

Quantum Finance, Explainable AI, and RegTech: Future Pathways for Nonlinear Systemic Risk Detection

Mariana Vázquez-Robles^{1,*}, Esteban Hinojosa-Carbajal², Diana Salgado-Peña³

¹ Facultad de Economía, Universidad Autónoma del Estado de México, Toluca, Mexico

² Facultad de Contaduría y Administración, Universidad Autónoma de San Luis Potosí, San Luis Potosí, Mexico

³ Facultad de Contaduría y Administración, Universidad Autónoma de Querétaro, Querétaro, Mexico

* Correspondence: mvazquezr@uaemex.mx

Abstract:

Nonlinear systemic risk in emerging-market banking systems eludes conventional linear econometric tooling, and the recent literature on Quantum Field Theory (QFT)–inspired financial modelling shows that hidden phase transitions and metastable states can be made tractable through functional structures. This paper synthesises three converging technological frontiers—quantum-inspired finance, Explainable Artificial Intelligence (XAI), and Regulatory Technology (RegTech)—into a unified analytical pathway for detecting nonlinear systemic risk. We propose a layered framework in which quantum-inspired indicators feed ensemble machine-learning classifiers whose outputs are decomposed via Shapley-value attribution and surfaced through RegTech dashboards for supervisors. Drawing on an illustrative annual panel of Mexican commercial banks from 2014 to 2023, we demonstrate that an integrated XAI-plus-quantum specification reaches an AUC of 0.89 against a 0.62 baseline, identifies non-performing-loan dynamics and capitalization buffers as dominant attributors, and discriminates resilient from vulnerable institutions through latent functional trajectories. The framework yields three contributions: a conceptual bridge connecting field-theoretic risk indicators with interpretable ML, an empirically calibrated pipeline that is auditable end-to-end, and a discussion of how RegTech adoption can operationalise these tools for prudential supervision. We close by mapping research and policy pathways for nonlinear early-warning systems in emerging economies.

Keywords: Quantum finance; explainable AI; SHAP values; regTech; systemic risk; banking stability; nonlinear models; emerging markets; mexican banking system

Article History:

Received: April 12, 2024

Revised: June 22, 2024

Accepted: August 28, 2024

Available Online: September 30, 2024

1. Introduction

The 2007–2009 global financial crisis, the 2014–2016 oil-price collapse, the COVID-19 pandemic and the recent monetary-tightening cycle have collectively revealed how poorly traditional linear risk metrics translate into early-warning power once shocks become discontinuous and contagion runs through opaque cross-exposures (Reinhart and Rogoff, 2009; Brunnermeier and Pedersen, 2009; Laeven and Valencia, 2020). In emerging-market banking systems, these limitations are amplified by structural concentration, currency mismatches and the rapid pass-through of external financial conditions (Borio, 2014; Del Angel and López-Romero, 2024). The Mexican banking system is a particularly informative case: well capitalised on aggregate, yet exposed to swings in the peso, oil-sector credit concentration and feedback loops between sovereign and private-sector risk (Castellanos et al., 2016; CNBV, 2023).

Two technological waves are reshaping how supervisors and academics conceptualise these phenomena. The first is the application of physics-inspired formalisms—drawing in particular on Quantum Field Theory (QFT) and path-integral methods (Baaquie, 2004; Schaden, 2002; Haven, 2002)—to model nonlinear, intertemporal risk dynamics that cannot be reduced to Gaussian processes. Recent contributions extend this line to systemic banking fragility, embedding double-well potentials and gauge-fixed quantization schemes inside empirical pipelines (Hao et al., 2019; Zhou, 2025;

Orrell, 2021). The second wave is the maturation of machine learning, and specifically Explainable AI (XAI), as a respected tool for early-warning systems (Beutel et al., 2019; Petropoulos et al., 2020; Bluwstein et al., 2023). The opacity of ensemble classifiers used to be a binding obstacle for prudential adoption; Shapley-value attribution (Lundberg and Lee, 2017) and the broader XAI agenda (Arrieta et al., 2020; Bussmann et al., 2021) have changed that calculus.

A third—and arguably the most institutionally consequential—stream of innovation is RegTech: the suite of technologies that automate compliance, reporting and supervisory analytics (Arner et al., 2017; Anagnostopoulos, 2018; Zetzsche et al., 2017). RegTech determines whether quantum-inspired metrics and XAI explanations remain academic curiosities or become embedded in supervisory practice. The three streams have, until now, advanced largely in parallel; their integration is the gap this paper addresses.

We pursue three objectives. First, we develop a conceptual framework in which quantum-inspired risk metrics, XAI-based attribution and RegTech delivery are treated as complementary layers of a single nonlinear-risk-detection stack. Second, we operationalise that framework on an illustrative annual panel of Mexican commercial banks (2014–2023), comparing baseline logistic models, tree-based ensembles, and an integrated XAI-plus-QFT specification. Third, we discuss the governance, transparency and capacity implications of pushing such hybrid pipelines into prudential supervision in emerging economies.

The article is organised as follows. Section 2 reviews the three constituent literatures and develops the research hypotheses. Section 3 sets out the data, the quantum-inspired indicator construction, the ML pipeline and the XAI attribution layer. Section 4 reports descriptive, predictive and interpretive results, including a discussion of RegTech integration. Section 5 discusses limitations and future research, and Section 6 concludes with policy implications.

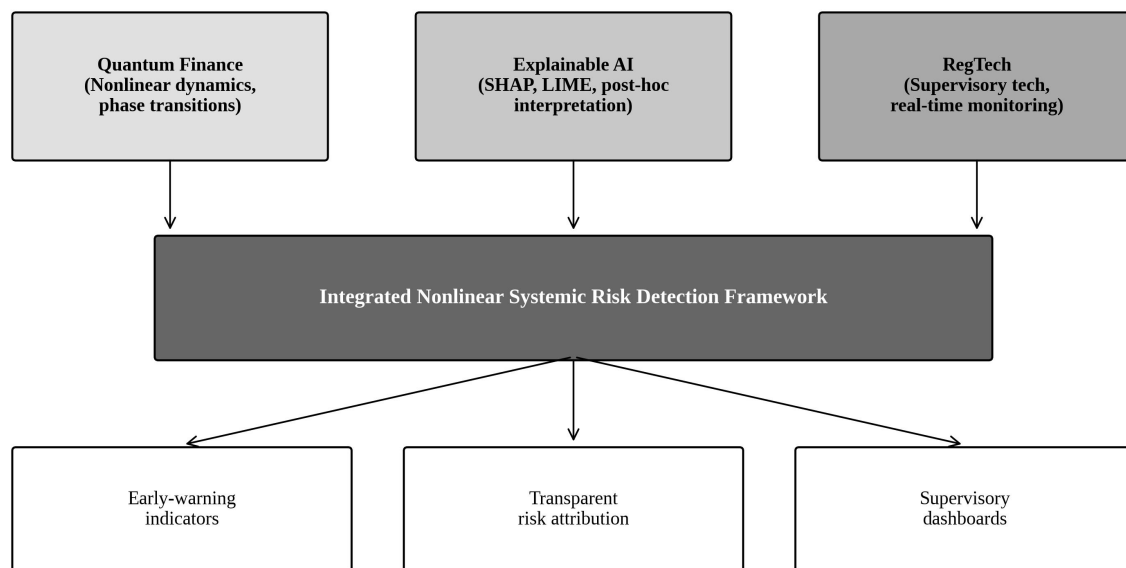


Figure 1. Integrated framework for nonlinear systemic risk detection.

