

Hypothesis

Not peer-reviewed version

National Financial Resilience: A Research Proposal for Financial Stability and ROI Measurement in AI Investment at Scale in US Banking

[Satyadhar Joshi](#)*

Posted Date: 30 April 2026

doi: 10.20944/preprints202604.2094.v1

Keywords: artificial intelligence; ROI measurement; capital allocation; risk management; US banking; financial resilience; model risk; real options; RAROC; causal inference; agentic AI; regulatory compliance; federal reserve; OCC; FDIC



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC, OpenAlex.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Hypothesis

National Financial Resilience: A Research Proposal for Financial Stability and ROI Measurement in AI Investment at Scale in US Banking

Satyadhar Joshi ^{1,2} 

¹ Independent Researcher, Alumnus, International MBA, Bar-Ilan University, Israel; satyadhar.joshi@gmail.com

² Alumnus, Touro College MSIT, NY, USA

Abstract

U.S. banks are investing unprecedented amounts in artificial intelligence, with annual spending at institutions like JPMorgan Chase, Bank of America, and Citigroup now exceeding \$2–\$4 billion each. Yet a critical national financial resilience problem persists: most U.S. banks cannot confidently determine whether these massive AI investments generate positive risk-adjusted returns, creating capital allocation inefficiency and potential systemic vulnerability. **This research proposal** outlines a comprehensive mixed-methods research design for investigating how senior executives in U.S. global banks govern enterprise AI investments, manage emerging financial risks, and measure return on investment when scaling AI across national banking operations. Drawing on the Resource-Based View, Paradox Theory, and the Technology-Organization-Environment framework, **this proposal develops** an integrated conceptual framework linking AI governance mechanisms, operating model configurations, and multi-dimensional ROI measurement specifically calibrated to the U.S. regulatory environment (Federal Reserve, OCC, FDIC). **The proposed study would employ** an embedded multiple-case design with semi-structured interviews of 30–40 C-Suite executives across 6–8 U.S.-headquartered global banks, supplemented by secondary analysis of SEC filings, FRED economic data, FDIC call reports, and Model Risk Management documentation. **We propose** a novel risk-adjusted ROI calculation framework incorporating direct financial benefits, indirect value creation, strategic option pricing, and probabilistic risk adjustments aligned with U.S. banking stress testing practices. **Anticipated methodological barriers** include organizational resistance, access constraints to senior executives, and causal attribution challenges—each addressed with specific mitigation strategies outlined in this proposal. **This proposal aims to contribute** empirically validated ROI measurement tools for executive decision-making at U.S. systemically important financial institutions and demonstrates a scholar-practitioner approach to bridging academic rigor with national financial stability priorities.

Keywords: artificial intelligence; ROI measurement; capital allocation; risk management; US banking; financial resilience; model risk; real options; RAROC; causal inference; agentic AI; regulatory compliance; federal reserve; OCC; FDIC

1. Introduction and Problem Statement

1.1. The Scale of AI Investment in Global Banking

Global banks are expending unprecedented financial resources on artificial intelligence infrastructure, talent, and deployment. JPMorgan Chase alone invests over \$2 billion annually in AI and cloud technologies [1]. Bank of America has prioritized large-scale AI initiatives, with annual spending on new technology increasing by 44% over the past decade, reaching approximately \$4 billion annually according to Chief Technology and Information Officer Hari Gopalkrishnan. Citigroup, Goldman Sachs, Wells Fargo, and Morgan Stanley have each committed comparable magnitudes, positioning financial services as one of the most AI-intensive sectors in the global economy.

As [2] document in their comprehensive systematic review, AI integration in financial services has evolved dramatically from 1989 to 2024. Their scientometric analysis reveals that global spending on analytics, AI, and big data platforms is projected to surpass \$300 billion by 2030, with machine learning, natural language processing, and blockchain technologies fundamentally reshaping financial operations and decision-making processes. The authors note that AI applications in credit scoring, fraud detection, digital insurance, robo-advisory services, and financial inclusion have transitioned from experimental prototypes to operational systems, yet the financial returns from these investments remain poorly understood and inconsistently measured across institutions.

1.2. The Business Problem: Unmeasured ROI and Capital Allocation Inefficiency

Despite staggering investment levels, a fundamental business problem remains: most banks cannot confidently determine whether their AI spending generates positive risk-adjusted returns, creating capital allocation inefficiency and shareholder value destruction. [3] conducted a quantitative study of 512 senior IT/IS managers in public and private organizations, identifying that relative advantage, top management support, cost-effectiveness, and competitive pressure positively influence AI adoption intentions, while government regulation and complexity negatively influence them. However, their research—like most existing literature—focuses on adoption intentions rather than post-adoption financial performance measurement and value realization.

Banks face a complex capital allocation dilemma: how to deploy billions of dollars in AI investments across competing use cases—fraud detection, credit scoring, customer personalization, algorithmic trading, risk modeling, compliance automation, and wealth management—while managing novel risk categories and maintaining regulatory compliance across multiple jurisdictions. Unlike traditional IT projects where ROI can be measured through direct cost savings or efficiency gains using standard capital budgeting techniques, AI initiatives generate complex, temporally delayed, and difficult-to-attribute financial outcomes that resist conventional evaluation methodologies.

Ref. [4] find through qualitative interviews with 29 AI developers, managers, and users that attitudes toward AI shift from negative or instrumental to positive as individuals gain knowledge and experience. This temporal dimension creates significant uncertainty in capital planning: executives may systematically undervalue AI investments because benefits accrue only after organizational learning occurs, or may overvalue them based on vendor promises that fail to materialize. Ref. [5] further emphasize that while AI value-creation mechanisms are increasingly documented, the organizational enablers and barriers to value realization remain undertheorized and empirically under-examined.

1.3. Regulatory and Risk Dimensions

The measurement problem is exacerbated by the emergence of Agentic AI—autonomous systems capable of independent reasoning and sequential action—which fundamentally challenges traditional human-in-the-loop control mechanisms and introduces novel financial risk categories. Ref. [6] emphasize that unlike traditional IT risks, AI introduces algorithmic bias, lack of explainability, model drift, and vulnerability to adversarial attacks—each capable of producing multimillion-dollar losses through failed deployments, regulatory fines, litigation, or reputational damage. Their strategic alignment framework for corporate and investment banking provides a governance-first approach to converting data and AI investments into measurable commercial value.

Ref. [7] provide a comparative analysis of how nine major jurisdictions—the European Union, United Kingdom, United States, Canada, Japan, Australia, India, China, and Islamic banking states—supervise AI in banking. Their ten-dimension comparative matrix identifies three archetypes of algorithmic governance: legal-normative (EU and UK), executive-fragmented (US and Canada), and state-centralized (China). The authors demonstrate that many jurisdictions adopt the language of "trustworthy" AI without establishing equivalent mechanisms of supervisory enforceability, creating regulatory arbitrage opportunities and cross-border compliance challenges for global banks operating across multiple regimes.

1.4. Trade Journal Perspectives on Practical Urgency

Trade publications highlight the practical urgency of this problem from multiple stakeholder perspectives. The Bank of England recently announced it will include AI risks in financial system stress tests, signaling that regulators view AI as a potential systemic risk requiring active monitoring, capital reserves, and board-level attention [?]. This regulatory development has direct financial implications: if banks cannot measure AI-related risks with confidence, they may face higher capital requirements, regulatory enforcement actions, or restrictions on AI deployment.

