A Multi-Variable Econometric Model of Cost Optimization for US Retail and Commercial Banking Institutions

Authors

Abey Litty, Daniel Klaus, Mia James

Date; November 24, 2026

Abstract

US retail and commercial banks face mounting pressure to optimize operating costs amid declining branch economics, rising technology expenditures, and intensifying competition from fintech entrants. Traditional cost-management approaches—static budgeting and single-equation stochastic frontier models—fail to capture the joint influence of fixed-cost structure, digital investment, regulatory burden, and scale economies on total operating expense. This study develops and validates a multi-variable econometric framework that integrates translog cost function estimation with machine-learning-based prediction to identify optimal cost configurations for US banks. Using Federal Reserve Call Report data from 2015–2024 for a stratified sample of 847 commercial banks, the study compares the proposed hybrid model against conventional cost-efficiency benchmarks. The hybrid model achieves 89.4% predictive accuracy in identifying banks that will outperform peer cost-efficiency medians in the subsequent year, a 14.2 percentage-point improvement over stochastic frontier baseline models. Feature importance analysis reveals that IT expenditure share, fixed-asset intensity, and branch-density ratios are the dominant predictors, collectively explaining 62.3% of cost-efficiency variance. The findings provide bank administrators with a replicable diagnostic framework and suggest that the optimal cost structure has shifted decisively toward technology-intensive configurations, with traditional branch-based models increasingly associated with efficiency penalties.

Keywords: bank cost optimization, econometric modeling, stochastic frontier analysis, fixed-cost structure, US commercial banking

1. Introduction

1.1 Background

The US banking industry has undergone a structural transformation in its cost architecture over the past two decades. The aggregate share of total operating expense attributable to physical fixed assets—branches, premises, and equipment—declined from approximately 17% in the early 1990s to 7% by the mid-2020s, while information technology and data-processing expenditures rose to approximately 10% of aggregate bank spending, with the average bank allocating 16% of its cost base to IT . This compositional shift has profound implications for how banks should conceptualize and manage cost optimization.

Concurrently, the academic literature on bank cost efficiency has evolved from neoclassical scale-and-scope estimation toward more granular examinations of fixed-versus-marginal cost tradeoffs. Federal Reserve research documents that fixed costs have become a dominant component of bank total cost, effectively increasing operational leverage and altering the relationship between competition, risk-taking, and lending efficiency . The conventional translog cost function, while still widely employed, assumes a degree of parameter stability that may no longer hold in an environment characterized by rapid technological substitution and regulatory change.

Recent methodological innovations offer promising avenues for refinement. Gupta, Puri, and Setia (2025) demonstrate that integrating directional distance function-based network data envelopment analysis with support vector regression substantially improves the precision of cost-efficiency forecasting in banking contexts . Similarly, blockchain-based infrastructure research suggests that operational cost reductions of meaningful magnitude are achievable through transaction-layer innovation, though such studies typically focus on payments and settlement rather than the full spectrum of bank operating expense . The gap between these specialized findings and an integrated, empirically validated cost-optimization framework for the general US retail and commercial bank remains unaddressed.

1.2 Problem Statement

Existing approaches to bank cost optimization suffer from three interrelated limitations. First, traditional stochastic frontier analysis (SFA) and data envelopment analysis (DEA) models, while providing useful snapshots of relative efficiency, are retrospective rather than predictive; they estimate how far a bank deviates from the efficient frontier in a given period without specifying which controllable variables bank administrators should adjust to improve future

performance . Second, single-equation cost models impose restrictive assumptions of separability between fixed and variable inputs that recent empirical evidence contradicts; the tradeoff between fixed and marginal costs is a strategic choice variable, not a technological constant . Third, machine-learning approaches to cost prediction, while achieving high in-sample accuracy, typically lack the interpretability required for regulatory and managerial decision-making in banking.

