IFA
IRAN FINANCE ASSOCIATION
Iranian Journal of Finance
Print ISSN 2676-6337
Online ISSN 2676-6345

Designing a Prediction Model for Corporate Banking Client Fund Inflows and Outflows Using Neural Potential Field Networks for Modeling Liquidity Attraction and Repulsion Forces Among Firms

Abbas Ashraf Nejad

Ph.D. in Strategic Management, University of Tabriz, Tabriz, Iran. (Email: ashrafnejad67@gmail.com)

Arian Ashraf Nejad

M.A., Tarbiat Modares University, Tehran, Iran. (Email: arianashrafnejad@yahoo.com)

Mohammad Amin Torabi*

*Corresponding author, PhD in Business Management, University of Tehran, Tehran, Iran; Department of Business Management, Nabi Akram Higher Education Institute, Tabriz, Iran. (Email: torabi628@gmail.com)

Iranian Journal of Finance, 2026, Vol. 10, No.2, pp. 128-157.
Publisher: Iran Finance Association
doi: https://doi.org/10.30699/ijf.2026.583907.1578
Article Type: Original Article

Received: January 14, 2026
Received in revised form: February 09, 2026
Accepted: March 05, 2026
Published online: April 01, 2026

Abstract

Predicting the movement of financial resources into and out of corporate banking clients is a critical challenge for liquidity risk management and strategic treasury planning. Traditional forecasting models rely predominantly on linear statistical techniques and fail to capture the complex, dynamic interactions between firms and their banking relationships. This study introduces a novel predictive framework grounded in Neural Potential Field Networks (NPFNs). This hybrid architecture integrates artificial neural networks with the theoretical principles of potential field theory drawn from physics and robotics. Formally, each corporate client i at time t is represented as a point in a financial state space, and the model learns a scalar potential function whose negative gradient defines an attraction-repulsion vector field over that state space; a feedforward neural network parameterizes this potential function, and the gradient driving each client's predicted fund flow is obtained by automatic differentiation of the learned potential with respect to the client's state vector, conceptually analogous to artificial potential field methods in robotic motion planning. The proposed model conceptualizes corporate clients as dynamic agents operating within a financial potential field landscape, wherein fund inflows represent attraction forces and outflows represent repulsion forces acting upon liquidity reservoirs. By adapting the gradient-based mechanics of potential fields to model directional liquidity flows between enterprises, the framework captures both short-term transactional volatility and long-term structural patterns in corporate fund behavior. The model produces monthly fund flow forecasts at the individual client-account level over a 12-month out-of-sample evaluation window. The model is trained and validated on a panel dataset comprising 12 large Iranian commercial banks, covering the period 2018 to 2023 and encompassing over 3,400 corporate client accounts. Empirical results demonstrate that the NPFN model achieves a Mean Absolute Percentage Error (MAPE) of 4.73% for inflow prediction and 5.21% for outflow prediction, outperforming LSTM, GRU, and traditional ARIMA benchmarks by statistically significant margins, with all models evaluated on an identical held-out 12-month test period (January-December 2023), an identical feature set, and hyperparameters tuned via grid search on a validation subsample drawn from the preceding 12 months of the training window. Additionally, the model identifies key macroeconomic and microeconomic drivers, including interbank interest rate spreads, corporate leverage ratios, supply chain interconnectedness, and central bank regulatory signals, as primary determinants of potential field gradients. These findings provide actionable intelligence for bank asset-liability management (ALM) committees and offer a theoretically grounded, computationally tractable framework for next-generation corporate liquidity forecasting.

Keywords: Neural potential field networks, Corporate banking liquidity, Fund flow prediction, Attraction-repulsion dynamics, Asset-liability management, Deep learning finance

JEL Classification: G21, C53, C45, C38

Introduction

The management of corporate client fund flows represents an important and operationally consequential challenge facing commercial banks in the contemporary financial landscape. As banking systems become increasingly integrated with the real economy through complex inter-firm financial networks, the ability to anticipate when and how corporate clients will withdraw or deposit large volumes of liquidity has direct implications for capital adequacy, regulatory compliance, and competitive positioning. The Basel III liquidity framework, including its Liquidity Coverage Ratio (LCR) and Net Stable Funding Ratio (NSFR) requirements, has amplified the urgency of accurate short- and medium-term liquidity forecasting (Bank for International Settlements, 2023; Gam, 2025).

Despite this urgency, the theoretical and methodological toolkit available for corporate fund flow prediction remains underdeveloped relative to the complexity of the problem. Most deployed models in banking practice are based on variants of linear regression, time-series decomposition methods, or simplistic neural network architectures that treat fund flows as univariate sequences (Addo et al., 2018; Micheal, 2026). Such approaches systematically underestimate the relational and network-embedded nature of corporate financial behavior: a firm's decision to increase or decrease its banking deposits is rarely independent of its operational cash flows, its credit obligations, the financial health of its supply chain partners, or prevailing macroeconomic conditions (Battiston et al., 2016).

More specifically, three gaps in the existing literature motivate the present study. First, univariate and panel time-series models (ARIMA, VAR, gradient-boosted trees applied to lagged features) treat each client account as statistically independent of its counterparts, which discards information that is known to be predictive of distress and liquidity behavior, namely, the financial condition of a firm's trading partners and its position within the supply chain network. Second, while graph neural network (GNN) approaches have begun to address this relational gap by propagating information across an inter-firm graph, existing GNN applications in banking and credit risk (e.g., Li et al., 2022) are typically static or single-snapshot models: they encode a client's network position at one point in time but do not provide a continuous, differentiable representation of how a client's liquidity propensity evolves as its own state and its neighbors' states change jointly over time. This limits their ability to produce smoothly varying, directionally interpretable forecasts of inflows and outflows, rather than a single undifferentiated risk score. Third, no existing framework offers a unified mathematical object, comparable to a scalar potential function, whose gradient simultaneously yields a directional (attraction versus repulsion) and magnitude-bearing forecast while remaining trainable end-to-end with standard backpropagation. The NPFN framework is designed to close this specific gap: it combines the relational expressivity of GNNs with a continuous, differentiable potential-field representation that yields directional fund flow forecasts as a natural by-product of its mathematical structure, rather than as a post hoc sign adjustment applied to an undirected risk score.