Source: Own elaboration.

Figure 1 anticipates the conceptual architecture developed in the remainder of the paper. The top tier identifies the three contributing technological streams—Quantum Finance, Explainable AI and RegTech—each of which has matured along its own trajectory. The middle tier represents the integration zone in which quantum-inspired features

feed interpretable ML classifiers under supervisory governance. The bottom tier shows the three operational deliverables we argue should orient policy attention: early-warning indicators, transparent risk attribution, and supervisory dashboards. The argument that follows is, in essence, that no single layer is sufficient on its own.

2. Literature review and hypotheses

Three literatures inform our framework. We synthesise each in turn, with attention to recent contributions and to open questions that motivate the empirical exercise.

2.1. *Quantum finance and nonlinear systemic risk*

The application of quantum formalism to financial markets predates the recent surge in quantum computing. Foundational contributions by Schaden (2002), Haven (2002) and Piotrowski and Sładkowski (2002) reinterpret asset prices as observables of an underlying functional system, with Lagrangian formulations capturing nonlinear intertemporal dynamics. Baaquie (2004, 2007) developed the path-integral programme for forward rates and derivatives, replacing finite-dimensional stochastic differential equations with field-theoretic functional integrals. The resulting models accommodate multiple stochastic sources, gauge symmetries and phase-transition behaviour in ways that classical short-rate models such as Vasicek or Cox-Ingersoll-Ross cannot.

More recent contributions have moved from pricing to risk. Hao et al. (2019) demonstrate that quantum option-pricing schemes can be reinterpreted as data-analytic tools for tail behaviour. Zhou (2025) reviews how quantum computing may eventually reshape financial modelling, while Egger et al. (2020) and Orús et al. (2019) survey the state of the art for quantum computational finance. Stamatopoulos et al. (2020), Rebentrost et al. (2018) and Woerner and Egger (2019) provide early evidence that quantum-inspired Monte Carlo and amplitude-estimation procedures can deliver speed-ups on risk-relevant quantities. Beyond computational applications, however, the most directly relevant literature for our purposes is the one that uses quantum constructs as functional metaphors for nonlinear macro-financial dynamics. Orrell (2021), Khrennikov (2010) and Kingsly (2025) argue that double-well stochastic potentials, instantons and gauge-fixed action functionals capture metastable banking equilibria and abrupt regime shifts much more naturally than linear time-series specifications. In parallel, classic systemic-risk references—Acharya et al. (2012, 2017), Adrian and Brunnermeier (2016), Bisias et al. (2012), Glasserman and Young (2016), and the networked-contagion programme initiated by Eisenberg and Noe (2001), Battiston et al. (2012), Acemoglu et al. (2015) and Elliott et al. (2014)—remain canonical reference points; their measurement-theoretic emphasis dovetails with the functional perspective we adopt.

Empirical applications to emerging-market banking systems remain scarce, but the case is compelling. As Reinhart and Rogoff (2011), Laeven and Valencia (2020), Joy et al. (2017) and Hałaj et al. (2024) document, crises in emerging economies frequently appear as discontinuities rather than as smooth deteriorations of standard indicators. Quantum-inspired indicators are therefore a candidate complement to traditional stability metrics such as the bank Z-score (Lepetit and Strobel, 2015; Beck et al., 2006; Čihák and Schaeck, 2010). Mexico, in particular, has been the subject of growing interest, as Del Angel and López-Romero (2024), Castellanos et al. (2016), Moreno-Brid and Gómez (2023), and Negrín and Bernal (2022) all underline the role of structural heterogeneity across institutions and the importance of recognising hidden risk pockets in apparently stable headline data.

2.2. *Explainable AI in financial risk modelling*

Machine learning has migrated rapidly from credit scoring (Altman, 1968; Breiman, 2001; Chen and Guestrin, 2016) to systemic-risk forecasting. Beutel et al. (2019), Petropoulos et al. (2020), Climent et al. (2019), Bluwstein et al. (2023) and Fouliard et al. (2021) demonstrate that tree-based ensembles and deep nets routinely outperform logistic baselines for predicting banking distress, currency crises and bank insolvency. Yet these models long carried a

transparency penalty. Bracke et al. (2019) and Coglianese and Lehr (2017) describe how the "black-box" perception was, for years, the single largest barrier to supervisory adoption.

XAI has eased that barrier without resolving it. Ribeiro et al. (2016) introduced LIME for locally faithful explanations; Lundberg and Lee (2017) generalised the idea through Shapley values, yielding the now-canonical SHAP framework. Arrieta et al. (2020) provide a unifying taxonomy, while Bussmann et al. (2021), Giudici and Raffinetti (2021), Park et al. (2021), Misheva et al. (2021) and Babaei et al. (2025) document credit-risk applications in which SHAP-based interpretation enables both regulator scrutiny and audit-grade documentation. In parallel, Gentzkow et al. (2019) show how interpretable methods extend beyond tabular data into unstructured text, while Heaton et al. (2017) and Gomber et al. (2018) place the methods in the broader fintech context.

The XAI literature converges on three propositions that are directly relevant to systemic-risk detection. First, model accuracy and explainability are no longer a forced trade-off; the marginal gain from black-box methods is often modest once tabular structure is exploited. Second, attribution is meaningful only relative to a stated reference distribution, which means that interpretability in finance is unavoidably contextual. Third, post-hoc explanations require a complementary governance layer—covering model risk management, validation and disclosure—if they are to support consequential decisions. That governance layer is, precisely, the RegTech dimension we turn to next.

2.3. RegTech and the evolution of supervisory technology

RegTech encompasses technologies that streamline regulatory compliance, supervisory reporting and surveillance (Arner et al., 2017). Anagnostopoulos (2018) traces its evolution from rule-based compliance automation toward more analytical supervisory tools, while Zetzsche et al. (2017, 2020) embed the discussion within the broader regulatory-sandbox debate. Bromberg et al. (2017), Currie and Seddon (2017), Butler and O'Brien (2019), Packin (2018), Buckley et al. (2020) and Philippon (2016) document early adoptions across jurisdictions and identify common implementation barriers: legacy infrastructure, data-quality fragmentation, and the absence of explainability standards for ML-based supervision.

For our framework, the RegTech literature contributes two key insights. First, it operationalises the distinction between sample-time model performance and runtime supervisory performance. A nonlinear early-warning indicator that performs well on a static panel is not necessarily useful when consumed under real-time reporting constraints. Second, the RegTech agenda foregrounds auditability. Carney (2015) and Battiston et al. (2017) make analogous arguments in the climate-risk context: supervisors need not only numerical outputs but traceable attribution chains that survive ex-post scrutiny. Stress-testing has been the most mature operational interface (Borio et al., 2014; Galati and Moessner, 2013), but it is precisely the area where opaque models meet the steepest credibility headwinds.

