Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Beyond the Median: Reserve Scarcity, Liquidity Buffers and Distributional Stress in US Repo Markets

Vamsidhar Ambatipudi

ambatipudi.v@pilani.bits-pilani.ac.in

Birla Institute of Technology and Science, Pilani

Research Article

Keywords: repo markets, bank reserves, money market dispersion, overnight reverse repo, financial intermediation, funding risk

Posted Date: November 12th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11444156/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Beyond the Median: Reserve Scarcity, Liquidity Buffers and Distributional Stress in US Repo Markets

Original Article

Abstract

How can banks and central banks recognise pressure in overnight funding markets when most transactions still take place at ordinary rates? This study examines the full published distribution of US secured overnight funding rates from April 2018 to October 2026. It combines 2,134 SOFR observation dates with reserve, bank-asset and overnight reverse repo information, using 438 weekly observations for the main reserve regressions. Funding costs are related to reserves throughout the rate distribution, with a larger estimated slope at the 99th percentile (−2.00 basis points per percentage-point increase in reserves relative to bank assets) than at the median (−1.21 basis points). Dispersion between the 75th and 25th percentiles also narrows as reserves rise. Interactions with the Federal Reserve's overnight reverse repo balance and quarter-end reporting dates identify conditions under which these associations differ. The upper-tail reserve threshold is not consistently earlier than the median threshold. In a separate 2023-onward forecast test, upper-tail information improves Brier scores relative to a median model at one and five days, but forecast probabilities require substantial recalibration. The results support a distribution-aware approach to liquidity monitoring that combines reserve levels, liquidity buffers, market segmentation and disciplined forecast verification.

Keywords: repo markets; bank reserves; money market dispersion; overnight reverse repo; financial intermediation; funding risk

JEL classifications: E42; E52; E58; G21; G23

1 Introduction

The overnight repo market is a basic part of the financial system. Banks, dealers and investors use it to lend and borrow cash against securities, often for only one day. When this market works smoothly, its rates are close together and funding remains available. When pressure develops, a few borrowers may face much higher rates even while the rate on a typical transaction remains stable. This difference is important for bank treasury teams, market participants and central banks responsible for smooth monetary-policy implementation.

The standard reserve-adequacy question asks how many bank reserves are needed before an additional dollar has little effect on overnight interest rates. Afonso et al. (2025) provide an influential time-varying reserve-demand analysis and place the broadly abundant region near 12% to 13% of bank assets. This approach is valuable, but a market median summarises the centre of a distribution. It need not summarise the experience of institutions obtaining funding at high rates. Earlier work on trading frictions, liquidity hoarding and short-term financing explains why a funding market may be uneven even with ample total liquidity (Afonso and Lagos 2015; Acharya and Skeie 2011; Brunnermeier and Pedersen 2009).

This paper asks three connected questions. First, do reserve conditions relate differently to the low, middle and high parts of the repo-rate distribution? Second, does that relationship depend on other cash buffers, particularly the Federal Reserve's overnight reverse repo facility, or on balance-sheet reporting dates? Third, does observing the high-rate tail improve the identification of future pressure, and are the resulting probabilities reliable enough to guide decisions? These questions link the pricing of a financial service to the distribution of intermediation costs and the design of public monitoring systems.

The study uses official US rate percentiles from April 2018 to October 2026, including the September 2019 funding-market disturbance, the period of unusually large liquidity balances, and later quantitative tightening. It

estimates reserve associations at five published percentiles, models nonlinear reserve thresholds, examines two-stock liquidity interactions, contrasts related repo markets, studies calendar pressures and evaluates genuinely later forecast observations. The empirical story is built around patterns that survive more than one comparison, rather than assuming that the most expensive transaction must always warn earlier than the median.

The findings add three elements to the debate. First, there is a meaningful distributional gradient in reserve sensitivity and a clearly estimated narrowing of the interquartile rate range as reserves increase. Second, the observed relationships depend on ON RRP conditions and market composition, suggesting that aggregate reserve levels alone provide an incomplete account of funding pressure. Third, formal forecast evaluation shows why the ability to rank vulnerable days must be separated from the ability to attach trustworthy probabilities to them. This distinction matters when moving from academic prediction to operational financial-stability decisions.

2 Literature and analytical framework

2.1 Reserves, trading frictions and unequal funding costs

Reserve balances permit banks to settle payments and meet internal liquidity needs. If every bank could exchange reserves freely and immediately, one aggregate reserve measure might be sufficient to understand overnight funding conditions. Actual markets operate with relationship networks, credit limits, balance-sheet costs and institutional differences. Afonso and Lagos (2015) show how search and bargaining influence interbank trades. Acharya and Skeie (2011) explain how liquidity-hoarding incentives can emerge under funding risk. The conceptual implication is that a constraint can bind for some participants even if others remain comfortable.

Repo markets add secured collateral and dealer intermediation to these frictions. Copeland et al. (2014) document how borrower-level experiences in the tri-party market can differ from aggregate market conditions. Krishnamurthy et al. (2014) show why the size and institutional composition of repo finance matter. Gorton and Metrick (2012) explain the significance of repo finance in periods of broader financial disruption. Together, these studies support the proposition that both the level and distribution of rates are relevant to an assessment of short-term financing conditions.

