributions. Definitions should prevent reserve releases, uncollected recoveries or unvalidated savings from inflating the calculation.

Information undertakings should focus on material evidence. A monthly or quarterly pack can include claim volumes, routing accuracy, exceptions, rework, cycle-time distribution, reopened claims, complaints, material reserve movement, recoveries ageing, model changes, incidents and fallback tests. Thresholds should reflect the portfolio and avoid false precision. The borrower should explain population or methodology changes.

Performance triggers can require a remediation plan, independent review, benefit exclusion or restriction of the automated route. A model failure should not automatically create an event of default when a tested manual process protects service and cash. Escalation should follow materiality, persistence and management response. Serious regulatory, data, fraud or reporting failures may require stronger consequences.

Definitions determine usefulness. Transferable cash should exclude restricted accounts and require evidence of the applicable distribution condition. Net recurring benefit should deduct all direct and

retained control cost. A material model change should cover threshold, training population, vendor model, feature, decision scope and fallback. These definitions reduce disputes and make reporting reproducible.

Headroom should reflect reporting lag. A quarterly covenant can remain compliant while a weekly claims deterioration consumes liquidity. Early-warning triggers can use queue age, recovery ageing, paid-severity movement and capital forecast. They should prompt information and action before a formal financial breach. The borrower retains operational decision authority, while the lender gains time to assess exposure and proposed cures.

Cash controls can include distribution conditions, minimum holding-company liquidity, debt-service reserve, limits on additional debt and restricted payments. The agreement should address acquisitions, disposals, portfolio transfers, reinsurance changes and material outsourcing. The lender needs notice when these actions alter the cash perimeter or risk profile.

Table 7. Illustrative covenant and information architecture

Protection	Illustrative design	Evidence	Cure path
debt-service coverage	minimum ratio using transferable cash	entity cash bridge and distribution evidence	cash retention, equity cure or amortisation
liquidity	minimum unrestricted holding-company cash	bank and treasury reconciliation	restricted payments and liquidity plan
capital distribution condition	no distribution below agreed headroom	regulatory capital return and forecast	retain cash in regulated entity
model-performance trigger	threshold for persistent error, rework or drift	governed cohort dashboard	restrict model and execute remediation
material incident notice	prompt notice of service, data or conduct event	incident register and impact assessment	fallback, customer protection and recovery
information undertaking	periodic claims-credit evidence pack	reconciled operating, actuarial and finance data	independent review if deficient

Draft terms require transaction-specific legal, regulatory, accounting and tax advice.

16 Build an information pack lenders can verify

The financing data room should contain the group structure, licences, capital policy, historical distributions, debt documents, reinsurance, claims strategy, delegated-authority arrangements, model inventory, validation, incidents, complaints and controlled financial reconciliations. A concise index should state owner, reporting date, version and limitations.

The operating pack should provide claim-level or aggregated cohort evidence suitable for the lender's diligence rights and confidentiality obligations. It should define each KPI and preserve numerator, denominator, exclusions and observation window. Management presentations should reconcile to the data extracts. Differences should be explained before credit approval.

The financial pack should bridge insurance service result, operating profit, capital generation, regulated cash, permitted distributions and holding-company cash. It should show historical forecast accuracy. Reinsurance receivables need counterparty, dispute and ageing detail. Large claims and catastrophe exposures need individual or event-based review where material.

The AI pack should include the decision statement, model card, data lineage, performance by cohort, validation, change history, human authority, monitoring, vendor terms, resilience and rollback. Source artefacts should be retained so that the borrower can reproduce the approved version. A screen demonstration is supplementary evidence.

The lender should record open items in an evidence ledger with decision consequence. Missing evidence can lead to a haircut, condition precedent, drawstop, covenant, reserve account or reduced commitment. This makes uncertainty explicit rather than burying it in a narrative qualification.

17 Execute a staged financing and implementation plan

A staged plan aligns funding with evidence. Stage one establishes the ledger, reconciliations and baseline. Stage two validates the challenger on frozen historical cohorts. Stage three runs a bounded pilot with human authority. Stage four expands only after operating, conduct, cash and resilience thresholds are met. Stage five incorporates the evidenced benefit into forecasts and covenant calculations.