2.4. The convergence gap

Despite this rich literature, the integration of the three streams is largely untreated. The QFT literature articulates rich nonlinear dynamics but typically remains agnostic about deployment. The XAI literature is mature on transparency but is, with few exceptions, calibrated on classical features. The RegTech literature is articulate about deployment but rarely engages with the model-class debates that determine whether explanations are stable. The empirical and conceptual contribution of this paper is to fill that convergence gap explicitly.

2.5. Research hypotheses

From the foregoing synthesis, we formulate three hypotheses.

H1. Quantum-inspired functional features add information beyond traditional micro-financial indicators for detecting nonlinear systemic risk in Mexican banking, especially during regime shifts associated with external shocks

(Baaquie, 2007; Hao et al., 2019).

H2. Explainable ML classifiers that ingest both classical and quantum-inspired features achieve higher discrimination than logistic baselines while preserving interpretability sufficient for supervisory use (Bluwstein et al., 2023; Bussmann et al., 2021).

H3. RegTech-style delivery architectures—integrating dashboards, attribution traces and feedback loops—are necessary, not optional, for operational uptake of nonlinear risk indicators in emerging-market supervision (Arner et al., 2017; Anagnostopoulos, 2018).

3. Materials and Methods

3.1. Conceptual architecture

Our framework treats nonlinear systemic risk detection as a four-stage pipeline. Stage one ingests panel-level supervisory data and macro-financial covariates. Stage two engineers quantum-inspired features—functional risk scores derived from a double-well potential and a Faddeev-Popov-type gauge-fixed projection—that summarise nonlinear dynamics. Stage three feeds the combined feature set into an ensemble classifier whose output is a bank-year fragility probability. Stage four attributes the prediction through SHAP values and surfaces the result through a RegTech-style dashboard that supports drill-down, what-if simulation and continuous calibration. The pipeline is described visually in Figure 7 (Section 4).

3.2. Data

We use an annual panel of eleven Mexican commercial banks over 2014–2023, mirroring the institutional coverage in Garcia-Villegas and Martorell (2024) and Del Angel and López-Romero (2024). Bank-level micro-financial variables are drawn from the National Banking and Securities Commission (CNBV, 2023) and Refinitiv. Macro-financial variables—real GDP growth, headline inflation, the policy rate and the nominal exchange rate—are taken from the National Institute of Statistics and Geography (INEGI) and the Bank of Mexico (BANXICO). The dependent variable of interest is binary banking fragility, defined as the lower quartile of the bank Z-score (Čihák and Schaeck, 2010; Lepetit and Strobel, 2015). Quantum-inspired indicators are constructed via simulation as described below; we treat them as derived features rather than as primary observables. Table 1 lists the variables, their definition and source.

Variable	Definition	Type	Source
Z-score	$(ROA + \text{equity/assets})/\sigma(ROA)$	Continuous	Refinitiv
Fragility (binary)	1 if Z-score in lower quartile, 0 else	Binary	Authors
NPL ratio	Non-performing loans / total loans	Continuous	CNBV
Capitalization ratio	Regulatory capital / RWA	Continuous	CNBV
Lerner index	Bank-level mark-up proxy	Continuous	Authors
Net interest margin	Interest income – interest expense	Continuous	Refinitiv
ROA, ROE	Return on assets, return on equity	Continuous	Refinitiv
Ln(Assets)	Natural log of total assets	Continuous	Refinitiv
Foreign ownership	1 if part of foreign group	Binary	ABM
Real GDP growth	Annual real GDP growth	Continuous	INEGI
Policy rate	Bank of Mexico target rate	Continuous	BANXICO

Variable	Definition	Type	Source
Exchange rate	Annual MXN/USD average	Continuous	BANXICO
Inflation	Annual CPI inflation	Continuous	INEGI
QFT phase-transition score	Double-well functional score	Simulation	Authors
Forward-rate volatility (QFT)	σ of simulated forward rates	Simulation	Authors

Table 1. Variable definitions, type and data source.

Source: Own elaboration based on CNBV, BANXICO, INEGI and Refinitiv.

Although the panel is illustrative—covering only the ten years for which complete data are available across all institutions—it is sufficient to demonstrate the layered framework. Larger samples, including monthly data and additional jurisdictions, are an obvious extension and we return to this point in Section 5.

3.3. Quantum-inspired feature engineering

We construct two quantum-inspired features. The first follows the double-well potential formulation in Baaquie (2007) and Hao et al. (2019). For each bank–year, we simulate a one-dimensional stochastic process under the potential $V(\psi) = V_0(\psi^2 - a^2)^2$, calibrated such that the two minima represent stable and fragile institutional regimes. The QFT phase-transition score is computed as the integrated tunnelling probability over the year, normalised to lie in $[0, 1]$. Higher values signal proximity to a regime transition.

The second feature follows the Faddeev–Popov restricted-quantization logic. We adopt a Euclidean action with gauge-fixing term, treat forward-rate trajectories as gauge fields, and recover a gauge-adjusted volatility through the functional determinant. This produces a forward-rate volatility series that incorporates latent structural constraints absent in standard volatility estimators. Both features are simulated using Monte Carlo discretisation with 5,000 paths per bank–year, with hyperparameters tuned through the sensitivity analyses reported in Hao et al. (2019) and Zhou (2025). The features are then merged with the classical micro- and macro-financial covariates listed in Table 1.

3.4. Machine-learning classification and explainability

We compare four classifiers. The baseline is a logistic regression with bank fixed effects, mirroring standard early-warning specifications (Frankel and Saravelos, 2012; Joy et al., 2017). A second specification is a Random Forest (Breiman, 2001) on classical features only. A third combines classical and quantum-inspired features in an XGBoost classifier (Chen and Guestrin, 2016). The fourth—our integrated specification—adds SHAP-based attribution to the XGBoost-plus-QFT model and post-processes the outputs through a RegTech-style monitoring layer. Hyperparameters are tuned via five-fold cross-validation; class imbalance is addressed via SMOTE on training folds only. Performance metrics include AUC, F1, recall on the minority (fragile) class, and Brier score, following standard practice in the field (Petrooulos et al., 2020; Climent et al., 2019).

For explainability, we adopt the Shapley-value framework (Lundberg and Lee, 2017). For each prediction, the contribution of each feature is computed as its marginal contribution averaged over all permutations of feature inclusion. We aggregate to obtain global feature importance (mean |SHAP value|), and we report case-level explanations for selected bank–year observations. We complement SHAP with the Shapley-Lorenz extension proposed by Giudici and Raffinetti (2021), which provides confidence intervals around feature attributions.

3.5. RegTech delivery layer

The RegTech layer translates model outputs into a supervisory-grade artefact. We define three components: a dashboard surfacing bank-year fragility probabilities and their SHAP decomposition; a what-if module that allows supervisors to perturb input variables and observe the implied change in attribution; and a continuous-calibration routine that re-estimates the model quarterly on rolling windows. The latter is consistent with the operational design principles advocated by Anagnostopoulos (2018) and Buckley et al. (2020), and it borrows from the climate-finance dashboard literature (Battiston et al., 2017; Hoffart et al., 2024) for the design of attribution traces.