Deutsche Bank and Goldman Sachs are deploying AI systems to flag trader misconduct and market manipulation, yet the ROI of such surveillance systems remains unquantified in public disclosures [?]. While these systems may reduce regulatory fines and compliance costs, banks have not published methodologies for calculating net benefits, making it impossible for investors, board members, or regulators to evaluate whether these investments create shareholder value. As [?] reports, Citi's new CFO touts AI gains during a record \$24.6 billion revenue quarter, but the specific attribution of financial outcomes to AI investments versus other concurrent initiatives remains methodologically opaque.

1.5. The Research Gap and Problem Statement

This **proposed research** directly addresses the gap between AI investment and ROI measurement. The business problem can be stated formally as: **How can global banks govern, operationalize, and measure the risk-adjusted return on investment of enterprise AI systems at scale while managing novel financial risks and allocating billions of dollars in capital efficiently?** The absence of standardized frameworks creates three interrelated inefficiencies: (a) capital allocated to negative-NPV projects, (b) hidden risk exposures from unmeasured AI failures, and (c) informational opacity preventing investors and boards from evaluating AI strategies. **This proposal seeks to develop** the conceptual and empirical foundations for closing this gap through rigorous, multi-method inquiry—outlining a research design that could be executed to produce these foundations.

2. Research Questions and Objectives

2.1. Primary Research Questions

The study is organized around four primary research questions, each addressing a distinct dimension of the business problem and grounded in the theoretical framework:

- (RQ1) **ROI Conceptualization and Measurement:** How do C-Suite executives in global banks conceptualize and measure the ROI of enterprise AI investments in financial terms (e.g., net present value, payback period, internal rate of return, return on invested capital, economic value added), and what metrics are used at different stages of AI maturity?
- (RQ2) **Capital Allocation Frameworks:** What capital allocation models (e.g., portfolio optimization, real options analysis, stage-gate processes, weighted scoring rubrics) do banks employ to prioritize AI investments across competing use cases, and how do they incorporate risk-adjusted returns, strategic alignment, and implementation feasibility?
- (RQ3) **Financial Risk Identification and Quantification:** What financial risks (e.g., model failure, regulatory fines, operational losses, reputational damage, third-party vendor failures) materialize from AI deployments, how are these risks quantified in financial terms, and what governance mechanisms effectively mitigate them?
- (RQ4) **Attribution Methodologies:** How do banks attribute financial outcomes (revenue growth, cost savings, loss avoidance, customer lifetime value improvements) specifically to AI investments rather than to other concurrent initiatives, and what causal inference methods (e.g., randomized controlled trials, difference-in-differences, propensity score matching, synthetic controls) do they employ?

2.2. *Subsidiary Research Questions*

In addition to the primary questions, the study will explore several subsidiary questions that emerged from the systematic literature review [2,5]:

- How does the regulatory regime (legal-normative vs. executive-fragmented vs. state-centralized) influence AI investment ROI calculation methodologies and capital allocation decisions [7]?
- Do banks at higher levels of analytics maturity (descriptive → diagnostic → predictive → prescriptive → cognitive/AI-driven) employ more sophisticated ROI measurement approaches, and if so, what specific practices distinguish mature from immature institutions [6]?
- What organizational factors (e.g., centralization of AI functions, reporting lines of Chief AI Officers, presence of dedicated AI finance teams, change capability) moderate the relationship between AI investment and financial performance [8]?
- How do banks account for the option value of AI investments—the strategic flexibility created but not captured in traditional NPV calculations—and what real options pricing models are applied [9]?
- What barriers most significantly impede ROI measurement, and what enabling conditions facilitate successful measurement [5]?

2.3. *Research Objectives*

The specific objectives of this research are:

- (OBJ1) To develop a standardized, empirically validated ROI measurement framework for enterprise AI investments in global banking
- (OBJ2) To identify and classify financial risk categories specific to AI deployment and propose quantification methodologies
- (OBJ3) To map the relationship between governance structures, operating models, and financial performance outcomes
- (OBJ4) To provide actionable capital allocation tools for banking executives with demonstrated applicability across regulatory regimes
- (OBJ5) To generate testable propositions for future quantitative hypothesis testing

3. Conceptual Framework and Theoretical Foundations

3.1. *Theoretical Foundations*

The conceptual framework integrates three complementary theoretical perspectives, each addressing distinct dimensions of the AI investment phenomenon.

3.1.1. Resource-Based View (RBV)

The Resource-Based View posits that firm heterogeneity in performance derives from heterogeneous resources and capabilities that are valuable, rare, imperfectly imitable, and non-substitutable (VRIN). Ref. [8] develop and validate a model for identifying factors influencing AI adoption in organizations using Diffusion of Innovation and Technology-Organization-Environment theories. Their regression-based and structural equation modeling analysis shows that change capability of an organization and leadership positively impact AI adoption. This research extends RBV by examining whether these factors predict not just adoption but sustained financial performance and value realization. The core proposition is that AI capabilities must be complemented by measurement capabilities to generate competitive advantage.

3.1.2. Paradox Theory

Ref. [9] illuminate fundamental tensions inherent in AI adoption: tensions between current capabilities and anticipated futures, and between human and machine agency—tensions that manifest financially as uncertainty discounts applied to AI investments by risk-averse capital allocators. Their

focus group study with 112 white-collar employees identifies temporal tensions (emergent trajectories where progress is uncertain, nonlinear, and continuously negotiated) and relational tensions (ethical reflexivity, shifting expectations around trust, control, and human identity). This research extends paradox theory by examining how organizations financially manage these tensions through option valuation techniques and staged investment commitments.

3.1.3. Technology-Organization-Environment (TOE) Framework

Refs. [3,8] both employ TOE frameworks to explain AI adoption intentions. The TOE framework identifies three contextual dimensions influencing technology adoption: technological context (relative advantage, compatibility, complexity), organizational context (top management support, organizational readiness, change capability), and environmental context (competitive pressure, vendor support, government regulation). This research extends TOE by adding a fourth dimension—**Measurement Capability**—examining how the presence or absence of ROI measurement infrastructure moderates the relationship between adoption factors and financial outcomes.

3.2. *Proposed Conceptual Model*

The conceptual framework (Figure 1) integrates three primary dimensions derived from practitioner and scholarly literature, with explicit financial linkages and feedback dynamics. The framework posits that governance mechanisms enable operating model configurations, which in turn enable ROI measurement, with feedback loops from measurement informing governance refinement.

