SQF – Schneegans Quantitative Finance

# Beyond the Liquidity Ladder

## Expected Time Under Stress as a New Treasury Risk Measure

A Stochastic Liquidity Modeling Framework
for Quantitative Treasury Analytics

Dr. Joachim Schneegans

July 2026

## Contents

## 1 Executive Summary

Treasury liquidity management begins with a liquidity maturity ladder: a time-ordered view of expected inflows, outflows, available liquidity and funding needs. The ladder remains indispensable, but a single planned path cannot show how long a future liquidity position may remain under pressure, whether breaches recur, or whether a stressed position recovers before the planning horizon. This whitepaper develops a stochastic extension of that established Treasury framework.

The model starts from planned position-level cash flows and their historical signed deviations. The least absolute shrinkage and selection operator (LASSO) estimates sparse predictive factor exposures for each cash-flow position [1]. The union of factors selected across positions defines the factor vector calibrated in a joint factor autoregression (JFA). Simulated factor paths and jointly simulated cash-flow residuals form two stochastic components within one Monte Carlo architecture. Each scenario produces a complete liquidity path rather than only an endpoint distribution.

Behavioural cash-flow modelling, penalised regression, multivariate scenario generation and Monte Carlo simulation are established upstream methods; stochastic cash-flow analysis has a well-established corporate-risk precedent [2]. The contribution developed here is the systematic application of occupation-time concepts and excursion-based path statistics to simulated Treasury liquidity paths. Together, these statistics quantify the duration, persistence, recurrence, recovery and severity of liquidity stress.

The principal duration measure is *Expected Time Under Stress* (ETUS): the expected calendar time spent in a predefined liquidity stress region. Conditional ETUS measures the average duration among scenarios in which stress actually occurs. These measures are complemented by first-passage time, excursion counts and durations, recovery statistics, maximum depth and integrated shortfall. The resulting metric set connects naturally to funding tenor, contingency preparation time, liquidity-buffer sizing and recovery planning.

> **The central Treasury insight is that two scenarios with the same terminal liquidity and the same minimum liquidity can require materially different responses when one produces a long uninterrupted shortfall and the other produces brief, recoverable episodes.**

## 2 From the Liquidity Ladder to Stochastic Paths

### 2.1 The Treasury foundation

The traditional liquidity ladder organises contractual and expected cash flows into maturity buckets and accumulates them into a funding profile. In practice, Treasury combines contractual cash flows with behavioural assumptions, contingent outflows, liquidity buffers and potential management actions. Liquidity gap analysis, survival-horizon analysis and stress testing provide complementary views of the institution's capacity to meet obligations [3, 4].

The framework proposed here does not replace the ladder. It treats the deterministic plan as a reference path: the institution's information set and planning assumptions at a specified vintage. It is neither a forecast with certainty nor a neutral baseline that can be revised retrospectively. The model asks how realised cash flows may deviate from that reference path and how those deviations accumulate into alternative liquidity trajectories.

This positioning matters operationally. Treasury continues to own the cash-flow taxonomy, behavioural assumptions, currencies, time buckets, liquidity perimeter, buffer definition and management-action policy. The stochastic layer adds a distribution of paths around those choices. It therefore extends established liquidity planning without asserting direct equivalence to LCR, NSFR or any prescribed regulatory calculation.

### 2.2 Cash-flow architecture

The cash-flow perimeter should preserve economically distinct sources rather than forcing all future flows into a single statistical series.

Table 1: Cash-flow classes in the Treasury reference plan.

| **Class** | Typical content | Treatment in the framework |
|---|---|---|
| **Contractual** | Debt maturities, coupons, committed payments and receivables | Retained explicitly according to contractual terms and data quality |
| **Behavioural** | Non-maturing deposits, prepayments, renewals and customer drawdowns | Represented by approved behavioural assumptions and modelled deviations |
| **Contingent** | Collateral calls, committed facilities, guarantees and stress-sensitive outflows | Included through scenario-dependent assumptions or dedicated positions |
| **Management actions** | Funding issuance, asset monetisation and discretionary balance-sheet measures | Modelled separately, with feasibility, timing and governance constraints |

For cash-flow position $i$ and bucket $t$, let $P_{i,t}$ denote the planned net cash flow recorded in the applicable plan vintage, and let $A_{i,t}$ denote the subsequently realised net cash flow. The historical signed plan deviation is

$$D_{i,t} = A_{i,t} - P_{i,t}. \tag{1}$$

The sign convention is preserved throughout: a positive deviation means that realised net cash flow is above plan; a negative deviation means that realised net cash flow is below plan. Plan vintages are essential. Each $D_{i,t}$ must compare the actual with the plan genuinely available before $t$, not with a later restatement.

