
**Graduate Institute of International and Development Studies
International Economics Department
Working Paper Series**

Working Paper No. HEIDWP17-2026

**Intraday Prediction of Operating-Rate Deviations
from the Policy Rate: Evidence from Peru**

Diego Franco

Central Reserve Bank of Peru

Delia Ruiz

Central Reserve Bank of Peru

Walter Cuba

Central Reserve Bank of Peru

Chemin Eugène-Rigot 2
P.O. Box 136
CH - 1211 Geneva 21
Switzerland

Intraday Prediction of Operating-Rate Deviations from the Policy Rate: Evidence from Peru

Diego Franco

Delia Ruiz

Walter Cuba

Banco Central de Reserva del Perú

Abstract

This paper develops an intraday early-warning framework to predict deviations of the volume-weighted overnight interbank rate from the BCRP policy rate after the close of the Central Bank's second intervention window. We study both upward deviations, associated with liquidity-scarcity episodes, and downward deviations, associated with liquidity-abundance episodes. Using a unique high-frequency dataset spanning 2015-2025, we evaluate whether morning liquidity indicators can anticipate rate deviations exceeding 5 basis points. We compare a regularized logistic regression with a nonlinear artificial neural network (ANN), estimating separate models for each direction of deviation. Both models are calibrated on a chronological development sample and evaluated on a held-out test period. The logit model outperforms the ANN in both cases, with a statistically significant ranking advantage (ROC-AUC of 0.95 vs. 0.88 for upward deviations; 0.77 vs. 0.75 for downward deviations).

Average marginal effects reveal an economically coherent asymmetry. Market concentration, measured by the HHI, and cross-bank dispersion in reserve requirement compliance reduce the probability of upward deviations and increase the probability of downward deviations. We interpret this as reflecting the presence of a small number of large, readily identifiable liquidity providers: their visibility reduces search frictions and prevents rate spikes when the market is short, while the same concentration shifts bargaining power toward borrowers, who become the scarce side of the negotiation, when the market is long. Overall, the findings support the feasibility of a simple, interpretable early-warning tool for BCRP money market operators

Keywords: Interbank money market, Monetary policy implementation, Early-warning models, Machine learning, Market concentration, Peru

JEL: E58, E43, G21, C53

The authors thank Cyril Monnet from University of Bern for the academic supervision of this paper. This research took place through the coaching program under the Bilateral Assistance and Capacity Building for Central Banks (BCC), financed by SECO, and the Graduate Institute in Geneva. The views expressed in this paper are solely those of the author(s) and do not necessarily reflect those of the Banco Central de Reserva del Perú.

1. Introduction

Peru has operated an inflation-targeting framework since 2002, with an annual inflation target of 2% and a tolerance band of 1%-3%. The Central Reserve Bank of Peru (BCRP) implements monetary policy through a corridor system, where the key policy reference rate – currently set at 4.25% (as of May 2026) – serves as the operational target for the overnight interbank rate. This rate anchors short-term market rates and guides liquidity conditions, supported by open market operations, reserve requirements, and occasional foreign exchange interventions to ensure effective transmission in a partially dollarized economy.

Ensuring that the interbank rate remains well-anchored to the policy rate is fundamental to effective monetary policy implementation. However, in practice, episodic deviations do occur. Two prominent episodes stand out: (i) the 2022 global inflation shock, which amplified supply-side pressures and caused the spread between the interbank and the policy rate to widen to +40 bps and -60 bps over several months; and (ii) the 2015 period of national currency depreciation against the US dollar, driven by falling commodity prices and market expectations of US monetary policy normalization, during which deviations persisted and in some cases reached more than +50 bps. Beyond these macro-driven episodes, more routine intraday deviations also occur. The BCRP intervenes in the morning (approximately 10:00-11:00) and, if liquidity imbalances persist, again in the early afternoon (approximately 12:00-13:30). After this window, the Bank's regular intervention is mostly due to few specific factors that have not been cleared in the previous hours.

This paper develops an early-warning model that uses information available during the morning to estimate the probability of a relevant rate deviation after 13:30, in either direction. The central hypothesis is that both upward deviations (liquidity scarcity) and downward deviations (liquidity abundance) reflect intraday liquidity conditions and can be anticipated from: (i) early market dynamics such as the morning spread and interbank rate volatility, (ii) the initial liquidity position of the banking system, (iii) the heterogeneity of reserve requirement advances across banks, and (iv) the BCRP's open market operations.

We model the two outcomes separately, as the mechanisms generating upward and downward pressure need not be symmetric – a possibility we examine directly in Section 6. To address this question, we estimate two binary classification models, one for each direction of deviation, in which the outcome equals 1 if the volume-weighted interbank rate after 13:30 deviates from policy rate by more than 5 bps. We use a regularized logistic regression (logit) as a transparent, interpretable benchmark, and compare it against an artificial neural network (ANN) capable of capturing potential nonlinearities. Both models are estimated on a unique dataset of high-frequency intraday transactions to which we have privileged access, and are validated using a walk-forward, chronological scheme designed to replicate the information available to a trading desk in real time.

The paper contributes to two related strands of literature. First, it extends the work on interbank rate dynamics and monetary policy transmission by focusing on intraday deviations rather than daily or monthly averages. Second, it adds to a growing body of applied machine learning in central banking, where neural networks have been used to predict changes in the policy rate, but not short-term operational deviations of the interbank rate from the policy target.

2. Literature review

The economic literature on monetary policy implementation has traditionally focused on the ability of central banks to align short-term market interest rates with their operational targets.

Within an interest rate corridor framework, this alignment relies critically on the daily management of aggregate structural liquidity and the predictability of reserve demand by financial institutions.

Evidence from different monetary systems shows that corridor frameworks can effectively anchor market rates. In a study of the Eurosystem, Nautz and Offermanns (2007) show that the overnight interbank rate exhibits a stable cointegrating relationship with the policy target, with short-term deviations quickly smoothed out through open market operations.

Similarly, Furfine (2002) highlights how the operational framework of the interbank market, together with central bank interventions, allows the system to absorb large credit and liquidity shocks without pushing market rates persistently away from the policy target. In the same spirit, Whitesell (2006) analyzes the efficiency of the corridor system and argues that it can generate flatter reserve demand curves, thereby anchoring the market rate more closely to the policy target. In particular, he shows that narrower corridor spreads make the reserve demand curve flatter.

In the Peruvian context, Pozo and Lupu (2011) use this framework to study the BCRP corridor through a cointegration analysis. Their findings confirm a high speed of convergence and show that liquidity management and open market operations are highly effective complementary tools for keeping the overnight interbank rate anchored. More recently, Velarde (2025) expands on this daily aggregate view of the Peruvian interbank market by analyzing the drivers of rate deviations from 2017 to 2024 using an ordinary least squares (OLS) framework with heteroskedasticity and autocorrelation consistent standard errors.

