

The Treasury General Account and the Fed's Balance Sheet

Roc Armenter

Federal Reserve Bank of Philadelphia and the Wharton School

WP 26-48

PUBLISHED

October 2026

ISSN: 1962-5361

Disclaimer: This Philadelphia Fed working paper represents preliminary research that is being circulated for discussion purposes. The views expressed in these papers are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Philadelphia or the Federal Reserve System. Any errors or omissions are the responsibility of the authors. Philadelphia Fed working papers are free to download at: <https://www.philadelphiafed.org/search-results/all-work?searchtype=working-papers>.

DOI: <https://doi.org/10.21799/frbp.wp.2026.48>

The Treasury General Account and the Fed’s Balance Sheet

Roc Armenter*

Federal Reserve Bank of Philadelphia

The Wharton School

September 23, 2026

Abstract

Changes in the Treasury’s cash management practices after the COVID pandemic increased the Fed’s balance sheet by as much as 2 percentage points of nominal GDP, including the larger buffer required to maintain ample reserve levels. A parsimonious time-series model shows that, once we account for the impact of the Treasury’s extraordinary measures during debt-ceiling episodes, the observed changes in the Treasury General Account since COVID are mostly driven by higher target balances. An effective sterilization program by the Fed would require timely and accurate information at high frequencies.

1 Introduction

In the last two decades, the Federal Reserve’s balance sheet has more than tripled in size, from around 6 percent of nominal GDP to over 21 percent in 2025. Granted, the balance sheet peaked at 35 percent of nominal GDP in 2022 and the FOMC has successfully reduced the Fed’s asset holdings by close to 2.5 trillion dollars. The normalization process, though, reached its end in December 2025, as reserves neared ample levels (Federal Open Market Committee 2025).

Why can’t the Fed’s balance sheet be reduced further? A lot of attention has been paid to the FOMC’s decision, formalized in 2019, to use an ample-reserves framework, which

*The views expressed here are mine and do not necessarily reflect those of the Federal Reserve Bank of Philadelphia or elsewhere in the Federal Reserve System. Contact: roc.armenter@phil.frb.org.

“ensures that control over the level of the federal funds rate and other short-term interest rates is exercised primarily through the setting of the Federal Reserve’s administered rates, and in which active management of the supply of reserves is not required.”¹ And indeed, the resulting increase in the supply of reserves accounts for more than half of the growth in the Fed’s balance sheet since 2007.

In this note, I turn instead to the lesser, but still significant, role of the Treasury General Account (TGA). In its service to the public, the Fed acts as the fiscal agent for the Treasury and hosts the Treasury’s cash balance as a liability in its balance sheet. Prior to 2008, the cash balance was small and stable, around \$5 billion, as most of the Treasury cash was deposited at commercial banks and other depository institutions through the Treasury Tax and Loan program. In the aftermath of the crisis, the Treasury altered its cash management practices (Santoro 2012), which resulted in a marked increase in the level and volatility of the TGA. A further change in the Treasury’s liquidity management mid-decade drove another increase in level and volatility; and after the TGA reached peak levels above a trillion during COVID, there was an apparent further jump in the average balance and volatility. A simple exercise suggests that the increase in volatility has also amplified the need for a larger buffer of reserves to maintain the ample-reserves framework. All told, a return to the pre-COVID level and volatility of the TGA could reduce the Fed’s balance sheet by up to 2 percentage points of nominal GDP.

A proper understanding of the Treasury’s cash management practices requires disentangling the effect of the extraordinary measures the Treasury enacts when the debt ceiling is binding. In particular, the Treasury runs down the TGA effectively to zero—though, to date, the debt ceiling has always been resolved before that happens. Quite obviously, the higher the initial TGA balance, the larger (and longer) the potential drawdown during a debt-ceiling episode. Once the effect of a larger target balance is accounted for, I show that the remaining variation in the TGA balance is, in log terms, roughly unchanged since the pre-COVID period—or perhaps has even decreased.