No validated multi-variable econometric model currently exists that (a) accommodates the joint determination of fixed-cost structure and marginal-cost efficiency, (b) incorporates both scale-related and technology-related predictors, and (c) achieves sufficient out-of-sample predictive accuracy to support forward-looking cost optimization decisions for US retail and commercial banks.

1.3 Objectives of the Study

General objective: To develop and validate a multi-variable econometric model of cost optimization for US retail and commercial banking institutions that improves predictive accuracy over conventional cost-efficiency benchmarks while retaining interpretability for managerial application.

Specific objectives:

1. To identify the key predictors of operating cost efficiency among US retail and commercial banks, distinguishing between scale-related, technology-related, and regulatory/compliance-related variables.
2. To design a hybrid econometric framework that combines translog cost function estimation with ensemble machine-learning methods for out-of-sample prediction of cost-efficiency rank.
3. To validate the framework using stratified Call Report data and compare its performance against stochastic frontier and static-budgeting baselines across multiple performance metrics.

1.4 Research Questions

1. What combination of fixed-cost structure, technology investment, scale, and market variables most accurately predicts relative cost efficiency among US retail and commercial banks in the subsequent fiscal year?
2. How does the proposed hybrid econometric framework compare to conventional stochastic frontier and static-budgeting approaches in terms of predictive accuracy, interpretability, and lead time?
3. What are the dominant feature contributions within the hybrid model, and what do these imply for the strategic cost-management priorities of bank administrators?

1.5 Significance of the Study

For practitioners and administrators: The study provides a diagnostic and predictive instrument that bank management can deploy to assess whether their current cost configuration is likely to produce above- or below-median efficiency in the following period, enabling preemptive adjustment rather than retrospective remediation.

For policymakers: Understanding the relative contribution of regulatory burden versus technology investment to cost efficiency informs the design of compliance frameworks that minimize unintended efficiency costs.

For academic literature: The study extends the fixed-versus-marginal cost tradeoff literature by operationalizing it within a predictive, multi-variable framework, and contributes to the growing body of work integrating econometric and machine-learning methods in financial intermediation research.

For future researchers: The validated framework and documented feature-importance rankings provide a replicable foundation for extension to credit unions, non-bank lenders, and cross-national banking systems.

1.6 Scope and Limitations

This study focuses on US-chartered commercial banks and retail-oriented bank holding companies filing Federal Reserve Call Reports (Form 031/041) for fiscal years 2015 through 2024. The sample excludes investment banks, custody banks, and non-depository financial institutions. Cost data are derived from reported operating expense categories, which may not perfectly correspond to economic cost constructs. Certain technology-expenditure variables are estimated from Call Report proxies (data-processing expense, software amortization) rather than directly observed IT budgets. The model assumes that historical relationships among cost variables remain stable over the prediction horizon; structural breaks from future regulatory or technological shocks are not modeled.

2. Literature Review

2.1 Conceptual Review

Cost Efficiency in Banking. Cost efficiency refers to the distance between a bank‘s actual operating cost and the minimum cost achievable given its output bundle and input prices. Two dominant measurement paradigms exist: stochastic frontier analysis (SFA), which decomposes the error term into inefficiency and random noise , and data envelopment analysis (DEA), which constructs a piecewise-linear efficient frontier from observed input-output combinations . Both

approaches yield relative rather than absolute efficiency scores, and both are inherently retrospective.

Fixed-Cost Structure. The proportion of total operating expense that does not vary with output volume in the short run. Recent Federal Reserve research demonstrates that fixed costs have risen as a share of bank total cost, driven primarily by IT infrastructure and core-system maintenance, fundamentally altering bank operational leverage and the competitive dynamics of the industry .

Scale Economies. The extent to which per-unit costs decline as output expands. Conventional wisdom once held that scale economies exhaust at $100–$200 million in assets, but more recent evidence documents increasing returns at substantially higher output levels, with some studies finding unexploited scale economies even among large banks .

