U.S. Treasury Buybacks

Market impact, counterfactual simulations, and policy effectiveness

An academic research report prepared for Ogdn Ames

Research cutoff: September 12, 2026 (UTC)
Market observations: through September 11 where available
Prepared by OpenClaw / Chip; AI-assisted research and original calculations

Useful market plumbing—not a stand-alone cure for rising yields.

The September 10 operation is verified. Treasury accepted $5.187 billion in face value against a $6 billion ceiling. The accepted basket implies approximately $3.684 billion in settlement cash, including accrued interest. The 10-year benchmark yield nevertheless rose 12.20 basis points on operation day.

That combination does not establish failure. Treasury's stated objective is regular, predictable off-the-run liquidity support—not a benchmark-yield target. Daily quotations show a relative liquidity signal, but selection bias, imperfect matching, spillovers, and a failing yield pretrend prevent a clean causal verdict.

The executed Monte Carlo model illustrates how yields can rise even when a buyback helps relative to the no-buyback counterfactual. Its intervention coefficients are assumptions, not estimates of the September operation's causal impact.

Table 1. Evidence classification

Evidence class	What this report supplies
Verified	Official schedule, announcements, results, and security-level accepted amounts.
Observed / descriptive	Bloomberg daily yields and quotes; FRED cross-check; matched contrasts.
Calculated	Cash-flow valuation, accrued interest, duration, convexity, and financing-adjusted DV01.
Illustrative	40,000 paired Monte Carlo paths; alternative size, stress, financing, and persistence assumptions.

Suggested citation: OpenClaw. (2026). U.S. Treasury buybacks: Market impact, counterfactual simulations, and policy effectiveness [AI-assisted research report prepared for Ogdn Ames].

Executive summary

Policy conclusion. Keep targeted buybacks as a complementary debt-management and liquidity tool. Do not make routine buybacks the primary response to an inflation shock, persistent fiscal-risk premium, or systemic funding crisis. Select the instrument that addresses the cause of the stress.

Table 2. Executive assessment

Question	Answer and confidence
Was the event real?	Yes. The revised schedule was published September 9; the operation ran September 10, 1:40–2:00 p.m. ET, with settlement scheduled September 11. High confidence.
How much was bought?	$5.187bn par from $10.489bn offered; 23 of 40 eligible issues. Capacity utilization 86.45%; offer acceptance 49.45%. High confidence.
Did yields fall?	No: Bloomberg 10-year yield rose 12.20bp and 20-year yield 10.07bp on September 10. Observed association, not a treatment effect.
Did liquidity improve?	Accepted securities had less quote widening than unaccepted eligible securities. A matched two-day contrast is -8.14 price-bp, but matching and treatment selection are problematic. Suggestive, not proven causal improvement.
Did taxpayers save money?	Not established. Cash paid is not face value retired; replacement funding, asset depletion, and future debt service must be counted.
What does simulation show?	For a hypothetical $6bn-par bill-funded operation, assumed mean day-one 10-year effect is -1.85bp. Yields still rise in 37.6% of modeled normal-regime paths. Not an empirical probability of success.

The strongest result is conceptual and accounting-based: a market-function intervention can help relative to its counterfactual while macroeconomic forces push the observed yield higher. The weakest result would be a precise causal claim about the September operation; this report does not make one.

Sources: U.S. Department of the Treasury (2026c–2026f, n.d.); Bloomberg L.P. (2026); author calculations. Simulated probabilities are conditional on the expressly assumed coefficient distributions.

1. Abstract, scope, and research questions

Abstract. This study evaluates U.S. Treasury buybacks as a market-function and debt-management instrument, anchored to the verified September 10, 2026 operation in the 10–20-year nominal coupon bucket. Official operation records are reconciled with daily security quotations and benchmark yields. A descriptive matched comparison is performed within the eligible universe, while diagnostic failures are used to limit causal interpretation. Bond cash flows are repriced to distinguish face value from cash expenditure and to measure financing-adjusted duration removal. A real-world Monte Carlo model combines historical Treasury yield covariance with explicitly illustrative, nonlinear policy effects. Results support buybacks as a targeted complement rather than a universal remedy: local liquidity support can coexist with a broad selloff, duration removal depends on funding, and a positive fiscal net present value cannot be inferred from retiring discounted bonds.

Questions answered

First, what was actually announced, accepted, and scheduled for settlement? Second, how did quoted spreads and yields behave around announcement and execution? Third, what financing-adjusted exposure did the purchase remove? Fourth, under what assumptions can a buyback change market outcomes? Finally, when should a different policy instrument lead?

Scope boundary

This is not a trading recommendation or a peer-reviewed causal evaluation. The principal information set stops September 12, 2026. Only one post-operation trading day is observed. Quarterly schedule entries after the cutoff are plans, not realized operations. No post-cutoff empirical outcomes are introduced.

No Treasury-wide fiscal savings estimate, reserve-balance estimate, or historical depth/transaction-impact estimate is invented. Daily Bloomberg quotes are not firm executable quotes for institutional size. The analysis explicitly reports these missing inputs.

Keywords: Treasury buybacks; market liquidity; off-the-run bonds; duration supply; Monte Carlo; debt management; causal identification.

2. Verified event chronology

Table 3. Verified event chronology

Date / time	Primary evidence	Interpretation
August 5	Quarterly refunding policy statement	Quarterly plan included up to $38bn liquidity-support and $25bn cash-management purchases. This is the original policy context, not a verified revised total.
September 9	Revised tentative schedule, labeled “For Publication September 9, 2026”	September 10 maximum shown as $6bn. August 11 same-bucket operation had a $2bn maximum.
September 10, 11:00 a.m. ET	Final announcement timestamp: 15:00 UTC	Final security list; 40 issues eligible. Preliminary release and final release are distinct events.
September 10, 1:40–2:00 p.m. ET	Operation timestamps: 17:40–18:00 UTC	Competitive offers evaluated; acceptance is not mechanically equal to the ceiling.
September 10, 2:01:34 p.m. ET	API result-update timestamp	Reported $10.489bn offered; $5.187bn accepted; 23 issues.
September 11	Settlement date in official results	Scheduled settlement; no independent delivery confirmation or penalties data were retrieved.

The historical-news assertion of a $2bn-to-$6bn increase is supported by the revised schedule and prior same-bucket ceiling. The public “Q3” schedule URL retrieved now contains the September revision too, so this report does not pretend it recovered an immutable August version of that same PDF.

Maturity bucket and actual eligibility differ: the schedule describes a 10–20-year range, while the final result gives eligible maturity dates February 15, 2037 through August 15, 2046. Subsequent calculations use the actual listed bonds, not a generic 15-year security.

Sources: U.S. Department of the Treasury (2026a–2026f); official buyback website/API snapshot. Raw PDFs, operation records, and retrieval hashes accompany the study.

3. What counts as success?

Treasury identifies two objectives: cash-management buybacks smooth cash balances and bill issuance, whereas liquidity-support buybacks provide a regular, predictable opportunity to sell off-the-run securities. The FAQ explicitly states that the current program is not intended to mitigate acute market stress (U.S. Department of the Treasury, n.d.).

Table 4. Prespecified success criteria

Objective	Relevant success measure	Not sufficient evidence
Execution	Eligible offers accepted at defensible prices and within stated constraints.	Spending the entire announced ceiling.
Liquidity support	Lower effective trading costs or price impact, better depth and intermediation, without excessive scarcity.	A lower benchmark yield alone.
Cash management	Reduced cash/bill-supply volatility net of financing costs.	Face-value debt retirement alone.
Fiscal value	Positive financing-adjusted discounted benefit under explicit assumptions.	Buying a low-coupon bond below par.
System stability	Reduced fire-sale amplification or funding dysfunction using an appropriate mandate.	A small routine auction during a large selloff.

This report prespecifies research thresholds before examining simulation outputs: a positive local liquidity improvement, a local liquidity-yield reduction of at least 1bp, and positive financing-adjusted fiscal NPV. These are analyst thresholds, not Treasury targets. Quoted-width improvements are descriptive empirical proxies, not definitive effective-cost measures.

Treasury buybacks are not QE

Treasury retires the bonds it purchases. Financing may add other debt or reduce Treasury’s cash asset. A central-bank purchase typically changes central-bank assets and reserve liabilities. Treasury cash flows can temporarily affect reserves, but their effect depends on payment and financing timing; a gross Treasury buyback cannot simply be labeled monetary easing (Garbade & Rutherford, 2007; Duffie & Keane, 2023).

