

7536884-global-liquidity-regimes-and-cross-asset-downside-risk

SUMMARY: This paper examines global liquidity as a regime-dependent state variable for cross-asset risk rather than as a standalone return predictor. A funding-based latent liquidity factor is estimated from seventeen monthly indicators over 2005-2025, covering official balance sheets, secured funding and money-market conditions, cross-border dollar funding and private credit. A dynamic factor model extracts common funding variation, and a two-state Markov-switching model identifies persistent liquidity regimes distinguished primarily by conditional volatility.

METHOD: The study uses a dynamic factor model to estimate a funding-based latent liquidity factor from seventeen monthly indicators over the period 2005-2025. A two-state Markov-switching model classifies the factor into stable and turbulent liquidity regimes. The analysis evaluates whether these regimes condition realized returns, liquidity exposures, out-of-sample forecasts, and downside-risk outcomes across thirteen asset-return series. The regimes are re-identified in real time using a causal design with periodically re-estimated parameters and regime probabilities filtered forward.

KEY FINDINGS:

- Global liquidity is better understood as a regime-dependent state variable rather than a standalone return predictor.
- The latent factor summarises global liquidity conditions.
- Liquidity regimes are persistent and distinguished mainly by conditional volatility.
- Credit, real estate, and equities underperform in turbulent funding regimes.
- Real-time regime probabilities preserve the drawdown reduction.
- Global liquidity informs risk-state identification, not point forecasts of returns.

IMPLICATION FOR TRADING: The study suggests that a funding-based global liquidity factor can be used to allocate assets in real-time, reducing maximum drawdowns relative to benchmarks. The real-time factor correlates at 0.82 with its full-sample counterpart, and the allocation based on real-time probabilities limits the maximum drawdown to 7.0% over 2015-2025, against 17.8% for a 60/40 portfolio. Global liquidity is therefore more informative as a real-time conditioning variable for cross-asset downside risk than as a mechanical return-forecasting signal.

Global Liquidity Regimes and Cross-Asset Downside Risk

Aisel Salakhova*, Maksim O. Zharkovsky

HSE University, Moscow, Russian Federation

* Corresponding author.

E-mail addresses: salakhovaaysel09@gmail.com (A. Salakhova), jarkom2007@gmail.com (M.O. Zharkovsky).

ORCID: 0009-0005-2321-062X (A. Salakhova); 0009-0002-8289-7765 (M.O. Zharkovsky).

JEL classification: C24; C38; C53; E44; G11; G15

Keywords: Global liquidity; Funding liquidity; Dynamic factor model; Markov switching; Real-time identification; Downside risk

CRedit authorship contribution statement

Aisel Salakhova: Conceptualization, Methodology, Software, Formal analysis, Data curation, Writing – original draft, Writing – review & editing, Visualization. Maksim O. Zharkovsky: Conceptualization, Supervision, Writing – review & editing.

Declarations of interest: none.

Word count (main text): 4,762.

Global Liquidity Regimes and Cross-Asset Downside Risk

Abstract

This paper examines global liquidity as a regime-dependent state variable for cross-asset risk rather than as a standalone return predictor. A funding-based latent liquidity factor is estimated from seventeen monthly indicators over 2005–2025, covering official balance sheets, secured funding and money-market conditions, cross-border dollar funding and private credit. A dynamic factor model extracts common funding variation, and a two-state Markov-switching model identifies persistent liquidity regimes that are distinguished primarily by conditional volatility rather than by the level of the factor. Risk-sensitive assets, particularly credit, real estate and equities, realise materially lower returns in the turbulent regime, while the US dollar behaves defensively. The factor adds little out-of-sample point-forecasting power relative to autoregressive benchmarks with financial-condition controls; its information is concentrated in risk-state identification. A regime-based allocation rule reduces maximum drawdowns relative to 60/40, equal-weight and momentum benchmarks. To address look-ahead concerns, the regimes are re-identified in real time with a causal design in which parameters are re-estimated periodically and regime probabilities are filtered forward. The real-time factor correlates at 0.82 with its full-sample counterpart, and the allocation based on real-time probabilities limits the maximum drawdown to 7.0% over 2015–2025, against 17.8% for a 60/40 portfolio. Global liquidity is therefore more informative as a real-time conditioning variable for cross-asset downside risk than as a mechanical return-forecasting signal.

Highlights

- A funding-based latent factor summarises global liquidity conditions.
- Liquidity regimes are persistent and separated mainly by volatility.
- Credit, real estate and equities underperform in turbulent funding regimes.
- Real-time regime probabilities preserve the drawdown reduction.
- Liquidity informs risk-state identification, not point forecasts of returns.

JEL classification: C24; C38; C53; E44; G11; G15

Keywords

Global liquidity; Funding liquidity; Dynamic factor model; Markov switching; Real-time identification; Downside risk

1. Introduction

Global liquidity is widely used to interpret asset prices, risk appetite and cross-border capital flows [1–3], but its empirical measurement remains unsettled. Conventional proxies such as central bank balance sheets, monetary aggregates or credit growth capture only partial dimensions of the liquidity environment [4,5]. This limitation has become more relevant since the global financial crisis, as funding conditions increasingly reflect the joint behaviour of

official balance sheets, secured funding markets, cross-border dollar intermediation and private credit creation [6,7].

The empirical challenge is not only measurement but also interpretation. Individual liquidity indicators are noisy and exposed to structural breaks in operating frameworks, benchmark conventions and market structure. Composite factor models reduce this measurement problem by extracting the common variation of a broad funding panel [4]. The relation between liquidity and asset returns is also unlikely to be linear or stable across states: funding constraints, collateral values and intermediary risk-bearing capacity should matter most when funding conditions become unstable [8,9]. The informational content of global liquidity may therefore lie in regime dependence rather than in average-return predictability.