Vissing-Jorgensen (2026) argues that the Fed could undertake a sterilization program to offset the impact of TGA fluctuations on the supply of reserves. I assess some of the challenges to implement such a program. Because so much of the volatility in the TGA, outside extraordinary measures, occurs and resolves within weeks, the Fed would face stringent information requirements. Using the time-series model, I find that the scope for effective sterilization vanishes quickly if information is lagged by two or three weeks. A model tuned for closer in-sample fit would require attenuating the program’s response to TGA fluctua-

¹See Federal Open Market Committee (2019) for the January 2019 statement of principles and Logan (2023) for a detailed discussion of “ample” as an operational concept.

tions. I also briefly discuss some governance and communication considerations related to such a program.

There are several proposals to shift the Treasury’s liquidity holdings to the private sector in order to reduce the Fed’s balance sheet. Anderson et al. (2026) consider a modernized version of the historical Treasury Tax and Loan (TT&L) program, which would shift a portion of the TGA back to commercial banks—however, as the authors themselves note, doing so could shift the demand for reserves outward. Their conservative estimate of the Fed’s balance-sheet reduction is between \$200 and \$400 billion. Logan and Schulhofer-Wohl (2026) discuss the aforementioned options as well as the possibility for the Treasury to invest in reverse repos with private-sector counterparties, which may avoid increasing the demand for reserves, as many counterparties would not be depository institutions.

2 Treasury General Account

Figure 1 plots the weekly TGA balance from 2010 through the end of the sample, displayed as a 2-week moving average. The COVID period (2020–mid 2021) is omitted due to the exceptional situation, which resulted in unusually high TGA balance levels. Shading indicates periods during which the Treasury was employing EM in response to a binding debt ceiling.




Figure 1: **Treasury General Account balance, 2010–present.** The COVID period (2020–mid 2021) is omitted. Source: Federal Reserve H.4.1 via FRED; Treasury announcements.

Several data features are immediate. First, there is clearly a drift in both level and volatility. Second, the TGA balance precipitously declines during the extraordinary-measures pe-

riods, reflecting the Treasury’s drawdowns on the account. And equally apparent are the sharp recoveries following these periods, as the Treasury replenishes the TGA balance. Of course, the drop in the TGA balance during EM periods can last longer, and the resulting dynamics have a larger amplitude, the higher the initial balance.

There have been several documented changes in the Treasury’s cash management approach. In May 2015 (U.S. Department of the Treasury 2015), the Treasury adopted the recommendation by the Treasury Borrowing Advisory Committee (TBAC) of a minimum operating balance of \$150 billion for the TGA balance. Even a casual look hints that the Treasury kept balances substantially above the minimum: Outside the EM episodes, the TGA balance appears to hover above \$250 bn. between 2016 and 2019. After the extraordinary COVID-era buildup, during which the TGA peaked near \$1.8 tn. (not shown), the Treasury appears to operate with a much higher target for the TGA balance. The latest available Treasury announcement assumes a \$950 bn. balance for the end of 2026 Q3 (U.S. Department of the Treasury 2026).

For the remainder of the paper, I split the data into four data *segments*, capturing the announced change in Treasury cash management practices in 2015, as well as the apparent change after the COVID episode. Table 1 lists the four calendar segments used throughout the note, swiftly followed by Table 2, which summarizes the TGA balance (Panel A) and its weekly changes (Panel B) by segment, alongside the full-sample statistics with and without the COVID period. Figure 2 makes these level shifts visible, marking the average balance within each calendar segment. For the purpose of this note, I will focus on the two segments that are most relevant: the PRE segment (2016–2019) and the POST segment (mid 2021–present).

Table 1: **Data segments.** The four calendar segments used throughout the note. Source: Federal Reserve H.4.1 via FRED.

Segment	Start	End	Weeks
EARLY	2010-01-06	2015-12-30	313
PRE	2016-01-06	2019-12-25	208
COVID	2020-01-01	2021-05-26	74
POST	2021-06-02	2026-02-04	245

The tables and figure confirm the initial visual inspection: Both the average balance and its volatility have been increasing across the segments. From the PRE to the POST segment, average and standard deviation more than doubled—and more than quadrupled relative to EARLY. Table 2 also reports the statistics for weekly changes. First, note that the average change is positive but very small—there is a drift *within* each segment, but it is negligible.

