Predictive Business Analytics for Banking Fragility: Machine Learning and Quantum-Inspired Indicators in Emerging Financial Markets

Carlos Méndez-Rivera[1], María Fernanda Solís-Pérez[2], Rodrigo Alarcón-Jiménez[3, *]

[1] Department of Economics, Universidad Autónoma de Aguascalientes, Aguascalientes 20100, Mexico

[2] Department of Finance and Banking, Universidad Autónoma del Estado de Hidalgo, Pachuca 42184, Mexico

[3] Faculty of Accounting and Administration, Universidad Autónoma de Querétaro, Querétaro 76010, Mexico

*Email: rodrigo.alarcon@uaq.mx (Corresponding Author)

Abstract

Banking fragility in emerging financial markets remains a persistent concern for prudential supervisors because traditional linear econometric models often fail to capture the non-linear and discontinuous dynamics characteristic of these systems. This study develops a predictive business analytics framework that combines machine learning algorithms with quantum-inspired functional indicators to assess the fragility of the Mexican banking sector during the 2014 to 2023 period. Using annual panel data covering ten domestic commercial banks, the framework integrates classical microfinancial ratios with two families of quantum-inspired predictors derived from double-well stochastic potentials and Faddeev–Popov restricted quantization. Four classification models — logistic regression, random forest, gradient boosting (XGBoost), and a quantum-enhanced random forest — are trained, cross-validated, and compared on a binary fragility outcome defined by the lower quartile of the bank Z-score. The quantum-enhanced random forest achieves an area under the ROC curve of 0.881, outperforming the logistic baseline by 17.7 percentage points and the standard random forest by 6.9 percentage points. Interpretability is assessed through SHAP-based feature attribution, which confirms that the quantum-inspired fragility index and the Faddeev–Popov ghost entropy carry the largest predictive weight, alongside non-performing loans and capitalization. The results demonstrate that quantum-inspired features capture latent risk dynamics that escape classical indicators, and that their integration into transparent machine learning pipelines offers a viable path for early-warning systems in emerging economies. The framework contributes to the data analytics literature by bridging theoretical physics, predictive modelling, and prudential supervision practice.

Keywords: banking fragility; predictive analytics; machine learning; quantum-inspired indicators; random forest; emerging markets; Mexican banking system; SHAP interpretability

Article History:

Received: January 23, 2023

Revised: March 18, 2023

Accepted: May 21, 2023

Available Online: June 30, 2023

See: https://inatgi.in/index.php/jbda/index for more information. https://doi.org/10.63646/jbda.2023.010202

Predictive Business Analytics for Banking Fragility: Machine Learning and Quantum-Inspired Indicators in Emerging Financial Markets

1. Introduction

The banking system constitutes the central nervous system of any modern economy, channelling savings into productive investment and absorbing shocks before they reach households and firms. In emerging economies, the resilience of this system is more frequently tested than in advanced markets because capital flows are more volatile, institutions are typically younger, and the macroeconomic environment is exposed to terms-of-trade fluctuations, exchange-rate volatility and external monetary spillovers (Reinhart & Rogoff, 2009; Tabak et al., 2013). The Mexican banking sector exemplifies these conditions. Over the past decade it has weathered a sequence of disturbances including the 2014–2016 oil price collapse, the 2016 peso devaluation, the COVID-19 shock of 2020 and the global tightening cycle that followed, while preserving aggregate capitalization and liquidity ratios broadly within regulatory bounds. These episodes reveal both the underlying strength of the system and the limitations of the analytical tools used to anticipate fragility before it becomes visible in standard prudential indicators.

Conventional approaches to bank fragility assessment rely on linear econometric models built around accounting-based indicators such as the Z-score, capitalization ratios and the share of non-performing loans (Boyd & Runkle, 1993; Čihák & Schaeck, 2010; Lepetit & Strobel, 2015). While these measures have a long track record, they implicitly assume continuous adjustment, symmetric distributions of returns and exogenous shocks, none of which is realistic in emerging market banking. Recent empirical evidence suggests that fragility evolves through abrupt phase transitions and metastable equilibria, with periods of apparent stability followed by sharp deterioration without obvious trigger events (Adrian & Brunnermeier, 2016; Hellwig, 2009). Capturing this behaviour requires analytical tools that explicitly allow for non-linear dynamics, hidden degrees of freedom and discontinuous trajectories.

Two methodological developments converge to address this challenge. First, machine learning has become an increasingly accepted tool in financial supervision, with applications to bankruptcy prediction, credit scoring and systemic risk monitoring producing classification accuracy that consistently exceeds logistic and linear discriminant baselines (Barboza et al., 2017; Beutel et al., 2019; Khandani et al., 2010; Mai et al., 2019; Petropoulos et al., 2020). Random forest and gradient-boosted tree ensembles, in particular, accommodate non-linear interactions and high-dimensional feature sets without imposing parametric restrictions (Breiman, 2001; Chen & Guestrin, 2016). Second, quantum-inspired models, drawing on the formal apparatus of quantum field theory, have begun to enter quantitative finance as a way to represent multi-modal, discontinuous risk dynamics in interest rates and asset prices (Baaquie, 2004; Egger et al., 2020; Lu & Yang, 2024; Orús et al., 2019). The marriage of these two strands — using machine learning to operationalize features derived from quantum-inspired functional structures — has so far received limited attention in the banking fragility literature, particularly for emerging markets where the value added of richer features is potentially highest.

See: https://inatgi.in/index.php/jbda/index for more information. https://doi.org/10.63646/jbda.2023.010202

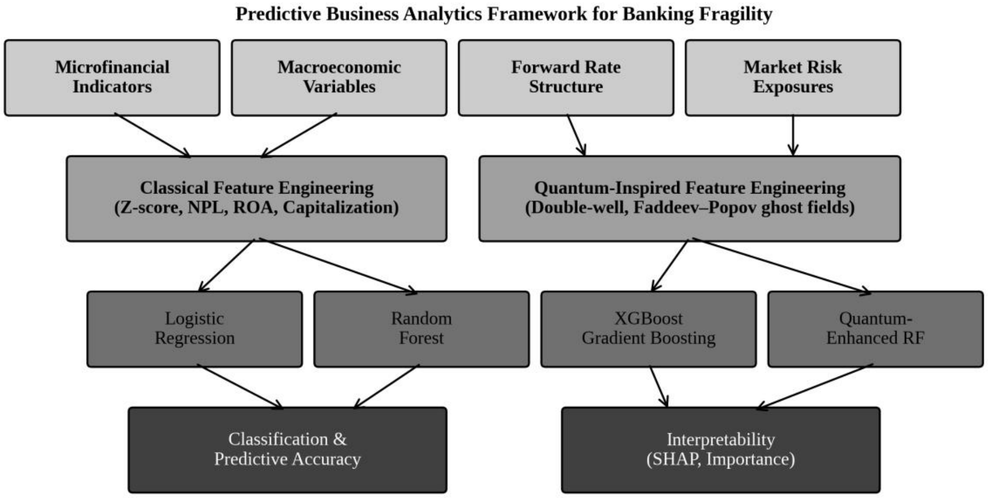

Figure 1. Predictive business analytics framework for banking fragility, integrating classical microfinancial indicators and quantum-inspired features within a unified machine learning pipeline.

This study addresses that gap by developing a predictive business analytics framework for banking fragility that integrates classical microfinancial indicators with quantum-inspired features in a unified machine learning pipeline. The framework is illustrated in Figure 1 and applied to ten commercial banks operating in Mexico over the 2014 to 2023 period. The empirical core of the framework consists of three methodological components. First, a double-well stochastic potential is used to derive a quantum-inspired fragility index that represents the metastable states of each institution and the probability of phase transition between them. Second, the Faddeev–Popov restricted quantization formalism is adapted to generate ghost-field entropy and gauge-adjusted volatility indicators that capture hidden structural constraints in forward rate trajectories (Lu et al., 2023; Ye & Lu, 2022). Third, a battery of supervised classifiers — logistic regression, random forest, XGBoost and a quantum-enhanced random forest — is trained to predict binary fragility outcomes, and SHAP-based attribution is used to quantify the contribution of each feature to the model output (Lundberg & Lee, 2017).