A regime-based reading raises a further, practical question. Latent states estimated with the full sample use information that was not available at the time, so a regime classification that is informative ex post may be unusable in real time. The same issue arises in real-time business-cycle analysis, where two-sided estimates of latent states are revised substantially at the end of the sample [10] and where the relevant test is whether causally filtered regime probabilities identify turning points with an acceptable delay [11,12]. A funding-state variable is useful for risk management only if it passes an analogous test.

This paper addresses both questions using a funding-based global liquidity factor estimated from seventeen monthly indicators over 2005–2025. A dynamic factor model extracts the common component of the panel, and a two-state Markov-switching model classifies the factor into stable and turbulent liquidity regimes. The analysis then evaluates whether these regimes condition realised returns, liquidity exposures, out-of-sample forecasts and downside-risk outcomes across thirteen asset-return series. Finally, the regimes are re-identified in real time: model parameters are estimated only on data available at each re-estimation date, and regime probabilities are filtered forward without look-ahead.

The results favour a risk-state interpretation. Liquidity regimes are highly persistent and are separated mainly by conditional volatility rather than by the level of the factor. Risk-sensitive assets realise weaker returns in the turbulent regime, with the clearest evidence for credit, real estate and equities, while the US dollar behaves defensively. The liquidity factor does not generate systematic out-of-sample point-forecasting gains relative to autoregressive models augmented with financial-condition controls. Its economic content is instead concentrated in downside-risk identification: a rule that shifts weight towards a defensive basket as the turbulent-regime probability rises reduces maximum drawdowns relative to benchmark strategies. This result survives real-time identification. With periodic re-estimation, the real-time factor correlates at 0.82 with the full-sample factor, and the real-time allocation attains a maximum drawdown of -7.0% over 2015–2025, close to the -6.3% obtained with full-sample parameters and well below the -17.8% of a 60/40 portfolio.

The paper makes three contributions. First, it shifts the empirical role of global liquidity from linear return prediction to regime-dependent risk conditioning; relative to factor-based liquidity measurement [4], the contribution is not a new index but a change in how the index is used. Second, it documents cross-asset return heterogeneity across funding regimes estimated from

funding data rather than from asset returns, which distinguishes the analysis from regime-switching allocation studies [13–15]. Third, it shows how a latent funding state can be identified in real time from an unbalanced and structurally changing panel, and why a naive real-time benchmark understates the reliability of this identification: the low correlation obtained from a recursively standardised principal-component factor reflects the change of estimator and the evolving composition of the panel rather than an intrinsic failure of real-time identification.

Section 2 presents the methodology, Section 3 describes the data and Section 4 reports the empirical results. Section 5 discusses robustness, Section 6 compares the findings with the related literature and Section 7 concludes.

2. Methodology

2.1 Global liquidity factor

The common component of the liquidity panel is estimated with a one-factor dynamic factor model. For each standardised liquidity indicator $x_{i,t}$,

$$x_{i,t} = \lambda_i f_t + \varepsilon_{i,t}, \quad (1)$$

where f_t is the latent global liquidity factor, λ_i is the factor loading and $\varepsilon_{i,t}$ is an idiosyncratic component. The factor and idiosyncratic terms follow first-order autoregressive processes, allowing common funding dynamics to be separated from series-specific persistence. The model is estimated by maximum likelihood in state-space form, and the Kalman filter and smoother handle the unbalanced panel, so that indicators with short or interrupted histories contribute when they are observed [16,17]. The sign of the factor is normalised so that higher values denote more accommodative funding conditions. Two versions of the factor are used: the smoothed factor conditions on the full sample and serves for ex post regime classification and the conditional-return analysis, whereas the filtered factor conditions only on data up to each month, given the estimated parameters, and underlies the allocation exercise. Alternative factor constructions are evaluated in Section 5.

2.2 Regime identification

Liquidity regimes are identified by applying a two-state Markov-switching model to the estimated factor [18]:

$$f_t = \mu(S_t) + e_t, \quad e_t \sim N(0, \sigma^2(S_t)), \quad S_t \in \{0,1\}. \quad (2)$$

Both the mean and the variance switch across regimes, and the latent state S_t follows a first-order Markov chain with transition probabilities p_{ij} . Persistence is summarised by the expected duration of each regime, $D_i = (1 - p_{ii})^{-1}$ [18]. The low- and high-volatility states are labelled stable and turbulent liquidity regimes. The turbulent label denotes elevated instability in the common funding component, not a crisis classification for every month in that state. Ex post classification assigns a month to the turbulent regime when its smoothed probability exceeds 0.5; the allocation exercise uses filtered probabilities.

2.3 Real-time identification

Full-sample estimation introduces look-ahead at three points: the moments used to standardise the indicators, the parameters of the factor and regime models, and the two-sided smoothing of the latent states. Following the real-time evaluation design of the business-cycle dating literature [11,12], the regimes are re-identified with a procedure that removes all three. At a re-estimation date τ , each indicator is standardised with moments computed from data up to τ , and an indicator enters the panel once it has at least 36 monthly observations. The factor model and the Markov-switching model are estimated on data up to τ . For every subsequent month t until the next re-estimation date, the Kalman filter and the Hamilton filter are run forward with the parameters fixed at their τ estimates, which yields the real-time turbulent-regime probability

$$\hat{p}_t = \Pr(S_t = 1 \mid x_1, \dots, x_t; \theta(\tau)), \quad \tau < t \leq \tau + k, \quad (3)$$

where $\theta(\tau)$ collects the parameters estimated on data up to τ and k is the re-estimation interval. The probability uses no information after month t . The sign of the factor and the labelling of the regimes follow the same conventions as in the full-sample model, which prevents label switching across re-estimations.