Table 2: **TGA summary statistics by segment.** Levels (Panel A) and weekly changes (Panel B), in billions of USD. Segment definitions in Table 1. Source: Federal Reserve H.4.1 via FRED, author's calculations.

Segment	Mean	Std. Dev.	Min	Max
<i>Panel A: TGA level (billions of USD)</i>				
EARLY	77.5	51.4	11.7	290.5
PRE	273.4	89.6	37.6	429.0
COVID	1,227.0	489.0	351.0	1,816.7
POST	593.1	228.2	44.8	958.0
Full sample	377.6	390.4	11.7	1,816.7
Full ex-COVID	295.6	260.4	11.7	958.0
<i>Panel B: Weekly change (billions of USD)</i>				
EARLY	0.4	24.9	-61.8	99.6
PRE	0.4	33.9	-90.5	97.4
COVID	6.0	70.0	-211.4	269.3
POST	0.4	59.4	-132.8	232.9
Full sample	0.9	44.4	-211.4	269.3
Full ex-COVID	0.4	41.1	-132.8	232.9

Treasury General Account by segment

Dashed lines mark the average balance within each segment.




Figure 2: **Treasury General Account by segment.** Weekly balance (two-week moving average, shown faint in the background) with the four calendar segments separated by dotted vertical lines; the dashed line within each segment marks its average balance. The COVID segment is omitted. Source: Federal Reserve H.4.1 via FRED.

Second, weekly changes are more volatile in later data segments, but the increase in the standard deviation is less pronounced than for the levels. Finally, we may return to visual inspection and Figure 2 to note that the segment averages appear to be lower than the TGA balance most of the time—and markedly so in the PRE and POST segments.

Now, the mechanical effect of drawing down under EM, and then rebuilding, is exacerbating the volatility but pushing down the average balance. And since it is not the same to draw down a balance of \$250 bn. as one of \$750 bn., there are good reasons to think that some of the differences across segments are brought about by the EM periods. In Section 3 I take on the challenge of disentangling the changes in the Treasury’s cash management practices (target balance, volatility...) and the dynamics of extraordinary measures with a time-series model, which will also assist the discussion of sterilization options in Section 4. Before doing so, I provide a back-of-the-envelope calculation on the contribution of changes in the TGA balance to the Fed’s balance sheet size.

2.1 Balance Sheet Reduction: Napkin Math

Here I offer a simple back-of-the-envelope calculation to gauge the implications for the Fed’s balance sheet size from the observed changes in TGA management, including the size of the buffer needed to maintain ample reserves in the face of elevated volatility.

Starting with the level shift, we settle on \$750 billion as a working assumption for the target for the TGA balance in the POST segment—above the segment average but below the latest Treasury estimates for 2026. For the PRE segment, we map the formal minimum operating balance of \$150 billion to a target balance of about \$250 billion. The difference in levels is therefore \$500 billion, which will be the napkin number for the level shift.

As the Federal Reserve Bank of New York (the “Desk”) is set to maintain ample reserves, permanent shifts in the TGA target will have a one-to-one effect on the size of the Fed’s balance sheet. That is, the Desk needs to add assets to restore reserves to the system, $d\text{Assets} = d\text{TGA}$, and the Fed’s balance sheet expands by the same amount. Thus, a \$500 billion increase in the TGA target level implies a \$500 billion increase in the Fed’s balance sheet.

Moreover, volatility also has effects on the balance sheet size. The TGA is a liability of the Fed, as are reserve balances. On the asset side, total Fed assets adjust only through SOMA operations and loan facilities. It follows that, at high frequencies or simply holding assets fixed, a dollar increase in the TGA drains a dollar of reserves: $d\text{TGA} = -d\text{Reserves}$. Thus, the dynamics of the TGA map one-to-one to variation in the supply of reserves. When operating in an ample-reserves regime, the Desk needs to maintain a buffer to ensure that

reserves remain ample at all times. A lower and less volatile TGA therefore reduces the variance of reserve supply, which in turn implies the Desk can tolerate a smaller buffer of reserves.