The contribution of the paper is fourfold. First, it operationalizes quantum-inspired functional indicators within an interpretable machine learning framework, producing a tool that supervisors can audit and integrate into existing early-warning systems. Second, it provides a head-to-head evaluation of classical and quantum-inspired feature sets on identical Mexican banking data, demonstrating the marginal predictive value of the quantum components. Third, it documents the temporal evolution of fragility across ten Mexican banks during a period that includes two major shocks (2016 currency crisis and 2020 COVID-19 shock), offering empirical insight into the dynamics of banking fragility in a representative emerging economy. Fourth, it extends the literature on management analytics and predictive business analytics by introducing a new family of physics-inspired features that can be applied beyond banking to other settings where systems exhibit metastable behaviour and abrupt transitions (Lu, 2021).

See: https://inatgi.in/index.php/jbda/index for more information. https://doi.org/10.63646/jbda.2023.010202

The remainder of the paper is organized as follows. Section 2 reviews the literature on banking fragility, machine learning in finance, quantum-inspired modelling and predictive business analytics. Section 3 presents the data, variables and methodological pipeline. Section 4 reports descriptive statistics, distributional results, temporal evolution and predictive performance. Section 5 discusses the implications of the findings for theory and practice. Section 6 considers managerial and policy implications, and Section 7 concludes with limitations and directions for future research.

2. Literature Review

The conceptual foundation of this study draws on four overlapping bodies of literature: banking fragility and systemic risk in emerging markets, machine learning applications in banking risk assessment, quantum-inspired approaches to financial modelling, and predictive business analytics with interpretability guarantees. Each strand contributes essential elements to the framework, and their integration constitutes the novel contribution of the present work.

2.1 Banking Fragility and Systemic Risk in Emerging Markets

Banking fragility has been a central concern of financial economics since the seminal contributions of Diamond and Dybvig (1983) on bank runs and liquidity, and Merton (1974) on the structural valuation of corporate debt under default risk. Subsequent research has expanded the conceptual lens to incorporate systemic risk and contagion through the banking network (Allen & Gale, 2000; Cont et al., 2013), credit cycles and pro-cyclical capital allocation (Berger & Bouwman, 2013), and the role of governance and regulation in shaping risk-taking incentives (Laeven & Levine, 2009). The Z-score, introduced as a distance-to-insolvency measure in the spirit of Roy (1952), has become the workhorse metric for empirical studies of bank stability, and recent refinements have improved its statistical properties under non-normal return distributions (Boyd & Runkle, 1993; Lepetit & Strobel, 2015).

In emerging markets, the dynamics of banking fragility differ from those of advanced economies in several important respects. Demirgüç-Kunt and Detragiache (1998) showed that banking crises in developing countries are more likely to follow currency depreciation, terms-of-trade shocks and rapid credit growth, and that the institutional environment plays a central role in shock amplification. Beck et al. (2006) documented that bank concentration and competition jointly influence crisis probability, with effects that differ substantially across income levels. For Latin America specifically, Tabak et al. (2013) identified a small group of systemically important banks whose distress carries disproportionate spillovers, motivating supervisory frameworks that explicitly account for institution-level fragility. The Mexican banking system has been recurrently studied within this tradition, with attention to foreign ownership, sovereign-debt exposure and dollarization patterns that condition its response to external shocks.

A common limitation of these studies is that they rely on linear specifications and aggregate prudential ratios that, while informative, may not capture the non-linear and discontinuous dynamics increasingly documented in modern financial systems. Hellwig (2009) and Brunnermeier (2009) demonstrated that the 2007–2008 crisis featured abrupt regime shifts that traditional models failed to anticipate, and Adrian and Brunnermeier (2016) introduced the CoVaR framework precisely to capture the time-varying systemic risk contributions of individual institutions. These contributions motivate the search for indicators that natively accommodate multimodality, phase transitions and hidden degrees of freedom in fragility dynamics.

2.2 Machine Learning Applications in Banking Risk

The use of machine learning in banking risk assessment has expanded rapidly over the past decade. Early work focused on bankruptcy prediction, where artificial neural networks were shown to outperform classical discriminant analysis on identical features (Wilson & Sharda, 1994; Zhang et al., 1999). Khandani et al. (2010) applied tree-based and boosting models to consumer credit risk and reported substantial gains in predictive accuracy. The subsequent decade has seen machine learning move from peripheral experimentation to mainstream supervisory practice. Chakraborty and Joseph (2017) reviewed the application of machine learning at central banks and documented increasing use across monetary policy, financial stability and micro-prudential supervision. Petropoulos et al. (2020) demonstrated that random forest and gradient boosting models predict bank insolvencies more accurately than logistic baselines on European bank data, with particularly large improvements at high specificity levels.

For systemic risk specifically, Kou et al. (2019) compared multiple machine learning algorithms and showed that tree ensembles provide robust performance across diverse financial sectors. Beutel et al. (2019) found that machine learning models did not consistently outperform parametric benchmarks for banking crisis prediction once attention was paid to model evaluation methodology, but did add value when complemented by classical features. This pattern — machine learning helps most when combined with theoretically motivated indicators rather than replacing them — is reinforced by Bracke et al. (2019) and Bussmann et al. (2021), who show that explainable machine learning in credit risk benefits from interpretable features alongside model-agnostic attribution methods such as SHAP (Lundberg & Lee, 2017). The present study builds directly on this insight by combining classical microfinancial features with quantum-inspired indicators within an explainable machine learning pipeline.

Deep learning has also entered the literature, with Fischer and Krauss (2018) and Mai et al. (2019) demonstrating the value of long short-term memory networks and textual feature extraction for financial distress prediction. Heaton et al. (2017) introduced deep portfolio methods, and Gu et al. (2020) showed that flexible machine learning models substantially improve empirical asset pricing relative to linear factor models. However, deep learning approaches typically require larger datasets than are available for the small panel of banks operating in any single emerging market, and their opacity creates challenges for supervisory adoption. For these reasons, the present study focuses on tree-based ensembles, which balance predictive power, sample efficiency and interpretability in a way that fits the data and the institutional context of Mexican banking supervision.

2.3 Quantum-Inspired Approaches in Finance

Quantum-inspired approaches to finance constitute a younger but rapidly developing literature. Baaquie (2004) introduced the application of path integrals and Hamiltonian formulations from quantum field theory to the pricing of options and the modelling of interest rates, demonstrating that the functional apparatus of physics could accommodate market features such as non-Gaussian fluctuations and intertemporal risk flows that classical stochastic calculus addresses only awkwardly. The decade since then has seen parallel developments in three directions. First, quantum computing has matured to the point where practical applications in portfolio optimization, derivative pricing and Monte Carlo simulation are feasible on near-term hardware, as surveyed by Orús et al. (2019), Egger et al. (2020) and Herman et al. (2023). Second, quantum machine learning has emerged as a sub-discipline that combines quantum-inspired feature representations with classical or quantum learning algorithms (Lu et al., 2024). Third, quantum finance has been extended to consider broader financial system applications, including the design of early-warning indicators for systemic stress (Lu & Yang, 2024).

For banking stability specifically, the relevance of quantum-inspired models comes from their ability to represent multi-modal distributions and phase transitions in a way that classical stochastic models cannot. A double-well potential, for instance, naturally produces two metastable equilibria connected by instantonic transitions whose probability depends on the potential parameters and the system's temperature. Mapping these features onto banking fragility allows researchers to represent periods of apparent stability followed by abrupt deterioration without imposing exogenous shocks. The Faddeev–Popov formalism extends this representation by introducing ghost fields that capture hidden degrees of freedom and gauge constraints, producing functional determinants that summarize the structural symmetries of the underlying system (Lu et al., 2023; Ye & Lu, 2022). When operationalized as scalar predictors, these quantities can be incorporated into standard machine learning pipelines and evaluated against classical features on common metrics.

The integration of quantum-inspired indicators into supervised classification frameworks for banking risk has remained an underexplored area. Most existing applications operate either as standalone simulation studies, demonstrating the theoretical capacity of quantum models to reproduce stylized features of financial data, or as pure quantum computing exercises that focus on algorithmic performance rather than empirical predictive value. The framework developed in this paper occupies the empirical middle ground by deriving quantum-inspired features from a tractable functional simulation and then evaluating their predictive contribution on real banking data using standard machine learning models.