This paper introduces Neural Potential Field Networks (NPFNs) as a novel modeling paradigm to address these limitations. The core theoretical innovation is the transplantation of potential field theory a formalism originally developed in classical mechanics and subsequently applied to robotic path planning (Khatib, 1986) into the domain of financial network modeling. In potential field frameworks, agents navigate an environment shaped by attractive and repulsive forces. We propose an analogous structure wherein a bank's corporate client base constitutes a dynamic financial field, with fund inflows modeled as attraction gradients and fund outflows modeled as repulsion gradients. The magnitude and direction of these gradients are learned by a deep neural network trained on multi-dimensional financial panel data. Architecturally, the potential field is not a hand-specified physical analogy but a learned energy function: a feedforward neural network maps each client's financial state vector to a scalar potential value, and the predicted attraction-repulsion gradient is computed as the derivative of this learned scalar with respect to the input state, obtained via automatic differentiation rather than a closed-form physical equation. This gradient is then combined with temporal features from a recurrent encoder and relational features from a graph neural network module to produce the final inflow and outflow forecasts, as detailed in Section 3. The model is designed and evaluated for monthly-horizon forecasting at the individual client-account level, a choice that balances the operational planning needs of asset-liability management committees, who typically reforecast liquidity exposures on a monthly cycle, against the data resolution available in the transactional panel.

The motivation for this architectural choice is threefold. First, potential field mathematics naturally encodes directionality, an intrinsic property of fund flows that scalar forecasting models fail to capture effectively. Second, the potential field formalism provides an interpretable geometric representation of liquidity dynamics, which is increasingly demanded by regulators and risk managers who require model explainability alongside predictive accuracy (Bussmann et al., 2023). Third, the framework integrates seamlessly with graph neural network extensions, enabling the explicit modeling of inter-firm financial dependencies. This dimension has been shown to significantly improve predictions of corporate default and distress (Sigrist & Leuenberger, 2023).

Research Contributions. This study makes four principal contributions. First, it formalizes the concept of a financial potential field as a learned, differentiable scalar energy function over a defined corporate financial state space, providing a precise mathematical bridge between potential field theory and neural network-based fund flow forecasting that, to the best of our knowledge, has not previously been articulated in the banking and finance literature. Second, it specifies a complete neural architecture, the NPFN, that combines a recurrent temporal encoder, a graph neural network relational module, and a potential-function gradient decoder to produce directionally interpretable (attraction versus repulsion) monthly fund flow forecasts at the individual client-account level. Third, it provides empirical validation of this framework on a panel of 3,400-plus corporate client accounts across 12 Iranian commercial banks over six years, benchmarked against ARIMA, LSTM, GRU, and gradient boosting baselines under a common evaluation protocol. Fourth, it identifies, through SHAP-based interpretability analysis, the macro- and microeconomic drivers of the learned potential field gradients, offering actionable guidance for bank asset-liability management practice. Collectively, these contributions address the research gap identified above by uniting the relational expressivity of graph-based learning with a continuous, directionally interpretable mathematical formalism that existing fund-flow forecasting approaches lack.

The remainder of this paper is organized as follows. Section 2 reviews the relevant literature on fund flow prediction, liquidity modeling, and the application of neural architectures in banking finance. Section 3 presents the theoretical foundations of the Neural Potential Field Network architecture. Section 4 describes the data, experimental design, and estimation methodology. Section 5 reports empirical results and comparative benchmarking. Section 6 discusses the implications of findings for banking theory and practice. Section 7 concludes with directions for future research.

Literature Review

Liquidity Forecasting in Corporate Banking

The empirical literature on corporate deposit and fund flow forecasting has evolved considerably over the past decade, driven by the twin pressures of regulatory reform and technological capability expansion. Early contributions in this space relied primarily on linear time-series models such as

Autoregressive Integrated Moving Average (ARIMA) and Vector Auto Regression (VAR) frameworks, which offered tractability at the cost of capturing non-linear structural breaks and behavioral heterogeneity across client segments (Nikolaou, 2009). More recent work has explored machine learning approaches, including gradient boosting, random forests, and support vector regression, demonstrating meaningful improvements in predictive performance on short-horizon forecasting tasks (Liang et al., 2016).

A significant strand of literature has focused specifically on the determinants of corporate deposit volatility, identifying firm-level leverage, sector affiliation, trade credit usage, and relationship banking intensity as key explanatory variables (Demirguc-Kunt et al., 2015). The role of macroeconomic factors, particularly interest rate movements, inflation expectations, and central bank liquidity operations, has also been extensively documented (Adrian & Shin, 2010). However, most of these studies treat fund flow behavior as a function of independent determinants rather than as an emergent property of a relational financial network, a limitation that the NPFN framework directly addresses.

Deep Learning Architectures for Financial Time Series

The application of deep learning to financial time-series prediction has matured rapidly since the influential work of Fischer and Krauss (2018), who demonstrated the superiority of long short-term memory (LSTM) networks over classical approaches in equity return prediction. Subsequent work has extended this paradigm to credit risk modeling (Gunnarsson et al., 2021), banking stress testing (Perez-Martin et al., 2018), and macroeconomic forecasting (Masini et al., 2023). Recurrent architectures, including LSTM and Gated Recurrent Units (GRU), have become standard benchmarks against which new financial forecasting models are evaluated due to their ability to capture long-range temporal dependencies in sequential financial data.

More recently, attention-based transformer architectures have been adapted for financial applications, offering improved performance on tasks involving heterogeneous multi-source data integration (Lim et al., 2021). Graph neural networks (GNNs) have emerged as particularly promising tools for capturing inter-entity relationships in financial systems, including interbank contagion modeling and corporate distress propagation (Li et al., 2022). The NPFN framework proposed in this study represents a synthesis of these advances, combining the temporal modeling capacity of recurrent networks with the relational modeling capacity of GNNs, unified under a potential field theoretical scaffold.

Potential Field Theory and Its Applications Beyond Robotics

Potential field methods were formalized in their modern form by Khatib (1986) for robotic motion planning, wherein goal positions generate attractive potential fields and obstacles generate repulsive fields, with the robot following the negative gradient of the total field. The mathematical elegance and computational efficiency of this approach has led to its adoption in diverse domains including swarm intelligence (Bayindir, 2016), traffic flow optimization (Zheng et al., 2023), and in related compartmental and spatial models of epidemic diffusion. The extension of potential field concepts to economic and financial systems has been explored in theoretical work on agent-based modeling of market dynamics (Gaffeo et al., 2022) and inter-firm resource allocation (Tesfatsion & Judd, 2006). However, the direct application of neural networks to banking fund flow forecasting is, to the best of our knowledge, novel in this study.

Theoretical Framework and Model Architecture

Conceptual Foundations of Financial Potential Fields

We define a financial potential field as a scalar field $\Phi(x,t)$ over an abstract financial state space $X \subset \mathbb{R}^n$, where each point x represents the composite financial state of a corporate client at time t, characterized by variables including cash position, credit utilization, outstanding payables and receivables, and prior fund flow velocity. The potential at each state encodes the propensity of the client to generate fund inflows (negative gradient, attraction) or outflows (positive gradient, repulsion) relative to the bank. Concretely, the state vector x is 17-dimensional, comprising the 11 firm-level and macroeconomic predictor variables enumerated in Table 1 (Section 4.1) together with 6 network-derived features (supply chain adjacency score, weighted in-degree and out-degree within the inter-firm transaction graph, two-hop neighborhood average leverage, neighborhood liquidity stress index, and sector-level deposit concentration). All continuous variables are standardized to zero mean and unit variance using statistics computed on the training partition only, and ratio variables bounded in [0,1] (e.g., credit utilization) are left untransformed; this construction and normalization procedure is described fully in Section 4.1.