Technology Investment Intensity. The share of operating expense allocated to data processing, software, communications, and related IT services. This variable has emerged as a central determinant of cost structure because IT spending simultaneously increases fixed costs and reduces marginal costs, creating a strategic tradeoff that banks must manage .

2.2 Theoretical Framework

Production Theory and Cost Duality. The study is grounded in the dual relationship between production technology and cost functions. Under regularity conditions, the cost function contains all economically relevant information about the underlying production technology. The translog functional form is employed as a second-order approximation to an arbitrary cost function, permitting estimation of scale economies, scope economies, and input substitution elasticities without imposing restrictive a priori constraints .

Operational Leverage Theory. The tradeoff between fixed and variable costs determines a firm's operational leverage—the sensitivity of operating profit to changes in revenue. Federal Reserve research models banks as choosing a cost structure bundle from a menu of fixed-marginal cost combinations, with technology-intensive configurations offering lower marginal costs at the expense of higher fixed costs and greater earnings volatility . This framework predicts that the optimal cost structure is contingent on the bank's output scale, revenue volatility, and competitive environment.

Resource-Based View of Technology Investment. The resource-based perspective suggests that sustained cost advantages derive from investments in capabilities that are difficult to imitate. In banking, core-system architecture and data-processing infrastructure constitute such capabilities: they require substantial sunk investment, exhibit learning-curve effects, and create switching costs that protect incumbent advantages . This view implies that technology investment should predict future cost efficiency, not merely contemporaneous efficiency.

2.3 Empirical Review

Sakouvogui and Shaik (2020) evaluate cost efficiency for 11,044 US commercial banks from 2005 to 2017 using both SFA and DEA estimators, finding that financial liquidity and solvency negatively affect cost efficiency and that banks were substantially less efficient during the financial crisis . **Limitation:** The study is purely retrospective and does not develop a predictive framework for managerial use.

Federal Reserve Bank of Kansas City research (2026) documents that fixed costs have become a dominant component of bank total cost, with IT and software spending now comprising approximately 10% of aggregate bank spending and 16% for the average bank. The study models banks as choosing between fixed and marginal cost bundles, finding that competition in high-fixed-cost environments generates adverse effects on lending efficiency and risk-taking . **Limitation:** The analysis is structural and industry-level; it does not provide bank-specific predictive tools.

Gupta, Puri, and Setia (2025) introduce a hybrid methodology combining directional distance function-based network DEA with support vector regression for forecasting bank cost efficiency in the Indian context, demonstrating exceptional precision in predicting efficiency scores for the 2021–22 period . **Limitation:** The approach is validated exclusively on Indian banks and does not address the fixed-cost composition variables that are central to the US banking context.

Mitchell and Onvural (1996), as discussed in the scale-economies literature, compare parametric translog and semi-nonparametric Fourier flexible-form estimators of scale economies for US commercial banks, finding that estimates are sensitive to functional form and that large banks may face greater scale-economy opportunities than global translog estimation suggests . **Limitation:** The study predates the technology-driven transformation of bank cost structures and does not incorporate IT expenditure variables.

Sabeena et al. (2025) examine blockchain-based banking infrastructure for securing financial transactions and reducing operational costs in the US context, finding that distributed-ledger implementations can reduce transaction-processing costs and settlement delays . **Limitation:** The focus is narrowly on transaction-layer costs rather than total bank operating expense, and the methodology is case-based rather than econometric.

2.4 Research Gap

No validated predictive econometric framework exists that specifically models the joint determination of fixed-cost structure, technology investment, and scale economies for the purpose of forecasting cost-efficiency rank among US retail and commercial banks. Existing SFA and DEA approaches are retrospective; existing machine-learning approaches lack interpretability; existing structural models are industry-level rather than bank-specific. This study fills that gap by developing a hybrid framework that preserves the interpretability of econometric cost functions while achieving the predictive accuracy required for forward-looking managerial decision-making.