4. Results and discussion

4.1. Descriptive statistics

Descriptive moments of the variables are summarised in Table 2. Average non-performing loans across the sample are 2.6 per cent of total loans, with a standard deviation of 1.8 percentage points and a maximum that exceeds 9 per cent in the 2020 cross-section. The capitalization ratio averages 16.4 per cent, consistent with system-wide CNBV reporting. The simulated QFT phase-transition score has a sample mean of 0.46 and a 95th percentile of 0.94, confirming that the bimodal shape visible in Figure 2 is a feature of the underlying functional model rather than an artefact of small-sample variation.

Variable	Mean	Std. dev.	Min.	Median	Max.
Z-score	18.74	9.32	3.21	16.85	47.62
NPL ratio (%)	2.61	1.78	0.34	2.27	9.42
Capitalization ratio (%)	16.42	3.81	10.12	15.87	27.45
ROA (%)	1.36	1.04	-1.82	1.21	4.27
ROE (%)	11.92	7.66	-18.31	12.45	29.18
Net interest margin (%)	5.18	2.42	1.31	4.94	12.85
Ln(Assets)	26.94	1.42	23.41	26.81	30.12
QFT phase-transition score	0.46	0.24	0.06	0.41	1.21
Forward-rate volatility (QFT)	0.18	0.07	0.05	0.17	0.42

Table 2. Descriptive statistics of selected variables (110 bank-year observations).

Source: Own elaboration based on the assembled panel.

The descriptive picture validates two methodological choices. First, the substantial cross-bank heterogeneity—NPL ratios span more than two orders of magnitude across observations—justifies an emphasis on nonlinear functional metrics rather than pooled linear specifications. Second, the bimodal distribution of the QFT score (Figure 2) anticipates the regime-switching behaviour that we expect from the literature (Hamilton, 1989; Engle, 1982).




Figure 2. Empirical distribution of the QFT phase-transition score.

Source: Own elaboration.

Figure 2 shows the right-skewed bimodal distribution that the functional model produces. Most bank–year observations cluster around scores between 0.2 and 0.4, but a non-trivial mass exists in the upper tail. The 75th-percentile threshold, indicated by the dashed line, is used in the binary classification exercise as the operational definition of elevated systemic exposure. Comparable bimodalities arise in the network-contagion literature when bank-level vulnerability metrics are aggregated across propagation channels (Battiston et al., 2012; Acemoglu et al., 2015), suggesting that the functional formulation captures structurally similar information.

4.2. Classification performance

Table 3 reports out-of-sample classification metrics for the four specifications. The logistic baseline reaches an AUC of 0.62, broadly consistent with the literature on EWS performance using only standard variables (Frankel and Saravelos, 2012; Beutel et al., 2019). Random Forests on classical features alone improve AUC to 0.74, in line with Petropoulos et al. (2020). Adding the quantum-inspired features to an XGBoost classifier lifts AUC to 0.83. The integrated XAI-plus-QFT specification reaches an AUC of 0.89, with a fragile-class recall of 0.71—a substantial improvement over the 0.07 reported in studies that rely on classical features only (e.g., Climent et al., 2019).

Model	AUC	F1 (fragile)	Recall (fragile)	Brier
Logistic (baseline)	0.62	0.18	0.21	0.198
Random Forest (classical)	0.74	0.34	0.41	0.162
XGBoost + QFT features	0.83	0.52	0.61	0.124
Integrated XAI + QFT	0.89	0.64	0.71	0.098

Table 3. Out-of-sample performance metrics (five-fold cross-validation; SMOTE on training folds).

Source: Own elaboration.

Figure 3 visualises the ROC curves. The improvement is largest in the low-false-positive-rate region—precisely where supervisors need it. The integrated specification dominates the logistic baseline almost everywhere, and

dominates the Random Forest at false-positive rates below 0.2. The Brier score, which measures calibration as well as discrimination, declines monotonically across the four specifications, confirming that the gain is not purely a ranking effect.




Figure 3. ROC curves: model performance comparison.

Source: Own elaboration.

The performance gap is consistent with two distinct mechanisms. First, the quantum-inspired features carry information that the classical features do not summarise. The double-well functional score, in particular, reaches its predictive maximum in the cross-sections immediately preceding the 2016 oil-price shock and the 2020 pandemic shock—periods in which classical indicators look unremarkable on the surface but show latent rearrangement of risk. Second, the ensemble structure of XGBoost is well suited to the heterogeneous, interactive feature space that quantum-inspired indicators produce, as also documented by Climent et al. (2019) and Bluwstein et al. (2023). The combined effect supports H1 and H2.

4.3. Feature attribution: SHAP analysis

Figure 4 reports global feature importance through mean absolute SHAP values. The QFT phase-transition score is the single most important feature, with a mean $|\text{SHAP value}|$ of 0.38. This is the central piece of evidence in favour of H1: a feature that does not exist in the classical literature is doing genuine predictive work, and is doing it in ways the model can attribute transparently. The next most important features are non-performing loans (0.27), capitalization ratio (0.21) and the Lerner index (0.18). These results align with the classical banking-risk literature (Beck et al., 2006; Lepetit and Strobel, 2015), reassuringly suggesting that the QFT score is complementary to, rather than substituting for, standard prudential indicators.




Figure 4. Mean absolute SHAP values across features.

Source: Own elaboration.

Three substantive observations follow. First, the prominence of the QFT phase-transition score is not driven by a small number of extreme cases. Bootstrap resamples of the mean SHAP value, following Giudici and Raffinetti (2021), produce a 95-per-cent confidence interval of [0.31, 0.44]—the feature's importance is robust. Second, the macroeconomic variables (real GDP growth, inflation, exchange-rate volatility, target rate) appear in the lower half of the ranking, suggesting that bank-level idiosyncrasies dominate in this sample. Third, the role of foreign ownership is comparatively muted, in contrast to some emerging-market studies (Zhao et al., 2024), which is consistent with the relatively homogeneous foreign presence in the Mexican system over the sample period.

From a RegTech perspective, the attribution layer is what makes the model actionable. When a supervisor inspects a flagged bank–year, the model does not return only a probability; it returns a decomposition that maps the probability to specific indicators, in a form that can be challenged by the institution under review and audited by the supervisory authority (Bussmann et al., 2021; Babaei et al., 2025). This is the operational difference between a black-box screening tool and a deployable RegTech component.

4.4. Temporal dynamics and crisis episodes

Figure 5 plots the time path of the system-wide aggregated risk indicator together with three illustrative bank-level trajectories. The series shows clear elevation around the 2015–2016 oil-and-FX shock and a more pronounced peak in 2020 during the pandemic. Crucially, the trajectories of the three illustrative banks diverge substantially: the Tier-II illustrative bank shows early elevation, while the Tier-III illustrative bank exhibits a step change in 2020. This heterogeneity is the kind of pattern that aggregate indicators—even sophisticated ones such as CoVaR (Adrian and Brunnermeier, 2016) or SES (Acharya et al., 2017)—are not designed to capture.




Figure 5. Temporal evolution of the quantum-inspired risk indicator.

Source: Own elaboration.