Two designs are considered. In Design A, the parameters are estimated once on 2005–2014 and held fixed over the entire evaluation period 2015–2025. In Design B, the parameters are re-estimated every 36 months on the expanding sample, so that indicators with shorter histories enter as data accumulate; intervals of 24 and 48 months are reported as sensitivity checks. The evaluation window coincides with the out-of-sample forecasting window. Real-time identification is assessed by the share of months in which the real-time regime call ($\hat{p}_t > 0.5$) agrees with the ex post classification, by the correlation of the real-time and full-sample factors in levels and in first differences, and by the lag that maximises the cross-correlation of the two regime sequences.

2.4 Return tests, forecasting and allocation

The liquidity regimes are linked to cross-asset returns and risk through four exercises. First, realised monthly returns are compared across stable and turbulent regimes, with inference based on regressions of returns on a turbulent-regime dummy with Newey–West standard errors (three lags) and on Welch two-sample tests. Second, regime-dependent liquidity exposures are estimated from forward-return regressions of cumulative returns over horizons of one, three and six months on the liquidity factor, a turbulent-regime indicator, their interaction and financial-condition controls. Third, expanding-window out-of-sample forecasts at the three-month horizon over 2015–2025 compare an autoregressive model, the same model augmented with financial-condition controls and a liquidity-augmented specification [19], assessed by the incremental out-of-sample R^2 [20], Clark–West tests [21] and directional accuracy [22] against the autoregressive-control and naive-positive benchmarks.

Fourth, turbulent-regime probabilities drive an allocation rule between a fixed risk-on basket (the S&P 500, MSCI World and Nasdaq-100 equity indices, investment-grade and high-yield credit, and REITs) and a defensive basket (aggregate bonds, gold and the US dollar). Both baskets are

equally weighted and defined ex ante from standard asset characteristics [13,14]. The defensive weight equals the turbulent-regime probability lagged by one month, so that the portfolio return is

$$r_t^{\text{GL}} = (1 - w_t) r_t^{\text{RO}} + w_t r_t^{\text{D}} - c |w_t - w_{t-1}|, \quad w_t = \hat{p}_{t-1}, \quad (4)$$

where r_t^{RO} and r_t^{D} are the returns of the risk-on and defensive baskets and c is a transaction cost of 10 basis points per unit of one-way turnover. In the full-sample exercise, $\hat{p}_t$ is the filtered probability from the full-sample model; in the real-time exercise, it is the probability in Eq. (3). The rule is compared with a static 60/40 mix of the two baskets, an equal-weight ($1/N$) portfolio of the nine allocation assets and a cross-sectional momentum strategy. Differences in Sharpe ratios are tested with the studentised circular block bootstrap of Ledoit and Wolf [23]. The exercise asks whether regime probabilities reduce downside exposure relative to benchmark strategies, not whether they generate tradable alpha [15,24].

3. Data

3.1 Global liquidity panel

The analysis uses an unbalanced monthly panel from January 2005 to December 2025. The liquidity dataset contains seventeen indicators grouped into four funding blocks: official liquidity (four series), secured funding and money-market conditions (six), cross-border US dollar funding (three) and private credit (four). Official liquidity is proxied by the balance sheets of the Federal Reserve, the European Central Bank, the Bank of Japan and the People's Bank of China. The secured-funding block contains policy-floor spreads, the TED spread and its post-LIBOR successor, tri-party repo volumes and overnight reverse repo (ON RRP) usage. The cross-border block contains the euro and yen cross-currency bases and the FRA–OIS spread, which measure dollar funding pressure in FX swap and interbank markets [25]. The private-credit block covers bank credit, commercial paper, money-market fund assets and credit to the private non-financial sector. Table 1 summarises the data structure, and Appendix Table A1 reports variable definitions, sources, transformations and coverage.

[Insert Table 1 about here]

Balance-sheet and credit aggregates enter as logarithmic first differences, whereas rates, spreads, basis measures and ON RRP usage are retained in levels. This separates quantity-based liquidity growth from price-based funding conditions and avoids unstable log changes for series with long spans of zero or near-zero values. Balance sheets are converted to US dollars at month-end exchange rates, and quarterly series are interpolated to monthly frequency. All series are standardised and winsorised at ± 4 standard deviations. Benchmark transitions are handled explicitly: pre- and post-LIBOR funding-spread proxies are kept as separate variables, and the cross-currency basis series are spliced only after subperiod standardisation. Missing observations, including the periods before an indicator becomes available, are handled within the state-space estimation.

Coverage is uneven. The cross-currency bases start in October 2013, repo volumes in August 2014 and the FRA–OIS spread in November 2015, and the post-reform secured-funding spreads are available only from 2021–2022. This evolving composition is immaterial for full-sample estimation, but it matters for real-time identification, because several of the indicators that carry the most information about dollar funding conditions after 2015 are unavailable, or have short histories, at the start of the evaluation window.

3.2 Asset returns and controls

The return panel consists of monthly US dollar log returns for thirteen asset-return series: four equity indices (S&P 500, MSCI World, MSCI Emerging Markets and Nasdaq-100), US aggregate bonds, US investment-grade and high-yield corporate bonds, a broad commodity index, gold, US REITs, the broad trade-weighted US dollar, an EM currency index and EM sovereign credit. Four financial-condition controls — the VIX, the two-year real Treasury yield, the 10-year minus 2-year term spread and the high-yield option-adjusted spread — enter the return and forecasting specifications to separate liquidity information from broader macro-financial conditions. Definitions and sources are reported in Appendix Table A2.

4. Empirical results

4.1 Liquidity factor and regimes

Figure 1 plots the liquidity factor and the smoothed probability of the turbulent regime, and Panel A of Table 2 reports the regime estimates. The factor is strongly persistent, consistent with funding conditions evolving as a slow-moving state rather than as a sequence of transitory shocks. Its largest movements occur around the global financial crisis, the March 2020 shock and the 2022–2023 policy normalisation. The Markov-switching model separates the sample into two persistent regimes, with persistence probabilities of 0.961 and 0.983 and expected durations of 25.6 and 59.8 months in the stable and turbulent states. The classification is driven by volatility: the conditional variance rises from 0.052 in the stable regime to 1.479 in the turbulent regime, a ratio of about 28, while the regime means are close to each other at -0.050 and 0.026 .

