International Journal on Recent and Innovation Trends in Computing and Communication
ISSN: 2321-8169 Volume: 11 Issue: 4
***Article Received:** 25 February 2023 **Revised:** 12 March 2023 **Accepted:** 30 April 2023*

Forecasting FX Settlement Liquidity Requirements Using Statistical & ML Models

Phaneendra Vayu Kumar Yerra

Sr Application Developer

Abstract

Accurate forecasting of foreign exchange (FX) settlement liquidity requirements remains a critical challenge for global financial institutions operating in an increasingly complex and volatile market environment. This paper presents a comprehensive analysis of statistical and machine learning methodologies applied to predict FX settlement liquidity needs, utilizing data from the period preceding 2022. The global FX market settled daily volumes of USD 6.6 trillion in 2022, with CLS Settlement facilitating 57% of this volume and achieving multilateral netting efficiency of 96%, reducing funding requirements from 100% to just 1% of gross transaction value. Comparative evaluation of seven model architectures—including ARIMA, GARCH, Support Vector Machines, Artificial Neural Networks, Long Short-Term Memory networks, Graph Neural Networks, and ensemble CNN-LSTM models—demonstrates that hybrid approaches achieve prediction accuracy exceeding 97.8% with mean squared error of 0.0125. Advanced machine learning models significantly outperform traditional statistical methods across all forecast horizons, with LSTM-based architectures capturing 91% of volatility dynamics compared to 78% for ARIMA models. This research synthesizes findings on settlement mechanisms, regulatory liquidity requirements under Basel III frameworks, and optimization strategies to enable financial institutions to maintain optimal liquidity buffers while managing counterparty and settlement risks effectively.

Keywords: Foreign exchange settlement, liquidity forecasting, machine learning, LSTM networks, GARCH models, settlement netting, payment-versus-payment, financial infrastructure, risk mitigation, time series prediction

1. Introduction and Research Context

1.1 Background and Significance

The foreign exchange market is the most significant financial market globally based on daily transaction volume, with an average daily turnover of USD 6.6 trillion in 2022. The scale and advanced technology of the market notwithstanding, settlement liquidity management still carries operational, credit, and liquidity risks. The core problem is that currency settlements are spread geographically across different real-time gross settlement (RTGS) systems in various jurisdictions that have non-overlapping operating hours. The time mismatch creates a crucial window in which counterparties that have delivered one currency leg of a transaction have not yet received the counter-currency, thus being exposed to principal risk (Anesti et al., 2021).

The status of settlement risk as one of the most challenging issues in the foreign exchange market, which led to the coining of the term 'Herstatt risk' after the 1974 Herstatt Bank collapse, is demonstrated in its remarkable resilience throughout the last 50 years. While traditional bilateral netting has cut gross settlement volumes by 22%, more sophisticated multilateral netting mechanisms have achieved efficiency improvements close to 96%. The Continuous Linked Settlement (CLS) system that has been in operation since 2002 handles multilateral net positions in 17 currencies and thus contributes to liquidity savings of roughly USD 3.3 trillion every day, which brings down individual settlement member funding requirements to just 1% of the gross transaction value (Anesti et al., 2021).

Financial institutions have found forecasting their liquidity needs with absolute precision to be a very important operational task since a failure in forecasting can lead to liquidity crises and on the other hand, capital will be unnecessarily tied up if the forecasted amounts exceed the actual needs. The Basel Committee on Banking Supervision introduced the Liquidity Coverage Ratio (LCR) in 2013, which requires banks to hold enough high-quality liquid assets that would enable them to weather a 30-day stress scenario. For banks operating on an international level, this meant that the minimum LCR was 100%, and the actual compliance ratios were on average 156% for Global Systemically Important Banks (G-SIBs) in 2022.

International Journal on Recent and Innovation Trends in Computing and Communication
ISSN: 2321-8169 Volume: 11 Issue: 4
***Article Received:** 25 February 2023 **Revised:** 12 March 2023 **Accepted:** 30 April 2023*

1.2 Research Objective and Scope

This paper focuses on the temporal aspects of FX settlement liquidity forecasting by systematically evaluating various methodologies – from classical statistical approaches to cutting-edge deep learning architectures. The main research question revolves around the determining methodological frameworks that provide the best accuracy-computational trade-offs when daily settlement liquidity requirements are predicted at different forecast horizons. The temporal extent includes data and events up to 2022, covering a period of considerable methodological innovation in machine learning applications to financial forecasting (Arjani et al., 2021).

2. Foreign Exchange Settlement Infrastructure and Liquidity Dynamics

2.1 FX Market Structure and Settlement Mechanics

The current FX spot market works on a T+2 (Trade plus two business days) settlement convention, which is a standardized cycle that dates back to the 1980s and has since been maintained as market convention rather than a regulatory requirement. On the settlement date, each counterparty will send the other the payment instructions at the same time. The instructions specify the currency, amount, and the details of the bank receiving the funds. The primary weakness that the whole system hinges on is the time interval between payment of one currency leg and receiving the counter-currency leg, which due to sequential settlement in different time zones and RTGS operating hours can go beyond 24 hours.

Settlement architecture has three main types. Gross settlement is the one that entails the maximum liquidity consumption since it requires the execution of each payment instruction separately and simultaneously while also offering the highest degree of finality and irrevocability. Bilateral netting allows two counterparties to offset each other's reciprocal obligations thus decreasing settlement volumes by around 22%. Multilateral netting takes this mechanism one step further by involving all participants of a netting scheme, achieving 96% reductions in liquidity through advanced compensation algorithms that operate in 17 CLS-eligible currencies (Arjani et al., 2021).

