

Claude AI for Investment Banking: Rebuilding Origination, Diligence and Pitch Work

A controlled operating model for source-grounded research, analysis and execution support

Chennakeshav (CK) Adya

Independent Researcher

July 2026

ABSTRACT

Background. Investment-banking teams produce research, financial analysis and transaction materials under demanding time, confidentiality and review constraints. Claude and connected financial-services tooling can support document analysis, modelling and pitchbook preparation, while official financial-sector evidence identifies data, model, governance and third-party risks that require active control. **Objective.** This paper develops an operating model for applying Claude to origination, diligence and pitch work while preserving source traceability, information barriers, professional judgement and accountable approval. **Approach.** The paper combines current Anthropic product documentation, the Model Context Protocol specification, UK financial-sector survey evidence, Financial Stability Institute analysis and the NIST Generative AI Profile. It decomposes each workflow into retrieval, extraction, calculation, drafting, checking, judgement and approval tasks. A worked mandate and productivity model are presented as Matchpoint scenario analysis; their inputs are illustrative and fully disclosed. **Findings.** The appropriate unit of adoption is the controlled task. Read-only retrieval, structured extraction and evidence-linked drafting can enter an early pilot. Valuation conclusions, client positioning, conflicts decisions and external release remain accountable human decisions. A viable architecture combines approved sources, least-privilege connectors, source-level citations, deterministic calculation checks, exception logs and named approval gates. The scenario analysis shows how a team can measure capacity effects without treating gross model output as realised productivity. **Implications.** Deal teams should judge deployment through source coverage, correction rates, exception closure and review effort. The operating model places auditable evidence and professional accountability inside the workflow.

Highlights

- The controlled task is the practical unit of AI adoption in investment banking.
- Source coverage, reviewer correction and unresolved exceptions are core operating metrics.
- Read-only retrieval and extraction provide a suitable starting point for a governed pilot.
- Every connected tool requires an explicit permission boundary, audit trail and approval rule.
- Capacity claims require disclosed baselines, review loads, adoption rates and sensitivities.

JEL Classification: G24, G28, G32, O32, O33

Keywords: investment banking, generative artificial intelligence, Claude, origination, due diligence, pitchbooks, model risk, data governance, human oversight

This paper is an educational and structural analysis prepared for research purposes. It is not investment advice, legal advice, an offer or a solicitation. It contains no client information. Product capabilities and policies are described from cited public sources current to 31 July 2026; each institution must verify its contract, deployment configuration and applicable obligations.

1. INTRODUCTION

Investment banking combines standardised production with concentrated professional judgement. A team repeatedly collects information, compares businesses, constructs financial models, prepares questions and assembles presentations. Each mandate also contains facts, conflicts, assumptions and positioning choices that require accountable judgement. Generative AI therefore enters a workflow with substantial technical potential and substantial control requirements.

Anthropic describes Claude for Financial Services as supporting financial analysis, due diligence, market research, modelling and presentation workflows through financial-data and internal-platform connections [1]. Its finance-agent templates extend this description to pitchbooks, market research, due diligence and modelling [3]. These publications establish product intent and available workflow patterns. They do not establish a universal productivity result for an investment bank. Customer statements in vendor publications remain customer-reported evidence and require attribution.

The financial-sector evidence explains why the control design matters. In the Bank of England and Financial Conduct Authority's 2024 survey, 75 per cent of responding firms reported current AI use and a further 10 per cent planned use within three years [9]. Respondents rated data privacy and protection, data quality and data security among the largest current risks. They expected third-party dependencies, model complexity and embedded models to increase most strongly [9]. The survey covers UK financial services and does not isolate investment-banking deal teams.

This paper asks how a corporate-finance or investment-banking team can redesign three high-volume workflows around Claude while maintaining the evidence and accountability expected of transaction work. The analysis addresses three research questions:

- **RQ1.** Which origination, diligence and pitch tasks are suitable for retrieval, extraction, calculation, drafting, checking or agentic execution?
- **RQ2.** Which controls preserve source traceability, confidentiality, information barriers and professional accountability?
- **RQ3.** How should a team measure capacity effects when model output creates additional review, exception handling and governance work?

The paper makes four contributions. First, it presents a task-level map for the three workflows. Second, it defines a reference architecture connecting approved data, model services and office-document tools through explicit permission and review boundaries. Third, it provides a transparent scenario model that separates gross time released from realised capacity. Fourth, it supplies a 90-day implementation sequence and an operational control checklist.

The unit of analysis is a task with a defined input, output, evidence standard, permission scope and approver. This framing supports selective adoption. It gives partners, risk functions and technology owners a shared vocabulary for deciding what enters a pilot and what stays outside it.