4. Historical and academic evidence

Table 5. Academic evidence and transfer limitations

Source	Relevant result / mechanism	Transfer limitation
Han et al. (2007)	The journal abstract reports $67.5bn of illiquid debt retired with small market-impact cost; auction selection favored relatively less-illiquid bonds within the eligible set.	Different fiscal regime; neither the cost nor the causal multiplier is transplanted into 2026. Abstract reviewed; full journal text not obtained.
Garbade & Rutherford (2007)	Full report explains buybacks for new-issue liquidity, cash smoothing, and debt management with and without a budget surplus.	Debt retirement financed by new issuance is a portfolio substitution, not free debt reduction.
D’Amico & King (2013)	Local supply and imperfect substitution can transmit purchase effects; accessible 2012 working-paper version reviewed.	Federal Reserve asset purchases are not the same balance-sheet operation as Treasury-funded buybacks.
Hu et al. (2013); Fleming (2003)	Yield-curve noise and bid–ask/depth measures capture different liquidity dimensions.	Daily indicative spreads and simple curve residuals cannot reproduce intraday institutional liquidity.
Vayanos & Vila (2021)	Preferred habitats and risk-averse arbitrageurs motivate maturity-specific demand effects.	Theory gives channels, not a verified $6bn policy coefficient.
Duffie & Keane (2023)	Fiscal and central-bank market-function purchases can be complementary; design and financing matter.	Working-paper policy analysis, not an evaluation of this operation.

The early-2000s experience occurred in a different fiscal setting. The modern operation panel contains small-value tests in April 2024; tests should not be mistaken for the regular program’s scale-up. In this archive, May 29, 2024 marks the first regular-sized liquidity-support operation. This distinction is retained in the reproducibility files.

Connolly and Struby’s (2024) published article is bibliographically verified; its earlier working-paper abstract was accessible. It reinforces the need to distinguish local supply effects from the overall bond return. No numerical coefficient from that abstract is used to calibrate this report.

5. Market transmission and yield decomposition

$$y_t^{(n)} = \tfrac{1}{n}\sum_{j=0}^{n-1} E_t[r_{t+j}] + TP_t^{(n)} + LP_t^{(n)}$$

(1) Conceptual yield decomposition. Expected future short rates, term premium TP, and liquidity/convenience component LP are not separately identified by a simple observed yield change.

A liquidity premium can fall because a dealer expects a reliable exit through the buyback facility. Local bond supply can also decline. At the same time, higher expected policy rates, inflation compensation, or term premia can push the benchmark yield upward. Convenience value may enter with a negative sign, so LP should not automatically be assumed positive.

$$\Delta y^{obs} = \Delta y^{macro} + \Delta y^{policy} + \Delta y^{other}$$

(2) Observed yield movements combine policy and non-policy forces. Without identification, these components cannot be recovered from their sum.

For illustration only: a +12bp observed move could combine a +14bp macro shock with a -2bp policy effect. It could also be a +12bp macro shock with no policy effect. The observed +12bp does not tell these stories apart.

Adrian et al. (2013) supply a formal term-premium estimation framework; this report does not relabel principal components as estimated term premia. The five-maturity factor model used below is a covariance model, not a structural decomposition. Gürkaynak et al. (2006) explain Treasury curve estimation; this report’s interpolated benchmark residual is materially simpler.

Potential adverse channels

Retiring too much free float can reduce a bond’s tradability or increase scarcity. Paying through fair value can transfer value to sellers. Replacement long-duration issuance can offset duration extraction. Short funding can move risk into future bill refinancing. These possibilities are retained in the evaluation rather than presumed away.

6. Data, provenance, and coverage

Table 6. Data inventory

Dataset	Coverage used	Important limitation
Official operations	162 returned records; 160 with results status, April 3, 2024–September 10, 2026.	Not every record is a regular liquidity-support operation; status/type filters matter.
Security quotations	60 CUSIPs; June 1–September 11, 2026; 4,500 raw and 4,320 holiday-cleaned observations.	40 September-10 eligible issues and 20 outside-bucket comparisons; vendor daily fields, not trade prints.
Bloomberg curve	2y, 5y, 10y, 20y, 30y; 924 clean observations, January 3, 2023–September 11, 2026.	Benchmark indices are not identical to the Federal Reserve constant-maturity series.
FRED / Federal Reserve H.15	Independent daily 2y/5y/10y/20y/30y par-yield cross-check through September 10.	September 11 was not present in retrieved FRED snapshot.
Official PDFs / scholarship	Primary announcement/result PDFs; Treasury FAQs and policy statements; academic full texts or documented abstracts.	A live web page can change; saved snapshots and hashes establish what was read.

Requested Bloomberg fields: PX_LAST, PX_BID, PX_ASK, and YLD_YTM_MID for securities; PX_LAST for benchmark indices. The response reported Bloomberg Desktop API, timestamp September 12, 2026, 05:00:27 UTC for security data. No security errors, inverted spreads, duplicates, or missing requested columns were found in the used panel.

Treasury holidays are removed using H.15 missing-value dates rather than retaining repeated vendor holiday observations. No nonholiday values are forward-filled. FRED's retrieval endpoint returned a longer history than requested; dates are filtered in code.

Not retrieved: historical executable depth, quote sizes, CUSIP-level dealer inventory, effective spreads, repo specialness, futures-deliverability controls, verified contemporaneous free float, and operation-specific replacement issuance. Missing variables constrain identification and fiscal conclusions.

7. Bond mathematics: face value is not cash

$$P_{dirty}(y) = \sum_{k=1}^{K} \frac{CF_k}{(1 + y/2)^{u+k-1}}$$

(3) Semiannual yield convention. u is the ACT/ACT fraction of a coupon period to the next payment. Dirty price equals clean price plus accrued interest.

$$\frac{\Delta P}{P} \simeq -D_{mod}\Delta y + \frac{1}{2}C(\Delta y)^2$$

(4) Duration/convexity approximation; yields are decimals. One basis point is 0.0001. Exact cash-flow repricing is used to check the approximation.

Table 7. Accepted-basket valuation

Accepted basket calculation	Result
Face value retired	$5.187bn
Clean-price consideration	$3.663353bn
Estimated accrued interest	$20.891m
Estimated settlement cash	$3.684244bn
Market-value-weighted modified duration	11.907 years
Market-value-weighted convexity	170.407
Gross DV01	$4.387m per bp

Accepted auction prices are rounded to three decimals in the source. Accrued interest is reconstructed from maturity-linked semiannual coupon dates using an ACT/ACT period fraction. Estimated settlement cash is therefore a calculation, not a Treasury cash-payment confirmation. Rounding of accepted prices alone implies a worst-case cash uncertainty of about ±$25,935. Yield-to-maturity and risk measures are solved from the reconstructed dirty prices.

The roughly $1.503bn gap between face value and settlement cash is not an NPV saving. Low coupons make future contractual payments worth less than face value at current rates; Treasury must still finance the repurchase or consume a cash asset. Source: U.S. Department of the Treasury (2026f); author calculations.

8. Financing-adjusted duration removal

$$DV01 = D_{mod} \, MV \, 10^{-4}$$

(5) Dollar value of a one-basis-point yield change; MV is market value, not face value.

$$DV01_{net} = DV01_{repurchased} - DV01_{replacement}$$

(6) Positive values represent net private-sector duration removal under the specified financing assumption.