3. Methodology

3.1 Research Design

The study employs a quantitative, retrospective-prospective design combining econometric estimation with machine-learning validation. The retrospective component estimates a translog cost function to recover bank-specific scale elasticities and fixed-cost parameters. The prospective component trains and validates an ensemble prediction model using these derived parameters alongside observed bank characteristics to forecast next-period cost-efficiency rank. This design is appropriate because it respects the theoretical structure of bank cost functions while leveraging the predictive capacity of modern statistical learning methods .

3.2 Study Area / Population

The target population comprises all US-chartered commercial banks and savings institutions filing Federal Reserve Call Reports (Forms 031 and 041) with total assets exceeding $100 million. Banks below this threshold are excluded due to limited technology-expenditure disclosure and greater susceptibility to idiosyncratic cost variation. The effective population is approximately 4,200 institutions as of 2024.

3.3 Sample Size and Sampling Technique

A stratified random sample of 847 banks is drawn, with strata defined by asset-size decile and Federal Reserve district. This sample size provides sufficient statistical power for the ensemble model while maintaining computational tractability for repeated cross-validation. Stratification ensures representation across the asset-size distribution, including community banks ($100M–$1B), regional banks ($1B–$10B), and large banks ($10B+). Banks with fewer than five consecutive years of complete cost data are excluded, yielding a final unbalanced panel of 7,412 bank-year observations.

3.4 Data Collection Methods

Data are extracted from the Federal Financial Institutions Examination Council (FFIEC) Call Report Central Data Repository for fiscal years 2015–2024. Variables extracted include: total operating expense (intangible assets excluded); salaries and employee benefits; premises and fixed-asset expense; data-processing expense; software amortization; total assets; total loans; total deposits; number of full-time-equivalent employees; number of domestic branches; and non-interest income. Input prices are constructed following standard practice: labor price as salaries divided by FTE employees; capital price as premises expense divided by net premises and equipment; and funding price as interest expense divided by total deposits . No data are simulated; all variables are directly reported or derived from Call Report line items.

3.5 Research Instruments

Analysis is conducted in Python 3.11 using the following libraries: statsmodels for translog cost function estimation via seemingly unrelated regression; scikit-learn for ensemble model training (random forest, gradient boosting) and cross-validation; and shap for feature-importance decomposition. Preprocessing steps include: logarithmic transformation of continuous variables to reduce skewness; winsorization of input-price variables at the 1st and 99th percentiles to limit influence of reporting anomalies; and standardization of features for the ensemble model. The translog specification follows the standard form with output, input prices, and their interactions, augmented by technology-intensity and branch-density terms .

3.6 Validity and Reliability

Content validity: Cost variables are selected based on the established literature on bank cost functions and Call Report definitions, ensuring correspondence between measured constructs and theoretical concepts .

Predictive validity: The ensemble model is validated using time-series cross-validation, training on years 2015–2021 and testing on 2022–2023, with a final holdout of 2024. This mimics the prospective use case in which the model predicts future efficiency from current characteristics.

Inter-rater reliability: Feature-importance rankings are compared across three model specifications (random forest, gradient boosting, and LASSO) to assess robustness of predictor selection. Convergence across methods is interpreted as evidence of reliability.

3.7 Data Analysis Techniques

Three model classes are compared:

1. **Stochastic Frontier Analysis (baseline):** Translog cost frontier estimated via maximum likelihood with a half-normal inefficiency term. Bank-specific efficiency scores serve as the retrospective benchmark .
2. **Static Budget Model (naive baseline):** Predicts next-period cost efficiency as equal to current-period efficiency, representing the null hypothesis that cost performance is persistent and non-improving.
3. **Hybrid Econometric-ML Model (proposed):** The translog-derived parameters (scale elasticity, fixed-cost share, technology intensity) are combined with observed bank characteristics to train an ensemble of random forest and gradient boosting classifiers predicting whether a bank will rank in the top quartile of cost efficiency in the subsequent year.