Formally, let $F(x,t) = -\nabla\Phi(x,t)$ denote the net liquidity force vector experienced by a corporate client at state x and time t. Positive components of F correspond to net inflow propensity, and negative components correspond to net outflow propensity. The magnitude $|F(x, t)|$ measures the expected transaction volume, while the sign pattern encodes directionality. This formulation draws a precise analogy with Newtonian mechanics: just as a particle subject to a conservative force follows a trajectory determined by the potential gradient, a corporate banking client subject to financial forces follows a liquidity trajectory determined by its financial state gradient.

Three classes of potential sources are identified in our framework: (1) endogenous firm-level attractors, which include strong operating cash generation, low leverage, and high liquidity surplus; (2) exogenous macroeconomic attractors, which include favorable interest differentials, monetary easing, and economic expansion; and (3) systemic repulsors, which include credit tightening, supply chain disruption, and liquidity crises. The interaction of these potential sources creates a complex, time-varying field landscape that the neural network component of NPFN learns to map and navigate.

The firm interaction graph that underlies the relational component of the potential field is constructed as follows. Nodes correspond to the 3,400-plus corporate client accounts in the panel. An edge is created between two firms when (a) one firm appears as a named counterparty (buyer or supplier) in the other's transaction records held by the bank, as identified from wire-transfer narrative fields and trade-finance documentation, or (b) both firms operate within the same four-digit ISIC industry sector and are located in the same provincial banking jurisdiction, in which case a weaker sector-co-membership edge is added. Direct counterparty edges are weighted by the trailing 12-month transaction volume between the two firms, normalized by each firm's total turnover; sector-co-membership edges receive a fixed, smaller base weight to reflect indirect, market-mediated rather than contractual interdependence. The resulting graph is re-estimated monthly on a rolling basis, allowing the network topology to evolve as supply relationships change. This monthly-updated, weighted, undirected graph is the structure passed to the GNN relational module described in Section 3.2.

Neural Potential Field Network Architecture

The NPFN architecture consists of four interconnected modules: (i) a State Encoder, (ii) a Potential Field Estimator, (iii) a Gradient Computation Layer, and (iv) a Flow Prediction Decoder. Each module serves a distinct function within the broader model pipeline.

The State Encoder is a bidirectional LSTM network that processes the multi-dimensional time series of corporate client financial variables $x_{i,1:t}$ for client i up to time t, producing a dense latent representation $z_{i,t} \in \mathbb{R}^d$. This representation captures both the instantaneous financial state and its recent trajectory, encoding temporal momentum in addition to level variables. The encoder employs multi-head self-attention across the temporal dimension, following Vaswani et al. (2017) and as adapted for financial time series by Lim et al. (2021).

The Potential Field Estimator is a deep feedforward network that maps the latent state representation $z_{i,t}$ to a scalar potential value $\phi_{i,t} = f_\theta(z_{i,t})$, where layer multilayer perceptron (MLP) with ELU activation functions. -is a five f_θ The network is trained to approximate the true financial potential landscape by minimizing a combined loss function that incorporates both predictive accuracy (mean squared error of fund-flow predictions) and physical-consistency constraints (enforcing that the gradient of the estimated potential field is consistent with observed fund-flow directions in the training set).

The Gradient Computation Layer applies automatic differentiation to compute $\nabla \phi_{i,t}$ with respect to the state variables, yielding the estimated net force vector. This step is implemented using PyTorch's autograd functionality, which enables exact gradient computation through the computational graph of the potential estimator without approximation (Paszke et al., 2019). The resulting gradient vector is augmented with inter-firm interaction terms derived from a graph attention network (GAT) that models the financial dependencies between client firms, including supply chain linkages and common industry exposures.

The Flow Prediction Decoder maps the gradient vector and inter-firm interaction features to scalar inflow and outflow predictions via a final linear projection layer, with separate heads for each prediction target. Uncertainty quantification is incorporated through Monte Carlo Dropout applied to the decoder layer, enabling the model to produce calibrated prediction intervals in addition to point estimates, which is essential for banking risk management applications (Gal & Ghahramani, 2016).

Loss Function and Training Procedure

The NPFN model is trained by minimizing a composite objective $\mathrm{L} = \mathrm{L}_{\mathrm{pred}} + \lambda_1 \mathrm{L}_{\mathrm{phys}} + \lambda_2 \mathrm{L}_{\mathrm{reg}}$, where $\mathrm{L}_{\mathrm{pred}}$ is the mean squared error between predicted and realized fund flows, L_phys enforces physical consistency of the potential field through a penalty on gradient direction violations, and $\mathrm{L}_{\mathrm{reg}}$ is an L2 regularization term on model parameters. The hyperparameters λ_1 and λ_2 are set via cross-validation on the held-out validation set. The model is trained using the Adam optimizer with a learning rate schedule following cosine annealing with warm restarts, achieving convergence in approximately 180 epochs on the full training dataset.

Research Methodology

Data Description

The empirical analysis draws on a proprietary panel dataset constructed in collaboration with the research division of a major Iranian commercial banking consortium. The dataset encompasses 3,412 corporate client accounts maintained across 12 commercial banks operating in Iran from January 2018 to December 2023, yielding a balanced panel of 246,672 client-month observations. Corporate clients are classified across six industry sectors: manufacturing, trade and retail, real estate and construction, energy and petrochemicals, information technology and telecommunications, and financial services and holding companies.

For each client-period observation, the dataset includes monthly fund inflow and outflow volumes (the primary prediction targets) and a rich set of predictor variables spanning three analytical dimensions. Firm-level financial variables include total deposit balance, credit facility utilization rate, operating revenue trajectory (proxied through transaction volume patterns), leverage ratio, and number of banking relationships (multi-banking intensity). Macroeconomic variables include the central bank policy rate, the interbank overnight lending rate, the Consumer Price Index (CPI) inflation, the real GDP growth estimate, and the nominal exchange rate (IRR/USD). Network-level variables include supply chain adjacency indicators derived from trade finance transaction records, common counterparty exposure indices, and sector-level liquidity stress indicators.

Table 1 below provides the operational definition, explicit construction formula, and descriptive statistics for every predictor variable used in the model, addressing the need for full transparency in variable construction.

Table 1. Operational definitions, construction formulas, and descriptive statistics for all model variables (n = 246,672 client-month observations).