His findings confirm that initial liquidity levels, reserve requirement advance deviations, and monetary policy expectations are the primary drivers of daily rate pressures. Crucially, Velarde quantifies these mechanics, showing that a 1 billion PEN increase in initial liquidity is associated with a negative 0.04 bps deviation in the interbank rate, and that deviations display moderate persistence: a positive 5 bps deviation generates a deviation of 1 bps after two days. While these findings underscore the daily effectiveness of BCRP in correcting liquidity imbalances, existing aggregate approaches assess rate dynamics using daily, weekly or monthly data. As a result, they exclude the high-frequency relationships that occur within the trading day but are inherently unobservable in end-of-day data.

Another strand of the literature argues that intraday dynamics and the distribution of liquidity across banks play a central role in shaping aggregate liquidity conditions and rate behavior in the interbank market. For example, Bech and Garratt (2003) formalize a game-theoretic framework showing that banks face a strategic trade-off between the cost of holding intraday liquidity and the delay costs associated with postponing payments. This trade-off can generate coordination failures, as banks may hoard liquidity early in the day and thereby trigger structural gridlocks.

Similarly, Acharya et al. (2012) develop a model of imperfect competition in interbank lending and show that aggregate liquidity sufficiency does not necessarily guarantee rate stability. When liquidity is concentrated among a small group of banks, surplus institutions can exploit their market power to extract rents from liquidity-deficient banks, leading to higher borrowing costs even when total systemic reserves are adequate. This structural friction explains why daily aggregate liquidity measures may obscure underlying market stress.

Empirical evidence reinforces the importance of these intraday and distributional frictions. Acharya and Merrouche (2013) document a 30% increase in liquidity demand among large settlement banks immediately following market freezes. This precautionary liquidity hoarding intensifies on days with high payment settlement activity and among institutions facing greater credit risk. The authors show that such behavior significantly increases overnight interbank rates.

In the Peruvian context, Velarde (2025) finds that liquidity concentration, measured by the Herfindahl-Hirschman Index (HHI), has a positive but statistically insignificant effect in his daily aggregate specification. This finding motivates a shift toward measuring cross-bank heterogeneity in reserve positions at higher frequencies, as distributional frictions may emerge within the trading day but remain hidden in end-of-day aggregate data.

Furthermore, Sichei et al. (2012) find that in their analysis of the Kenyan interbank market, large banks engage in price discrimination based on counterparty size, consistently charging higher interest rates to smaller banks than to their larger counterparts. Also, the study highlights that large banks also discriminate regarding the volume of credit extended.

In contrast, Bech & Monnet (2016) employ a search-based model of the interbank money market to demonstrate that the market functions as a decentralized mechanism where banks, constrained by reserve requirements, engage in strategic negotiations to redistribute liquidity. The authors argue that deviations of the interbank interest rate from the policy target are not necessarily evidence of market abuse; rather, they represent the equilibrium outcome of banks maximizing their trade surplus. Crucially, this outcome relies on the assumption of equal bargaining power between counterparties.

Despite the depth of these research streams, a critical gap persists at their intersection. Traditional cointegration, error-correction, and daily regression models (e.g., Nautz & Offermanns, 2007; Pozo & Lupu, 2011; Velarde, 2025) suffer from severe structural limitations when applied to short-term money markets. By aggregating data into daily or weekly frequencies, these linear models inherently assume constant and symmetric relationships between liquidity indicators and interest rate movements. Interbank liquidity distress does not develop smoothly; instead, it manifests through abrupt, non-linear spikes and localized frictions (Bech and Garrat, 2003; Acharya et al., 2012).

These regressions capture end-of-day averages, so they obscure the fact that market stress is highly time-dependent, typically concentrating within specific operational windows, such as the late-afternoon closing session when routine central bank interventions are structurally constrained. Consequently, linear aggregate specifications fail to capture the asymmetric threshold effects where a minor intraday reserve shortfall can trigger a disproportionate spike in rates. This masking effect strongly justifies our shift toward an intraday, high-frequency analytical framework capable of isolating late-afternoon dynamics.

Conversely, existing market microstructure and strategic behavior studies (e.g., Bech & Garratt, 2003; Baglioni & Monticini, 2008; Acharya et al., 2012; Garcia et al., 2025) diagnose intraday stress and expectation shocks conceptually or descriptively, but they lack predictive translation for operational policy. Their empirical methodologies typically rely on contemporaneous equilibrium equations or descriptive event studies that classify market friction after its materialization. Consequently, they fail to translate high-frequency microstructure signals into real-time, forward-looking probability distributions.

Furthermore, a critical dimension of this gap lies in the structural objective of existing empirical frameworks. The current literature provides valuable retrospective and diagnostic explanations of liquidity crises and interbank rate volatility, as it helps identify the drivers of market stress after these episodes have materialized. However, much less attention has been given to predictive, real-time Early Warning Systems (EWS) designed for live central banking operations. Most empirical strategies rely on contemporaneous or lagged daily aggregates, which makes them more suitable as explanatory tools than as anticipatory instruments.

This limitation is particularly relevant for central bank trading desks. In practice, policymakers need tools capable of processing intraday microstructure signals before the interbank market

closes, when there is still room to take corrective action. By not translating the theoretical mechanisms of liquidity frictions into an ex-ante operational framework, the existing literature leaves an important gap for monetary authorities seeking to detect and respond to stress during critical intraday windows.

This paper contributes to bridging this gap in two ways. First, departing from traditional low-frequency transmission analysis, we shift the analytical focus toward intraday rate deviations in the Peruvian interbank market. In particular, we isolate after central bank intervention dynamics, a window in which liquidity pressures may become more relevant for daily interbank operations. This allows us to assess whether cross-bank dispersion in reserve requirement advances contains more predictive information for rate volatility than aggregate liquidity balances.

Second, we extend the application of machine learning tools to central bank operations by developing a binary classification early warning model focused on the operational stability of the policy target. By incorporating intraday spreads and volatility metrics into a neural network, the framework is designed to capture nonlinear rate outcomes triggered by localized liquidity shortfalls. In this sense, the model translates high-frequency market signals into a predictive tool for identifying potential deviations before the market closes.

Methodologically, framing the problem as a binary classification task under severe class imbalance—where significant interest rate deviations are rare but operationally costly events—represents a substantial departure from traditional econometric approaches. While standard macroeconomic regressions minimize average continuous errors, they tend to wash out tail-risk events, leading to poor predictive power during crises. In contrast, by leveraging non-linear network architecture tuned specifically to maximize operational metrics such as the Area Under the Precision-Recall Curve (AUPRC) and critical operational recall, our framework ensures that the central bank trading desk can detect a high proportion of actual liquidity distress episodes. This optimization capability bridges a crucial gap, providing an actionable, real-time warning tool that conventional linear models structurally cannot deliver.

3. Monetary Policy Implementation at the BCRP

BCRP operates a corridor system in which the overnight interbank rate is guided toward the policy reference rate. The corridor is defined by a lending facility rate (ceiling) and a deposit facility rate (floor), both set symmetrically around the policy rate. The Bank intervenes daily to manage system-wide liquidity and keep the interbank rate within this corridor.