The proposed framework is implemented and numerically validated in a Python prototype using a synthetic Treasury portfolio. The implementation provides the computational basis for the illustrative results presented in this paper.

# 3 Position-Level Deviation Model

## 3.1 Predictive factor exposures

Let $Z_t$ denote the broad, information-safe candidate factor universe. For each position $i$, the conditional deviation model is

$$D_{i,t} = \alpha_i + Z_t^\top \beta_i + \epsilon_{i,t}, \tag{2}$$

where $\alpha_i$ is an intercept, $\beta_i$ is a vector of predictive factor sensitivities and $\epsilon_{i,t}$ is the unexplained cash-flow residual. Candidate drivers can include yield-curve factors, credit spreads, inflation, activity measures, foreign-exchange rates, volatility and institution-specific behavioural indicators. Publication lags and data revisions must be respected.

The position-level LASSO estimator [1] is

$$(\widehat{\alpha}_i, \widehat{\beta}_i) = \arg\min_{\alpha,\beta} \left\{ \frac{1}{2n} \sum_{t=1}^{n} \left( D_{i,t} - \alpha - Z_t^\top \beta \right)^2 + \lambda_i \|\beta\|_1 \right\}. \tag{3}$$

The penalty is selected through time-respecting validation, and preprocessing is fitted on training observations only. The estimated coefficients represent predictive exposure, not causal impact.

## 3.2 Factor selection before joint dynamics

For each position, define the selected factor set

$$\mathcal{F}_i = \{j : \widehat{\beta}_{i,j} \neq 0\}. \tag{4}$$

The joint scenario generator is calibrated only on the union

$$\mathcal{F} = \bigcup_i \mathcal{F}_i. \tag{5}$$

Let $X_t = Z_{t,\mathcal{F}}$ denote the resulting selected-factor vector. This explicit selection step keeps the multivariate dynamics aligned with the drivers retained by the position models and avoids calibrating a high- dimensional joint process on irrelevant candidates.

# 4 One Monte Carlo Architecture, Two Stochastic Blocks

## 4.1 Factor innovation block

After LASSO selection, the retained factors are simulated jointly by a parsimonious first-order vector autoregression [5], termed the *joint factor autoregression* (JFA) here.

$$\begin{aligned}
X_t &= c + \Phi X_{t-1} + u_t, && \mathbb{E}[u_t] = 0, \qquad \mathrm{Var}(u_t) = \Sigma_u, \\
(\widehat{c}, \widehat{\Phi}) &= \arg\min_{c,\Phi} \sum_{t=2}^{T} \|X_t - c - \Phi X_{t-1}\|_2^2, \\
X_t^{(s)} &= \widehat{c} + \widehat{\Phi} X_{t-1}^{(s)} + u_t^{(s)}, && u_t^{(s)} \sim \mathcal{N}(0, \widehat{\Sigma}_u).
\end{aligned} \tag{6}$$

The parameters $\widehat{c}$ and $\widehat{\Phi}$ are estimated by multivariate least squares. The resulting fitted differences are historical JFA estimation residuals and are used to estimate $\widehat{\Sigma}_u$. In the Monte Carlo simulation, new Gaussian innovations $u_t^{(s)}$, rather than historical residuals, drive the future factor paths. They are independent across future dates and scenarios, while $\widehat{\Sigma}_u$ captures their contemporaneous cross-factor dependence. The matrix $\widehat{\Phi}$ generates own-factor persistence and lagged cross-factor dependence, so the simulated factor paths are not iid.

The JFA is a baseline scenario-generator specification rather than a requirement of the downstream methodology; alternative joint factor models can be used without changing Occupation Time, ETUS, Conditional ETUS, excursion-based analysis or severity measures.

### 4.2 Cash-flow residual block

The residual vector

$$\epsilon_t = (\epsilon_{1,t}, \dots, \epsilon_{m,t})^\top \tag{7}$$

collects shocks that remain after conditioning on the selected factors. These are different objects from $u_t$: they affect unexplained cash-flow deviations rather than factor dynamics. Their contemporaneous dependence can be represented by

$$\mathrm{Var}(\epsilon_t) = \Sigma_\epsilon, \tag{8}$$

or empirically by resampling complete residual rows. The baseline may assume $\mathrm{Cov}(u_t, \epsilon_t) = 0$, but that assumption requires empirical and economic challenge.