Facility structure can mirror these stages. An initial refinancing tranche may rely on existing cash generation. A delayed-draw or incremental tranche for technology can require validated milestones. Amortisation can step up after implementation. A pricing ratchet may respond to demonstrated coverage rather than a self-reported AI score.

The first 100 days should name owners and dates. Claims operations complete data and process mapping. Finance produces the entity cash bridge. Actuarial teams reconcile liability views. Risk and compliance approve conduct testing. Technology completes resilience and exit tests. Treasury integrates the evidence into liquidity forecasts. The board reviews risk appetite and the lender pack.

Benefits should enter budgets when responsibility for delivery is accepted. Staff, vendor and control costs should remain visible. A quarterly benefits committee can compare observed cohorts with the approved case, decide whether to release contingency and explain changes. The same evidence supports lender reporting and board oversight.

Implementation sequencing should protect the highest-consequence claims. A borrower may begin with document completeness, duplicate detection and queue prioritisation before automating more material recommendations. Early stages can improve data and workflow while generating evidence for later decisions. Expansion should follow the control environment and observed outcome, rather than a fixed date promised in the financing presentation.

The lender's technical adviser, actuarial adviser and financial adviser should use one issue register. Separate workstreams can otherwise apply different populations or valuation dates. The register should link each finding to the cash model, capital case, condition, covenant or disclosure. Closed items need evidence and an accountable approver.

Exit criteria are as important as expansion criteria. The borrower should know when to restrict a product, revert to human triage, change a vendor or retire the model. Funding documents should support orderly correction rather than create pressure to continue an unsafe workflow to preserve a ratio.

18 Apply lender red flags and decision rules

The first red flag is a benefit claim without a reconciled population. The second is a dashboard that reports averages without maturity, mix or exception detail. The third is a model trained with information that would not have been available at first notification. The fourth is a reserve or cash benefit that cannot be reconciled to controlled accounts.

Governance red flags include no named owner, no independent validation, unclear settlement authority, undocumented overrides, silent vendor change and no rollback. Conduct red flags include rising complaints, excluded vulnerable cohorts, unexplained outcome differences and weak appeal. Resilience red flags include no manual capacity, untested recovery, incomplete queue reconstruction and concentration in one provider.

Financial red flags include debt service sized to consolidated profit while cash is trapped, dividend forecasts without headroom, recoveries treated as cash before collection, reserve releases inside EBITDA, and recurring implementation cost labelled exceptional. A lender should also question benefit that appears only during a favourable claims period.

Decision rules should be simple. Include observed recurring cash after costs. Haircut benefits with limited maturity. Exclude timing movements from permanent savings. Stress distributions and capital separately. Require a fallback for material automation. Escalate material customer, data or regulatory failures. Preserve the ability to reduce exposure before liquidity is exhausted.

A weak causal claim should be replaced with a sensitivity. If the team cannot isolate how much severity improvement came from triage, it can model a range from zero to management's estimate and size debt to the supported case. If distribution timing is uncertain, it can test delay. If vendor exit cost is unknown, it can reserve liquidity. This approach keeps uncertainty visible and converts open questions into bounded financing consequences.

Documentation quality is itself a signal. Inconsistent KPI definitions, missing approval records and unreconciled presentations indicate that management may struggle to report after closing. A lender can require a pre-closing data remediation plan, independent validation or a shorter reporting cycle. Persistent inability to reproduce the case can justify lower leverage even when current results appear strong.

These rules do not eliminate judgement. They make judgement traceable. The credit memorandum should show which assumptions were evidenced, management-estimated or excluded and identify the resulting sizing change.

19 Recognise limitations and boundaries

Claims portfolios differ materially. A short-tail property portfolio, cyber book, professional-liability book and delegated specialty programme will not share one triage design or development pattern. The framework must be adapted to product wording, legal environment, distribution, claims authority, data maturity and reinsurance.

The worked case is hypothetical. Its amounts, ratios, thresholds, costs and scenarios are analytical assumptions rather than observed company results or forecasts. It does not establish that AI triage will reduce loss, expense or capital for any insurer. A borrower-specific conclusion requires current financial, actuarial, operational, legal, regulatory, tax and technology diligence.

Public disclosures provide market context but cannot substitute for private claim files, bordereaux, reinsurance, capital and governance evidence. Company reports apply their own definitions and reporting periods. Cross-company comparison requires reconciliation. Regulatory requirements also change and should be confirmed at the transaction date.