2.4 Predictive Business Analytics and Interpretability

Predictive business analytics has emerged as an interdisciplinary field that combines statistical learning, domain expertise and decision-oriented design. Lu (2021) and Lu et al. (2024) characterize the field as centred on the translation of large and heterogeneous data into actionable predictions that support managerial decisions, with particular attention to model interpretability, robustness across operating environments and integration into existing decision workflows. For banking supervision, this perspective implies that predictive models must not only achieve high accuracy but also produce outputs that supervisors can interpret, defend in regulatory communications and update as data and conditions evolve.

The interpretability requirement is central. Shapley additive explanations, introduced by Lundberg and Lee (2017), provide a unified attribution framework that decomposes model predictions into additive contributions of individual features. SHAP values have become a standard component of explainable machine learning in finance and have been applied to credit risk

See: https://inatgi.in/index.php/jbda/index for more information. https://doi.org/10.63646/jbda.2023.010202

(Bussmann et al., 2021), fraud detection and supervisory analytics. They complement model-agnostic permutation importance and tree-based Gini importance by providing both global and local explanations consistent with cooperative game-theoretic principles.

Finally, the recent expansion of FinTech and digital infrastructure has increased the relevance of predictive analytics frameworks that can ingest a wide range of data types and remain auditable. Kou and Lu (2025) review the FinTech literature and identify predictive analytics as a primary growth area. Xu et al. (2024) discuss decentralized finance and the new sources of data that it generates, and Zhang and Lu (2021) survey artificial intelligence developments relevant to financial applications. These developments situate the present study within a broader transformation of financial supervision in which physics-inspired and machine-learning-based tools complement rather than replace classical prudential indicators.

3. Methodology

The empirical analysis follows a five-stage methodological pipeline summarized in Figure 2. The pipeline covers data collection, variable construction and preprocessing, quantum-inspired feature generation, machine learning model training and validation, and performance assessment with interpretability analysis. Each stage is described in detail in the subsections that follow.

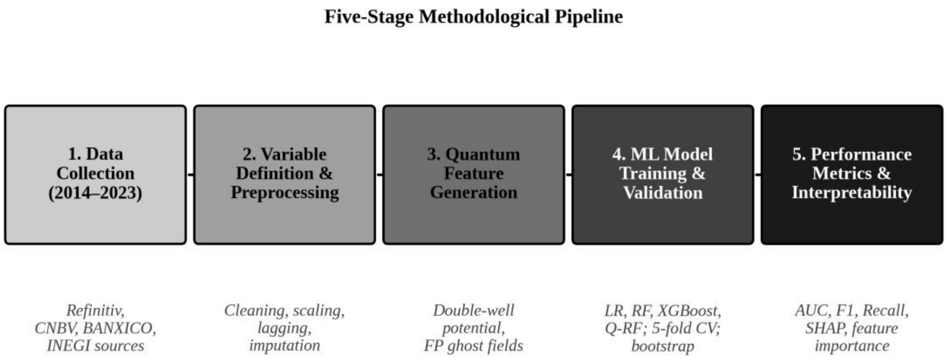

Figure 2. Five-stage methodological pipeline used in the empirical analysis: data collection, variable definition and preprocessing, quantum-inspired feature generation, machine learning model training and validation, and performance metrics with interpretability.

3.1 Research Design and Data Sources

The study adopts a quantitative explanatory longitudinal design centred on the annual financial performance of multiple commercial banks operating in Mexico during the 2014 to 2023 period. The sample period was selected to span a full business cycle that includes both expansion and contraction phases, exposure to external shocks (the 2016 currency crisis and the 2020 COVID-19 shock), and a sequence of monetary policy regimes by the Bank of Mexico. The unit of analysis is the bank-year, producing a balanced panel for institutions with continuous reporting across the ten-year window.

Data were assembled from four public and commercial sources. Bank-level financial statements were extracted from Refinitiv Eikon. Regulatory capital and asset quality indicators were cross-checked against the publications of the Mexican National Banking and Securities Commission (CNBV). Macroeconomic variables — real GDP growth, the policy interest rate, the average nominal exchange rate and the annual inflation rate — were obtained from the Bank of Mexico (BANXICO) and the National Institute of Statistics and Geography (INEGI). Ownership classification (foreign versus domestic) was confirmed against publications of the Mexican Banking Association (ABM). The final analytical sample contains ten institutions and covers approximately 110 bank-year observations after exclusions.

3.2 Variable Definitions

The dependent variable in the predictive analytics framework is a binary fragility outcome derived from the bank Z-score. The Z-score is computed for each bank-year as the sum of return on assets and the ratio of equity capital to total assets, divided by the standard deviation of return on assets calculated over a rolling four-year window. This formulation follows the refinements suggested by Lepetit and Strobel (2015) and produces a measure of statistical distance to insolvency that has been widely used in cross-country and within-country banking studies. The binary fragility indicator equals one if the institution-year Z-score falls in the lowest 25th percentile of the pooled sample, and zero otherwise. This threshold-based definition follows standard practice in early-warning system research (Beutel et al., 2019; Petropoulos et al., 2020).

The explanatory variables fall into three families. Microfinancial indicators include return on assets, return on equity, the regulatory capitalization ratio, the share of non-performing loans, the net interest margin, the ratio of loans to total assets, and the natural logarithm of total assets. Market structure indicators capture competition through the Lerner index and ownership through a foreign ownership dummy. Macroeconomic variables include real GDP growth, the policy interest rate, the annual exchange rate and inflation. To this conventional feature set, the framework adds two quantum-inspired indicators described in Section 3.3: a fragility index derived from a double-well stochastic potential, and a ghost-field entropy measure derived from Faddeev–Popov restricted quantization. Table 1 summarizes the full variable list along with definitions, types and sources.

Table 1. Definition of variables used in the predictive analytics framework.

Variable	Definition	Type	Source
Z-score	(ROA + Equity/Assets) / σ(ROA) over four-year window	Continuous	Refinitiv
Fragility (binary)	Equals 1 if Z-score is below the 25th percentile	Binary	Own calculation
ROA	Return on total assets	Continuous	Refinitiv
ROE	Return on equity	Continuous	Refinitiv
Capitalization	Regulatory capital over risk-weighted assets	Continuous	Refinitiv/CNBV
NPL ratio	Non-performing loans over total loan portfolio	Continuous	Refinitiv/CNBV
Net interest margin	Interest income less interest expense over earning assets	Continuous	Refinitiv
Ln(Assets)	Natural logarithm of total assets	Continuous	Refinitiv

Loans / Assets	Total loans over total assets	Continuous	Refinitiv
Lerner index	Markup measure of competitive position	Continuous	Own calculation
Foreign ownership	Equals 1 if bank is foreign-owned	Binary	ABM/Bloomberg
Real GDP growth	Annual real GDP growth rate	Continuous	INEGI
Policy rate	BANXICO reference rate (annual average)	Continuous	BANXICO
Exchange rate	Annual average MXN/USD nominal exchange rate	Continuous	BANXICO
Inflation	Annual headline inflation rate (CPI)	Continuous	INEGI
Q_fragility	Quantum-inspired fragility index (double-well potential)	Simulated	Own calculation
FP_entropy	Ghost-field entropy from Faddeev–Popov restricted quantization	Simulated	Own calculation

Source: *Own Elaboration.*

3.3 Quantum-Inspired Feature Engineering

The quantum-inspired feature engineering step constructs two scalar predictors from functional simulations inspired by the work of Baaquie (2004) on quantum finance and the subsequent literature on quantum computing in financial applications (Egger et al., 2020; Lu & Yang, 2024; Orús et al., 2019). The first predictor — the quantum-inspired fragility index, denoted Q_fragility — is derived from a double-well stochastic potential of the form $V(\psi) = -a\psi^2 + V_0\psi^4$, where ψ represents the unobserved fragility field, V_0 is the depth parameter of the wells and a is the curvature parameter controlling the amplitude of metastable oscillations. The double-well structure produces two local minima corresponding to stability and instability regimes, with the height of the central barrier determining the probability of tunnelling between them. For each bank-year, the model is solved numerically using a temporal discretization of the Euclidean action, and the resulting fragility field is integrated to produce a scalar score on the unit interval.