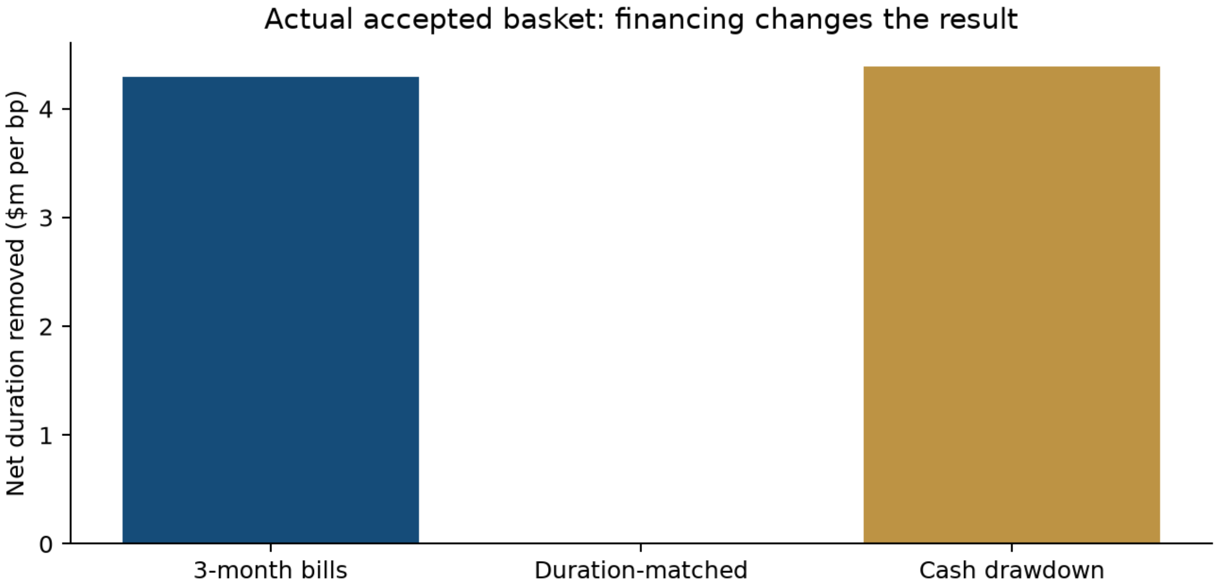

Figure 1. Author calculations from the actual accepted basket. Bill duration is assumed to be 0.25 years. Cash financing consumes a fiscal asset; this chart is not a measure of fiscal profit.

Table 8. Financing-adjusted duration removal

Funding assumption	Replacement DV01	Net removed
3-month bills	$0.092m/bp	$4.295m/bp
Duration-matched coupons	$4.387m/bp	$0.000m/bp
Existing cash	$0.000m/bp	$4.387m/bp

There is no evidence here identifying a particular new issue as the actual marginal financing for September 10. These are counterfactual financing cases. Bill funding preserves most of the immediate duration removal but adds rollover exposure. Duration-matched issuance removes essentially no net duration, although liquidity benefits from swapping old debt into more liquid new debt may remain.

9. Fiscal NPV: the correct accounting

$$NPV = PV(CF_{avoided}) - PV(CF_{new}) - C_{execution} - PV(C_{cash})$$

(7) For debt financing, C_cash is zero and new debt funds the purchase. For cash financing, include the cash asset consumed instead of new-debt cash flows. Do not count both the initial repurchase outlay and the debt principal that funds it without also including the financing proceeds.

At fair value with duration-matched refinancing and the same discount curve, the avoided debt service and replacement debt service have equal present values before friction. Retiring debt below par does not change that identity. A benefit must come from lower future financing costs, improved market functioning, risk management, or other demonstrable advantages.

$$G = K\delta\, a_{15}(r) - Kc, \quad a_{15}(r) = \frac{1-(1+r/2)^{-30}}{r}$$

(8) Normalized fixed-funding example. K is cash financed; delta is a persistent refinancing-yield concession in decimal units; c is an all-in cost as a fraction of cash. This is a sensitivity calculation, not an estimated Treasury saving.

The accompanying CSV has a normalized 15-year funding-equivalent NPV column. It deliberately values coupon-concession benefits at an unchanged discount benchmark. When a scenario is labeled bill-funded or cash-funded, this column is only a fixed-funding-equivalent sensitivity—not the actual NPV of bills or cash. It must not be added to the bill rollover-risk metric.

The separate one-year bill metric measures a financing-side rate surprise on rolling short debt; it is not a complete one-year policy NPV. Its approximate 0.375 exposure assumes three later quarterly resets under a linearly changing short rate, with terminal principal unchanged. It omits a full dynamic funding strategy.

A definitive realized fiscal estimate remains unavailable because the marginal funding plan, execution benchmark, longer-run issuance concession, and counterfactual cash strategy are not identified. Cash drawdown also forgoes optionality and buffer capacity (Garbade & Rutherford, 2007; Duffie & Keane, 2023).

10. Empirical design and identification limits

The operation panel establishes eligible and purchased securities. Within the September eligible set, 23 issues were accepted and 17 were not. A minimum-distance assignment pairs 17 accepted issues to the 17 unaccepted issues using standardized pre-event maturity, coupon, quote width, and curve-relative yield. Six accepted issues remain unmatched.

$$\widehat{\theta} = \frac{1}{J}\sum_{j=1}^{J}\left[(x^{A}_{j,post} - x^{U}_{j,post}) - (x^{A}_{j,pre} - x^{U}_{j,pre})\right]$$

(9) Descriptive matched difference-in-differences. A = accepted, U = unaccepted but eligible. Baseline: September 1–8 (five trading days); post: September 10–11. The contrast is not identified as a causal treatment effect.

Acceptance is endogenous: Treasury considers prevailing prices and relative value. Eligible-but-unaccepted issues may also benefit from expected future purchases, and comparisons just outside the maturity bucket do not overlap the treated maturity range. Therefore, neither group is a clean untreated counterfactual.

The outside-bucket sample comprises ten shorter and ten longer securities drawn from the nearest available adjacent-bucket operation lists. It is used for descriptive plots only—not to claim a maturity-controlled treatment effect. The 160-result panel is not fed into a naïve staggered two-way-fixed-effects estimator: repeated exposure, spillovers, and sparse identification would make that misleading.

Bootstrap intervals resample matched pairs 5,000 times, holding the match fixed. They reflect cross-pair dispersion, not common-shock uncertainty, matching uncertainty, or exogenous assignment. A formal intraday design with executable liquidity and controls for simultaneous macro news would be stronger, but those data were not obtained.

The historical selection problem is consistent with Han et al. (2007). Absence of those data is reported, not repaired by pretending an observational contrast is randomized.

11. Observed quote behavior

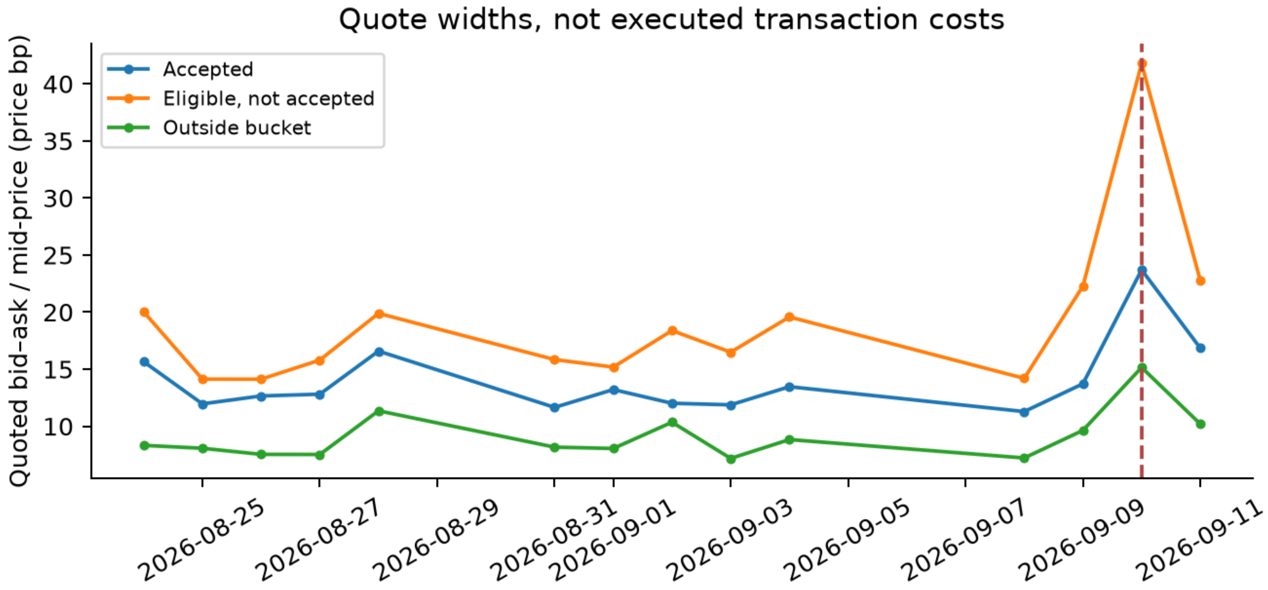

Figure 2. Daily mean quoted bid–ask width in price-basis-points of mid-price. Source: Bloomberg L.P. (2026); author calculations. An increase means wider quotes, not improvement.