### 4.3 From innovations to liquidity paths

For scenario $s$, the systematic deviation component is

$$\widehat{D}^{(s)}_{\mathrm{sys},i,t} = \widehat{\alpha}_i + \left(X^{(s)}_{\mathcal{F}_i,t}\right)^\top \widehat{\beta}_{i,\mathcal{F}_i}. \tag{9}$$

Adding the residual shock gives

$$D^{(s)}_{i,t} = \widehat{D}^{(s)}_{\mathrm{sys},i,t} + \epsilon^{(s)}_{i,t}, \qquad C^{(s)}_{i,t} = P_{i,t} + D^{(s)}_{i,t}. \tag{10}$$

Position cash flows aggregate to $C_t^{(s)} = \sum_i C^{(s)}_{i,t}$, and the liquidity path is

$$L_t^{(s)} = L_0 + \sum_{r \le t} C_r^{(s)}. \tag{11}$$

For irregular grids, all duration measures use actual bucket widths.

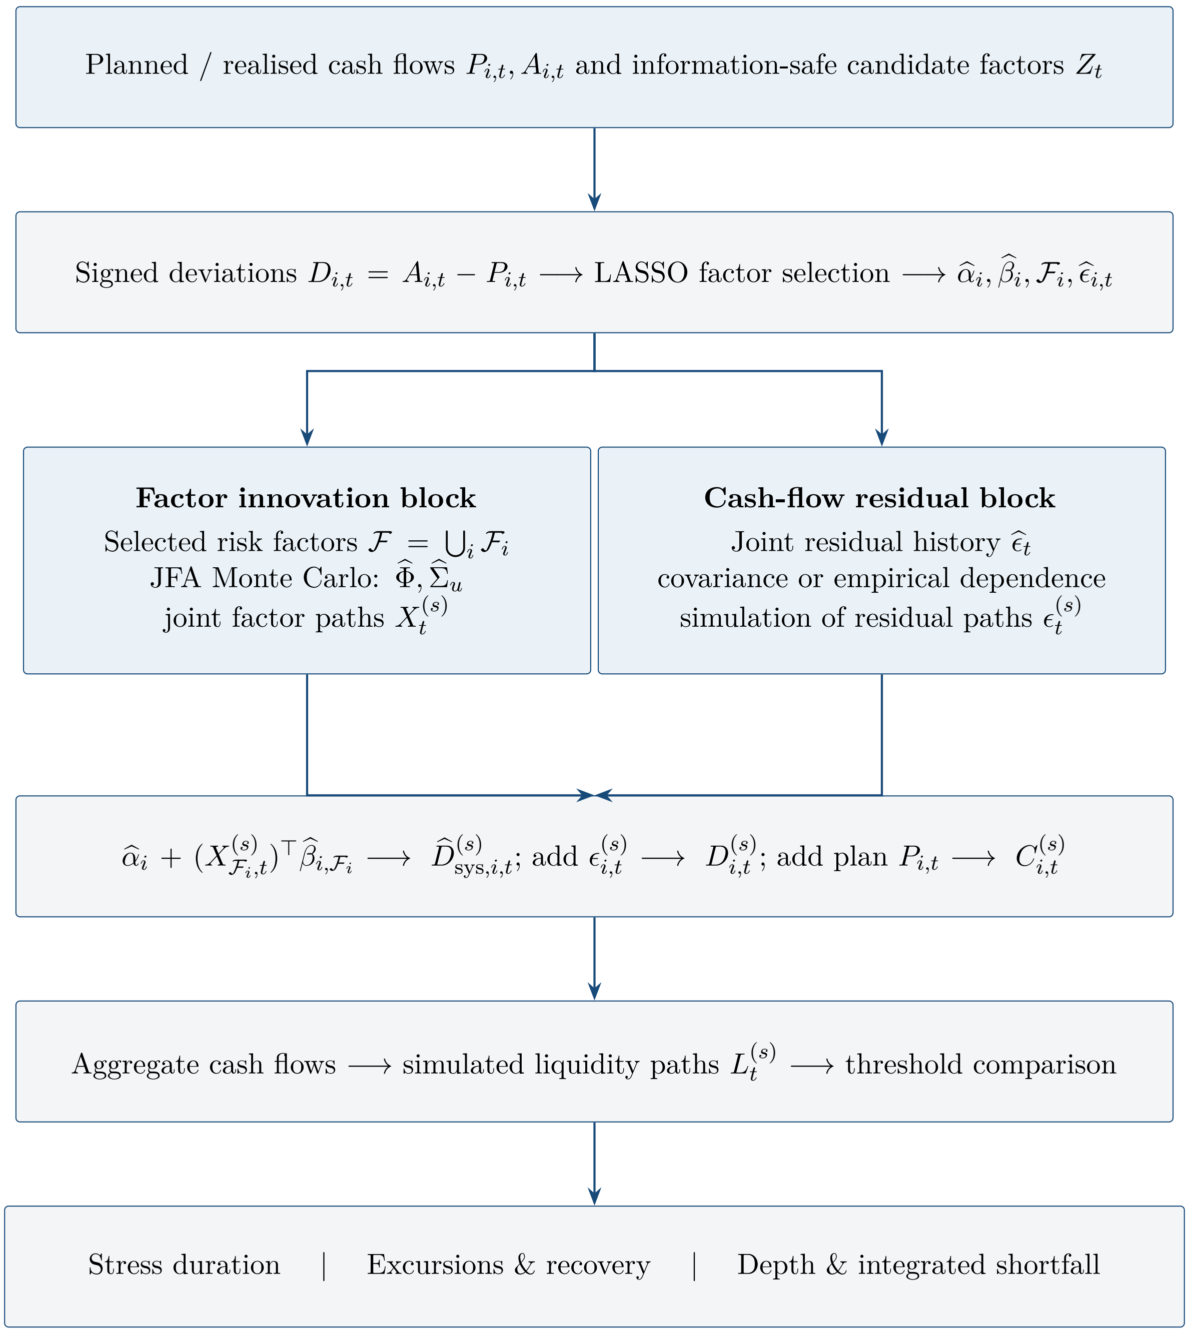


Figure 1: Integrated calibration, selected-factor scenario generation and liquidity-path analytics. The two stochastic components belong to one scenario architecture and merge before liquidity-path construction.

# 5 A Modular Factor Scenario Generator

## 5.1 Scenario alignment

Each simulated selected-factor path carries a consistent scenario and future-day index into every position model. Using only the position's selected coordinates and estimated coefficients produces systematic cash-flow deviations that remain mutually consistent with the common macro-financial scenario.

The independently generated residual vector $\epsilon_t^{(s)}$ represents cash-flow uncertainty not explained by the selected risk factors. For each Monte Carlo scenario $s$, this residual component is added to the factor-driven cash-flow deviation generated from $X_t^{(s)}$, yielding the total simulated cash-flow deviation. The factor innovations $u_t^{(s)}$ and cash-flow residuals $\epsilon_t^{(s)}$ therefore represent distinct sources of uncertainty within the same simulated liquidity scenario.

## 5.2 Modularity and alternative generators

Occupation and excursion analysis begins only after liquidity paths have been constructed. The path methodology therefore does not depend on the JFA baseline. Historical block bootstrap, multivariate mean-reverting processes and regime-switching models can supply alternative joint factor paths without changing the downstream definitions of liquidity stress.