2. EVIDENCE BASE AND PROPOSITIONS

2.1 Product and interface evidence

Anthropic's public financial-services materials describe a connected environment for market data, internal knowledge and office-document workflows [1, 2, 3]. The 2025 release described finance skills for comparable-company analysis, discounted cash-flow modelling, due-diligence data packs, company teasers, earnings analysis and coverage reports [2]. The May 2026 finance-agent release describes templates that combine instructions, governed data connectors and subagents [3]. These are vendor descriptions of capability and product architecture.

The Claude API citations feature can tie response text to passages in user-provided documents [4]. A citation proves that the cited passage was supplied and linked to the output. Reviewers still assess the source's

authority, date, context and analytical relevance. Investment-banking controls therefore need two separate tests: citation integrity and source quality.

The Model Context Protocol distinguishes prompts, resources and tools [6]. Resources supply context. Tools can retrieve information or perform actions [7, 8]. The specification recommends a human in the loop who can deny tool invocations and interfaces that reveal available and invoked tools [8]. These distinctions support a risk-based permission model. Read access to an approved research library has a different operational consequence from write access to a client file, CRM record or outbound communication channel.

Public retention documentation differentiates product routes and contractual arrangements [5]. A bank cannot determine confidentiality treatment from a product name alone. The controlling evidence includes the executed contract, workspace configuration, deployment route, retention setting, geographic arrangement and subprocessor position. The control owner should record these elements before any confidential mandate material enters the system.

2.2 Financial-sector risk evidence

The 2024 Bank of England and FCA survey reports broad use of AI across respondent firms and identifies data, security, model and third-party risks [9]. The Financial Stability Institute describes governance, skills, model-risk management, data governance and third-party providers as areas requiring attention [10]. It also identifies hallucination and anthropomorphism as risks associated with generative AI [10]. The FSI authors state that their views do not necessarily represent the BIS or its members.

NIST's Generative AI Profile identifies confabulation, privacy, information-security, intellectual-property and human-AI configuration risks [11]. It organises actions through Govern, Map, Measure and Manage. This structure fits a deal-team implementation: governance defines accountability; mapping establishes task context and impact; measurement records performance and failure; management constrains or retires the workflow.

The Bank of England's February 2026 summary of industry roundtables records participant support for principles- and outcomes-based guardrails and continued caution among second-line functions [12]. Participants also discussed explainability and data-governance constraints. The summary records roundtable views under the Chatham House Rule; it is not new policy.

2.3 Task decomposition

A transaction workflow can be decomposed into seven task classes.

Task class	Typical activity	Evidence requirement	Default control position
Retrieval	Find filings, research, CRM history or data-room documents	Provenance, date and access entitlement	Read-only approved sources
Extraction	Capture figures, clauses, dates and entities	Page or cell citation; schema validation	Sampling and exception review
Calculation	Apply formulas, bridge values and run scenarios	Deterministic reconciliation	Formula and unit tests
Drafting	Create narrative, questions and slide text	Claim-level source links	Named reviewer
Checking	Compare model, narrative and presentation	Defined rules and materiality	Independent re-performance
Judgement	Select comparables, position valuation or assess risk	Rationale and accountable owner	Human decision
Approval	Release information internally or externally	Authority and audit record	Human sign-off

The decomposition separates linguistic fluency from decision authority. A model can draft a comparable-company rationale. The sector lead owns the comparable set. A model can reconcile figures across a spreadsheet and presentation. The engagement lead owns the released valuation and client positioning.

2.4 Propositions

Four propositions follow from the evidence and task model.

Proposition 1. Source-bound tasks with explicit schemas and deterministic checks will reach an acceptable control state earlier than open-ended judgement tasks.

Proposition 2. Realised capacity depends on reviewer correction and exception-handling effort as well as gross generation speed.

Proposition 3. Connected tools increase utility and operational exposure together; least privilege and approval gates become more important as write authority expands.

Proposition 4. Adoption quality will be visible in source coverage, correction rates, exception closure, audit completeness and user adherence before it is visible in headline productivity.

These propositions are operational hypotheses. The paper does not claim that they have been tested across a representative sample of banks.

3. CONTROLLED REFERENCE ARCHITECTURE

3.1 Architecture layers

The reference architecture contains six layers. The first is identity and policy: named users, role-based entitlements, mandate membership and information barriers. The second is the approved-source layer: market-data services, public filings, internal research, CRM history and deal documents. The third is the retrieval and citation layer, which provides bounded context with source identifiers. The fourth is the model and orchestration layer, containing approved models, task instructions and validation routines. The fifth is the controlled-tool layer for spreadsheets, documents, presentations and approved internal systems. The sixth is the evidence and approval layer: logs, exceptions, reviewer decisions, released versions and monitoring metrics.




Figure 1. Controlled investment-banking AI reference architecture

Source: Matchpoint framework, drawing on the MCP control model [6, 7, 8].

The information path should carry a mandate identifier, user identity, source identifier, document classification, task version and output status. A released output should be reproducible from an archived source set, recorded instruction version and named review record. Model reproducibility may remain limited when providers update services. Version and date records preserve the available audit trail.