[Insert Figure 1 about here]

[Insert Table 2 about here]

The turbulent regime is therefore not a crisis dummy but a persistent high-instability funding state that accounts for about two-thirds of the sample (166 of 252 months). The smoothed probabilities identify three long turbulent spells — from the start of the sample to mid-2009, from late 2014 to mid-2019 and from 2020 to early 2024 — and a short spell from mid-2025. These spells contain the acute stress episodes but extend well beyond them, which is consistent with the interpretation of the regime as a state of the funding system rather than as a sequence of crises. The regime evidence supports reading the factor as a funding-state variable rather than as a linear easing indicator.

4.2 Conditional cross-asset returns

Panel B of Table 2 and Figure 2 show that the return effects of liquidity regimes concentrate in funding-sensitive assets. Mean monthly S&P 500 returns fall from 1.52% in the stable regime to 0.53% in the turbulent regime, high-yield credit from 0.99% to 0.29% and REITs from 1.20% to -0.32%. The largest and most statistically reliable gaps are in credit and real estate: under HAC inference, the differentials are significant at the 5% level for investment-grade credit, high yield and REITs, and at the 10% level for the S&P 500 and EM credit. Welch tests give the same pattern of significance (Appendix Table A3). Twelve of the thirteen differentials are positive; the exception is the US dollar.

[Insert Figure 2 about here]

This ordering is consistent with a funding-based mechanism. Credit and real estate are directly exposed to refinancing conditions, spread compression and intermediary balance-sheet capacity, while equities also weaken in turbulent regimes, though less uniformly across regions and styles. Commodities and gold show small, statistically insignificant differences, indicating weaker links to the funding state over the sample.

The US dollar is the only asset with a higher mean return in the turbulent regime. This is consistent with its dual role as funding currency and reserve asset [7,26,27]: when funding conditions become unstable, dollar exposure behaves defensively rather than cyclically. Because the composition of the allocation baskets is fixed ex ante (Section 2.4), this pattern is not an ex post justification of the defensive basket; it reinforces the channel-specific nature of the evidence. A natural concern is that differentials of this size are driven by the two acute crises in the sample. Section 5 shows that excluding the global financial crisis and the COVID-19 shock attenuates the differentials but leaves twelve of the thirteen positive and four significant, all in fixed income.

Regime-dependent liquidity betas provide complementary but weaker evidence (Appendix Table A4). The strongest patterns appear at the three- and six-month horizons rather than at one month: in the turbulent regime, the liquidity slope at the three-month horizon is positive and significant for the S&P 500, MSCI World, Nasdaq-100 and REITs. This is consistent with liquidity effects being transmitted through slower adjustments in leverage, risk premia and balance-sheet constraints. Because the beta estimates are less uniform than the conditional-return results, they are treated as supporting evidence.

4.3 Forecasting evidence

The regime results do not imply strong point predictability. Panel A of Table 3 shows that, over the 2015–2025 out-of-sample window, adding the filtered liquidity factor to an autoregressive model with financial-condition controls produces a negative incremental out-of-sample R^2 for ten of the thirteen assets. Clark–West tests are significant at the 10% level for only two assets. Directional accuracy is more favourable, exceeding 70% for six assets and improving on the autoregressive-control benchmark for ten, but it beats a naive positive-return forecast for only six. Asset-level results are reported in Appendix Table A5. The factor therefore carries some directional information but no systematic point-forecasting gain, in line with the broader difficulty of predicting returns out of sample [19,20].

[Insert Table 3 about here]

4.4 Allocation and real-time identification

The allocation exercise evaluates a narrower implication of the same information set: whether the probability of the turbulent regime identifies periods of elevated downside risk. Panel B of Table 3 reports the full-sample results. Net of transaction costs, the regime-based rule earns an annualised return of 9.25% with volatility of 7.63%, a Sharpe ratio of 1.21 and a maximum drawdown of −11.8%. The maximum drawdowns of the 60/40, equal-weight and momentum benchmarks are −27.3%, −31.0% and −34.9%. The rule has an annualised one-way turnover of about 35%, and its performance is essentially insensitive to costs: the Sharpe ratio moves from 1.22 to 1.20 as costs rise from 0 to 50 basis points. The Sharpe improvement over the benchmarks ranges from 0.37 to 0.43 and is significant at the 10% level against the one-sided alternative under Ledoit–Wolf tests, but not at conventional two-sided levels (Appendix Table A6). Figure S1 in the Appendix plots the cumulative wealth and drawdown paths. The results are therefore best read as evidence of drawdown mitigation rather than alpha generation [28].

These full-sample results use filtered probabilities, but the parameters of the factor and regime models and the standardisation of the indicators are estimated over the full sample. Table 4 therefore reports the real-time identification of Section 2.3 over 2015–2025.

[Insert Table 4 about here]

Design A, which fixes all parameters at their 2005–2014 estimates, performs poorly: the real-time regime call agrees with the ex post classification in only 35% of months, and the real-time factor correlates at 0.29 with the full-sample factor. The reason is the composition of the panel. At the end of 2014 only ten of the seventeen indicators have sufficient history to enter the model, so the cross-currency bases, the FRA–OIS spread, repo volumes and the post-reform secured-funding spreads — which carry much of the information about dollar funding conditions after 2015 — are excluded by construction. A model frozen in 2014 cannot track a funding system whose relevant indicators it does not observe.