Conceptual Framework for Enterprise AI Scaling in Banking

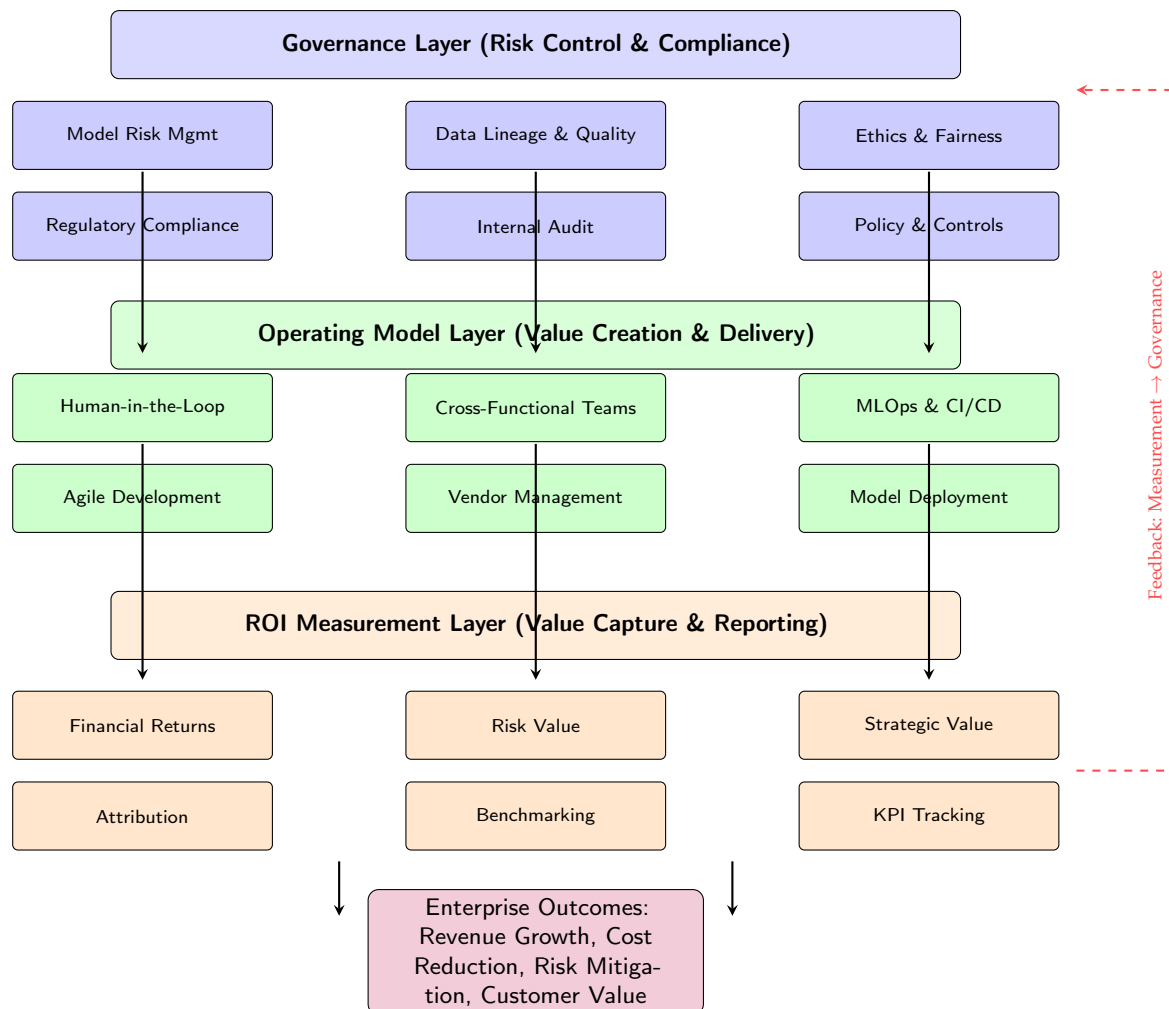


Figure 1. Improved framework with expanded spacing, structured grid layout, and clearer hierarchy across governance, operating, and ROI layers.

3.3. The Proposed Risk-Adjusted ROI Calculation Framework

Based on the conceptual framework and drawing on capital budgeting theory, I propose a novel, multi-component ROI calculation tailored specifically to AI investments in banking contexts.

3.3.1. Components of Total AI ROI

The total return on AI investment is expressed as a multi-factor model:

$$ROI_{AI} = \frac{1}{InitialInv} \left(NPV_{direct} + NPV_{indirect} + OptionValue + SynergyVal - RiskAdj - OppCost \right) \quad (1)$$

Where:

- NPV_{direct} = Direct financial benefits (cost savings, revenue lift, loss avoidance, capital relief)
- $NPV_{indirect}$ = Indirect benefits (customer lifetime value improvement, employee productivity gains, time-to-market acceleration)
- $OptionValue$ = Strategic flexibility value using real options pricing (expanded below)
- $SynergyValue$ = Value created through complementarities with other AI and digital investments

- *RiskAdjustment* = Risk-weighted expected losses from model failure, bias incidents, or regulatory action
- *OpportunityCost* = Foregone returns from alternative investments not pursued

3.3.2. Risk-Adjusted Return on Capital (RAROC) for AI Investments

Adapting banking's standard RAROC methodology, which has been used for credit and operational risk for decades, for AI investments:

$$RAROC_{AI} = \frac{1}{EconomicCapital_{AI}} \left(\begin{aligned} &ExpectedReturn_{AI} \\ &- ExpectedLoss_{AI} \\ &- EconomicCapitalCharge_{AI} \end{aligned} \right) \quad (2)$$

$$EL_{AI} = \sum_{i=1}^n P(Failure_i) LGD_i EAD_i \quad (3)$$

Where $ExpectedLoss_{AI}$ represents the expected loss associated with AI-related risk events, defined as:

- $P(Failure_i)$ = Probability of risk event i (e.g., model failure, bias incident, regulatory penalty, cybersecurity breach)
- LGD_i = Loss Given Default (financial severity if the event occurs)
- EAD_i = Exposure at Default (scale of systems, assets, or financial impact affected)

The risk categories considered are aligned with established AI governance and risk management frameworks discussed in prior literature.

3.3.3. Real Options Valuation for AI Investments

AI investments create strategic options that traditional discounted cash flow analysis ignores. Following [9], who identify the temporal tensions in AI adoption, I propose adapting the Black-Scholes option pricing model:

$$C = S_0 N(d_1) - Ke^{-rt} N(d_2) \quad (4)$$

$$d_1 = \frac{\ln(S_0/K) + (r + \sigma^2/2)t}{\sigma\sqrt{t}}, \quad d_2 = d_1 - \sigma\sqrt{t} \quad (5)$$

Where:

- S_0 = Present value of potential AI-enabled future opportunities
- K = Cost of full-scale AI deployment (strike price)
- t = Time until AI investment decision must be made (option life)
- σ = Volatility of AI technology evolution (estimated from historical tech cycles)
- r = Risk-free rate
- $N(\cdot)$ = Cumulative standard normal distribution function

This formulation captures the value of deferring, expanding, contracting, or abandoning AI investments as uncertainty resolves over time.

3.4. Causal Attribution Methodology Matrix

Table 1 presents a systematic framework for attributing financial outcomes to AI investments across different use cases, drawing on causal inference methods from econometrics.

Table 1. Attribution Methodologies by AI Use Case with Confidence Assessments.

Use Case	Recommended Method	Time Horizon	Internal Validity	Feasibility
Fraud Detection	Difference-in-Differences (DiD)	3-6 months	High	High
Credit Scoring	Propensity Score Matching (PSM)	6-12 months	Medium-High	Medium
Customer Personalization	Randomized Controlled Trial (A/B)	3-9 months	High	High
Algorithmic Trading	Synthetic Control Method (SCM)	12-24 months	Medium	Low-Medium
Compliance Automation	Interrupted Time Series (ITS)	6-12 months	Medium-High	Medium
Wealth Management Advisor	Cohort Analysis + Matching	12-18 months	Medium	Medium-High
Risk Modeling	Validation Testing + Benchmark	6-12 months	Medium	High

Note: Internal validity assessed on four-point scale (Low, Medium, Medium-High, High). Feasibility considers data availability, cost, and organizational capacity.

