MatchPoint Research Working Paper No. 2026/T06

AI for Private Credit Underwriting: Faster, Sharper Risk Decisions

A controlled evidence, modelling and decision architecture for private lenders and investment teams

Chennakeshav (CK) Adya
Independent Researcher
July 2026

ABSTRACT

Background. Private-credit underwriting combines fragmented evidence, bespoke structures, judgement-based adjustments and material decision authority. AI can assist extraction, financial analysis, risk modelling, drafting and monitoring while creating provenance, model, confidentiality and authority risks. **Objective.** This paper develops a controlled architecture for applying AI across the private-credit underwriting workflow. **Approach.** The analysis reviews 18 official, supervisory and primary research sources. Market estimates, regulatory scope and empirical dataset boundaries are retained. **Findings.** The Federal Reserve estimated private-credit loans at about USD 1.4 trillion in the second half of 2025, equal to about 10% of US non-financial corporate debt under its stated definition. The FSB estimated USD 1.5 trillion to USD 2.0 trillion globally at end-2024 across contributing jurisdictions and identified data, borrower, valuation and interconnection vulnerabilities. The proposed credit evidence ledger links sources, facts, adjustments, assumptions, calculations, models, scenarios, terms and approval. Deterministic services calculate; validated models estimate; language models retrieve and draft; named people approve material decisions. **Implications.** A lender can begin with bounded extraction and calculation tasks, test complete underwriting packets in shadow mode and expand use only after evidence, fidelity, review effort, incidents and cost meet approved thresholds.

Highlights

- The credit evidence ledger is the operating record from source intake to committee decision.
- Deterministic calculations, statistical models and language models receive separate controls.
- Every adjustment, scenario, covenant and term remains linked to evidence and authority.
- Model performance is evaluated through discrimination, calibration, stability and intended use.
- Productivity claims require accepted packets, reviewer effort, cost and incident evidence.

JEL Classification: G21, G23, G24, G32, M15, O33

Keywords: artificial intelligence, private credit, direct lending, credit underwriting, machine learning, credit risk, financial spreading, covenants, explainability, model risk, human oversight

1. INTRODUCTION

Private credit underwriting combines speed, judgement and bespoke structure. A lender may move from an initial teaser to a binding term sheet through borrower financials, management evidence, sponsor materials, legal documents, market information, collateral records, downside cases and a negotiated covenant package. The information is frequently incomplete, differently defined and revised during the process. The credit decision therefore depends on both disciplined analysis and a visible record of how evidence, assumptions, calculations and judgement became an approved exposure.

The market context makes this control problem material. The Federal Reserve's May 2026 Financial Stability Report estimated that private credit loans represented about USD 1.4 trillion in the second half of 2025, equal to about 10% of total debt of US non-financial corporations and about one-third of below-investment-grade debt excluding bank loans. These figures use the report's stated market definition and source set [1]. The Financial Stability Board estimated global private-credit assets at USD 1.5 trillion to USD 2.0 trillion at end-2024 across contributing jurisdictions, while noting definitional and data gaps [2]. The estimates are not directly interchangeable.

Artificial intelligence can assist the repeated, evidence-heavy components of underwriting. Document models can classify and extract. Statistical or machine-learning models can identify patterns in historical outcomes. Language models can compare terms, retrieve evidence, draft cited analysis and organise exceptions. Deterministic code can spread accounts, calculate ratios, build cash-flow cases and test covenants. Human underwriters and investment committees retain responsibility for perimeter, adjustments, structure, price, exceptions and the final credit decision.

The central question is therefore not whether an AI system can produce a credit memo. The relevant question is whether an underwriting team can produce a more complete, timely and reviewable decision record while preserving source identity, calculation integrity, independent challenge and explicit authority.

This paper develops such an operating model for private-credit managers, direct lenders, business development companies, family-office credit teams and related investment functions. It addresses deal triage, document intake, financial spreading, quality-of-earnings adjustments, risk analysis, structure, covenants, pricing support, credit-memo preparation, committee review and the transition into monitoring.

The paper is a research synthesis and design framework. It does not report a Matchpoint production deployment, an observed reduction in underwriting time, an observed change in default or recovery, or a realised investment return.

2. DEFINITIONS, SCOPE AND METHOD

2.1 Private credit and underwriting scope

Private credit is used here for non-publicly traded debt provided primarily by non-bank lenders to businesses through bilateral or club-style negotiation. Strategies may include direct lending, unitranche, mezzanine, asset-based, special-situations and other forms of private debt. The Federal Reserve describes direct lending as bilateral origination by a single lender or small group and notes that many private-credit loans are senior secured and floating rate [3]. The FSB's May 2026 report uses a market-wide definition and records material variation across jurisdictions [2].

The underwriting scope begins when a potential opportunity enters the approved pipeline and ends when an authorised decision and its conditions are recorded. It includes:

- eligibility and mandate fit;
- borrower, sponsor and group identity;
- source registration and document control;
- historical financial spreading and normalisation;
- cash-flow, leverage, coverage and liquidity analysis;

- business, industry, management and sponsor assessment;
- collateral, guarantee and structural analysis;
- base, downside and break-case scenarios;
- covenant and documentation design;
- pricing and risk-return support;
- credit-memo preparation;
- independent review and investment-committee decision; and
- transfer of approved assumptions, terms and monitoring obligations into the portfolio record.

The final decision remains a human governance event. AI-assisted evidence, calculations and model outputs are inputs to that event.

2.2 Four technology classes

The phrase AI can hide materially different controls. This paper separates four classes.

Deterministic services apply fixed rules or formulas. Examples include accounting spreads, leverage ratios, debt-service coverage, covenant headroom, borrowing-base eligibility and waterfall calculations. A deterministic service should produce the same result from the same approved inputs and code version.

Statistical and machine-learning models estimate a quantity or class, such as default probability, loss severity, risk grade or anomaly score. These models require development evidence, validation, calibration, monitoring and a defined use boundary.

Language and document models classify text, extract proposed facts, retrieve passages, compare clauses and draft analysis. Their outputs require source-level verification and should not be treated as calculations or verified facts merely because the prose is fluent.

Workflow agents coordinate retrieval, tool calls, state transitions and review tasks. Their authority comes from policy, permissions and approval gates. An agent does not acquire credit authority through model capability.

This separation matters because the April 2026 interagency US model-risk guidance applies to traditional statistical and quantitative models and to non-generative, non-agentic AI models. It explicitly states that generative and agentic AI models are outside its scope, while noting that governance practices should determine appropriate controls for tools and systems outside that document [10]. The paper therefore uses that guidance only for covered model classes and as a comparative control reference. NIST, FSB and other AI-specific material informs the wider control architecture [13,14].

2.3 Evidence method

The analysis uses 18 sources grouped by purpose.

Market and financial-stability evidence. Federal Reserve, FSB, IMF and IOSCO sources describe market size, structure, borrower risk, valuation, liquidity, interconnection and data gaps [1-5].

Credit-process and valuation references. Basel Committee, OCC, EBA and FCA material describes sound credit-granting, data, independent review, exceptions, model use, monitoring and private-market valuation practices [6-9]. Most of this material applies to banks or regulated firms within the stated jurisdiction. It is used as a comparative design reference and is not represented as a legal obligation for every private-credit manager.

AI governance evidence. The revised US model-risk guidance, EBA machine-learning report, Bank of England and FCA survey, NIST profile, FSB consultation and CBUAE guidance inform model inventory, validation, data, explainability, human oversight, third-party and incident controls [10-15]. The FSB material is a consultation report and does not create a binding international standard [14]. The CBUAE guidance applies to licensed financial institutions and focuses on uses that may bear on consumers [15].