Design B, which re-estimates the parameters every 36 months and lets indicators enter as their histories accumulate, recovers the funding state. The number of indicators in the model rises from ten to fifteen across re-estimations; the real-time factor correlates at 0.82 with the full-sample factor in levels and at 0.69 in first differences, and the real-time regime call agrees with the ex post classification in 63% of months. The disagreement follows the pattern documented in real-time recession dating [11,12]. Real-time calls lag the ex post regime by about two months, and filtered probabilities are more conservative than smoothed ones: the real-time rule classifies about half of the evaluation-window months as turbulent, against roughly four-fifths ex post. As with real-time estimates of the output gap [10], the gap between one-sided and two-sided estimates is largely a property of the estimation problem — the smoother uses future information to date the regime — rather than evidence that the state cannot be identified in real time.

The economic implication is what matters for risk management. Panel B of Table 4 shows that the allocation based on Design B probabilities earns a Sharpe ratio of 1.26 with a maximum

drawdown of -7.0% over 2015–2025, compared with 1.19 and -6.3% for the same rule using full-sample parameters, and with maximum drawdowns of -17.8% , -19.2% and -15.1% for the 60/40, equal-weight and momentum benchmarks. The more conservative real-time classification therefore does not translate into larger drawdowns, which suggests that the real-time rule captures the episodes that dominate the downside of the risk-on basket. Design A, despite its weak classification, still limits the maximum drawdown to -15.6% , but with a Sharpe ratio of 0.99. Section 5 shows that the Design B result is robust to the re-estimation interval and to transaction costs.

Taken together, the real-time evidence changes the interpretation of the allocation exercise. The drawdown reduction is not an artefact of full-sample estimation: a causal procedure that uses only information available at each date delivers a comparable reduction, provided that the model is re-estimated as the funding panel evolves.

5. Robustness

Appendix Table A7 summarises the robustness checks. The first group concerns the construction of the factor. Re-estimating the model on the non-winsorised panel leaves the factor essentially unchanged (correlations of 1.00 in levels and 0.98 in first differences) and preserves the high-volatility regime. Removing ON RRP usage, excluding the central-bank balance-sheet block and replacing the dynamic factor with the first principal component each yield factors that remain positively correlated with the baseline (0.70, 0.82 and 0.64) and retain a high-volatility state, but with weaker variance separation. The balance-sheet exclusion in particular compresses the turbulent-regime share to 31%, suggesting that official liquidity is a material contributor to, though not the sole driver of, the regime structure.

Because ON RRP usage has large swings, a natural concern is that the factor merely tracks this series. The loadings in Appendix Table A8 do not support this view. The cross-border dollar-funding block has the highest average communality (0.22), followed by secured funding and money-market conditions (0.17), and ON RRP usage ranks third among the seventeen indicators (0.32), behind the SOFR–T-bill spread and the EUR/USD basis (0.35 each). Re-entering ON RRP as a share of the Federal Reserve balance sheet, which preserves its information while removing its dollar scale, leaves the factor essentially unchanged (correlation of 1.00; 96% regime agreement) and retains the conditional-return ordering. The weaker results when ON RRP is dropped therefore reflect a loss of information rather than dependence on the scale of a single series.

The second group concerns the regimes and the return evidence. A three-state Markov-switching model recovers a high-volatility state that overlaps the baseline turbulent regime in 68% of months, so the stable–turbulent interpretation is not an artefact of imposing two states. Excluding the global financial crisis (December 2007–June 2009) and the COVID-19 shock (February–May 2020) attenuates the conditional-return differentials but leaves twelve of the thirteen positive and four significant at the 10% level — aggregate bonds, investment-grade and high-yield credit and EM credit (Appendix Table A9). The differentials are therefore not an artefact of acute crisis episodes and are most durable in funding-sensitive fixed income. The medium-horizon pattern in

liquidity betas persists at the twelve-month horizon and under alternative control sets, and the weak point-forecasting result is unchanged under alternative specifications.

The third group concerns the allocation exercise. Threshold and capped variants of the rule retain maximum drawdowns of -11.9% and -12.9% over the full sample, against -27.3% for the 60/40 benchmark. For the real-time rule, Appendix Table A10 varies the re-estimation interval and transaction costs. With re-estimation every 24 months, regime agreement rises to 67% and the maximum drawdown is -7.8% ; with a 48-month interval, agreement falls to 43%, again reflecting the delayed entry of new indicators, but the maximum drawdown of -10.8% remains well below the -17.8% of the 60/40 portfolio over the same window. Real-time turnover is higher than in the full-sample case, at 68–89% a year, yet the Sharpe ratio of the 36-month design declines only from 1.28 to 1.20 as costs rise from 0 to 50 basis points, and its maximum drawdown is unchanged.

Finally, an earlier real-time diagnostic reported in Table A7 (check 9) replaced the dynamic factor model with a principal-component factor computed from a recursively standardised panel, with missing observations filled by series means, and obtained a correlation of only 0.36 with the full-sample filtered factor. This diagnostic combines two changes: recursive standardisation and a switch of estimator that, even over the full sample, lowers the correlation with the baseline factor to 0.64 (check 4), because principal components cannot exploit the state-space treatment of missing data. The procedure of Section 2.3 isolates the effect of look-ahead and shows that it is moderate once the model is re-estimated as data accumulate.

6. Comparison with the literature

The findings relate to three strands of literature. The first concerns the measurement of global liquidity. Eickmeier et al. [4] decompose global liquidity into global monetary policy, credit supply and credit demand factors, and the global financial cycle literature documents a common factor in risky asset prices worldwide that is closely linked to US monetary conditions [1,2]. The factor in this paper is narrower, being built from funding-market indicators, and its information lies in its conditional volatility rather than in its level: the regimes are separated by a variance ratio of about 28 with nearly identical means. This supports the view that funding liquidity matters through episodes of instability, as in the liquidity spirals of Brunnermeier and Pedersen [8] and the procyclical intermediary leverage documented by Adrian and Shin [6], rather than through a linear relation with the stance of monetary policy.

