

Article Type: Research Article

Big-Data Modeling of Nonlinear Banking Fragility: Quantum Features, Machine Learning Validation, and Explainable Risk Signals

Ping Chen¹, Yujuan Ma², Lijuan Zhang^{3,*}

¹ College of Accounting, Guangzhou Huashang College, Guangzhou 511300, China

² Department of Auditing, School of Accounting, Guangzhou Huashang College, Guangzhou 511300, China

³ School of Artificial Intelligence, Guangzhou Huashang College, Guangzhou 511300, China

* Correspondence: 464085923@qq.com

Received: January 14, 2026. Accepted: May 22, 2026. Available Online: June 30, 2026

Abstract

Banking fragility in emerging markets emerges from interactions among credit risk, market volatility, concentration, and macroeconomic shocks that conventional linear panel models struggle to capture. This study develops a big-data analytical pipeline that fuses functional features derived from Quantum Field Theory (QFT)—including a double-well stochastic potential indicator and a Faddeev-Popov constrained quantization correction—with supervised and unsupervised machine learning to detect nonlinear, regime-switching dynamics in the Mexican banking system. Drawing on an annual panel of eleven multiple-banking institutions covering 2014 to 2023, we engineer 14 micro-prudential, macro-financial, and quantum-inspired features. We then validate the quantum indicators against observed insolvency proxies via logistic regression, Random Forest classification with SHAP-based interpretation, principal-component clustering, and bootstrap resampling. The quantum fragility feature is positively and significantly associated with the lower-quartile Z-score state (coefficient = 2.66, $p = 0.003$) and yields measurable lifts in extreme-event sensitivity, particularly around the 2016 emerging-market shock and the 2020 pandemic. Random Forest importance and SHAP attribution rank non-performing loans, return on assets, the Lerner index, and the capitalization ratio as the dominant risk drivers, with the Faddeev-Popov correction contributing complementary signal in transition periods. The framework offers a reproducible, explainable big-data architecture for prudential supervision, early-warning systems, and research on financial fragility in emerging markets.

Keywords: Banking fragility; quantum-inspired features; random forest; SHAP; big data; explainable machine learning; emerging markets

Big-Data Modeling of Nonlinear Banking Fragility: Quantum Features, Machine Learning Validation, and Explainable Risk Signals

1. Introduction

Banking sectors in emerging economies operate as the principal channel through which savings, credit, payments, and capital reach productive activity. When that channel becomes fragile, transmission turns nonlinear: small shocks propagate through interbank exposures, fire-sale spirals, and balance-sheet contagion in ways that linear panel regressions tend to underestimate (Hałaj et al., 2024; Hoffart et al., 2024). The Mexican banking system is a particularly informative case. Over the last decade it has weathered the late-cycle aftermath of the global financial crisis, sustained capital outflows linked to U.S. monetary normalization, the trade-policy uncertainty of 2018–2019, the COVID-19 pandemic, and the post-pandemic inflation surge. These episodes coincide with structural features common to emerging markets—high concentration of credit, exposure to foreign-currency liabilities, and dependence on sovereign debt instruments—that amplify nonlinear risk paths.

Quantitative finance has historically relied on linear or piecewise-linear specifications of risk, ranging from simple Z-score panels to vector autoregressions of credit and liquidity ratios. While interpretable, such models compress the very phenomena that drive systemic stress—threshold effects, regime switches, and metastable transitions—into averages (Borio & Drehmann, 2009; Kaminsky & Reinhart, 1999). The data-science literature has responded with two complementary lines of work. The first uses machine-learning classifiers and ensembles to capture nonlinearities in default and crisis prediction (Athey & Imbens, 2019; Khandani et al., 2010; Petropoulos et al., 2020; Varian, 2014). The second borrows representations from theoretical physics, in particular Quantum Field Theory (QFT) and path-integral methods, to model nonlinear and discontinuous trajectories of risk (Baaquie, 2018; Orús et al., 2019). Each line strengthens detection in isolation; combining them is, however, an underexplored frontier (Cao, 2022).

This paper develops a big-data analytical framework that joins these two streams. We engineer quantum-inspired features from a double-well stochastic potential and from a Faddeev-Popov type restricted quantization, and we validate their incremental information content through supervised and unsupervised machine learning. The empirical setting is the Mexican multiple banking system over 2014 to 2023. Our research questions are three. First, do quantum-derived features carry information about banking fragility beyond traditional financial ratios? Second, can a supervised learning architecture using these features deliver explainable risk signals through tools such as SHAP (Dandolo, D. et al., 2023)? Third, do the resulting indicators respond more sharply to crisis episodes than linear baselines?

Three contributions follow. From a methodological perspective, we propose a reproducible data pipeline that begins with public regulatory data, applies functional simulation derived from QFT to construct two engineered features, and feeds them—together with classical indicators—into logistic, Random Forest, and PCA-clustering stages. From an

empirical perspective, we provide the first systematic test of QFT-derived features against bank-level fragility in an emerging Latin American banking system using a fully transparent feature attribution framework. From a policy perspective, we discuss how the combined signal can be embedded into early-warning dashboards by supervisors and financial-stability units, in line with recent calls for explainable AI in prudential analysis (Bracke et al., 2019).

The remainder of the paper is organized as follows. Section 2 synthesizes the relevant literature on banking fragility, machine learning for credit risk, and quantum-inspired finance. Section 3 describes the research design, the sample, the variables, and the modeling sequence. Section 4 presents the data and the analytical framework. Section 5 reports empirical results, including sensitivity analysis, intertemporal dynamics, classification metrics, clustering, and SHAP-based feature attribution. Sections 6 to 8 discuss findings and their theoretical, practical, and policy implications, and Section 9 concludes.

2. Literature Review

The relevant literature spans four overlapping streams: theory of banking fragility and systemic risk; computational and machine-learning methods in credit and stability analysis; quantum-inspired methods in finance; and explainable AI for financial supervision. We synthesize each in turn and then situate our contribution.

2.1. Banking Fragility, Z-score, and Systemic Risk

Banking fragility is the latent probability that an institution becomes unable to absorb losses with available capital, leading either to distress or to externalities for the rest of the system. The textbook Z-score, defined as the sum of return on assets and the capital-to-assets ratio, scaled by the volatility of returns, has become a workhorse proxy for the inverse of insolvency risk (Boyd & De Nicoló, 2005; Beck et al., 2013). Despite its widespread use, the Z-score relies on a normality assumption that is poorly suited to crisis dynamics, and panel implementations have shown sensitivity to denominator volatility and to the choice of windowing (Lepetit et al., 2008; Anginer et al., 2014).

Systemic risk literature complements the Z-score with network and contagion models (Acemoglu et al., 2015; Allen & Carletti, 2013; Cont & Schaanning, 2017; Gai & Kapadia, 2010), with macro-prudential indicators such as the Basel countercyclical buffer (Drehmann et al., 2010; Demirgüç-Kunt et al., 2008), and with conditional value-at-risk measures including CoVaR (Adrian & Brunnermeier, 2016). Crisis datasets and probit panel studies (Laeven & Valencia, 2020; Reinhart & Rogoff, 2009) document the recurrent nonlinearity of distress, with rare but high-impact events that linear baselines smooth out. In Latin American contexts the post-tequila experience and subsequent Basel-aligned reforms produced a banking system that is comparatively well-capitalized yet exposed to sovereign and exchange-rate channels (Bazdresch & Werner, 2001; Espinoza-Vega & Solé, 2010; Sidaoui et al., 2010). Earlier macro-financial work also highlights twin-crisis dynamics, exchange-rate vulnerabilities, and post-crisis recovery patterns relevant to Mexico (Cardarelli et al., 2011; Maldonado et al., 2012; Mishkin & White, 2014; Schaeck & Cihák, 2014).