The empirical reserve-demand framework of Afonso et al. (2025) estimates how the marginal funding rate responds as aggregate reserves change. The present analysis complements that approach by asking how the relationship varies across the published transaction-weighted percentile ladder. It does not equate a high-rate percentile with a named bank or assume that each percentile corresponds to the same borrower over time.

2.2 Liquidity buffers, reporting dates and market measurement

Central bank reserves are only one component of the broader cash environment. Money-market investors can also place cash in the Federal Reserve's overnight reverse repo facility. A stock of cash at this facility can be relevant to the supply of funding available elsewhere, although its effects depend on institutions, incentives and market conditions. The broader theoretical interpretation is a two-stock liquidity environment: reserves held by banks and cash positioned at the reverse repo facility jointly describe an important part of monetary-policy implementation (Sunderam 2015; Afonso et al. 2021).

Dealer balance-sheet constraints and quarter-end reporting incentives offer a second explanation for unusually high rates. Intermediaries may face a particularly costly expansion of balance sheets at reporting dates, regardless of changes in aggregate reserves. Such concerns are consistent with models connecting market liquidity and funding liquidity (Brunnermeier and Pedersen 2009) and evidence on US Treasury intermediation capacity (Duffie 2020). The September 2019 episode illustrates how reserve distribution, payment flows and intermediation frictions may interact rather than represent separate explanations (Afonso et al. 2021).

The choice of reference rate also matters. SOFR covers a wider set of Treasury secured funding transactions than tri-party-only measures. Different transaction samples can respond differently to market conditions. Müller and Paulick (2026) demonstrate the importance of the measurement approach when studying overnight money markets, while Norden and Weber (2010) show why changes in banks' funding modes matter for intermediation and stability. These insights motivate separate comparisons of SOFR, TGCR and BGCR rather than treating them as identical indicators.

2.3 Hypotheses and competing interpretations

The analysis evaluates three claims. The distributional claim is that reserve sensitivity varies across rate percentiles and can be seen in the gaps between percentiles. The conditional-liquidity claim is that reserves and ON RRP balances jointly characterise observed pricing, while reporting-date pressures can also affect the high-rate tail. The information claim is that lagged upper-tail measures can improve future-stress prediction beyond lagged median information. A positive predictive comparison, however, is not enough to establish useful event probabilities. Calibration and comparison with simple reference forecasts are separate requirements (Brier 1950; Gneiting and Raftery 2007).

3 Data and empirical design

3.1 Market coverage and measurement

The US setting is observed at business-day frequency from 2 April 2018 to 16 October 2026. The New York Fed's reference-rate data provide SOFR, TGCR and BGCR overnight secured funding rates and the effective federal funds rate (EFFR). For SOFR, there are 2,134 reported dates; the first and 99th percentiles are populated on 2,132 of these dates. For every market and percentile, the reported annualised rate is compared with interest on reserve balances (IORB), or its earlier counterpart, interest on excess reserves (IOER). All spreads are multiplied by 100 to express them in basis points. The upper-tail premium is the 99th-percentile spread minus the median spread. Full dispersion is the 99th percentile minus the 1st percentile, while middle dispersion is the 75th minus the 25th percentile.

Weekly reserve balances are obtained from the Federal Reserve's WRESBAL series and scaled by the H.8 measure of commercial-bank assets. Reserve balances are converted from US dollars in millions to billions, while H.8 assets are in billions. The principal reserve indicator equals reserves divided by bank assets, commonly reported in percentage points. Daily ON RRP balances and publicly reported secured funding volumes provide complementary liquidity and trading-activity measures. The data sources, units and frequency are summarised in Table 1. Daily Treasury balances cover a smaller portion of the panel; therefore the associated daily regression uses a restricted sample. Treasury settlement values could not support a separate empirical test and are not used to claim identification.

Table 1. Official sources, coverage and variable definitions

Series / construct	Reported sample	Role in analysis
SOFR / TGCR / BGCR rate percentiles	2,134 / 2,134 / 2,134 dates, Apr 2018–Oct 2026	Funding prices across transaction-weighted percentiles
Effective federal funds rate	2,146 dates, Apr 2018–Oct 2026	Unsecured-market comparison
WRESBAL bank reserves	446 weekly values within main sample window	Weekly reserve stock, scaled by bank assets
H.8 total commercial-bank assets	445 weekly values within main sample window	Denominator of reserve ratio
ON RRP, RRPONTSYD	2,133 values in study window	Nonbank cash-buffer indicator
IORB / earlier IOER	Policy series joined by effective date	Common rate anchor for basis-point spreads
Weekly estimation sample	438 weeks	Percentile elasticities, thresholds and two-stock regressions
Daily Treasury-change subset	1,008 SOFR-tail regression days	Calendar interactions with lagged Treasury information

Note: Sample sizes refer to published inputs or complete-case regressions, not identical observations for every variable. Source: New York Fed Markets API, Federal Reserve Economic Data and US Treasury Fiscal Data; authors own calculations.

3.2 Identification, time alignment and inference

The weekly analysis is used for reserve-response estimation so that repeatedly carrying a weekly reserve number into daily records does not create false independent information. The central model relates each weekly SOFR percentile spread to the reserve-to-assets ratio, quarter-end and month-end indicators, log funding volume and an indicator for the later policy regime. Heteroskedasticity- and autocorrelation-consistent uncertainty estimates are used to accommodate serial dependence. The central descriptive parameter is the change in a given rate spread, in basis points, associated with a one-percentage-point change in the reserve-to-assets ratio, holding the listed covariates constant.