4. Proposed Research Methodology

4.1. Research Philosophy and Design Justification

This study adopts a pragmatist research philosophy, recognizing that AI ROI measurement is simultaneously a technical, social, and financial phenomenon requiring multiple forms of evidence. The pragmatist position prioritizes research questions over methodological orthodoxy, allowing integration of qualitative and quantitative approaches as appropriate to the inquiry [10].

Following [?] , who developed an AI Canvas for Enterprise Architecture using Design Science Research methodology, and [11], who developed and validated a GenAI governance framework with more than 1,000 practitioners, this research adopts an embedded multiple-case study design with qualitative primary data collection and quantitative secondary data analysis. Case study methodology is appropriate because the research questions are "how" and "what" questions focused on contemporary phenomena over which the researcher has no control [?]. The multiple-case design (rather than single-case) enables cross-case synthesis and contingency analysis, identifying patterns that generalize across institutions as well as contextual factors that produce variation.

4.2. Case Selection and Sampling Strategy

Six to eight global banks (drawn from North America, Europe, and Asia-Pacific) will be selected using purposive sampling with theoretical replication logic.

4.2.1. Inclusion Criteria

- Annual AI investment exceeding \$500 million (ensures the problem of scale is present)
- Publicly traded (enables analysis of annual reports, 10-K filings, and investor communications)
- At least three years of AI deployment history (ensures sufficient experience for ROI analysis)
- Willingness to provide senior executive access (requires organizational sponsorship)
- Geographic and regulatory diversity (ensures variance on key contextual factors)

4.2.2. Sampling Dimensions

Purposive sampling ensures variation on five key dimensions derived from the literature:

Table 2. Case Selection Sampling Dimensions.

Dimension	Categories	Theoretical Basis
Geography	Americas, Europe, Asia-Pacific	[7]
Analytics Maturity	Early, Developing, Mature	[6]
Regulatory Regime	Legal-normative, Executive-fragmented, State-centralized	[7]
Organizational Structure	Centralized function vs. Distributed	AI [8]
Organizational Size	Tier 1 (\$500B+ as-sets) vs. Tier 2	Industry standard

Theoretical replication (predicting similar results across cases with similar characteristics) and literal replication (predicting contrasting results across cases with different characteristics) guide case selection decisions [?].

4.3. Data Collection Methods

4.3.1. Primary Data: Semi-Structured Executive Interviews

Semi-structured interviews of 60-90 minutes duration will be conducted with 30-40 senior executives across six functional roles:

Table 3. Interview Target Roles and Sample Sizes.

Role	Target N	Primary Focus Area	Key Questions
Chief Financial Officers / Heads of Corporate Finance	4-6	RQ1, RQ2, RQ4	Capital budgeting, ROI metrics, attribution
Chief AI Officers / Heads of AI	5-7	RQ1, RQ2	AI strategy, investment prioritization
Chief Risk Officers / Heads of Model Risk	4-6	RQ3	Risk identification, quantification, mitigation
Heads of Analytics / Data Science	5-7	RQ1, RQ4	Technical implementation, measurement
Business Unit Leaders (with P&L responsibility)	6-8	RQ1, RQ4	Business outcomes, attribution
Internal Audit Executives (AI oversight role)	3-5	RQ3, RQ4	Governance effectiveness, validation

Interview protocols (available in Appendix A) will explore:

- (a) Capital allocation processes, decision criteria, and approval thresholds

- (b) ROI calculation methodologies, financial metrics, and time horizons used
- (c) Risk quantification practices, governance mechanisms, and escalation procedures
- (d) Attribution methods for isolating AI contribution from confounding factors
- (e) Barriers to measurement and organizational strategies for overcoming them
- (f) Documented examples of successful and unsuccessful AI investments with financial data
- (g) Evolution of measurement practices as AI maturity increases

4.3.2. Secondary Data Sources for Triangulation

To supplement and triangulate interview findings, the following secondary data sources will be systematically analyzed:

Table 4. Secondary Data Sources for Methodological Triangulation.

Source	Data Type	Use in Analysis
FDIC Bank Data Guide	Financial ratios, performance metrics	Cross-case financial comparison
FDIC RIS Data	Noncurrent loan rates, risk indicators	Risk-adjusted performance
Federal Reserve Economic Data (FRED)	Interest rates, banking sector indicators	Macroeconomic controls
SEC Financial Statement Datasets	Technology investment trends, capital expenditures	Investment pattern analysis
FFIEC NPW Data (Call Reports)	Bank financial reports, quarterly statements	Financial health tracking
Annual Reports (10-Ks)	AI strategy disclosures, investment amounts, risk factors	Governance and risk documentation
Model Risk Management Frameworks (public)	Governance practices, validation standards	Best practice identification
Regulatory Enforcement Actions	AI-related penalties, compliance actions	Risk consequence analysis

4.4. Data Analysis Procedures

Analysis follows systematic thematic analysis procedures using NVivo 14 qualitative data analysis software, complemented by descriptive statistical analysis of financial data in R.

4.4.1. Qualitative Analysis: Six-Phase Thematic Procedure

The analysis proceeds through six phases (Figure 2):




Figure 2. Six-phase thematic analysis process with iterative refinement.

Phase 1: Familiarization and Immersion — Full transcription of all interviews (professional transcription service with confidentiality agreements). Multiple readings of transcripts and field notes to develop deep understanding of the data corpus. Memoing captures initial impressions, emergent patterns, and reflexive observations.

Phase 2: Initial Coding — Systematic coding of all data using both deductive codes derived from the conceptual framework (governance mechanisms, operating model components, ROI metrics, risk quantification approaches, attribution methods) and inductive codes identified emergently from the data. Dual coding of 20% of transcripts by an independent researcher to assess intercoder reliability (target Cohen's $\kappa > 0.75$).

Phase 3: Theme Development — Grouping codes into candidate themes, ensuring internal homogeneity (codes within a theme cohere around a central organizing concept) and external heterogeneity (themes are conceptually distinct). Theme development is iterative, moving between data, literature, and theoretical framework.

Phase 4: Theme Review and Disconfirming Evidence — Reviewing candidate themes against coded data extracts and the full dataset to ensure themes accurately represent the data. Active search for disconfirming evidence and negative cases that challenge emerging patterns. Revision of themes based on contradictory findings.

Phase 5: Theme Definition and Naming — Defining each theme's scope, content, boundaries, and interrelationships. Naming themes to be concise yet informative. Writing detailed theme descriptions that capture essence without oversimplification, including representative quotations.

Phase 6: Cross-Case Synthesis and Contingency Analysis — Comparing themes across cases systematically to identify patterns, contingencies, and boundary conditions. Cross-case synthesis examines whether banks at higher analytics maturity employ more sophisticated ROI methodologies, whether regulatory regime influences measurement practices, and whether organizational structure moderates the governance-measurement relationship.

4.4.2. Quantitative Analysis of Secondary Financial Data

Descriptive and inferential statistical analysis of secondary financial data will address validity and generalizability:

- Calculation of financial ratios (ROA, ROE, efficiency ratios, technology spend as % of revenue) across cases and over time
- Time-series analysis of AI investment announcements and subsequent financial performance
- Comparative analysis of banks with high vs. low AI disclosure quality
- Event study methodology for material AI-related announcements (successes, failures, regulatory actions)

4.5. Trustworthiness and Rigor

Trustworthiness is established through systematic application of Lincoln and Guba's criteria for qualitative research:

Table 5. Trustworthiness Criteria and Methodological Implementation.