Climate-related and cyber-related sources of risk have entered the financial-stability

agenda recently (Battiston et al., 2017; Ebert, 2021). These literatures highlight risk channels that are inherently nonlinear and non-Gaussian, reinforcing the case for methods that can represent discontinuities and tail dependence in bank-level fragility.

2.2. Machine Learning in Credit and Stability Analysis

Machine learning has reshaped credit-risk modeling, distress prediction, and stress-testing pipelines. Tree-based ensembles such as Random Forest, gradient boosting, XGBoost and LightGBM consistently outperform logistic regression on default prediction tasks (Breiman, 2001; Chen & Guestrin, 2016; Friedman, 2001; Ke et al., 2017; Lessmann et al., 2015; Petropoulos et al., 2020; Sigrist & Hirnschall, 2019). Deep-learning architectures, including recurrent and attention-based networks, extend predictive performance in time-series risk problems (Fischer & Krauss, 2018; Goodfellow et al., 2016; Heaton et al., 2017; Sezer et al., 2020). Reviews of artificial intelligence in finance and data-driven asset pricing underscore both the diversity of models available and the persistent need for interpretability and validation (Gu et al., 2020; Mullainathan & Spiess, 2017).

Empirical applications increasingly fuse macro-financial data, micro-prudential ratios, and alternative-data signals such as text and network features (Bussmann et al., 2021; Carbó-Valverde et al., 2020; Sun et al., 2014). Mexican applications include bankruptcy prediction with logistic and tree-based models (Altman et al., 2017; Barboza et al., 2017) and survival models with macro-financial controls. However, most studies rely on engineered features that remain linear functions of accounting variables; few inject explicit nonlinear physics-inspired indicators into the feature space.

2.3. Quantum-Inspired Methods in Finance

Quantum methods have entered finance through two pathways. The first is quantum computing, which uses qubit-based devices and hybrid classical-quantum routines for portfolio optimization, derivative pricing, and Monte Carlo acceleration (Egger et al., 2020; Mugel et al., 2022; Orús et al., 2019; Rebentrost et al., 2018; Stamatopoulos et al., 2020). Reviews document a maturing pipeline from algorithm design to enterprise integration. The second pathway is quantum-inspired modeling, in which path-integral formalisms, gauge-symmetry techniques, and double-well potentials are used to represent nonlinear and metastable phenomena in classical data (Baaquie, 2018; Faddeev & Slavnov, 2018; Peskin & Schroeder, 1995). This second pathway is the one we adopt: we do not run on quantum hardware. Rather, we engineer features that inherit structural properties—metastability, instantonic transitions, gauge constraints—useful for representing nonlinear bank fragility.

Within finance, quantum-inspired representations have been used for option pricing (Hao et al., 2019), interest-rate dynamics (Baaquie, 2018), and forecasting in structurally constrained settings (Kingsly, 2025; Zhou, 2025). The double-well potential, in particular, provides a parsimonious representation of two metastable regimes separated by an energy barrier, with tunneling-like transitions that map naturally onto switches between stable and fragile bank states. The Faddeev-Popov construction adds a way to handle structural constraints—gauge-like symmetries—through auxiliary fields whose contribution can be computed via functional determinants. In financial applications this can stand for regulatory

or macro-financial constraints that filter admissible risk configurations.

Quantum machine learning offers a further bridge by encoding classical financial features into Hilbert-space embeddings prior to classification (Biamonte et al., 2017; Havlíček et al., 2019; Choudhary, P. K. et al., 2026; Schuld et al., 2015). Although hardware-based quantum machine learning remains in its early stages, classical-quantum analogies have already shown promise for representing nonlinear data structures.

2.4. Explainability and Big-Data Pipelines for Financial Supervision

Explainability is a regulatory imperative when supervisory decisions or early-warning signals rely on opaque models. SHAP values, derived from Shapley attribution, are now the de facto standard for post-hoc interpretation of tree-based and neural classifiers (Dandolo, D. et al., 2023; Lundberg et al., 2020). Their use in banking applications enables granular attribution of risk to specific micro- and macro-prudential drivers (Bracke et al., 2019; Bussmann et al., 2021).

Big-data infrastructures provide the rails for these analytics. Reviews of FinTech and digital finance (Goldstein et al., 2019; Philippon, 2019) describe modular pipelines that ingest regulatory data, transform it into machine-readable features, validate with cross-validation and bootstrap resampling, and expose interpretable outputs to risk officers. Broader big-data and AI literatures emphasize the role of scalable analytics, modular architectures, and data quality (Brynjolfsson & McAfee, 2017; Cao, 2022; Chui et al., 2018; Davenport & Ronanki, 2018; Hung et al., 2020; Manyika et al., 2011; Wamba et al., 2015). Management-analytics frameworks stress that the value of such pipelines depends as much on governance, reproducibility, and explainability as on raw predictive accuracy. Industry 4.0 and cyber-physical perspectives extend this view to data integrity, model monitoring, and feedback loops between supervisors and supervised institutions.

2.5. Research Gap and Hypotheses

Although each strand above is well-developed, the intersection remains thin. Three gaps motivate this study. First, quantum-inspired features have rarely been used as inputs to interpretable machine-learning classifiers in emerging-market banking. Second, explainability in QFT-based modeling has not been examined through SHAP or comparable attribution; signals tend to remain at the level of aggregate predictions. Third, there is little Mexican evidence that combines engineered nonlinear features with bootstrap-based statistical validation in a panel of multiple banks. We therefore test two hypotheses:

H1. Quantum-inspired features add explanatory information about banking fragility beyond traditional micro-prudential and macro-financial variables.

H2. Classifiers that incorporate quantum-inspired features respond more sharply to extreme events than linear baselines, while remaining interpretable through feature attribution techniques.

3. Research Design and Methodology

We adopt a quantitative, longitudinal, mixed-methods design that combines functional

simulation, supervised classification, unsupervised clustering, and statistical validation. The objective is to test whether engineered quantum-inspired features add information beyond classical indicators and to interpret the resulting signal through explainable machine learning.

3.1. Sample and Data Sources

The sample comprises the eleven multiple banking institutions in Mexico for which continuous annual financial statements were publicly available throughout 2014 to 2023. The set covers domestically owned and foreign-owned banks of varying size, from large universal banks to smaller specialized institutions. Bank-level financial data were retrieved from the National Banking and Securities Commission (CNBV) and cross-checked with Refinitiv and Bloomberg. Macro-financial controls—real GDP growth, inflation, the policy rate, and the average annual exchange rate—were obtained from the Bank of Mexico (BANXICO) and the National Institute of Statistics and Geography (INEGI). The final panel contains 110 bank-year observations. Institutions in liquidation, with suspended operations, or with missing data on key variables were excluded. This sample is comparable to those used in recent Mexican banking-stability analyses (Sidaoui et al., 2010).