Each trading day, BCRP interventions follow a structured sequence. In the morning (approximately 10:00-11:00), the bank conducts open market operations—primarily repurchase agreements (repos) or fixed-term deposits—to address aggregate excess or shortfalls in system liquidity. After the morning session, a second round of operations is conducted in the early afternoon (approximately 12:00-13:30). After that cutoff, the bank no longer conducts routine market-wide interventions. At that stage, liquidity support is limited to direct emergency repos with individual institutions, which are activated generally when a commercial bank records a negative current account balance.

The operational target variable is the volume-weighted overnight interbank rate, which ideally should track the policy rate closely. The interbank market brings together commercial banks that trade short-term liquidity in local currency. Banks with excess reserves lend to banks with shortfalls, either through unsecured interbank loans or through secured repo transactions.

The reserve requirement framework plays an important role in this context. Banks are required to maintain minimum reserve levels, computed as a fraction of their liabilities. Reserve requirement advances allow banks to draw on future reserve positions to meet current obligations, providing a buffer against short-term liquidity shocks. The distribution of reserve requirement advances across banks —not just their aggregate level— can affect whether individual institutions face funding pressures that are not fully offset by market transactions.

This operational framework motivates the early warning model developed in this paper. As the market approaches the afternoon intervention cutoff, routine liquidity management becomes increasingly constrained. Therefore, having a reliable probability estimate of later rate deviations before that window closes would allow money market operators to adjust their interventions accordingly.

4. Data

4.1. Overview

The empirical analysis relies on a unique high-frequency dataset assembled from internal BCRP sources. The dataset covers the period from 2015 to 2025 at a daily frequency, with intraday granularity for most key variables. Access to this data is privileged and not publicly available; it has been obtained through the authors' institutional roles at the BCRP.

The dataset covers 2015-2025 period with complete intraday data. We integrate information from four primary sources:

- Interbank market transactions: every overnight interbank trade (unsecured loans) recorded during the trading day, including transaction time, agreed rate, counterparties, and traded amount. This allows us to compute volume-weighted average rates at different intraday cutoffs (before and after intervention cutoff).
- Open market operations: all BCRP intervention operations (repos, fixed-term deposits and certificate of deposits), including the time the auction was conducted, the amount offered, the bids received, and the amounts allotted.
- Reserve requirement data: bank-level reserve requirement positions and advances for each business day, allowing computation of cross-bank dispersion in reserve positions.
- Current account balances: individual bank balances in their current accounts at the BCRP at the start of each business day, which serve as the initial liquidity stock for the system.

On average, during this period, there are 38 interbank operations on a regular day, with a standard deviation of 19,7. The volume is on average 1 096,4 million PEN and the interbank interest rate is closely on average, standard deviation, minimum and maximum to the policy rate.

Table 1: Daily Descriptive Statistics of Intraday Interbank Market Operations

Variable	Mean	Std. Dev.	Min	Max
Number of operations	37,5	19,7	1	128
Volumen (MM PEN)	1 096,4	566,2	0,1	4 325
Interbank interest rate (%)	3,99	1,95	0,07	7,85
Policy rate (%)	3,97	1,94	0,25	7,75

* The sample covers N = 2 582 days

Analyzing the intraday transaction dynamics before and after cutoff reveals that most of the market activity occurs in the earlier session. Specifically, operations executed before afternoon

intervention cutoff account for 74.8% of the total transaction volume and 74.5% of the total value transacted, with both periods exhibiting similar average transaction sizes. However, notable differences emerge regarding pricing and volatility. The interest rate spread —defined as the transaction rate minus the policy rate in absolute value— is 4 bps higher in the afternoon session.

Furthermore, the standard deviation of these spread reaches 20,8 bps after the afternoon intervention cutoff, representing a 46.3% increase relative to the morning session. This heightened volatility suggests that interbank operations in the latter part of the day carry a higher premium and a significant higher risk.

Table 2: Intraday Segmentation (before and after second intervention window)

Intraday Segmentation		
Variable	Before	After
Number of operations	72 320	24 405
Average transaction size (MM PEN)	29,2	29,5
Sum of transactions (MM PEN)	2 111 060	719 972
Average interest rate (%)	4,2	4,0
Average interest rate spread (bps)	3,7	7,7
Standard deviation interest rate spread (bps)	14,2	20,8

* The interest rate spread is defined as the absolute difference between the transaction's interest rate and the BCRP policy rate.

Examining the daily distribution of the interest rate spread across different intraday periods reveals that, on average, afternoon spreads (after second window) are lower than those observed during the morning session. This pattern contrasts with the transaction-level analysis in the previous table because these daily rates are volume-weighted by the size of each transaction. This finding may reflect the efficacy of the open market operations of BCRP in aligning the operational market rate with the policy target, particularly for larger transaction amounts.

However, both the standard deviation and the extreme percentiles are substantially higher in the afternoon session. This implies that while deviations are generally well-contained by central bank interventions under normal conditions, certain operations —likely driven by liquidity stress scenarios on specific dates— cause afternoon rates to deviate significantly from the policy target. Exploring the determinants of these extreme intraday deviations constitutes the core focus of this research.

Table 3: Distribution of Daily Interbank Interest Rate Spread by Intraday Period

Statistic	Full sample	Before	After
Mean	1,9	2,4	0,4
Std. Dev.	17,1	14,6	23,0
Percentiles:			
5th	-6,2	-4,5	-17,5
10th	-3,1	-1,1	-7,3
25th	-0,1	0,0	0,0
50th	0,0	0,0	0,0
75th	0,2	0,0	0,2
90th	3,2	2,9	3,0
95th	18,8	16,4	22,3

4.2. Variable Construction

The dependent variable is a binary indicator defined as:

$$y_t^+ = 1 [ind_t > 0.05]$$

Where ind_t is the difference between the volume-weighted interbank rate after approximately 13:30 and the BCRP policy rate on day t . Since rates are expressed in percentage points, 0.05 corresponds to 5 bps. The event of interest is therefore an upward deviation greater than 5 bps, a threshold that in the view of BCRP money market operators, represents a materially relevant departure from the policy target.

To capture the symmetric phenomenon of excess-liquidity pressure, we define a second complementary outcome for downward deviations:

$$y_t^- = 1[ind_t < -0.05]$$

Which equals one when the volume-weighted interbank rate after the second intervention window cutoff falls more than 5 bps below the policy rate. The two outcomes are modelled separately throughout the paper — y_t^+ reflecting liquidity-scarcity episodes in which the BCRP would need to inject reserves, and y_t^- reflecting liquidity-abundance episodes in which the BCRP would need to absorb (sterilize) reserves— because the economic mechanisms generating upward and downward pressure are conceptually distinct and, as Section 6 shows, are driven by different combinations of explanatory variables.

The explanatory variables, organized into 5 economic blocks, are common to both specifications with one exception: the block capturing BCRP intervention operations differs by direction, reflecting the operational tool used by the Bank under each type of pressure (Table 1).