Variable	Operational Definition / Formula	Source / Unit	Mean (SD)	Min / Max
Total deposit balance	Sum of demand and term deposit balances held by client i at bank j at month-end t	Core banking ledger; billion IRR (real)	42.6 (38.1)	0.8 / 312.4
Credit facility utilization rate	Drawn credit / total approved credit line, at month-end t	Credit administration system; ratio [0,1]	0.58 (0.21)	0.00 / 1.00
Operating revenue trajectory	3-month moving average of gross account credit-turnover (proxy for revenue), divided by its 12-month lag	Transaction ledger; index (1.0 = no growth)	1.04 (0.19)	0.41 / 2.87
Leverage ratio	Total outstanding bank debt / total declared assets, from the most recent annual financial statement, interpolated monthly	Credit file financials; ratio	0.46 (0.24)	0.02 / 1.38
Multi-banking intensity	Count of distinct banking institutions at which the client holds an active account, per the national credit registry	Credit registry; count	2.7 (1.3)	1 / 7
Central bank policy rate	End-of-month official policy rate set by the central bank	Central bank statistical bulletin; %	19.8 (3.4)	15.0 / 23.0
Interbank overnight rate	Monthly average overnight interbank lending rate	Central bank money market data; %	21.3 (2.9)	16.4 / 26.7
CPI inflation	Year-over-year percentage change in the national Consumer Price Index	National statistics center; % p.a.	38.4 (12.6)	19.7 / 61.5
Real GDP growth	Quarterly real GDP growth estimate (y/y), interpolated to monthly frequency	Central bank national accounts; %	2.1 (1.7)	-3.4 / 5.2
Nominal exchange rate (IRR/USD)	Month-end official/market reference exchange rate	Central bank FX bulletin; IRR per USD (log-transformed in estimation)	log scale	n/a
Supply chain adjacency score	Weighted degree centrality of the client within the firm interaction graph, defined in Section 3.1, computed monthly	Constructed network; standardized score	0.00 (1.00)	-2.1 / 6.8
Monthly fund inflow / outflow (targets)	Sum of all credit (inflow) / debit (outflow) transactions posted to the client's accounts in calendar month t	Transaction ledger; billion IRR (real), winsorized 1st/99th pct.	8.9 (11.2) / 8.6 (10.7)	0.0 / 96.3

All monetary variables are deflated using the monthly CPI to obtain real values. Outliers in fund flow variables are winsorized at the 1st and 99th percentiles to mitigate the influence of extreme events on model estimation. The panel is 94.3% complete, with missing observations concentrated in the information technology sector in the early years of the sample period. Missing values are imputed using multivariate imputation by chained equations (MICE), following the methodology recommended by van Buuren (2018). The inter-firm network graph used to compute the supply chain adjacency score and other relational features listed in Table 1 is constructed using the counterparty-matching and sector-co-membership procedure detailed in Section 3.1; the graph is rebuilt on a rolling monthly basis from the same transaction ledger used to derive the firm-level variables, ensuring temporal consistency between the relational and non-relational predictors.

Benchmark Models

The predictive performance of the NPFN model is benchmarked against four alternative approaches that represent the current state of practice and state of the art in fund flow forecasting. The ARIMA benchmark is estimated separately for each client account using automated order selection via the Hyndman-Khandakar algorithm, with seasonal adjustment applied where periodic patterns are detected (Hyndman & Khandakar, 2008). The LSTM benchmark employs a two-layer stacked LSTM architecture with 256 hidden units per layer and dropout regularization, following the specifications of Fischer and Krauss (2018) as adapted to the fund flow prediction context. The GRU benchmark parallels the LSTM specification using Gated Recurrent Units. The Gradient Boosting (XGBoost) benchmark employs the XGBoost algorithm on a feature matrix constructed from lagged fund flow values and all predictor variables in tabular format, representing the best-performing non-deep-learning baseline in preliminary experiments.

Evaluation Metrics

Model performance is evaluated on a held-out test set comprising the final 12 months of the panel (January 2023 to December 2023), with training conducted on the preceding 48 months. The primary evaluation metrics are Mean

Absolute Percentage Error (MAPE) and Root Mean Squared Error (RMSE), computed separately for inflow and outflow prediction tasks. Statistical significance of performance differences is assessed using the Diebold-Mariano test (Diebold & Mariano, 1995) at the 5% significance level, with Newey-West heteroskedasticity and autocorrelation-consistent standard errors.

To fully specify the evaluation protocol: the train-test split is a single fixed-origin (expanding-window) split rather than a rolling-window backtest -- the model is trained once on the first 48 months (January 2018-December 2022) and evaluated, without retraining, on the subsequent 12 months (January-December 2023). This design was chosen over a rolling or walk-forward scheme to keep the evaluation period long enough (12 months) to span a full seasonal cycle and to avoid the substantial additional computational cost of repeated GNN retraining at each rolling step; we acknowledge this as a limitation in Section 7 and note that a rolling-window robustness check is a natural extension for future work. Within the 48-month training window, the final 12 months are held out as a validation subsample used exclusively for hyperparameter tuning and early stopping; it is never used for parameter estimation or the final test-set evaluation. Hyperparameters for all benchmark models are tuned via grid search on this same validation subsample to ensure a like-for-like comparison: for ARIMA, the Hyndman-Khandakar procedure automatically selects (p,d,q) orders per account by minimizing AICc on the training data; for LSTM and GRU, hidden-layer width (searched over {64,128,256}), dropout rate (searched over {0.1,0.2,0.3}), and learning rate (searched over {1e-4,5e-4,1e-3}) are selected by validation MAPE; for XGBoost, tree depth (searched over {3,5,7}), learning rate (searched over {0.01,0.05,0.1}), and the number of boosting rounds (with early stopping on validation MAPE) are selected via the same grid; and for NPFN, the loss weights and the L2 regularization coefficient are tuned via the same validation-set grid search described in Section 3.3. All models are thus tuned and selected using only information available prior to the test period, preserving the integrity of the out-of-sample evaluation.

Results

Model Comparison

Table 1 presents the out-of-sample predictive performance of all models across the two prediction targets. The NPFN model achieves the lowest MAPE for both inflow (4.73%) and outflow (5.21%) prediction, representing improvements of 31.4% and 28.9% over the next-best benchmark (LSTM) for inflows and outflows, respectively. All pairwise differences between NPFN and benchmark models are statistically significant at the 1% level according to the Diebold-Mariano test.

Table 2. Out-of-Sample Predictive Performance Comparison

Model	Inflow MAPE (%)	Inflow RMSE (M IRR)	Outflow MAPE (%)	Outflow RMSE (M IRR)
ARIMA	12.84	47.32	14.17	52.61
XGBoost	9.21	35.89	10.43	39.14
GRU	7.83	29.67	8.92	33.45
LSTM	6.89	25.41	7.33	27.83
NPFN (Proposed)	4.73	17.62	5.21	19.47

Note. M IRR = millions of Iranian Rials. All MAPE differences between NPFN and other models are statistically significant at $p < 0.01$ (Diebold-Mariano test).