3.2 Source registry

Every source entering the workflow should be registered with ownership, authority, permitted use, update frequency and retention treatment. Public filings and executed client documents occupy different evidence classes. Broker research and licensed market data carry contractual use restrictions. CRM notes may contain personal or relationship-sensitive information. A source registry makes those distinctions visible to the orchestration layer and reviewer.

A useful source hierarchy has four levels. Level A contains executed agreements, audited accounts, regulator filings and data obtained directly from an authorised provider. Level B contains management information and client-provided working data, clearly labelled by provenance and review status. Level C contains reputable secondary research. Level D contains unverified web content or generated material. Decision-critical claims should trace to the highest available level. Generated text never becomes a source through repetition.

3.3 Permission design

Each connector should begin with the smallest permission needed for the task. A market-mapping pilot may require read-only access to public filings and a controlled research folder. A pitch-production pilot may require read access to an approved model and write access to a specific working-copy directory. CRM write access, outbound email, external publishing and transaction execution require separate authority and explicit confirmation.

The MCP specification treats tools as model-controlled functions and recommends human control over invocation [8]. An institutional implementation should supplement this with allowlists, scoped credentials, output-path restrictions, rate limits and immutable logs. Tool descriptions and annotations require validation because the specification states that clients must treat them as untrusted unless they come from trusted servers [8].

3.4 Evidence objects and review gates

Every material output should contain or reference an evidence object. For a diligence finding, the object includes the source document, page or cell, extracted text, issue category, materiality, confidence marker, reviewer and disposition. For a valuation slide, it includes the data date, source, currency, units, calculation lineage, selected peer set and reviewer approval.

Four review gates cover the lifecycle:

1. **Input gate:** classification, entitlement, malware control and source registration.
2. **Analysis gate:** citation coverage, calculation reconciliation and exception review.
3. **Mandate gate:** conflicts, client facts, positioning and senior judgement.
4. **Release gate:** version control, disclosure, credentials, formatting and authorised delivery.

The gates create stopping points. An output remains a working draft until all applicable gates are recorded.

4. ORIGATION WORKFLOW

4.1 Market mapping and target formation

Origination begins with a defined thesis. The team specifies geography, sector, size, ownership, transaction trigger and exclusion criteria. Retrieval can then collect source-bound facts from approved databases and filings. Extraction can place companies into a standard schema. Drafting can create short hypotheses for senior review.

A target record should separate verified facts from hypotheses. Verified fields can include legal name, reporting currency, most recent filing date and disclosed ownership. Hypothesis fields can include likely capital need, strategic trigger and relevance to a buyer or investor. Every hypothesis should carry a rationale and review date. This structure avoids presenting model-generated opportunity language as fact.

The model can also identify missing evidence. A complete origination brief contains an evidence map and an uncertainty register. Missing ownership information, stale financials and inconsistent revenue measures are findings. They should remain visible through the coverage decision.

4.2 Account and meeting preparation

Meeting preparation often requires assembling public information, relationship history, recent correspondence, prior proposals and sector context. A controlled workflow retrieves these materials from entitled systems and produces a dated brief. The brief should label public facts, internal relationship notes and analyst hypotheses separately.

The relationship owner approves the meeting objective, participants, sensitivities and outreach language. The model can draft questions and identify inconsistencies. It does not determine conflicts clearance or the suitability of a proposed engagement. Those decisions remain with authorised professionals and the firm's established processes.

4.3 Pipeline control

A coverage pipeline benefits from consistent stage definitions and evidence requirements. AI can flag missing fields, overdue follow-ups and unsupported stage changes. It can prepare a daily exception list. The relationship owner remains accountable for client contact, stage assessment and forecast judgement.




Figure 2. Origination workflow with controlled augmentation

Source: Matchpoint framework.

The redesigned sequence is thesis definition; approved-source retrieval; schema extraction; evidence and uncertainty review; senior hypothesis selection; conflicts process; authorised outreach; CRM recording. The main control improvement is an explicit boundary between collected fact and commercial hypothesis.

4.4 Origination control matrix

Activity	Claude role	Human owner	Required evidence	Release condition
Build long list	Retrieve and extract	Coverage analyst	Source and date for each factual field	Sampling complete
Draft target thesis	Draft	Sector lead	Linked facts and labelled hypotheses	Sector lead approval
Prepare meeting brief	Retrieve, reconcile and draft	Relationship owner	Public and internal sources separated	Sensitivities reviewed
Draft questions	Draft	Deal lead	Question rationale	Deal lead selection
Update pipeline	Validate fields or prepare	Relationship owner	Stage definition and	Authorised CRM action

Activity	Claude role	Human owner	Required evidence	Release condition
	change		activity evidence	
Draft outreach	Draft	Relationship owner and compliance process	Approved claims and credentials	Human release

5. DILIGENCE WORKFLOW

5.1 Data-room inventory

Diligence begins with the document population. The system should inventory filenames, versions, dates, owners, page counts and document types before substantive extraction. Duplicate, corrupt, password-protected and scanned files enter an exception queue. The inventory gives reviewers a denominator for completeness.