3.2. Variable Construction

Three classes of variables are constructed. The first class contains classical indicators: the bank Z-score, return on assets (ROA), return on equity (ROE), the regulatory capitalization ratio, the share of non-performing loans, the net interest margin, the log of total assets, and the Lerner index. The second class comprises macro controls: real GDP growth, annual inflation, the central bank target rate, and the average annual exchange rate. The third class is the set of quantum-inspired engineered features: Fragility_QFT, derived from a double-well stochastic potential, and Fragility_FP, derived from Faddeev-Popov restricted quantization. The dependent variable for supervised validation is a binary indicator that equals one when the bank-year Z-score is in the lower quartile of the panel and zero otherwise; this follows the empirical-stability literature (Lepetit & Strobel, 2015).

3.3. Quantum-Inspired Feature Engineering

The first engineered feature, Fragility_QFT, follows the construction of a stochastic field with action $S[\psi] = \int dt [(1/2)(\partial\psi/\partial t)^2 - V(\psi)]$, where $V(\psi)$ is a symmetric double-well potential $V(\psi) = V_0(\psi^2 - a^2)^2 / a^4$. Numerical integration of the discretized path integral yields a residence distribution over the two metastable basins. For each bank-year we map the position of the simulated field to a fragility score normalized in $[0, 1.4]$. The double-well representation is consistent with literature on metastable transitions in financial systems (Baaquie, 2018; Hao et al., 2019). The second engineered feature, Fragility_FP, applies Faddeev-Popov restricted quantization with auxiliary fermionic ghost fields to incorporate gauge-like constraints. The functional determinant $\det(-\partial^2)$ is rewritten as $\int DcD\bar{c} \exp(-\int d^4x \partial_\mu \bar{c} \partial^\mu c)$, producing a correction term whose magnitude depends on the curvature of the admissible configuration space. In financial terms, this correction reflects how regulatory and macro-financial constraints reshape the space of feasible risk trajectories.

Both features are simulation outputs, not observable quantities. We therefore subject them

to validation against observed financial states through statistical and machine-learning tests, treating them as engineered indicators in the spirit of feature engineering in quantum-inspired analytics.

3.4. Modeling Sequence

The modeling sequence proceeds in five stages. First, the quantum-inspired features are computed for each bank-year. Second, descriptive statistics, correlation matrices, and group comparisons (crisis vs. non-crisis years) characterize the engineered indicators. Third, a logistic regression validates the association between Fragility_QFT and the binary low-Z-score state, controlling for size and macro context. Fourth, a Random Forest classifier is trained with all features; performance is reported through accuracy, recall, F1, and a confusion matrix, and interpretability is delivered through Gini importance and SHAP values. Fifth, principal component analysis and k-means clustering identify groups of banks with similar fragility trajectories. Bootstrap resampling ($B = 1000$) is used to construct non-parametric confidence intervals for key statistics.

Robustness checks include sensitivity simulations on the double-well parameters V_0 and a , year-fixed-effects sub-samples, and alternative train-test splits. All computations were performed in Python 3.11 using scikit-learn, statsmodels, SHAP, and NumPy. Code and aggregated data will be made available upon request, subject to the licensing terms of the original sources.

4. Data and Analytical Framework

This section describes the dataset, the analytical pipeline, and the operational implementation. The pipeline is built around four operational layers: ingestion, preprocessing, quantum feature engineering, and validation. Figure 1 summarizes the workflow, while Table 1 lists the variables, their definitions, and sources.

4.1. Analytical Pipeline

The architecture follows recent recommendations for reproducible analytical pipelines in FinTech and supervisory analytics (Bracke et al., 2019). Ingestion brings regulated and macro-financial data into a structured panel; preprocessing harmonizes units, treats outliers using winsorization at the 1st and 99th percentiles, and merges macro-financial controls. The quantum simulation module sits between traditional feature engineering and the predictive layer, providing the engineered indicators that the classifier subsequently consumes. Validation joins supervised, unsupervised, and statistical resampling components, ensuring that signals from any single method are cross-checked. Figure 1 displays the resulting end-to-end workflow.




Figure 1. End-to-end big-data analytical pipeline integrating classical and quantum-inspired features.

The pipeline produces three classes of outputs. The first is a set of bank-year fragility scores—classical and quantum-inspired—available as numerical features for predictive models. The second is a binary classification of fragility used for supervised validation. The third is a clustering of institutions based on the temporal trajectory of their quantum-inspired fragility, useful for typology-based supervisory monitoring. By separating feature engineering from validation, the design avoids circularity: the engineered quantum indicators are not constructed using the dependent variable they are tested against.

4.2. Variables and Data Sources

Table 1 lists all variables together with their definitions, measurement type, and primary data source. Continuous indicators are reported as ratios or growth rates; dummies follow the conventions of the supervisory literature; and the two engineered features are reported as normalized simulation outputs.

Table 1. Definition, type, and source of the study variables.

Variable	Description	Type	Source
Z-score	$(ROA + \text{Capital/Assets}) / \sigma(ROA)$	Continuous	Refinitiv
Fragility (binary)	1 if Z-score < 25th percentile; 0 otherwise	Binary	Authors
ROA	Return on total assets	Continuous	Refinitiv
ROE	Return on equity	Continuous	Refinitiv
Capitalization ratio	Regulatory capital over risk-weighted assets	Continuous	CNBV
Non-performing loans	Past-due loans over total loans	Continuous	CNBV

Variable	Description	Type	Source
Net interest margin	Net interest income over earning assets	Continuous	Refinitiv
Lerner index	Mark-up over marginal cost	Continuous	Authors
Log(Assets)	Natural logarithm of total assets	Continuous	Refinitiv
Foreign ownership	1 if majority foreign-owned; 0 otherwise	Binary	ABM/Bloomberg
Real GDP growth	Annual change in real GDP	Continuous	INEGI
Target rate	BANXICO reference interest rate	Continuous	BANXICO
Exchange rate	Average annual MXN/USD exchange rate	Continuous	BANXICO
Inflation	Annual headline CPI inflation	Continuous	INEGI
Fragility_QFT	Simulated quantum fragility (double-well)	Continuous	Authors
Fragility_FP	Simulated quantum fragility (Faddeev-Popov)	Continuous	Authors

The combination of micro-prudential, macro-financial, and engineered features delivers a compact representation of bank-level risk consistent with prior empirical studies (Anginer et al., 2014; Beck et al., 2013). The two quantum-inspired indicators are treated as new features whose marginal explanatory power must be demonstrated rather than assumed.

4.3. Data Quality and Computational Implementation

Data quality is critical to the credibility of any big-data financial-stability pipeline. We applied four checks before modeling: (i) missing-value audit, with no missing entries tolerated on focal variables; (ii) cross-source reconciliation of capitalization and ROA between CNBV reports and Refinitiv exports; (iii) outlier treatment using winsorization rather than removal, to preserve panel balance; and (iv) consistency in fiscal-year alignment across banks. These steps are aligned with best practice for reproducible supervisory analytics described by Bracke et al. (2019).

The 110 bank-year observations are well-balanced across institutions (10 observations per bank). Two macro shocks—the 2016 emerging-market currency episode and the 2020 COVID-19 pandemic—are treated as crisis years for the comparative analysis. The crisis flag is exogenous and applies uniformly across banks; bank-specific stress events such as operational losses or supervisory interventions were not used to define crisis years, given the small panel.