Primary research. BIS and peer-reviewed studies examine machine learning, transaction data, text and explainability in credit scoring [16-18]. Their datasets concern fintech, small-business or SME lending rather than institutional private-credit underwriting. The paper retains that boundary and does not extrapolate reported predictive performance into private-credit outcomes.

Evidence class	Permitted use	Boundary retained
Official market analysis	Market structure, scale, risk and data context	Definitions, dates, jurisdictions and underlying data differ
Supervisory or regulatory material	Control objectives and process design	Scope and legal applicability stated
Consultation	Emerging governance themes	Not binding and subject to revision
Empirical credit research	Capability and evaluation evidence	Dataset and lending segment retained
Matchpoint synthesis	Workflow, control, scorecard and roadmap	Requires organisation-specific validation
Matchpoint scenario	Economic calculation structure	Unverified management assumptions only

3. THE UNDERWRITING PROBLEM

3.1 A market with limited common data

Private credit is negotiated, illiquid and heterogeneous. The FSB identifies significant differences in market definition and limited granular fund and loan-level data across jurisdictions [2]. IOSCO similarly highlights limited transparency and valuation consistency across private finance [5]. These conditions constrain market-wide benchmarking and also appear inside individual underwriting files: private companies may have short reporting histories, uneven management information and borrower-specific definitions.

The Federal Reserve's 2024 FEDS Note studied approximately 17,000 unique US private-credit loans originated by 718 private-debt funds and business development companies during 2013-2023. The note reports that the average loan size exceeded USD 80 million from 2022 in that sample and that more than two-thirds of the loans were term loans [3]. The sample does not represent every strategy or jurisdiction. Its value for this paper is methodological: underwriting technology must preserve facility, tranche, lender, borrower, sponsor and time dimensions rather than flatten a bespoke deal into a generic score.

The May 2026 FSB report identifies borrower credit quality, valuation opacity, private ratings, sector concentration, leverage, liquidity features and interconnections as areas for continuing surveillance [2]. The Federal Reserve's May 2026 report also notes concerns about asset quality and serviceability among some private-credit borrowers [1]. An underwriting system needs to preserve the uncertainty and exception record behind these risks.

3.2 Financial statements are only one evidence layer

An underwriter may receive audited accounts, monthly management accounts, budgets, customer data, contracts, debt schedules, tax records, bank statements, quality-of-earnings work, legal diligence, insurance material, collateral appraisals and management representations. Each source has an owner, date, perimeter, currency, accounting basis and status.

The same label may carry different meanings. EBITDA may be statutory, management-reported, covenant-defined, sponsor-adjusted or lender-adjusted. Debt may include drawn facilities, leases, factoring, guarantees, earn-outs, deferred consideration, shareholder instruments or permitted exclusions. Cash may be unrestricted, trapped, pledged or operationally unavailable. Revenue may be reported, billed, contracted or collected.

An AI system that extracts a value without these attributes can create false precision. The operating model therefore treats every proposed fact as a source-bound record rather than a free-standing number.

3.3 The decision combines prediction and structure

Private-credit underwriting asks at least four different questions.

- **Capacity:** can the borrower service and repay the proposed debt under a specified operating and rate path?
- **Resilience:** what happens when revenue, margin, working capital, rates, currency, capex or exit conditions move adversely?
- **Protection:** how do collateral, guarantees, covenants, cash controls, information rights, amortisation and remedies change loss exposure and response time?
- **Economics:** does spread, fee, call protection, equity participation or other consideration compensate for the retained risk and portfolio use?

A default model addresses only part of this system. It does not select a covenant package, resolve an EBITDA adjustment, determine collateral enforceability or decide whether a sponsor's behaviour supports an exception. Technology must therefore support judgement without collapsing the decision into a single probability.

3.4 Valuation and monitoring begin at entry

Entry underwriting creates the baseline for future valuation and monitoring. The FCA's 2025 multi-firm review covered 36 firms with approximately GBP 3 trillion of global private assets under management in aggregate, including private debt and other private-asset strategies. The FCA identified independence, expertise, transparency, consistency, conflicts, backtesting and defined ad hoc valuation processes as important features [6]. The sample and findings cover private markets broadly.

For a lender, the approved underwriting record should therefore retain the initial enterprise-value case, coverage tests, yield evidence, assumptions, adjustment rationales and triggers for reassessment. A memo that cannot be decomposed into its source and assumption records weakens later valuation challenge.

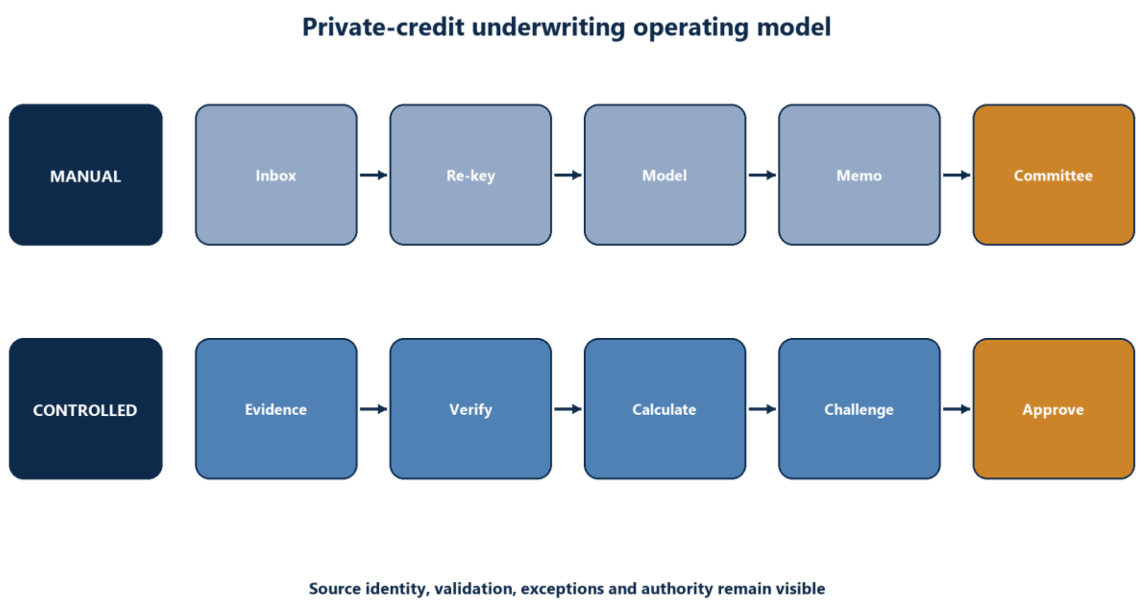

Figure 1. Manual underwriting hand-offs compared with a controlled evidence-to-decision workflow
Source: Matchpoint framework.

4. THE CREDIT EVIDENCE LEDGER

4.1 Record types

The proposed operating unit is a credit evidence ledger. It connects the following record types.

Record	Minimum fields	Control purpose
Entity	Legal name, identifier, jurisdiction, role, ownership	Prevent borrower, guarantor and sponsor confusion
Source	File or system ID, owner, date, version, hash, permission	Establish provenance and currentness
Fact	Value, unit, period, perimeter, source location, extraction status	Make proposed facts reviewable
Adjustment	Base value, amount, rationale, evidence, owner, approval	Separate reported and adjusted measures
Assumption	Variable, value, range, basis, scenario, approver	Expose model judgement
Calculation	Formula or code version, inputs, output, validation	Make ratios reproducible
Model output	Model, version, intended use, score, explanation, limits	Preserve model lineage
Exception	Rule, observed breach, severity, owner, resolution	Prevent silent policy drift
Term	Facility, amount, price, maturity, covenant, condition	Connect analysis to structure
Decision	Forum, date, vote, conditions, dissent, authority	Record accountable approval
Monitoring trigger	Metric, threshold, frequency, source, escalation	Carry underwriting into portfolio management

The ledger does not require a single software product. It requires stable identifiers and relationships across the document store, spreading system, calculation service, model registry, deal-management platform and approval record.