The chosen generator must preserve economically relevant dependence, avoid look-ahead bias and expose its innovation assumptions. Its adequacy is judged both by factor diagnostics and by the sensitivity of downstream liquidity-path measures.

# 6 Occupation Time: Measuring Duration of Stress

## 6.1 Stress region and path functional

Let $(L_t)_{0\leq t\leq T}$ be a liquidity process with deterministic, possibly time-dependent threshold $h_t$. Define the liquidity stress region

$$\mathcal{S}_t = \{\ell : \ell < h_t\}. \tag{12}$$

The occupation time is

$$\mathrm{OT}_{\mathcal{S}}(T) = \int_0^T \mathbf{1}_{\{L_t \in \mathcal{S}_t\}}\, dt = \int_0^T \mathbf{1}_{\{L_t < h_t\}}\, dt. \tag{13}$$

On a grid $0 = t_0 < t_1 < \cdots < t_N = T$, with bucket widths $\Delta t_j = t_j - t_{j-1}$, the scenario-wise estimator is

$$\mathrm{OT}_{\mathcal{S}}^{(s)}(T) \approx \sum_{j=1}^{N} \Delta t_j \mathbf{1}_{\{L_{t_j}^{(s)} < h_{t_j}\}}. \tag{14}$$

The endpoint convention must be fixed. With daily buckets, $\mathrm{OT}_{\mathcal{S}}$ is measured in days.

### 6.2 Expected Time Under Stress

The principal duration measure is

$$\mathrm{ETUS}_{\mathcal{S}}(T) = \mathbb{E}[\mathrm{OT}_{\mathcal{S}}(T)]. \tag{15}$$

For $M$ equally weighted scenarios, it is estimated pathwise by the sample mean of Equation (14). A useful companion is

$$\mathrm{CETUS}_{\mathcal{S}}(T) = \mathbb{E}[\mathrm{OT}_{\mathcal{S}}(T) \mid \mathrm{OT}_{\mathcal{S}}(T) > 0]. \tag{16}$$

ETUS answers how much stress time is expected across all scenarios; conditional ETUS answers how long stress lasts on average when it occurs.

**Probability representation.** On the finite horizon, Tonelli's theorem applied to the non-negative indicator gives the equivalent representation

$$\mathrm{ETUS}_{\mathcal{S}}(T) = \int_0^T \mathbb{P}(L_t \in \mathcal{S}_t)\, dt. \tag{17}$$

This representation shows that ETUS can equivalently be interpreted as the time integral of the marginal probability of being in the stress region. In the Monte Carlo implementation, however, ETUS is estimated directly from the simulated liquidity paths by averaging their scenario-wise occupation times.