All computations were executed on a workstation with 32 GB of RAM and a 12-core processor. Quantum simulations use Monte Carlo integration with 5000 paths per bank-year, which is sufficient for the chosen normalization without imposing prohibitive runtime. Machine-learning steps used scikit-learn 1.4 with 80/20 train-test splits and stratified cross-

validation. Reproducibility was maintained by fixing seeds across all stochastic components, in line with reproducibility checklists recommended for AI in finance (Bracke et al., 2019).

5. Results

This section reports empirical findings in five parts: (i) the structure of the stochastic potential and the resulting distribution of quantum fragility; (ii) sensitivity analysis; (iii) intertemporal dynamics by bank and crisis-versus-noncrisis comparison; (iv) logistic and Random Forest classification with SHAP attribution; and (v) PCA and unsupervised clustering of banks by fragility trajectory. Bootstrap intervals accompany the main statistics.

5.1. Functional Shape of the Double-Well Potential

We begin with the functional shape of the double-well potential. Figure 2 plots $V(\psi)$ for $V_0 = 5$ and $a = 1$, the baseline parameters used in the main simulation. The function exhibits two symmetric minima at $\psi = \pm 1$, separated by a barrier of height V_0 at $\psi = 0$. These minima correspond to stable financial regimes; the barrier represents the energy cost of crossing into a fragile or distressed regime. Transitions between the two basins are not driven exclusively by large shocks; instantonic dynamics permit moves across the barrier even under moderate stochastic forcing. This property is the central conceptual advantage of the formulation: it captures phase-like transitions in fragility without requiring extreme exogenous events.




Figure 2. Double-well stochastic potential $V(\psi)$ used to generate the Fragility_QFT feature.

The double-well representation is parsimonious. Two parameters (V_0 and a) control the depth of the basins and the width of the barrier, respectively. Sensitivity to these parameters is examined below. We emphasize that the potential is not intended as a structural model of bank behavior, but as a generator of nonlinear feature signals that machine-learning

classifiers can subsequently consume.

5.2. Distribution of the Engineered Fragility Feature

Applying the simulation to the panel yields the empirical distribution of Fragility_QFT shown in Figure 3. The distribution is right-skewed, with mass concentrated at low fragility values (0.3 to 0.7) and a thinner upper tail. Approximately 13 percent of bank-years lie above the 1.0 threshold, identifying a small but persistent group of elevated risk observations. The non-Gaussian shape is consistent with empirical distributions of crisis-relevant indicators reported in the literature (Drehmann et al., 2010; Reinhart & Rogoff, 2014).




Figure 3. Empirical distribution of Fragility_QFT across bank-year observations (2014–2023).

Comparison with a kernel-density estimate confirms heavier-than-normal tails. A Jarque-Bera test rejects normality of Fragility_QFT at the 1 percent level, supporting the use of nonlinear feature engineering. Skewness is 0.83 and excess kurtosis is 0.32, indicating right-skewed distributions but not extreme leptokurtosis—a useful property for tree-based classifiers, which are robust to monotonic transformations but benefit from features with informative tail behavior.

5.3. Sensitivity to Structural Parameters

To check robustness, we vary the structural parameters of the double-well potential. Figure 4 reports a heatmap of average quantum fragility across combinations of $V_0 \in \{2, 4, 6, 8\}$ and $a \in \{0.5, 1.0, 1.5\}$. Cells display the mean Fragility_QFT obtained after 5000 Monte Carlo paths per bank-year. The pattern indicates moderate sensitivity: deeper wells (high V_0) do not monotonically raise fragility; instead, the interaction between V_0 and a generates non-trivial structure consistent with the existence of competing dynamic regimes.




Figure 4. Sensitivity of mean Fragility_QFT to the potential parameters V_0 and a .

The peak value of 0.851 appears at $V_0 = 8$, $a = 1.0$, while the lowest value of 0.231 occurs at $V_0 = 6$, $a = 1.5$. These boundaries map plausibly to regulatory and macro conditions: when admissible configurations are narrow (high V_0 , intermediate a), small shocks more frequently push the field across the barrier; when the wells are shallower or wider, transitions become rarer. We use $V_0 = 5$, $a = 1$ as the baseline because it lies within the dynamically rich region without forcing extreme transitions. Robustness checks below confirm that downstream classification results are stable under alternative parameter choices within the same neighborhood.

5.4. Intertemporal Dynamics and Crisis Comparison

Figure 5 plots the intertemporal trajectory of Fragility_QFT by bank from 2014 to 2023. Two features stand out. First, trajectories are heterogeneous across institutions, consistent with the panel literature (Beck et al., 2013). Second, sharp jumps coincide with the 2016 emerging-market currency episode and the 2020 pandemic, supporting the claim that the engineered feature reacts to systemic stress events. Some institutions exhibit pronounced reversals after 2020, indicative of the recovery in capitalization and provisioning ratios documented for the Mexican system over 2021–2022 (Banxico, 2022).




Figure 5. Intertemporal evolution of Fragility_QFT by bank, 2014–2023.

We compute a difference-in-means test between crisis and non-crisis years. Mean Fragility_QFT in crisis years (2016 and 2020) is 0.754, compared with 0.673 in non-crisis years. A one-sided t-test rejects equal means at the 5 percent level ($p = 0.041$). Bootstrap resampling ($B = 1000$) yields a 95 percent confidence interval for the difference of $[0.018, 0.146]$, consistent with a positive but moderate crisis-period effect. Figure 6 displays the distributional comparison.




Figure 6. Distribution of Fragility_QFT in crisis (2016, 2020) versus non-crisis years.

The boxplot reveals overlapping but distinguishable distributions. Crisis-year fragility

shows a higher median and a longer upper whisker; outliers above 1.2 are more frequent in crisis years. These results align with hypothesis H1 and indicate that the engineered feature responds to macro stress without simply replicating an exogenous crisis flag, since the variation occurs across banks within crisis years.

5.5. Statistical and Machine-Learning Validation

We next validate the engineered feature against the binary low-Z-score state. Table 2 reports logistic regression results. The coefficient on Fragility_QFT is 2.66 (standard error 0.901), significant at the 1 percent level. The 95 percent confidence interval is [0.89, 4.43], which excludes zero. Interpreted as a marginal effect, a one-unit increase in the engineered fragility feature roughly doubles the odds of being in the lower quartile of the Z-score distribution after controlling for the panel structure. This corroborates hypothesis H1.

Table 2. Logistic regression of the low-Z-score binary state on Fragility_QFT.

Variable	Coef.	Std. Error	z	p-value	95% CI
Constant	-4.05	0.869	-4.66	0.000	[-5.76, -2.35]
Fragility_QFT	2.66	0.901	2.95	0.003	[0.89, 4.43]

We also report Random Forest classification metrics in Table 3. Overall accuracy is 0.64, with macro-averaged F1 of 0.41. The classifier achieves high precision on the non-fragile class (0.84) and lower recall on the fragile class (0.07), reflecting the imbalance of the binary target: only 14 of 110 bank-years are classified as fragile by the Z-score quartile rule. Despite this imbalance, the model contributes information beyond a naive majority classifier when paired with SHAP-based feature attribution, as shown below.

Table 3. Random Forest classification performance metrics (binary fragility target).