To account for this effect, assume that there is a level of reserves supplied ex-ante by the Desk, but the effective level of reserves varies (negatively) with TGA balances. The Desk is given a fault tolerance $\epsilon\Delta > 0$ in determining the ex-ante level of reserves: That is, the Desk chooses a buffer such that the probability that effective reserves fall below some critical value during interval Δ is at most $\epsilon\Delta$.

With the right amount of heroic assumptions (see Appendix A for details), we can call on the reflection principle to obtain an expression for the buffer chosen by the Desk:

$$B = \sigma \sqrt{\Delta} \left(-\Phi^{-1}(\epsilon\Delta/2) \right), \quad (1)$$

where the length of a period Δ is two weeks, σ is the standard deviation of changes in the TGA, and Φ^{-1} is the inverse of the standard normal cumulative distribution function. The last term is simply the usual critical z value for one-sided tests.

To populate this formula, I assume a fault tolerance of about once every 200 months, which implies $-\Phi^{-1}(\epsilon/2)$ is between 3 and 4. The standard deviation of TGA changes in the last segment is $\sigma \approx \$60$ billion, which delivers a reasonable bracket for the buffer, somewhere between $B \approx 180$ billion and $B \approx 240$ billion. Given parameter and other sources of uncertainty, perhaps a bracket of \$200–\$300 billion is a more reasonable guess.

Critically, the buffer formula scales linearly with σ . There is a factor of a bit less than two in volatility of TGA changes across the PRE and POST segments, so the buffer guess-estimated for the PRE segment is \$100–\$150 billion, which seems plausible given the data. The difference in buffers is therefore on the order of \$100 billion.

Adding up, \$600 billion is less than 2 percentage points of nominal GDP in the Fed’s balance sheet. I view this number as justifying further analysis, even if it is nowhere close to the total growth in the Fed’s balance sheet from pre-2008 levels.

3 Modeling TGA balances

Next I introduce a simple time-series model that will allow disentangling the impact of the extraordinary measures from the underlying cash management practices. The model will be put to work further to provide a proof-of-concept for sterilization options in the next section.

3.1 Extraordinary Measures

The Treasury triggers the so-called “extraordinary measures” (EM) to avoid defaulting on financial obligations after the statutory debt limit, known as the “debt ceiling,” is reached. There is a host of those measures, but the one that is most relevant here is the practice of drawing down the TGA balance. Indeed, the infamous “X date” is defined as the date at which the Treasury would run out of cash in the TGA.

EM were in force in nearly a quarter of all weeks from 2010 to the present, across 11 distinct debt-ceiling episodes. About two-thirds of the EM episodes lasted more than a quarter of a year. The resulting TGA movements are large: within a single episode, the balance has swung by as much as \$522 billion peak to trough.

Indeed, EM episodes, and the following rapid TGA rebuilding phase, dominate the tail of the distribution of TGA changes, particularly at monthly frequencies. Figure 3 shows the distribution of 5-week TGA changes, separated by regimes (Normal, that is, no EM; Under EM; and Post-EM, given by the 10 weeks immediately following). The two key data segments, PRE and POST, are shown as top and bottom panels.

Two features stand out. First, the POST segment exhibits a substantially wider distribution of 5-week changes than the PRE segment, consistent with the increase in both average TGA level and its volatility, documented above. Second, EM periods are associated with large negative changes in both segments, while Post-EM periods contribute disproportionately to the right tail, reflecting the Treasury’s pattern of rapidly replenishing the TGA after the debt ceiling is raised.

Focusing on the bottom panel, showing the POST segment, it is clear that EM episodes dominate both tails at the monthly frequency. We should not assign the long tails to business-as-usual fluctuations in the TGA: In the next subsection, a model is proposed that accounts for the distinct behavior of TGA changes during EM episodes—and the following rebuilding phase.