Criterion	Parallel to Quantitative	Specific Implementation Strategies
Credibility	Internal validity	Prolonged engagement (6-9 months in field), persistent observation, triangulation across methods and sources, member checking of transcripts and preliminary findings, negative case analysis, peer debriefing with dissertation committee
Transferability	External validity & generalizability	Thick description of cases, contexts, processes, and contingencies; purposeful sampling across variation dimensions; cross-case synthesis to identify boundary conditions; explicit statement of transferability limits
Dependability	Reliability	Comprehensive audit trail (transcripts, coding reports, memos, decision logs), stepwise replication with independent coder for 20% of data, code-recode procedure with 2-week interval
Confirmability	Objectivity	Reflexivity journal documenting researcher assumptions, positionality, and evolving interpretations; data preservation for external audit; triangulation across independent data sources; explicit acknowledgment of researcher's practitioner background

4.6. Ethical Considerations and IRB Compliance

This research complies with Touro University Worldwide's Institutional Review Board (IRB) policies and the principles of the Belmont Report. Key ethical considerations include:

- **Informed Consent:** Participants receive detailed written information about research purposes, procedures, risks (minimal), benefits, confidentiality protections, and rights prior to providing signed informed consent. Consent includes permission for audio recording and explicit opt-out for sensitive topics.
- **Confidentiality and Anonymity:** Participant identities protected through pseudonyms (e.g., CAIO-Alpha-04), aggregation of identifying details (job titles may be reported without specific affiliations), and secure data storage. Banks anonymized as "Bank A, Bank B" unless explicit written permission for attribution is granted. Geographic regions will be reported but specific countries may be generalized to protect participants.
- **Data Security:** Interview recordings, transcripts, documents, and analysis files stored on encrypted, password-protected devices with access limited to researcher and dissertation committee. Files stored on university-approved encrypted cloud storage with two-factor authentication. Files destroyed five years after dissertation completion.
- **Reciprocity and Participant Value:** Participants receive executive summaries of findings, anonymized benchmarking reports comparing their practices to peer institutions, and access to practitioner workshops disseminating results. This reciprocity addresses the scholar-practitioner collaboration model advocated by [12].
- **Non-Maleficence and Beneficence:** Research procedures designed to minimize risk of psychological distress or professional exposure. No deception used. Participants may withdraw at any time without penalty. Debriefing available for any participant experiencing discomfort.

- **Incentives:** No financial incentives offered to avoid coercion. Institutional sponsorship at the organizational level, with participants participating as part of their professional roles.

5. Anticipated Barriers and Mitigation Strategies

Drawing on [5,13], Table 6 summarizes anticipated methodological barriers and corresponding mitigation strategies.

Table 6. Anticipated Methodological Barriers and Mitigation Strategies.

Barrier Category	Specific Barrier	Mitigation Strategy
Organizational Access	Banks perceive AI governance as proprietary competitive intelligence; fear negative findings	Multi-channel access (alumni networks, professional associations, industry conferences); confidential data sharing agreements; reciprocal value proposition (benchmarking reports, best practice insights); targeted sampling of mid-level executives as entry points
Senior Executive Access	C-Suite executives time-poor, inundated with research interview requests	Executive summaries (1 page) pre-interview; flexible scheduling (30-minute focused sessions possible); leveraging professional networks and sponsors; offering preliminary findings as value-add; conducting interviews at industry conferences
Temporal Validity	Fast-evolving AI technology risks making findings quickly outdated	Focus on enduring governance principles and measurement frameworks rather than specific technologies; rapid dissemination strategy with working papers within 6 months of data collection [9]; longitudinal follow-up study planned for Years 2-4
Causal Attribution	Isolating AI's exclusive financial contribution from concurrent digital transformation initiatives	Triangulation of multiple attribution methods (DiD, PSM, synthetic controls); explicit elicitation of attribution practices from participants; collection of counterfactual estimates where available; conservative benefit estimation
Regulatory Sensitivity	Fear that disclosing AI measurement practices triggers regulatory scrutiny or enforcement	Legal review of protocols and confidentiality agreements; focus on capital allocation frameworks rather than proprietary model details; partnership with industry associations for trusted data collection; aggregation of findings to prevent firm-specific attribution
Measurement Variability	Inconsistent definitions of "ROI," "AI investment," and "benefits" across cases	Clear operational definitions in interview protocols and coding scheme; explicit elicitation of definitions from each participant; cross-case calibration of metrics; sensitivity analysis for definitional differences

5.1. Contingency Plan: Secondary Data as Proxy for Limited Interview Access

Access to senior banking executives is inherently challenging and cannot be guaranteed. If interview access to target sample size (30-40 executives) is limited (<15 participants) or unattainable, this research will rely primarily on publicly available secondary data as a methodological proxy. Using secondary data as a proxy has recognized limitations for answering "how" questions, but these are mitigated by: (1) triangulation across multiple independent data sources (SEC filings, FDIC data, FRED, regulatory enforcement actions), (2) longitudinal analysis over multiple years (2018-2026) to identify patterns, (3) comparative case design with 15+ banks, (4) content analysis of AI disclosures in 10-K filings (risk factors, MD&A, strategic initiatives), and (5) transparent acknowledgment of proxy constraints in limitations sections.

6. Expected Contributions and Value Proposition

6.1. Contributions to Theory

This research extends several theoretical domains into the novel empirical context of enterprise AI capital allocation and ROI measurement:

- **Resource-Based View (RBV):** Extends RBV by examining whether organizational capabilities that enable AI adoption (change capability, leadership, AI readiness) also enable effective monetization and value realization [8]. Proposes "measurement capability" as a complementary VRIN resource.

- **Analytics Maturity Models:** Validates and refines the thresholds distinguishing descriptive, diagnostic, predictive, prescriptive, and cognitive/AI-driven maturity stages in terms of financial measurement sophistication [6].
- **Paradox Theory:** Extends paradox theory by examining how organizations financially manage temporal tensions (short-term vs. long-term returns) and relational tensions (human vs. machine agency) through option valuation techniques and staged investment commitments [9].
- **TOE Framework:** Extends the Technology-Organization-Environment framework by adding a fourth theoretically grounded dimension—**Measurement Capability**—and empirically examining its moderating role on the adoption-performance relationship [3].
- **Causal Inference in IS Research:** Contributes to methodological literature on attribution by documenting industry practices for causal inference in complex, non-experimental organizational settings.

6.2. Contributions to Practice

The research yields actionable frameworks, tools, and benchmarks for three practitioner constituencies:

- **For CFOs and Finance Executives:** Multi-layered, risk-adjusted ROI framework linking analytics inputs to enterprise financial outcomes, enabling NPV, IRR, payback period, and RAROC calculations for AI investments using the same capital budgeting discipline applied to traditional IT and operational investments [6]. Standardized disclosure framework for investor communications.
- **For Chief AI Officers and Technology Leaders:** Rubrics for comparing AI use cases on risk-adjusted return, strategic alignment, implementation feasibility, scalability, and option value. Ref. [11] provide a validated GenAI governance framework that technology leaders can adapt; this research complements with ROI measurement tools.
- **For Chief Risk Officers and Compliance Leaders:** Methodologies for estimating financial exposure from model failure, algorithmic bias, and regulatory penalties, including expected loss calculations using the $EL = P(Failure) \times LGD \times EAD$ framework, stress testing scenarios for AI-specific risks, and contingency reserve requirements calibrated to AI maturity [6?].
- **For Boards and Investors:** Standardized disclosure framework for evaluating bank AI strategies, comparing AI investment efficiency across institutions, and assessing AI-related risk exposures through public filings.