The second strand concerns regime switching in asset returns. Ang and Bekaert [13] and Guidolin and Timmermann [14] identify regimes in which risky assets combine lower returns with higher volatility and in which regime-dependent allocation improves on static portfolios, and Kritzman et al. [15] show that regime signals can improve dynamic strategies. The present paper reaches a similar conclusion with regimes estimated from funding data rather than from the returns themselves, which reduces the risk that the regimes merely restate the return patterns they are used to explain. The defensive behaviour of the dollar in the turbulent regime is consistent with its safe-asset and funding-currency roles [7,26], and the high communalities of the cross-currency bases are consistent with the evidence that deviations from covered interest parity have become a persistent indicator of dollar funding constraints since the global financial crisis [25].

The third strand concerns predictability and real-time inference. The forecasting results align with the evidence that return predictors rarely deliver systematic out-of-sample gains [19], while even small statistical gains can be economically meaningful [20]. The allocation results fit this pattern: the liquidity factor fails as a point predictor but improves the downside profile of a simple rule, consistent with studies that evaluate dynamic allocation rules by their risk properties rather than by alpha [24,28]. The real-time evidence parallels the business-cycle dating literature, in which Markov-switching probabilities identify turning points in real time with some delay [11,12] and two-sided estimates of latent states are revised substantially at the end of the sample [10]. The liquidity regimes behave in the same way: real-time calls lag the ex post classification by about two months and are more conservative, but they preserve most of the economic value of the regime information.

7. Conclusion

This paper evaluates global liquidity as a regime-dependent funding-state variable rather than as a linear return predictor. The evidence is asymmetric across the two roles that liquidity is often expected to play. As a characterisation of the funding environment, the construct is informative: the estimated regimes are persistent, separated mainly by funding volatility rather than by the level of liquidity, and associated with economically coherent cross-asset return differentials concentrated in credit, real estate and equities. As a return predictor, it is weak: the factor delivers no systematic out-of-sample point-forecasting gain over autoregressive-control benchmarks.

The economic value of liquidity information therefore lies in risk-state identification. Turbulent-regime probabilities reduce drawdowns relative to benchmark strategies, and the reduction survives a real-time design in which the model uses only information available at each date and is re-estimated periodically. The Sharpe improvement is positive but not statistically decisive at conventional two-sided levels, so the evidence supports the use of global liquidity as a conditioning variable for cross-asset risk management rather than as a source of alpha.

Three qualifications bound the contribution. First, the cross-asset evidence is channel-specific, applying most clearly to assets exposed to leverage, refinancing risk and intermediary balance-sheet constraints. Second, real-time identification requires periodic re-estimation: a model frozen early in the sample fails because the indicators that carry information about the post-2015 funding system are not yet available, and real-time regime calls lag and understate turbulence relative to the ex post classification. Third, the evaluation window contains a limited number of distinct funding episodes, so the real-time allocation results should be confirmed as further episodes accumulate. Higher-frequency funding indicators and applications to individual markets would sharpen the real-time signal and test the external validity of the regime structure.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Data availability

The data were obtained from Bloomberg, the Federal Reserve Economic Data (FRED) database and the Bank for International Settlements (BIS). The FRED and BIS series are publicly available; the Bloomberg data are available to Bloomberg subscribers and cannot be redistributed by the authors. The code used to construct the factor, estimate the models and produce the results is available from the authors on reasonable request.

Declaration of generative AI and AI-assisted technologies in the writing process

During the preparation of this work the authors used Claude (Anthropic) in order to improve the clarity, structure and language of the manuscript and to assist with editing. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

References

- [1] H. Rey, Dilemma not trilemma: the global financial cycle and monetary policy independence, in: Proceedings of the Jackson Hole Economic Policy Symposium, Federal Reserve Bank of Kansas City, 2013, pp. 285–333.
- [2] S. Miranda-Agrippino, H. Rey, U.S. monetary policy and the global financial cycle, *Review of Economic Studies* 87 (6) (2020) 2754–2776.
- [3] O. Jeanne, D. Sandri, Global financial cycle and liquidity management, *Journal of International Economics* 146 (2023) 103736.
- [4] S. Eickmeier, L. Gambacorta, B. Hofmann, Understanding global liquidity, *European Economic Review* 68 (2014) 1–18.
- [5] L.S. Goldberg, Global liquidity: drivers, volatility and toolkits, *IMF Economic Review* 72 (1) (2024) 1–31.
- [6] T. Adrian, H.S. Shin, Liquidity and leverage, *Journal of Financial Intermediation* 19 (3) (2010) 418–437.
- [7] V. Bruno, H.S. Shin, Cross-border banking and global liquidity, *Review of Economic Studies* 82 (2) (2015) 535–564.
- [8] M.K. Brunnermeier, L.H. Pedersen, Market liquidity and funding liquidity, *Review of Financial Studies* 22 (6) (2009) 2201–2238.