4.2 Source identity and cut-off

Every source should be registered before it supports a material conclusion. Registration records the file hash or system version, reporting period, receipt date, owner, confidentiality class, permission and supersession status. A cut-off timestamp defines the evidence population used for the decision.

Late-arriving evidence creates a controlled change event. The system identifies affected facts, adjustments, calculations, scenarios, paragraphs and terms. The underwriter decides whether to reopen analysis or document immateriality. This approach supports forward traceability from a changed source to every dependent decision component.

4.3 Proposed facts and verified facts

Document AI should create proposed facts. A proposed fact records the extracted value and exact source location, together with unit, sign, period and perimeter. Verification occurs through one or more of:

- exact agreement with a structured source;
- deterministic cross-foot and reconciliation;
- agreement between independent documents;
- reviewer confirmation; or
- documented acceptance of an unresolved difference.

The status is retained. Generated narrative must not convert an unverified proposed fact into a verified fact.

4.4 Adjustment lineage

Adjustment analysis often carries more risk than extraction. The ledger distinguishes reported, sponsor-adjusted, third-party-adjusted and lender-adjusted measures. Each adjustment connects to evidence, recurrence, cash effect, accounting treatment, period, owner and approval.

For example, a cost saving may be contracted, initiated, partly realised or proposed. An acquisition contribution may be historical, run-rate or forecast. A one-off expense may be genuinely non-recurring or a regular feature described differently each year. The system can collect comparable adjustments and flag repeated classifications. The underwriter retains the conclusion.

4.5 Decision traceability

Backward traceability begins with an approved conclusion and returns to the term, scenario, calculation, assumption, fact and source that support it. Forward traceability begins with a new source or changed assumption and identifies every affected output.

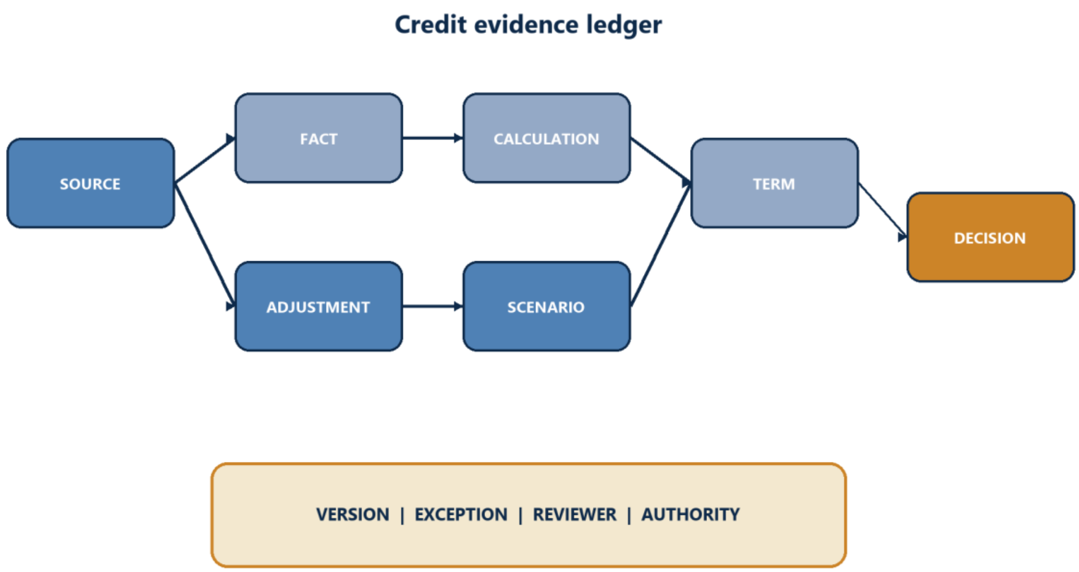

Figure 2. Credit evidence graph connecting source, fact, adjustment, calculation, scenario, term and decision
Source: Matchpoint framework; traceability objectives informed by [7-10].

5. CONTROLLED UNDERWRITING ARCHITECTURE

5.1 Layered design

The architecture separates data, deterministic analysis, statistical models, language models and authority.

Layer	Function	Release condition
Source and identity	Register entities, files, systems, permissions and versions	Approved source perimeter
Evidence services	Classify, extract, cite, reconcile and surface contradictions	Proposed facts remain labelled
Deterministic finance	Spread, calculate, scenario-test and covenant-test	Independent recomputation passes

Layer	Function	Release condition
Statistical models	Estimate risk or anomaly within intended use	Validation and monitoring thresholds pass
Language layer	Retrieve, compare, explain and draft	Citation and claim checks pass
Workflow policy	Route tasks, exceptions, budgets and stop conditions	Policy state is valid
Human authority	Approve adjustments, exceptions, terms and decision	Named authority recorded
Audit and monitoring	Retain versions, tests, overrides, incidents and outcomes	Record complete and immutable

This design makes the official ledger, calculation and decision systems authoritative. The language model is an interface and synthesis layer over approved evidence. It is not the book of record.

5.2 Typed tools and least privilege

An underwriting agent should use narrow tools with explicit input and output schemas. Examples include `register_source`, `extract_table`, `reconcile_statement`, `calculate_leverage`, `run_scenario`, `test_covenant`, `retrieve_clause`, `create_exception` and `prepare_memo_section`.

Each tool call carries deal identity, user authority, permitted sources, calculation version and purpose. Read access is separated from write access. Draft creation is separated from approval. External transmission remains behind a recipient and authority check.

5.3 Transaction state machine

The deal moves through explicit states such as intake, source registration, preliminary screen, full underwriting, independent review, committee ready, approved with conditions, declined, withdrawn and monitoring setup. Required records and permissions vary by state.

A state transition should be blocked when mandatory evidence, validations or approvals are missing. The block produces an exception rather than a silent workaround.

5.4 Model inventory

The inventory records every statistical model, document model, language model, embedding service, rules engine and externally supplied score used in underwriting. It includes owner, provider, version, intended use, prohibited use, input data, validation status, limitations, monitoring, fallback and downstream dependencies.

The revised US interagency guidance describes a risk-based approach for covered quantitative and non-generative models, including intended use, testing, validation, monitoring, inventory, documentation and vendor oversight [10]. Generative and agentic systems remain outside that guidance's formal scope. The wider inventory is therefore a Matchpoint control proposal informed by NIST and FSB AI governance material [13,14].