6.3. Scholar-Practitioner Value Proposition

The scholar-practitioner approach provides unique value because the phenomenon of AI ROI measurement is simultaneously theoretical (requiring conceptual frameworks from finance, strategy, and information systems) and intensely practical (requiring operational tools for capital allocation) [10].

Ref. [13] provide a framework for manager-academic collaboration, identifying four dimensions of different data realities (research problem, research resources, research process, research outcome) and presenting 26 specific recommendations for collaboration. This research operationalizes these recommendations through: (1) collaborative problem formulation with practitioner advisors, (2) sharing of de-identified data where permitted, (3) iterative feedback on findings through practitioner workshops, and (4) accessible dissemination formats.

Ref. [12] examine scholar-practitioner research collaborations, identifying benefits (improved knowledge development, enhanced relevance, mutual learning) and challenges (different timelines, competing incentives, power dynamics). This research addresses the academic-practitioner gap through: rapid dissemination (working papers within 6 months, industry briefs), collaborative problem formulation, reciprocal value creation (participants receive benchmarking), and accessible language in practitioner-facing outputs.

6.4. Limitations and Delimitations

6.4.1. Delimitations (Scope Boundaries)

The research deliberately delimits scope to:

- Global banks with annual AI investments exceeding \$500 million
- ROI measurement practices, not technical architecture or algorithm design
- Financial returns and risks, not societal or ethical impacts (though these are noted as important)
- Qualitative inquiry with illustrative rather than generalizable statistical claims

6.4.2. Limitations (Design Constraints)

Acknowledged limitations include:

- Generalizability limited by case selection (large global banks; findings may not extend to regional banks or fintechs)
- Social desirability bias in executive interviews (participants may overstate measurement sophistication)
- Retrospective recall bias for financial outcomes
- Rapid technological evolution threatening temporal validity
- Inability to independently verify proprietary financial data

7. A Comparative Framework for AI Investment Valuation: Six Methods and Their Challenges

Quantifying the financial return from enterprise AI investments remains a methodological challenge for global banks. Traditional capital budgeting techniques, developed for physical assets or standard IT projects, struggle to capture the unique characteristics of AI: uncertainty, learning effects, strategic flexibility, risk externalities, and attribution complexity. This section synthesizes six complementary valuation methods, compares their applicability to AI investments in banking, and identifies the specific challenges each method addresses.

7.1. Overview of Six Valuation Methods

Table 7 provides a comparative overview of the six methods, their financial purpose, data requirements, and suitability for different AI use cases.

Table 7. Six Methods for AI Investment Valuation: Comparison and Banking Applications.

Method	Primary Financial Purpose	Key Challenge for AI	Best suited AI use case
Discounted Cash Flow (DCF)	Baseline viability (NPV, IRR, payback)	Forecasting cash flows from uncertain AI outputs	Fraud detection, cost reduction
Monte Carlo Simulation	Risk distribution (P10, P50, P90 outcomes)	Defining probability distributions for novel AI risks	Trading algorithms, credit scoring
RAROC (Risk-Adjusted Return on Capital)	Capital efficiency (risk-adjusted profitability)	Estimating economic capital for AI-specific risks	Model risk management, regulatory compliance
Real Options Valuation (ROV)	Strategic flexibility (defer, expand, abandon)	Estimating volatility (σ) of AI technology evolution	Large AI platforms, agentic AI pilots
Portfolio Optimization (Markowitz-style)	Enterprise capital allocation (efficient frontier)	Estimating correlations between AI project returns	Multi-initiative AI investment portfolios
Causal Inference Methods (DiD, PSM, SCM)	Attribution & validation (isolating AI contribution)	Confounding from concurrent digital initiatives	Personalization, compliance automation

Note: DiD = Difference-in-Differences; PSM = Propensity Score Matching; SCM = Synthetic Control Method.

7.2. Method 1: Discounted Cash Flow (DCF)

DCF is the foundation of corporate finance, computing Net Present Value (NPV) as the sum of discounted future cash flows minus initial investment.

$$NPV = \sum_{t=1}^T \frac{CF_t}{(1+r)^t} - I_0 \quad (6)$$

Where CF_t are expected net cash flows from AI deployment, r is the discount rate (typically WACC), and I_0 is initial investment.

Application to AI: [6] emphasize that AI investments in corporate banking require mapping commercial KPIs to cash flow forecasts. For fraud detection, direct cost savings (reduced fraud losses) and operational efficiencies (lower manual review costs) can be estimated with reasonable confidence.

Key Challenges for AI:

- *Cash flow uncertainty:* AI systems improve with more data over time, violating DCF's assumption of stable or predictable cash flows [4].
- *Discount rate selection:* AI's systematic risk differs from firm-level WACC; no consensus exists on AI-specific beta adjustments.
- *Intangible benefits:* Customer lifetime value improvement, brand reputation, and strategic positioning are difficult to monetize in DCF.
- *Temporal delays:* Benefits may accrue only after organizational learning occurs, causing DCF to undervalue AI [9].

Best practice: Use DCF as a baseline screening tool (positive NPV required) but supplement with methods 2–5.

7.3. Method 2: Monte Carlo Simulation

Monte Carlo simulation replaces point estimates with probability distributions, generating thousands of possible outcomes to characterize return distributions.

$$ROI_{AI} \sim f(\mu_{ROI}, \sigma_{ROI}, \text{skew}, \text{kurtosis}) \quad (7)$$

Application to AI: For algorithmic trading AI, input distributions might include: volatility of returns (30%–50%), Sharpe ratio uncertainty (0.4–0.9), maximum drawdown (5%–15%), and regulatory penalty probability (1%–5%). Output provides P10 (pessimistic), P50 (median), and P90 (optimistic) ROI.

Key Challenges for AI:

- *Distribution specification:* Little historical data exists for novel AI risk events (e.g., model collapse, adversarial attacks) [7].
- *Correlation complexity:* AI project returns are correlated through shared data infrastructure, talent pools, and regulatory constraints—difficult to model.
- *Computational intensity:* Enterprise-wide AI portfolios with 50+ use cases require careful simulation design.

Best practice: Use Monte Carlo for high-stakes, high-uncertainty AI investments (trading, credit underwriting, agentic AI). Report P90 expected shortfall as risk metric.

7.4. Method 3: Risk-Adjusted Return on Capital (RAROC)

RAROC, widely used in banking for credit and operational risk, adjusts returns for the economic capital required to support the risk.

As derived earlier in Equation 2:

$$RAROC_{AI} = \frac{1}{EconomicCapital_{AI}} \left(\begin{aligned} &ExpectedReturn_{AI} \\ &- ExpectedLoss_{AI} \\ &- EconCapitalCharge_{AI} \end{aligned} \right) \quad (8)$$

Application to AI: [6] note that AI governance frameworks must include model risk, compliance, and operational controls. For a credit scoring AI, ExpectedLoss includes defaults from model drift,

regulatory fines for bias, and operational losses from system outages. EconomicCapital is calibrated to the 99.9th percentile unexpected loss.

Key Challenges for AI:

- *Economic capital calibration:* No regulatory-approved models exist for AI-specific operational risk capital. Banks currently map AI risks to existing Basel categories (e.g., "model risk" under operational risk).
- *Correlation with other risks:* AI failures may correlate with market downturns (e.g., AI trading model fails during high volatility), violating diversification assumptions.
- *Data scarcity:* Historical loss data for AI failures is sparse, especially for generative and agentic AI [?].