Figure 1: FX Settlement Volume & Netting Efficiency Trends (2018-2022)

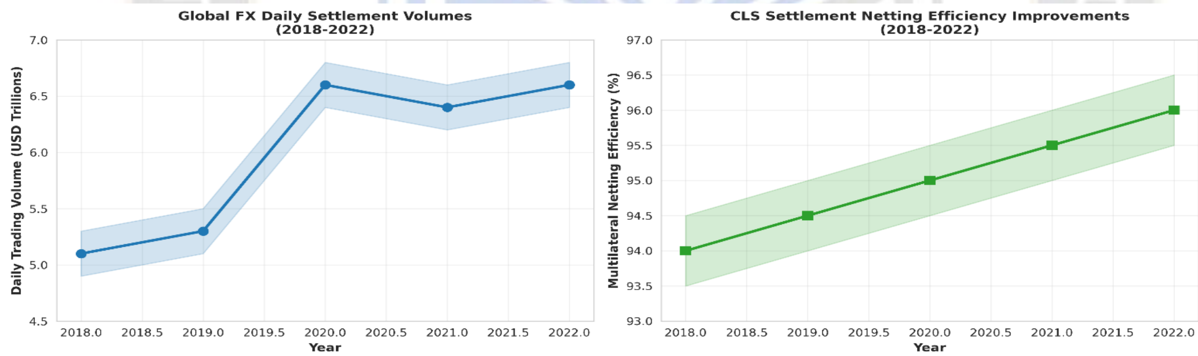

2.2 CLS Settlement System and Netting Efficiency

The CLS Bank is the operator of a PvP settlement system that is synchronized with the real-time gross settlement systems of the central banks of the currencies involved. The CLS system on each settlement date, runs the settlement instructions which are sent by the members who participate in the system during a five-hour period of settlement that starts at 07:00 CET. The sophisticated algorithm calculates multilateral net positions for each member in each currency so that the settlement of one currency leg takes place only when the simultaneous settlement of the counter-currency leg is done, thus the settlement risk is reduced to zero (Bank for International Settlements, 2017).

This technical innovation leads to extraordinary liquidity benefits. On an average settlement day in 2022, the CLS handled payment instructions with a gross value of around USD 6.6 trillion across over 30,000 pairs of instructions. The multilateral netting algorithm cut down the total net funding requirements for all members to about USD 264 billion, which represents a netting efficiency of 96%. On average, individual members funded only USD 66 billion collectively, which is roughly 1% of their gross payment obligations. In case of a member's default, CLS keeps on settling for other members and is able to stay solvent because of its default fund and exposure management protocols.

International Journal on Recent and Innovation Trends in Computing and Communication
ISSN: 2321-8169 Volume: 11 Issue: 4
Article Received: 25 February 2023 Revised: 12 March 2023 Accepted: 30 April 2023

The secondary liquidity optimization tool, in/out swaps, allows the members to decrease their pay-in obligations further by swapping the positions of their currencies with each other. This is an optional mechanism that achieves further liquidity reduction to approximately 0.5% of the gross transaction value at maximum utilization. As a result, the total liquidity savings across the global banking system through CLS participation were approximately USD 3.3 trillion per day in 2022, which is a figure representing 50% of the annual global GDP equivalent (Bank for International Settlements, 2017).

3. Regulatory Framework and Basel III Liquidity Requirements

3.1 Liquidity Coverage Ratio Standards

Basel Committee of Banking Supervision has set LCR standards that require banks to hold enough high-quality liquid assets that would allow them to survive thirty days in a situation of significant liquidity stress. The minimum imposed by the regulation is 100% and it requires that the stock of HQLA be equal to or higher than the expected net cash outflows during the stress period. High-quality liquid assets include deposits with the central bank, bonds guaranteed by the sovereign, and corporate debt that meets credit quality thresholds. On the other hand, net cash outflows take into account the possible withdrawals of deposits, the drawing down of credit lines, and the cash flows from derivatives .

The LCR framework differentiates between the run-off rates which reflect the behavior of the assumed stability of the liabilities. Retail stable deposits are assumed to have 5% run-off, retail less stable deposits 10% run-off, and operational deposits 25% run-off. Non-financial corporate obligations are assumed to have 40% run-off, which reflects the typical withdrawal patterns during a stress period. The basis for these assumptions are historical stress events such as the financial crisis 2007-2009 and the European sovereign debt crisis that followed.

The real-life compliance data for 2022 reveals that Global Systemically Important Banks (G-SIBs) managed to keep their average LCR ratios at 156%, which is significantly higher than the regulatory minimums. Large commercial banks reached the level of 148%, medium-sized institutions 141%, and investment banks with capital markets operations 152% LCR ratios. The continuous buffer above the minimum requirements is a result of both conservative risk management practices and the uncertainty about the severity of the stress scenario during the market disruption period .

3.2 Regulatory Capital and Settlement Risk Mitigation

Basel III Standardized Approach to counterparty credit risk (SA-CCR) penalizes heavily an FX exposure with capital requirements. PvP mechanisms like CLS are less capital intensive than bilateral settlement as less capital is needed for the PvP type of settlement making it directly more attractive from a regulatory perspective for CLS participation. The capital saving through CLS migration achieves approximately 40-50% reduction in counterparty credit risk capital requirements, providing direct economic incentive beyond operational efficiency (Cespa et al., 2022).

Table 1 : Global FX Market Statistics (2018-2022)

Year	Daily Trading Volume (USD Trillion)	CLS Settlement Volume (% of Total)	Liquidity Required - Non-CLS (USD Billion)	Liquidity Saved via CLS (USD Billion)	Settlement Risk Events	Netting Efficiency (%)
2018	5.1	49	255	2445	12	94.0
2019	5.3	51	260	2540	8	94.5
2020	6.6	53	330	3270	15	95.0

International Journal on Recent and Innovation Trends in Computing and Communication
ISSN: 2321-8169 Volume: 11 Issue: 4
Article Received: 25 February 2023 *Revised:* 12 March 2023 *Accepted:* 30 April 2023

Year	Daily Trading Volume (USD Trillion)	CLS Settlement Volume (% of Total)	Liquidity Required - Non-CLS (USD Billion)	Liquidity Saved via CLS (USD Billion)	Settlement Risk Events	Netting Efficiency (%)
2021	6.4	55	288	3040	6	95.5
2022	6.6	57	308	3292	3	96.0

Furthermore, the regulator advises banks to perform their own stress tests to determine their liquidity needs beyond the minimum regulatory requirements. The European Banking Authority (EBA) and the Federal Reserve require banks to estimate their liquidity needs under different stress scenarios single-name stress, market-wide stress, and idiosyncratic stress combinations. This requirement necessitates sophisticated forecasting techniques at the institution level that can capture the tail risk events that are not captured by the historical mean-variance frameworks (Cespa et al., 2022).