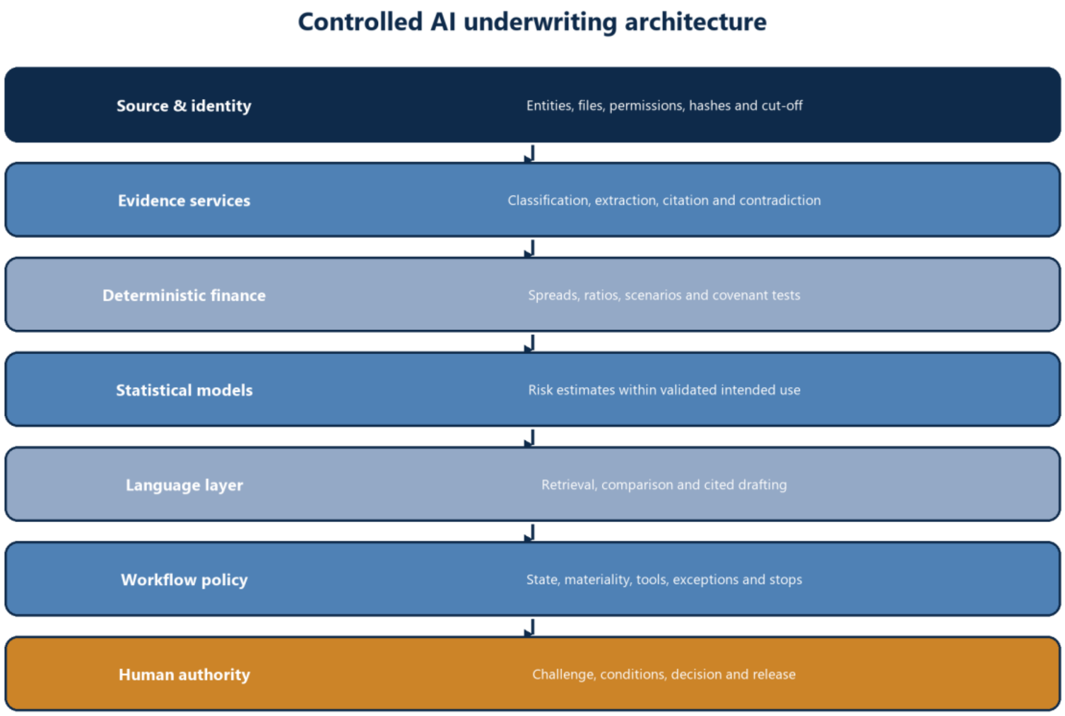

Figure 3. Layered AI underwriting architecture with separate evidence, calculation, model, language and authority planes

Source: Matchpoint architecture; model and AI control references [10,13-15].

6. AI ACROSS THE UNDERWRITING WORKFLOW

6.1 Intake and mandate fit

The intake service structures the opportunity against approved strategy, geography, sector, ticket, instrument, return, concentration and exclusion rules. It resolves the borrower, sponsor and group and checks existing exposure and conflicts.

AI can classify materials and propose missing-information requests. A deterministic policy service tests explicit limits. The investment team decides whether the opportunity merits resources and whether any exception proceeds.

The output is an intake packet containing the source list, entity map, eligibility results, exceptions and named decision.

6.2 Document and data-room intelligence

Document AI can inventory folders, classify documents, identify versions, locate requested items and extract proposed fields. Language models can compare representations across teasers, lender presentations, diligence reports and legal drafts. They can surface conflicting dates, definitions and amounts.

The control objective is completeness and contradiction visibility. The model must cite the exact source location and preserve the distinction between missing, unreadable, contradictory and not applicable.

Prompt injection is relevant when untrusted documents can influence a language model. Document text should be treated as evidence rather than instruction. Tool permissions, retrieval boundaries and system policies should remain outside document control.

6.3 Financial spreading and normalisation

Extraction proposes the financial statement line, period, currency, sign, accounting basis and source. Deterministic code cross-foots statements, tests opening and closing balances, reconciles cash flow, checks units and records unresolved differences.

Mapping logic should preserve both source labels and approved underwriting taxonomy. Manual reclassifications require rationale and reviewer identity. Consolidated, divisional, legal-entity and covenant perimeters remain distinct.

AI may assist account mapping and anomaly detection. It should not silently plug differences or infer missing values. Where a value is estimated, the record begins with an explicit estimate label and an approved basis.

6.4 Earnings, cash conversion and working capital

The system can group proposed adjustments, detect repetition across periods and compare management, sponsor and third-party positions. It can calculate conversion measures under multiple definitions and trace working-capital seasonality.

Judgement remains central. Recurrence, control, timing, cash effect and downside treatment require an underwriter. The system should show how each adjustment changes leverage, coverage and covenant headroom.

6.5 Business and industry risk

AI-assisted retrieval can organise customer concentration, contract terms, retention, price-volume-mix, supplier dependencies, regulatory exposures and market evidence. The output should separate borrower-provided evidence, independent public evidence and analyst judgement.

External data requires provenance and an approved licensing basis. Time-sensitive market evidence carries an as-of date. A retrieval result without stable source identity does not support a material conclusion.

6.6 Structure, collateral and guarantees

The structure module records the complete capital stack, priority, maturity, amortisation, cash-pay and payment-in-kind terms, security, guarantees, intercreditor rights, permitted debt and leakage. It connects collateral values to valuation date, appraiser, method, haircut, jurisdiction and enforceability review.

AI can compare term sheets and documents, extract proposed clauses and flag deviations from an approved playbook. Legal conclusions and enforceability remain with qualified counsel and the authorised lender.

6.7 Scenarios, covenants and pricing support

Scenario models should run through deterministic formulas using approved assumptions. The base, downside and break case identify the driver, shock path, mitigation, liquidity requirement, covenant effect and repayment consequence.

The covenant engine should test historical and projected periods under the negotiated definition. It records baskets, cure rights, holidays, testing dates and reporting obligations. Language models can explain results and compare draft clauses; the calculation service determines headroom.

Pricing support combines exposure, loss analysis, liquidity, capital use, concentration, optionality, execution cost and fund economics. A model output informs the discussion. The authorised team approves the price and structure.

6.8 Credit memorandum and committee pack

The drafting service assembles only approved facts, calculations, scenarios and cited evidence. Each paragraph has a claim record. Material statements carry source or analysis lineage. Open exceptions and unresolved contradictions appear in a visible section.

The committee pack includes:

- decision requested;
- exposure and structure;
- evidence cut-off;
- risk thesis and mitigants;
- financial and scenario outputs;
- adjustments and assumptions;
- policy and model exceptions;
- proposed covenants and monitoring;
- open conditions;
- independent-review findings; and
- named recommendation and dissent.

The committee may approve, approve with conditions, defer, reject or request further work. The decision record is not generated by the model.

6.9 Transition to monitoring

Approved terms, covenant definitions, reporting dates, information rights, base-case assumptions and early-warning indicators transfer into monitoring without re-keying. The monitoring system retains the underwriting baseline and measures change.

AI can classify new borrower reports, update proposed facts and surface deviations. Deterministic services recalculate approved metrics. Material changes route to the portfolio manager, independent credit function or committee according to policy.

7. MODEL AND OUTPUT EVALUATION

7.1 Evaluation by component

One blended accuracy score is insufficient. Each component needs a metric aligned to its failure mode.

Component	Primary measures	Release evidence
Document classification	Precision, recall, abstention, version errors	Stratified file set
Field extraction	Exact match, unit accuracy, period accuracy, citation validity	Human-labelled fields
Financial spreading	Cross-foot rate, reconciliation rate, mapping exceptions	Deterministic recomputation
Statistical risk model	Discrimination, calibration, stability, out-of-time performance	Independent validation
Language analysis	Claim support, contradiction recall, unsupported-claim rate	Source-linked review set
Covenant engine	Clause match, formula fidelity, date and basket tests	Legal and numerical cases
Memo drafting	Accepted paragraphs, material omission, review time, rework	Blind reviewer comparison
Workflow agent	Tool precision, state compliance, stop accuracy, authority violations	End-to-end packet tests

7.2 Historical splits and leakage

Credit models should be tested out of time and, where relevant, across borrower cohorts, sectors, sponsor status, facility type and macro conditions. Random splits can allow information from the same borrower or period to appear in both development and test sets. The evaluation design records the unit of observation and prevents that leakage.