Accepted bonds’ quoted costs need not decline in absolute terms for their relative behavior to be better. On operation day, the unaccepted eligible group widened more. This can be consistent with selective liquidity support, but also with pre-existing differences and selection into acceptance.

Table 9. Descriptive matched quote-width contrasts

Window	Matched quote contrast	95% pair-bootstrap interval
announcement	-3.62 price-bp	[-5.77, -1.41]
operation	-14.58 price-bp	[-18.15, -11.29]
settlement	-1.69 price-bp	[-3.73, 0.18]
operation and settlement	-8.14 price-bp	[-10.76, -5.78]

The two-day estimate is -8.14 price-bp. This is a relative bid–ask-width contrast, not an 8.14bp fall in bond yields, nor a reduction in an institution’s actual implementation shortfall. On settlement day alone, its bootstrap interval includes zero. Sustained improvement cannot be established from this short follow-up.

12. Yield and matching diagnostics

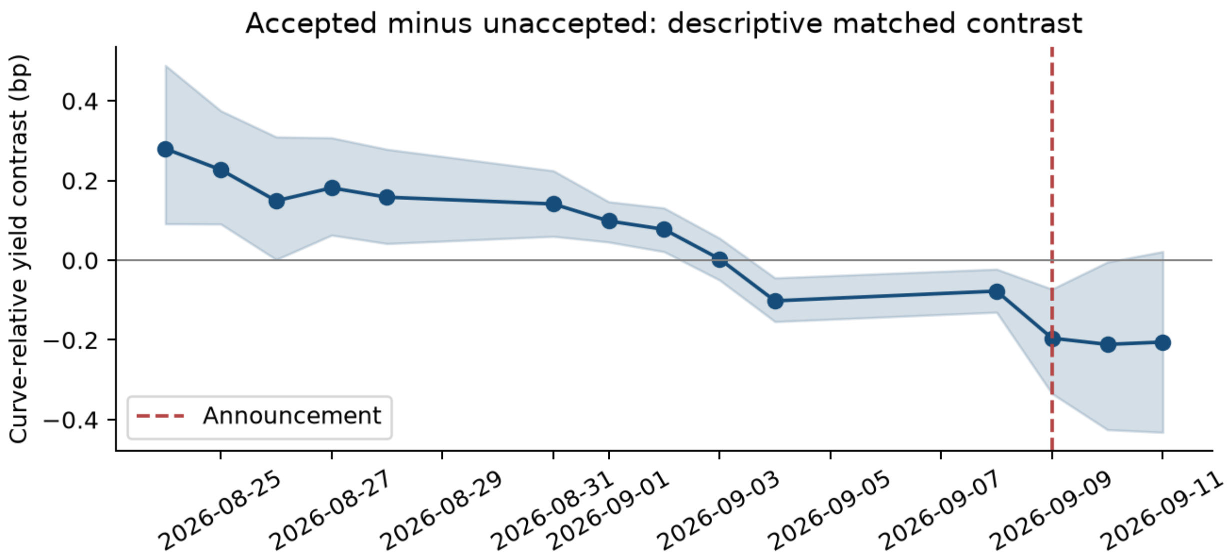

Figure 3. Curve-relative yield contrast, accepted minus unaccepted, centered on the five-day baseline. Shading is a pair-bootstrap interval, not a causal confidence band. The clear pre-event trend undermines a yield-effect claim.

The operation-plus-settlement curve-relative yield contrast is -0.208bp, with a 95% pair-bootstrap interval spanning zero. The pre-event slope is -0.053bp per trading day. A pre-existing trend of that size can explain much of the apparent post-event advantage.

Matching remains imbalanced. Standardized accepted-minus-unaccepted differences are -0.68 for remaining maturity, +0.42 for coupon, -0.75 for pre-event quote width, and +0.19 for the yield residual. These are not close balance for several important variables. A statistically narrow bootstrap interval cannot cure this design failure.

A strict squared-distance caliper of 1 retains 2 pairs; the resulting small-sample sensitivity is recorded in results.json. Across 1-, 5-, 10-, and 20-day baselines, the quote contrast is directionally stable, whereas the yield contrast changes materially. Placebo yield contrasts as large in absolute value occur in 63 pre-event windows with frequency 23.8%.

Conclusion: there is descriptive evidence of differential quote behavior, but no defensible isolated causal yield effect. No treatment coefficient from these contrasts is passed into the Monte Carlo model.

13. Why the market still sold off

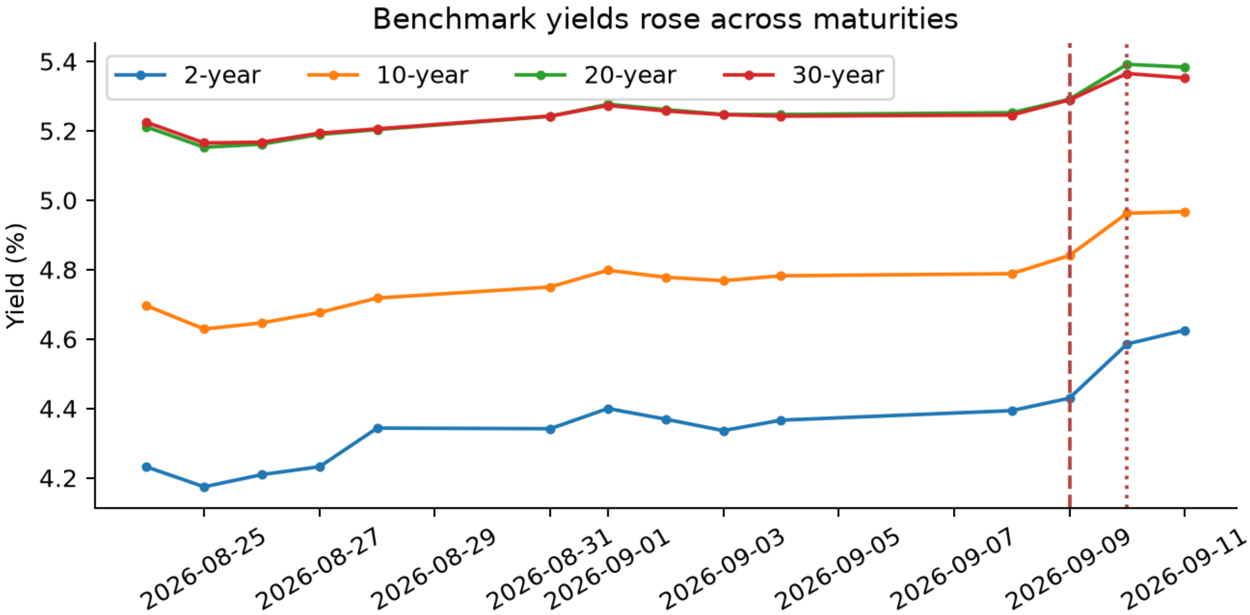

Figure 4. Bloomberg benchmark yields. Dashed and dotted lines mark announcement and operation dates. These are temporal markers, not identified causes. Source: Bloomberg L.P. (2026).

Table 10. Observed benchmark-yield changes

Maturity	Sept. 9 yield	Sept. 10 yield	Change
2 years	4.4295%	4.5856%	+15.61bp
5 years	4.6163%	4.7618%	+14.55bp
10 years	4.8406%	4.9626%	+12.20bp
20 years	5.2908%	5.3915%	+10.07bp
30 years	5.2889%	5.3652%	+7.63bp

The move was not confined to the eligible bucket. That is consistent with a common macro/rates shock, although it does not identify its cause. The FRED/H.15 cross-check likewise shows rising rates: +12bp at ten years and +11bp at twenty years. Small differences from Bloomberg reflect different series and observations.

Earlier archive reporting described higher oil prices and inflation concerns. That context is not used as a verified numerical shock instrument. The empirical design cannot apportion the selloff among inflation, policy-rate expectations, term premia, issuance, or the buyback. It can reject the simplistic claim that an observed yield increase alone proves operational failure.

14. Monte Carlo model and calibration

$$\Delta \mathbf{y}_t = L\mathbf{z}_t + \boldsymbol{\varepsilon}_t + \boldsymbol{\mu}_{s,t}$$

(10) Five maturity changes, in basis points. Three principal-component shocks plus diagonal idiosyncratic residual reproduce historical covariance; scenario drift is specified separately.

The real-world model uses 920 holiday-cleaned daily changes through September 8, 2026, excluding the event from calibration. Three principal components explain 99.82% of sample covariance. The model is positive definite; its minimum covariance eigenvalue is 0.0772 bp-squared.