Three additional approaches examine whether the relationship is stable and informative. The threshold logic follows the distinction between slope and cutoff uncertainty emphasised by Hansen (1999). Serial-dependence-aware resampling draws on the wider time-series bootstrap literature (Politis and Romano 1994). First, percentile-specific threshold regressions fit separate slopes on either side of a reserve-ratio cutoff, and a blocked resampling procedure assesses the difference between the 99th-percentile and median cutoffs. Second, reserve balances are interacted with ON RRP balances or a low-buffer indicator; the models also account for volume and calendar periods. Third, daily models evaluate interactions with quarter-end and month-end dates, while alternative models compare the upper parts of the SOFR, TGCR and BGCR distributions. An extreme-value fit, a common-factor analysis and structural-stability tests provide supplementary interpretation, not stand-alone causal identification.

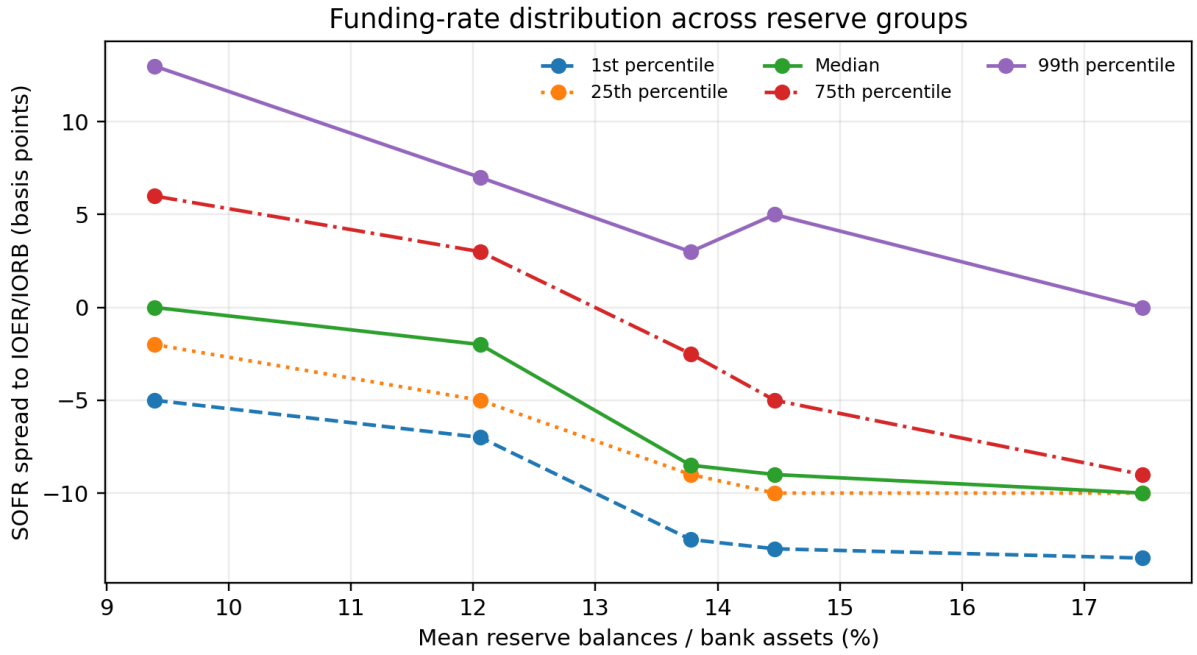
For prediction, monetary controls are aligned with conservative publication delays and lagged information is used where required. The notebook constructs point-in-time proxies by delaying weekly reserve values by two days and bank-asset readings by twelve days; these are conservative publication-lag conventions, not historical data vintages. Forecast comparison follows the principle of assessing paired loss differences on common observations (Diebold and Mariano 1995). Forecast parameters are trained on observations ending in 2022 and are assessed on separate observations from 2023 onwards. One-, five- and ten-day horizons are considered. For the principal event, the future 99th-percentile SOFR spread exceeds 25 basis points above the reserve remuneration rate. Model comparisons use the area under the ROC curve (AUC), precision-recall AUC, Brier score and reliability grouping. The paired Brier loss is evaluated on identical test observations with serial-dependence-aware uncertainty.

4 Empirical results

4.1 A reserve gradient across funding-rate percentiles

Figure 1 shows how the complete SOFR rate distribution changes across five groups of reserve availability. In the lowest-reserve group, the median spread is approximately zero and the 99th percentile is 13 basis points above the reserve remuneration rate. In the highest-reserve group, the median spread is −10 basis points and the 99th percentile is approximately zero. Full dispersion is about 18 basis points in the lowest-reserve group and 15 basis points in the highest-reserve group. These descriptive comparisons show that movements in the centre of the market do not exhaust movements across participants.

Figure 1. Funding-rate fan across observed reserve-to-assets groups



Note: Each plotted point is a recorded summary from the five reserve groups. It is not an estimated causal response or a fitted daily reserve-demand curve.