Rejected or unobserved loans create selection constraints. Defaults and recoveries may be sparse, delayed or differently defined. The model card should state coverage and uncertainty rather than imply precision beyond the data.

7.3 Discrimination and calibration

Discrimination tests whether the model ranks risk. Calibration tests whether predicted probabilities align with observed rates over the stated horizon and definition. A model can rank well and calibrate poorly. Both matter when model outputs affect structure, price or limits.

The evaluation should include Brier score or another proper scoring rule, calibration curves, area under the receiver operating characteristic curve, precision-recall measures where defaults are rare, stability, segment results and threshold consequences. The selected metrics depend on intended use.

7.4 Explainability and monotonicity

Wang and co-authors compared a scorecard with several machine-learning models using tens of thousands of small-business loans and millions of transactions. Their XGBoost model with monotonic constraints improved the K-S statistic by 7% relative to the scorecard in that specific dataset and was deployed with a SHAP-based explanation approach [17]. This is small-business application evidence; it is not a private-credit result.

Explanation methods can show feature contribution without proving causality or model validity. They should accompany conceptual review, data assessment, outcome analysis and challenge. Monotonic constraints can align selected relationships with policy, while interactions and data drift remain subject to testing.

7.5 Text evidence

Stevenson, Mues and Bravo studied 60,000 textual loan-officer assessments from a specialised micro and SME lender. Text alone produced reported AUC values between 0.55 and 0.71 across their tests, while combining text with traditional data did not produce an aggregate performance lift [18]. The result supports direct evaluation of incremental value. It does not justify using narrative as an unverified substitute for financial evidence.

7.6 Stress and domain shift

The BIS study by Gambacorta and co-authors used proprietary transaction-level data from a Chinese fintech firm and found that its machine-learning and non-traditional-data model performed better than traditional models in predicting losses and defaults during the analysed negative credit-supply shock. The advantage declined for borrowers with longer credit histories [16]. The consumer and fintech environment differs from institutional private credit.

An underwriting model should be tested under interest-rate, margin, revenue, working-capital, collateral and refinancing stress. Domain-shift indicators include changes in deal source, borrower size, sector, structure, sponsor behaviour, accounting quality, covenants and macro conditions.

7.7 Override and outcome analysis

Every model output and human override should be retained. Override analysis asks whether overrides cluster by team, sector, sponsor or direction and whether later outcomes support the change. It should not create a mechanical rule that discourages appropriate judgement.

The April 2026 US interagency guidance emphasises conceptual soundness, outcome analysis, ongoing monitoring and effective challenge for covered models [10]. The Basel Committee's 2025 credit-risk principles also emphasise robust data, periodic validation, independent review and complete credit-life-cycle controls for banks [7]. Private-credit managers can adopt analogous internal disciplines without representing themselves as subject to bank-specific guidance.

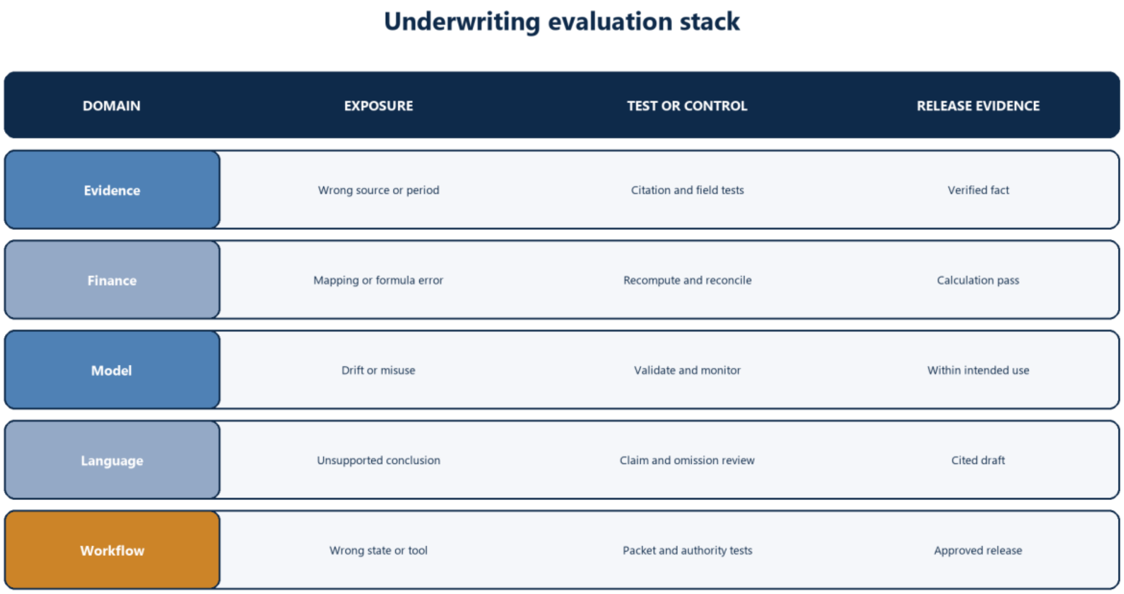

Figure 4. Evaluation stack from source and calculation fidelity through model, memo and workflow outcomes
Source: Matchpoint framework; empirical and model-evaluation evidence [10,16-18].

8. DECISION RIGHTS AND INDEPENDENT CHALLENGE

8.1 Authority matrix

Activity	AI or system	Deal team	Independent credit or risk	Investment committee
Register source and extract proposed fact	Prepare	Verify	Sample or challenge	Informed
Spread accounts and calculate ratios	Execute controlled service	Review exceptions	Validate method and sample	Informed
Propose adjustment	Organise evidence	Recommend	Challenge	Approve if material
Produce model output	Execute approved model	Interpret	Validate and monitor	Consider within boundary
Draft credit analysis	Prepare cited draft	Own conclusion	Challenge material risks	Review
Propose structure and covenants	Compare and test	Recommend	Challenge	Approve
Approve policy exception	Record and route	Request	Recommend or challenge	Approve under policy
Release term sheet or commitment	Prepare	Confirm	Confirm conditions as required	Named authority releases

The exact allocation depends on fund governance. The invariant is that system capability does not change delegated authority.

8.2 Independent review

Independent review focuses on evidence quality, policy compliance, model use, assumptions, adjustments, downside severity, structure, concentration and conditions. It can reproduce key calculations from the ledger and inspect unresolved exceptions.

The OCC's June 2026 handbook discusses underwriting standards, exception and override reporting, independent credit review and model-risk management within its banking scope [8]. The EBA loan-origination guidelines address governance, creditworthiness, pricing, collateral and monitoring for covered institutions [9]. These sources support a general design principle: origination should not be the sole source of challenge.

8.3 Abstention and stop conditions

The system should abstain when evidence is missing, the source is unreadable, entities conflict, the calculation does not reconcile, the request exceeds intended use, the model is outside its population, the data cut-off is stale or a tool would exceed authority.

Abstention produces an exception with an owner. It should not produce a plausible completion.

9. RISK AND CONTROL FRAMEWORK

9.1 Evidence and calculation risk

Material facts require source and version. Numerical outputs require deterministic recomputation where a formula can be specified. Unit, sign, currency, period and perimeter tests run before downstream use. Unsupported claims remain blocked from the committee pack.

9.2 Model risk

Covered statistical models receive intended-use statements, development records, independent validation, limits, monitoring and inventory. Generative and agentic components receive task-specific evaluations, source tests, tool tests, authority tests, incident monitoring and change control under the wider AI framework.