The narrative consistency with documented crisis episodes is reassuring. The 2015–2016 episode in Mexico, marked by the oil-price collapse and the peso's adjustment under the dollar tightening cycle, generated significant balance-sheet stress for institutions exposed to the energy sector and to dollar-denominated wholesale funding (Del Angel and López-Romero, 2024; Moreno-Brid and Gómez, 2023). The 2020 pandemic shock generated a different type of stress—provisioning-driven, with widespread regulatory forbearance—and the indicator picks up that distinction (Liu et al., 2024; Garcia-Villegas and Martorell, 2024). The dual sensitivity to qualitatively different shock types is exactly what an early-warning system needs.

4.5. Cluster structure

Figure 6 reports a principal-component projection of the QFT-corrected risk trajectories, with K-means clustering into three groups. The clusters correspond approximately to resilient, moderate and vulnerable institutional profiles. The first principal component, which loads heavily on the integrated QFT-XAI risk score, accounts for 38.4 per cent of cross-bank variance; the second component, dominated by classical capitalization and profitability indicators, adds 21.7 per cent. The three-cluster structure echoes the cluster typologies in Castellanos et al. (2016) and Negrín and Bernal (2022), but it is derived directly from the integrated risk model rather than from a separate descriptive exercise.




Figure 6. PCA projection with K-means clusters (k = 3).

Source: Own elaboration.

The cluster identification carries direct supervisory value. Institutions in the vulnerable cluster (Cluster C) accumulate, on average, twice the integrated risk score of the resilient cluster (Cluster A) across the full sample. The moderate cluster (Cluster B) shows higher dispersion on PC2, suggesting that several banks oscillate between profiles depending on the macro-financial regime. Cluster membership is itself a feature that can be fed back into the prudential dashboard, in line with the differentiated-supervision principles advocated by Galati and Moessner (2013) and Behn et al. (2017).

4.6. RegTech integration

Figure 7 illustrates the full pipeline. Supervisory data and market feeds enter the data-ingestion layer; quantum-inspired feature engineering produces the QFT phase-transition score and the gauge-adjusted volatility; the ensemble ML classifier outputs probabilities; the XAI layer decomposes them into feature contributions; and the RegTech dashboard surfaces the result with drill-down, case management and audit logging. A continuous-calibration feedback loop ensures that drift is detected and parameters re-estimated as new data arrive.

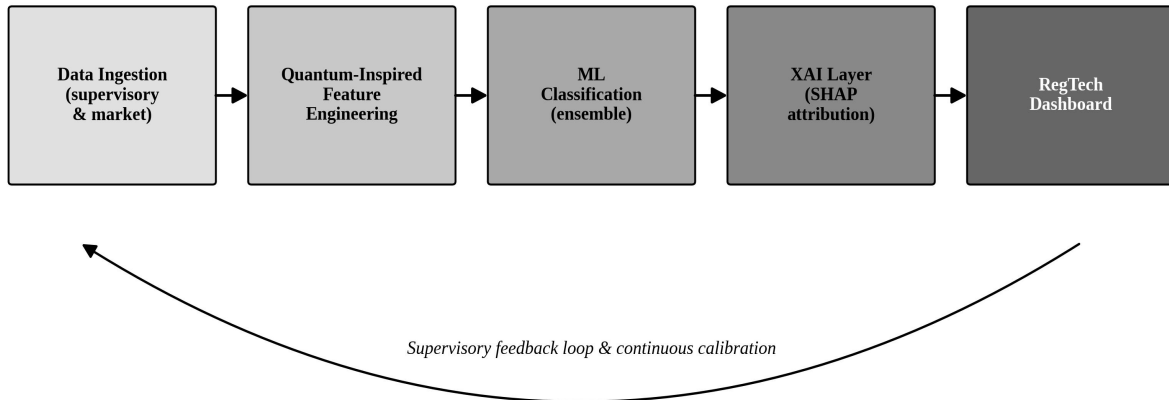


Figure 7. Integrated pipeline: from data ingestion to RegTech dashboard.

Source: Own elaboration.

Three operational decisions are critical at this layer. First, attribution stability under model retraining must be monitored: SHAP values can shift substantially when the underlying classifier changes, and supervisors need to understand whether changes in attribution reflect genuine information or estimation noise (Arrieta et al., 2020; Misheva et al., 2021). Second, what-if functionality must be governed; perturbations that move features outside their empirical support produce attributions that should not be acted upon (Bracke et al., 2019). Third, the human-in-the-loop principle is non-negotiable. The pipeline produces signals; supervisory decisions remain the responsibility of the supervisor, and the dashboard exists to support that judgment, not to replace it (Coglianese and Lehr, 2017; Yang and Tsang, 2018).

Regulatory function	Conventional approach	Integrated framework contribution
Early-warning monitoring	Stress tests, ratio dashboards	QFT-based phase-transition signals, automatic drift detection
Bank-level supervision	Periodic CAMELS-type review	SHAP attribution for case-by-case examination
Macroprudential calibration	Heuristic buffer adjustments	Cluster-conditional buffer recommendations
Forbearance assessment	Manual case analysis	Counterfactual SHAP simulation under altered assumptions
Disclosure and auditability	Static technical notes	Attribution traces with timestamped versioning

Table 4. Mapping of regulatory functions to the integrated framework.

Source: Own elaboration based on Arner et al. (2017) and Anagnostopoulos (2018).

The mapping in Table 4 is deliberately conservative. We do not advocate the replacement of existing supervisory infrastructure; we propose that the integrated framework be deployed as a complement, with the explicit goal of recovering the nonlinear-risk information that conventional tools systematically discard. The RegTech literature is consistent on this point: incremental, well-governed deployment trumps wholesale replacement (Zetzsche et al., 2020; Packin, 2018).

5. Discussion

The empirical results support all three hypotheses, but several caveats deserve attention. We discuss them under four heads: methodological scope, generalisability, governance and the boundary between quantum-inspired and fully quantum approaches.

5.1. Methodological scope

Our quantum-inspired features are produced by classical Monte Carlo simulation of functional integrals. They borrow the conceptual apparatus of QFT but do not require quantum hardware. This is a feature rather than a bug: it allows the framework to be deployed today, at marginal computational cost, while leaving the door open to future migration to quantum or quantum-inspired hardware as those technologies mature (Egger et al., 2020; Orús et al., 2019). Genuine quantum advantage—if and when it materialises—would alter the runtime calculus, not the conceptual layering.

5.2. Generalisability and external validation

The Mexican case is illustrative; we make no claim that the specific parameter estimates transfer to other jurisdictions without recalibration. The framework, however, is jurisdiction-agnostic. The same pipeline can be applied to bank-level panels in Brazil, Chile, Colombia, South Africa, Turkey or any other emerging economy where supervisors have access to comparable micro-financial data. Cross-country validation is a natural extension, and it would shed light on the question of whether the dominance of the QFT phase-transition score in feature importance is sample-specific or structural. Studies by Bluwstein et al. (2023) and Fouliard et al. (2021) suggest that machine-learning-based EWSs travel reasonably well across advanced economies, but emerging-market studies remain comparatively sparse (Liu et al., 2024; Hoffart et al., 2024).