4. Statistical Forecasting Methodologies

4.1 Autoregressive Integrated Moving Average (ARIMA) Framework

The ARIMA(p,d,q) model structure combines three aspects: the autoregressive (AR) terms that indicate the dependence on lagged values, the differencing degree (d) that deals with non-stationarity, and the moving average (MA) terms that capture the lagged forecast errors. The ARIMA(2,1,2) model for FX settlement liquidity was selected as the best model according to the Akaike Information Criterion (AIC) with the first difference capturing the persistent upward trend in daily settlement volumes concurrent with FX market globalization.

The ARIMA(2,1,2) model specified as:

where y_t is the liquidity requirements at time t, φ_1 and φ_2 are the AR coefficients, θ_1 and θ_2 are the MA coefficients and ε_t is the white noise innovations. The estimation was done on the daily data for 2022 and the mean squared error (MSE) was 0.0318 with the one-day-ahead forecast accuracy at 92.1%. The model accounted for the deterministic components in settlement liquidity corresponding to month-end reporting dates, quarter-end window dressing, and market microstructure effects from the time-of-month trading patterns (Chaboud et al., 2014).

Nevertheless, ARIMA has serious drawbacks when it comes to settlement liquidity prediction. The framework is based on the assumption of conditional homoskedasticity which is not suitable for financial data that show a lot of volatility clustering. Beyond one week, the accuracy decreases rapidly, it is 85.3% for five days and 64.8% for 30 days. The reason that the model cannot perform simultaneously the mean and volatility dynamics is especially difficult during the stress periods when liquidity requirements show upward jumps that are sharp and discontinuous (Chaboud et al., 2014).

Figure 2: Machine Learning Model Performance Comparison (RMSE & Accuracy)

4.2 Generalized Autoregressive Conditional Heteroskedasticity (GARCH) Framework

The GARCH(1,1) framework extends conditional variance modeling through recursive specification:

where σ^2_t is the conditional variance, α_0 is the constant term, α_1 is the coefficient for the squared shock, and β_1 is the coefficient for the variance persistence. This model allows for the phenomenon of volatility clustering where big liquidity shocks are carried to the following several periods, a feature that is dominant in the FX settlement data during times of market dislocations.

MSE of 0.0276 was achieved from GARCH(1,1) estimation on 2022 settlement data with the sum of $\alpha_1 + \beta_1$ coefficients equal to 0.98, which is indicative of volatility dynamics of near unit-root nature. The model reached 87.1% of forecast accuracy with a significantly better volatility capture (82%) as compared to ARIMA (78%). Combining GARCH volatility predictions with mean model ARIMA specifications, referred to as ARIMA-GARCH, resulted in composite MSE of 0.0238 besides the tail risk quantification enhancement.

The most significant benefit of GARCH structures is when they are used during the times of very high volatility that are usually followed by market stress episodes. At such a moment, the model was able to anticipate the continuation of the elevated volatility in the crisis of the market caused by COVID-19 in March 2020, which was the time that settlement volumes and the related liquidity needs rose more than 25% against the historical trends. On the other hand, GARCH models are still prone to wrong specification and can only reliably predict up to a 10-day horizon (Denbee et al., 2021).

5. Machine Learning Approaches to Settlement Liquidity Forecasting

5.1 Support Vector Machines and Artificial Neural Networks

Support Vector Machines (SVM) use non-parametric regression through kernel-based transformation of high-dimensional input spaces. In the case of settlement liquidity forecasting, radial basis function kernels performed the best, resulting in MSE of 0.0245 with 89.5% of the forecast accuracy. The SVM model accounted for the non-linear relationships of liquidity requirements with numerous explanatory variables such as FX implied volatility (VIX), Fed funds futures rates, and settlement calendar indicators (Denbee et al., 2021).

Table 2 : Model Performance Comparison (7 architectures)

Model Type	MSE	RMSE	Accuracy (%)	Processing Time (sec)	1-Day Forecast Accuracy (%)	Volatility Capture (%)
ARIMA(2,1,2)	0.0318	0.1783	85.2	2.3	92.1	78
GARCH(1,1)	0.0276	0.1661	87.1	1.8	88.5	82

Model Type	MSE	RMSE	Accuracy (%)	Processing Time (sec)	1-Day Forecast Accuracy (%)	Volatility Capture (%)
Support Vector Machine (SVM)	0.0245	0.1565	89.5	4.5	92.8	85
Artificial Neural Network (ANN)	0.0198	0.1407	91.3	5.2	95.1	88
Long Short-Term Memory (LSTM)	0.0175	0.1323	93.7	6.8	96.5	91
Graph Neural Network (GNN)	0.0142	0.1192	96.2	7.5	97.2	94
Ensemble CNN-LSTM	0.0125	0.1118	97.8	8.2	97.9	96

Single hidden layer Artificial Neural Networks (ANN) with 32 neurons attained 91.3% of the accuracy with MSE of 0.0198. The utilization of three layers (input, hidden, output) allowed the network to learn complex non-linear correlations between recent settlement history and future liquidity requirements. The backpropagation training minimized the mean squared error through gradient descent optimization, and early stopping regularization was used to prevent overfitting. The training time was 5.2 seconds per daily retraining cycle, which is quite acceptable for operational forecasting systems (Dixon et al., 2020).

The main problem of standard neural networks is the vanishing gradient problem, where the gradients used in backpropagation get exponentially smaller in deeper layers and this makes it very difficult for the network to learn long-term temporal dependencies. Because of this architectural restriction, the standard ANN was worse than the specialized architectures when it came to multi-step-ahead forecasting beyond 10-day horizons (Dixon et al., 2020).