### 6.3 Occupation time and local time

Occupation time is calendar time spent inside a region. Local time describes the intensity with which a continuous stochastic process visits a particular level and is connected to occupation time through occupation-density theory [6]. ETUS is based on occupation time. The discrete bucket-counting estimator in Equation (14) is not a local-time estimator.

## 7 Excursion-Based Analysis and Severity

For a discrete liquidity path observed at end of day, a stress excursion is a completed threshold-return cycle. It starts when consecutive end-of-day observations record a crossing from $L_t^{(s)} > h_t$ to $L_t^{(s)} < h_t$, comprises the ensuing maximal sequence below the threshold, and ends with the first recorded crossing back to $L_t^{(s)} > h_t$. A below-threshold sequence still open at the horizon is reported separately as a right-censored stress episode and is not counted as a completed excursion. Both completed excursions and right-censored stress episodes contribute to occupation time. This operational discrete-time definition requires neither interpolation nor local time and is not a claim to a full implementation of Itô excursion theory.

Occupation time alone cannot distinguish one uninterrupted episode from several shorter episodes with the same total duration. First-passage time, excursion count, individual and longest duration, recovery time and the terminal-breach indicator describe the temporal organisation of stress. A recovery rule must state whether one non-stress bucket is sufficient or a sustained exit is required. The numerical example uses the first end-of-day value at or above the threshold.

Duration does not measure the amount of missing liquidity. Define the instantaneous shortfall

$$G_t^{(s)} = (h_t - L_t^{(s)})^+, \qquad x^+ = \max(x, 0). \tag{18}$$

Maximum depth and integrated shortfall are

$$\mathrm{MD}^{(s)}(T) = \sup_{0 \le t \le T} G_t^{(s)}, \tag{19}$$

$$\mathrm{IS}^{(s)}(T) = \int_0^T G_t^{(s)}\, dt \approx \sum_{j=1}^{N} \Delta t_j (h_{t_j} - L_{t_j}^{(s)})^+. \tag{20}$$

Occupation measures duration, excursions describe temporal structure, and severity measures depth. Integrated shortfall combines depth and duration in currency-time units.

## 8 Path-Level Illustration: Why Full-Path Analysis Changes the Treasury View

Figure 2 compares two persisted simulations with nearly identical terminal and minimum liquidity but different temporal structures. Scenario 4119 has one completed threshold-return cycle over future days 2–23 and returns above the threshold on day 24. Scenario 8484 has four completed excursions and 33 occupation days. Markers identify entry and return crossings; shading shows the integrated-shortfall contribution for $L_t < 10{,}000$.

The paths are actual Monte Carlo scenarios. Attribution reconstructed from their simulated risk-factor and cash-flow paths identifies credit spreads, interest-rate levels and volatility as the strongest adverse factor contributions; wholesale funding, committed facilities and margin/collateral provide the largest negative cash-flow contributions during stress. Their timing generates materially different occupation and excursion structures despite nearly identical terminal and minimum liquidity. These model-based path attributions should not be interpreted as causal effects.

## 9 Practical Treasury Interpretation

Applications include liquidity and funding planning, contingency funding, collateral management, ILAAP, internal stress testing, risk appetite and early-warning indicators. Thresholds may represent an internal reserve, management buffer or time-varying operational requirement and require independent approval with an explicit currency, perimeter and horizon.

Maximum depth sizes peak support, while conditional ETUS and the longest excursion indicate its potential tenor. Endpoint diagnostics remain useful secondary views but omit duration, recurrence and recovery.