Table 11. Historical daily yield volatility

Maturity	Daily yield volatility
2 years	6.78bp
5 years	6.54bp
10 years	5.87bp
20 years	5.43bp
30 years	5.20bp

Normal-regime drift is set to zero rather than extrapolating a sample-average rise in rates indefinitely. The resulting yield process is a finite-horizon risk scenario engine, not an arbitrage-free pricing model or an interest-rate forecast. One-year fan charts are especially assumption-dependent.

Forty thousand seeded paths are simulated; only 4,000 are used for plotting fan quantiles. Every policy/no-policy comparison uses the same market shocks. Parameter uncertainty in liquidity strength, persistence, and duration transmission is drawn independently of market shocks. This independence is a simplifying assumption that may understate adverse state dependence.

A five-day historical vector-block bootstrap provides a non-Gaussian robustness comparison. A separate 300-resample moving-block calculation quantifies volatility-estimation uncertainty; it is not combined into the principal outcome intervals. Monte Carlo methodology follows the framework discussed by Glasserman (2003).

15. Nonlinear intervention assumptions

$$g(q) = \frac{1 - e^{-q/6}}{1 - e^{-1}}, \quad \ell_h = A_s g(q) 2^{-(h-1)/H}$$

(11) q is billions of par, g(6)=1. A is the assumed signed liquidity-yield benefit at $6bn; H is its half-life in trading days. The function imposes diminishing one-off liquidity returns.

$$\Delta y_{m,h}^{policy} = -w_m \left[\ell_h + B \frac{q}{6} \frac{D - D_f}{D} \right]$$

(12) w is a bell-shaped maturity loading centered on 15 years with width 7 years. D_f is replacement-debt duration. The persistent term-supply coefficient B is assumed, not estimated.

Table 12. Illustrative intervention parameters

Parameter	Distribution / rule	Evidence status
A, normal liquidity effect	Triangular(-0.5, 2, 5) local yield bp at $6bn	Illustrative; permits adverse effects but favors benefits by construction.
B, persistent duration channel	Triangular(0, 0.15, 0.5) local yield bp at $6bn	Illustrative, nonnegative; scale must not be read as literature-calibrated.
H, persistence	5, 20, or 60 trading days; probabilities 25%, 50%, 25%	Illustrative parameter uncertainty.
Quoted-width mapping	20% improvement per 2bp local benefit, capped at ±50%	Illustrative; not an estimated effective-spread response.
Execution friction	Uniform 0.1–1.0 price-bp of cash used	Assumed all-in sensitivity, not measured auction slippage.

The assumed distribution assigns positive liquidity effects in most paths. Consequently, a high “probability of improvement” largely reflects that input choice, not proof from the data. The negative-effect support is limited; this is not a worst-case prior. Setting A and B to zero makes the intervention mechanically ineffective, while negative A makes local yields and quoted costs worse.

Local effects are capped for repeated programs. Without such limits, summing operations indefinitely would create implausible liquidity gains. Preferred-habitat theory motivates the shape, not its numerical scale (Vayanos & Vila, 2021).

16. Scenarios, funding, and success criteria

Table 13. Conditional market regimes

Regime	Market shocks	Intervention assumptions
Normal	Historical covariance; zero expected drift.	Baseline coefficient and persistence distributions.
Inflation repricing	1.25× volatility; an assumed +50bp ten-year macro shift builds over 63 trading days.	Same local liquidity efficacy as normal; buyback does not remove inflation.
Liquidity stress	2.5× volatility plus a +20bp local shock, decaying with a 20-day half-life.	2.5× local efficacy; still not a claim that the routine program is designed for acute crises.

No probabilities are assigned across regimes. The inflation and liquidity shocks are analyst scenarios, not historical estimates of September 10. Correlation between intervention strength and market shocks is not estimated. The stress case must be read as a mechanism illustration, not a policy authorization.

Table 14. Purchase-size scenarios

Purchase size	Treatment
$0, $2bn, $6bn, $10bn par	One-off comparisons; $10bn is outside the verified operation ceiling.
$5.187bn par	Actual accepted size, using the same hypothetical coefficient model.
$6bn every five trading days	26 operations by trading day 126: $156bn par. Illustrative scale-up, not an announced Treasury program.

Each size is evaluated with bill financing, duration-matched refinancing, and cash drawdown. Horizons are 1, 21, 63, and 252 trading days. Cash consumption for the full repeated program would be about $110.8bn at the observed cash/par ratio, before operational and market-price changes; feasibility cannot be presumed.

Success indicators are conditional probability of a positive local improvement; at least 1bp local benefit; the change in quoted-width proxy; investor repricing; and financing-side cost/risk. A reduced yield relative to baseline is not the same as a negative observed yield change.

17. Simulated immediate effects

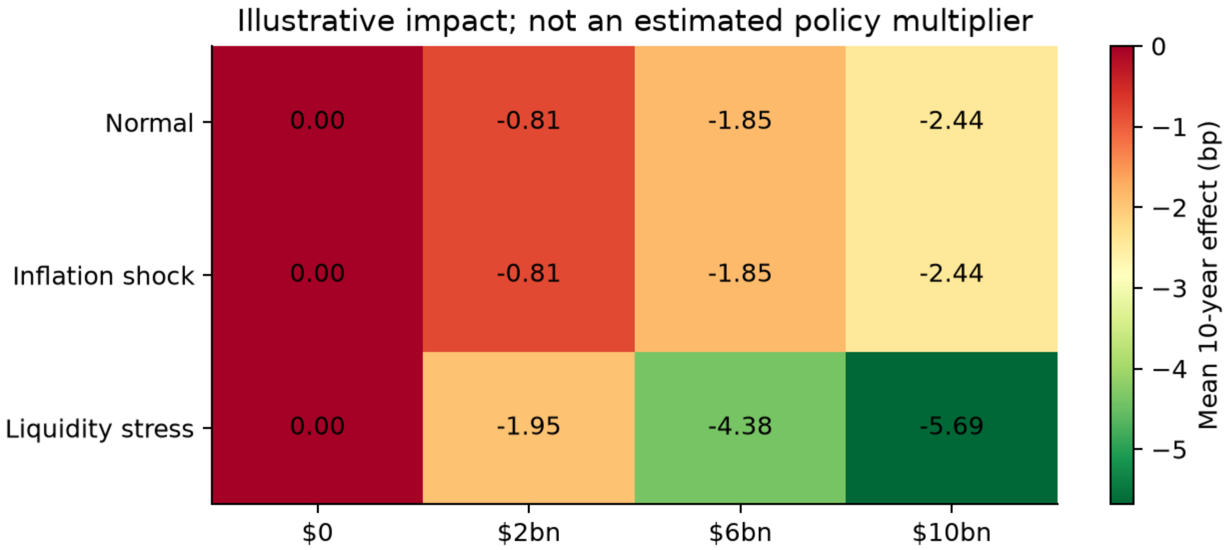

Figure 5. Mean day-one 10-year yield effect relative to the same-shock no-buyback baseline. All policy multipliers are illustrative. Dollar labels refer to face value, not cash expenditure.

Table 15. Immediate simulated policy effects

Size / regime	Mean 10y effect	5th–95th effect	P(yield rises)
Actual $5.187bn / normal	-1.69bp	[-3.05, -0.38]bp	38.6%
$6bn / normal	-1.85bp	[-3.34, -0.42]bp	37.6%
$6bn / inflation	-1.85bp	[-3.34, -0.42]bp	44.3%
$6bn / liquidity stress	-4.38bp	[-8.11, -0.81]bp	77.3%

In the $6bn normal-regime illustration, yields rise despite positive local liquidity improvement in 36.7% of all paths. The probability of at least 1bp local improvement is 84.0%, conditional on the assumed A distribution. Neither statistic estimates the probability that the actual operation succeeded.

At the actual accepted size, the numerical effect differs from the headline-size scenario. This is why a reported maximum must not be treated as the amount purchased. The observed +12.20bp move remains outside this model's causal interpretation: we do not subtract the assumed effect and call the difference an estimated macro shock.