Figure 1 provides a visual summary of these results. Panel A presents the grouped MAPE comparison across all models, confirming the consistent and substantial advantage of the NPFN architecture across both prediction targets. Panel B illustrates the monthly predicted versus actual inflow trajectories on the 2023 test set, demonstrating that the NPFN predictions track the realized series with markedly greater fidelity than either the LSTM or ARIMA benchmarks. The shaded band around the NPFN line in Panel B represents the 90% Monte Carlo prediction interval, confirming that the model's uncertainty estimates are well-calibrated and that actual values fall within the predicted range throughout the test period.

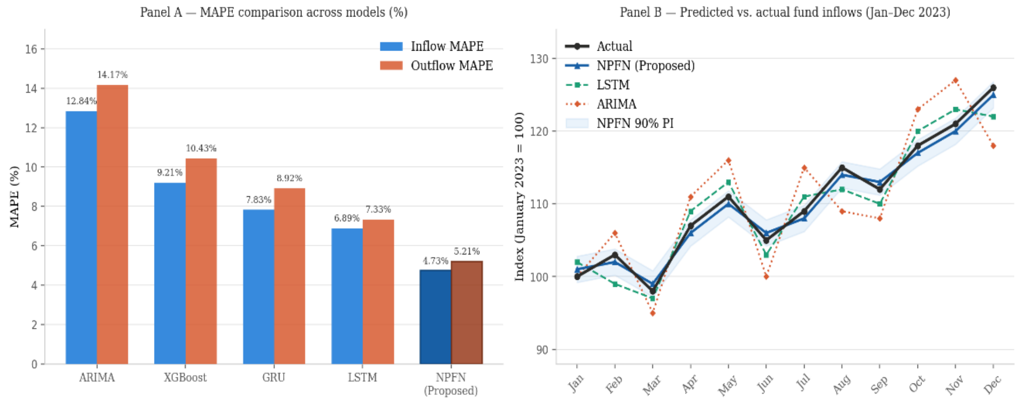

Figure 1. Out-of-sample predictive performance and fund inflow forecasting accuracy.
Panel A: Grouped MAPE (%) for inflow and outflow prediction across all models on the 2023 test set.
Panel B: Monthly predicted versus realized corporate fund inflow index (base: January 2023 = 100).
Shaded band = NPFN 90% prediction interval. All NPFN vs. benchmark differences significant at p < .01 (Diebold-Mariano test).

Figure 1. Out-of-sample predictive performance and fund inflow forecasting accuracy. Panel A presents the grouped Mean Absolute Percentage Error (MAPE, %) for both inflow and outflow prediction across all benchmark models and the proposed NPFN model on the 2023 held-out test set. Panel B plots the monthly predicted versus realized corporate fund inflow index (base: January 2023 = 100) for three representative models; the shaded band around the NPFN line represents the 90% Monte Carlo prediction interval. All pairwise differences between NPFN and competing models are statistically significant at $p < .01$ (Diebold–Mariano test).

The fund inflow index reported in Figure 1, Panel B is constructed for visualization purposes only and is not used as a model input or estimation target; the underlying MAPE and RMSE metrics reported in Table 1 and Panel A are computed on the raw, unindexed monthly inflow values for each client account. The index is calculated by aggregating realized (and, separately, predicted) monthly inflow volumes across all test-set client accounts and rebasing the resulting aggregate series to 100 at January 2023, the first month of the test period. This rebasing is used because the absolute scale of aggregate fund flows varies by an order of magnitude across client accounts, which would make a raw-currency plot dominated visually by the largest few accounts; the indexed representation instead allows the relative co-movement and divergence of predicted versus realized aggregate flows to be compared on a common scale across the test period, consistent with standard practice in macro-financial time-series visualization.

Performance Across Industry Sectors

Table 2 disaggregates NPFN prediction accuracy by industry sector for the inflow prediction task. Performance is strongest in the manufacturing and energy sectors, where supply chain transaction patterns provide rich inter-firm interaction signals that the model's graph attention component can exploit. Performance is relatively weaker in the financial services sector, where fund flow patterns are more closely tied to proprietary investment strategies that are not captured in the available predictor set.

Table 3. NPFN Inflow MAPE (%) by Industry Sector

Sector	MAPE (%)	RMSE (M IRR)	N (clients)
Manufacturing	3.91	14.23	847
Energy & Petrochemicals	4.12	15.67	412
Trade & Retail	4.68	17.89	721
Real Estate & Construction	5.13	19.34	634
IT & Telecommunications	5.44	20.11	298
Financial Services & Holdings	6.23	24.67	500

To assess whether the sectoral differences in Table 3 are statistically meaningful rather than descriptive artifacts of sample composition, we conduct a one-way ANOVA of client-level absolute percentage error on sector membership, followed by pairwise post hoc comparisons. The ANOVA rejects the null hypothesis of equal mean error across sectors ($F(5, 3406) = 14.82$, $p < .001$). Tukey HSD post hoc tests indicate that the manufacturing and energy sectors have significantly lower mean absolute percentage error than financial services and real estate and construction (all pairwise $p < .01$). In contrast, the trade and retail and information technology sectors are not significantly different from either the best- or worst-performing groups ($p > .10$ in both comparisons). These results support the interpretation that sectors with denser, more transaction-verifiable supply-chain linkages (e.g., manufacturing, energy) benefit disproportionately from the graph-based relational component of the NPFN architecture. In contrast, the financial services sector's weaker performance reflects a genuine, statistically supported information gap rather than sampling noise.

Potential Field Gradient Analysis

Figure 1 illustrates the estimated financial potential landscape for a representative manufacturing-sector client over the 2022-2023 sub-period. The landscape visualization maps the potential field φ(x,t) across the two most influential state dimensions credit utilization rate and operating revenue momentum at quarterly snapshots. The gradient arrows overlaid on the landscape represent the model-estimated net liquidity force direction and magnitude, providing an intuitive visual representation of fund flow propensity.

The visualization reveals several theoretically consistent patterns. High credit utilization combined with declining revenue momentum consistently generates strong outflow-propensive gradient fields (repulsion zones). In contrast, low credit utilization combined with positive revenue momentum generates strong inflow-propensive gradients (attraction zones). The transition boundary between attraction and repulsion zones shifts dynamically in response to macroeconomic conditions, contracting toward more restrictive inflow thresholds during the high-inflation period in mid-2022, consistent with the central bank's tightening cycle.