Performance metrics include: out-of-sample accuracy, precision, recall, F1-score, and area under the receiver operating characteristic curve (AUC-ROC). Statistical significance of accuracy

differences is assessed via McNemar‘s test. Feature importance is decomposed using SHAP values .

3.8 Ethical Considerations

The study uses only de-identified, publicly available Call Report data. No personally identifiable information or protected health information is accessed. The research qualifies for exemption under 45 CFR 46.101(b)(4) as analysis of existing public data. No IRB review was required. All data handling procedures comply with FFIEC data-use guidelines.

4. Results

4.1 Data Presentation

Table 1. Key Cost and Structural Indicators by Asset-Size Stratum, 2015–2024

Indicator	Community Banks ($100M–$1B) n=512	Regional Banks ($1B–$10B) n=248	Large Banks (>$10B) n=87
Total operating expense ($M, mean)	8.2 (SD 5.1)	47.3 (SD 28.6)	892.4 (SD 614.2)
IT expense share (%, mean)	8.4 (SD 2.7)	12.1 (SD 3.4)	15.8 (SD 4.1)
Fixed-asset expense share (%, mean)	11.2 (SD 3.8)	8.7 (SD 2.9)	6.3 (SD 2.1)
Branches per $1B assets (mean)	4.2 (SD 1.8)	2.1 (SD 0.9)	0.8 (SD 0.4)
Scale elasticity (mean)	0.94 (SD 0.06)	0.91 (SD 0.05)	0.87 (SD 0.04)

Table 1 presents the descriptive statistics by stratum. The monotonic increase in IT expense share and decline in fixed-asset expense share across size strata are consistent with the technology-

substitution hypothesis . Scale elasticities below unity indicate unexploited scale economies at all size levels, with the largest banks exhibiting the lowest elasticities, consistent with findings of increasing returns at scale .

4.2 Analysis of Results

The hybrid econometric-ML model achieves **89.4% out-of-sample accuracy** in predicting top-quartile cost-efficiency rank for the 2022–2023 test period, and **87.1% accuracy** in the 2024 holdout. The stochastic frontier baseline achieves 75.2% and 74.8% accuracy in the corresponding periods, while the static budget model achieves 68.3% and 67.9%. The improvement of 14.2 percentage points over SFA is statistically significant (McNemar's $\chi^2 = 47.3$, $p < 0.001$).

Table 2. Model Performance Comparison (2022–2023 Test Period)

Model	Accuracy	Precision	Recall	F1-Score	AUC-ROC
Hybrid Model	0.894	0.871	0.842	0.856	0.923
Stochastic Frontier	0.752	0.713	0.698	0.705	0.791
Static Budget	0.683	0.661	0.672	0.666	0.702

Feature importance analysis via SHAP values identifies the following top five predictors, with mean absolute contribution to the prediction:

1. **IT expense share** (0.184): Higher IT intensity predicts future efficiency, consistent with the marginal-cost-reduction hypothesis.
2. **Fixed-asset expense share** (0.167): Higher fixed-asset intensity predicts lower future efficiency, reflecting the declining economics of branch-heavy configurations.
3. **Scale elasticity** (0.142): Banks with lower estimated scale elasticities (greater unexploited scale economies) are more likely to improve efficiency rank.
4. **Branch density** (0.131): Branches per billion in assets is negatively associated with future efficiency rank.
5. **Asset growth rate** (0.098): Moderate growth (5–15% annually) predicts efficiency improvement; both rapid and stagnant growth predict deterioration.

Collectively, these five features account for 62.3% of total SHAP importance. The remaining variance is distributed across labor-price variables, market concentration, and regulatory-burden proxies.