9.3 Confidentiality and data rights

Private-credit files may contain confidential borrower, sponsor, employee, customer and counterparty information. The architecture applies classification, purpose limitation, least privilege, encryption, retention, jurisdiction and vendor-use controls. Training or provider retention of deal information requires an explicit approved basis.

The CBUAE guidance discusses transparency, bias, accountability, explainability, data privacy, human oversight and immediate intervention for AI and machine learning used by licensed financial institutions within its scope [15]. A private fund outside that scope still needs its own legal and contractual analysis.

9.4 Third-party risk

The provider inventory includes document processors, model vendors, data providers, cloud platforms and valuation services. It records ownership, sub-processors, data location, model change, service levels, testing evidence, concentration, continuity and exit.

The Bank of England and FCA 2024 survey reported that one third of AI use cases were third-party implementations and that the three most frequently named model providers represented 44% of named model providers [12]. These are respondent-reported financial-sector figures rather than private-credit-only measures.

9.5 Bias, fairness and explainability

Institutional corporate lending differs from consumer credit. Applicable discrimination, data-protection and explanation requirements depend on jurisdiction, borrower, guarantor and use. The inventory identifies when personal data or individual decisions enter the workflow. Legal analysis defines the applicable obligations.

The EBA's 2023 follow-up report on machine learning in internal-ratings-based models discusses prudent use and interactions with GDPR and the EU AI framework within its stated banking context [11]. Explainability is treated as one part of governance rather than a substitute for validation.

9.6 Security and prompt injection

Untrusted documents cannot alter system policy or tool authority. The platform separates instructions from evidence, sanitises content, restricts tool calls, validates output schemas and requires approval for material writes or external release. Security tests include malicious documents, hidden text, contradictory instructions, oversized inputs and data-exfiltration attempts.

NIST AI 600-1 supplies voluntary generative-AI risk-management considerations, including confabulation, information integrity, privacy, security, human-AI configuration and third-party risk [13]. The paper applies these as control references rather than outcome guarantees.

9.7 Records and incidents

The record includes source versions, prompts or task specifications where appropriate, model and tool versions, outputs, validations, reviewer actions, overrides, approvals and final release. Incident classes include unsupported material claim, wrong entity, wrong calculation, unauthorised source, data leakage, improper tool action, missed stop, incorrect recipient and monitoring failure.

Private-credit AI risk and control map

DOMAIN	EXPOSURE	TEST OR CONTROL	RELEASE EVIDENCE
Evidence	Missing or stale	Manifest and cut-off	Current source set
Numbers	Unit or formula error	Deterministic service	Recomputed result
Models	Bias or deterioration	Validation and drift	Approved use
Data	Leakage or reuse	Classification and rights	Permitted processing
Authority	Wrong term or recipient	Named approval gate	Release record

Figure 5. Underwriting risk-control map across evidence, calculation, models, data, security, third parties and authority

Source: Matchpoint synthesis based on [7-15].

10. PRODUCTIVITY AND ECONOMIC CASE

10.1 Measurement unit

The economic unit is an accepted underwriting packet or completed stage, not a generated document. A packet is accepted when required evidence, calculations, exceptions, review and authority records pass the defined gate.

The baseline should measure:

- analyst active time;
- senior and independent-review time;
- total elapsed time;
- sources and pages reviewed;
- rework and exception cycles;
- calculation and citation defects;
- late information and version changes;
- external data and adviser cost;
- opportunities declined before full underwriting;
- committee deferrals; and
- post-close correction or monitoring setup failures.

Released time has economic value only when redeployed to accepted throughput, deeper coverage, higher-value judgement or lower cost.

10.2 Value equation

Let:

- (N) be attempted packets per year;
- (a) be the accepted-output rate;
- (H_b) be baseline active hours per accepted packet;
- (H_p) be post-pilot active hours per accepted packet;
- (c_h) be loaded hourly cost;
- (C_t) be annual model, data and tool cost;
- (C_f) be annualised build, validation and control cost;
- (V_q) be observed quality or risk value that finance approves; and
- (C_i) be observed incident and remediation cost.

The illustrative structure is:

Annual operating value = $N \times a \times (H_b - H_p) \times c_h + V_q - C_t - C_f - C_i$

Every input requires a defined source and observation period. (V_q) should remain zero unless the organisation approves a measured method.

10.3 Illustrative scenario

The following scenario is unverified and uses illustrative management assumptions solely to show the calculation structure.

Input	Bounded pilot	Mid-size platform	Scaled platform	Status
Attempted underwriting packets per year	120	240	480	Illustrative management assumption
Accepted-output rate	70%	82%	90%	Illustrative management assumption
Baseline active hours per accepted packet	90	120	150	Illustrative management assumption
Post-pilot active hours per accepted packet	72	84	102	Illustrative management assumption
Loaded blended hourly cost	USD 120	USD 160	USD 200	Illustrative management assumption
Model, data and tool cost per attempted packet	USD 250	USD 400	USD 650	Illustrative management assumption
Annualised build, validation and control cost	USD 250,000	USD 450,000	USD 750,000	Illustrative management assumption
Approved quality value	USD 0	USD 0	USD 0	No observed basis assumed

The table does not establish a business case. It omits any forecast of lower defaults, higher recoveries, improved pricing or increased investment returns. Those outcomes require a longer performance window, comparable portfolios and a defensible attribution method.

10.4 Productivity tree

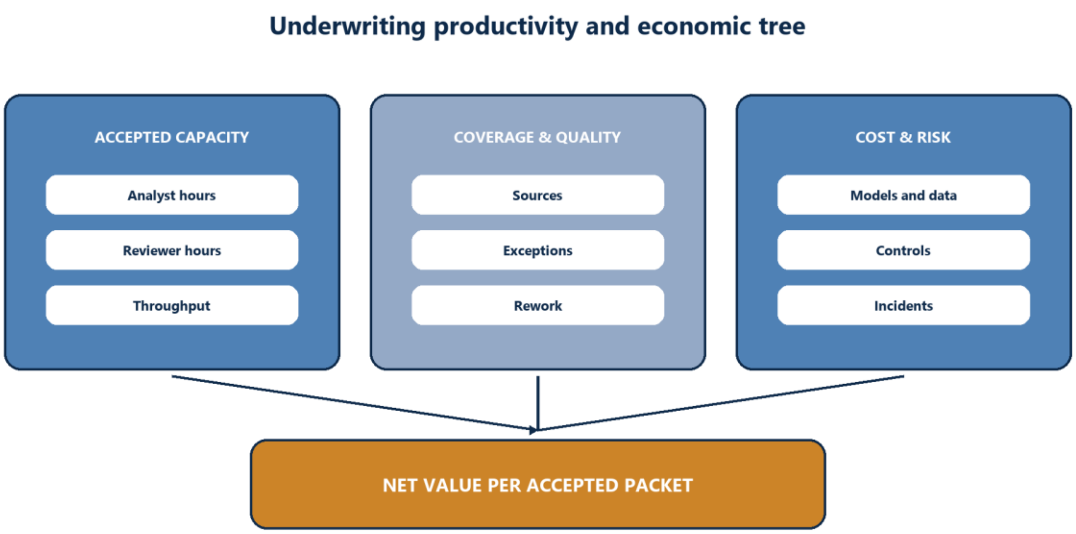

Figure 6. Productivity and economic tree for accepted capacity, coverage, quality, cost and incidents
Source: Matchpoint scenario analysis. Operating inputs require observed baseline and pilot evidence.

11. IMPLEMENTATION ROADMAP

11.1 Stage 1: mandate and inventory

Define strategy, borrower and instrument scope, underwriting stages, materiality, authority, data classes, jurisdictions, models, vendors and current systems. Select a bounded use case with sufficient historical examples.