18. Persistence and the broad market

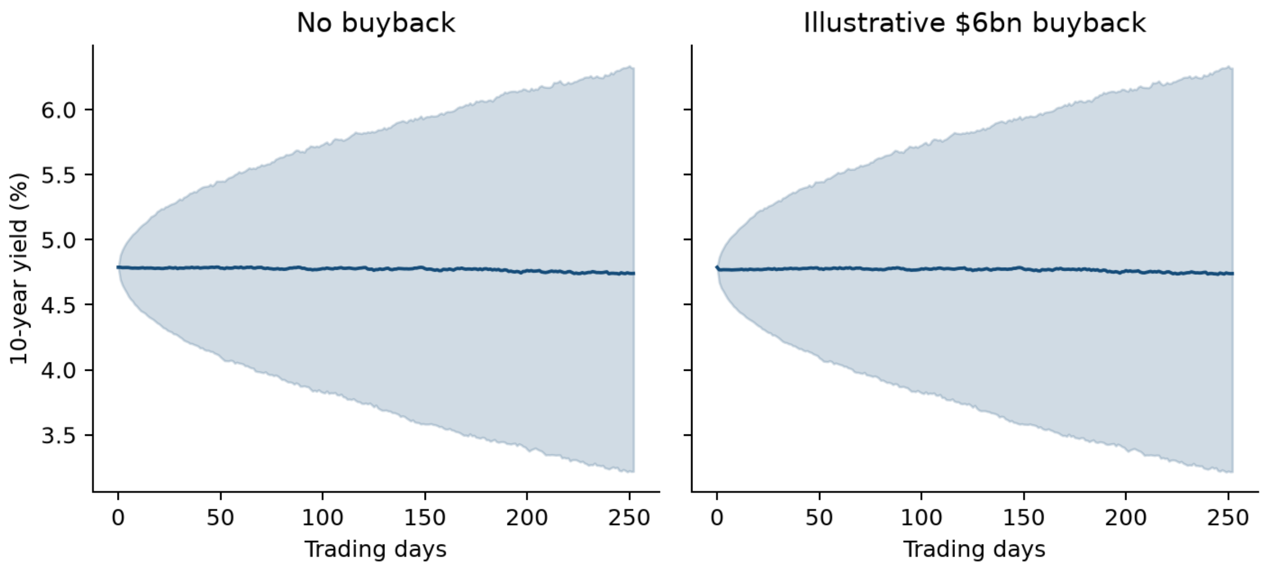

Figure 6. Normal-regime 10-year yield fan: median and 5th–95th percentiles. $6bn par, bill-funded; 4,000 plotting paths from a 40,000-path simulation. Wide market-risk bands dwarf the assumed policy shift.

Table 16. Persistence of simulated effects

Horizon	Mean paired effect	P(yield rises)	P(local gain ≥1bp)
1 trading day	-1.85bp	37.6%	84.0%
1 month	-0.95bp	48.6%	46.8%
3 months	-0.47bp	49.9%	13.2%
1 year	-0.19bp	49.8%	0.0%

The one-off liquidity benefit decays. A small assumed persistent duration component remains when financing is short. By one year, normal-regime market uncertainty overwhelms that shift. These are risk scenarios, not yield forecasts.

In the inflation scenario at three months, the 10-year yield still rises in 80.4% of paths. The model is not saying inflation has that probability; it is saying that, conditional on the specified +50bp macro scenario, the modeled buyback is too small to dominate it.

The full CSV includes all five maturities, financing cases, sizes, and horizons. Permanent rate suppression would require a different scale and policy mechanism than a transient local-liquidity intervention.

19. Investor price impact and exact repricing

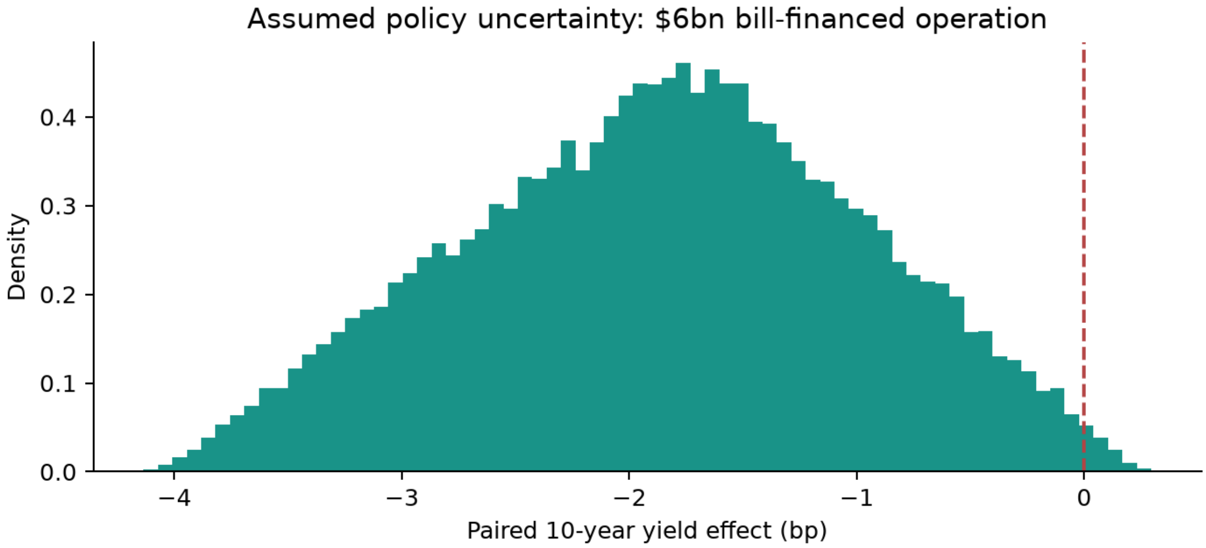

Figure 7. Distribution of the paired day-one 10-year yield effect under assumed policy-parameter uncertainty. This is not a distribution of empirically estimated September-10 effects.

Exact repricing of an illustrative 15-year, 4% coupon bond gives a mean policy-minus-no-policy return difference of 0.202%, with a 5th–95th interval of [0.046%, 0.365%]. This is an instantaneous valuation comparison under a modeled shift, not a 15-year total-return forecast.

Table 17. Exact and approximate basket repricing

Parallel yield shock	Actual basket: exact price return	Duration/convexity approximation
-100bp	12.805%	12.759%
-25bp	3.031%	3.030%
+25bp	-2.924%	-2.924%
+100bp	-11.097%	-11.055%

The actual basket calculations use each purchased bond’s own cash flows and fitted yield. The approximation is close for small shocks and less accurate for 100bp changes. Coupon income, roll-down, reinvestment, and nonparallel curve shifts are not represented by this instantaneous price table.

Investor mark-to-market gains are not Treasury fiscal savings. Gains may reflect lower discount rates, transfers from taxpayers through execution prices, or changing liquidity services. The fiscal question requires a separate financing-aware accounting.

20. Cost–benefit frontier, not a savings claim

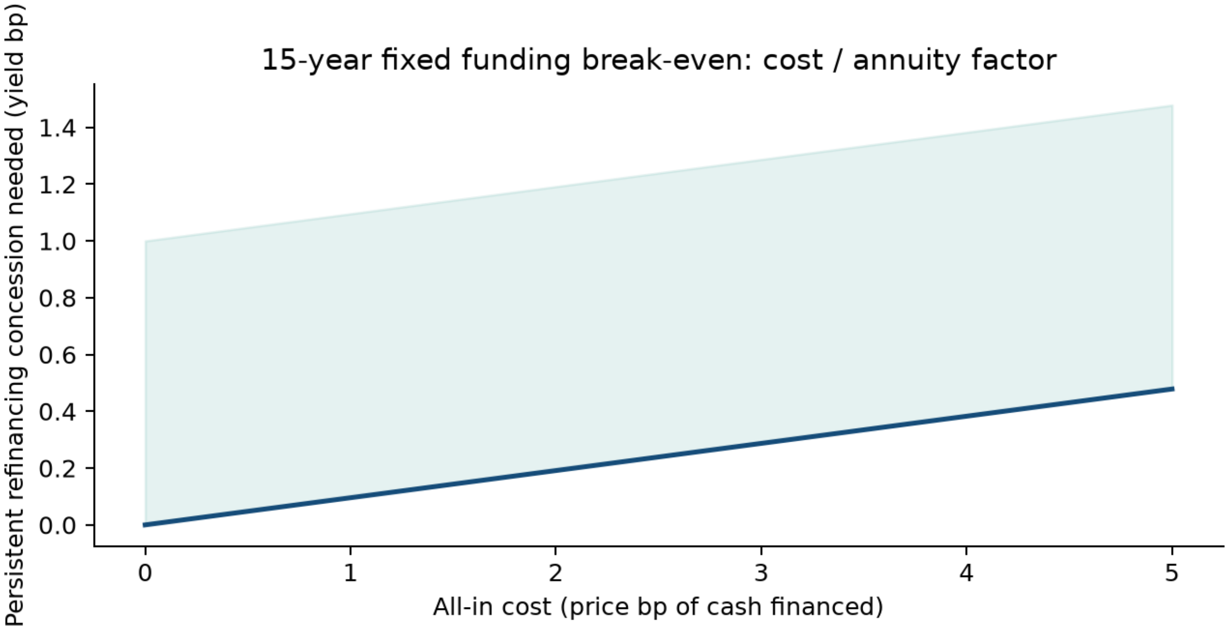

Figure 8. Persistent refinancing concession required to offset an assumed all-in cost, using a normalized 15-year fixed-funding annuity. This is a break-even sensitivity, not realized fiscal evidence.