Table 4. Variables

Block	Variable	Description
Market Signals	Morning Deviation	Volume-weighted deviation of the interbank rate from the policy rate before intervention cutoff. Captures early morning misalignment.
	Morning Intraday Volatility	Standard deviation of Interbank transaction rates before intervention cutoff. Proxy for morning market volatility.
Autonomous Factors	Projected BN (National Bank) Injection	BCRP's projected injection of liquidity into the system via <i>Banco de la Nación</i> operations (as percentage of initial account balance).
	Projected BN (National Bank) Absorption	BCRP's projected absorption of liquidity into the system via <i>Banco de la Nación</i> operations (as percentage of initial account balance).
Liquidity / Structure	Initial Balance	Initial balance of bank current accounts at the BCRP at start of the day. Proxy for aggregate system liquidity.
	HHI Concentration Index	HHI index computed over the cross-bank distribution of initial current account balances. Proxy for the concentration of system liquidity across institutions.
	Reserve Requirement Deviation (lag 1)	Cross-sectional standard deviation of reserve requirement advances across banks, lagged one day to respect information timing.
BCRP Operations	Injection / Sterilization Amount Awarded (% of Initial Liquidity)	Amount awarded in same-day BCRP injection auctions (y_t^+ model) or sterilization auctions (y_t^- model), expressed as a share of system initial liquidity.
	Injection / Sterilization Unmet Demand Share	Share of auction demand left unsatisfied by the BCRP. Proxy for rationing on the injection side (y_t^+ model) or sterilization side (y_t^- model).
Interbank Volume	Interbank Volume Stress (Dummy)	Indicator equal to 1 on days with abnormally high interbank transaction volume.
	Interbank Volume Freeze (Dummy)	Indicator equal to 1 on days with abnormally low interbank transaction volume.

(*) The BCRP Operations block is direction-specific: the positive-deviation model includes injection-side auction variables, while the negative-deviation model includes sterilization-side auction variables, reflecting the operational tool the BCRP would use under each type of pressure.

(**) The predictive models (logit and ANN) were estimated with a wider set of regressors; some were excluded from the average marginal effects table due to severe collinearity or high correlation with a retained regressor, which does not affect predictive performance but does impair the identification of individual coefficients.

5. Empirical Strategy


5.1. Variable Interaction Analysis

Before estimating formal models, we examine the bivariate relationships between the explanatory variables and the outcome to verify that the hypothesized signs hold in the data and to detect potential nonlinearities.

Initial liquidity and deviations

Sorting observations by quintiles of the initial current account balance reveals a clear negative relationship: lower initial liquidity is associated with larger positive afternoon deviations, while abundant liquidity is associated with near-zero or negative deviations. This supports the hypothesis that aggregate liquidity shortfalls translate into upward rate pressure.

Figure 1: Initial Liquidity and Afternoon Rate Deviation by Volatility Regime




The relationship between initial liquidity and afternoon deviations is substantially stronger when morning interbank rate volatility is high. Under low-volatility regimes, the level of initial liquidity has little marginal effect. Under high-volatility regimes, low initial liquidity is associated with much larger deviations. This pattern motivates the inclusion of both variables.

Reserve requirement heterogeneity

The highest quintile of cross-bank reserve requirement dispersion is associated with the largest average afternoon deviation (4.8 bps), consistent with the idea that liquidity fragmentation creates market frictions. Even if aggregate liquidity is sufficient, a heterogeneous distribution may leave some banks in acute need of funding, generating upward pressure on the interbank rate. The relationship is not perfectly monotonic at intermediate quintiles, suggesting a possible nonlinear pattern.


Figure 2: Reserve Requirement Dispersion and Afternoon Rate Deviations



Repo operations:

Days with repo operations show positive average deviations even when demand was fully satisfied, suggesting that the presence of repos is itself a signal of pre-existing liquidity stress. The most striking pattern concerns *RepoUnmet*: the highest quintile of unsatisfied demand is associated with an average deviation of approximately 13.1 bps. Importantly, this should not be interpreted as a causal effect; the BCRP conducts repos precisely when conditions are already tight, so repo variables are best understood as signals of liquidity stress rather than independent drivers of deviations.

Figure 3: Repo Operations, Unmet Demand and Afternoon Rate Deviations



5.2. Logit Benchmark Model

We model the two outcomes —upward deviations and downward deviations— through two independently estimated logistic regressions that share a common general specification:

$$Pr(y_t = 1 | X_t) = \Lambda(\alpha + \beta' M_t + \gamma' A_t + \delta' L_t + \theta' S_t + \phi' V_t)$$

where $\Lambda(\cdot)$ is the logistic function and the five blocks correspond to the variables defined in Table 4 (Section 4.2):

- M_t : Morning Market signals (Morning Deviation, Morning Intraday Volatility).
- A_t : Autonomous Liquidity factors (Projected BN injection, Projected BN Absorption).
- L_t : Liquidity and structural conditions (Initial Balance, HHI Concentration Index, Reserve Deviation lag 1).
- S_t : BCRP operations (Injections and Sterilizations variables).
- V_t : Interbank volume stress indicators (Interbank Volume Stress, Interbank Volume Freeze).

For prediction, each logit is estimated with L2 regularization. The regularization strength and the relative weight assigned to deviation events are selected via walk-forward cross-validation over the development sample described in Section 5.5, choosing the combination that maximizes the F2-score, a precision-recall metric that weights recall four times more heavily than precision, consistent with the operational cost of a missed deviation relative to a false alarm. The decision threshold is subsequently calibrated using out-of-fold predicted probabilities from the same development sample, and applied, without further adjustment, to the held-out test period.

For economic interpretation, each specification is re-estimated without regularization on the reduced regressor set of Table 4, using heteroskedasticity and autocorrelation-consistent (HAC, Newey-West) standard errors to account for the serial correlation expected in daily financial data. We interpret each specification through average marginal effects (AME): for a continuous regressor, the AME measures the average change in the predicted probability of the event for a one-unit increase in that regressor, holding the remaining covariates fixed at their observed values; for dummy regressors, it measures the average discrete change in probability when the indicator switches from 0 to 1.

Unlike the predictive specification, the AME specification does not reweight observations by event rarity, since the marginal effects should reflect the population-average relationship implied by the correctly specified likelihood, rather than a decision-oriented reweighting. This interpretation device is used throughout Section 6 to discuss the economic content of each coefficient.

5.3. Neural Network Model

To capture potential nonlinear relationships between liquidity conditions, volatility, and interbank rate deviations, we also estimate an artificial neural network (ANN). The architecture consists of an input layer, a first dense hidden layer of 12 neurons with ReLU activation, a dropout layer (rate = 0.10), a second dense hidden layer of 6 neurons with ReLU activation, and an output layer of 1 neuron with sigmoid activation. The model is trained using binary cross-entropy loss with the Adam optimizer.

As with the logit, the relative weight assigned to deviation events and an L2 weight-regularization coefficient are selected via walk-forward cross-validation within the training portion of the development sample, to reduce the risk of overfitting given the limited number of deviation events available. Early stopping, monitored on a held-out validation slice, is used to determine the training horizon. The decision threshold is calibrated on the same validation slice to maximize the F2-score, mirroring the criterion used for the logit. The use of a neural network is motivated by the descriptive evidence in Section 5.1, which suggests that the relationship between liquidity conditions and deviations may be nonlinear and regime dependent.