Exit gate: approved charter, owner, source perimeter, risk classification and baseline plan.

11.2 Stage 2: evidence and deterministic foundation

Create entity and source identifiers, approved financial taxonomy, adjustment records, calculation services, scenario definitions, covenant definitions and exception types. Build a labelled evaluation set.

Exit gate: source-to-output traceability and deterministic recomputation pass on the evaluation set.

11.3 Stage 3: extraction and retrieval pilot

Run document classification, proposed fact extraction, citation, contradiction and retrieval over historical cases. Measure exact-match, unit, period, perimeter, citation and abstention results.

Exit gate: task thresholds pass for the intended document population; unresolved classes remain excluded.

11.4 Stage 4: model and memo shadowing

Run approved statistical models and cited memo drafting in parallel with the current process. Reviewers remain blind to origin where practical. Compare accepted output, omissions, rework, review time and model performance.

Exit gate: quality is non-inferior under the approved test and reviewer effort is measured.

11.5 Stage 5: controlled production

Use the system for selected deals with approval-bound outputs. Monitor evidence defects, calculation failures, unsupported claims, exceptions, overrides, incidents, cost and elapsed time.

Exit gate: the control committee approves continued use based on observed task and workflow evidence.

11.6 Stage 6: portfolio integration

Transfer the approved underwriting ledger into monitoring, valuation and amendment workflows. Test change impact, covenant surveillance, rating review and event escalation.

Exit gate: underwriting assumptions and approved terms remain traceable through the credit life cycle.

11.7 Stage 7: expansion

Expand to new strategies, sectors or jurisdictions only after new validation. Model and policy changes trigger regression testing. Provider or data changes receive controlled review.

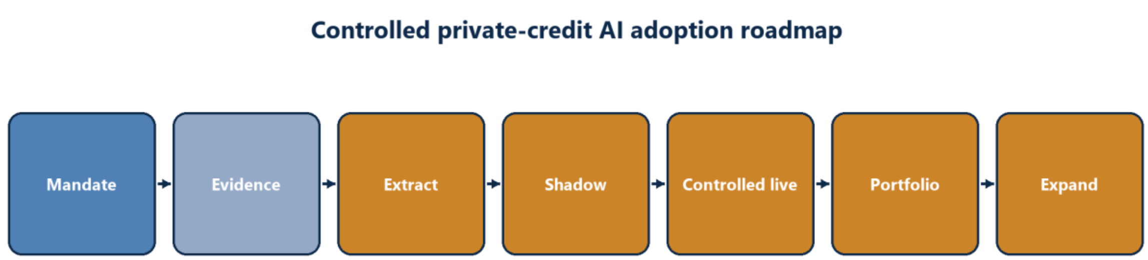

Figure 7. Controlled adoption roadmap from mandate and evidence foundation through portfolio integration and expansion

Source: Matchpoint framework.

12. APPLICATION TO PRIVATE-CREDIT MANAGERS AND FAMILY OFFICES

12.1 Private-credit manager or direct lender

An A3 manager can use the framework to standardise evidence intake across origination teams while retaining strategy-specific scorecards and committee authority. The strongest initial modules are likely to be source registration, financial spreading, adjustment lineage, covenant extraction, cited memo preparation and monitoring setup.

The operating benefit should be assessed through accepted throughput, coverage, review effort and exception timing. A larger number of memos is not a sufficient outcome.

12.2 Family-office investment team

An A2 family office may encounter private credit through direct deals, club transactions, co-investments or manager allocations. For direct deals, the framework supports underwriting and concentration analysis. For fund selection, it creates diligence questions about the manager's evidence, models, overrides, valuation, monitoring and incident controls.

The family office can ask a manager to demonstrate:

- how reported and adjusted EBITDA are separated;
- how model outputs influence decisions;
- who validates models and challenges underwriting;
- how exceptions and overrides are tracked;
- how covenants transfer into monitoring;
- how private valuations are backtested;
- how confidential data is handled by AI vendors; and
- what authority any agent or automated workflow possesses.

12.3 Shared control opportunity

A manager and allocator can align around a common evidence vocabulary without sharing confidential borrower data. Portfolio reporting can expose approved metrics, definitions, exceptions and change records. The underlying deal ledger remains permissioned.

13. LIMITATIONS AND RESEARCH AGENDA

First, public private-credit data remains incomplete and definitions vary. Market-size estimates in the Federal Reserve, FSB and IMF sources use different coverage and dates [1,2,4].

Second, empirical AI studies cited here do not evaluate institutional private-credit underwriting. They cover fintech or small-business lending and use distinct borrower populations, horizons and outcome definitions [16-18].

Third, default and recovery outcomes occur with delay and may be affected by amendments, payment-in-kind interest, restructurings and lender intervention. A short pilot cannot establish improved credit performance.

Fourth, a model that performs on historical data may fail under a new sector, sponsor mix, structure, accounting environment or macro regime. Out-of-time and domain-specific validation remains necessary.

Fifth, bank guidance does not automatically apply to private funds. Each organisation needs legal and regulatory analysis by jurisdiction, vehicle, investor, borrower and activity.

Sixth, the numerical scenario is unverified and uses illustrative management assumptions. It supports calculation design only.

Future research should build a permissioned private-credit benchmark containing realistic deal documents, version changes, accounting perimeters, adjustments, scenarios, structures, covenants and committee outcomes. The benchmark should test extraction, lineage, calculation, calibration, contradiction, abstention, override, monitoring and long-horizon state. Multi-year field studies should compare accepted throughput, reviewer time, risk selection, amendments, defaults, recoveries and incidents with appropriate controls for portfolio composition.

14. CONCLUSION

AI can increase the analytical capacity of a private-credit team when the operating model preserves the structure of the credit decision. Sources remain identifiable. Proposed facts remain distinguishable from verified facts. Financial calculations remain reproducible. Statistical models remain within validated use. Language-model output remains cited. Exceptions remain visible. Material terms and commitments remain subject to named human authority.

The credit evidence ledger connects these disciplines. It links each conclusion to the borrower, source, fact, adjustment, assumption, calculation, model, scenario, term, reviewer and decision that support it. It also carries the approved underwriting baseline into monitoring and valuation.

The implementation sequence begins with a bounded task and a baseline. It tests component fidelity, complete packets, reviewer effort, end-to-end state and authority. Expansion follows observed evidence. This creates a practical route to faster and sharper analysis while keeping underwriting rigour measurable.

DECLARATIONS

Author. Chennakeshav (CK) Adya.

Affiliation. Independent Researcher.

Conflict of interest. The author is the Founder and Managing Partner of Matchpoint Partners. The paper describes a Matchpoint framework and does not report a client engagement or product performance.

Funding. No external funding was received for this paper.

Data availability. The paper uses public sources listed below. The numerical economic scenario is generated for framework demonstration and is not observed firm data.

Use of AI. AI tools assisted research organisation, drafting, document production and quality checks under author direction. The author retains responsibility for the paper.