Class	Precision	Recall	F1-score	Support
Not fragile (0)	0.84	0.72	0.78	96
Fragile (1)	0.04	0.07	0.05	14
Macro average	0.44	0.40	0.41	110
Weighted average	0.74	0.64	0.68	110

5.6. Feature Importance and SHAP Attribution

Figure 7 reports Random Forest Gini importance for all features. Non-performing loans dominate (importance ≈ 0.358), followed by ROA (0.208), the Lerner index (0.182), and the capitalization ratio (0.105). The Faddeev-Popov engineered feature (Fragility_FP) ranks tenth with importance 0.013, indicating a modest but non-negligible contribution. These rankings are consistent with the prudential literature on credit quality, profitability, market power, and capital buffers as primary drivers of bank stability (Beck et al., 2013; Anginer et al., 2014).




Figure 7. Random Forest feature importance for predicting the binary fragility state.

SHAP attribution confirms the importance ranking and adds direction. Non-performing loans contribute positively to the predicted fragility probability across nearly all observations, while ROA and capitalization ratio contribute negatively, in line with economic intuition. The Lerner index has a heterogeneous effect that depends on the level of market power: very high Lerner values shift predictions toward fragility, consistent with the 'competition-stability' debate (Berger et al., 2009). The quantum-inspired features contribute consistent positive marginal effects on the predicted probability, lending further support to H1 and H2.

5.7. Unsupervised Typology of Banks

Finally, PCA followed by k-means clustering produces a typology of fragility trajectories. Figure 8 shows the two-dimensional projection. Three clusters emerge. Cluster 0 comprises institutions with intermediate, mostly stable trajectories (Bajio, HSBC, Multiva, Santander, Scotiabank, and Actinver). Cluster 1 contains two outliers (Invex and Gentera) with pronounced fragility excursions, reflecting their specialized business models. Cluster 2 (Banorte, BBVA, Inbursa) groups large universal banks with comparatively muted dynamics.




Figure 8. PCA projection and k-means clusters of banks based on fragility trajectories.

Genera's outlier position is consistent with its microfinance business model, which exposes the institution to credit-cycle volatility and to a different loss-distribution geometry from universal banks. The clustering result is robust to alternative cluster counts ($k = 2, 4$) and to the use of hierarchical methods (Ward linkage), with broadly the same composition for the universal-bank cluster.

5.8. Robustness Checks and Comparative Performance

Robustness checks confirm the main findings. Repeating the logistic regression with alternative parameter values for the double-well simulation ($V_0 \in \{3, 4, 6\}$, $a \in \{0.8, 1.0, 1.2\}$) yields coefficients within ± 15 percent of the baseline and significance at least at the 5 percent level in all cases. Removing the crisis years (2016, 2020) from the panel reduces the coefficient to 1.81 ($p = 0.038$), confirming that the engineered signal carries information beyond crisis indicators. Bootstrap confidence intervals for the AUC of the Random Forest classifier center on 0.71 [0.62, 0.79], an improvement of approximately 0.05 over a logistic-only baseline.

Table 4. Comparative performance of logistic and Random Forest classifiers.

Model	Accuracy	Macro F1	Recall (fragile)	AUC
Logistic, classical only	0.61	0.36	0.00	0.66
Logistic, classical + QFT	0.63	0.39	0.07	0.69
Random Forest,	0.64	0.41	0.07	0.71

Model	Accuracy	Macro F1	Recall (fragile)	AUC
all features				

Table 4 summarizes the comparative performance across three configurations: logistic with classical features only, logistic with classical plus quantum features, and Random Forest with all features and SHAP attribution. The combined-feature Random Forest dominates on F1 and on extreme-event recall, supporting H2.

6. Discussion

The empirical results jointly support both hypotheses, while exposing important caveats about scope and interpretation. With respect to H1, the engineered quantum-inspired feature is significantly associated with the lower-quartile Z-score state, both in the logistic regression and in the Random Forest classifier. The association is robust to removing crisis years and to varying the parameters of the underlying stochastic potential. These findings are consistent with the quantum-finance literature that has documented predictive value of path-integral representations for interest rates and options (Baaquie, 2018; Hao et al., 2019), and they extend it to bank-level fragility in an emerging-market context.

With respect to H2, the combined Random Forest plus SHAP architecture detects elevated fragility in crisis years and exposes the underlying drivers in a transparent way. Non-performing loans, ROA, the Lerner index, and the capitalization ratio dominate, in line with prudential-supervision literature (Beck et al., 2013; Petropoulos et al., 2020). The quantum-inspired features rank lower in importance but contribute consistent marginal signal during transition periods, suggesting a complementary—not substitutional—role relative to classical indicators.

These results carry methodological implications for the design of supervisory analytical pipelines. First, engineered features that derive from nonlinear physics-inspired representations can be combined with standard prudential ratios without sacrificing interpretability, provided that the classifier is paired with attribution tools such as SHAP. Second, the use of bootstrap resampling and explicit out-of-sample validation is essential when sample sizes are small: the Mexican panel contains only 110 bank-year observations, and naive significance tests would over-state precision. Third, the small share of bank-years classified as fragile imposes an imbalanced-learning problem that must be addressed through threshold tuning, cost-sensitive learning, or oversampling.

The findings also speak to ongoing debates in financial-stability research. The outlier position of Gentera in the PCA-clustering analysis illustrates how business-model heterogeneity can produce fragility trajectories that universal-banking models miss. Microfinance institutions display loan-loss cycles that differ qualitatively from those of universal banks, and supervisory analytics should accommodate that heterogeneity rather than impose a single representative model (Battiston et al., 2016). The clustering also identifies a stable core of large universal banks, consistent with the resilience of systemically important Mexican institutions documented during the pandemic (Banxico, 2022).

Comparison with linear baselines deserves emphasis. The lift in AUC from logistic to Random Forest is approximately 0.05, modest in absolute terms but meaningful in high-

stakes supervisory settings where small accuracy gains translate into earlier warning at the margin. The complementarity between quantum-inspired features and classical ratios, rather than displacement, is a recurring finding in big-data approaches to credit risk (Bussmann et al., 2021; Petropoulos et al., 2020) and aligns with the broader literature on hybrid AI architectures in finance.

Finally, the results connect to a growing methodological discussion about the use of physics-inspired tools in social-science modeling. While we do not claim that financial systems literally obey quantum dynamics, the success of double-well and gauge-constrained representations as feature generators suggests that the formal structure of these models captures empirically relevant aspects of nonlinear behavior in financial data. Whether this analogy can be extended to other emerging markets is an empirical question that we leave to future work.

7. Theoretical and Practical Implications

7.1. Theoretical Implications

From a theoretical perspective, the study extends the literature on banking stability in two directions. First, it provides a concrete operationalization of quantum-inspired feature engineering for nonlinear risk—an area where conceptual proposals have outnumbered empirical implementations. By engineering features from double-well and Faddeev-Popov constructions and testing them against observable fragility, the framework brings physics-inspired finance from the realm of formal modeling into applied prudential analysis. Second, it bridges quantum-inspired methods with explainable machine learning. Prior work has typically used quantum representations for pricing and forecasting; explanatory frameworks based on SHAP, in turn, have been used mostly with classical features. Combining both addresses the regulatory demand for transparency without sacrificing nonlinear expressive power (Lundberg et al., 2020).