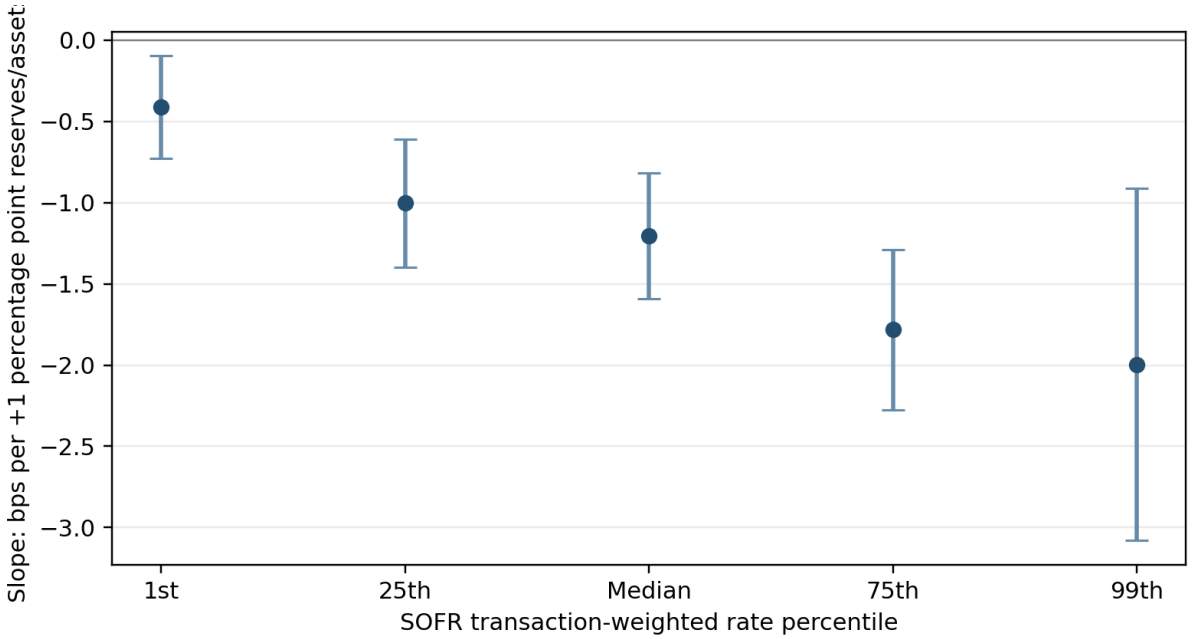
The weekly controlled regressions in Table 2 and Figure 2 sharpen this pattern. A one-percentage-point increase in the reserve-to-assets ratio is associated with a fall of approximately 0.41 basis points at the first percentile, 1.21 basis points at the median and 2.00 basis points at the 99th percentile. Each of these slopes is statistically distinguishable from zero at the five-per-cent level. The estimate at the upper end is about 66% larger in magnitude than at the median. However, the direct difference between the upper-tail and median slopes is not statistically established at the five-per-cent level ($p = 0.069$). The firm finding is a broad gradient in estimated associations, not a conclusive proof that the 99th percentile is always more responsive.

Table 2. Reserve associations across SOFR rate percentiles

Outcome	Slope (bps per +1 pp)	95% confidence interval	p value	Weeks
1st	-0.414	[-0.729, -0.098]	0.01019	438
25th	-1.003	[-1.396, -0.610]	5.628e-07	438
Median	-1.206	[-1.594, -0.818]	1.13e-09	438
75th	-1.783	[-2.275, -1.290]	1.35e-12	438
99th	-1.996	[-3.079, -0.914]	0.0003004	438
75th − 25th	-0.779	[-0.934, -0.625]	5.132e-23	438
99th − median	-0.791	[-1.643, 0.061]	0.06896	438

Note: Weekly models control for quarter-end, month-end, SOFR log volume and the later-policy-regime indicator. HAC confidence intervals. A negative slope indicates lower spreads when reserves/assets are higher.

Figure 2. Reserve-response coefficient ladder and 95% confidence intervals



Note: Points and intervals are the same controlled weekly estimates reported in Table 2. Figure is constructed directly from the evidence workbook.

A particularly clear result concerns the middle of the distribution. Table 2 shows that the 75th-to-25th-percentile gap narrows by approximately 0.78 basis points for a one-percentage-point increase in reserves relative to bank assets (95% confidence interval −0.934 to −0.625). This is a direct dispersion result and does not depend on treating an individual extreme transaction as representative. It offers a simple operational measure that adds information beyond the median.

4.2 Nonlinear reserve relationships

Reserve-response models need not have the same slope when reserves are already plentiful as when banks have smaller cushions. Table 3 reports the estimated cutoff for each percentile in a segmented regression. The median cutoff is 13.50% of bank assets and the 99th-percentile cutoff is 11.13%. Both piecewise specifications improve their information criterion compared with a single linear slope, but the cutoffs should not be interpreted as universal policy targets. Their difference, −2.37 percentage points, has a blocked bootstrap interval from −4.08 to

0.00 percentage points. Removing September 2019 reduces the gap to −0.52 percentage points. Thus, the data support nonlinearity but not a stable claim that high-rate tails turn sharply upwards earlier than the median.