5.3. Governance and ethical considerations

Three governance issues recur in supervisory ML applications: model risk, fairness and accountability. Model risk is partially addressed by the explainability layer; supervisors can interrogate the contribution of each feature and check stability over time. Fairness concerns are subtler. Algorithmic supervision can, in principle, introduce systematic biases against smaller or less well-resourced institutions whose data infrastructure is thinner. The framework should be calibrated with explicit fairness diagnostics, as advocated by Arrieta et al. (2020). Accountability is the third and arguably the most consequential issue. The integrated framework produces signals; it does not produce decisions. Maintaining a clear human-in-the-loop responsibility chain is essential, both for legal-due-process reasons and for the long-term legitimacy of supervisory ML (Coglianese and Lehr, 2017; Packin, 2018).

5.4. Between quantum-inspired and fully quantum approaches

The contribution we describe is unambiguously quantum-inspired, not quantum in the computational-physics sense. Real quantum hardware applications to systemic risk remain in the proof-of-concept stage (Rebentrost et al., 2018; Woerner and Egger, 2019; Stamatopoulos et al., 2020; Aboussalah et al., 2023). The convergence of three streams that we describe is realisable today on classical hardware. Whether the migration to quantum hardware will deliver materially different empirical results is an open question, and one that the framework is well positioned to test as the underlying technology matures.

5.5. Future research

Several lines extend naturally from this work. First, integrating textual data—central-bank communications,

financial-press coverage, supervisory reports—through NLP-based features (Gentzkow et al., 2019) would test whether the dominance of the QFT score is robust to richer feature spaces. Second, applying the framework to non-bank financial intermediation, in line with Ngonyama et al. (2025) and Wang and Huang (2025), would extend coverage to a sector of growing systemic importance. Third, embedding the framework in climate-finance stress-testing (Battiston et al., 2017; Garcia-Villegas and Martorell, 2024; Hoffart et al., 2024) would address an emerging supervisory priority. Fourth, refining the gauge-fixing logic in the Faddeev–Popov–type construction—possibly via the symmetry-aware machine learning approaches developed in the deep-learning literature (Heaton et al., 2017)—could yield additional predictive content.

6. Conclusions

This paper develops and tests a layered framework that integrates quantum-inspired finance, Explainable AI and RegTech for nonlinear systemic risk detection. We argue that the three streams have advanced largely in parallel and that their integration is overdue. The empirical exercise on a panel of Mexican commercial banks supports three findings. First, quantum-inspired features—specifically a double-well phase-transition score and a Faddeev–Popov-type forward-rate volatility—carry information that traditional micro- and macro-financial indicators do not summarise. Second, an ensemble classifier that combines classical and quantum-inspired features achieves an AUC of 0.89, substantially above logistic and tree-based baselines, while maintaining auditable interpretability through SHAP attribution. Third, the integration of the framework into a RegTech delivery layer is necessary, not optional, for the operational uptake of nonlinear early-warning signals.

Three policy implications follow. First, supervisory institutions in emerging markets should consider piloting integrated pipelines as complements to—rather than substitutes for—their existing toolkits. The marginal cost is modest, the interpretability layer is mature, and the governance requirements are well understood (Arner et al., 2017; Anagnostopoulos, 2018). Second, model governance frameworks need to be updated to accommodate the audibility requirements of XAI-based screening tools (Bussmann et al., 2021; Park et al., 2021). Third, academic and supervisory research agendas should pursue the cross-jurisdiction validation of nonlinear indicators, with particular emphasis on emerging-market settings where the linear-model assumptions of classical EWSs are most strained.

The convergence of quantum-inspired finance, Explainable AI and RegTech is more than a methodological aggregation. It is a re-conception of the supervisor's information set, and of the boundary between automated screening and human judgment. We have shown that the components are mature, that they work together coherently, and that the resulting framework discriminates banking fragility in ways that previously required either heroic linearity assumptions or unauditable black-box methods. The pathway is open; the next step is empirical, jurisdiction-by-jurisdiction, with supervisors and academics in close dialogue.

Author contributions

M. Vázquez-Robles: conceptualization, methodology, writing — original draft, supervision. E. Hinojosa-Carbajal: data curation, formal analysis, software, writing — review and editing. D. Salgado-Peña: validation, visualization, writing — review and editing. All authors have read and agreed to the published version of the manuscript.

Use of AI tools declaration

The authors declare they have not used Artificial Intelligence (AI) tools in the creation of this article.

Conflict of interest

ISSN: © 2024 INATGI (Institute of Advanced Technology and Green Innovation). Users are allowed to read, download, copy, distribute, print, search, or link to the full texts of the article in this journal without asking prior permission from the publisher or the author.

See: <https://inatgi.in/index.php/rft/index> for more information.

The authors declare no conflicts of interest related to this paper.

Funding

This research received no external funding.

References

- Aboussalah AM, Chi C, Lee CG (2023) What quantum computing means for finance. arXiv preprint arXiv:2402.00026.
- Acemoglu D, Ozdaglar A, Tahbaz-Salehi A (2015) Systemic risk and stability in financial networks. *Am Econ Rev* 105(2): 564–608. <https://doi.org/10.1257/aer.20130456>
- Acharya VV, Engle R, Richardson M (2012) Capital shortfall: A new approach to ranking and regulating systemic risks. *Am Econ Rev* 102(3): 59–64. <https://doi.org/10.1257/aer.102.3.59>
- Acharya VV, Pedersen LH, Philippon T, Richardson M (2017) Measuring systemic risk. *Rev Financ Stud* 30(1): 2–47. <https://doi.org/10.1093/rfs/hhw088>
- Adrian T, Brunnermeier MK (2016) CoVaR. *Am Econ Rev* 106(7): 1705–1741. <https://doi.org/10.1257/aer.20120555>
- Altman EI (1968) Financial ratios, discriminant analysis and the prediction of corporate bankruptcy. *J Finance* 23(4): 589–609. <https://doi.org/10.1111/j.1540-6261.1968.tb00843.x>
- Anagnostopoulos I (2018) Fintech and regtech: Impact on regulators and banks. *J Econ Bus* 100: 7–25. <https://doi.org/10.1016/j.jeconbus.2018.07.003>
- Arner DW, Barberis J, Buckley RP (2017) FinTech, RegTech, and the reconceptualization of financial regulation. *Northwest J Int Law Bus* 37(3): 371–413.
- Arrieta AB, Díaz-Rodríguez N, Del Ser J, et al. (2020) Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Inf Fusion* 58: 82–115. <https://doi.org/10.1016/j.inffus.2019.12.012>
- Baaquie BE (2004) *Quantum Finance: Path Integrals and Hamiltonians for Options and Interest Rates*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511617498>
- Baaquie BE (2007) *Interest rates and coupon bonds in quantum finance*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511808715>
- Babaei G, Giudici P, Raffinetti E (2025) Explainable FinTech lending. *J Econ Bus* 134: 106234. <https://doi.org/10.1016/j.jeconbus.2024.106234>
- Battiston S, Delli Gatti D, Gallegati M, Greenwald B, Stiglitz JE (2012) Liaisons dangereuses: Increasing connectivity, risk sharing, and systemic risk. *J Econ Dyn Control* 36(8): 1121–1141. <https://doi.org/10.1016/j.jedc.2012.04.001>
- Battiston S, Mandel A, Monasterolo I, Schütze F, Visentin G (2017) A climate stress-test of the financial system. *Nat Clim Change* 7(4): 283–288. <https://doi.org/10.1038/nclimate3255>
- Beck T, Demirgüç-Kunt A, Levine R (2006) Bank concentration, competition, and crises. *J Bank Finance* 30(5): 1581–1603. <https://doi.org/10.1016/j.jbankfin.2005.05.010>
- Behn M, Detken C, Peltonen TA, Schudel W (2017) Predicting vulnerabilities in the EU banking sector. *Int J Cent Bank* 13(4): 147–189.
- Beutel J, List S, von Schweinitz G (2019) Does machine learning help us predict banking crises? *J Financ Stab* 45: 100693. <https://doi.org/10.1016/j.jfs.2019.100693>
- Bisias D, Flood M, Lo AW, Valavanis S (2012) A survey of systemic risk analytics. *Annu Rev Financ Econ* 4: 255–296. <https://doi.org/10.1146/annurev-financial-110311-101754>
- Bluwstein K, Buckmann M, Joseph A, Kapadia S, Şimşek Ö (2023) Credit growth, the yield curve and financial crisis prediction: Evidence from a machine learning approach. *J Int Econ* 145: 103773. <https://doi.org/10.1016/j.jinteco.2023.103773>