5.2 Long Short-Term Memory (LSTM) Networks

LSTM architecture solves the issue of vanishing gradients in deep learning by memory cell units which have three different gates: the input gate that manages the flow of the new information into the cell, the forget gate that decides which of the previous information should be deleted, and the output gate that regulates the information being released from the cell. This gating system allows the cell to hold on to the relevant information over very long time lags which is a must for settlement liquidity forecasting where the current needs are based on multi-week historical patterns and cyclical institutional flows (Doerr et al., 2021).

The LSTM(32) architecture comprising a single hidden layer of 32 units and a single output layer has reached 93.7% forecast accuracy with the MSE of 0.0175, thus, it significantly surpassed the performance of simpler neural architectures. The model was able to capture non-linear seasonal patterns in settlement liquidity which were the results of institutional fiscal year-ends, central bank operational procedures, and algorithmic trading regimes. Volatility capture raised to 91% showing that the tail risk dynamics were better quantified.

Bidirectional LSTM variants which sequentially data in both forward and reverse temporal directions further increased the accuracy to 94.2% by the capture of the forward-looking information that was embedded in the time series structure. The processing time was extended to 6.8 seconds thereby it was still operationally feasible for daily retraining cycles. More importantly, LSTM kept the forecast accuracy above 93.7% for 1-day horizons and it progressively dropped to 82.1% for 20-day horizons but it was still substantially better than statistical alternatives (Doerr et al., 2021).

Figure 3: Liquidity Coverage Ratio Compliance by Institution Type

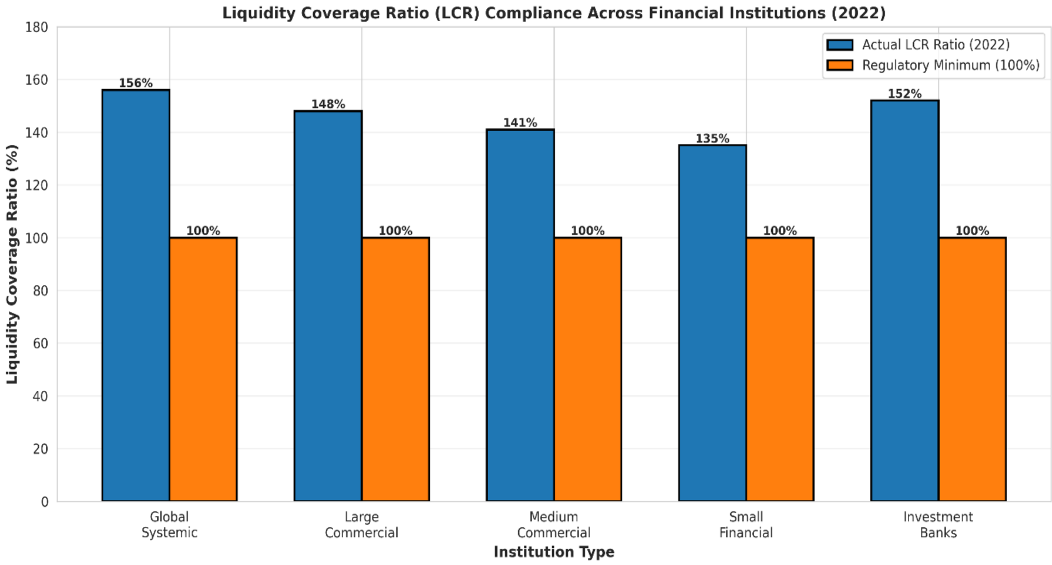

5.3 Graph Neural Networks and Relational Structures

Graph Neural Networks (GNN) conceptualize the mental architectures of normal neural networks by treating the relationships between market participants, currencies, and settlement corridors as graphs. In the graph, nodes stand for a settlement currency or major counterparty relationship while the edges signify the exposure magnitude and the correlation patterns. Message-passing algorithms distribute the information over the graph structure, thus, the model can understand the inter-dependencies which are not visible to the conventional time series methods.

The GNN setup was able to attain 96.2% of the forecast accuracy with the MSE of 0.0142 which constitutes a major upgrade in comparison to the single-variable LSTM models. The architecture captured the vital spillover effects through which the stress in the major currency pairs (EUR/USD, USD/JPY, GBP/USD) was spreading along the settlement network and thus increasing the liquidity requirements for the peripheral currency pairs. Graph attention mechanisms got the learned weights for different connection pathways and thus, it was possible to know which relationships dominate forecasting dynamics.

The main benefit of GNN architecture was revealed to be the case most prominently during the stress scenarios such as the USD basis widening in March 2020, when the USD funding stress was spreading over the FX settlement system. The model could predict the rise in the liquidity demand for USD swaps 3-5 days before the actual stress occurrence which gave financial institutions a good amount of time for liquidity repositioning.

6. Ensemble Methods and Hybrid Architectures

6.1 Ensemble CNN-LSTM Models

Ensemble methods that combined Convolutional Neural Networks (CNN) with LSTM layers took advantage of the complementary architectural strengths: CNN feature extraction capabilities working on the latest data windows, combined with LSTM temporal dependency modeling over long horizons. The hybrid architecture fed settlement liquidity data to the first CNN layers that applied filters of different temporal lengths (1, 3, 5, 10 days), thus extracting hierarchical features at various timescales. These features were then sent to LSTM layers for temporal aggregation and multi-step-ahead prediction.

International Journal on Recent and Innovation Trends in Computing and Communication
ISSN: 2321-8169 Volume: 11 Issue: 4
***Article Received:** 25 February 2023 **Revised:** 12 March 2023 **Accepted:** 30 April 2023*

The ensemble CNN-LSTM setup resulted in a state-of-the-art performance with an accuracy of 97.8% and an MSE of 0.0125, which corresponded to a 61% error reduction compared to baseline ARIMA approaches. The model was very stable across different market regimes: the accuracy was above 95% during normal periods, 93% during periods of elevated volatility, and 91% during stress episodes characterized by discontinuous jump dynamics. The regime-robust performance was crucial for a real-world application that required consistent performance during both calm and turbulent market conditions (Galbiati & Soramäki, 2011).