Table 3. Estimated percentile-specific reserve thresholds and threshold sensitivity

Measure	Estimated value	Model comparison / uncertainty	Coverage
1st	13.50%	53.09	438
25th	13.04%	48.49	438
Median	13.50%	37.94	438
75th	11.13%	23.24	438
99th	11.13%	26.79	438
99th minus median	−2.37 pp	95% bootstrap interval: [−4.08, 0.00]	80 resamples
Gap without September 2019	−0.52 pp	Sensitivity specification	435 weeks

Note: Delta BIC is improvement of segmented over linear fit; positive values favour segmented fit. Bootstrap threshold inference is exploratory because cutoff estimation is nonregular. The final row is an exclusion sensitivity, not a separate main-sample estimate.

4.3 Reserves and reverse repo cash balances are jointly relevant

Table 4 examines the central two-stock proposition. When reserves and ON RRP are entered together with their interaction, the interaction coefficient is positive and statistically significant for the median (0.130; $p < 0.001$) and for the 99th percentile (0.192; $p = 0.008$). In the specified model, ON RRP is measured in units of US$100 billion and reserves in percentage points of bank assets. The estimates imply that the reserve-spread association varies across the observed ON RRP environment. A second specification compares periods with a low ON RRP buffer and also reports statistically significant interactions at the median and 99th percentile. This pattern supports conditional liquidity monitoring, although it does not identify a causal effect of shifting money into or out of the facility.

Table 4. Two-stock liquidity specification: reserves, ON RRP and their interaction

Outcome	Term	Estimate	95% confidence interval	p value
Median	Reserves/assets	-1.622	[-2.051, -1.193]	1.3e-13
Median	ON RRP per US$100bn	-1.820	[-2.362, -1.279]	4.37e-11
Median	Reserves × ON RRP	0.130	[0.084, 0.176]	2.87e-08
99th percentile	Reserves/assets	-2.462	[-3.805, -1.119]	0.000328
99th percentile	ON RRP per US$100bn	-2.306	[-3.917, -0.696]	0.005
99th percentile	Reserves × ON RRP	0.192	[0.050, 0.334]	0.00815
Full spread	Reserves/assets	-1.691	[-2.986, -0.396]	0.0105
Full spread	ON RRP per US$100bn	-0.873	[-2.445, 0.698]	0.276
Full spread	Reserves × ON RRP	0.102	[-0.033, 0.238]	0.14

Note: 438 weekly observations per model. Each model also contains calendar controls, log SOFR volume and a post-2021 indicator. These coefficients describe conditional historical associations and cannot independently identify policy intervention effects.

The descriptive state cells further illustrate the interaction. In the lowest observed reserve group and the lowest ON RRP group, the median recorded 99th-percentile spread is 10 basis points. At intermediate reserves and the lowest ON RRP group, the corresponding median is 8 basis points, compared with 2 basis points in the higher ON RRP group available for that reserve band. These cell comparisons combine changes in economic conditions

and sample composition; the controlled interaction models are more appropriate for inference. The key finding is the value of observing the two stocks jointly.

4.4 Quarter-end pressures and market-segment differences

Funding markets are also shaped by time-specific capacity constraints. Table 5 summarises the calendar and market-segment results. The daily regression with Treasury-change availability shows a negative quarter-end × reserve-ratio coefficient for full SOFR dispersion (−9.80; p = 0.028). A related model for the 99th-percentile-minus-median gap gives −9.03 (p = 0.032). These coefficients should be read as an interaction: quarter-end effects vary with the reserve ratio. They are not unconditional increases of 9 or 10 basis points on every quarter-end day. A useful interpretation is that reserve conditions and balance-sheet reporting dates reinforce one another in the observed rate distribution. This is consistent with, but does not separately establish, a dealer balance-sheet mechanism.

The comparison of nearby secured markets gives a second type of evidence. The reserve coefficient for SOFR upper-spread minus TGCR upper-spread is −0.973 (p < 0.001); for SOFR upper-spread minus BGCR upper-spread it is −0.270 (p = 0.031). A volume-share control is also significant in the BGCR comparison. These results indicate that expanding the basket of eligible transactions changes how the upper part of the measured rate distribution relates to reserves. They do not identify individual dealer actions because the benchmarks contain overlapping transactions.

Table 5. Calendar and market-segment evidence

Dependent outcome and term	Coefficient	95% confidence interval	p value	Days
Full dispersion: quarter-end interaction	-9.804	-18.54 to -1.07	0.028	1008
Upper dispersion: quarter-end interaction	-9.029	-17.26 to -0.80	0.032	1008
99th spread: quarter-end interaction	-10.727	-20.52 to -0.94	0.032	1008
SOFR − TGCR upper-spread reserve coefficient	-0.973	-1.49 to -0.46	0.000	2132
SOFR − BGCR upper-spread reserve coefficient	-0.270	-0.52 to -0.02	0.031	2132

Note: First three rows: daily model including lagged Treasury change and qend × reserve/assets; last two rows: difference between markets in published upper-tail spreads. Calendar and composition covariates differ by specification. All are observational associations.