File content should be treated as untrusted input. A malicious or irrelevant instruction embedded in a document can affect an agentic workflow. NIST identifies information-security risks in generative AI systems [11]. Connected implementations should isolate document content from system instructions, restrict tool authority and require approval for state-changing actions.

5.2 Extraction and reconciliation

Extraction should use a task-specific schema. A debt schedule schema can require facility, lender, currency, commitment, drawn amount, rate, maturity, security and change-of-control terms. A customer-contract schema can require counterparty, term, renewal, price, volume, termination and assignment. Empty fields remain empty and become exceptions. The system should not complete missing facts through prediction.

Reconciliation compares the same fact across sources. Revenue can differ across audited accounts, management reporting and the operating model because of period, scope, accounting treatment or error. The workflow should display the values together with source and date. The reviewer determines the explanation and materiality.

The Claude citations feature can return passages linked to supplied documents [4]. A controlled diligence design should retain the cited passage and location with the extracted field. Citation coverage can then be measured as the share of material findings that contain a valid source pointer. Reviewers still confirm that the passage supports the finding.

5.3 Question matrices and issue logs

The system can convert missing, conflicting or unusual observations into draft questions. Each question should carry the triggering evidence, issue category, materiality and owner. The deal team selects which questions go to the client or counterparty. The issue log records open, answered, partially answered and closed states with supporting evidence.

This structure reduces a common diligence failure: an answer arrives in an email or revised file, yet the original issue stays open or loses its evidence trail. A linked issue register preserves chronology and reviewer accountability.

5.4 Diligence task-control matrix

Task	AI function	Primary failure mode	Required control	Accountable reviewer
File inventory	Classification	Missing or duplicate files	Count reconciliation and exception queue	Data-room manager
Clause extraction	Extraction	Wrong clause or omitted qualifier	Page citation and sample review	Relevant adviser
Financial extraction	Extraction	Units, sign or period error	Cell/page source and totals check	Financial analyst

Task	AI function	Primary failure mode	Required control	Accountable reviewer
Cross-document check	Comparison	False match or missed inconsistency	Materiality rules and reviewer disposition	Workstream lead
Question drafting	Drafting	Unsupported or poorly framed question	Evidence link and owner review	Workstream lead
Issue prioritisation	Assistance	Misstated materiality	Human materiality decision	Deal lead
Final diligence conclusion	Draft support	Incomplete reasoning	Source pack, open-issue list and senior sign-off	Engagement lead

Evidence-linked diligence workflow




Figure 3. Evidence-linked diligence workflow

Source: Matchpoint framework.

The sequence is inventory; classification; schema extraction; citation attachment; deterministic reconciliation; exception queue; reviewer disposition; question matrix; evidence-backed closure. The issue register forms the control spine.

6. PITCH AND EXECUTION WORKFLOW

6.1 Financial analysis

Financial work should separate data movement, formula construction, assumption selection and valuation judgement. Claude can assist with mapping source data, explaining formulas, generating checks and documenting assumptions. The approved spreadsheet remains subject to deterministic tests: balance-sheet balance, cash-flow reconciliation, sign and unit checks, period consistency, circularity controls and scenario boundaries.

Comparable-company analysis requires judgement in peer selection, metric definition and treatment of outliers. The model can assemble candidate data and draft a rationale. The sector lead approves the peer set, dates, adjustments and interpretation. Discounted cash-flow analysis also requires human approval of forecast logic, discount rate, terminal value and sensitivities.

6.2 Narrative and presentation production

A pitchbook should be generated from an approved fact and assumption register. Slide text can be drafted from that register with claim-level source links. The presentation workflow should check that figures, dates and units agree with the financial model and source register. Credentials, transaction claims and market statements require separate verification.

The slide is an output view of controlled underlying objects: chart data, narrative claim, source, calculation, reviewer and approval status. This design supports corrections at the source object and reduces inconsistent manual updates across pages.

6.3 Approval controls

Output	Minimum control	Human decision
Comparable-company table	Data date, source, currency, metric definition and formula check	Peer selection and exclusions
DCF output	Model reconciliation, assumption register and sensitivity checks	Forecast, discount rate and terminal value
Market overview	Claim-level sources and date cut-off	Interpretation and client relevance
Credentials page	Approved credentials library and engagement-status check	Permission to use each credential
Transaction structure	Sources and explicit assumptions	Feasibility, positioning and advice
Executive summary	Cross-reference to evidence and open issues	Core message and recommendation
Client-ready deck	Version, formatting, disclosures and release record	Final external release

6.4 Release discipline

The release workflow should generate a frozen package containing the approved presentation, supporting model, source register, exception log and sign-off record. Draft outputs should carry a visible status. External delivery permissions should sit outside the model's default authority. The engagement lead or authorised delegate initiates release after the recorded gate.