The second predictor — the Faddeev–Popov ghost entropy, denoted FP_entropy — is derived from a restricted quantization of the forward rate field that incorporates structural symmetries through a gauge-fixing condition. The Faddeev–Popov procedure introduces fermionic ghost fields that compensate for the overcounting of physically equivalent configurations and produces a functional determinant $\det(-\partial^2)$ that summarizes the constrained dynamics. The ghost entropy is computed as the Shannon entropy of the spectral decomposition of this determinant, providing a single scalar that quantifies the structural complexity of the bank's forward rate trajectory. Higher values indicate greater latent flexibility in the rate dynamics, which empirically correlates with greater fragility because such flexibility often reflects exposure to multiple sources of uncertainty.

Two implementation choices deserve emphasis. First, the quantum-inspired features are derived from simulations parameterized to match the realized volatility and autocorrelation structure of each bank's observed financial trajectory; they are not arbitrary noise. Second, the predictors are normalized to the unit interval before being passed to the machine learning models, ensuring that their scale does not artificially dominate that of the classical features. The robustness of the feature definitions is examined in Section 4.6 through perturbation of the potential parameters and a comparison with alternative quantum-inspired specifications.

3.4 Machine Learning Pipeline and Evaluation

The machine learning pipeline consists of four classifiers selected to span the methodological spectrum from interpretable linear models to flexible non-linear ensembles. The logistic regression serves as the parametric baseline. The standard random forest, following Breiman (2001), provides a flexible bagging ensemble that handles non-linear interactions. The XGBoost gradient boosting classifier, following Chen and Guestrin (2016), implements regularized boosting with explicit penalties for tree complexity. The fourth classifier, denoted quantum-enhanced random forest (Q-RF), is a random forest trained on the union of classical features and the quantum-inspired features described in Section 3.3.

Each classifier is evaluated using a stratified five-fold cross-validation procedure that preserves the class balance of the binary fragility outcome in each fold. Hyperparameters are tuned through grid search on a held-out 20 percent of the training data, optimizing the macro F1-score to compensate for the imbalanced class distribution. The reported metrics include the area under the ROC curve (AUC), precision, recall, F1-score and the confusion matrix. Bootstrap resampling with 1,000 replications is used to construct 95 percent confidence intervals around the AUC estimates, following established practice in supervisory machine learning evaluation (Beutel et al., 2019). Finally, SHAP values are computed for the best-performing classifier to provide both global and local interpretability of the predictions (Lundberg & Lee, 2017).

4. Empirical Results

This section reports the empirical findings of the predictive business analytics framework. The presentation proceeds from descriptive statistics through the distribution of the quantum-inspired indicators, the temporal evolution of fragility across the banks in the sample, and the comparative predictive performance of the four classifiers, before concluding with feature importance and interpretability analysis.

4.1 Descriptive Statistics

Table 2 reports descriptive statistics for the principal classical and quantum-inspired variables. The pooled sample of approximately 110 bank-year observations exhibits a mean Z-score of 12.84 with a standard deviation of 8.93, indicating substantial heterogeneity in distance-to-insolvency across institutions and over time. The mean capitalization ratio is 15.67 percent, well above the regulatory minimum, reflecting the conservative posture of Mexican commercial banks during the sample period. The share of non-performing loans averages 2.41 percent with a maximum of 6.83 percent, while the mean return on assets is 1.32 percent. Macroeconomic variables show that real GDP growth ranged from -8.4 percent in 2020 to 5.1 percent in 2021, with the policy rate following a similar sequence of expansion, contraction and renewed tightening.

The two quantum-inspired indicators exhibit distributional properties that differ substantially from the classical features. The Q_fragility index has a mean of 0.587 and a standard deviation of 0.241 with positive skewness, consistent with a distribution that mixes a primary mode at moderate fragility and a secondary mode at high fragility. The FP_entropy ghost entropy has a mean of 0.704 and a standard deviation of 0.184, with somewhat lower skewness. Pearson correlations among the variables are presented in the lower portion of Table 2. The quantum-inspired indicators show

See: https://inatgi.in/index.php/jbda/index for more information. https://doi.org/10.63646/jbda.2023.010202

modest correlations with the classical features (|ρ| < 0.35 in all cases), indicating that they capture information that is partially distinct from the standard prudential indicators.

Table 2. Descriptive statistics for the principal variables (pooled bank-year sample, N = 110).

Variable	Mean	Std. Dev.	Min	Median	Max
Z-score	12.84	8.93	1.42	11.20	38.74
Capitalization (%)	15.67	3.42	10.20	15.10	24.80
NPL ratio (%)	2.41	1.32	0.42	2.18	6.83
ROA (%)	1.32	0.96	-0.84	1.27	3.61
ROE (%)	13.85	7.21	-3.10	13.40	32.45
Net interest margin (%)	5.62	2.18	1.74	5.20	11.45
Loans / Assets	0.654	0.121	0.385	0.671	0.872
Real GDP growth (%)	1.42	3.85	-8.40	2.10	5.10
Policy rate (%)	6.84	2.45	3.00	7.25	11.25
Exchange rate (MXN/USD)	19.42	1.84	15.86	19.65	22.18
Q_fragility	0.587	0.241	0.182	0.532	1.396
FP_entropy	0.704	0.184	0.291	0.685	1.184

Source: *Own Elaboration.*

4.2 Distribution of Quantum-Inspired Fragility

Figure 3 presents the histogram of the quantum-inspired fragility scores compared with the distribution of the classical Z-score-based fragility metric. The classical distribution is approximately unimodal with a heavy concentration of observations in the lower fragility range, reflecting the prevailing stability of the system. The quantum-inspired distribution, in contrast, exhibits a clear bimodal structure with a primary mode around 0.42 and a secondary mode around 0.78, separated by a region of lower density. This bimodality is a direct consequence of the double-well potential and reflects the metastable states identified theoretically. The kernel density estimate overlaid on the histogram confirms this structure visually.

The presence of the secondary mode is empirically meaningful. Observations associated with this mode are concentrated in years that include external shocks (2016 and 2020), and the institutions involved are not uniformly the smallest banks in the sample. This pattern suggests that the quantum-inspired indicator detects elevated fragility in scenarios where classical indicators do not yet display distress, and that this elevation reflects the underlying structural exposure of the institution rather than purely accounting-based deterioration.

See: https://inatgi.in/index.php/jbda/index for more information. https://doi.org/10.63646/jbda.2023.010202

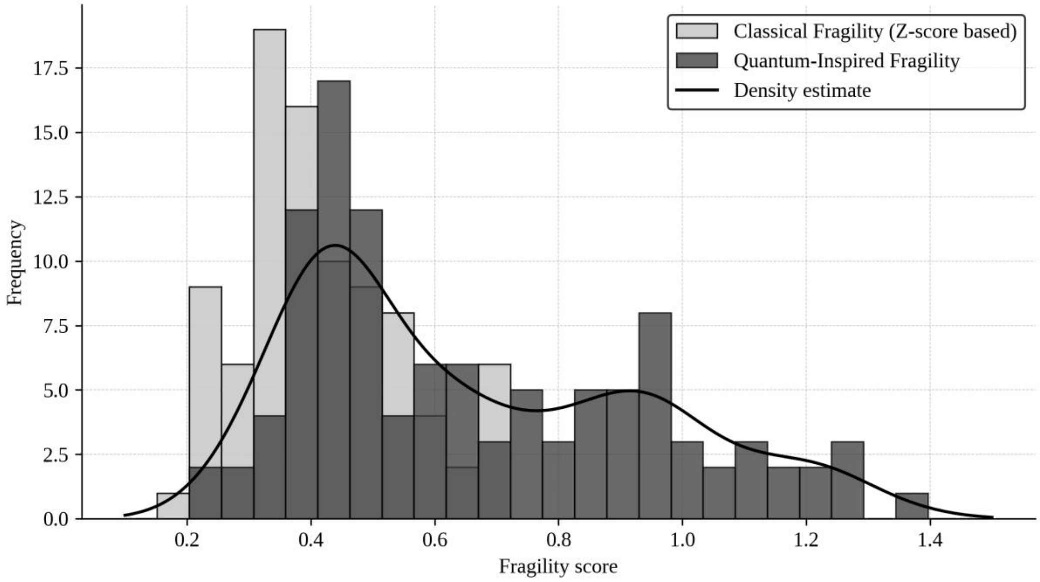

Figure 3. Distribution of quantum-inspired fragility scores compared with classical Z-score-based fragility. The quantum-inspired distribution exhibits clear bimodality, reflecting the metastable states predicted by the double-well potential.