- [9] C. Borio, H. Zhu, Capital regulation, risk-taking and monetary policy: a missing link in the transmission mechanism?, *Journal of Financial Stability* 8 (4) (2012) 236–251.
- [10] A. Orphanides, S. van Norden, The unreliability of output-gap estimates in real time, *Review of Economics and Statistics* 84 (4) (2002) 569–583.
- [11] M. Chauvet, J. Piger, A comparison of the real-time performance of business cycle dating methods, *Journal of Business & Economic Statistics* 26 (1) (2008) 42–49.
- [12] J.D. Hamilton, Calling recessions in real time, *International Journal of Forecasting* 27 (4) (2011) 1006–1026.
- [13] A. Ang, G. Bekaert, International asset allocation with regime shifts, *Review of Financial Studies* 15 (4) (2002) 1137–1187.
- [14] M. Guidolin, A. Timmermann, Asset allocation under multivariate regime switching, *Journal of Economic Dynamics and Control* 31 (11) (2007) 3503–3544.
- [15] M. Kritzman, S. Page, D. Turkington, Regime shifts: implications for dynamic strategies, *Financial Analysts Journal* 68 (3) (2012) 22–39.
- [16] C. Doz, D. Giannone, L. Reichlin, A quasi-maximum likelihood approach for large, approximate dynamic factor models, *Review of Economics and Statistics* 94 (4) (2012) 1014–1024.
- [17] M. Bańbura, M. Modugno, Maximum likelihood estimation of factor models on datasets with arbitrary pattern of missing data, *Journal of Applied Econometrics* 29 (1) (2014) 133–160.
- [18] J.D. Hamilton, A new approach to the economic analysis of nonstationary time series and the business cycle, *Econometrica* 57 (2) (1989) 357–384.
- [19] I. Welch, A. Goyal, A comprehensive look at the empirical performance of equity premium prediction, *Review of Financial Studies* 21 (4) (2008) 1455–1508.
- [20] J.Y. Campbell, S.B. Thompson, Predicting excess stock returns out of sample: can anything beat the historical average?, *Review of Financial Studies* 21 (4) (2008) 1509–1531.
- [21] T.E. Clark, K.D. West, Approximately normal tests for equal predictive accuracy in nested models, *Journal of Econometrics* 138 (1) (2007) 291–311.
- [22] M.H. Pesaran, A. Timmermann, A simple nonparametric test of predictive performance, *Journal of Business & Economic Statistics* 10 (4) (1992) 461–465.
- [23] O. Ledoit, M. Wolf, Robust performance hypothesis testing with the Sharpe ratio, *Journal of Empirical Finance* 15 (5) (2008) 850–859.

[24] M.D. Braga, C.R. Nava, M.G. Zoia, Kurtosis-based vs volatility-based asset allocation strategies: do they share the same properties? A first empirical investigation, *Finance Research Letters* 54 (2023) 103797.

[25] W. Du, A. Tepper, A. Verdelhan, Deviations from covered interest rate parity, *Journal of Finance* 73 (3) (2018) 915–957.

[26] Z. Jiang, A. Krishnamurthy, H. Lustig, Foreign safe asset demand and the dollar exchange rate, *Journal of Finance* 76 (3) (2021) 1049–1089.

[27] M. Tronzano, Safe-haven currencies as defensive assets in global stock portfolios: a reassessment of the empirical evidence (1999–2022), *Journal of Risk and Financial Management* 16 (5) (2023) 273.

[28] G. Bellu, C. Conversano, Protected adaptive asset allocation, *Finance Research Letters* 32 (2020) 101095.

Table 1

Structure of the data: liquidity panel, asset universe and controls

Block	No. of series	Representative indicators	Transformation	Captures
Official liquidity	4	Fed, ECB, BoJ, PBoC balance sheets	USD conversion; log differences	Official balance-sheet liquidity
Secured funding and money-market conditions	6	Policy-spread proxies, TED, SOFR–T-bill, ON RRP	Levels, including ON RRP	Short-term funding stress and liquidity absorption
Cross-border US dollar funding	3	Cross-currency bases (EUR, JPY), FRA–OIS spread	Levels; spliced post-standardisation	Dollar funding pressure
Private credit and bank-based liquidity	4	US private credit and market-based funding aggregates	Log differences	Private credit creation
Asset returns	13	Equities, fixed income, credit, commodities, REITs, FX	Monthly USD log returns	Dependent variables and allocation universe
Financial-condition controls	4	VIX, real 2Y yield, term spread, HY OAS	Levels	Controls

Notes: The table summarises the monthly dataset, January 2005 to December 2025. The global liquidity panel comprises the first four blocks (seventeen indicators); the remaining rows list the dependent asset-return series and the financial-condition controls. Balance-sheet and credit aggregates enter as logarithmic first differences; rates, spreads, basis measures and overnight reverse repo (ON RRP) usage enter in levels. ON RRP is kept in levels because long spans of zero or near-zero usage make log differences mechanically unstable. All series are standardised and winsorised at ± 4 standard deviations before factor estimation; missing observations are handled within the state-space model. Full definitions, sources, tickers and coverage are reported in Appendix Tables A1 (liquidity indicators) and A2 (returns and controls).

Table 2

Liquidity-regime characteristics and conditional cross-asset returns

Panel A: Two-state Markov-switching regime estimates

Statistic	Stable regime	Turbulent regime
Conditional mean	−0.0505	0.0260
Conditional variance	0.0521	1.4793
Persistence	0.9610	0.9833
Expected duration (months)	25.6	59.8
Observations	86	166
Share of sample	0.34	0.66

Panel B: Mean monthly returns by regime (percent)

Asset	Stable	Turbulent	Stable – Turbulent	HAC t-stat
S&P 500	1.52	0.53	1.00	1.91*
MSCI World	1.27	0.41	0.86	1.57
MSCI EM	0.99	0.37	0.63	0.81
Nasdaq 100	1.81	0.89	0.92	1.48
US Agg	0.41	0.19	0.22	1.54
US Corp IG	0.64	0.18	0.46	2.31**
US High Yield	0.99	0.29	0.69	2.17**
BCOM	0.27	−0.07	0.34	0.56
Gold	1.14	0.83	0.31	0.52
REIT	1.20	−0.32	1.53	2.14**
Dollar	−0.03	0.14	−0.17	−0.77
EM FX	0.31	0.14	0.17	0.73
EM Credit	0.81	0.28	0.53	1.79*

Notes: Estimates are from a two-state Markov-switching model fitted by maximum likelihood to the smoothed global liquidity factor (standardised units), monthly over 2005–2025. In Panel A, "persistence" is the probability of remaining in the regime (the diagonal of the transition matrix) and the expected duration equals $(1 - p_{ii})^{-1}$ months; the low- and high-volatility states are labelled stable and turbulent. In Panel B, returns are monthly log returns in percent and the difference is stable minus turbulent; HAC t-statistics use Newey–West standard errors with maxlags = 3. ***, **, and * denote significance at the 1%, 5% and 10% levels, respectively. Welch and HAC inference for all thirteen assets is reported in Appendix Table A3.