7. PRODUCTIVITY ECONOMICS

7.1 Measurement model

Productivity analysis should begin with observed task baselines. A team records elapsed and active hours for a representative sample, broken down by retrieval, extraction, calculation, drafting, checking, judgement, approval and rework. The AI-enabled pilot records the same categories plus prompt preparation, exception handling and model review.

The following model is **Matchpoint scenario analysis**. It is illustrative and does not represent observed performance at Matchpoint, Anthropic or any investment bank.

For task class i :

Gross hours released = baseline hours $\times$ eligible share $\times$ adoption rate $\times$ task-level time reduction.

Incremental control hours = AI-processed hours $\times$ review load + exception hours + governance hours.

Realised capacity = gross hours released minus incremental control hours.

The model excludes revenue effects and quality benefits. Those outcomes require separate evidence.

7.2 Base scenario

The illustrative mandate uses 420 baseline team hours across origination, diligence and pitch production. Eligibility is limited to retrieval, structured extraction, first-draft preparation and defined checks. Judgement and approval hours are excluded from time-reduction assumptions.

Workflow	Baseline hours	Eligible share	Adoption rate	Time reduction on adopted eligible work	Review and exception hours	Realised hours released
Origination	70	55%	70%	45%	7	5.1

Workflow	Baseline hours	Eligible share	Adoption rate	Time reduction on adopted eligible work	Review and exception hours	Realised hours released
Diligence	210	60%	65%	40%	18	14.8
Pitch and execution	140	50%	65%	35%	13	2.9
Total	420				38	22.8

The arithmetic is rounded to one decimal place. For origination, gross hours released are $70 \times 55\% \times 70\% \times 45\% = 12.1$ hours; subtracting seven control hours produces 5.1 realised hours. Diligence produces 32.8 gross hours and 14.8 realised hours after 18 control hours. Pitch work produces 15.9 gross hours and 2.9 realised hours after 13 control hours.

The scenario yields 60.8 gross hours released and 22.8 realised hours released, equivalent to 5.4 per cent of the 420-hour baseline. These figures are model outputs from disclosed assumptions. They are not forecasts or measured results.




Figure 4. Matchpoint scenario analysis: baseline, gross release and control load

Source: Matchpoint scenario analysis. Values are illustrative assumptions, not observed performance.

7.3 Sensitivity and decision rules

The model is most sensitive to task eligibility, reviewer correction and exception handling. A fast draft with extensive correction can consume more senior time than a slower controlled workflow. A pilot should therefore measure correction minutes per output, material error rate, exception volume and reviewer seniority.

Management can set an exit rule for each task. A task progresses when citation coverage, material error rate, reviewer correction time and audit completeness remain within approved thresholds across a defined sample. A task pauses when exceptions exceed tolerance or the source environment changes. Thresholds are institution-specific management decisions.

Capacity should be assigned deliberately. Released analyst time can support deeper client preparation, more complete diligence, sector work or reduced peak-hour pressure. The measurement record should identify the destination of released time if management wants to test commercial or quality effects.

8. RISK AND GOVERNANCE

8.1 Governance model

The governance model assigns an accountable executive sponsor, a workflow owner, a technology owner, a risk owner, data owners and named reviewers. The workflow owner defines the intended use, excluded use,

evidence standard and exit criteria. The technology owner maintains model, connector and instruction versions. The risk owner reviews controls and incidents. Data owners approve sources and access.

NIST's Govern, Map, Measure and Manage functions offer a useful operating cycle [11]. The firm governs through ownership and policy; maps the task, users, data and consequences; measures accuracy, correction, security and drift; manages through restrictions, remediation or retirement.

8.2 Risk-control matrix

Risk	Exposure in deal work	Preventive control	Detective control	Response owner
Confabulation	Unsupported facts, clauses or explanations	Source-bound prompts and required citations	Claim-level review and source sampling	Workstream lead
Weak source authority	Correctly cited low-quality material	Approved-source hierarchy	Source-quality review	Research owner
Data leakage	Confidential information reaches an unapproved service or user	Contract, configuration, classification and access controls	Access logs and incident monitoring	Data and security owners
Prompt injection	Data-room content attempts to redirect tools or expose data	Content isolation, least privilege and tool allowlists	Tool-call and anomaly logs	Security owner
Spreadsheet error	Unit, sign, formula or linkage error	Controlled templates and calculation boundaries	Reconciliation and formula tests	Model owner
Permission escalation	Tool performs an unauthorised write or release	Scoped credentials and approval gates	Immutable action logs	Technology owner
Model or vendor change	Output behaviour changes	Version registry and change assessment	Benchmark rerun and drift monitoring	Model owner
Automation bias	Reviewer accepts fluent output without adequate challenge	Training, visible evidence and accountable sign-off	Correction analysis and review sampling	Engagement lead
Record failure	Source, instruction or decision cannot be reconstructed	Output package and retention rule	Audit-completeness checks	Records owner
Third-party concentration	Dependence on a model, cloud or data provider	Exit plan and architecture portability	Service and concentration review	Vendor owner

8.3 Confidentiality and retention

The deployment decision should be documented at contract and configuration level. Anthropic publishes separate retention information for API and commercial product routes [5]. The institution should record its actual route, agreed retention, training treatment, subprocessor chain, access model and deletion process. A public vendor statement cannot substitute for the executed terms.