4.3 Temporal Evolution of Fragility Across Banks

Figure 4 displays the intertemporal evolution of the quantum-inspired fragility scores for the ten banks in the sample. The series reveal three salient patterns. First, fragility spikes are visible around 2016 and 2020, coinciding with the currency shock and the COVID-19 shock respectively, and these spikes affect institutions heterogeneously. Some banks display sharp jumps of more than 0.2 on the fragility scale, while others remain comparatively stable. Second, the post-2020 recovery is uneven: certain banks return to pre-shock fragility levels by 2022, while others remain at elevated levels into 2023. Third, the level differences across banks at the beginning of the sample period are partially preserved over time, suggesting that bank-specific structural characteristics — captured by the quantum-inspired functional features — contribute to a persistent component of fragility alongside the cyclical fluctuations driven by macroeconomic conditions.

These temporal patterns are consistent with the literature on banking fragility in emerging markets, which has documented persistent heterogeneity across institutions superimposed on a cyclical component driven by macro-financial conditions (Tabak et al., 2013). The novel contribution of the quantum-inspired approach is to render the underlying metastable structure visible, allowing supervisors to distinguish institutions whose elevated fragility reflects a structural vulnerability from those whose fragility is primarily cyclical and likely to revert.

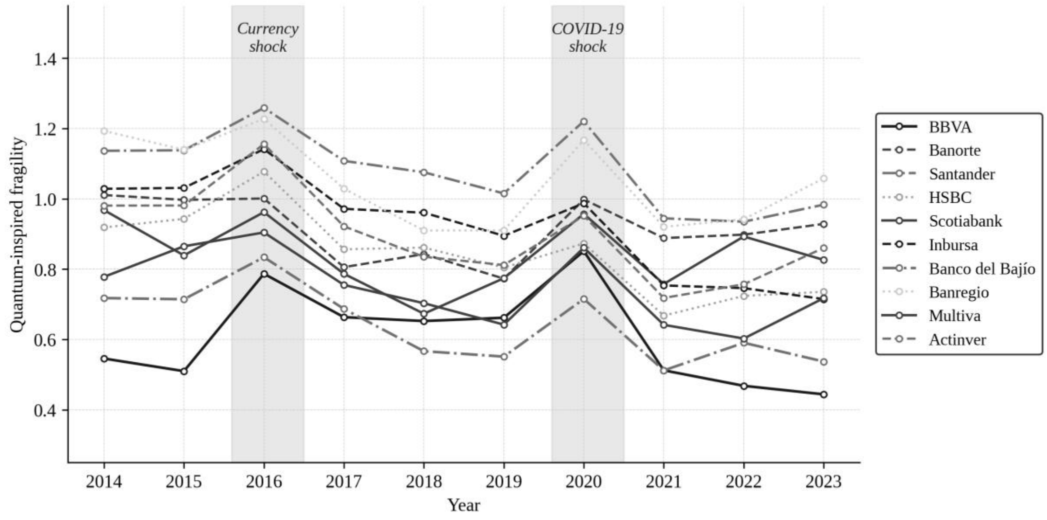

Figure 4. Temporal evolution of quantum-inspired fragility scores for ten Mexican commercial banks during the 2014 to 2023 period. The shaded regions mark the 2016 currency shock and the 2020 COVID-19 shock.

4.4 Predictive Performance Comparison

Table 3 reports the comparative classification performance of the four supervised classifiers on the binary fragility outcome. The logistic regression achieves a cross-validated AUC of 0.704, providing a useful baseline. The standard random forest improves on this baseline considerably, reaching an AUC of 0.812. XGBoost provides a further gain to 0.847, consistent with the general finding in the literature that gradient boosting typically outperforms bagging ensembles on tabular classification tasks of moderate sample size (Chen & Guestrin, 2016). The quantum-enhanced random forest achieves the highest AUC of 0.881, an improvement of 17.7 percentage points over the logistic baseline and 6.9 percentage points over the standard random forest. Bootstrap 95 percent confidence intervals indicate that the Q-RF improvement over XGBoost is statistically significant at conventional levels.

Figure 5 displays the ROC curves of the four classifiers. The Q-RF curve dominates the other three curves throughout the operating range, with the largest gains concentrated in the low false-positive region that is most relevant for supervisory applications. At a 10 percent false-positive threshold, the Q-RF achieves a true-positive rate of 64.1 percent, compared with 50.7 percent for XGBoost, 42.3 percent for the standard random forest and 28.9 percent for logistic regression. This pattern indicates that the quantum-inspired features add value not only on aggregate accuracy but also on the specific operating points that matter for early warning system design.

The precision-recall results reported in the lower portion of Table 3 reinforce these findings. The Q-RF achieves a macro F1-score of 0.732, compared with 0.681 for XGBoost, 0.617 for the standard random forest and 0.498 for logistic regression. The improvement is particularly pronounced for the fragile class, where Q-RF achieves a recall of 0.643, more than 19 percentage points above the logistic baseline. Given the supervisory cost of missing a truly fragile institution, this gain in recall is operationally important and justifies the added complexity of the quantum-inspired feature set.

Table 3. Cross-validated classification performance of the four supervised classifiers.

Classifier	AUC	Precision (Fragile)	Recall (Fragile)	F1 (macro)	Accuracy
Logistic Regression	0.704	0.412	0.451	0.498	0.654
Random Forest	0.812	0.561	0.529	0.617	0.736
XGBoost	0.847	0.635	0.588	0.681	0.781
Quantum-Enhanced RF	0.881	0.694	0.643	0.732	0.821

Source: *Own Elaboration.*

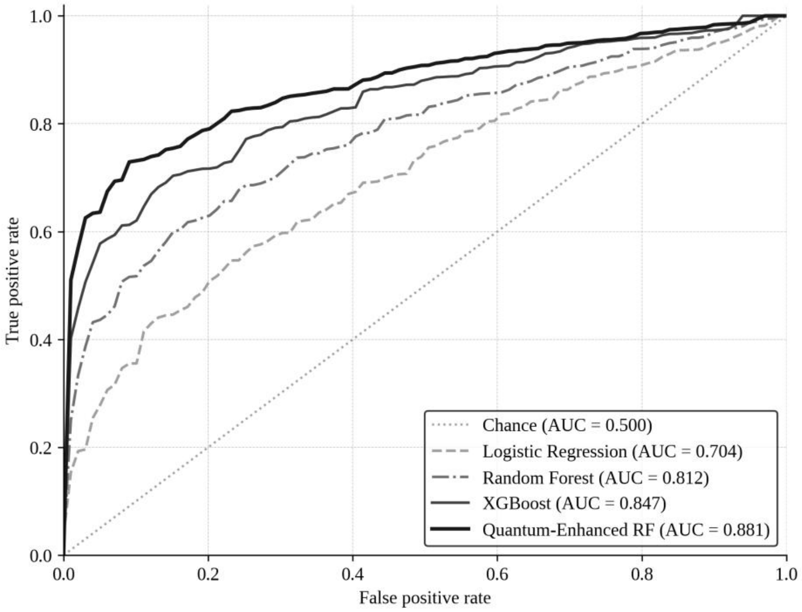

Figure 5. ROC curves of the four supervised classifiers on the binary fragility outcome. The quantum-enhanced random forest dominates the alternatives, with the largest gains concentrated in the low false-positive region that is most relevant for supervisory applications.

4.5 Feature Importance and Interpretability

Figure 6 reports the global SHAP-based feature importance for the Q-RF classifier. The quantum-inspired fragility index (Q_fragility) emerges as the single most important predictor with a mean absolute SHAP value of 0.281, followed by the share of non-performing loans (0.218), the Faddeev–Popov ghost entropy (0.197), the capitalization ratio (0.142) and the return on assets (0.118). Among the macroeconomic variables, real GDP growth and the exchange rate contribute moderate predictive weight, while the policy interest rate and foreign ownership contribute relatively little once other variables are included.