5.4. Evaluation Metrics

Given that the event of interest occurs in a small share of observations (approximately 4% for upward deviations and 17% for downward deviations in the respective test periods), accuracy alone is not informative. We focus on the following metrics:

- ROC-AUC: measures the ability of the model to rank days with deviations above days without deviations. Values near 1 indicate near-perfect discrimination.
- AUPRC (Average Precision): more appropriate than ROC-AUC when the positive class is rare, since it penalizes false alarms directly relative to the (small) number of alerts issued, rather than relative to the (large) number of non-events.
- F2-score: the weighted harmonic mean of precision and recall, with recall weighted four times more heavily than precision. We adopt the F2-score, rather than the F1-score, as our primary threshold-selection criterion, reflecting the operational view that a missed deviation (false negative) is materially more costly to a trading desk than an additional false alarm.
- Recall for class 1: the share of actual deviation days correctly flagged by the model. Operationally, this is the most important metric since missed deviations (false negatives) have direct consequences for liquidity management.
- Precision for class 1: the share of model alerts that correspond to actual deviations. Relevant for avoiding alert fatigue.

To assess whether the observed performance gap between the logit and the ANN is statistically meaningful rather than a feature of a single test sample, we additionally report, for each direction of deviation, a DeLong test for the paired difference in ROC-AUC, and a paired bootstrap procedure (5,000 resamples) for the difference in AUPRC, for which no closed-form test is available. Both tests exploit the fact that the two models are evaluated on the identical set of test-period observations.

5.5. Data Splits

Both models are evaluated out-of-sample using a chronological holdout rather than a random partition of the data, replicating the information a trading desk would have available in real time. For each direction of deviation, we distinguish between a development sample, used exclusively for model calibration, and a test sample, held out entirely until the final evaluation reported in Section 6.

Within the development sample, hyperparameters (regularization strength and event weighting for the logit; event weighting and L2 penalty for the ANN) are selected via walk-forward cross-validation: the development sample is split into a sequence of chronologically ordered folds, each using only past observations to validate on a subsequent period, and candidate hyperparameters are scored by their average out-of-fold performance. This procedure avoids relying on a single, potentially unrepresentative validation window—a meaningful concern in our setting, since deviation events are unevenly distributed over time—while still respecting the chronological ordering of the data at every step. The final decision threshold is likewise calibrated using out-of-fold predicted probabilities pooled across the development sample, rather than a single held-out year, before being applied unchanged to the test period.

The two specifications use different development and test windows, reflecting the periods in which each type of pressure has been economically most relevant. The positive-deviation model is developed on data from 2015 to 2021 (1,585 observations, 176-177 deviation events) and evaluated out-of-sample over 2022-2025 (980 observations, 36 events), a period that captures the 2022 global inflation episode together with the subsequent lower-frequency years. The negative-

deviation model is developed on data from 2015 to 2023 (2,078 observations, 169 events) and evaluated over 2024-2025 (487 observations, 85 events), corresponding to the recent easing cycle in which excess-liquidity conditions have been most prevalent.

6. Results

This section presents the empirical results for the two early-warning models: the positive-deviation outcome (y_t^+), flagging afternoon interbank rates closing more than 5 bps above the policy rate (Section 6.1), and the negative-deviation outcome (y_t^-), flagging rates closing more than 5 bps below it (Section 6.2). For each outcome we report, in turn, the logit benchmark and the neural network, before closing with a formal statistical comparison of the two. Section 6.3 discusses the asymmetries that emerge between the two directions of deviation.

6.1. Positive Deviations

6.1.1. Logit Benchmark


Out-of-sample performance over the 2022-2025 test window is strong for a rare event: the model achieves a ROC-AUC of 0.95 and an AUPRC of 0.49, far above the 0.04 baseline implied by the roughly 4% prevalence of the event in the test sample.

Table 5. Out-of-sample performance – Logit, Positive Deviations (Test 2022-2025)

Metric	Positive, Logit
Observations	980
Positive deviation events	36
Decision threshold	0.69
ROC-AUC	0.95
AUPRC	0.49
Baseline AUPRC	0.04
Precision, class 1	0.41
Recall, class 1	0.64
F1-score, class 1	0.50
F2-score, class 1	0.58
Accuracy	0.95

Notes: Class 1 corresponds to positive afternoon deviations greater than 5 basis points.
The F2-score gives more weight to recall than precision. Baseline AUPRC equals event prevalence.

Figure 4. ROC and Precision-Recall curves - Logit, Positive Deviations (test 2022-2025)



The logit is interpreted through average marginal effects (AME), computed on the 2015-2021 development sample (Table 6). The results are broadly consistent with the economic mechanism

we expect to drive upward pressure but also reveal a more nuanced role for market structure than a simple market-power story would predict.

Table 6. Average Marginal Effects – Logit, Positive Deviations

Explanatory Variables	M1 (Positive Logit, HAC SE)
Morning Deviation	0.1125*** (4.982)
Morning Intraday Volatility	0.0262*** (3.894)
Projected BN Absorption	-0.0379*** (-2.800)
Projected BN Injection	-0.0013 (-0.354)
Initial Balance	-0.0750*** (-2.660)
HHI Concentration Index	-0.0197** (-2.443)
Reserve Deviation (Lag 1)	-0.0418** (-2.401)
Injection: Unmet Demand Share	0.0071 (1.142)
Injection Amount Awarded (% of Initial Liquidity)	0.0128** (1.992)
Interbank Volume Stress (Dummy)	0.0068 (0.387)
Interbank Volume Freeze (Dummy)	0.0174 (0.832)
Observations (Development Set 2015–2021)	1585
Pseudo R-squared	0.5509
HAC maxlags	7

Notes: HAC (Newey–West) z-statistics reported in parentheses below the AME.

* p<0.10; ** p<0.05; *** p<0.01

Consistent with expectations, a larger morning deviation and higher morning intraday volatility both raise the probability of an afternoon upward deviation, confirming that early market signals carry forward into the closing window. Initial balance is strongly negative, confirming that a more liquid system at the start of the day is less prone to afternoon scarcity. The HHI Concentration Index and Reserve Deviation (Lag 1) are both negative and significant, a result that deserves particular attention. A purely market-power interpretation of concentration (Acharya et al., 2012) would predict that a more concentrated system makes scarcity episodes more likely, as dominant lenders extract rents from liquidity-short banks.

Our results point in the opposite direction for scarcity episodes: when liquidity is concentrated among a few large institutions, the probability of an afternoon shortfall falls, not rises. We interpret this as evidence of an intermediation, rather than an extraction, channel. Large banks holding a disproportionate share of system liquidity are natural counterparties for banks in temporary need of funds, and in a concentrated system there are, by construction, fewer of these natural liquidity providers, making them easier for a constrained bank to locate and negotiate with.