Mandate-level controls should apply information barriers and need-to-know access. Cross-mandate retrieval should be disabled unless explicitly authorised. Logs should avoid unnecessary reproduction of confidential content. Working files should remain in approved repositories with established retention and legal-hold processes.

8.4 Model risk and human review

Generative AI can produce confident, incorrect content. NIST calls this confabulation [11]. The FSI identifies hallucination as a generative-AI risk relevant to financial-sector oversight [10]. Controls should therefore rely on source evidence, deterministic checks and accountable review. A confidence score generated by the same model is not sufficient evidence of correctness.

Review intensity should reflect consequence. A public market brief may allow sampled factual review. A valuation conclusion, fairness analysis, client recommendation or external statement requires comprehensive professional review under the firm's applicable process. The system should record the reviewer and the evidence considered.

8.5 Incident and change management

An incident process should cover data exposure, unauthorised action, material factual error, missing records and control bypass. The response includes containment, affected-output identification, notification under applicable procedures, root-cause analysis and benchmark updates.

Changes to models, instructions, connectors, source schemas and office-document integrations can alter performance. Material changes should trigger a scoped validation before production use. The benchmark set should contain representative easy cases, ambiguous cases, adversarial content and known failure examples.

9. NINETY-DAY IMPLEMENTATION ROADMAP

9.1 Phase 1: govern and map, days 1 to 20

The firm selects one workflow with bounded data and an accountable sponsor. It defines intended use, exclusions, users, sources, permissions, outputs and review rules. The team establishes a benchmark from completed or public-material cases. Contract, retention and security evidence is recorded.

Exit criteria are an approved use-case card, source registry, data classification, permission map, reviewer matrix, benchmark and incident route.

9.2 Phase 2: read-only pilot, days 21 to 45

The pilot uses read-only sources and creates working drafts in a segregated environment. Suitable tasks include document inventory, schema extraction, evidence-linked summaries and draft question lists. Every output is reviewed. The team records source coverage, material errors, correction time, exceptions and user adherence.

Exit criteria are a completed sample, documented failure modes, threshold results, closed high-severity exceptions and approval for the next scope.

9.3 Phase 3: governed production, days 46 to 70

Approved tasks move into a controlled live workflow. The system produces evidence objects, exception registers and versioned outputs. Spreadsheet and presentation assistance can enter scope where formula and release controls exist. Access remains mandate-specific.

Exit criteria are stable operating metrics, complete audit records, trained users, support ownership and a tested rollback process.

9.4 Phase 4: connected tools, days 71 to 90

The firm considers tightly scoped write actions. Each action receives an allowlist, destination restriction, confirmation rule and log. External communications and publication remain behind explicit human release. The team reruns benchmarks after model or connector changes.

Phase	Primary owner	Deliverable	Core metric	Gate
Govern and map	Executive and workflow owners	Approved use-case card	Control completeness	Risk approval
Read-only pilot	Workflow owner	Reviewed pilot sample	Source coverage and corrections	Threshold decision
Governed production	Operations owner	Live controlled task	Exceptions and audit completeness	Production approval

Phase	Primary owner	Deliverable	Core metric	Gate
Connected tools	Technology and risk owners	Scoped action workflow	Authorised action rate and incidents	Action approval

10. WORKED MANDATE

10.1 Scenario and evidence labels

This worked example is **Matchpoint scenario analysis**. It describes an illustrative mid-market capital-raise mandate. The company, numbers and events are fictional. They do not represent a client or transaction.

The illustrative company is Meridian Systems, a business-services software provider seeking growth capital. The fictional working data contains three years of accounts, a management model, 25 customer contracts, a debt schedule and a product plan. The workflow labels each value as scenario fact, management assumption or analytical output.

10.2 Origination brief

The system inventories approved public-style material and prepares a one-page brief. It records the fictional company's product, ownership, reported figures and potential trigger. It creates an uncertainty register for customer concentration, renewal terms and financing need. The coverage director reviews the facts and selects the meeting hypothesis. The conflicts process and outreach approval remain outside the model workflow.

10.3 Diligence pack

The system inventories the fictional data room and finds 32 files. It extracts customer term, renewal and termination fields into a schema. Four contracts are scanned and enter an exception queue. Two contracts contain inconsistent renewal dates between the signed agreement and management schedule. The system links both passages and drafts questions. A reviewer determines which date controls and whether the difference is material.