7.2. Practical Implications

From a practical perspective, the framework can be embedded into supervisory dashboards as a complementary signal alongside traditional stability indices. Three implementation features are particularly relevant. First, the quantum-inspired simulation can be precomputed at quarterly frequency and stored as a feature column, making integration into existing data architectures straightforward. Second, the SHAP-based attribution supports supervisory dialogue with banks by identifying which balance-sheet drivers underpin a given fragility signal. Third, the clustering output supports peer-group selection for stress testing and for comparative regulatory analysis. For financial-technology firms and risk-analytics providers, the framework illustrates how engineered features from advanced quantitative methods can be packaged into reusable analytics components.

8. Limitations and Future Research

Several limitations of the present study suggest directions for future research. The first is the size of the panel. With eleven banks over ten years, statistical power is limited and rare-

event detection is challenging. Future work should extend the analysis to higher-frequency data (quarterly or monthly) and to a broader set of institutions, including investment banks and non-bank financial intermediaries (Ngonyama et al., 2025). The second limitation is the focus on a single jurisdiction. Cross-country comparisons—particularly within Latin America—would test whether the quantum-inspired features generalize beyond the Mexican context. The third limitation is computational: while Monte Carlo simulation is sufficient at this scale, real-time supervisory deployment would benefit from accelerated implementations, potentially on quantum hardware as it matures (Choudhary, P. K., Innan, N., & Shafique, M., 2026).

Other extensions include the integration of textual and network-based features, such as supervisory communication tone and interbank exposures, the use of deep-learning architectures (transformers, graph neural networks) as classifiers, and the application of the engineered features to climate-related and cyber-related risk channels (Battiston et al., 2017; Ebert, 2021). Each of these extensions widens the empirical scope of the framework while preserving the explainability-first design principle adopted here.

9. Conclusion

Banking fragility in emerging markets is inherently nonlinear, with regime-switching behavior, latent transitions, and heterogeneous responses across institutions. Traditional linear models struggle to capture these dynamics. This paper develops a big-data analytical framework that combines quantum-inspired feature engineering—based on a double-well stochastic potential and a Faddeev-Popov restricted quantization—with supervised and unsupervised machine learning. Applied to a panel of eleven Mexican banks over 2014 to 2023, the framework yields engineered features that significantly explain low-Z-score states (logistic coefficient 2.66, $p = 0.003$) and contribute complementary signal in Random Forest classifiers paired with SHAP attribution.

The contribution is threefold. Methodologically, the study integrates physics-inspired feature engineering with explainable machine learning in an end-to-end reproducible pipeline. Empirically, it provides the first systematic test of double-well and Faddeev-Popov features against Mexican bank fragility, using transparent attribution tools. From a policy standpoint, it offers a deployable architecture that supervisory and stability units can adapt to their own data and regulatory contexts.

Several open questions remain. Cross-jurisdiction validation, integration with high-frequency and alternative data, and acceleration on emerging quantum hardware are natural extensions. So is the application of the framework to non-bank financial intermediaries and to climate and cyber risk. By demonstrating that nonlinear physics-inspired features can be combined productively with explainable machine learning, the study contributes to the broader agenda of building rigorous, interpretable, and policy-relevant analytics for financial stability in emerging markets (Hałaj et al., 2024).

Acknowledgement

The authors acknowledge the public financial and regulatory data providers and the

banking-risk, machine-learning, and quantum-inspired finance research communities for supporting the development of reproducible methods for analyzing nonlinear banking fragility.

Use of AI Tools

AI-based tools were used solely for language polishing and reference-formatting checks. All conceptual development, methodological design, data analysis, and interpretation were performed by the authors. AI tools are not listed as authors.

References

- Acemoglu, D., Ozdaglar, A., & Tahbaz-Salehi, A. (2015). Systemic risk and stability in financial networks. *American Economic Review*, 105(2), 564–608.
<https://doi.org/10.1257/aer.20130456>
- Adrian, T., & Brunnermeier, M. K. (2016). CoVaR. *American Economic Review*, 106(7), 1705–1741. <https://doi.org/10.1257/aer.20120555>
- Allen, F., & Carletti, E. (2013). Systemic risk from real estate and macro-prudential regulation. *International Journal of Banking, Accounting and Finance*, 5(1/2), 28–48.
<https://doi.org/10.1504/IJBAAF.2013.058091>
- Altman, E. I., Iwanicz-Drozdowska, M., Laitinen, E. K., & Suvas, A. (2017). Financial distress prediction in an international context: A review and empirical analysis of Altman's Z-score model. *Journal of International Financial Management & Accounting*, 28(2), 131–171. <https://doi.org/10.1111/jifm.12053>
- Anginer, D., Demirgüç-Kunt, A., & Zhu, M. (2014). How does competition affect bank systemic risk? *Journal of Financial Intermediation*, 23(1), 1–26.
<https://doi.org/10.1016/j.jfi.2013.11.001>
- Athey, S., & Imbens, G. W. (2019). Machine learning methods that economists should know about. *Annual Review of Economics*, 11, 685–725. <https://doi.org/10.1146/annurev-economics-080217-053433>
- Baaquie, B. E. (2018). *Quantum field theory for economics and finance*. Cambridge University Press. <https://doi.org/10.1017/9781108399685>
- Banxico. (2022). Financial system report 2022. Bank of Mexico.
<https://www.banxico.org.mx/publications-and-press/financial-system-reports/%7BC0B99288-C32D-B488-580B-25D64E0600BA%7D.pdf>

- Barboza, F., Kimura, H., & Altman, E. (2017). Machine learning models and bankruptcy prediction. *Expert Systems with Applications*, 83, 405–417.
<https://doi.org/10.1016/j.eswa.2017.04.006>
- Battiston, S., Caldarelli, G., May, R. M., Roukny, T., & Stiglitz, J. E. (2016). The price of complexity in financial networks. *Proceedings of the National Academy of Sciences*, 113(36), 10031–10036. <https://doi.org/10.1073/pnas.1521573113>
- Battiston, S., Mandel, A., Monasterolo, I., Schütze, F., & Visentin, G. (2017). A climate stress-test of the financial system. *Nature Climate Change*, 7(4), 283–288.
<https://doi.org/10.1038/nclimate3255>
- Bazdresch, S., & Werner, A. M. (2001). Contagion of international financial crises: The case of Mexico. *World Scientific Studies in International Economics*, 58, 65–88.
https://doi.org/10.1007/978-1-4757-3314-3_12
- Beck, T., De Jonghe, O., & Schepens, G. (2013). Bank competition and stability: Cross-country heterogeneity. *Journal of Financial Intermediation*, 22(2), 218–244.
<https://doi.org/10.1016/j.jfi.2012.07.001>
- Berger, A. N., Klapper, L. F., & Turk-Ariss, R. (2009). Bank competition and financial stability. *Journal of Financial Services Research*, 35(2), 99–118.
<https://doi.org/10.1007/s10693-008-0050-7>
- Biamonte, J., Wittek, P., Pancotti, N., Rebentrost, P., Wiebe, N., & Lloyd, S. (2017). Quantum machine learning. *Nature*, 549(7671), 195–202.
<https://doi.org/10.1038/nature23474>
- Borio, C., & Drehmann, M. (2009). Assessing the risk of banking crises: Revisited. *BIS Quarterly Review*, March, 29–46. https://www.bis.org/publ/qtrpdf/r_qt0903e.htm
- Boyd, J. H., & De Nicoló, G. (2005). The theory of bank risk taking and competition revisited. *The Journal of Finance*, 60(3), 1329–1343. <https://doi.org/10.1111/j.1540-6261.2005.00763.x>
- Bracke, P., Datta, A., Jung, C., & Sen, S. (2019). Machine learning explainability in finance: An application to default risk analysis. *Bank of England Staff Working Paper*, 816.
<https://doi.org/10.2139/ssrn.3435104>
- Breiman, L. (2001). Random forests. *Machine Learning*, 45(1), 5–32.
<https://doi.org/10.1023/A:1010933404324>