3.2 Time-Series Model

Let $y_t = \log(\text{TGA}_t)$ denote the log of the weekly TGA balance. The *regime* $s_t \in \{0, 1\}$ is *observed*: $s_t = 0$ during extraordinary measures (EM) periods and $s_t = 1$ otherwise (Normal). Separately, the full sample is partitioned into $K = 4$ exogenous *segments* with known break dates, corresponding to distinct calendar periods previously discussed, with segment indices $k(t) \in \{1, \dots, 4\}$ and labels EARLY, PRE, COVID, and POST. Regime and segment are distinct constructs, and the model needs both. The *regime* s_t switches at high frequency with the debt-ceiling calendar and selects which of two data-generating




Figure 3: **Distribution of TGA changes, by regime and segment.** Five-week TGA changes, split by extraordinary-measures regime (Normal, Under EM, Post-EM), for the PRE (top) and POST (bottom) segments. Source: Federal Reserve H.4.1 via FRED, author's calculations.

processes governs the balance in a given week. The *segment* $k(t)$ is a fixed, low-frequency calendar partition that leaves the dynamics unchanged and shifts only the *level* to which the Normal regime reverts, capturing the Treasury’s occasional changes in its target balance.

The dynamics differ by regime. In the Normal regime, the log balance follows an autoregressive process with a segment-specific intercept $\mu_{k(t)}$,

$$y_t = \mu_{k(t)} + \Phi(t) + \sum_{j=1}^4 \psi_j y_{t-j} + \varepsilon_t, \quad \varepsilon_t \sim N(0, \sigma^2), \quad (2)$$

while in the EM regime, the TGA balance follows a unit root with a downward drift $\gamma > 0$,

$$y_t = -\gamma + \Phi(t) + y_{t-1} + \varepsilon_t, \quad \text{equivalently} \quad \Delta y_t = -\gamma + \Phi(t) + \varepsilon_t. \quad (3)$$

Note that the log specification prevents negative balances: In the model, an EM episode can go on indefinitely as the remaining balance is always reduced by a constant fraction each week, so it asymptotically approaches zero without ever becoming negative.

Across segments, the only parameters changing are the Normal-regime intercepts μ_k ; shifting μ_k moves the implied target balance for segment k . Everything else is common across segments and regimes: the AR coefficients $\psi_1, \dots, \psi_4$, the EM drift γ (with $\gamma > 0$ a weekly drawdown), the innovation variance σ^2 , and the seasonal controls $\Phi(t)$.² We use quarterly rather than monthly effects here for parsimony: While finer controls would capture more seasonality, including calendar-day effects, they also tend to absorb some EM dynamics, which cluster in particular months. For example, close to 40 percent of all business days in May and June fall within an extraordinary-measures drawdown period. The unconditional mean of the Normal regime for segment k , absent seasonal effects and emergency measures, is $y_k^* = \mu_k / (1 - \sum_j \psi_j)$.

The model is estimated at weekly frequency by maximum likelihood, conditioning on the first four observations. It has $K + 9 = 13$ free parameters estimated jointly—the $K = 4$ segment intercepts, the drift γ , three quarter fixed effects, four AR coefficients, and σ^2 . Because the regime sequence s_t is observed, the transition probabilities p_{11} and p_{00} are simply equated to their sample estimates.

Table 3 reports the parameter estimates. The autoregression is dominated by the first lag, $\hat{\psi}_1 = 0.930$, with the remaining lags small; the coefficients sum to $\sum_j \hat{\psi}_j = 0.814$, implying

²This last is a deterministic quarterly seasonal,

$$\Phi(t) = \alpha_2 \mathbf{1}\{Q2(t)\} + \alpha_3 \mathbf{1}\{Q3(t)\} + \alpha_4 \mathbf{1}\{Q4(t)\},$$

with the first quarter the omitted baseline and each week assigned to the quarter holding the majority of its days.

rapid mean reversion within the Normal regime (a shock has a half-life below a month). The innovation standard deviation is $\hat{\sigma} = \sqrt{0.068} \approx 0.261$ in log units. The EM drift is positive and significant, $\hat{\gamma} = 0.056$: during extraordinary measures, the TGA is drawn down at roughly 5.6% per week in log terms.³

The segment intercepts imply different target balances, which can now be defined as the level the TGA would settle at were the Normal regime to persist indefinitely. In dollars, the target for the PRE period is a bit less than \$290 bn., which is modestly above the raw average for that segment. In the POST period, the gap is substantially larger: the model returns a target of about \$750 bn., while the raw average for the segment is near \$590 bn.