Causal attribution remains difficult. Mix, inflation, catastrophe, staffing, wording, repair markets and settlement strategy can change alongside the model. Longer-tail outcomes may not mature within the financing timetable. The lender should therefore use staged inclusion, sensitivity and reporting rather than a permanent conclusion from an early pilot.

The framework is not legal, regulatory, actuarial, accounting, tax, investment or credit advice. Decision makers should obtain qualified advice and approvals for the relevant entities and jurisdictions.

20 Conclude with a financeable evidence standard

AI claims triage becomes relevant to debt capacity when it produces controlled evidence through the full chain from notification to cash. Faster routing is a starting metric. Financeable value requires correct

routing, preserved customer outcomes, lower rework or leakage, reliable reserving, resilient service and transferable cash after capital constraints.

The legal-entity map prevents consolidated profit from being mistaken for debt-service cash. The event ledger makes populations and timestamps reproducible. Cohorts distinguish mix from performance. The champion-challenger design tests incremental value. Governance, conduct and resilience controls limit downside. The cash bridge and scenarios translate operating evidence into sizing. Covenants preserve visibility after closing.

The strongest credit case may recognise less benefit at opening than management expects. That conservatism creates room for evidence to earn recognition through performance. A staged facility can fund necessary investment while protecting liquidity and policyholder obligations. Transparent milestones can also improve dialogue among claims, actuarial, finance, risk, boards and lenders.

The decision standard is practical: include only benefit that can be traced, reconciled, governed and sustained. Treat model uncertainty, cash timing and transferability as distinct questions. Keep rollback and remediation funded. Under that discipline, claims automation can support a better financing decision without asking creditors to underwrite an unvalidated technology promise.

Frequently asked questions

Frequently asked questions

Q: Does faster claims triage automatically increase debt capacity? A: No. Faster routing supports debt capacity only when stable cohorts show better operating outcomes and the resulting net cash is available for debt service after implementation cost, claims timing, capital, liquidity and legal-entity constraints.

Q: Which claims metric matters most to a lender? A: No single metric is sufficient. Lenders need a connected set covering correct routing, cycle-time distribution, rework, reopening, complaints, reserve movement, paid loss, expense, recoveries, cash and model control.

Q: Can reserve releases be treated as recurring cash benefit? A: A reserve release is an accounting or actuarial movement and does not by itself create cash. The lender should reconcile it to ultimate-loss evidence, paid claims, capital effects and distributions before assigning credit value.

Q: How should an early pilot enter the financing case? A: The base case should include only observed and repeatable net benefit. Limited-maturity evidence can support a haircut, milestone or delayed-draw condition. Management-estimated benefit should remain outside the lender case until verified.

Q: What happens if the model fails after closing? A: The borrower should use a tested fallback, preserve decision logs, reconcile affected queues, protect customers and apply defined restriction, rollback and remediation actions. Financing terms can require notice and a cure plan for material failures.

Q: Should the lender receive claim-level personal data? A: Access should follow purpose, confidentiality, data-minimisation and applicable law. Aggregated cohorts, controlled samples and independent assurance may provide sufficient evidence without unrestricted personal data.

Q: How should reinsurance benefit be reflected? A: The model should distinguish recognised recovery from cash collected and consider notice, coverage, dispute, counterparty and timing. A recovery should receive a haircut when collectability or timing is uncertain.

Q: Which covenants are most relevant? A: Debt-service coverage, minimum liquidity, distribution and capital conditions, material-incident notice, information undertakings and model-performance triggers can be relevant. Definitions and thresholds require transaction-specific design.