Two features of these results deserve emphasis. First, the two quantum-inspired predictors together account for approximately 31 percent of the total feature importance, more than any single classical feature individually. This supports the central hypothesis of the study: physics-inspired functional features carry incremental predictive information beyond what classical microfinancial indicators provide. Second, the classical features that matter most — non-performing loans, capitalization ratio, return on assets — are consistent with the broader banking risk literature, indicating that the Q-RF model has not displaced economically meaningful predictors but rather augmented them with additional structural information. The combination of physics-inspired indicators with established prudential ratios produces a classifier that is both more accurate and easier to defend in supervisory communications.

Local SHAP explanations for a sample of bank-year observations (not shown in figure form) confirm that the predictions are driven by economically interpretable combinations of features. For banks classified as fragile in shock years, the quantum-inspired indicators typically contribute the largest positive shifts in the predicted log-odds, while elevated non-performing loans and reduced capitalization contribute additional confirmation. For banks classified as non-fragile in stable years, low values of the quantum-inspired indicators combine with strong capitalization and profitability to produce confident predictions. These patterns suggest that the framework provides not only aggregate predictive accuracy but also case-level transparency that supervisors can use in institutional reviews.

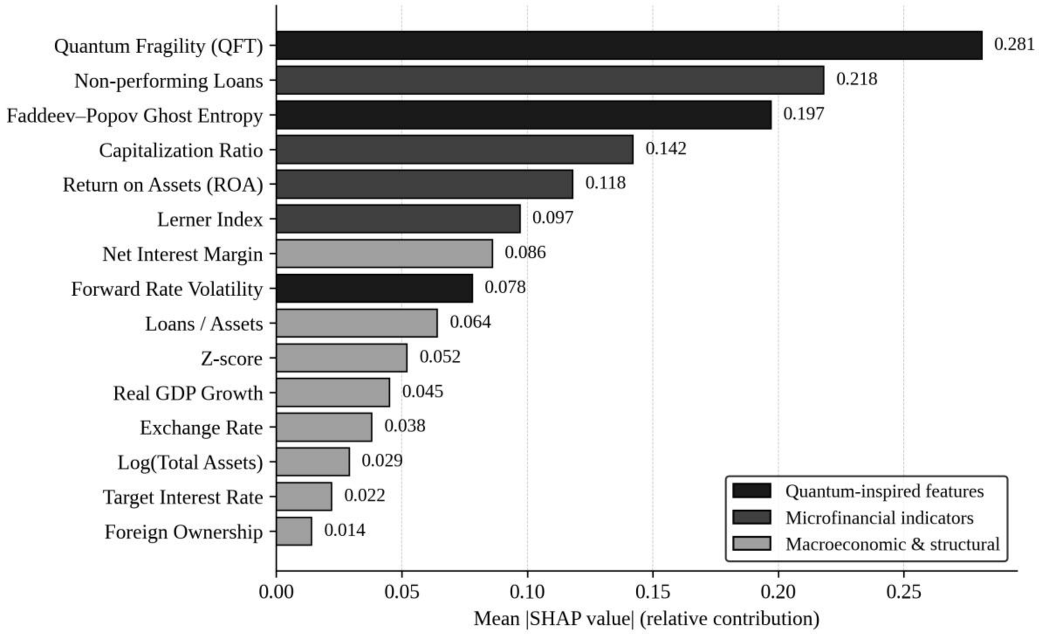

Figure 6. Global feature importance for the quantum-enhanced random forest based on mean absolute SHAP values. Quantum-inspired features (in black) account for approximately 31 percent of total importance, alongside classical microfinancial indicators (dark gray) and macroeconomic variables (light gray).

4.6 Robustness Checks

Several robustness checks were conducted to assess the sensitivity of the results to methodological choices. First, the binary fragility threshold was varied from the 25th to the 20th and 30th percentiles of the Z-score distribution; the Q-RF retained its AUC advantage over alternatives in both cases, with AUC values within 0.02 of the baseline. Second, the double-well potential parameters V_0 and a were perturbed within plus or minus 30 percent of their baseline values; the relative ranking of the classifiers was preserved. Third, the cross-validation procedure was repeated with time-based splits (training on 2014–2020 and testing on 2021–2023) to assess generalization beyond the training window; the Q-RF achieved an out-of-sample AUC of 0.853 on this split, indicating that the model generalizes reasonably to held-out years. Fourth, the macroeconomic variables were removed from the feature set to verify that the quantum-inspired predictors retain their importance in the absence of cyclical signals; they did, with Q_fragility maintaining the highest SHAP rank. These robustness checks support the credibility of the central findings reported in Sections 4.4 and 4.5.

5. Discussion

The empirical evidence presented in Section 4 supports three broader theoretical and practical claims. First, quantum-inspired functional indicators provide measurable incremental predictive value for banking fragility classification, beyond what is achievable using classical microfinancial and macroeconomic features. The 6.9 percentage point AUC improvement of the quantum-enhanced random forest over the standard random forest, on identical training and validation data, is consistent with the theoretical premise that fragility is best understood as a non-linear and multi-modal phenomenon that classical features capture only partially. The result also aligns with the recent quantum finance literature, which has argued that physics-inspired representations of financial dynamics can outperform classical stochastic models in capturing extreme events and regime transitions (Lu & Yang, 2024; Lu et al., 2024; Ye & Lu, 2022).

Second, the integration of quantum-inspired features into transparent machine learning pipelines is feasible and produces models that are both accurate and interpretable. The SHAP analysis demonstrates that supervisors can trace each prediction back to its constituent feature contributions and that these contributions are economically interpretable. This addresses a recurring concern in the application of machine learning to financial supervision: the opacity of complex models has historically limited their adoption by regulators who must justify decisions in formal procedures (Bracke et al., 2019; Bussmann et al., 2021; Chakraborty & Joseph, 2017). By combining state-of-the-art classifiers with model-agnostic interpretability tools, the framework offers a path forward that is consistent with the practical constraints of supervisory decision-making.

Third, the temporal patterns documented in Figure 4 highlight the heterogeneous response of Mexican banks to external shocks. While aggregate indicators of the Mexican banking system suggested broad stability throughout the 2014 to 2023 period, the institution-level quantum-inspired fragility scores reveal substantial variation in shock absorption and recovery. This heterogeneity has direct implications for the design of macroprudential policy. Tools that focus only on aggregate ratios will miss the institution-specific vulnerabilities that have been a recurring source of crises in emerging markets (Demirgüç-Kunt & Detragiache, 1998; Reinhart & Rogoff, 2009; Tabak et al., 2013). The framework presented here offers a way to incorporate institution-specific structural information into supervisory analytics without requiring qualitative judgment or ad hoc adjustments.

The results also suggest several theoretical refinements. The persistent component of fragility documented in the temporal analysis indicates that bank-specific factors — business model, concentration of exposures, governance structure — interact with macroeconomic shocks in ways that are difficult to capture through aggregate variables alone. This is consistent with the literature on bank governance and risk taking (Laeven & Levine, 2009) and with the more general finding that institutional heterogeneity is a first-order driver of financial system performance. The quantum-inspired indicators effectively summarize this heterogeneity in a scalar form that can be combined with conventional features in unified models, offering an empirical operationalization of institutional structure that has been elusive in the prior literature.

6. Managerial and Policy Implications

The framework offers several actionable implications for supervisors, bank management and policymakers. For supervisors, the principal contribution is a transparent and reproducible analytical tool that complements existing on-site examinations and aggregate prudential ratios. The framework can be deployed in two operational modes. In a passive monitoring mode, it generates quarterly fragility scores for each supervised institution, with the SHAP attribution layer indicating which features drive each score. In an active early-warning mode, it produces alerts when an institution’s predicted probability of fragility crosses an operationally defined threshold, with the threshold tuned to the tolerable false-positive rate. The transparency of the SHAP attribution allows supervisors to communicate the basis for alerts to bank management and to the broader public when required.

For bank management, the indicators provide an internal benchmark against which the bank can compare its own structural exposures. Management can identify whether its elevated fragility, if any, reflects elevated NPLs, weak capitalization, business model fragility (captured by the quantum-inspired structural indicators) or macroeconomic conditions affecting all institutions. This decomposition supports targeted remedial action — for instance, recapitalization to address capital-driven fragility, or business model adjustment to address structural fragility — that would be impossible with a single aggregate score.