The financial extraction identifies three revenue measures with different scope. The system lists each value, period and source. The analyst documents the reconciliation. Missing answers remain open. No value is chosen automatically.

10.4 Pitch production

The approved fact register feeds a draft market and company section. The model creates a comparable-company candidate list from the defined sector criteria. The sector lead selects the peer set. The spreadsheet workflow checks currency, units, enterprise-value bridge and displayed multiples. The valuation owner approves all assumptions and conclusions.

The presentation workflow links every material chart and claim to its data object. A cross-document check compares the deck against the model and fact register. The release package contains the approved deck, model, fact register, open-issue register and sign-off record.

Illustrative mandate control flow



Figure 5. Worked mandate from source registration to release

Source: Matchpoint illustrative scenario.

10.5 Worked-mandate control outcomes

The example illustrates five control outcomes. First, the document inventory creates a completeness denominator. Second, schema extraction leaves missing information visible. Third, reconciliation presents conflicting facts for human disposition. Fourth, valuation judgement remains with named professionals. Fifth, release occurs from a frozen evidence package.

No productivity conclusion follows from the example. A live pilot would need observed timings, corrections and exceptions across a defined sample.

11. LIMITATIONS AND FUTURE RESEARCH

This paper is a design study. It does not report a controlled trial, representative bank survey or audited implementation. The productivity section is an illustrative scenario with disclosed assumptions. The worked mandate is fictional. These features limit empirical conclusions.

The product evidence is largely supplied by Anthropic. Capability descriptions can change, and customer-reported performance may not transfer to another institution or workflow. A production decision requires current product documentation, executed contractual terms and institution-specific testing.

The regulatory evidence is concentrated in UK and cross-sector sources. Investment banks operate under jurisdiction-, entity- and activity-specific requirements. Each implementation requires advice and approval under its applicable obligations. This paper does not provide legal or regulatory advice.

Future research should test the propositions with task-level operational data. A useful study would pre-register task definitions, sample selection, baseline methods, accuracy thresholds and review standards. It would record active time, elapsed time, correction effort, material errors, exceptions and source coverage across human-only and AI-assisted conditions. Further work should also examine permission design, prompt-injection resilience, model-change effects and reviewer behaviour.

12. CONCLUSION

Claude can support investment-banking origination, diligence and pitch production through retrieval, structured extraction, calculation assistance, evidence-linked drafting and defined checks. Current Anthropic materials describe these workflow capabilities [1, 2, 3]. Current financial-sector and NIST sources identify data, model, governance, security and third-party risks that require active management [9, 10, 11, 12].

The controlled task provides a practical adoption unit. Each task receives approved inputs, bounded permissions, an evidence standard, deterministic checks where available, exception handling and a named approver. Judgement and release remain accountable human actions.

The implementation sequence begins with read-only, source-bound work. Measurement focuses on source coverage, correction effort, material errors, exceptions, audit completeness and user adherence. Connected

actions enter scope after the control evidence supports them. This sequence gives management a traceable basis for deciding which workflows to expand, restrict or retire.

DECLARATIONS

Funding. The author received no external funding for this research.

Conflicts of interest. The author is the managing partner of an independent capital-advisory firm. No client information or live mandate is used in this paper.

Data availability. The paper uses public sources and an illustrative scenario. It uses no proprietary dataset.

Use of generative AI. Generative AI assisted research organisation, drafting and document production. The author retains responsibility for source selection, analysis and final approval.

Disclaimer. This paper is an educational and structural analysis. It is not investment advice, legal advice, an offer or a solicitation.