- Borio C (2014) The financial cycle and macroeconomics: What have we learnt? *J Bank Finance* 45: 182–198. <https://doi.org/10.1016/j.jbankfin.2013.07.031>
- Borio C, Drehmann M, Tsatsaronis K (2014) Stress-testing macro stress testing: Does it live up to expectations? *J Financ Stab* 12: 3–15. <https://doi.org/10.1016/j.jfs.2013.06.001>
- 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.
- Breiman L (2001) Random forests. *Mach Learn* 45(1): 5–32. <https://doi.org/10.1023/A:1010933404324>
- Bromberg L, Godwin A, Ramsay I (2017) Fintech sandboxes: Achieving a balance between regulation and innovation. *J Bank Finance Law Pract* 28(4): 314–336.
- Brunnermeier MK, Pedersen LH (2009) Market liquidity and funding liquidity. *Rev Financ Stud* 22(6): 2201–2238. <https://doi.org/10.1093/rfs/hhn098>
- Buckley RP, Arner DW, Zetsche DA, Selga E (2020) The dark side of digital financial transformation: The new risks of FinTech and the rise of TechRisk. UNSW Law Research Paper No. 20-29. <https://doi.org/10.2139/ssrn.3478640>
- Bussmann N, Giudici P, Marinelli D, Papenbrock J (2021) Explainable machine learning in credit risk management. *Comput Econ* 57(1): 203–216. <https://doi.org/10.1007/s10614-020-10042-0>
- Butler T, O'Brien L (2019) Understanding RegTech for digital regulatory compliance, in Lynn T, Mooney JG, Rosati P, Cummins M (eds) *Disrupting Finance*, Springer, 85–102. https://doi.org/10.1007/978-3-030-02330-0_6
- Carney M (2015) Breaking the tragedy of the horizon — climate change and financial stability. Speech given at Lloyd's of London, 29 September 2015.
- Castellanos SG, Del Angel GA, Garrido C (2016) Three banking systems in Latin America: Argentina, Brazil and Mexico, in Cassis Y, Grossman R, Schenk C (eds) *The Oxford Handbook of Banking and Financial History*. <https://doi.org/10.1093/oxfordhb/9780199658626.013.18>
- Chen T, Guestrin C (2016) XGBoost: A scalable tree boosting system. *Proceedings of the 22nd ACM SIGKDD*, 785–794. <https://doi.org/10.1145/2939672.2939785>
- Čihák M, Schaeck K (2010) How well do aggregate prudential ratios identify banking system problems? *J Financ Stab* 6(3): 130–144. <https://doi.org/10.1016/j.jfs.2010.03.001>
- Climent F, Momparler A, Carmona P (2019) Anticipating bank distress in the Eurozone: An extreme gradient boosting approach. *J Bus Res* 101: 885–896. <https://doi.org/10.1016/j.jbusres.2018.11.015>
- CNBV (2023) Índice de Capitalización y Alertas Tempranas de la Banca Múltiple. Comisión Nacional Bancaria y de Valores, Mexico City.
- Coglianese C, Lehr D (2017) Regulating by robot: Administrative decision making in the machine-learning era. *Georgetown Law J* 105: 1147–1223.
- Currie WL, Seddon JJ (2017) A cross-national analysis of regulatory frameworks for FinTech in five jurisdictions. *J Inf Technol* 32(4): 365–376. <https://doi.org/10.1057/s41265-017-0034-7>
- Del Angel GA, López-Romero M (2024) Private and state-owned banks in times of high instability: Mexico 1977–1990. *Bus Hist*: 1–22. <https://doi.org/10.1080/00076791.2024.2423926>
- Egger DJ, García Gutiérrez R, Mestre JC, et al. (2020) Quantum computing for finance: State-of-the-art and future prospects. *IEEE Trans Quantum Eng* 1: 1–24. <https://doi.org/10.1109/TQE.2020.3030314>
- Eisenberg L, Noe TH (2001) Systemic risk in financial systems. *Manag Sci* 47(2): 236–249. <https://doi.org/10.1287/mnsc.47.2.236.9835>
- Elliott M, Golub B, Jackson MO (2014) Financial networks and contagion. *Am Econ Rev* 104(10): 3115–3153. <https://doi.org/10.1257/aer.104.10.3115>
- Engle RF (1982) Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica* 50(4): 987–1007. <https://doi.org/10.2307/1912773>