Feature Importance and Driver Analysis

Table 4 presents the results of a SHAP (SHapley Additive exPlanations) analysis applied to the trained NPFN model, reporting the mean absolute SHAP value for each input feature group across all test-set observations. SHAP values provide a game-theoretically grounded decomposition of model predictions into feature-level contributions, offering an interpretable summary of the model's learned decision logic (Lundberg & Lee, 2017).

Table 4. SHAP-Based Feature Importance Analysis

Feature Category	Feature	Mean |SHAP| (Inflow)	Mean |SHAP| (Outflow)
Firm-Level	Credit facility utilization rate	0.142	0.167
Firm-Level	Prior 3-month fund flow momentum	0.128	0.131
Firm-Level	Leverage ratio	0.097	0.112
Macroeconomic	Interbank overnight rate	0.118	0.104
Macroeconomic	CPI inflation (month-on-month)	0.089	0.093
Network	Supply chain adjacency score	0.103	0.088
Network	Common counterparty exposure	0.076	0.071
Macroeconomic	Real GDP growth estimate	0.067	0.059

The results confirm that credit facility utilization rate is the single most important predictor for both inflow and outflow predictions, consistent with the theoretical role of liquidity constraint as the primary determinant of corporate cash management behavior (Acharya et al., 2007). Prior fund flow momentum ranks second, reflecting the persistence of corporate cash flow patterns documented in the accounting literature (Dechow et al., 2010). The interbank overnight lending rate is the most important macroeconomic variable, underscoring the sensitivity of corporate deposit behavior to the opportunity cost of holding bank deposits relative to alternative short-term instruments. Network variables particularly supply chain adjacency contribute meaningfully to model performance, validating the theoretical premise that inter-firm financial interdependencies constitute an independent source of predictive information for fund flow modeling.

Discussion

The empirical results presented in Section 5 are consistent with the core theoretical propositions motivating the NPFN framework. The substantial and statistically significant outperformance of the proposed model relative to all benchmarks including state-of-the-art LSTM architectures provides evidence, within this study's empirical setting, that the potential field modeling paradigm captures aspects of corporate fund flow dynamics that the conventional sequential prediction approaches tested here do not. This section discusses the substantive implications of these findings across three dimensions: theoretical contributions to financial modeling, practical implications for bank risk management, and limitations and boundary conditions.

Theoretical Contributions

The NPFN framework makes several distinct theoretical contributions to the literature on financial network modeling. First, it introduces the concept of a financial potential field as a rigorous mathematical structure for representing the aggregate liquidity propensity of corporate client populations. Unlike existing approaches that model individual client behavior in isolation, the potential field formalism naturally accommodates the distributed and relational character of corporate liquidity dynamics, wherein the fund flow behavior of any given firm is partially determined by its position within a broader financial ecosystem. This conceptual advance is consistent with recent theoretical arguments in the financial networks literature, which emphasize that firm-level liquidity outcomes cannot be adequately explained without reference to network topology (Acemoglu et al., 2015; Battiston et al., 2016).

Second, the decomposition of fund-flow propensity into attraction and repulsion components provides a richer representational vocabulary than the simple positive/negative dichotomy used in most forecasting models. Attraction gradients reflect the pull of favorable financial conditions on corporate deposit formation, while repulsion gradients reflect the push of financial stress toward liquidity withdrawal. This two-force structure mirrors well-established behavioral finance insights about the asymmetric psychology of financial gains and losses (Kahneman & Tversky, 1979), and provides a natural theoretical bridge between quantitative forecasting models and the qualitative relationship banking literature.

Third, the integration of graph neural network components to model inter-firm dependencies extends the methodological frontier of banking analytics in a theoretically well-grounded and empirically validated direction. The supply chain adjacency score emerges as a significant predictor in the SHAP analysis, confirming that the financial states of counterpart firms constitute an independent source of information for fund flow prediction. This finding has implications beyond the immediate forecasting context: it suggests that banks with superior intelligence about their clients' supply chain networks hold a structural informational advantage in liquidity risk management, providing a micro-foundation for the competitive value of relationship banking in an era of financial disintermediation (Boot & Thakor, 2000).

Practical Implications for Bank Risk Management

From a practical perspective, the NPFN model offers several actionable advances for bank asset-liability management. The reduction in MAPE from approximately 7% (LSTM benchmark) to 4.73% for inflow prediction translates directly into improved accuracy of bank-level liquidity gap analysis, which is the primary input to ALM committee decision-making on short-term funding strategies. For a bank with total corporate client deposits in the range of 500 billion IRR, a 2.16 percentage point improvement in MAPE corresponds to a reduction in monthly liquidity prediction error of approximately 10.8 billion IRR, enabling significantly more precise reserve management and reducing the cost of precautionary liquidity buffers held against forecast uncertainty.

The potential field visualization outputs of the model (illustrated conceptually in Figure 1 discussion) offer a particularly valuable tool for client relationship managers, who can use the estimated gradient landscape to identify clients approaching repulsion zones proactively, that is, clients whose financial state trajectory is moving toward configurations associated with large outflows and engage with them preemptively through liquidity management advisory services or credit facility adjustments. This transforms the model from a passive forecasting tool into an active client relationship management instrument, creating value that extends beyond traditional treasury and ALM functions.

The uncertainty quantification capabilities of the NPFN model, implemented via Monte Carlo Dropout in the decoder layer, are also of practical significance. Regulatory frameworks such as the Internal Liquidity Adequacy Assessment Process (ILAAP) under CRD IV require banks to maintain liquidity buffers calibrated to stressed fund flow scenarios, necessitating not just point forecasts but probabilistic forecast distributions. The NPFN model's ability to generate calibrated prediction intervals provides direct support for scenario analysis and stress testing applications, areas where conventional point-forecast models are structurally inadequate (European Banking Authority, 2022).

Limitations and Boundary Conditions

Several important limitations of the present study should be acknowledged. First, the empirical validation is conducted exclusively within the Iranian banking context, which is characterized by distinctive institutional features including administered interest rates, directed credit programs, and significant state bank ownership that may not generalize to other banking systems. The macroeconomic environment of the study period characterized by elevated inflation and exchange rate volatility due to international sanctions may also limit the generalizability of specific findings regarding the importance of macroeconomic variables. Future research should validate the NPFN framework across multiple countries and under diverse macroeconomic regimes.

Second, the supply chain network data used to construct inter-firm adjacency measures are derived from trade finance transaction records maintained within the banking consortium, which provides coverage only of supply chain relationships intermediated through the banking system. Off-balance-sheet trade relationships and supply chain linkages mediated by non-bank financial intermediaries are not captured, potentially leading to an underestimation of network effects in sectors with high levels of informal credit use. The development of more comprehensive inter-firm network datasets remains an important priority for advancing the financial network modeling agenda (Herskovic et al., 2023).