The annuity factor is 10.451. An all-in cost of one price-basis-point of financed cash requires a persistent refinancing concession of approximately 0.096 yield-basis-points to break even in this simplified calculation. Temporary quote improvement alone does not establish that concession.

The $6bn-par hypothetical basket requires about $4.262bn cash at the actual cash/par ratio. Under the model's assumed term-supply concession, the normalized fixed-funding-equivalent NPV averages about $0.706m. This number is an assumption-driven sensitivity, not a Treasury saving estimate. It excludes a separately quantified value of better market functioning and any actual issuance concession.

The bill-financing-side one-year rate-surprise metric ranges from approximately -$27.2m to +$26.7m at the 5th–95th percentiles in the normal scenario. These are costs/benefits relative to a fixed short-rate baseline, not the full net benefit of the buyback. They cannot be added to the normalized 15-year metric.

The comparison shows why a small modeled static benefit can be dominated by funding risk. Duration-matched financing avoids the modeled duration extraction but may still improve liquidity. Cash drawdown uses a scarce buffer. A comprehensive cost–benefit study would value all three strategies with a common dynamic discounting and issuance framework.

21. Robustness and model validation

Table 18. Monte Carlo convergence

Path count	Mean day-one 10y effect	Monte Carlo standard error
10000	-1.8668bp	0.0088bp
20000	-1.8526bp	0.0062bp
40000	-1.8509bp	0.0044bp

The 20,000-to-40,000 change in the mean is approximately 0.0017bp. At 40,000 paths, Monte Carlo standard error is 0.0044bp. Numerical precision is therefore adequate for the shown two-decimal figures; it says nothing about whether the assumed policy coefficient is correct.

Table 19. Market-risk distribution robustness

Horizon	Gaussian P(10y rises)	Historical-block P(10y rises)
1 days	37.67%	37.77%
21 days	48.45%	49.08%
63 days	49.81%	50.32%
252 days	49.43%	51.82%

The robustness table uses 20,000 paths for both displayed calculations; the main tables use 40,000. Five-day vector blocks preserve some serial dependence and non-Gaussian shocks. Intervention coefficients are held fixed across methods, so this check tests market-risk distribution sensitivity, not policy validity.

A 2026 holdout check places 99.4% of daily 10-year changes inside the pre-2026 $\pm 1.96\sigma$ band, above nominal 95% coverage. This signals overdispersion in that holdout, not a perfectly calibrated model. Historical-volatility uncertainty spans 5.45–6.30bp across the block-resampled 5th–95th range.

Checks passed: PSD covariance; zero-size policy effect equals zero; duration-matched net DV01 equals zero; CUSIP accepted amounts reconcile; no crossed quotes or duplicate security dates; exact bond-price identity; small-shock duration/convexity error bound. Identification and calibration limitations remain despite these numerical checks.

22. Should buybacks be the main solution?

Table 20. Policy tools by problem

Problem / regime	Lead response	Role for buybacks
Normal off-the-run friction	Predictable liquidity support plus efficient issuance and intermediation.	Useful targeted tool; monitor costs, depth, and free-float scarcity.
Temporary funding squeeze	Appropriate repo/liquidity facilities and collateral/intermediation capacity, subject to eligibility.	Complement, not substitute for financing when the core problem is funding.
Persistent inflation or fiscal-risk repricing	Credible monetary/fiscal policy and sustainable financing strategy.	Cannot eliminate expected inflation or debt-supply risk merely by swapping bonds.
Systemic fire sales / dysfunctional market	A coordinated, mandate-consistent market-function response with triggers and exit plans.	Potential specially designed complement; the current routine program is not an acute-stress facility.
Long-run structural fragility	Clearing, transparency, settlement, and dealer-capacity improvements with attention to margin/liquidity trade-offs.	Helpful recurring outlet, not a substitute for market infrastructure.

QE can remove duration at scale, but it is a central-bank monetary-policy instrument with different balance-sheet effects and objectives. Issuance changes can influence duration supply without repurchasing old securities. Repo facilities primarily address financing. Central clearing can improve netting while introducing margin demands. None is a universally superior substitute in every regime.

Treasury’s current liquidity-support FAQ expressly excludes an acute-stress mandate. Duffie and Keane (2023) discuss how fiscal purchases could be designed for market-function contingencies, but that is not the same as claiming the September operation was such a facility.

Recommendation: use buybacks as a complement. They can be the primary operational instrument for a narrowly defined off-the-run liquidity objective, but not the main macroeconomic solution to elevated Treasury yields.

23. Implications for markets and decision rules

Treasuries and relative-value investors

A targeted operation may benefit eligible or closely substitutable bonds relative to the curve, but acceptance is endogenous and future eligibility can spread benefits beyond purchases. A duration-neutral relative-value position tests a different thesis from a directional long-duration trade. This report does not establish a tradable mispricing or recommend either position.

Equities, credit, mortgages, and the dollar

If a program genuinely reduces benchmark discount rates or volatility, it can influence other asset prices. A small local Treasury liquidity gain need not transmit materially to corporate spreads or equities. Mortgage rates depend on Treasury yields, mortgage spreads, and convexity hedging; the dollar depends on international relative rates and risk conditions. No estimated cross-asset beta is supplied, so no numerical equity, FX, or mortgage forecast is asserted.

What would change the conclusion?

Evidence of lower effective spreads and price impact in a credible matched intraday design, with sustained improvement and stable free float, would strengthen the liquidity case. Evidence that new issuance consistently achieves a measurable concession larger than purchase and funding costs would strengthen the fiscal case. Conversely, persistent quote deterioration, scarcity, weak execution, or offsetting duration issuance would weaken it.

Practical evaluation rule

Before scaling, specify the problem, benchmark execution, measure both eligible and comparable securities, track spillovers, and reconcile funding. Evaluate several operations rather than one favorable day. Use verified changes in trading costs and debt-service NPV rather than gross purchase size as the scorecard.

The defensible answer to “Did this fail?” is: it did not prevent a broad selloff, but preventing that selloff was not its stated objective, and this dataset cannot establish whether it causally succeeded or failed at liquidity support.

24. Limitations, glossary, and reproducibility

Table 21. Glossary

Term	Meaning
Basis point	0.01 percentage point of yield; 0.0001 in decimal yield.
Price-bp spread	Bid–ask width divided by mid-price × 10,000; not yield bp.
DV01	Market-value dollar sensitivity to one yield basis point.
Convexity	Second-order curvature of price versus yield.
Off-the-run	Previously issued security, no longer the most recent benchmark.
Term premium	Compensation for bearing interest-rate risk, distinct from expected short rates.
Counterfactual	Outcome under the same conditions without the policy.
Pair-bootstrap interval	Uncertainty from resampling matched pairs; not proof of causal identification.

Core limitations: one short event window; imperfect common support; endogenous acceptance; simple curve-relative yield proxy; no intraday depth, effective spreads, repo specialness, historical free float, or identified funding package. The quote contrast is suggestive but cannot establish realized liquidity savings. Monte Carlo effects are deliberately not fitted to that contrast.

The accompanying bundle contains analyze.py, build_report.py, dependency versions, source manifest with hashes, cleaned CSV inputs, results.json, full scenario outputs, matched pairs, and valuation checks. Static analysis can be rerun offline from saved inputs. Refresh scripts are optional and depend on source availability and existing Bloomberg access.

Reproduction: create a virtual environment, install requirements.txt, run **python analyze.py**, then **python build_report.py**. Run these in the extracted study directory. The seed is 20260912; named offsets separate market-path and robustness draws. Figures and tables are generated from stored data and results, not manually typed estimates.