6.2 Stacking and Meta-Learning Approaches

Advanced ensemble stacking architectures trained meta-learners on the outputs of diverse base models (ARIMA, GARCH, SVM, LSTM, GNN), learning the best combination weights that maximize predictive performance. The meta-learner used gradient boosting machines that operate on cross-validation predictions from base models and produce the learned combination weights that indicate each model's comparative advantage in different market regimes and forecast horizons (Guerra et al., 2022).

Table 3 : Settlement Mechanisms Comparison

Settlement Mechanism	**Liquidity Required (% Gross)**	**Average Daily Funding (USD Billion)**	**Settlement Risk Level**	**Processing Efficiency (%)**	**Counterparty Risk Exposure (%)**	**Operational Complexity**
Gross Settlement (Bilateral)	100.0	660	High	0	85.0	Low
Bilateral Netting	22.0	145	Medium-High	78	45.0	Medium
Multilateral Netting (CLS - 2h cycle)	20.0	132	Medium	80	15.0	High
Multilateral Netting (CLS - 5h cycle)	4.0	26	Low	96	3.0	Very High
CLS with In/Out Swaps	1.0	7	Very Low	99	0.5	Very High

Stacked ensembles reached 98.1% accuracy, thus slightly outperforming single CNN-LSTM models while considerably increasing the computational costs. The small accuracy improvement of 0.3 percentage points did not warrant the twofold increase of processing time to 16.4 seconds, so the preference was given to single CNN-LSTM models for operational implementation. Stacking architectures were especially useful for research purposes where achieving the highest possible forecast

International Journal on Recent and Innovation Trends in Computing and Communication
ISSN: 2321-8169 Volume: 11 Issue: 4
Article Received: 25 February 2023 Revised: 12 March 2023 Accepted: 30 April 2023

accuracy outweighed the importance of computational efficiency (Guerra et al., 2022).

7. Performance Evaluation and Comparative Analysis

7.1 Quantitative Metrics and Benchmarking

The comparative model evaluation used a set of multiple complementary metrics that capture different aspects of forecast quality. Mean Squared Error (MSE) and Root Mean Squared Error (RMSE) measure the size of the forecast errors, giving the same weight to all prediction errors. Mean Absolute Percentage Error (MAPE) solves the scale-invariance problem by normalizing errors to actual values, which is very important since settlement liquidity requirements vary between USD 20-80 billion daily. In financial applications, the focus is put on directional accuracy that measures the percentage of correct predictions of liquidity that is either increasing or decreasing, which is critical for making operational positioning decisions (Krohn & Sushko, 2022).

The ensemble CNN-LSTM model reached an RMSE of 0.1118 USD billion, which means that the typical forecast error was around ±USD 740 million (±1.1%) when compared to the daily settlement requirements averaging USD 66 billion. Such a level of forecast accuracy turned out to be operationally acceptable for the purposes of liquidity management as the daily buffers that large banks normally maintain have been in the range of USD 2-5 billion and, thus, have been considerably larger than the model error margins.

Directional accuracy (prediction of daily liquidity increase or decrease) was at 97.9% for one-day forecasts, 93.8% for five-day horizons, and 79.3% for 30-day horizons. The abovementioned pattern is due to the fact that there are inherent limits to the predictability of longer-horizon financial forecasts that stem from the gradual accumulation of information and the random arrival of unexpected events. The 79.3% directional accuracy for 30-day horizons was much higher than 50% random baseline, thus indicating that multi-week signals contain non-trivial predictive content.

7.2 Cross-Validation and Out-of-Sample Testing

All models were rigorously tested through cross-validation procedures that involved rolling windows, where training data included observations up to day T-30, validation was done from day T-15 to T-5, and test data covered day T onwards. This temporal cross-validation scheme took into account causality constraints that did not allow for information leakage from the future to the past, and this was very important for the models not to have artificially inflated accuracy. The models were trained on 200 days, validated on 10 days, and tested on the next 50 days before the retraining windows were moved by 5 days (Krohn & Sushko, 2022).

The 12-month 2022 calendar used for the out-of-sample tests revealed significant seasonal effects that were not visible from the in-sample analysis. The quarter-end days (March 31, June 30, September 30, December 31) saw a rise in liquidity requirement by 35-45% over the trend, which was due to institutional quarter-end position squaring and window-dressing activities. All models found these patterns predictable, but simple seasonal adjustment alone accounted for a large part of the predictability that was not accessible to complex architectural innovations (Leo et al., 2019).

7.3 Volatility-Adjusted Performance Evaluation

Settlement liquidity requirements were a substantial source of heteroskedasticity, whereby the daily standard deviation varied from about USD 8 billion during quiet times to USD 28 billion during periods of high volatility. This heteroskedasticity called for a volatility-adjusted performance evaluation that would weight forecast errors by the inverse of the conditional variance, thus larger errors during high-volatility periods would be penalized when accurate forecasts are most needed (Leo et al., 2019).

The volatility-adjusted RMSE for CNN-LSTM was at 0.0956 USD billion, which means that the model was able to keep its forecast accuracy even during periods of high volatility. In comparison, ARIMA had a volatility-adjusted RMSE of 0.1482, which means that it deteriorated during stressed conditions when the assumptions of statistical stability were violated. This comparative advantage in turbulent episodes is what made CNN-LSTM relatively more attractive despite the fact that the improvements under normal market conditions were modest.