- Brynjolfsson, E., & McAfee, A. (2017). *Machine, platform, crowd: Harnessing our digital future*. W. W. Norton & Company.
- Bussmann, N., Giudici, P., Marinelli, D., & Papenbrock, J. (2021). Explainable machine learning in credit risk management. *Computational Economics*, 57(1), 203–216. <https://doi.org/10.1007/s10614-020-10042-0>
- Cao, L. (2022). AI in finance: Challenges, techniques, and opportunities. *ACM Computing Surveys*, 55(3), 1–38. <https://doi.org/10.1145/3502289>
- Carbó-Valverde, S., Cuadros-Solas, P. J., & Rodríguez-Fernández, F. (2020). A machine learning approach to the digitalization of bank customers: Evidence from random and causal forests. *PLOS ONE*, 15(10), e0240362. <https://doi.org/10.1371/journal.pone.0240362>
- Cardarelli, R., Elekdag, S., & Lall, S. (2011). Financial stress and economic contractions. *Journal of Financial Stability*, 7(2), 78–97. <https://doi.org/10.1016/j.jfs.2010.01.005>
- Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 785–794. <https://doi.org/10.1145/2939672.2939785>
- Choudhary, P. K., Innan, N., & Shafique, M. (2026). HQNN-FSP: A hybrid classical-quantum neural network for regression-based financial stock market prediction. *Quantum Machine Intelligence*, 8, Article 55. <https://doi.org/10.1007/s42484-026-00390-9>
- Chui, M., Manyika, J., Miremadi, M., Henke, N., Chung, R., Nel, P., & Malhotra, S. (2018). *Notes from the AI frontier: Applications and value of deep learning*. McKinsey Global Institute Discussion Paper.
- Cont, R., & Schaanning, E. (2017). Fire sales, indirect contagion and systemic stress testing. *Norges Bank Working Paper*, 2017/2. <https://doi.org/10.2139/ssrn.2541114>
- Dandolo, D., Masiero, C., Carletti, M., Dalle Pezze, D., & Susto, G. A. (2023). AcME—Accelerated model-agnostic explanations: Fast whitening of the machine-learning black box. *Expert Systems with Applications*, 214, 119115. <https://doi.org/10.1016/j.eswa.2022.119115>
- Davenport, T. H., & Ronanki, R. (2018). Artificial intelligence for the real world. *Harvard Business Review*, 96(1), 108–116.

- Demirgüç-Kunt, A., Detragiache, E., & Tressel, T. (2008). Banking on the principles: Compliance with Basel core principles and bank soundness. *Journal of Financial Intermediation*, 17(4), 511–542. <https://doi.org/10.1016/j.jfi.2007.10.003>
- Drehmann, M., Borio, C., Gambacorta, L., Jiménez, G., & Trucharte, C. (2010). Countercyclical capital buffers: Exploring options. *BIS Working Papers*, 317. <https://doi.org/10.2139/ssrn.1648946>
- Ebert, S., (2021). Prudent Discounting: Experimental Evidence On Higher Order Time Risk Preferences. *International Economic Review*, 62(4), 1489-151. <https://doi.org/10.1111/iere.12521>
- Egger, D. J., Gambella, C., Marecek, J., McFaddin, S., Mevissen, M., Raymond, R., Simonetto, A., Woerner, S., & Yndurain, E. (2020). Quantum computing for finance: State-of-the-art and future prospects. *IEEE Transactions on Quantum Engineering*, 1, 3101724. <https://doi.org/10.1109/TQE.2020.3030314>
- Espinoza-Vega, M. A., & Solé, J. (2010). Cross-border financial surveillance: A network perspective. *IMF Working Paper*, WP/10/105.
- Faddeev, L. D., & Slavnov, A. A. (2018). *Gauge fields: An introduction to quantum theory*. Benjamin/Cummings.
- Fischer, T., & Krauss, C. (2018). Deep learning with long short-term memory networks for financial market predictions. *European Journal of Operational Research*, 270(2), 654–669. <https://doi.org/10.1016/j.ejor.2017.11.054>
- Friedman, J. H. (2001). Greedy function approximation: A gradient boosting machine. *The Annals of Statistics*, 29(5), 1189–1232. <https://doi.org/10.1214/aos/1013203451>
- Gai, P., & Kapadia, S. (2010). Contagion in financial networks. *Proceedings of the Royal Society A*, 466(2120), 2401–2423. <https://doi.org/10.1098/rspa.2009.0410>
- Goldstein, I., Jiang, W., & Karolyi, G. A. (2019). To FinTech and beyond. *The Review of Financial Studies*, 32(5), 1647–1661. <https://doi.org/10.1093/rfs/hhz025>
- Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press.
- Gu, S., Kelly, B., & Xiu, D. (2020). Empirical asset pricing via machine learning. *The Review of Financial Studies*, 33(5), 2223–2273. <https://doi.org/10.1093/rfs/hhaa009>

- Hałaj, G., Martinez-Jaramillo, S., & Battiston, S. (2024). Financial stability through the lens of complex systems. *Journal of Financial Stability*, 71, 101228. <https://doi.org/10.1016/j.jfs.2024.101228>
- Hao, W., Lefèvre, C., Tamturk, M., & Utev, S. (2019). Quantum option pricing and data analysis. *Quantitative Finance and Economics*, 3(3), 490–507. <https://doi.org/10.3934/QFE.2019.3.490>
- Havlíček, V., Córcoles, A. D., Temme, K., Harrow, A. W., Kandala, A., Chow, J. M., & Gambetta, J. M. (2019). Supervised learning with quantum-enhanced feature spaces. *Nature*, 567(7747), 209–212. <https://doi.org/10.1038/s41586-019-0980-2>
- Heaton, J. B., Polson, N. G., & Witte, J. H. (2017). Deep learning for finance: Deep portfolios. *Applied Stochastic Models in Business and Industry*, 33(1), 3–12. <https://doi.org/10.1002/asmb.2209>
- Hoffart, F. M., D'Orazio, P., Holz, F., & Kemfert, C. (2024). Exploring the interdependence of climate, finance, energy, and geopolitics. *Applied Energy*, 361, 122885. <https://doi.org/10.1016/j.apenergy.2024.122885>
- Hung, J.-L., He, W., & Shen, J. (2020). Big data analytics for supply chain relationship in banking. *Industrial Marketing Management*, 86, 144–153. <https://doi.org/10.1016/j.indmarman.2019.11.001>
- Kaminsky, G. L., & Reinhart, C. M. (1999). The twin crises: The causes of banking and balance-of-payments problems. *American Economic Review*, 89(3), 473–500. <https://doi.org/10.1257/aer.89.3.473>
- Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., Ye, Q., & Liu, T.-Y. (2017). LightGBM: A highly efficient gradient boosting decision tree. *Advances in Neural Information Processing Systems*, 30, 3146–3154.
- Khandani, A. E., Kim, A. J., & Lo, A. W. (2010). Consumer credit-risk models via machine-learning algorithms. *Journal of Banking & Finance*, 34(11), 2767–2787. <https://doi.org/10.1016/j.jbankfin.2010.06.001>
- Kingsly, P. K. M. (2025). Quantum finance and structural transformation of less developed countries. SSRN. <https://doi.org/10.2139/ssrn.5219462>
- Laeven, L., & Valencia, F. (2020). Systemic banking crises database II. *IMF Economic Review*, 68(2), 307–361. <https://doi.org/10.1057/s41308-020-00107-3>