The negative coefficient on Reserve Deviation (Lag 1) —a measure of how unevenly banks have advanced toward their reserve requirement compliance— is consistent with the same underlying mechanism: on days when some banks are further ahead in their reserve compliance than others, those banks have spare capacity to lend intraday, which likewise reduces the probability of a system-wide shortfall. In the search-based framework of Bech and Monnet (2016), this is naturally read as a reduction in search friction: whether the source of heterogeneity is structural (bank size) or transitory (the daily state of reserve compliance), a market with identifiable

counterparties holding spare liquidity resolves shortfalls more efficiently than one in which liquidity needs and surpluses are evenly, and therefore anonymously, spread across many similarly sized banks.

Projected BN Absorption is negative and significant, while Projected BN Injection is not. This asymmetry is best understood in the light of the institutional channel through which this variable becomes available. *Banco de la Nación* reports its projected net operations with commercial banks to the BCRP at midday; this information is not observed by commercial banks themselves, who see only the outcome of BCRP's own subsequent operations. This suggests that the negative coefficient on Projected BN Absorption reflects a BCRP-side anticipatory response, the Bank adjusting its own operations in light of a privileged signal of an incoming liquidity withdrawal, rather than a market-wide reaction, since commercial banks have no direct visibility into this variable.

A projected injection, being unambiguously liquidity-easing news, does not require the same corrective adjustment, consistent with its coefficient being statistically indistinguishable from zero. The positive coefficient on Injection Amount Awarded should be read in the same spirit: larger injection auctions tend to coincide with days the BCRP has already identified, through this same privileged channel, as liquidity-tight, so the variable captures the magnitude of the imbalance the Bank is responding to, rather than a causal effect of the injection itself. Because the interbank market is already under pressure before the BCRP's own operations take place, our design cannot speak to the counterfactual of how much larger these deviations would have been absent the Bank's intervention, a natural question for future research.

6.1.2 Neural Network


The ANN is calibrated following the procedure described in Sections 5.3 and 5.5, with hyperparameters chosen via walk-forward cross-validation and the decision threshold set to maximize the F2-score on the validation slice of the development sample.

Table 7. Out-of-sample performance – Neural Network, Positive Deviations (Test 2022-2025)

Metric	Positive, ANN
Observations	980
Positive deviation events	36
Decision threshold	0.11
ROC-AUC	0.88
AUPRC	0.44
Baseline AUPRC	0.04
Precision, class 1	0.29
Recall, class 1	0.61
F1-score, class 1	0.4
F2-score, class 1	0.50
Accuracy	0.93
Notes: Class 1 corresponds to positive afternoon deviations greater than 5 basis points. The F2-score gives more weight to recall than precision. Baseline AUPRC equals event prevalence.	

The neural network attains lower ROC-AUC and AUPRC than the logit benchmark reported in Table 5. This runs against the common prior that a nonlinear model should weakly dominate a linear one when the underlying relationship is regime-dependent, and is examined formally in Section 6.1.3.

Figure 5. ROC and Precision-Recall curves – Neural Network, Positive Deviations (test 2022-2025)



6.1.2. Statistical Comparison

Table 8. Logit vs. Neural Network – Paired Comparison, Positive Deviations

Metric	Logit	ANN	Difference	Test	p-value
ROC-AUC	0.946	0.882	+0.064	DeLong	0.035
AUPRC	0.490	0.445	+0.045	Paired bootstrap	0.336
AUPRC	0.490	0.445	+0.045	Paired permutation	0.710

Notes: Both models are evaluated on the identical test set (2022-2025, 980 observations, 36 events), which validates the use of paired tests.

The DeLong test evaluates the null hypothesis of equal ROC-AUC; the bootstrap and permutation tests (5,000 resamples) evaluate the null hypothesis of equal AUPRC.

The logit's advantage in ranking ability (ROC-AUC) is statistically significant at the 5% level. The advantage in AUPRC, while present, does not reach conventional significance under either test, consistent with the limited statistical power available with only 36 deviation events in the test period.

6.2. Negative Deviations

6.2.1. Logit Benchmark

Out-of-sample performance over the 2024-2025 test window remains strong: the model achieves a ROC-AUC of 0.77 and an AUPRC of 0.43, well above the 0.17 baseline implied by the roughly 17% prevalence of the event – a much more frequent event than upward deviations, consistent with the recent monetary easing cycle.

Table 9. Out-of-sample performance – Logit, Negative Deviations (Test 2024-2025)


Metric	Negative, Logit
Observations	487
Positive deviation events	85
Decision threshold	0.45
ROC-AUC	0.77
AUPRC	0.43
Baseline AUPRC	0.17
Precision, class 1	0.30
Recall, class 1	0.87
F1-score, class 1	0.45
F2-score, class 1	0.63

Notes: Class 1 corresponds to positive afternoon deviations more than 5 basis points below the policy rate.

The F2-score gives more weight to recall than precision. Baseline AUPRC equals event prevalence.

Consistent with the F2-oriented calibration, the model attains high recall (0.87), correctly flagging the large majority of excess-liquidity episodes at the cost of a comparatively high false-alarm rate, most visible in 2025.

Figure 6. ROC and Precision-Recall curves – Logit, Negative Deviations (test 2024-2025)



The logit is interpreted through average marginal effects, computed on the 2015-2023 development sample (Table 10).

Table 10. Average Marginal Effects – Logit, Negative Deviations

Explanatory Variables	M1 (Negative Logit, HAC SE)
Morning Deviation	-0.1014*** (-2.778)
Morning Intraday Volatility	0.0303*** (2.695)
Projected BN Absorption	-0.0252 (-1.208)
Projected BN Injection	-0.0369 (-1.506)
Initial Balance	0.0192** (2.576)
HHI Concentration Index	0.0185*** (3.365)
Reserve Deviation (Lag 1)	0.0060* (1.795)
Sterilization: Unmet Demand Share	0.0198*** (2.681)
Sterilization Amount Awarded (% of Initial Liquidity)	0.0133*** (3.196)
Interbank Volume Stress (Dummy)	0.0330* (1.898)
Interbank Volume Freeze (Dummy)	0.0399* (1.870)
Observations (Development Set 2015–2023)	2078
Pseudo R-squared	0.1726
HAC maxlags	7

Notes: HAC (Newey–West) z-statistics reported in parentheses below the AME.

* p<0.10; ** p<0.05; *** p<0.01

The most striking result is that the HHI Concentration Index and Reserve Deviation (Lag 1) — both negative and significant in the positive-deviation model (Table 6)— are positive here. Initial balance is also positive, as expected: a more liquid system is, mechanically, more likely to see the afternoon rate settle below the policy rate. The sign reversal on concentration and reserve dispersion is economically meaningful rather than a statistical artifact; we discuss this pattern jointly with its positive-model counterpart in Section 6.3, where we show that both results are consistent with a single underlying mechanism rather than two unrelated effects.