References

- [1] Prudential Regulation Authority, Insurance Supervision: 2026 priorities. <https://www.bankofengland.co.uk/prudential-regulation/publication/2026/january/letter-from-gareth-truran-and-shoib-khan-insurance-supervision-2026-priorities>
- [2] Prudential Regulation Authority, Annual Report 2025/26. <https://www.bankofengland.co.uk/prudential-regulation/publication/2026/june/pru-annual-report-2025-26>
- [3] Bank of England and Financial Conduct Authority, DP5/22 Artificial Intelligence and Machine Learning. <https://www.bankofengland.co.uk/prudential-regulation/publication/2022/october/artificial-intelligence>
- [4] Prudential Regulation Authority, SS1/21 Operational resilience: Impact tolerances for important business services. <https://www.bankofengland.co.uk/prudential-regulation/publication/2021/march/operational-resilience-impact-tolerances-for-important-business-services-ss>
- [5] Prudential Regulation Authority, PS7/26 Operational incident and third-party reporting. <https://www.bankofengland.co.uk/prudential-regulation/publication/2026/march/operational-incident-and-third-party-reporting-policy-statement>
- [6] Financial Conduct Authority, ICOBS 8 Claims handling. <https://handbook.fca.org.uk/handbook/icobs8>
- [7] Financial Conduct Authority, FG22/5 Finalised guidance on the Consumer Duty. <https://www.fca.org.uk/publications/finalised-guidance/fg22-5-final-non-handbook-guidance-firms-consumer-duty>
- [8] Financial Conduct Authority, Motor insurance compensation following claims-process improvements. <https://www.fca.org.uk/news/press-releases/over-270000-motorists-receive-200m-motor-insurance-compensation-insurers>
- [9] Lloyd's, Full Year Results 2025. <https://www.lloyds.com/about-lloyds/investor-relations/financial-results/full-year-results-2025>
- [10] Lloyd's, Annual Report 2025. <https://assets.lloyds.com/media-651c0e64-c1d0-4f97-90f7-883c69fe2ef2/42102f78-bc6d-40be-8d43-52ef92775992/2025%20Lloyd%27s%20Annual%20Report.pdf>
- [11] Hiscox Group, Report and Accounts 2025. <https://www.hiscoxgroup.com/fyr2025/>
- [12] Beazley Management Limited, Syndicate 2623 Annual Report and Accounts 2025. https://assets.lloyds.com/media-651c0e64-c1d0-4f97-90f7-883c69fe2ef2/05e38ee5-acdb-452f-8f69-cb6d807a9170/2623_Accounts_2025Q4_20260220_190842.html
- [13] Lancashire Holdings Limited, Annual Report and Accounts 2025. <https://www.lancashiregroup.com/content/dam/lancashire/corporate2020/OAR/oar-2025/documents/Lancashire-Annual-Report-and-Accounts-2025.pdf>
- [14] IFRS Foundation, IFRS 17 Insurance Contracts. <https://www.ifrs.org/issued-standards/list-of-standards/ifrs-17-insurance-contracts/>
- [15] International Association of Insurance Supervisors, Application Paper on the supervision of artificial intelligence. <https://www.iaisweb.org/uploads/2024/11/Draft-Application-Paper-on-the-supervision-of-artificial-intelligence.pdf>
- [16] International Association of Insurance Supervisors, Regulation and supervision of artificial intelligence and machine learning in insurance: a thematic review. <https://www.iaisweb.org/uploads/2023/12/AI-ML-thematic-review-report.pdf>
- [17] Information Commissioner's Office, AI and data protection risk toolkit. <https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/artificial-intelligence/guidance-on-ai-and-data-protection/ai-and-data-protection-risk-toolkit/>
- [18] Bank of England and Financial Conduct Authority, FS2/23 Artificial Intelligence and Machine Learning. <https://www.bankofengland.co.uk/prudential-regulation/publication/2023/october/artificial-intelligence-and-machine-learning>
- [19] Prudential Regulation Authority, Insurance regulatory reporting. <https://www.bankofengland.co.uk/prudential-regulation/regulatory-reporting/insurance>
- [20] Prudential Regulation Authority, Insights from thematic review of general insurance reserving and capital modelling. <https://www.bankofengland.co.uk/prudential-regulation/letter/2023/june/insights-from-thematic-review>
- [21] Prudential Regulation Authority, SS2/21 Outsourcing and third party risk management. <https://www.bankofengland.co.uk/prudential-regulation/publication/2021/march/outsourcing-and-third-party-risk-management-ss>
- [22] Lloyd's, Principle 4: Claims Management. <https://www.lloyds.com/market-resources/market-oversight/principles-for-doing-business-at-lloyds/claims-management>

About the Author

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 closure.

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 and leads 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.

<https://www.linkedin.com/in/ckadya/>

<https://www.matchpoint-partners.com/team/ck-adya.html>

This paper is part of a continuing series on the structure of private and alternative markets. The views expressed are the author's own. The paper is for information only, describes market structure in general terms, and does not constitute investment, legal, tax or regulatory advice or a recommendation in respect of any security, vehicle or counterparty.