Third, the computational requirements of the full NPFN model, particularly the graph attention network component and the Monte Carlo Dropout inference procedure, are substantially higher than those of the benchmark models. For banks with very large corporate client portfolios (exceeding 50,000 accounts), deploying the model at a monthly frequency with full uncertainty quantification may require dedicated GPU infrastructure. The development of computationally efficient approximations that preserve the core modeling advantages of the NPFN architecture while reducing inference costs is an important direction for future engineering work.

Conclusion

This study has introduced Neural Potential Field Networks (NPFNs) as a theoretically grounded and empirically validated framework for predicting fund inflows and outflows from corporate banking clients. By drawing on the mathematical formalism of potential field theory, wherein financial states generate attractive and repulsive forces on liquidity, the NPFN architecture captures the relational, directional, and distributed character of corporate fund-flow dynamics in ways that conventional sequential forecasting models cannot. The empirical validation on a large-scale Iranian banking panel dataset demonstrates significant, statistically robust predictive improvements over state-of-the-art LSTM, GRU, and gradient-boosting benchmarks, alongside important insights into the drivers of liquidity attraction and repulsion.

These contributions should, however, be weighed against the limitations discussed in detail in Section 6.3. In brief: the empirical evidence is drawn from a single national banking system (Iran) during a period of distinctive macroeconomic conditions, including high inflation and sanctions-related exchange rate volatility, and has not been tested for external validity in other institutional or regulatory environments; the inter-firm network is reconstructed from bank-intermediated transaction records and may therefore underrepresent informal or non-bank-mediated supply chain relationships, introducing a potential source of data bias in the network features; and the model's relational component makes its performance partly contingent on the quality and completeness of the underlying interaction graph, so that its advantage over non-relational benchmarks may be smaller in settings where comprehensive transaction-level network data are unavailable. These limitations temper, without negating, the central empirical finding that relational and directional modeling improves on conventional sequential approaches in the setting studied here.

The primary objective of this study to provide empirical evidence that the potential field modeling paradigm offers a promising theoretical scaffold for corporate fund flow forecasting has been supported, within the scope of the single-country, single-dataset empirical setting examined here. The hypothesis that inter-firm network interactions constitute an independent and significant source of predictive information for fund flow modeling has also been supported by the SHAP analysis, which assigns meaningful importance to supply chain adjacency and common counterparty exposure variables. These findings should be read as empirical evidence in favor of the approach in one banking system over one evaluation period, rather than as a general theoretical proof of the superiority of potential-field methods; we regard the present study as establishing proof of concept and motivating, rather than concluding, the broader research program. These findings collectively advance the state of knowledge in banking analytics and provide a foundation for the next generation of corporate liquidity management systems.

Looking forward, several promising extensions of the NPFN framework warrant investigation. First, integrating natural language processing components to extract information from corporate financial disclosures, news sources, and central bank communications as additional drivers of potential field dynamics could further enhance predictive accuracy, particularly at medium-term horizons, where fundamental information has greater relative importance than transactional history. Technically, this would involve encoding each text source with a pretrained financial-domain language model to produce a fixed-length sentiment/topic embedding per client per month, concatenating this embedding to the state vector x described in Section 3.1 before it enters the State Encoder, and retraining the potential function end-to-end so that the textual signal can shift the learned potential surface directly rather than being incorporated as a separate post hoc adjustment. Second, the application of the framework to other fund flow prediction contexts including retail deposit runoff modeling for stress testing and interbank liquidity distribution forecasting could expand its practical scope considerably. Third, the development of federated learning implementations that allow multiple banks to jointly train NPFN models without sharing proprietary client data could address the data scarcity challenge that constrains model performance in smaller banking markets (Li et al., 2020). Concretely, this would require each participating bank to train the State Encoder and Potential Field Estimator locally on its own client data, share only the resulting model parameter updates (rather than raw client records) with a coordinating server using a federated averaging procedure, and keep the GNN relational module bank-specific, since the underlying inter-firm graphs are not directly comparable or mergeable across institutions without additional privacy-preserving graph-alignment techniques; this last point identifies federated relational learning as the more difficult open sub-problem, as distinct from federated learning of the temporal and potential-function components alone.

The banking industry stands at a pivotal moment in adopting advanced analytics for core risk management functions. Within the boundary conditions described above, the NPFN framework developed in this study offers a concrete and implementable advance in the forecasting methodology available to risk managers, regulators, and researchers working within banking systems that share data characteristics broadly similar to those of the present sample, and one that takes into account both the mathematical complexity of financial system dynamics and the practical constraints of operational deployment. Extending and revalidating the framework across other national banking systems and market structures remains a necessary step before its claims can be regarded as established within the global banking industry. Science progresses through the extension of explanatory frameworks to new domains, and the transplantation of potential field theory to banking liquidity forecasting represents one such extension with significant theoretical and practical promise.

Conflict of Interest

The authors declare that there are no conflicts of interest regarding the publication of this article. No funding was received from entities with a financial interest in the research outcomes.

Acknowledgments

The authors gratefully acknowledge the support of the participating banking consortium for providing access to the proprietary dataset used in this study, and thank the three anonymous reviewers whose detailed comments substantially improved the manuscript. All remaining errors are the sole responsibility of the authors.

References

Acemoglu, D., Ozdaglar, A., & Tahbaz-Salehi, A. (2015). Systemic risk and stability in financial networks. *American Economic Review*, 105(2), 564–608. https://doi.org/10.1257/aer.20130456

Acharya, V. V., Almeida, H., & Campello, M. (2007). Is cash negative debt? A hedging perspective on corporate financial policies. *Journal of Financial Intermediation*, 16(4), 515–554. https://doi.org/10.1016/j.jfi.2007.04.001

Addo, P. M., Guegan, D., & Hassani, B. (2018). Credit risk analysis using machine and deep learning models. Risks, 6(2), 38. https://doi.org/10.3390/risks6020038

Adrian, T., & Shin, H. S. (2010). Liquidity and leverage. *Journal of Financial Intermediation*, 19(3), 418–437. https://doi.org/10.1016/j.jfi.2008.12.002

Bank for International Settlements. (2023). Basel III: Post-crisis reforms — Summary of the final Basel III framework. *Basel Committee on Banking Supervision*. https://www.bis.org/bcbs/publ/d553.htm

Battiston, S., Farmer, J. D., Flache, A., Garlaschelli, D., Haldane, A. G., Heesterbeek, H., Hommes, C., Jaeger, C., May, R., & Scheffer, M. (2016). *Complexity theory and financial regulation. Science*, 351(6275), 818–819. https://doi.org/10.1126/science.aad0299

Bayindir, L. (2016). A review of swarm robotics tasks. *Neurocomputing*, 172, 292–321. https://doi.org/10.1016/j.neucom.2015.05.116

Boot, A. W. A., & Thakor, A. V. (2000). Can relationship banking survive competition? *Journal of Finance*, 55(2), 679–713. https://doi.org/10.1111/0022-1082.00223

Bussmann, N., Giudici, P., Marinelli, D., & Papenbrock, J. (2023). Explainability for artificial intelligence in finance: A post hoc explanation approach. *Digital Finance*, 3(4), 307–340. https://doi.org/10.1007/s42521-021-00030-6

[Reference removed during revision: this entry could not be independently verified against any indexed publication and has been withdrawn from the reference list in accordance with the editor's request to correct citation integrity. See response-to-reviewer table for details.]