4.5 Common market stress, extremes and structural variation

Table 6 brings together the common-factor, extreme-event and structural-change checks. Funding pressure has a shared market component but is not uniform. A principal-component analysis of upper-rate dispersion across SOFR, BGCR, TGCR and EFFR finds that the first common factor accounts for 61.1% of the selected measures' variation. At the same time, the part of SOFR upper-rate variation left after removing that common factor still shows associations with reserve conditions and volume. This is compatible with both system-wide and market-specific sources of funding pressure. It also strengthens the case for monitoring several related benchmarks.

An extreme-value description records 151 days above the selected 18-basis-point upper-spread threshold, grouped into 57 clusters. Its estimated shape parameter is 0.75. A separate regime model identifies low- and high-dispersion states with means of about 10.26 and 22.54 basis points. Because the smoothed regime probabilities use information from later dates, they are used only to characterise historical patterns. Finally, a CUSUM test

shows instability in the median spread model ($p < 0.001$), while the analogous upper-dispersion test is not significant ($p = 0.415$). These results support the practical value of monitoring changing market conditions rather than treating one historical rate relationship as permanent.

Table 6. Complementary evidence on fragmentation, extremes and temporal change

Indicator	Result	Interpretation
Explained variation, first common factor	61.13%	SOFR, BGCR, TGCR, EFFR upper spread components
High-dispersion regime mean	22.54 bps	Descriptive Markov state model
Low-dispersion regime mean	10.26 bps	Descriptive Markov state model
Extreme-value event clusters	57	151 exceeding days, 18 bps threshold
Median-model CUSUM p value	< 0.001	Model stability rejected
Upper-dispersion CUSUM p value	0.415	No rejection in this test

Note: Regime means are smoothed in-sample descriptions. The extreme-value fit is calculated on published daily percentiles rather than underlying transaction-level extreme observations.

4.6 Forecast evaluation: ranking days versus reliable probabilities

Table 7 presents the corresponding forecast scores, including the constant-frequency reference for the one-day horizon. The out-of-sample experiment measures whether information from the upper end of the rate distribution helps to recognise days when funding pressure is more likely. In the one-day 2023-onward test, 17 events occur in 941 assessed observations, an event frequency of 1.81%. A median-only model has an ROC AUC of 0.829 and a Brier score of 0.409. Adding lagged upper-tail information increases ROC AUC to 0.841 and lowers the Brier score to 0.343. The paired loss gain is 0.0657 with a blocked 95% confidence interval of 0.0534 to 0.0769. At five days, the Brier gain is 0.0426, also positive, while at ten days the tail model is slightly worse and the difference is inconclusive. These comparisons establish relative value against the particular median model at selected horizons.

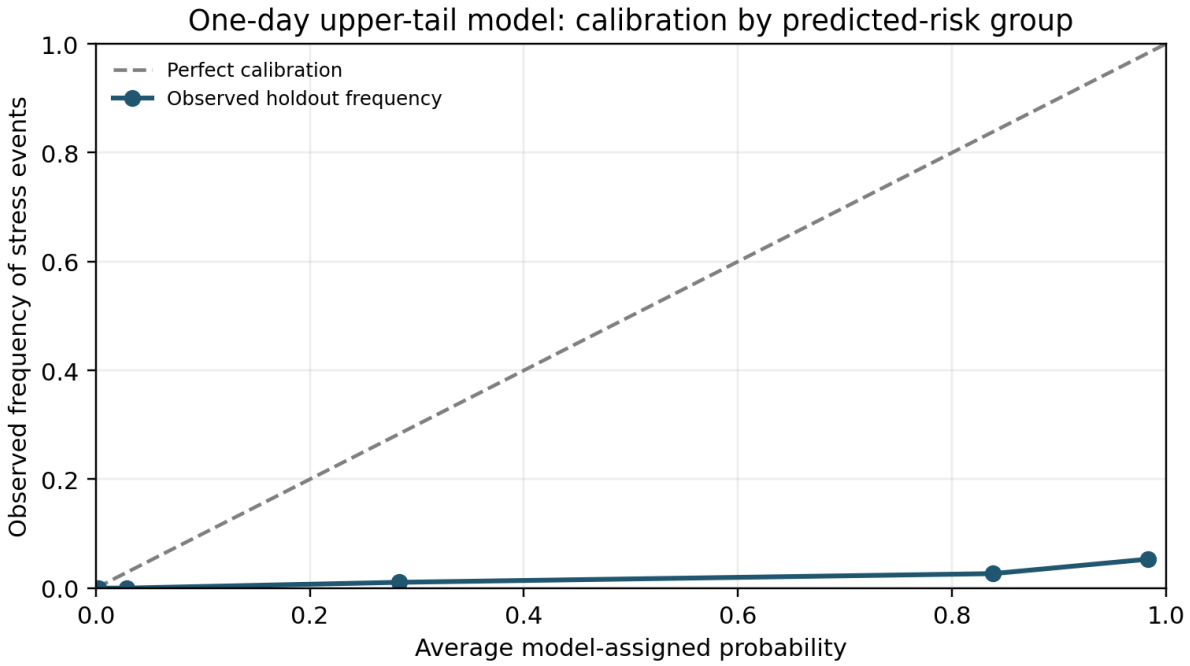
The forecast study also provides a more demanding finding for practice. A probability should be close to the actual event frequency among similarly predicted days. Figure 3 demonstrates that the upper-tail model substantially overstates one-day event frequencies in high predicted-risk groups. The mean predicted one-day probability is approximately 42.7%, whereas realised events occur on about 1.8% of the assessed days. A constant forecast using the 3.60% training event rate achieves an approximately 0.018 Brier score on the one-day test, far below either complex model. The upper-tail model may therefore rank relatively risky days, but its reported event probabilities are not suitable for direct decision thresholds. This distinction is as important as the relative improvement against the median model.