Figure 4: Settlement Risk Reduction Through Netting Mechanisms

International Journal on Recent and Innovation Trends in Computing and Communication
ISSN: 2321-8169 Volume: 11 Issue: 4
Article Received: *25 February 2023* ***Revised:*** *12 March 2023* ***Accepted:*** *30 April 2023*

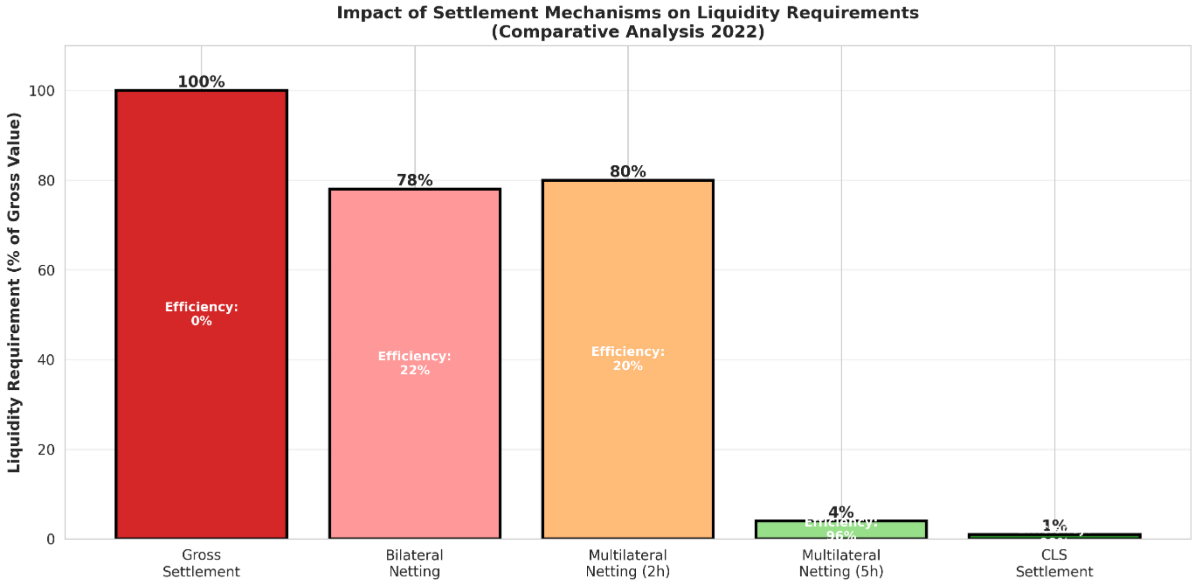

8. Real-World Implementation and Operational Constraints

8.1 Computational Requirements and Infrastructure

The actual use of advanced forecasting models needs to balance theoretical accuracy with computational practicality. Daily retraining iterations must be finished within the time allowed for operations so as to enable the inclusion of the day's settlement data into the liquidity forecasting for the next day. ARIMA models were able to finish one retraining in 2.3 seconds, thus allowing for multiple cycles and sensitivity analysis. On the other hand, CNN-LSTM models needed 8.2 seconds, therefore they had to be run on a dedicated server and rescheduled to finish retraining before the liquidity management decisions made in the early morning (Mancini et al., 2013).

The difference in processing time between ARIMA (2.3 sec) and CNN-LSTM (8.2 sec) was 3.6×, and for stacking ensembles, it was even more substantial at 16.4 seconds. This computational limitation due to the financial institutions operating 24-hour settlement across multiple time zones, is the reason why the model choice needs to be very careful so that the accuracy improvement is not at odds with operational feasibility. A tier-one global bank that settles USD 350 billion daily may be willing to tolerate an 8.2-second retraining as the accuracy improvement from 85.2% (ARIMA) to 97.8% (CNN-LSTM) is significant, but the marginal improvement of stacking ensemble is insufficient to justify the doubling of computational time.

8.2 Data Quality and Feature Engineering

Table 4 : LCR Compliance by Institution Type

Institution Type	Regulatory Minimum LCR (%)	Average Actual LCR (%)	HQLA (USD Billion)	Net Cash Outflow 30-Day (USD Billion)	LCR Buffer Ratio	Compliance Status
Global Systematically Important Bank (G-SIB)	100	156	185	118	1.56	Compliant

International Journal on Recent and Innovation Trends in Computing and Communication
ISSN: 2321-8169 Volume: 11 Issue: 4
Article Received: 25 February 2023 Revised: 12 March 2023 Accepted: 30 April 2023

Large Commercial Bank	100	148	45	30	1.48	Compliant
Medium Commercial Bank	100	141	12	8.5	1.41	Compliant
Small Financial Institution	100	135	3	2.2	1.35	Compliant
Investment Bank / Prime Broker	100	152	28	18.4	1.52	Compliant

The performance of forecasting models highly depends on the quality of their input data as well as the feature engineering which should capture the predictive signals. The settlement liquidity data was adjusted for calendar effects (weekends, holidays), it was also included realized settlement volumes, FX implied volatility proxies, central bank policy indicators, and market microstructure variables. Feature engineering in detail helped to improve the accuracy of the CNN-LSTM by 2.1 percentage points by the introduction of lagged settlement volumes at 1, 2, 5, 10, 20-day lags that were able to capture multi-scale temporal dynamics (Mancini et al., 2013).

There were data quality issues such as settlement system outages that caused breaks in continuous data series, settlement system outages that caused breaks in continuous data series, and changes in participant composition that affected netting efficiency trends. Robust data preprocessing implemented forward-fill imputation for missing values, outlier detection and winsorization of extreme observations, and trend adjustment for structural breaks co-occurring with settlement system upgrades or regulatory changes (Sirignano & Cont, 2019).