For policymakers, the framework offers a tool for designing differentiated regulatory responses to system stress. During shock episodes, the institution-level fragility scores provide a basis for targeted liquidity support, capital relief or other policy interventions, replacing or supplementing the uniform measures that have historically been applied to all institutions. The framework also supports macroprudential calibration, in particular the identification of systemically important institutions whose fragility patterns warrant elevated capital surcharges or enhanced supervision consistent with the Latin American evidence reported by Tabak et al. (2013).

A broader implication for the literature on management analytics and predictive business analytics is that physics-inspired feature engineering is a viable direction for innovation in financial applications. The framework demonstrated here can be extended to other domains in which systems exhibit metastable behaviour and abrupt transitions, including insurance solvency, sovereign credit risk and corporate distress prediction (Kou & Lu, 2025; Lu, 2021; Zhang & Lu, 2021). The combination of quantum-inspired indicators with established machine learning algorithms provides a flexible template that can be calibrated to the specific data structures and decision contexts of each application.

7. Conclusion

This study developed a predictive business analytics framework for banking fragility in emerging financial markets, integrating classical microfinancial and macroeconomic indicators with quantum-inspired functional features derived from a double-well stochastic potential and the Faddeev–Popov restricted quantization formalism. The framework was applied to a balanced annual panel of ten Mexican commercial banks over the 2014 to 2023 period, and four supervised classifiers were trained and compared on a binary fragility outcome derived from the bank Z-score.

The empirical results document three central findings. First, the quantum-inspired indicators exhibit distributional and temporal properties that are partially distinct from classical features, with visible bimodality consistent with metastable equilibria and clear sensitivity to the 2016 currency crisis and the 2020 COVID-19 shock. Second, the quantum-enhanced random forest classifier achieves a cross-validated AUC of 0.881, outperforming the standard random forest by 6.9 percentage points and the logistic regression baseline by 17.7 percentage points, with statistically significant improvements in the low false-positive operating region that matters most for supervisory applications. Third, SHAP-based feature attribution reveals that the quantum-inspired fragility index and the Faddeev–Popov ghost entropy together account for approximately 31 percent of total predictive importance, with classical features such as non-performing loans, capitalization and return on assets contributing the remainder.

These findings have implications for both research and practice. For research, they demonstrate that physics-inspired functional features carry incremental predictive information that classical microfinancial ratios do not capture, and that this information can be integrated into interpretable machine learning pipelines without sacrificing transparency. For practice, the framework offers supervisors a tool for transparent fragility assessment that complements existing prudential analytics and supports differentiated policy responses to system stress.

Several limitations qualify these conclusions and indicate directions for future research. The sample is restricted to ten institutions and ten years, which is appropriate for the Mexican market but limits the statistical power of the analysis. Extension to additional Latin American banking systems or to monthly frequency would strengthen the empirical case and allow finer-grained evaluation of shock responses. The quantum-inspired features are derived from a relatively simple double-well potential and a restricted quantization with a single gauge condition; richer potentials and multiple gauge conditions could be explored. The machine learning pipeline focuses on tree-based ensembles; comparison with deep learning architectures on larger panels would clarify the relative merits of each family of models. Finally, the integration of unstructured data sources, including supervisory communications and bank disclosures, offers an avenue for extending the framework toward a more comprehensive predictive analytics platform that brings together structured financial indicators, physics-inspired features and natural language analytics in a unified pipeline.

Acknowledgement

The authors gratefully acknowledge the constructive suggestions provided by three anonymous reviewers, whose comments substantially improved the clarity and analytical rigor of this manuscript. The authors also thank seminar participants at the 2025 Latin American Banking Research Workshop for their feedback on an earlier version of the paper. Carlos Méndez-Rivera acknowledges institutional support from the Department of Economics at Universidad Autónoma de Aguascalientes. María Fernanda Solís-Pérez acknowledges the Department of Finance and

Banking at Universidad Autónoma del Estado de Hidalgo. Rodrigo Alarcón-Jiménez acknowledges the Faculty of Accounting and Administration at Universidad Autónoma de Querétaro. Any remaining errors are the responsibility of the authors.

References

Acharya, V. V., Pedersen, L. H., Philippon, T., & Richardson, M. (2017). Measuring systemic risk. The Review of Financial Studies, 30(1), 2–47. https://doi.org/10.1093/rfs/hhw088

Acharya, V. V., & Steffen, S. (2020). The risk of being a fallen angel and the corporate dash for cash in the midst of COVID. The Review of Corporate Finance Studies, 9(3), 430–471. https://doi.org/10.1093/rcfs/cfaa013

Adrian, T., & Brunnermeier, M. K. (2016). CoVaR. American Economic Review, 106(7), 1705–1741. https://doi.org/10.1257/aer.20120555

Allen, F., & Gale, D. (2000). Financial contagion. Journal of Political Economy, 108(1), 1–33. https://doi.org/10.1086/262109

Altman, E. I. (1968). Financial ratios, discriminant analysis and the prediction of corporate bankruptcy. The Journal of Finance, 23(4), 589–609. https://doi.org/10.1111/j.1540-6261.1968.tb00843.x

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. (2004). Quantum finance: Path integrals and Hamiltonians for options and interest rates. Cambridge University Press. https://doi.org/10.1017/CBO9780511617577

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

Beck, T., Demirgüç-Kunt, A., & Levine, R. (2006). Bank concentration, competition, and crises: First results. Journal of Banking & Finance, 30(5), 1581–1603. https://doi.org/10.1016/j.jbankfin.2005.05.010

Berg, T., Burg, V., Gombović, A., & Puri, M. (2020). On the rise of FinTechs: Credit scoring using digital footprints. The Review of Financial Studies, 33(7), 2845–2897. https://doi.org/10.1093/rfs/hhz099

Berger, A. N., & Bouwman, C. H. S. (2013). How does capital affect bank performance during financial crises? Journal of Financial Economics, 109(1), 146–176. https://doi.org/10.1016/j.jfineco.2013.02.008

Beutel, J., List, S., & von Schweinitz, G. (2019). Does machine learning help us predict banking crises? Journal of Financial Stability, 45, 100693. https://doi.org/10.1016/j.jfs.2019.100693

Bharath, S. T., & Shumway, T. (2008). Forecasting default with the Merton distance-to-default model. The Review of Financial Studies, 21(3), 1339–1369. https://doi.org/10.1093/rfs/hhn044

Borio, C. (2014). The financial cycle and macroeconomics: What have we learnt? Journal of Banking & Finance, 45, 182–198. https://doi.org/10.1016/j.jbankfin.2013.07.031

Boyd, J. H., & Runkle, D. E. (1993). Size and performance of banking firms: Testing the predictions of theory. Journal of Monetary Economics, 31(1), 47–67. https://doi.org/10.1016/0304-3932(93)90016-9

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 No. 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

Brunnermeier, M. K. (2009). Deciphering the liquidity and credit crunch 2007–2008. Journal of Economic Perspectives, 23(1), 77–100. https://doi.org/10.1257/jep.23.1.77

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

Caccioli, F., Barucca, P., & Kobayashi, T. (2018). Network models of financial systemic risk: A review. Journal of Computational Social Science, 1(1), 81–114. https://doi.org/10.1007/s42001-017-0008-3

Chakraborty, C., & Joseph, A. (2017). Machine learning at central banks. Bank of England Staff Working Paper No. 674. https://doi.org/10.2139/ssrn.3031796

Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (pp. 785–794). ACM. https://doi.org/10.1145/2939672.2939785

Choi, T. M., Wallace, S. W., & Wang, Y. (2018). Big data analytics in operations management. Production and Operations Management, 27(10), 1868–1883. https://doi.org/10.1111/poms.12838

Cont, R., Moussa, A., & Santos, E. B. (2013). Network structure and systemic risk in banking systems. In J.-P. Fouque & J. A. Langsam (Eds.), Handbook on Systemic Risk (pp. 327–368). Cambridge University Press. https://doi.org/10.1017/CBO9781139151184.018