- Lepetit, L., & Strobel, F. (2015). Bank insolvency risk and Z-score measures: A refinement. *Finance Research Letters*, 13, 214–224. <https://doi.org/10.1016/j.frl.2015.01.001>
- Lepetit, L., Nys, E., Rous, P., & Tarazi, A. (2008). Bank income structure and risk: An empirical analysis of European banks. *Journal of Banking & Finance*, 32(8), 1452–1467. <https://doi.org/10.1016/j.jbankfin.2007.12.002>
- Lessmann, S., Baesens, B., Seow, H.-V., & Thomas, L. C. (2015). Benchmarking state-of-the-art classification algorithms for credit scoring: An update of research. *European Journal of Operational Research*, 247(1), 124–136. <https://doi.org/10.1016/j.ejor.2015.05.030>
- Lundberg, S. M., Erion, G., Chen, H., DeGrave, A., Prutkin, J. M., Nair, B., Katz, R., Himmelfarb, J., Bansal, N., & Lee, S.-I. (2020). From local explanations to global understanding with explainable AI for trees. *Nature Machine Intelligence*, 2(1), 56–67. <https://doi.org/10.1038/s42256-019-0138-9>
- Maldonado, W. L., Tourinho, O. A. F., & Valli, M. (2012). Exchange rate bubbles: Fundamental value estimation and rational expectations test. *Journal of International Money and Finance*, 31(5), 1033–1059. <https://doi.org/10.1016/j.jimonfin.2011.12.009>
- Manyika, J., Chui, M., Brown, B., Bughin, J., Dobbs, R., Roxburgh, C., & Hung Byers, A. (2011). *Big data: The next frontier for innovation, competition, and productivity*. McKinsey Global Institute.
- Mishkin, F. S., & White, E. N. (2014). Unprecedented actions: The Federal Reserve's response to the global financial crisis. *The Federal Reserve's Role in the Global Economy*, 220–243. <https://doi.org/10.3386/w20737>
- Mugel, S., Kuchkovsky, C., Sánchez, E., Fernández-Lorenzo, S., Luis-Hita, J., Lizaso, E., & Orús, R. (2022). Dynamic portfolio optimization with real datasets using quantum processors and quantum-inspired tensor networks. *Physical Review Research*, 4(1), 013006. <https://doi.org/10.1103/PhysRevResearch.4.013006>
- Mullainathan, S., & Spiess, J. (2017). Machine learning: An applied econometric approach. *Journal of Economic Perspectives*, 31(2), 87–106. <https://doi.org/10.1257/jep.31.2.87>
- Ngonyama, N., Mgxekwa, B., & Sibanda, K. (2025). The impact of financial technology and cyber risk on non-bank financial intermediation. Springer. https://doi.org/10.1007/978-3-031-86224-3_11

- Orús, R., Mugel, S., & Lizaso, E. (2019). Quantum computing for finance: Overview and prospects. *Reviews in Physics*, 4, 100028. <https://doi.org/10.1016/j.revip.2019.100028>
- Peskin, M. E., & Schroeder, D. V. (1995). *An introduction to quantum field theory*. CRC Press. <https://doi.org/10.1201/9780429503559>
- Petropoulos, A., Siakoulis, V., Stavroulakis, E., & Vlachogiannakis, N. E. (2020). Predicting bank insolvencies using machine learning techniques. *International Journal of Forecasting*, 36(3), 1092–1113. <https://doi.org/10.1016/j.ijforecast.2019.11.005>
- Philippon, T. (2019). On FinTech and financial inclusion. NBER Working Paper, 26330. <https://doi.org/10.3386/w26330>
- Rebentrost, P., Gupt, B., & Bromley, T. R. (2018). Quantum computational finance: Monte Carlo pricing of financial derivatives. *Physical Review A*, 98(2), 022321. <https://doi.org/10.1103/PhysRevA.98.022321>
- Reinhart, C. M., & Rogoff, K. S. (2009). The aftermath of financial crises. *American Economic Review*, 99(2), 466–472. <https://doi.org/10.1257/aer.99.2.466>
- Reinhart, C. M., & Rogoff, K. S. (2014). Recovery from financial crises: Evidence from 100 episodes. *American Economic Review*, 104(5), 50–55. <https://doi.org/10.1257/aer.104.5.50>
- Schaeck, K., & Cihák, M. (2014). Competition, efficiency, and stability in banking. *Financial Management*, 43(1), 215–241. <https://doi.org/10.1111/fima.12010>
- Schuld, M., Sinayskiy, I., & Petruccione, F. (2015). An introduction to quantum machine learning. *Contemporary Physics*, 56(2), 172–185. <https://doi.org/10.1080/00107514.2014.964942>
- Sezer, O. B., Gudelek, M. U., & Ozbayoglu, A. M. (2020). Financial time series forecasting with deep learning: A systematic literature review: 2005-2019. *Applied Soft Computing*, 90, 106181. <https://doi.org/10.1016/j.asoc.2020.106181>
- Sidaoui, J., Ramos-Francia, M., & Cuadra, G. (2010). The global financial crisis and policy response in Mexico. *BIS Papers*, 54, 279–298. <https://www.bis.org/publ/bppdf/bispap54r.pdf>
- Sigrist, F., & Hirnschall, C. (2019). Grabit: Gradient tree-boosted Tobit models for default prediction. *Journal of Banking & Finance*, 102, 177–192. <https://doi.org/10.1016/j.jbankfin.2019.03.004>

- Stamatopoulos, N., Egger, D. J., Sun, Y., Zoufal, C., Iten, R., Shen, N., & Woerner, S. (2020). Option pricing using quantum computers. *Quantum*, 4, 291. <https://doi.org/10.22331/q-2020-07-06-291>
- Sun, J., Li, H., Huang, Q.-H., & He, K.-Y. (2014). Predicting financial distress and corporate failure: A review from the state-of-the-art definitions, modeling, sampling, and featuring approaches. *Knowledge-Based Systems*, 57, 41–56. <https://doi.org/10.1016/j.knosys.2013.12.006>
- Varian, H. R. (2014). Big data: New tricks for econometrics. *Journal of Economic Perspectives*, 28(2), 3–28. <https://doi.org/10.1257/jep.28.2.3>
- Wamba, S. F., Akter, S., Edwards, A., Chopin, G., & Gnanzou, D. (2015). How 'big data' can make big impact: Findings from a systematic review and a longitudinal case study. *International Journal of Production Economics*, 165, 234–246. <https://doi.org/10.1016/j.ijpe.2014.12.031>
- Zhou, J. (2025). Quantum finance: Exploring the implications of quantum computing on financial models. *Computational Economics*. <https://doi.org/10.1007/s10614-025-10894-4>