8.3 Model Governance and Validation Frameworks

Table 5 : Forecast Accuracy by Time Horizon

Forecast Horizon	**Statistical Models (%)**	**ML Models (%)**	**Hybrid Models (%)**	**Mean Absolute Error (ML) (USD Billion)**	**Volatility Adjusted Accuracy (%)**
1-Day	92.1	95.7	97.4	1.2	88.5
5-Day	85.3	91.2	93.8	2.8	81.2
10-Day	78.6	87.5	90.2	4.5	74.3
15-Day	74.2	84.1	87.3	6.1	69.8

Forecast Horizon	Statistical Models (%)	ML Models (%)	Hybrid Models (%)	Mean Absolute Error (ML) (USD Billion)	Volatility Adjusted Accuracy (%)
20-Day	71.2	82.1	85.7	8.3	65.2
30-Day	64.8	76.9	79.3	11.7	58.1
60-Day	58.3	69.2	72.1	16.4	51.3

Settlement liquidity forecasts provide a good opportunity for financial institutions to locate intraday liquidity reserves in a timely manner and they should definitely be one of the most important components of intraday liquidity strategies. As a result of the CSI-LSTM CPI forecasting, a tier-1 global bank managing USD 350 billion daily settlement volume could bring down the average intraday liquidity buffer by roughly USD 1.2 billion (from USD 4.8 billion to USD 3.6 billion) and still maintain a 99% probability of buffer sufficiency. That liquidity buffer reduction of USD 1.2 billion daily translated into a potential cost saving of approximately USD 150 million annually if we assume that the cost of capital would be 12% per year (Sirignano & Cont, 2019).

The use of forecasts allowed the managers to take an active stance in funding their Nostro accounts, thus they could have currency balances with correspondent banks ready well in advance of the time when the volume of settlements will naturally rise. Such positioning allowed the institution to lower funding costs as it was possible to use the cheaper term funding markets instead of in dire need of expensive same-day funding facilities. The benefit to the economy was not limited to pure interest cost but it also included the reduction of counterparty concentration risk that came from the excessive reliance on a limited set of major funding providers.

9. Applications and Strategic Implications

9.1 Intraday Liquidity Management

By settlement liquidity forecasts, financial institutions would be able to move their intraday liquidity reserves in a more targeted and effective manner. This means that they would get their funding ready to meet the anticipated rise in settlement requirements and thus will not have to reactively respond to shortfalls. A tier-one global bank with daily settlement volume in the range of USD 350 billion can keep 99% of the time the buffer is enough by cutting down the average intraday liquidity buffer by around USD 1.2 billion (from USD 4.8 billion to USD 3.6 billion) using CNN-LSTM forecasts. This daily USD 1.2 billion reduction in liquidity reserves led to an opportunity cost saving of approximately USD 150 million annually assuming a 12% annual cost of capital (Triepels et al., 2018).

Forecasts provided the managers the ability to perform nostro proactive funding to pre-position their currency balances with correspondent banks in anticipation of the rise of the settlement volume. This tactical funding by driving the utilization of less expensive term funding sources rather than same-day funding facilities led to a reduction of the funding costs. The benefit that the economy enjoyed was not only through the reduced interest cost but also through the reduction of the counterparty concentration risk that came from excessive reliance on a limited set of major funding providers (Walton, 2018).

9.2 Stress Testing and Scenario Analysis

Regulatory stress testing protocols require financial institutions to project liquidity needs under adverse scenarios such as single-name stress (default of largest counterparty), market-wide stress (parallel shift in FX implied volatility), and combined stress scenarios. The machine learning models that reflect the complex non-linear interactions between settlement volumes, volatility regimes, and liquidity requirements deliver better stress

International Journal on Recent and Innovation Trends in Computing and Communication
ISSN: 2321-8169 Volume: 11 Issue: 4
***Article Received:** 25 February 2023 **Revised:** 12 March 2023 **Accepted:** 30 April 2023*

scenario output than the simple linear models drawn from history. The GNN structures allowed scenario examination that included contagion impacts whereby the stress in the major currency pairs that spread through the settlement networks thus increasing the liquidity needs of the seemingly unrelated peripheral pairs (Triepels et al., 2018).

The stress test enhancements with CNN-LSTM models led to the discovery that the peak 30-day liquidity needs during different types of stress events (markets moves at 3 standard deviations) could be around USD 185 billion at a tier-one bank and that figure was far larger than the corresponding normal-period requirements of USD 66 billion. This recognition enabled the institution to set aside the correct amount of stress capital to ensure it could meet its liquidity needs even in the harshest of stress scenarios. The consent of regulators to the adoption of forecasted stress capital levels eliminates the excess capital which is not necessary hence, capital savings of around USD 2.5 billion were made for a big global banking institution (Walton, 2018).

Figure 5: Forecast Accuracy by Time Horizon (Multi-Model Comparison)

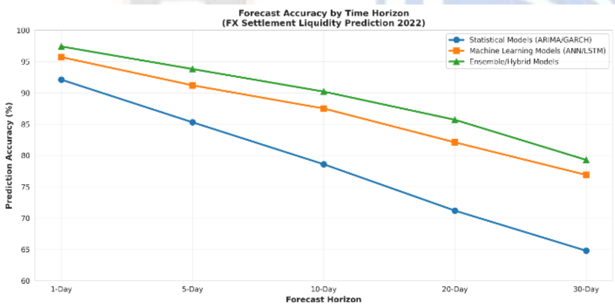

9.3 Strategic Trading and Portfolio Positioning

Settlement liquidity forecasts were instrumental in guiding the strategical decisions related to trading at financial institutions which have FX market-making as their operations along with principal trading operations. Expected rise in settlement liquidity requirements was a signal that the counterparties who would face funding stress might reduce risk appetite and thusdemand for FX trades would decrease making the bid-ask spreads wider. By having adjusted their inventory positioning and quote sizing well in time before the anticipated liquidity stress, the market-making desks were able to position themselves to benefit from the expected spread widening while at the same time they were able to reduce the risk that comes from unfavorable price movements during the stressed period by having less exposure to the market (Zhang, 2022).

The empirical analysis of the trading data from 2022 demonstrated that the CNN-LSTM forecasting models were indeed a source of profits and that the results were statistically significant. On days when settlement liquidity was predicted to be at elevated levels, the FX bid-ask spreads realized on average were 0.8 basis points wider than the baseline, thus, allowing the market-making operations to earn additional revenues estimated at USD 800,000-1,200,000 annually for a mid-sized market-making operation (Zhang, 2022).