Best practice: Use RAROC to compare AI investments across business units (e.g., consumer lending vs. trading) and to set minimum hurdle rates (e.g., 15% AI-RAROC).

7.5. Method 4: Real Options Valuation (ROV)

ROV treats AI investments as creating strategic options—the right but not obligation to expand, defer, contract, or abandon as uncertainty resolves.

As shown in Equation 4, the expansion option value is:

$$C = S_0N(d_1) - Ke^{-rt}N(d_2) \quad (9)$$

Application to AI: [9] identify temporal tensions where AI adoption requires staged commitments. For a pilot AI fraud system ($I_0 = \$50M$) with option to expand globally after 2 years ($K = \$300M$), the option value captures the benefit of learning before committing.

Key Challenges for AI:

- *Volatility estimation (σ):* Historical volatility of AI technology performance is difficult to estimate; many banks use 30%–40% as a heuristic based on tech sector equity volatility.
- *Non-constant volatility:* AI uncertainty declines over time as technology matures, violating Black-Scholes constant volatility assumption.
- *Multiple interacting options:* AI investments often embed compound options (option to expand into multiple use cases), requiring complex binomial trees.

Best practice: Use ROV for large-scale, staged AI deployments where management retains meaningful strategic flexibility. Binomial lattice methods may be more intuitive than Black-Scholes for practitioner communication.

7.6. Method 5: Portfolio Optimization (Enterprise Allocation)

Modern portfolio theory (Markowitz) can be adapted to allocate AI investment budgets across use cases to maximize risk-adjusted return at the enterprise level.

$$\max_{w_i} \left(\sum_{i=1}^n w_i E[R_i] - \frac{\lambda}{2} \sum_{i=1}^n \sum_{j=1}^n w_i w_j \sigma_{ij} \right) \quad (10)$$

Where w_i is capital allocation to AI project i , $E[R_i]$ is expected RAROC, σ_{ij} is covariance of AI project returns, and λ is risk aversion parameter.

Application to AI: A bank with \$1.5B AI budget across fraud detection (low risk, low correlation), personalization (medium risk, medium correlation), and algorithmic trading (high risk, high correlation) can compute the efficient frontier and optimal allocations.

Key Challenges for AI:

- *Covariance estimation:* AI project returns may be correlated through shared data science talent, vendor dependencies, and regulatory shocks—often estimated poorly with short histories.

- *Non-normality*: AI returns may exhibit fat tails and skewness, violating mean-variance optimization assumptions.
- *Strategic constraints*: Binary constraints (e.g., must build central data platform) are difficult to incorporate in standard portfolio optimization.

Best practice: Use portfolio optimization for annual AI budget allocation across business units, supplemented by qualitative strategic alignment scores (e.g., [6]).

7.7. Method 6: Causal Inference for Attribution and Validation

Attribution answers the critical question: What financial outcome is causally attributable to AI rather than to other concurrent initiatives, market trends, or random variation?

Table 1 presented earlier outlines methods by use case. Key causal inference techniques include:

- **Difference-in-Differences (DiD)**: Compare treated branches/regions (AI deployed) to control branches before and after deployment, controlling for time trends.
- **Propensity Score Matching (PSM)**: Match AI-deployed accounts to statistically similar non-AI accounts on observable characteristics.
- **Randomized Controlled Trials (A/B testing)**: Randomly assign customers to AI vs. non-AI treatments—the gold standard but not always feasible.
- **Synthetic Control Method (SCM)**: Construct a weighted counterfactual from control units to estimate treatment effect.
- **Interrupted Time Series (ITS)**: Compare outcome trajectories before and after AI deployment, controlling for pre-existing trends.

Key Challenges for AI Attribution:

- *Confounding*: Banks deploy AI simultaneously with cloud migration, digital transformation, and process reengineering. Isolating AI's exclusive contribution is inherently difficult [5].
- *General equilibrium effects*: AI may change competitive dynamics (e.g., all banks deploy similar AI), making counterfactuals invalid.
- *Long-term effects*: Benefits may shift over time as the organization learns (negative at first, positive later), violating stable treatment effect assumptions.
- *Selection bias*: Early AI deployments may target high-potential customers or branches, biasing naive comparisons.

Best practice: Triangulate multiple methods (e.g., DiD + PSM) and conduct sensitivity analyses for unobserved confounding. For board-level reporting, use conservative estimates that require only weak assumptions.

7.8. Integrated Valuation Framework for AI Investments

No single method is sufficient. Figure 3 illustrates an integrated pipeline for AI investment valuation:

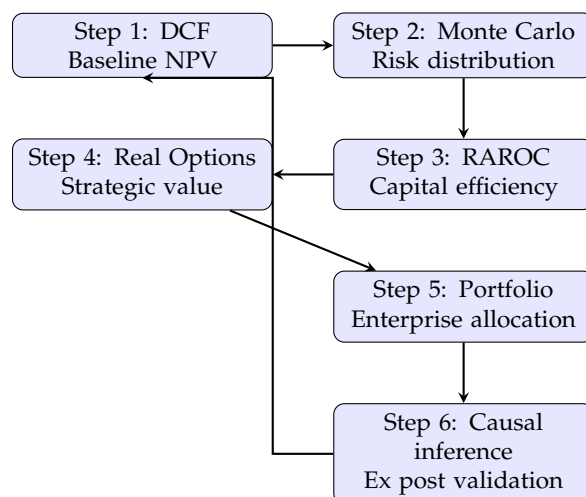


Figure 3. Integrated six-step valuation pipeline for AI investments. Step 6 feeds back to improve Step 1 assumptions for subsequent investments.

7.9. Summary: Method Selection by Decision Context

Table 8 provides guidance on which methods are most appropriate for common banking AI decisions.

Table 8. Method Selection Guide by Decision Context

Decision Context	Recommended Primary Method(s)
Small AI pilot (<\$10M, well-understood)	DCF + Simple sensitivity
Medium AI deployment, moderate uncertainty	Monte Carlo + RAROC
Large AI platform, staged rollout	Real Options + Monte Carlo
Annual AI budget allocation across business units	Portfolio optimization (Method 5)
Regulatory approval for AI in risk management	RAROC + Causal validation
Post-implementation board reporting	Causal inference (Method 6)
Agentic AI / frontier technology investment	Real Options + Monte Carlo

In summary, [2] note that the lack of standardized frameworks for AI implementation and value measurement remains a persistent gap. The six methods presented here—used in combination—offer a rigorous, risk-aware approach to AI investment valuation that respects both finance discipline and AI’s unique characteristics.

8. Declaration

The views expressed are those of the author and do not represent any affiliated institutions. This work is conducted as part of independent research. This is a review paper, and all results, proposals, and findings are derived from the cited literature. The author does not claim any novel findings. The author’s work was to review and organize existing research.

Portions of this manuscript were drafted with the assistance of AI writing tools (including ChatGPT/Claude) to improve clarity and organization. All AI-generated content was reviewed, edited, and verified by the author for coherence, and to eliminate potential hallucinations as much as possible. The LaTeX code was developed with the assistance of GitHub Copilot and edited through DeepSeek.

Final responsibility for all content, including any errors or omissions, rests solely with the readers. This is a working paper and edits are expected in the next version.