REFERENCES

[1] Board of Governors of the Federal Reserve System. (2026). *Financial Stability Report: May 2026*. https://www.federalreserve.gov/publications/files/financial-stability-report-20260508.pdf

[2] Financial Stability Board. (2026). *Report on Vulnerabilities in Private Credit*. 6 May 2026. https://www.fsb.org/uploads/P060526.pdf

[3] Cai, F., and Haque, S. (2024). Private credit: Characteristics and risks. *FEDS Notes*. Board of Governors of the Federal Reserve System. https://doi.org/10.17016/2380-7172.3462

[4] International Monetary Fund. (2024). The rise and risks of private credit. In *Global Financial Stability Report: The Last Mile: Financial Vulnerabilities and Risks*, Chapter 2. https://www.imf.org/en/Publications/GFSR/Issues/2024/04/16/global-financial-stability-report-april-2024

[5] International Organization of Securities Commissions. (2023). *Thematic Analysis: Emerging Risks in Private Finance*. IOSCO FR10/23. https://www.iosco.org/library/pubdocs/pdf/IOSCOPD745.pdf

[6] Financial Conduct Authority. (2025). *Private market valuation practices*. Multi-firm review, 5 March 2025. https://www.fca.org.uk/publications/multi-firm-reviews/private-market-valuation-practices

[7] Basel Committee on Banking Supervision. (2025). *Principles for the Management of Credit Risk*. https://www.bis.org/bcbs/publ/d595.pdf

[8] Office of the Comptroller of the Currency. (2026). *Lending and Loan Portfolio Risk Management*. Comptroller's Handbook, Version 1.0, June 2026. https://www.occ.treas.gov/publications-and-resources/publications/comptrollers-handbook/files/lending-loan-portfolio-risk-management/pub-ch-lending-loan-portfolio.pdf

[9] European Banking Authority. (2020). *Guidelines on loan origination and monitoring*. EBA/GL/2020/06. https://www.eba.europa.eu/activities/single-rulebook/regulatory-activities/credit-risk/guidelines-loan-origination-and-monitoring

[10] Board of Governors of the Federal Reserve System, Federal Deposit Insurance Corporation, and Office of the Comptroller of the Currency. (2026). *Supervisory Guidance on Model Risk Management*. SR 26-2 attachment. https://www.federalreserve.gov/supervisionreg/srletters/SR2602a1.pdf

[11] European Banking Authority. (2023). *Follow-up report on the use of machine learning for internal ratings-based models*. https://www.eba.europa.eu/publications-and-media/press-releases/eba-publishes-follow-report-use-machine-learning-internal

[12] Bank of England and Financial Conduct Authority. (2024). *Artificial intelligence in UK financial services: 2024*. Survey of 118 responding firms. https://www.bankofengland.co.uk/report/2024/artificial-intelligence-in-uk-financial-services-2024

[13] National Institute of Standards and Technology. (2024). *Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile*. NIST AI 600-1. https://doi.org/10.6028/NIST.AI.600-1

[14] Financial Stability Board. (2026). *Sound Practices for Responsible Adoption of Artificial Intelligence: Consultation Report*. 10 June 2026. https://www.fsb.org/uploads/P100626.pdf

[15] Central Bank of the UAE. (2026). *Guidance Note on the Consumer Protection and Responsible Adoption and Use of Artificial Intelligence and Machine Learning by Licensed Financial Institutions in the U.A.E.* https://rulebook.centralbank.ae/en/rulebook/guidance-note-consumer-protection-and-responsible-adoption-and-use-artificial-intelligence

[16] Gambacorta, L., Huang, Y., Qiu, H., and Wang, J. (2019). How do machine learning and non-traditional data affect credit scoring? New evidence from a Chinese fintech firm. *BIS Working Papers*, No. 834. https://www.bis.org/publ/work834.htm

[17] Wang, W., Lesner, C., Ran, A., Rukonic, M., Xue, J., and Shiu, E. (2020). Using small business banking data for explainable credit risk scoring. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(08), 13396-13401. https://doi.org/10.1609/aaai.v34i08.7055

[18] Stevenson, M., Mues, C., and Bravo, C. (2021). The value of text for small business default prediction: A deep learning approach. *European Journal of Operational Research*, 295(2), 758-771. https://doi.org/10.1016/j.ejor.2021.03.008

APPENDIX A. CREDIT EVIDENCE RECORD

Field	Example requirement
Deal and entity ID	Stable internal identifier and verified legal entity
Source ID	File hash or system version, date, owner and permission
Fact	Value, unit, sign, period, perimeter and source location
Verification	Proposed, reconciled, corroborated, reviewer-verified or unresolved
Adjustment	Reported value, amount, rationale, cash effect and approval
Assumption	Variable, scenario, range, basis and approver
Calculation	Formula or code version, inputs, output and validation
Model output	Model, version, intended use, score, explanation and limits
Exception	Policy or validation breach, owner, severity and resolution
Term	Facility, price, maturity, covenant, security and condition
Decision	Forum, authority, outcome, dissent, date and conditions
Monitoring	Metric, threshold, frequency, source and escalation

APPENDIX B. UNDERWRITING APPROVAL MATRIX

Output or action	Preparer	Validator	Approver	Release constraint
Source manifest	System or analyst	Deal-team reviewer	Deal lead	Approved perimeter
Financial spread	Controlled service	Analyst and reconciliation	Deal lead	No unresolved material difference
Adjustment schedule	Analyst with AI support	Independent credit	Committee if material	Evidence and rationale visible
Risk-model output	Approved model	Model owner or validator	Designated risk authority	Intended use only
Scenario analysis	Controlled service	Deal team and independent credit	Committee	Assumptions approved
Covenant package	Deal team and counsel	Independent credit and legal	Committee or delegated authority	Draft and executed definitions reconciled
Credit memorandum	Deal team	Independent credit	Committee	Sources, exceptions and conditions complete
Term sheet	Deal team and counsel	Legal and control functions	Named authority	Recipient and version check
Monitoring setup	Portfolio team	Operations and independent credit	Portfolio authority	Approved terms transferred

APPENDIX C. PILOT SCORECARD

Dimension	Measure	Definition
Evidence	Source coverage	Required sources registered and current
Evidence	Citation validity	Material claims supported by exact source locations

Dimension	Measure	Definition
Extraction	Field exact match	Value, unit, period and perimeter all correct
Finance	Reconciliation pass	Statements and schedules satisfy deterministic checks
Finance	Calculation fidelity	Independent recomputation matches approved result
Model	Calibration	Predicted and observed outcomes align within approved bands
Model	Stability	Performance and data drift remain within limits
Language	Unsupported-claim rate	Material claims without adequate support
Workflow	Tool precision	Permitted tool calls with correct deal and parameters
Workflow	Stop accuracy	Required abstentions and escalations occur
Authority	Approval violation	Material or external action without named approval
Quality	First-pass acceptance	Packets accepted without material correction
Productivity	Reviewer time	Median active reviewer hours per accepted packet
Productivity	Accepted throughput	Accepted packets per defined period
Risk	Incident rate	Classified incidents per production packet
Cost	Total operating cost	Models, data, tools, people, validation and controls

APPENDIX D. INCIDENT RECORD

Minimum fields are incident ID, deal, date, stage, system and model version, affected source or output, severity, detection method, immediate containment, authority notified, borrower or recipient impact, root cause, remediation, retest, owner and closure approval.

APPENDIX E. GLOSSARY

Abstention. A controlled decision by the system not to produce or release an output because evidence, confidence, authority or intended-use requirements are not met.

Credit evidence ledger. The linked record of entities, sources, facts, adjustments, assumptions, calculations, models, exceptions, terms, decisions and monitoring triggers.

Deterministic service. A rule or calculation that produces the same output from the same approved inputs and version.

Intended use. The approved decision, population, data and operating context for a model or system.

Model override. A documented human change to or departure from a model output.

Proposed fact. A value or statement extracted or generated for verification and not yet accepted as verified.

Private credit. Non-publicly traded debt provided primarily by non-bank lenders through bilateral or club-style negotiation, as used within this paper's stated scope.

Work packet. A bounded underwriting task with defined evidence, tools, output, checks, exceptions, owner and authority.