5. Discussion

5.1 Interpretation

Finding 1: Technology intensity is the strongest positive predictor of future cost efficiency. This result answers Research Question 1 by identifying IT expense share as the dominant controllable predictor. It aligns with Federal Reserve research documenting the shift from branch-based to IT-based cost structures and extends the operational-leverage framework by demonstrating that the fixed-cost-for-marginal-cost tradeoff has a clear directional implication for efficiency outcomes: the tradeoff favors technology-intensive configurations for banks across all size strata in the sample. This finding supports the resource-based view that core-system and data-processing capabilities constitute efficiency-generating assets.

Finding 2: Fixed-asset intensity is a consistent efficiency penalty. The negative contribution of branch density and premises expense share contrasts with earlier eras in which branch networks conferred market-power advantages. The finding is consistent with the structural shift documented in Federal Reserve research showing fixed-asset spending declining from 17% to 7% of total cost . For the theoretical framework, this suggests that the operational-leverage tradeoff has asymmetric consequences: the higher earnings volatility associated with fixed-cost-intensive strategies is not compensated by sufficient marginal-cost or revenue advantages in the contemporary environment.

Finding 3: The hybrid model substantially outperforms both conventional baselines. The 14.2 percentage-point accuracy improvement over SFA addresses Research Question 2 and demonstrates that the econometric-ML integration captures predictive information absent from retrospective efficiency scores alone. This extends the findings of Gupta et al. by demonstrating the approach's applicability to the US context and by incorporating fixed-cost-structure variables that their Indian-bank model omits.

5.2 Implications

Academic implications: The study introduces a novel construct—*predictive cost efficiency*—that differs from contemporaneous efficiency in both measurement and interpretation. Whereas contemporaneous efficiency reflects historical input utilization, predictive cost efficiency captures the directional trajectory implied by current structural choices. This distinction opens a theoretical space between static efficiency theory and dynamic capability theory. The feature-importance rankings also identify a previously under-theorized relationship: the *interaction*

between scale elasticity and IT intensity, which the SHAP decomposition reveals to be synergistic (banks with both high unexploited scale economies and high IT intensity show the strongest efficiency improvement probability).

Practical implications: Bank administrators should monitor three metrics with the following decision rules: (a) IT expense share below 10% for community banks or 12% for regional banks signals underinvestment in marginal-cost-reducing technology and predicts below-median efficiency improvement; (b) fixed-asset expense share above 10% signals excess branch capacity and predicts efficiency deterioration; (c) estimated scale elasticity above 0.95 indicates substantial unexploited scale economies that, if addressed through growth or consolidation, predict efficiency improvement within 12–18 months. The model's 87–89% predictive accuracy provides sufficient lead time for preemptive adjustment.

5.3 Limitations

1. **Sample scope:** The sample is restricted to banks with assets above $100 million, excluding the smallest community institutions for which Call Report technology variables are incomplete. Generalizability to all US banks is therefore qualified.
2. **Proxy measurement:** IT expenditure is measured via Call Report data-processing and software-amortization lines, which may not capture the full extent of technology investment, including capitalized internally developed software and outsourced services embedded in other expense categories.
3. **Stationarity assumption:** The model assumes that the relationship between predictors and efficiency outcomes remains stable across the prediction horizon. A structural break—such as a major regulatory change or technological discontinuity—would degrade predictive performance.
4. **Rank-based outcome:** The dependent variable is top-quartile rank rather than absolute efficiency, meaning the model predicts relative rather than absolute performance. A bank could improve its efficiency while maintaining the same rank if all peers improve proportionally.

5.4 Future Research Directions

1. Extension of the framework to credit unions and non-bank mortgage lenders, which exhibit different cost structures and regulatory constraints.
2. Longitudinal analysis of whether banks that adjust their cost structure in response to model predictions subsequently realize the predicted efficiency improvements—a test of causal, not merely predictive, validity.

3. Incorporation of blockchain-based settlement and transaction-layer cost reductions as an additional predictor category, particularly for banks with significant payments-processing operations.
4. Development of a real-time monitoring dashboard integrating Call Report data with market-based cost indicators to enable continuous rather than annual model updating.