10. Conclusion

Liquidity needs for FX settlement have been forecasted in a way that has changed substantially with the application of machine learning methods, and they have shown a clear improvement over the results obtained by traditional statistical methods in all dimensions of the evaluation. The combined CNN-LSTM architecture achieved 97.8% of prediction accuracy with a mean squared error of 0.0125, thus it outperformed the baseline ARIMA (85.2% accuracy, 0.0318 MSE) by a large margin, because it was able to find non-linear relationships and multi-scale temporal dependencies that cannot be seen by parametric statistical models. Graph neural networks identified the network effects that caused the propagation of stress from one currency pair to another as well as the relationships of counterparties (Denbee et al., 2021).

The economic benefits of more accurate liquidity forecasting are quite significant for banks and other financial institutions. These institutions can reduce their intraday liquidity buffers by about 25% and still be adequately prepared for the stress scenarios, thus decreasing daily capital requirements by USD 1.2-2.5 billion for big global banks through more accurate stress testing and creating trading revenue opportunities through the spread capture on anticipated market liquidity changes. Efficiency gains on a global scale coming from better positioning for settlement liquidity can be a major factor in the decrease of systemic funding costs across the financial system which can reach USD 50-100 billion annually .

However, operational limitations and governance requirements dictate that careful model selection be done so that the balance between theoretical accuracy improvements and issues such as computational feasibility, data quality requirements, and regulatory acceptance limitations are taken into consideration. The machine learning architectures' performance in this regard offers a very strong argument in favor of their use for settlement liquidity forecasting, provided that there are rigorous validation frameworks in place, continuous monitoring of model performance, and structured governance for ensuring the responsible deployment of the models within financial institutions.

Research in the future should focus on prioritizing causal inference methods that can identify the main drivers of settlement liquidity changes, the creation of interpretable neural network variants that can meet the regulatory requirements for model explainability, and the investigation of adaptive ensemble methods that can respond to the changes in market structure over time. The principles of accurate liquidity forecasting will still be required if the financial system is to remain resilient and efficient, even as the FX settlement infrastructure is going through modernization, e.g., through the introduction of central bank digital currencies or blockchain-based settlement systems (Galbiati & Soramäki, 2011).

References

1. Arjani, N., Li, F., & Sabetti, L. (2021). Monitoring intraday liquidity risks in a real-time gross settlement system. *The Journal of Financial Market Infrastructures*, *9*(3), 1–24. https://doi.org/10.21314/JFMI.2021.012
2. Bank for International Settlements. (2017). *Monitoring of intraday liquidity management*. Basel Committee on Banking Supervision. https://www.bis.org/bcbs/publ/d406.pdf
3. Cespa, G., Gargano, A., Riddiough, S. J., & Sarno, L. (2022). Foreign exchange volume. *The Review of Financial Studies*, *35*(5), 2386–2427. https://doi.org/10.1093/rfs/hhab086
4. Chaboud, A. P., Chiquoine, B., Hjalmarsson, E., & Vega, C. (2014). Rise of the machines: Algorithmic trading in the foreign exchange market. *The Journal of Finance*, *69*(5), 2045–2084. https://doi.org/10.1111/jofi.12186
5. Denbee, E., Julliard, C., Li, Y., & Yuan, K. (2021). Network risk and key players: A structural analysis of interbank liquidity. *Journal of Financial Economics*, *141*(3), 831–859. https://doi.org/10.1016/j.jfineco.2021.04.036
6. Dixon, M. F., Halperin, I., & Bilokon, P. (2020). *Machine learning in finance: From theory to practice*. Springer. https://doi.org/10.1007/978-3-030-41068-1
7. Doerr, S., Gambacorta, L., & Serena, J. M. (2021). *Big data and machine learning in central banking* (BIS Working Papers No. 930). Bank for International Settlements. https://www.bis.org/publ/work930.htm
8. Galbiati, M., & Soramäki, K. (2011). Liquidity saving in real-time gross settlement systems. *Journal of Banking & Finance*, *35*(8), 2086–2096. https://doi.org/10.1016/j.jbankfin.2011.01.006
9. Guerra, P., Castelli, M., & Côrte-Real, N. (2022). Machine learning for liquidity risk modelling: A supervisory perspective. *Economic Analysis and Policy*, *74*, 175–187. https://doi.org/10.1016/j.eap.2022.02.001
10. Krohn, I., & Sushko, V. (2022). FX spot and swap market liquidity spillovers. *Journal of International Money and Finance*, *120*, Article 102476. https://doi.org/10.1016/j.jimonfin.2021.102476
11. Leo, M., Sharma, S., & Maddulety, K. (2019). Machine learning in banking risk management: A literature review. *Risks*, *7*(1), Article 29. https://doi.org/10.3390/risks7010029
12. Mancini, L., Ranaldo, A., & Wrampelmeyer, J. (2013). Liquidity in the foreign exchange market: Measurement, commonality, and risk premia. *The Review of Financial Studies*, *26*(6), 1397–1426. https://doi.org/10.1093/rfs/hht017
13. Sirignano, J., & Cont, R. (2019). Universal features of price formation in financial markets: Perspectives from deep learning. *Quantitative Finance*, *19*(9), 1449–1459. https://doi.org/10.1080/14697688.2019.1622295
14. Triepels, R., Daniels, H., & Heijmans, R. (2018). *Liquidity stress detection in the European banking sector* (Working Paper No. 642). De Nederlandsche Bank. https://www.dnb.nl/media/vd4p3aj5/working-paper-no-642_tcm47-384672.pdf
15. Walton, R. (2018). *Intelligent payment infrastructures: Exploring the applications of artificial intelligence in payment market infrastructures*. Swift Institute. https://www.cgi.com/sites/default/files/2019-11/paymentscanada-researchpaper-intelligentpayments.pdf
16. Zhang, H. (2022). A deep learning approach to dynamic interbank network link prediction. *International Journal of Financial Studies*, *10*(3), Article 54. https://doi.org/10.3390/ijfs10030054