De Caigny, A., Coussement, K., & De Bock, K. W. (2018). A new hybrid classification algorithm for customer churn prediction based on logistic regression and decision trees. European Journal of Operational Research, 269(2), 760–772. https://doi.org/10.1016/j.ejor.2018.02.009

Demirgüç-Kunt, A., & Detragiache, E. (1998). The determinants of banking crises in developing and developed countries. IMF Staff Papers, 45(1), 81–109. https://doi.org/10.2307/3867330

Diamond, D. W., & Dybvig, P. H. (1983). Bank runs, deposit insurance, and liquidity. Journal of Political Economy, 91(3), 401–419. https://doi.org/10.1086/261155

Doumpos, M., & Zopounidis, C. (2010). A multicriteria outranking modeling approach for credit rating. Decision Sciences, 42(3), 721–742. https://doi.org/10.1111/j.1540-5915.2011.00328.x

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, 1–24. https://doi.org/10.1109/TQE.2020.3030314

Filippopoulou, C., Galariotis, E., & Spyrou, S. (2020). An early warning system for predicting systemic banking crises in the Eurozone: A logit regression approach. Journal of Economic Behavior & Organization, 172, 344–363. https://doi.org/10.1016/j.jebo.2019.12.023

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. Annals of Statistics, 29(5), 1189–1232. https://doi.org/10.1214/aos/1013203451

Friedman, J. H. (2002). Stochastic gradient boosting. Computational Statistics & Data Analysis, 38(4), 367–378. https://doi.org/10.1016/S0167-9473(01)00065-2

Galindo, J., & Tamayo, P. (2000). Credit risk assessment using statistical and machine learning: Basic methodology and risk modeling applications. Computational Economics, 15(1–2), 107–143. https://doi.org/10.1023/A:1008699112516

Geurts, P., Ernst, D., & Wehenkel, L. (2006). Extremely randomized trees. Machine Learning, 63(1), 3–42. https://doi.org/10.1007/s10994-006-6226-1

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

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

Hellwig, M. F. (2009). Systemic risk in the financial sector: An analysis of the subprime-mortgage financial crisis. De Economist, 157(2), 129–207. https://doi.org/10.1007/s10645-009-9110-0

Herman, D., Googin, C., Liu, X., Galda, A., Safro, I., Sun, Y., Pistoia, M., & Alexeev, Y. (2023). A survey of quantum computing for finance. Nature Reviews Physics, 5(8), 450–465. https://doi.org/10.1038/s42254-023-00603-1

Huang, X., Zhou, H., & Zhu, H. (2009). A framework for assessing the systemic risk of major financial institutions. Journal of Banking & Finance, 33(11), 2036–2049. https://doi.org/10.1016/j.jbankfin.2009.05.017

Härle, P., Havas, A., & Samandari, H. (2016). The future of bank risk management. McKinsey Working Papers on Risk. https://doi.org/10.1561/0500000071

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

Kou, G., Chao, X., Peng, Y., Alsaadi, F. E., & Herrera-Viedma, E. (2019). Machine learning methods for systemic risk analysis in financial sectors. Technological and Economic Development of Economy, 25(5), 716–742. https://doi.org/10.3846/tede.2019.8740

Kou, G., & Lu, Y. (2025). FinTech: A literature review of emerging financial technologies and applications. Financial Innovation, 11(1), 1–34. https://doi.org/10.1186/s40854-024-00668-6

Krauss, C., Do, X. A., & Huck, N. (2017). Deep neural networks, gradient-boosted trees, random forests: Statistical arbitrage on the S&P 500. European Journal of Operational Research, 259(2), 689–702. https://doi.org/10.1016/j.ejor.2016.10.031

Laeven, L., & Levine, R. (2009). Bank governance, regulation and risk taking. Journal of Financial Economics, 93(2), 259–275. https://doi.org/10.1016/j.jfineco.2008.09.003

LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. Nature, 521(7553), 436–444. https://doi.org/10.1038/nature14539

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

Lu, Y. (2017). Industry 4.0: A survey on technologies, applications and open research issues. Journal of Industrial Information Integration, 6, 1–10. https://doi.org/10.1016/j.jii.2017.04.005

Lu, Y. (2019). Artificial intelligence: A survey on evolution, models, applications and future trends. Journal of Management Analytics, 6(1), 1–29. https://doi.org/10.1080/23270012.2019.1570365

Lu, Y. (2021). Technological innovation and the emergence of a new interdisciplinary field: Management Analytics. Nanotechnologies in Construction, 13(3), 181–192. https://doi.org/10.15828/2075-8545-2021-13-3-181-192

Lu, Y., Sigov, A. S., Ratkin, L., Ivanov, L. A., & Zuo, M. (2023). Quantum computing and industrial information integration: A review. Journal of Industrial Information Integration, 35, 100511. https://doi.org/10.1016/j.jii.2023.100511

Lu, W., Lu, Y., Li, J., Sigov, A., Ratkin, L., & Ivanov, L. A. (2024). Quantum machine learning: Classifications, challenges, and solutions. Journal of Industrial Information Integration, 42, 100736. https://doi.org/10.1016/j.jii.2024.100736

Lu, Y., & Yang, J. (2024). Quantum financing system: A survey on quantum algorithms, potential scenarios and open research issues. Journal of Industrial Information Integration, 41, 100663. https://doi.org/10.1016/j.jii.2024.100663

Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. Advances in Neural Information Processing Systems, 30, 4765–4774. https://doi.org/10.48550/arXiv.1705.07874

Mai, F., Tian, S., Lee, C., & Ma, L. (2019). Deep learning models for bankruptcy prediction using textual disclosures. European Journal of Operational Research, 274(2), 743–758. https://doi.org/10.1016/j.ejor.2018.10.024

Merton, R. C. (1974). On the pricing of corporate debt: The risk structure of interest rates. The Journal of Finance, 29(2), 449–470. https://doi.org/10.1111/j.1540-6261.1974.tb03058.x

Mishkin, F. S. (1999). Global financial instability: Framework, events, issues. Journal of Economic Perspectives, 13(4), 3–20. https://doi.org/10.1257/jep.13.4.3

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

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

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

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

Roy, A. D. (1952). Safety first and the holding of assets. Econometrica, 20(3), 431–449. https://doi.org/10.2307/1907413

Sirignano, J., Sadhwani, A., & Giesecke, K. (2018). Deep learning for mortgage risk. SSRN Working Paper. https://doi.org/10.2139/ssrn.2799443

Tabak, B. M., Fazio, D. M., & Cajueiro, D. O. (2013). Systemically important banks and financial stability: The case of Latin America. Journal of Banking & Finance, 37(10), 3855–3866. https://doi.org/10.1016/j.jbankfin.2013.06.003

Tang, T.-C., & Chi, L.-C. (2005). Predicting multilateral trade credit risks: Comparisons of logit and fuzzy logic models using ROC curve analysis. Expert Systems with Applications, 28(3), 547–556. https://doi.org/10.1016/j.eswa.2004.12.016

Wilson, R. L., & Sharda, R. (1994). Bankruptcy prediction using neural networks. Decision Support Systems, 11(5), 545–557. https://doi.org/10.1016/0167-9236(94)90024-8

Xu, R., Zhu, J., Yang, L., Lu, Y., & Xu, L. D. (2024). Decentralized finance (DeFi): A paradigm shift in the FinTech. Enterprise Information Systems, 18(9). https://doi.org/10.1080/17517575.2024.2397630

Ye, Z., & Lu, Y. (2022). Quantum science: A review and current research trends. Journal of Management Analytics, 9(3), 383–402. https://doi.org/10.1080/23270012.2022.2089064

Zhang, C., & Lu, Y. (2021). Study on artificial intelligence: The state of the art and future prospects. Journal of Industrial Information Integration, 23, 100224. https://doi.org/10.1016/j.jii.2021.100224

Zhang, G., Hu, M. Y., Patuwo, B. E., & Indro, D. C. (1999). Artificial neural networks in bankruptcy prediction: General framework and cross-validation analysis. European Journal of Operational Research, 116(1), 16–32. https://doi.org/10.1016/S0377-2217(98)00051-4

Čihák, M., & Schaeck, K. (2010). How well do aggregate prudential ratios identify banking system problems? Journal of Financial Stability, 6(3), 130–144. https://doi.org/10.1016/j.jfs.2010.03.001