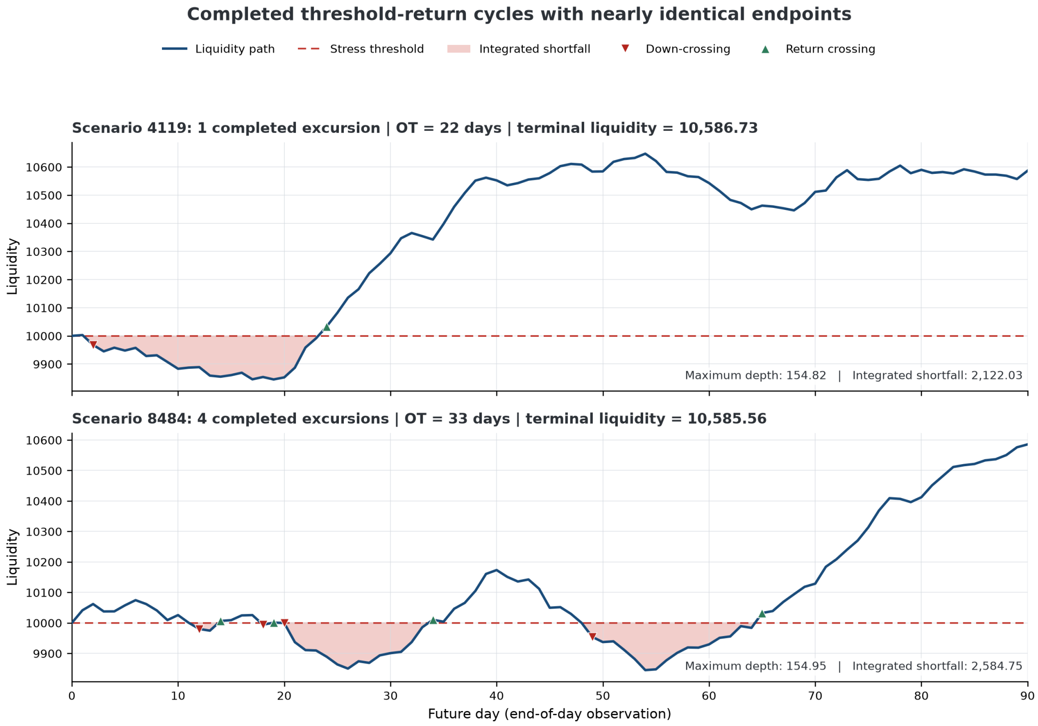


Figure 2: Actual simulated paths at $h = 10{,}000$. Scenario 4119 and Scenario 8484 have nearly identical terminal liquidity (10,586.73 versus 10,585.56) and minima (9,845.18 versus 9,845.05), yet their completed return-cycle structure differs. Scenario 4119 has one 22-day excursion; Scenario 8484 has four excursions lasting 2, 1, 14 and 16 days, for 33 occupation days in total. Integrated shortfall is 2,122.03 versus 2,584.75.

Table 2: Path measures and Treasury decisions.

| **Measure** | Treasury interpretation |
|---|---|
| **Any-breach probability** | Likelihood that the path enters the stress region |
| **Occupation time / ETUS** | Expected duration for which liquidity support may be required |
| **Conditional ETUS** | Expected support tenor if stress actually occurs |
| **First-passage time** | Preparation time before the first projected stress episode |
| **Longest excursion** | Potential uninterrupted tenor of a facility or buffer deployment |
| **Excursion count** | Recurrence and repeated execution or governance burden |
| **Maximum depth** | Peak liquidity amount and input to buffer sizing |
| **Integrated shortfall** | Combined depth-duration exposure in currency-time units |
| **Terminal breach** | Unresolved shortfall at the planning horizon |

# 10 Portfolio-Level Monte Carlo Results

While Section 8 illustrates path-level differences using two individual Monte Carlo scenarios, this section summarizes the full distribution across 10,000 simulated liquidity paths.

## 10.1 Synthetic design and path distribution

**Illustrative synthetic example.** The calculation uses 10,000 scenarios over 90 future days with $L_0 = 10{,}000$. Its six cash-flow categories are non-maturing deposits, loan repayments, wholesale funding, margin and collateral, committed facilities and operating cash flows. The moderately positive reference plan ranges from 4.16 to 11.84 units daily, averages 8.51 and cumulates to 766.00 units. Positions remain signed to preserve planned inflows and outflows.

The eight LASSO-retained factors are interest-rate level, credit spread, inflation, GDP growth, FX factor, sector factors 1 and 2, and volatility; the JFA models them jointly. This synthetic selection is not a universal Treasury factor specification. Margin and collateral links to credit spreads, FX and volatility, while other categories retain different subsets. The prototype-generated values are not empirical bank results.