Table 3

Out-of-sample forecasting performance and the regime-based allocation

Panel A: Out-of-sample forecasting summary (13 assets)

Statistic	Count
Negative $\Delta R^2(\text{OOS})$ vs AR + controls	10/13
Positive $\Delta R^2(\text{OOS})$ vs AR + controls	3/13
Clark–West significant at 10%	2/13
Directional accuracy > 70%	6/13
Directional accuracy improves on AR + controls	10/13
Directional accuracy improves on the naive-positive benchmark	6/13

Panel B: Portfolio performance, 2005–2025

Strategy	Ann. return	Ann. volatility	Sharpe	Max drawdown
GL Continuous (10 bps)	9.25	7.63	1.21	–11.80
60/40	6.75	8.00	0.84	–27.32
1/N	6.88	8.74	0.79	–30.96
Momentum	9.64	12.11	0.80	–34.93

Notes: Panel A summarises expanding-window out-of-sample forecasts of monthly returns over 2015–2025 at the three-month horizon, using the filtered global liquidity factor. The models are nested: an AR(1); an AR(1) augmented with financial-condition controls (AR + controls); and the same specification augmented with the liquidity factor (GL + controls). $\Delta R^2(\text{OOS})$ is the Campbell–Thompson incremental out-of-sample R^2 of GL + controls relative to AR + controls; "naive positive" denotes an always-positive return forecast. Counts are out of the thirteen asset-return series. Panel B reports annualised performance (percent, except the Sharpe ratio) for the continuous regime-based rule ("GL Continuous"), which shifts weight toward a fixed defensive basket as the lagged filtered probability of the turbulent regime rises, net of 10 bps transaction costs, against 60/40, equal-weight (1/N) and cross-sectional momentum benchmarks. Parameters and standardisation are estimated over the full sample; real-time results are reported in Table 4. Full asset-level forecasts (Table A5), regime-dependent liquidity betas (Table A4) and allocation diagnostics — Ledoit–Wolf Sharpe-difference tests (Table A6) — are reported in the Appendix.

Table 4

Real-time regime identification and allocation, 2015–2025

Panel A: Real-time identification quality

Design	Regime agreement (%)	Factor corr., levels	Lag	Factor corr., differences
A: parameters fixed at 2005–2014 estimates	34.8	0.29	–1	0.30
B: re-estimation every 36 months	62.9	0.82	2	0.69

Panel B: Allocation performance, 2015–2025

Strategy	Ann. return	Ann. volatility	Sharpe	Max drawdown
GL real-time, Design B	7.24	5.74	1.26	–7.03
GL real-time, Design A	9.27	9.35	0.99	–15.58
GL, full-sample parameters	7.17	6.02	1.19	–6.33
60/40	6.91	7.52	0.92	–17.81
1/N	7.12	8.22	0.87	–19.16
Momentum	9.82	11.35	0.87	–15.10

Notes: The evaluation window is January 2015 to December 2025. In Design A, the standardisation moments and the parameters of the dynamic factor and Markov-switching models are estimated once on 2005–2014, and the Kalman and Hamilton filters are run forward with fixed parameters. In Design B, the parameters are re-estimated every 36 months on the expanding sample (at the end of 2014, 2017, 2020 and 2023), indicators enter once they have 36 observations, and probabilities are filtered forward between re-estimations. No quantity at month t uses data after t . Regime agreement is the share of months in which the real-time call (turbulent-regime probability above 0.5) coincides with the ex post classification based on full-sample smoothed probabilities. Factor correlations compare the real-time filtered factor with the full-sample smoothed factor. The lag is the shift, in months, that maximises the cross-correlation of the real-time and ex post regime sequences; positive values indicate that real-time calls lag. The correlation between Design B real-time probabilities and full-sample filtered probabilities is 0.20, reflecting the more conservative real-time classification. Panel B reports annualised performance (percent, except the Sharpe ratio) of the allocation rule of Eq. (4), net of 10 bps transaction costs; the defensive weight equals the one-month-lagged real-time turbulent-regime probability. The full-sample row uses the probabilities underlying Table 3 over the same window. Sensitivity to the re-estimation interval and to transaction costs is reported in Appendix Table A10.




Fig. 1. The global liquidity factor and the probability of the turbulent regime, 2005–2025

Notes: Panel A plots the smoothed (solid) and filtered (grey) estimates of the latent global liquidity factor, in standardised units; higher values denote more accommodative common funding conditions. Panel B plots the smoothed probability of the high-volatility ("turbulent") regime from a two-state Markov-switching model estimated on the smoothed factor. The sample is monthly, January 2005 to December 2025. Shaded bands mark major funding-stress episodes: the global financial crisis (2008–2009), the euro-area sovereign-debt crisis (2011–2012), the COVID-19 shock (2020) and the 2022–2023 monetary tightening.




Fig. 2. Cross-asset return differentials between stable and turbulent liquidity regimes

Notes: Bars show the difference in mean monthly log return (percent) between the stable and turbulent regimes, computed as stable minus turbulent, for each of the thirteen asset-return series. Positive values indicate higher average returns in the stable regime. Assets are ordered from the largest to the smallest differential. Dark bars denote differentials significant at the 10% level or better under Newey–West (HAC) inference (maxlags = 3); full inference is reported in Table 2, Panel B. Sample: monthly, 2005–2025.