The BCRP intervention variables display a complementary asymmetry to the positive-deviation model. Neither Projected BN absorption nor Projected BN injection is significant here, whereas BN absorption was strongly significant for scarcity episodes; this is consistent with an anticipatory response that is specific to scarcity threats, since a projected liquidity injection carries no urgency for pre-emptive correction.

By contrast, both sterilization-related variables —unmet demand share and amount awarded— are positive and statistically significant. This result should not be interpreted as evidence that BCRP sterilization is ineffective. Rather, because afternoon operations are implemented after liquidity imbalances have already materialized and are sized in response to the conditions observed by the Bank on the same day, these variables are better understood as indicators of the severity of the underlying excess-liquidity episode.

In other words, larger unmet demand or larger awarded amounts reflect days with deeper liquidity pressures, which are also more likely to exhibit afternoon rate deviations. Assessing whether, and by how much, BCRP intervention mitigates those deviations relative to a no-intervention counterfactual would require a causal design beyond the scope of our predictive framework and is left for future research.

6.2.2. Neural Network

Table 11. Out-of-sample performance – Neural Network, Negative Deviations (Test 2024-2025)


Metric	Negative, Logit
Observations	487
Positive deviation events	85
Decision threshold	0.14
ROC-AUC	0.75
AUPRC	0.41
Baseline AUPRC	0.17
Precision, class 1	0.27
Recall, class 1	0.81
F1-score, class 1	0.41
F2-score, class 1	0.58

Notes: Class 1 corresponds to positive afternoon deviations more than 5 basis points below the policy rate.

The F2-score gives more weight to recall than precision. Baseline AUPRC equals event prevalence.

As with positive-deviation model, the ANN underperforms the logit on both ranking metrics, reinforcing the pattern examined statistically below.

Figure 7. ROC and Precision-Recall curves – Neural Network, Negative Deviations (test 2024-2025)



6.2.3. Statistical Comparison

As in the positive-deviation case, the logit's ranking advantage is statistically significant, though the margin is smaller in absolute terms; the larger number of test-period events (85 vs. 36) gives the test enough statistical power to detect this smaller difference. The AUPRC advantage is again not statistically significant under either test.

Table 12. Logit vs. Neural Network – Paired Comparison, Negative Deviations

Metric	Logit	ANN	Difference	Test	p-value
ROC-AUC	0.775	0.749	+0.026	DeLong	0.046
AUPRC	0.429	0.407	+0.021	Paired bootstrap	0.112
AUPRC	0.429	0.407	+0.021	Paired permutation	0.752

Notes: Both models are evaluated on the identical test set (2024-2025, 487 observations, 85 events).

6.3. Discussion: Asymmetries Between Scarcity and Excess

Two robust patterns emerge from comparing the positive- and negative-deviation results. First, the logit outperforms the ANN in ranking ability for both directions of deviation, with statistically significant differences in ROC-AUC in both cases (Tables 8 and 12). At first glance, this result may seem counterintuitive. A neural network's main theoretical advantage over a linear model is its ability to capture nonlinear relationships, and Section 5.1 documents regime-dependent patterns in the underlying liquidity variables that could, in principle, favor a more flexible model.

Table 13. Summary of Out-of-sample Performance, All four models

	Logit, Positive	ANN, Positive	Logit, Negative	ANN, Negative
ROC-AUC	0.95	0.88	0.78	0.75
AUPRC	0.49	0.45	0.43	0.41
F2-score	0.58	0.50	0.63	0.58
Precision	0.41	0.29	0.30	0.27
Recall	0.64	0.61	0.87	0.81

In this setting, however, that theoretical advantage does not translate into better out-of-sample performance. The main reason is that both outcomes are rare events: the test periods contain only 36 and 85 deviations, respectively. With so few positive cases, reliably estimating a flexible nonlinear decision boundary is difficult. The ANN's additional flexibility may reduce bias, but it also substantially increases variance. In our application, this bias-variance trade-off appears unfavourable: the more flexible model fits the limited development sample less robustly than the simpler logit.

A second reason is related to regularization. The L2 penalty needed to prevent the ANN from overfitting also reduces the extent to which the network can exploit its nonlinear structure. By shrinking large weights, regularization makes the mapping learned by the ANN smoother and closer to a linear approximation. This penalty was necessary, since an unregularized network trained on such few events would likely capture idiosyncratic patterns in the development period rather than generalizable relationships. However, it also weakens precisely the flexibility that would otherwise distinguish the ANN from the logit.

Taken together, these results suggest that, in a rare-event setting with limited positive observations, a well-specified linear model can generalize more reliably than a more flexible alternative. The ANN's structural flexibility is therefore not sufficient, by itself, to overcome the additional estimation variance introduced by the model.

The second pattern is more central to the economic interpretation of the paper, market concentration (HHI) and cross-bank reserve requirement dispersion display opposite signs across the positive- and negative-deviation models. Rather than interpreting this sign reversal as two unrelated effects, we read it as evidence of a common underlying mechanism. In both specifications, the variables point to the same directional pressure on the interbank rate: higher market concentration and greater dispersion in reserve requirement compliance are associated with a lower likelihood of upward deviations and a higher likelihood of downward deviations.

This pattern suggests that both variables capture the presence of banks with spare liquidity capacity to supply to the market. In the case of HHI, this capacity comes from banks that are structurally large relative to the rest of the system. In the case of Reserve Deviation, it comes from banks that are further ahead in their reserve requirement compliance on a given day. These banks appear able and willing to place liquidity generously. When the system is short, their supply helps prevent the rate from spiking; when the system is already long on liquidity, the same supply can push the rate below the policy rate. Thus, concentration and reserve dispersion are not associated with offsetting forces, but with a persistent downward pressure on the equilibrium interbank rate.

This mechanism is particularly relevant in the Peruvian interbank market, which operates over the counter and through bilateral negotiation. In line with the search-based framework of Bech and Monnet (2016), the absence of a centralized order book means that transaction rates depend on which counterparty a bank finds and on the bargaining position in that specific match. A smaller number of identifiable liquidity suppliers can therefore shift bargaining power in a way that reinforces downward pressure on the equilibrium rate. When the market is long, borrowers become the scarce side of the negotiation, allowing them to secure funds on increasingly favourable terms as surplus liquidity concentrates among fewer suppliers; when it is short, the same concentrated suppliers' visibility and ready availability as counterparties reduces search frictions and prevents rate spikes. Both channels operate through the same underlying force – bargaining leverage shifting toward whichever side of the market is scarce – and reinforce downward pressure on the equilibrium rate across both regimes.

This tendency toward downward pressure may be further reinforced by the asymmetric structure of the BCRP's standing facilities corridor: the deposit facility floor sits 200 basis points below the policy rate, while the lending facility ceiling sits only 50 basis points above it. This asymmetry creates unequal outside options for the two sides of the market – a bank with surplus liquidity that fails to find an interbank counterparty faces a considerably harsher penalty than a bank with a liquidity need, giving surplus banks a stronger incentive to compete for placement in the interbank market than deficit banks have to compete for funds. This institutional feature is consistent with the markedly higher prevalence of downward deviations relative to upward deviations in our sample (17% vs. 3.67% of test-period observations) and may help explain why the bargaining-power mechanism described above manifests more forcefully when the market is long on liquidity than when it is short. It is worth noting that this interpretation reflects our reading of the observed pattern rather than a mechanism directly tested by our design: the result could also reflect market microstructure factors or bank-specific bilateral dynamics not explicitly captured by our variables, which merit further investigation.