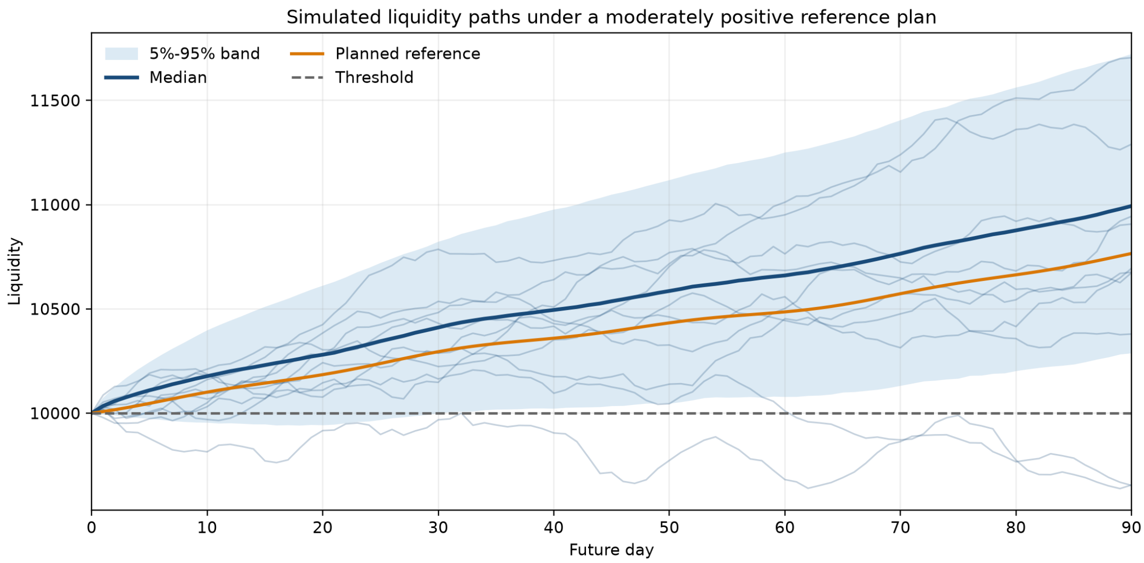


Figure 3: Selected simulated liquidity paths, the 5%–95% band, median, planned reference and illustrative threshold $h = 10{,}000$. The horizontal axis represents future days. The positive reference plan avoids an artificial deterministic decline while the stochastic deviations still produce stressed paths.

The median terminal liquidity is 10,993.74 and the mean is 10,999.01, with a terminal standard deviation of 435.94. The planned reference ends at 10,766.00. The baseline illustrative threshold is $h = 10{,}000$. It produces a non-trivial breach probability, a non-degenerate occupation-time distribution and observable variation in recovery and excursion structure. It is a transparent analytical threshold, not an economically or regulatorily calibrated limit.

### 10.2 Baseline result profile

Table 3: Path-risk results for the illustrative threshold $h = 10{,}000$.

| Metric | Persisted result |
|---|---|
| Breach probability | 30.76% |
| ETUS | 4.0808 days |
| Conditional ETUS | 13.2666 days |
| Median occupation time | 0 days |
| 95% occupation time | 25 days |
| Mean excursion count \| breach | 1.6268 |
| 95% longest excursion \| breach | 39 days |
| Recovery probability \| breach | 96.42% |
| Terminal-breach probability | 1.10% |
| 95% maximum depth \| breach | 273.88 |
| 95% integrated shortfall \| breach | 5,923.36 |

The unconditional median occupation time is zero because most scenarios do not breach. Conditional ETUS is therefore the clearer measure of duration among stressed paths, while the breach probability retains the frequency dimension. The low terminal-breach probability relative to the any-breach probability shows that most breached paths recover before day 90. The 95% maximum depth and integrated shortfall quantify severity among breached scenarios.

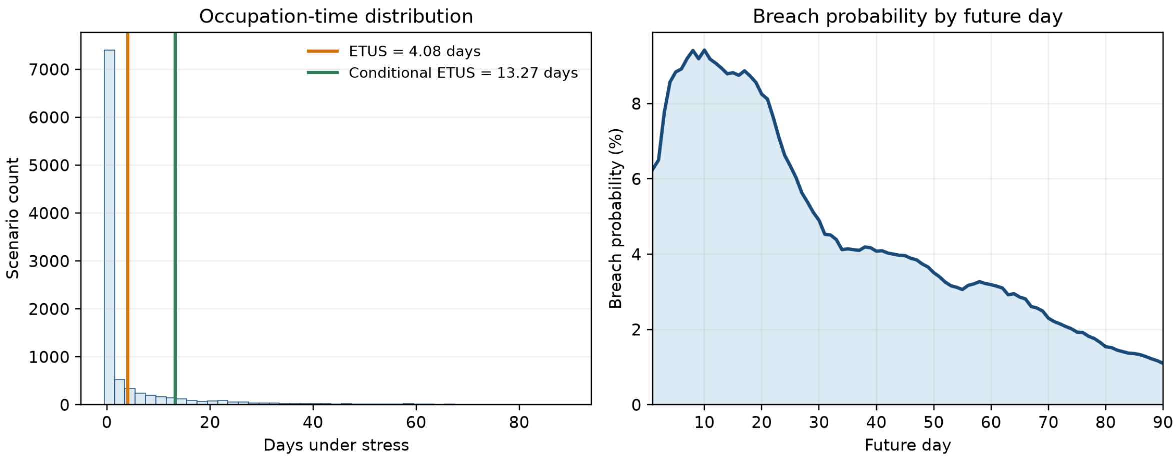


Figure 4: Duration diagnostics at $h = 10{,}000$. Left: the unconditional occupation-time distribution with ETUS and conditional ETUS. Right: breach probability by future day; its discrete time sum equals ETUS up to floating-point error.