REFERENCES

- [1] Anthropic (2026), 'Claude for Financial Services', updated 10 April 2026, <https://www.anthropic.com/news/claude-for-financial-services>.
- [2] Anthropic (2025), 'Advancing Claude for Financial Services', 27 October 2025, <https://www.anthropic.com/news/advancing-claude-for-financial-services>.
- [3] Anthropic (2026), 'Agents for Financial Services', 5 May 2026, <https://www.anthropic.com/news/finance-agents>.
- [4] Anthropic (2026), 'Citations', Claude Platform documentation, accessed 31 July 2026, <https://platform.claude.com/docs/en/build-with-claude/citations>.
- [5] Anthropic (2026), 'API and data retention', Claude Platform documentation, accessed 31 July 2026, <https://platform.claude.com/docs/en/manage-claude/api-and-data-retention>.
- [6] Model Context Protocol (2025), 'Overview: server features', protocol revision 18 June 2025, <https://modelcontextprotocol.io/specification/2025-06-18/server/index>.
- [7] Model Context Protocol (2025), 'Specification', protocol revision 18 June 2025, <https://modelcontextprotocol.io/specification/2025-06-18/index>.
- [8] Model Context Protocol (2025), 'Tools', protocol revision 18 June 2025, <https://modelcontextprotocol.io/specification/2025-06-18/server/tools>.
- [9] Bank of England and Financial Conduct Authority (2024), 'Artificial intelligence in UK financial services; 2024', <https://www.bankofengland.co.uk/report/2024/artificial-intelligence-in-uk-financial-services-2024>.
- [10] Crisanto, J.C., Leuterio, C.B., Prenio, J. and Yong, J. (2024), 'Regulating AI in the financial sector: recent developments and main challenges', FSI Insights No. 63, Financial Stability Institute, <https://www.bis.org/fsi/publ/insights63.htm>.
- [11] Autio, C., Schwartz, R., Dunietz, J., Jain, S., Stanley, M., Tabassi, E., Hall, P. and Roberts, K. (2024), 'Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile', NIST AI 600-1, updated 8 April 2026, <https://doi.org/10.6028/NIST.AI.600-1>.
- [12] Bank of England (2026), 'Summary of AI roundtables; February 2026', 16 February 2026, <https://www.bankofengland.co.uk/minutes/2026/february/summary-of-ai-roundtables-feb-2026>.
- [13] Anthropic (2026), 'Learning Curves', Anthropic Economic Index, 24 March 2026, <https://www.anthropic.com/research/economic-index-march-2026-report>.
- [14] Anthropic (2026), 'Cadences', Anthropic Economic Index, 26 June 2026, <https://www.anthropic.com/research/economic-index-june-2026-report>.
- [15] Anthropic (2026), 'Model system cards', accessed 31 July 2026, <https://www.anthropic.com/system-cards>.

APPENDIX A. SCENARIO ASSUMPTION REGISTER

Assumption	Base value	Status	Purpose
Total mandate team hours	420	Matchpoint scenario assumption	Denominator for illustration
Origination hours	70	Matchpoint scenario assumption	Workflow baseline
Diligence hours	210	Matchpoint scenario assumption	Workflow baseline
Pitch and execution hours	140	Matchpoint scenario assumption	Workflow baseline
Origination eligible share	55%	Matchpoint scenario assumption	Retrieval, extraction and drafting scope
Diligence eligible share	60%	Matchpoint scenario assumption	Inventory, extraction and reconciliation scope
Pitch eligible share	50%	Matchpoint scenario assumption	Data preparation, drafting and checking scope
Adoption rate	65% to 70%	Matchpoint scenario assumption	Share of eligible work processed in the pilot
Task-level time reduction	35% to 45%	Matchpoint scenario assumption	Gross processing effect on adopted eligible work
Review and exception hours	38	Matchpoint scenario assumption	Incremental control workload

These values are illustrative. They are included to make the arithmetic reproducible. They should be replaced with observed institutional baselines and pilot results before management uses the model for planning.

APPENDIX B. CONTROL CHECKLIST**Governance and scope**

- Named executive sponsor, workflow owner, technology owner and risk owner.
- Written intended use, users, outputs, excluded uses and decision boundaries.
- Current contract, deployment, retention and training-treatment evidence.
- Applicable legal, regulatory, information-barrier and records requirements mapped.
- Exit, suspension and incident criteria approved.

Data and access

- Approved-source register with owner, authority, licence and update cycle.
- Data classification and mandate membership applied.
- Least-privilege connector scopes documented and tested.
- Cross-mandate retrieval disabled unless explicitly authorised.
- Tool destinations, rate limits and approval requirements configured.

Output and review

- Citations retained for material extracted facts and drafted claims.
- Source authority reviewed independently of citation validity.
- Deterministic tests cover calculations, units, signs, periods and reconciliations.
- Missing and conflicting information enters an exception register.

- Named reviewers approve judgement and external release.
- Frozen release package includes sources, model, output, exceptions and sign-offs.

Monitoring and change

- Metrics record source coverage, corrections, material errors, exceptions and audit completeness.
- Representative benchmark set includes ambiguous and adversarial cases.
- Material model, connector, instruction or schema changes trigger validation.
- Incident response and rollback process tested.
- Periodic access, vendor and concentration reviews scheduled.

APPENDIX C. GLOSSARY

Citation integrity. Confirmation that an output points to a real location in a supplied source and that the cited passage supports the claim made.

Confabulation. Generative output that confidently presents erroneous or false content; NIST uses this term for behaviour also commonly described as hallucination [11].

Controlled task. A bounded activity with defined inputs, outputs, evidence requirements, permissions, checks and approval.

Evidence object. A structured record linking an output or finding to its source, location, classification, reviewer and disposition.

Exception register. A record of missing, conflicting, failed or out-of-tolerance items requiring human disposition.

Human approval gate. A recorded point at which an authorised person accepts, rejects or returns an output before a consequential action or release.

Least privilege. A permission design that grants only the access needed for the defined task and duration.

Matchpoint scenario analysis. Illustrative modelling created for this paper from disclosed assumptions; it is not observed performance data.

Source coverage. The proportion of material output claims or extracted fields that contain a valid link to an approved source.

Tool. An executable function exposed to a model application for retrieval, computation or action [8].