This asymmetry also extends to variables capturing liquidity management by Banco de la Nación (BN) and by the BCRP's own intervention operations. Projected BN Absorption – net withdrawals reported by BN, whose aggregate operations with commercial banks are reported to the BCRP at midday – is significant in the scarcity model, but with a negative sign: greater projected absorption is associated with a lower, not higher, probability of scarcity deviation. This is consistent with an anticipatory, corrective response by the BCRP ahead of its own afternoon operations, which offsets the withdrawal before it materializes in the afternoon rate. Projected BN Injection is not significant in either model, consistent with the absence of urgency when the projected flow is liquidity-easing. The BCRP's own intervention variables capture a different stage of the same process. Injection Amount Awarded reflects an ex-post outcome: the size of the auction needed to address whatever residual imbalance persists into the early afternoon, regardless of its source. Because injection auctions are not necessarily driven by BN flows a larger auction is better read as a symptom of the overall severity of that day's imbalance than as evidence that the BCRP's anticipatory response to BN flows specifically was insufficient. This distinguishes it from Projected BN Absorption, which reflects a specific, ex-ante identifiable threat rather than

the cumulative imbalance the Bank ultimately has to address. Unmet demand share, by contrast, is significant only for excess-liquidity episodes, suggesting that sterilization auctions more frequently fall short of fully absorbing the day's surplus than injection auctions fall short of meeting scarcity-driven demand. In all cases, our design characterizes the interbank market as one that is already under pressure by the time these operations take place, so the coefficients should not be read as evidence on the effectiveness of BCRP intervention; we return to this point, and to how it might be addressed in future work, in Section 7.

7. Conclusion

This paper develops and evaluates an intraday early-warning framework to anticipate, before the BCRP's afternoon intervention window closes, whether the overnight interbank rate is likely to deviate materially from the policy rate later that day. Using a unique high-frequency dataset for 2015–2025, we show that a small set of morning indicators carries meaningful predictive content for afternoon rate pressures. These indicators include the morning rate spread and its volatility, the initial liquidity position of the banking system, the concentration and dispersion of liquidity across banks, and signals from the Bank's own open market operations.

Across both upward and downward deviations, a regularized logistic regression outperforms a more flexible neural network out-of-sample, with statistically significant gains in ROC-AUC in both cases. This is useful from an operational perspective. The simpler model is not only more accurate in ranking high-risk days, but also more transparent, easier to interpret, and therefore more suitable for use by a money market operations desk.

The paper's central economic finding is that market concentration and cross-bank heterogeneity in reserve requirement compliance have direction-dependent implications for rate stability. Both variables reduce the probability of scarcity episodes, consistent with large or well-positioned banks acting as readily identifiable suppliers of liquidity that reduce search frictions in a bilateral market. At the same time, both variables increase the probability of excess-liquidity episodes, consistent with a shift in bargaining power toward borrowers, who become the scarce side of the market when liquidity is abundant. This asymmetry reconciles two mechanisms that may otherwise appear to be in tension: a search-friction channel, through which concentrated liquidity supply helps prevent rate spikes, and a bargaining-power channel, through which the same structure can place downward pressure on rates when liquidity is abundant.

This result is relevant for the BCRP's assessment of interbank market conditions. It suggests that the same structural features of the market can have stabilizing or destabilizing effects depending on the direction of the underlying liquidity imbalance. It also highlights the value of monitoring cross-bank reserve requirement dispersion in real time, since this variable appears to contain information beyond aggregate liquidity balances and market concentration alone.

The paper also has several limitations that point to natural directions for future research. First, the framework is predictive rather than causal. While the results are consistent with BCRP intervention variables acting as contemporaneous indicators of liquidity stress, the empirical design cannot identify how large afternoon deviations would have been in the absence of intervention. Since the market appears to be already imbalanced by the time these operations take place, it is plausible that BCRP intervention attenuates the size of the resulting deviation. Establishing this formally would require a causal design capable of isolating the effect of intervention, for example by exploiting institutional or timing discontinuities in how and when the Bank intervenes.

Second, the analysis is based on symmetric 5-basis-point thresholds. Extending the framework to predict the severity of deviations, rather than only their occurrence, could provide additional

operational value. A related extension would be to model the persistence of stress episodes across multiple days.

Third, the reduced-form approach identifies the predictive role of market concentration and reserve dispersion, but does not explicitly model the bilateral, over-the-counter matching process through which these structural features translate into equilibrium rates. A structural search-and-bargaining model in the tradition of Bech and Monnet (2016), calibrated to the Peruvian interbank market, could formalize the mechanism documented here and generate sharper quantitative predictions about how changes in market structure or reserve requirement design would affect rate stability.

8. References

- Acharya, V. V., & Merrouche, O. (2013). Precautionary hoarding of liquidity and interbank markets: Evidence from the subprime crisis. *Review of finance*, 17(1), 107-160.
- Acharya, V. V., Gromb, D., & Yorulmazer, T. (2012). Imperfect competition in the interbank market for liquidity as a rationale for central banking. *American Economic Journal: Macroeconomics*, 4(2), 184-217.
- Baglioni, A., & Monticini, A. (2008). The intraday price of money: Evidence from the e-MID interbank market. *Journal of Money, Credit and Banking*, 40(7), 1533-1540.
- Bech, M. L., & Garratt, R. (2003). The intraday liquidity management game. *Journal of economic theory*, 109(2), 198-219.
- Bech, M., & Monnet, C. (2016). A search-based model of the interbank money market and monetary policy implementation. *Journal of Economic Theory*, 164, 32-67.
- Furfine, C. (2002). The interbank market during a crisis. *European Economic Review*, 46(4-5), 809-820.
- Garcia, P., Pozo, J., & Velarde, R. (2025). Policy interest rate expectations and the behavior of the interbank money market. *Documento de Trabajo*, (019-2025).
- Nautz, D., & Offermanns, C. J. (2007). The dynamic relationship between the euro overnight rate, the ECB's policy rate and the term spread. *International Journal of Finance & Economics*, 12(3), 287-300.
- Pozo, J., & Lupú, J. (2011, octubre). Dinámica del Mercado Interbancario en el Perú [Presentación de ponencia]. XXIX Encuentro de Economistas, Lima, Perú.
- Sichei, M. M., Tiriongo, S., & Shimba, C. (2012). *Segmentation and efficiency of the interbank market in Kenya* (No. 1). KBA Centre for Research on Financial Markets and Policy Working Paper Series.
- Velarde, R. (2025). Determinantes de presiones en el mercado interbancario. *Revista Moneda*, (204), 9-15.
- Whitesell, W. (2006). Interest rate corridors and reserves. *Journal of Monetary economics*, 53(6), 1177-1195.