### 10.3 Threshold sensitivity

The sensitivity table shows why probability and duration should be reported together. A low threshold produces rare stress, while a higher threshold classifies most paths as stressed and

Table 4: Sensitivity to the illustrative liquidity threshold.

| Threshold | Breach probability | ETUS | Conditional ETUS | Terminal breach |
|---|---|---|---|---|
| 9,900 | 9.63% | 1.3503 days | 14.0218 days | 0.61% |
| 10,000 | 30.76% | 4.0808 days | 13.2666 days | 1.10% |
| 10,050 | 81.92% | 7.7197 days | 9.4235 days | 1.37% |

substantially increases expected duration. Threshold choice is therefore a governed Treasury input, not a quantity inferred from a preferred result.

**Limitation of the synthetic example.** The moderately positive reference plan is an illustrative design choice, not an estimate of a bank's expected funding surplus. It removes the mechanically declining baseline of the previous example and permits both stressed and recovering paths. The resulting recovery frequency is useful for demonstrating the metric set, but it must not be interpreted as evidence about an actual Treasury portfolio.

## 11 Validation and Model Governance

Validation must follow the complete chain from plan vintages to reported path measures. Historical plans should be versioned and reconciled to realised cash flows under stable sign, currency and aggregation conventions. Factor observations must reflect the information available at each date.

LASSO validation should assess time-ordered out-of-sample performance, selection stability and residual diagnostics. The selected-factor union and factor order must be persisted before JFA calibration. For the JFA baseline, validation covers stability of $\Phi$, innovation diagnostics, parameter sensitivity and robustness of $\Sigma_u$. Residual diagnostics challenge $\Sigma_\epsilon$ or the empirical resampling design and the assumption $\mathrm{Cov}(u_t, \epsilon_t) = 0$.

Monte Carlo convergence should be assessed for ETUS, conditional ETUS, breach probability and tail quantiles of excursion duration and severity. Grid sensitivity matters because coarse buckets can miss short crossings or misstate first passage and recovery. Excursion detection should be checked against deterministic edge cases, including no breach, boundary equality, multiple episodes and terminal censoring.

The numerical checks confirm the occupation-time/excursion-duration identity exactly on the discrete grid and the ETUS/probability-sum identity to a maximum error of $8.88 \times 10^{-16}$ across the three reported thresholds. These are implementation consistency checks for the synthetic example, not empirical model validation.

### 11.1 Core assumptions and limitations

The framework assumes that historical plan-deviation relationships are informative for the simulation horizon, scenario dynamics are sufficiently stable for their intended use, and the plan and threshold are internally coherent. These assumptions may fail during structural breaks, balance-sheet transformations or unprecedented liquidity events.

Model risk is concentrated not only in statistical estimation but also in perimeter choice, behavioural assumptions, plan-vintage integrity, factor information timing, threshold design and recovery rules. Governance should retain both model artifacts and scenario paths so that Treasury can explain which positions, selected factors and residual shocks drive reported stress duration.

## 12 Conclusion

The liquidity ladder remains the natural starting point for Treasury analysis. Extending it into a distribution of stochastic liquidity paths makes it possible to answer questions that a single plan, endpoint quantile or minimum-liquidity statistic cannot answer.

Expected Time Under Stress provides an interpretable duration measure in units of calendar time. Conditional ETUS separates duration among stressed scenarios from breach frequency. Excursion-based statistics reveal onset, persistence, recurrence and recovery, while maximum depth and integrated shortfall add severity. The contribution is an integrated occupation-and-excursion framework that makes the temporal structure of simulated liquidity stress visible and actionable for quantitative Treasury.

The architecture separates planned cash flows, position-level factor selection, selected-factor scenario generation and residual cash-flow simulation. The first-order JFA is a transparent baseline, but the path analytics remain independent of that choice. The numerical illustration demonstrates the operational consequence: almost identical endpoint and minimum liquidity can coexist with materially different stress duration, excursion structure and integrated shortfall.

## References

[1] Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1):267–288, 1996.

[2] RiskMetrics Group. Corporatemetrics technical document. Technical report, RiskMetrics Group, April 1999.

[3] Basel Committee on Banking Supervision. Principles for sound liquidity risk management and supervision. Technical Report bcbs144, Bank for International Settlements, September 2008.

[4] Leonard Matz and Peter Neu, editors. *Liquidity Risk Measurement and Management: A Practitioner's Guide to Global Best Practices.* John Wiley & Sons (Asia) Pte Ltd, 2007.

[5] Christopher A. Sims. Macroeconomics and reality. *Econometrica*, 48(1):1–48, 1980.

[6] Daniel Revuz and Marc Yor. *Continuous Martingales and Brownian Motion*, volume 293 of *Grundlehren der mathematischen Wissenschaften.* Springer Berlin, Heidelberg, 3 edition, 1999.