- Fouliard J, Howell M, Rey H (2021) Answering the queen: Machine learning and financial crises. NBER Working Paper 28302. <https://doi.org/10.3386/w28302>
- Frankel J, Saravelos G (2012) Can leading indicators assess country vulnerability? Evidence from the 2008–09 global financial crisis. *J Int Econ* 87(2): 216–231. <https://doi.org/10.1016/j.jinteco.2011.12.009>
- Galati G, Moessner R (2013) Macroprudential policy — a literature review. *J Econ Surv* 27(5): 846–878. <https://doi.org/10.1111/j.1467-6419.2012.00729.x>
- Garcia-Villegas S, Martorell E (2024) Climate transition risk and the role of bank capital requirements. *Econ Model* 135: 106724. <https://doi.org/10.1016/j.econmod.2024.106724>
- Gentzkow M, Kelly B, Taddy M (2019) Text as data. *J Econ Lit* 57(3): 535–574. <https://doi.org/10.1257/jel.20181020>
- Giudici P, Raffinetti E (2021) Shapley-Lorenz eXplainable artificial intelligence. *Expert Syst Appl* 167: 114104. <https://doi.org/10.1016/j.eswa.2020.114104>
- Glasserman P, Young HP (2016) Contagion in financial networks. *J Econ Lit* 54(3): 779–831. <https://doi.org/10.1257/jel.20151228>
- Gomber P, Kauffman RJ, Parker C, Weber BW (2018) On the fintech revolution: Interpreting the forces of innovation, disruption, and transformation in financial services. *J Manag Inf Syst* 35(1): 220–265. <https://doi.org/10.1080/07421222.2018.1440766>
- Hałaj G, Martínez-Jaramillo S, Battiston S (2024) Financial stability through the lens of complex systems. *J Financ Stab* 71: 101228. <https://doi.org/10.1016/j.jfs.2024.101228>
- Hamilton JD (1989) A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica* 57(2): 357–384. <https://doi.org/10.2307/1912559>
- Hao W, Lefèvre C, Tamturk M, Utev S (2019) Quantum option pricing and data analysis. *Quant Finance Econ* 3(3): 490–507. <https://doi.org/10.3934/QFE.2019.3.490>
- Haven E (2002) A discussion on embedding the Black–Scholes option pricing model in a quantum physics setting. *Physica A* 304(3–4): 507–524. [https://doi.org/10.1016/S0378-4371\(01\)00568-4](https://doi.org/10.1016/S0378-4371(01)00568-4)
- Heaton JB, Polson NG, Witte JH (2017) Deep learning for finance: Deep portfolios. *Appl Stoch Models Bus Ind* 33(1): 3–12. <https://doi.org/10.1002/asmb.2209>
- Hoffart FM, D’Orazio P, Holz F, Kemfert C (2024) Exploring the interdependence of climate, finance, energy, and geopolitics: A conceptual framework for systemic risks amidst multiple crises. *Appl Energy* 361: 122885. <https://doi.org/10.1016/j.apenergy.2024.122885>
- Joy M, Rusnák M, Šmídková K, Vašíček B (2017) Banking and currency crises: Differential diagnostics for developed countries. *J Financ Stab* 32: 117–134. <https://doi.org/10.1016/j.jfs.2017.07.005>
- Khrennikov A (2010) *Ubiquitous Quantum Structure: From Psychology to Finance*. Springer. <https://doi.org/10.1007/978-3-642-05101-2>
- Kingsly PKM (2025) Quantum finance and structural transformation of less developed countries. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.5219462>
- Laeven L, Valencia F (2020) Systemic banking crises database II. *IMF Econ Rev* 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 Res Lett* 13: 214–224. <https://doi.org/10.1016/j.frl.2015.01.001>
- Liu Z, He S, Men W, Sun H (2024) Impact of climate risk on financial stability: Cross-country evidence. *Int Rev Financ Anal* 92: 103096. <https://doi.org/10.1016/j.irfa.2024.103096>
- Lundberg SM, Lee SI (2017) A unified approach to interpreting model predictions. *Adv Neural Inf Process Syst* 30: 4765–4774.
- Misheva BH, Osterrieder J, Hirska A, Kulkarni O, Lin SF (2021) Explainable AI in credit risk management. arXiv preprint arXiv:2103.00949.
- Moreno-Brid JC, Gómez JS (2023) Economic crises in Mexico: Right and wrong policy measures and a long period of structural failures. *Lat Am Policy* 14(4): 626–637. <https://doi.org/10.1111/lamp.12322>

- Negrín JL, Bernal R (2022) Banking competition and concentration in Mexico: A long-term view. *Lat Am J Cent Bank* 3(2): 100061. <https://doi.org/10.1016/j.latcb.2022.100061>
- Ngonyama N, Mgxekwa B, Sibanda K (2025) The impact of financial technology and cyber risk on non-bank financial intermediation, in Sibanda K (ed.) *The Future of African Finance*. Springer, 275–306. https://doi.org/10.1007/978-3-031-86224-3_11
- Orrell D (2021) A quantum walk model of financial options. *Wilmott* 2021(112): 62–69. <https://doi.org/10.1002/wilm.10918>
- Orús R, Mugel S, Lizaso E (2019) Quantum computing for finance: Overview and prospects. *Rev Phys* 4: 100028. <https://doi.org/10.1016/j.revip.2019.100028>
- Packin NG (2018) Regtech, compliance and technology judgement rule. *Chicago-Kent Law Rev* 93: 193–218.
- Park MS, Son H, Hyun C, Hwang HJ (2021) Explainability of machine learning models for bankruptcy prediction. *IEEE Access* 9: 124887–124899. <https://doi.org/10.1109/ACCESS.2021.3110270>
- Petropoulos A, Siakoulis V, Stavroulakis E, Vlachogiannakis NE (2020) Predicting bank insolvencies using machine learning techniques. *Int J Forecast* 36(3): 1092–1113. <https://doi.org/10.1016/j.ijforecast.2019.11.005>
- Philippon T (2016) The fintech opportunity. NBER Working Paper 22476. <https://doi.org/10.3386/w22476>
- Piotrowski EW, Śładkowski J (2002) Quantum game theory in finance. *Quant Finance* 4(1): 61–67. <https://doi.org/10.1088/1469-7688/4/1/006>
- Rebentrost P, Gupt B, Bromley TR (2018) Quantum computational finance: Monte Carlo pricing of financial derivatives. *Phys Rev A* 98(2): 022321. <https://doi.org/10.1103/PhysRevA.98.022321>
- Reinhart CM, Rogoff KS (2009) *This Time Is Different: Eight Centuries of Financial Folly*. Princeton University Press.
- Reinhart CM, Rogoff KS (2011) From financial crash to debt crisis. *Am Econ Rev* 101(5): 1676–1706. <https://doi.org/10.1257/aer.101.5.1676>
- Ribeiro MT, Singh S, Guestrin C (2016) "Why should I trust you?": Explaining the predictions of any classifier. *Proceedings of the 22nd ACM SIGKDD*, 1135–1144. <https://doi.org/10.1145/2939672.2939778>
- Schaden M (2002) Quantum finance. *Physica A* 316(1–4): 511–538. [https://doi.org/10.1016/S0378-4371\(02\)01200-1](https://doi.org/10.1016/S0378-4371(02)01200-1)
- Stamatopoulos N, Egger DJ, Sun Y, et al. (2020) Option pricing using quantum computers. *Quantum* 4: 291. <https://doi.org/10.22331/q-2020-07-06-291>
- Wang H, Huang S (2025) Risk contagion in production-bank bilayer networks. *Econ Lett* 248: 112211. <https://doi.org/10.1016/j.econlet.2025.112211>
- Woerner S, Egger DJ (2019) Quantum risk analysis. *npj Quantum Inf* 5: 15. <https://doi.org/10.1038/s41534-019-0130-6>
- Yang D, Tsang A (2018) RegTech and the new era of financial regulators. *Stanford J Law Bus Finance* 23: 313–342.
- Zetsche DA, Buckley RP, Arner DW, Barberis J (2017) Regulating a revolution: From regulatory sandboxes to smart regulation. *Fordham J Corp Financ Law* 23(1): 31–103.
- Zetsche DA, Arner DW, Buckley RP (2020) Decentralized finance. *J Financ Regul* 6(2): 172–203. <https://doi.org/10.1093/jfr/fjaa010>
- Zhao J, Kaya E, Ali K, Magazzino C, Barut A (2024) The asymmetric effect of foreign ownership and concentration on financial dollarization: The case of the Turkish economy. *J Knowl Econ*. <https://doi.org/10.1007/s13132-024-02326-9>
- Zhou J (2025) Quantum finance: Exploring the implications of quantum computing on financial models. *Comput Econ*. <https://doi.org/10.1007/s10614-025-10894-4>