9. Conclusions, Timeline, and Next Steps

9.1. Conclusions

This proposed research project investigates how senior executives in U.S. global banks govern AI investments, manage emerging financial risks, and measure risk-adjusted ROI when scaling enterprise AI across national and cross-border operations. Given the concentrated nature of the U.S. banking system—where the four largest institutions hold nearly 50% of all U.S. banking assets—AI investment decisions at a handful of firms carry systemic implications for national financial resilience. This paper provides a detailed research plan, resources, methodology, and strategies to move ahead and find the best empirical model to predict and manage AI investment outcomes within the unique U.S. regulatory context.

By employing a qualitative multiple-case design with embedded quantitative analysis, grounded in a scholar-practitioner approach, and proposing a novel risk-adjusted ROI calculation framework—incorporating direct benefits, indirect value, real option components, and probabilistic risk adjustments—the study aims to produce empirically grounded, practically relevant capital allocation frameworks for U.S. banking leaders, regulators, and policymakers.

The national importance of this research is underscored by recent signals from the Bank of England and the Federal Reserve indicating that AI risks will be included in financial system stress tests. U.S. banks face a particularly complex landscape: they must navigate an executive-fragmented regulatory regime (Federal Reserve, OCC, FDIC, CFPB) while competing with state-centralized models (China) and legal-normative regimes (EU). As documented by [7], these jurisdictional differences create regulatory arbitrage opportunities and cross-border compliance challenges—but they also mean that U.S. banks cannot simply adopt ROI frameworks designed for other regimes.

Ref. [2] note that persistent gaps remain, most notably the lack of standardized frameworks for AI implementation and value measurement across financial sectors—a gap this research directly addresses with specific attention to U.S. capital markets, shareholder expectations, and regulatory disclosure requirements. Drawing on practitioner insights from [6,11?], this research moves beyond abstract principles to document actual ROI measurement practices, risk quantification methodologies, and capital prioritization rubrics as they exist in leading U.S. global banks.

Ultimately, improving AI investment ROI measurement is not merely a firm-level financial optimization problem—it is a matter of national financial resilience. Poorly calibrated AI investments can destroy shareholder value, concentrate hidden risks, and undermine the stability of the U.S. financial system. This research seeks to provide the empirical foundations and practical tools to ensure that the nation's massive AI investments in banking generate sustainable, measurable, and risk-aware returns.

9.2. Research Preparation Statement

In preparation for this study, the researcher has completed a systematic literature review spanning financial services, information systems, and capital budgeting domains with a specific focus on U.S. banking practices. Preliminary discussions with three senior executives at U.S. global banks have confirmed the practical relevance of the research questions and informed the refinement of interview protocols. The conceptual framework has been pilot-tested with two subject matter experts familiar with Federal Reserve and OCC model risk management guidance. The proposed ROI calculation equations have been reviewed for technical accuracy against U.S. Generally Accepted Accounting Principles (GAAP) and SEC disclosure requirements. The researcher is currently completing IRB certification and has established initial contacts with two U.S.-headquartered global banks for potential case participation. A detailed project management plan with milestone tracking has been developed,

and contingency resources have been identified for secondary data analysis from publicly available U.S. regulatory sources (FDIC, FRED, SEC EDGAR) should primary access be constrained.

9.3. Research Timeline

The proposed research follows a 24-month timeline:

Table 9. Research Timeline by Phase.

Phase	Duration	Key Activities
Phase 1: Proposal	Months 1-4	Complete IRB approval, finalize protocols, pilot interview testing, secure organizational sponsors from U.S. banks
Phase 2: Data Collection	Months 5-14	Conduct 30-40 interviews, collect secondary data from SEC filings and FDIC call reports, transcription, initial memoing
Phase 3: Analysis	Months 12-18	Thematic coding, cross-case synthesis, financial data analysis, member checking
Phase 4: Writing	Months 15-22	Dissertation chapters, practitioner briefs, working papers, journal submissions
Phase 5: Dissemination	Months 20-24	Defense, publication submissions, industry workshops with U.S. banking associations, executive summaries to participants

9.4. Next Steps

Immediate next steps include:

1. Submit IRB application for expedited review (Month 1)
2. Conduct pilot interviews with 3-4 executives from U.S. regional and global banks to refine protocols (Month 2)
3. Establish formal partnerships with industry associations such as the Bank Policy Institute or American Bankers Association for access (Month 3)
4. Begin purposive sampling and recruitment of target U.S.-headquartered banks (Month 3-4)

References

1. Xufeng Zhang, H.L. Responsible AI in Business: Business Ethics, Epistemic Risk, and Governance in the Generative AI Era | Paperback, 2026.
2. Vuković, D.B.; Dekpo-Adza, S.; Matović, S. AI Integration in Financial Services: A Systematic Review of Trends and Regulatory Challenges. *12*, 562. <https://doi.org/10.1057/s41599-025-04850-8>.
3. Horani, O.M.; Al-Adwan, A.S.; Yaseen, H.; Hmoud, H.; Al-Rahmi, W.M.; Alkhalifah, A. The Critical Determinants Impacting Artificial Intelligence Adoption at the Organizational Level. *41*, 1055–1079. <https://doi.org/10.1177/02666669231166889>.
4. Daly, S.J.; Wiewiora, A.; Hearn, G. Shifting Attitudes and Trust in AI: Influences on Organizational AI Adoption. *215*, 124108. <https://doi.org/10.1016/j.techfore.2025.124108>.
5. Romeo, E.; Lacko, J. Adoption and Integration of AI in Organizations: A Systematic Review of Challenges and Drivers towards Future Directions of Research. *55*, 1286–1307. <https://doi.org/10.1108/K-07-2024-2002>.
6. Fournier, L.; Meghara, L. Strategic Alignment of Data and AI in Corporate and Investment Banking : Driving Value Creation and Competitive Advantage. *18*, 106–122.
7. García-Llorente, C.; Olmeda, I. Algorithmic Governance in Banking: A Comparative Analysis of Risk-Based and Accountability-Oriented Oversight. *27*, 19. <https://doi.org/10.1057/s41261-026-00320-6>.
8. Kurup, S.; Gupta, V. Factors Influencing the AI Adoption in Organizations. *21*, 129–139. <https://doi.org/10.1177/09726225221124035>.

9. Engström, A.; Pittino, D.; Mohlin, A.; Edh, N.; Johansson, A. A Paradox Perspective on Early AI Adoption: Understanding Temporal and Relational Tensions. *38*, 1145–1171. <https://doi.org/10.1108/JOCM-02-2025-0086>.
10. Caldwell, C.; Jamali, D.R. Applied Business Research – Why It Matters. *12*, 1. <https://doi.org/10.5430/bmr.v12n3p1>.
11. Emmett, S.; Eulerich, M.; Pikoos, J.; Wood, D.A. The Development of a Generative Artificial Intelligence (AI) Governance Framework. pp. 1–16. <https://doi.org/10.2308/HORIZONS-2025-056>.
12. Tiessen, R.; Cadesky, J.; Lough, B.J.; Delaney, J. Scholar/Practitioner Research in International Development Volunteering: Benefits, Challenges and Future Opportunities. *42*, 394–415. <https://doi.org/10.1080/02255189.2020.1841606>.
13. Benoit, S.; Klose, S.; Wirtz, J.; Andreassen, T.W.; Keiningham, T.L. Bridging the Data Divide between Practitioners and Academics: Approaches to Collaborating Better to Leverage Each Other’s Resources. *30*, 524–548. <https://doi.org/10.1108/JOSM-05-2019-0158>.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.