Table 7. Out-of-sample forecasts against the median and a constant-probability reference

Days ahead	Model	Test days	Event rate	ROC AUC	PR AUC	Brier
1	Median	941	1.81%	0.829	0.281	0.4085
1	Upper-tail	941	1.81%	0.841	0.274	0.3428
1	Training-rate-only reference	941	1.81%	Not defined	Not defined	0.0181
5	Median	917	6.00%	0.723	0.197	0.1310
5	Upper-tail	917	6.00%	0.717	0.174	0.0884
10	Median	887	10.71%	0.726	0.278	0.0905
10	Upper-tail	887	10.71%	0.727	0.260	0.0934

Note: The constant reference predicts the 1-day training event frequency (3.60%) at every date; its 1-day score is recalculated from the test outcome count. For the 5-day comparison, paired Brier gain is 0.0426 [0.0326, 0.0520]. The 10-day gain is −0.0028 [−0.0089, 0.0020]. Lower Brier is better.

Figure 3. One-day stress forecast reliability compared with perfect calibration



Note: Five equal-frequency groups of saved 2023-onward upper-tail predictions. The 45-degree line indicates a forecast whose assigned probabilities match realised event frequencies.

Severity sensitivity provides a useful qualification. For a more frequent 15-basis-point event, the upper-tail model reports an ROC AUC of 0.910 and a Brier score of 0.0331, compared with 0.899 and 0.0385 for the median model. Its performance on rarer and more severe events is less dependable in absolute probability terms. Accordingly, additional distributional information is promising for alert ranking, while probability calibration and simple forecast benchmarks must remain core parts of an operational assessment.

5 Discussion and contribution to research

The central result is that funding conditions cannot be reduced to a single representative overnight rate. Even within the same secured funding benchmark, percentile slopes differ and middle dispersion is clearly associated with reserve abundance. This finding develops the reserve-demand discussion from an analysis of the representative funding price towards an analysis of the whole distribution of funding costs. The extension is economically meaningful because a borrowing institution's experience may diverge from the reported market median.

The conditional models refine the usual scarcity interpretation. Interactions between reserves and ON RRP are observable in both median and 99th-percentile models, and calendar effects are particularly important in the daily distributional specification. A plausible analytical view is therefore one of layered funding conditions: the quantity of reserves, the stock of other cash available to money markets, the balance-sheet capacity of intermediaries and the transaction mix together shape the observed rate distribution. This is a synthesis of competing explanations, not a claim that the study has causally isolated each channel.

The analysis also establishes an important boundary to the original tail-first hypothesis. Although upper-end rates are associated with reserve conditions, the segmented threshold at the 99th percentile is lower than the median threshold in the full sample and the gap is sensitive to September 2019. The evidence is more informative when read as support for heterogeneous and nonlinear pricing than as proof of a universal sequence in which high-rate borrowers encounter scarcity first. This conclusion is compatible with reserve distribution and dealer-capacity accounts operating together.

The forecasting evidence adds a methodological distinction to the literature on financial-stability indicators. AUC assesses the ordering of risky days, whereas the Brier score also assesses the correctness of stated probabilities. A relative gain between two fitted models does not, by itself, demonstrate useful probabilities. When a training-rate-only forecast dominates both fitted models in absolute Brier score, reported probabilities cannot be translated mechanically into policy thresholds. Hence the contribution is not a claimed operational probability forecast but a clear demonstration of what careful validation should test.

Finally, measurement is part of the economic argument. Repo benchmarks overlap but do not contain identical trades, and the market-specific coefficients show that using one aggregate distribution can conceal relevant differences. The methodological contribution is to combine official percentile ladders, matched weekly reserve regressions, conditional liquidity interactions, cross-market contrasts and forecast evaluation within one accountable design. It makes clear which findings are robust associations, which are explorations of market mechanisms and which are specific to the forecasting experiment.

6 Implications for financial services and policy

For central bank implementation teams, the evidence supports a monitoring dashboard that reports the median SOFR spread, the 75th-to-25th-percentile gap, the 99th-percentile spread and the corresponding TGCR and BGCR distributions. Reserve balances should be displayed relative to the banking system's size and alongside ON RRP outstanding balances. Such a dashboard would not replace reserve-demand estimation, but it could reveal pricing differences that a single median series does not display.

For bank treasury departments, the results support stress reviews based on a range of funding costs rather than a single published overnight benchmark. The stronger dispersion relationship at low reserves and the quarter-end interactions suggest that liquidity planning should consider both aggregate conditions and the timing of balance-sheet reporting. Scenario design can vary the reserve environment and calendar conditions separately, while recognising that the model estimates describe historical associations rather than tested interventions.

For repo dealers and money-market investors, differences across SOFR, TGCR and BGCR are informative about the transactions included in a benchmark. Comparing spreads across market segments can help distinguish broad changes in overnight rates from changes concentrated in particular transaction baskets. This is relevant for trading limits, internal market-risk reporting and oversight of secured funding exposure.