Bibliographic access notes: full texts were read for Garbade–Rutherford, Duffie–Keane, Fleming, the D’Amico–King working paper, and the Federal Reserve curve paper. Some journal references were verified through metadata/abstracts; no inaccessible full-text quotation is fabricated. Incorrect or blocked retrievals are marked as excluded in the manifest.

References

Adrian, T., Crump, R. K., & Moench, E. (2013). Pricing the term structure with linear regressions. *Journal of Financial Economics, 110*(1), 110–138. https://doi.org/10.1016/j.jfineco.2013.04.009

Bloomberg L.P. (2026). *U.S. Treasury benchmark yields and CUSIP-level daily prices, bid/ask quotes, and yields* [Data set]. Bloomberg Desktop API via BluAPI. Retrieved September 12, 2026.

Board of Governors of the Federal Reserve System. (n.d.). *Market yield on U.S. Treasury securities at constant maturity: 2, 5, 10, 20, and 30 years* [Data set]. FRED, Federal Reserve Bank of St. Louis. Retrieved September 12, 2026. https://fred.stlouisfed.org/series/DGS10

Connolly, M. F., & Struby, E. (2024). Treasury buybacks, the Federal Reserve's portfolio, and changes in local supply. *Journal of Banking & Finance, 168*, Article 107286. https://doi.org/10.1016/j.jbankfin.2024.107286

D'Amico, S., & King, T. B. (2012). *Flow and stock effects of large-scale asset purchases: Evidence on the importance of local supply* (Finance and Economics Discussion Series 2012-44). Board of Governors of the Federal Reserve System. https://www.federalreserve.gov/pubs/feds/2012/201244/201244pap.pdf

D'Amico, S., & King, T. B. (2013). Flow and stock effects of large-scale treasury purchases: Evidence on the importance of local supply. *Journal of Financial Economics, 108*(2), 425–448. https://doi.org/10.1016/j.jfineco.2012.11.007

Duffie, D., & Keane, F. (2023). *Market-function asset purchases* (Staff Report No. 1054). Federal Reserve Bank of New York. https://www.newyorkfed.org/medialibrary/media/research/staff_reports/sr1054.pdf

Federal Reserve Bank of New York. (n.d.). *Treasury debt auctions and buybacks as fiscal agent.* Retrieved September 12, 2026. https://www.newyorkfed.org/markets/treasury-debt-auctions-and-buybacks-as-fiscal-agent

Fleming, M. J. (2003). Measuring Treasury market liquidity. *Economic Policy Review, 9*(3), 83–108. https://www.newyorkfed.org/medialibrary/media/research/epr/03v09n3/0309flempdf.pdf

Garbade, K. D., & Rutherford, M. (2007). *Buybacks in Treasury cash and debt management* (Staff Report No. 304). Federal Reserve Bank of New York. https://www.newyorkfed.org/medialibrary/media/research/staff_reports/sr304.pdf

Glasserman, P. (2003). *Monte Carlo methods in financial engineering.* Springer. https://doi.org/10.1007/978-0-387-21617-1

Gürkaynak, R. S., Sack, B., & Wright, J. H. (2006). *The U.S. Treasury yield curve: 1961 to the present* (Finance and Economics Discussion Series 2006-28). Board of Governors of the Federal Reserve System. https://www.federalreserve.gov/pubs/feds/2006/200628/200628pap.pdf

Han, B., Longstaff, F. A., & Merrill, C. (2007). The U.S. Treasury buyback auctions: The cost of retiring illiquid bonds. *The Journal of Finance, 62*(6), 2673–2693. https://doi.org/10.1111/j.1540-6261.2007.01289.x

Hu, G. X., Pan, J., & Wang, J. (2013). Noise as information for illiquidity. *The Journal of Finance, 68*(6), 2341–2382. https://doi.org/10.1111/jofi.12083

U.S. Department of the Treasury. (n.d.). *FAQs about Treasury securities buybacks.* Retrieved September 12, 2026. https://www.treasurydirect.gov/help-center/faqs/buyback-faqs/

U.S. Department of the Treasury. (2026a, August 5). *Minutes of the meeting of the Treasury Borrowing Advisory Committee.* https://home.treasury.gov/news/press-releases/sb0592

U.S. Department of the Treasury. (2026b, August 5). *Quarterly refunding statement.* https://home.treasury.gov/news/press-releases/sb0590

U.S. Department of the Treasury. (2026c, September 9). *Tentative schedule of Treasury buyback operations: August 2026 refunding quarter* [Revised schedule]. https://home.treasury.gov/system/files/221/Tentative-Buyback-Schedule.pdf

U.S. Department of the Treasury. (2026d, September 9). *Treasury debt buyback operation: Preliminary announcement for September 10, 2026.* https://www.treasurydirect.gov/instit/annceresult/press/preanre/2026/BBPA_20260910174000.pdf

U.S. Department of the Treasury. (2026e, September 10). *Treasury debt buyback operation: Final announcement.* https://www.treasurydirect.gov/instit/annceresult/press/preanre/2026/BBA_20260910174000.pdf

U.S. Department of the Treasury. (2026f, September 10). *Treasury debt buyback operation: Results.* https://www.treasurydirect.gov/instit/annceresult/press/preanre/2026/BBR_20260910174000.pdf

Vayanos, D., & Vila, J.-L. (2021). A preferred-habitat model of the term structure of interest rates. *Econometrica, 89*(1), 77–112. https://doi.org/10.3982/ECTA17440

Appendix A. Actual accepted-security calculations

Table 22. Accepted security-level risk calculations

CUSIP	Par ($m)	Clean price	Duration	DV01 ($k/bp)
912810SR0	2	60.615	11.96	1.5
912810SQ2	2	59.949	12.20	1.5
912810ST6	1,703	61.840	12.07	1,280.8
912810SW9	706	66.367	11.86	557.1
912810SY5	163	69.743	11.63	133.5
912810TA6	12	64.043	12.29	9.5
912810TC2	13	66.116	12.13	10.5
912810TF5	1,294	69.609	12.05	1,088.1
912810TH1	3	78.555	11.44	2.7
912810TK4	581	79.555	11.60	537.7
912810TM0	3	86.089	11.22	2.9
912810TQ1	3	84.452	11.53	2.9
912810TS7	28	84.216	11.51	27.5
912810TU2	3	89.566	11.47	3.1
912810RE0	51	80.609	12.12	50.0
912810UB2	303	91.953	11.50	325.7
912810UD8	3	85.945	12.02	3.1
912810UF3	19	91.646	11.69	20.7
912810UJ5	202	92.996	11.87	223.8
912810UL0	3	95.840	11.69	3.4
912810UN6	3	94.246	11.99	3.4
912810UQ9	78	91.187	12.06	87.2
912810UT3	9	91.098	12.30	10.1

Source: U.S. Department of the Treasury (2026f), rounded accepted prices; author ACT/ACT accrued-interest and semiannual cash-flow calculations. Duration and DV01 use reconstructed dirty prices and fitted yields. These are not official Treasury risk measures. The two largest purchases account for 57.78% of accepted par.

Appendix B. Transparent coefficient sensitivity

A scenario probability is not robust if changing an unsupported coefficient changes the conclusion. This appendix deliberately replaces the probability distribution with fixed alternatives. At day one, the 10-year maturity loading is approximately 0.775. With no duration-supply coefficient, the results follow directly from the policy equation.

Table 23. Fixed-coefficient alternatives

Assumed $6bn local liquidity effect	10-year policy effect	Meaning
-2bp (adverse)	+1.55bp	The operation worsens the local component.
0bp; duration coefficient also zero	0.00bp	No modeled effect.
+2bp	-1.55bp	Moderate assumed local improvement.
+5bp	-3.87bp	Stronger assumed improvement.

These alternatives are not estimated probabilities. They show that model outputs inherit assumptions. A different policymaker risk appetite, funding strategy, market depth, or free-float constraint could justify a very different range. No cross-sectional price result from this study identifies these coefficients.

Final answer

Treasury buybacks can be effective at a targeted market-function objective. They are not a stand-alone solution for macro-driven increases in yields. For September 10, the operation and acceptance amounts are verified; differential quote behavior is suggestive; a clean causal liquidity effect and a realized fiscal benefit are not established. The appropriate role is complementary, with effectiveness judged against explicit liquidity, execution, and financing objectives—not whether the 10-year yield happened to fall.

The PDF is the research report. The supporting bundle contains executable code, full numerical results, and source/audit records. No simulation is presented as proof of a historical causal effect.