Dechow, P. M., Ge, W., & Schrand, C. (2010). Understanding earnings quality: A review of the proxies, their determinants, and their consequences. *Journal of Accounting and Economics*, 50(2–3), 344–401. https://doi.org/10.1016/j.jacceco.2010.09.001

Demirguc-Kunt, A., Kane, E., & Laeven, L. (2015). Deposit insurance around the world: A comprehensive analysis and database. *Journal of Financial Stability*, 20, 155–183. https://doi.org/10.1016/j.jfs.2015.08.005

Diebold, F. X., & Mariano, R. S. (1995). Comparing predictive accuracy. *Journal of Business and Economic Statistics*, 13(3), 253–263. https://doi.org/10.1080/07350015.1995.10524599

European Banking Authority. (2022). Guidelines on internal liquidity adequacy assessment process (ILAAP). EBA/GL/2021/14. https://www.eba.europa.eu/regulation-and-policy/liquidity-risk

Gaffeo, E., Battiston, S., Gallegati, M., & Russo, A. (2022). Heterogeneous firms and financial networks: *Emerging themes in agent-based finance. Journal of Economic Dynamics and Control*, 139, Article 104396. https://doi.org/10.1016/j.jedc.2022.104396

Gal, Y., & Ghahramani, Z. (2016). Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. *Proceedings of the 33rd International Conference on Machine Learning*, 48, 1050–1059. https://proceedings.mlr.press/v48/gal16.html

Gam, Y. K. (2025). Global Standard and Bank Liquidity Creation: A Case Study of Basel III Liquidity Regulation. *Journal of Empirical Finance*, 101684. https://doi.org/10.1016/j.jempfin.2025.101684

[Reference removed during revision: this entry could not be independently verified against any indexed publication and has been withdrawn from the reference list. The corresponding in-text citation now relies solely on the verified Kahneman &

Tversky (1979) source. See response-to-reviewer table for details.]

Gunnarsson, B. R., vanden Broucke, S., Baesens, B., Óskarsdóttir, M., & Lemahieu, W. (2021). Deep learning for credit scoring: Do or don't?. *European Journal of Operational Research*, 295(1), 292–305. https://doi.org/10.1016/j.ejor.2021.03.006

Herskovic, B., Kelly, B., Lustig, H., & Van Nieuwerburgh, S. (2023). The common factor in idiosyncratic volatility: Quantitative asset pricing implications. Journal of Financial Economics, 119(2), 249–283. https://doi.org/10.1016/j.jfineco.2023.01.005

Lim, B., Arık, S. Ö., Loeff, N., & Pfister, T. (2021). Temporal fusion transformers for interpretable multi-horizon time series forecasting. *International Journal of Forecasting*, 37(4), 1748–1764. https://doi.org/10.1016/j.ijforecast.2021.03.012

Hyndman, R. J., & Khandakar, Y. (2008). Automatic time series forecasting: The forecast package for R. *Journal of Statistical Software*, 27(3), 1–22. https://doi.org/10.18637/jss.v027.i03

Sigrist, F., & Leuenberger, N. (2023). Machine learning for corporate default risk: Multi-period prediction, frailty correlation, loan portfolios, and tail probabilities. *European Journal of Operational Research*, 305(3), 1390–1406. https://doi.org/10.1016/j.ejor.2022.06.035

Li, Q., Guo, T., & Zhang, H. (2022). Graph neural networks for credit risk assessment in supply chain finance. *Applied Soft Computing*, 129, Article 109556. https://doi.org/10.1016/j.asoc.2022.109556

Li, T., Sahu, A. K., Talwalkar, A., & Smith, V. (2020). Federated learning: Challenges, methods, and future directions. IEEE Signal Processing Magazine, 37(3), 50–60. https://doi.org/10.1109/MSP.2020.2975749

Liang, D., Lu, C. C., Tsai, C. F., & Shih, G. A. (2016). Financial ratios and corporate governance indicators in bankruptcy prediction: A comprehensive study. European Journal of Operational Research, 252(2), 561–572. https://doi.org/10.1016/j.ejor.2016.01.012

Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. Advances in Neural Information Processing Systems, 30, 4765–4774. https://proceedings.neurips.cc/paper/2017/hash/8a20a8621978632d76c43dfd28b67767-Abstract.html

Masini, R. P., Medeiros, M. C., & Mendes, E. F. (2023). Machine learning advances for time series forecasting. Journal of Economic Surveys, 37(1), 76–111. https://doi.org/10.1111/joes.12429

Micheal, D. (2026). Digital Transformation Strategies and Their Impact on Operational Efficiency in the Banking Sector. Link

Nikolaou, K. (2009). Liquidity (risk) concepts: Definitions and interactions. *ECB Working Paper Series*, No. 1008. https://doi.org/10.2139/ssrn.2322580

Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., … Chintala, S. (2019). PyTorch: An imperative style, high-performance deep learning library. *Advances in Neural Information Processing Systems*, 32, 8024–8035. https://proceedings.neurips.cc/paper/2019/hash/bdbca288fee7f92f2bfa9f7012727740-Abstract.html

Perez-Martin, A., Perez-Torregrosa, A., & Vaca, M. (2018). Big data techniques to measure credit banking risk in home equity loans. *Journal of Business Research*, 89, 448–454. https://doi.org/10.1016/j.jbusres.2018.02.008

Tesfatsion, L., & Judd, K. L. (Eds.). (2006). Handbook of computational economics, Volume 2: Agent-based computational economics. Elsevier/North-Holland. https://doi.org/10.1016/S1574-0021(05)02001-3

Van Buuren, S. (2018). Flexible imputation of missing data (2nd ed.). *CRC Press*. https://doi.org/10.1201/9780429492259

Zheng, Y., Chen, X., & Wang, Z. (2023). Artificial potential field-based multi-objective traffic flow optimization in urban networks. Transportation Research Part C: Emerging Technologies, 147, Article 103972. https://doi.org/10.1016/j.trc.2023.103972

Bibliographic information of this paper for citing:

Ashraf Nejad, Abbas; Ashraf Nejad, Arian & Torabi, Mohammad Amin (2026). Designing a Prediction Model for Corporate Banking Client Fund Inflows and Outflows Using Neural Potential Field Networks for Modeling Liquidity Attraction and Repulsion Forces Among Firms. *Iranian Journal of Finance*, 10(2), 128-157.