For financial-stability supervisors, the forecasting exercise supplies a practical quality-control rule. Before an early-warning score is presented as an event probability, its calibration must be checked on data not used to develop the model and compared with a simple historically observed event frequency. When calibration is weak, the score may still rank days for review, but cannot reliably quantify the chance of stress. This distinction prevents a statistically attractive model from being mistaken for a dependable reserve-policy rule.

7 Limitations and future research

Two directions can deepen this evidence. First, a future transaction-level or institution-linked dataset could separate changes in borrower composition from changes in prices faced by comparable borrowers, and extend the evidence to dealer balance-sheet positions and Treasury settlement flows. This would permit a more direct test of

the channels suggested by the ON RRP and quarter-end associations. Second, genuine historical information vintages and more stress episodes would permit probability recalibration, independent policy-threshold validation and stronger causal inference. These extensions follow naturally from the present evidence, which establishes a detailed public-data foundation without claiming that observed associations identify the effects of specific interventions.

8 Conclusion

An apparently calm median overnight funding rate can coexist with meaningful differences in what different transactions pay. In US secured funding data from 2018 to 2026, the reserve-to-assets ratio is related to costs along the published rate distribution, and the middle of that distribution becomes more compressed when reserves are more abundant. ON RRP balances, quarter-end conditions and differences across repo-market benchmarks help explain why a single stock or a single rate is not a complete measure of funding conditions. The evidence supports measuring the distribution of funding pressure, not assuming that an extreme percentile always announces reserve scarcity first. It also demonstrates that a model's ability to rank stressful days does not guarantee accurate probabilities. A practical monitoring system should therefore combine several rate percentiles, both cash stocks, calendar information and forecast-calibration checks. This approach offers a more complete and transparent basis for interpreting short-term funding markets and assessing the resilience of financial intermediation.

References

Acharya VV, Skeie D (2011) A model of liquidity hoarding and term premia in inter-bank markets. Journal of Monetary Economics 58:436–447. https://doi.org/10.1016/j.jmoneco.2011.05.006

Afonso G, Cipriani M, Copeland A, Kovner A, La Spada G, Martin A (2021) The market events of mid-September 2019. Federal Reserve Bank of New York Economic Policy Review 27(2). https://www.newyorkfed.org/research/epr/2021/epr_2021_market-events_afonso

Afonso G, Giannone D, La Spada G, Williams JC (2025) Scarce, abundant, or ample? A time-varying model of the reserve demand curve. Federal Reserve Bank of New York Staff Report 1019, revised November 2025. https://www.newyorkfed.org/research/staff_reports/sr1019.html

Afonso G, Lagos R (2015) Trade dynamics in the market for federal funds. Econometrica 83:263–313. https://doi.org/10.3982/ECTA10586

Brier GW (1950) Verification of forecasts expressed in terms of probability. Monthly Weather Review 78:1–3. https://doi.org/10.1175/1520-0493(1950)078<0001:VOFEIT>2.0.CO;2

Brunnermeier MK, Pedersen LH (2009) Market liquidity and funding liquidity. Review of Financial Studies 22:2201–2238. https://doi.org/10.1093/rfs/hhn098

Copeland A, Martin A, Walker M (2014) Repo runs: evidence from the tri-party repo market. Journal of Finance 69:2343–2380. https://doi.org/10.1111/jofi.12205

Diebold FX, Mariano RS (1995) Comparing predictive accuracy. Journal of Business & Economic Statistics 13:253–263. https://doi.org/10.1080/07350015.1995.10524599

Duffie D (2020) Still the world’s safe haven? Redesigning the U.S. Treasury market after the COVID-19 crisis. Hutchins Center Working Paper 62. https://www.brookings.edu/articles/still-the-worlds-safe-haven/

Gneiting T, Raftery AE (2007) Strictly proper scoring rules, prediction, and estimation. Journal of the American Statistical Association 102:359–378. https://doi.org/10.1198/016214506000001437

Gorton G, Metrick A (2012) Securitized banking and the run on repo. Journal of Financial Economics 104:425–451. https://doi.org/10.1016/j.jfineco.2011.03.016

Hansen BE (1999) Threshold effects in non-dynamic panels: estimation, testing, and inference. Journal of Econometrics 93:345–368. https://doi.org/10.1016/S0304-4076(99)00025-1

Krishnamurthy A, Nagel S, Orlov D (2014) Sizing up repo. Journal of Finance 69:2381–2417. https://doi.org/10.1111/jofi.12168

Müller A, Paulick J (2026) "The devil is in the details, but so is salvation": different approaches in money market measurement. Journal of Financial Services Research 69:285–317. https://doi.org/10.1007/s10693-025-00450-4

Norden L, Weber M (2010) Funding modes of German banks: structural changes and their implications. Journal of Financial Services Research 38:69–93. https://doi.org/10.1007/s10693-010-0084-5

Politis DN, Romano JP (1994) The stationary bootstrap. Journal of the American Statistical Association 89:1303–1313. https://doi.org/10.1080/01621459.1994.10476870

Sunderam A (2015) Money creation and the shadow banking system. Review of Financial Studies 28:939–977. https://doi.org/10.1093/rfs/hhu083