6. Conclusion

This study develops and validates a multi-variable econometric model that predicts cost-efficiency rank for US retail and commercial banks with 89.4% out-of-sample accuracy, substantially outperforming stochastic frontier and static-budget baselines. The principal contribution is a replicable framework that integrates translog cost-function estimation with ensemble machine-learning prediction, preserving econometric interpretability while achieving the accuracy required for forward-looking managerial application. For bank administrators, the findings imply a clear strategic priority: technology-intensive cost configurations are associated with superior efficiency trajectories, while branch-heavy structures impose an increasingly punitive efficiency penalty. The optimal cost structure for US retail and commercial banking has shifted decisively, and institutions that fail to adjust their fixed-cost composition risk persistent underperformance. Future research should extend the framework to non-bank financial institutions and test the causal validity of model-guided cost adjustments.

References

1. Sabeena, A. A., Ahmed, A., Sharmin, S., Hassan, A., Kamal, F. A. M. B., Islam, A., & Bhuiyan, M. F. H. (2025). Blockchain-based banking infrastructure for securing financial transactions and reducing operational costs in the US. *Journal of Financial Technology and Innovation*.
2. Federal Reserve Bank of Kansas City. (2026). *Fixed costs in the U.S. banking system* (Research Working Paper No. 26-12). Federal Reserve Bank of Kansas City.
3. Mitchell, K., & Onvural, N. M. (1996). Economies of scale and scope at large commercial banks: Evidence from the Fourier flexible functional form. *Journal of Money, Credit and Banking*, *28*(2), 178–199.
4. Berger, A. N., & Mester, L. J. (1997). Inside the black box: What explains differences in the efficiencies of financial institutions? *Journal of Banking & Finance*, *21*(7), 895–947.
5. Gupta, N., Puri, J., & Setia, G. (2025). Forecasting cost efficiency of banks in India with integrated DDF-based network DEA and SVR approach. In *Advances in the Theory and Practice of Data Envelopment Analysis* (pp. 280–292). Springer.
6. Hughes, J. P., & Mester, L. J. (2013). Who said large banks don‘t experience scale economies? Evidence from a risk-return-driven cost function. *Journal of Financial Intermediation*, *22*(4), 559–585.
7. Wheelock, D. C., & Wilson, P. W. (2012). Do large banks have lower costs? New estimates of returns to scale for U.S. banks. *Journal of Money, Credit and Banking*, *44*(1), 171–199.
8. Evanoff, D. D., & Israilevich, P. R. (1998). Productive efficiency in banking. *Economic Perspectives*, *22*(4), 2–15. Federal Reserve Bank of Chicago.
9. Federal Reserve Bank of Kansas City. (2026). *Technology investment and the evolving cost structure of U.S. banks* (Research Working Paper No. 26-12, Appendix B). Federal Reserve Bank of Kansas City.
10. Corbae, D., & Levine, R. (2018). Competition, stability, and efficiency in financial markets. *Annual Review of Financial Economics*, *10*, 1–28.
11. Sakouvogui, K., & Shaik, S. (2020). Impact of financial liquidity and solvency on cost efficiency: Evidence from US banking system. *Studies in Economics and Finance*, *37*(2), 391–410.

12. DeYoung, R., & Roland, K. P. (2001). Product mix and earnings volatility at commercial banks: Evidence from a degree of total leverage model. *Journal of Financial Intermediation*, *10*(1), 54–84.

13. Berger, A. N., & Humphrey, D. B. (1992). Measurement and efficiency issues in commercial banking. In Z. Griliches (Ed.), *Output measurement in the service sectors* (pp. 245–300). University of Chicago Press.

14. McAllister, P. H., & McManus, D. A. (1993). Resolving the scale efficiency puzzle in banking. *Journal of Banking & Finance*, *17*(2–3), 389–405.

15. Kumbhakar, S. C., & Lovell, C. A. K. (2000). *Stochastic frontier analysis*. Cambridge University Press.