Note that a single parameter governs the innovation volatility across segments and regimes. Because the model is log-normal, though, the volatility in levels depends on the level itself. Indeed, the standard deviation in levels scales with the level: Across all four segments, the standard deviation in levels is roughly half the average in levels.

The estimated transition probabilities, $\hat{p}_{11} = 0.983$ and $\hat{p}_{00} = 0.946$, imply expected durations of about 58 weeks in the Normal regime and 18 weeks in the EM regime.

Turning to the model fit, Figure 4 plots the observed against the one-step-ahead fitted value in log units, colored by segment. The points cluster on the 45-degree line—tightly for the PRE and POST segments.⁴ The fit is looser for the EARLY period (2010–2015), as the log specification struggles with the lower dollar amounts. It is important to note that, in dollars, the heteroskedasticity visible in Figure 4 *reverses*, with residuals being larger in later segments.

The model reproduces the central levels well—the segment means line up—but the single pooled innovation variance shows its limits on dispersion: because σ is common across segments, the model *understates* the volatility of the turbulent EARLY period and *overstates* that of the later data segments. Note as well that the model views EM episodes as completely exogenous and unpredictable: Given the short data segments, one more or less occurrence of EM, or a more or less lengthy episode, can have a sizable effect on the sample moments.

For the exercise at hand, we want to check that indeed the distribution of TGA changes, outside EM episodes, is similar across the PRE and POST segments—with the observed differences stemming mostly from the different base levels of the TGA rather than from changes in the underlying dynamics. Figure 5 plots the one-step-ahead residuals for the PRE and POST segments on Normal-regime weeks, in log units (top) and dollars (bottom). On the log scale, the two distributions are close and centered at zero: restricted to business-as-

³While significant, the drift parameter is quite fragile to the time-series specification. For example, monthly fixed effects can absorb too much of the drawdown and turn the drift parameter γ insignificant.

⁴The COVID segment (2020–mid 2021) is omitted from the figure, as elsewhere in the note, because its exceptional balances would compress the rest of the series; the segment is retained in estimation.

Table 3: **Model parameter estimates.** AR(4) with quarter fixed effects. Standard errors in parentheses.

	Estimate	SE
NORMAL REGIME		
<i>Panel A: Segment intercepts</i>		
μ_1	2.027	(0.259)
μ_2	2.320	(0.298)
μ_3	2.574	(0.329)
μ_4	2.495	(0.316)
<i>Panel B: AR coefficients</i>		
ψ_1	0.930	(0.041)
ψ_2	-0.172	(0.056)
ψ_3	-0.061	(0.056)
ψ_4	0.117	(0.039)
EM REGIME		
<i>Panel C: Drift</i>		
γ	0.056	(0.025)
COMMON ACROSS REGIMES		
<i>Panel D: Quarter fixed effects</i>		
Q2 (Apr–Jun)	0.056	(0.026)
Q3 (Jul–Sep)	-0.008	(0.026)
Q4 (Oct–Dec)	0.037	(0.026)
<i>Panel E: Variance and transitions</i>		
σ^2	0.068	(0.003)
p_{11} (Normal $\rightarrow$ Normal)	0.983	—
p_{00} (EM $\rightarrow$ EM)	0.946	—
Log-likelihood	704.65	
Parameters	13	
Observations	840	
Effective obs.	836	
AIC	-1383.3	
BIC	-1321.8	

Notes: Standard errors are reported in parentheses. Transition probabilities for the regimes are estimated by counting occurrences. The model is estimated at weekly frequency over the sample period 2010–2026. Source: Federal Reserve H.4.1 via FRED, author’s calculations.





Figure 4: **Observed vs. one-step fitted TGA, by segment.** The COVID segment is omitted. Source: Federal Reserve H.4.1 via FRED, author's calculations.


usual weeks, the PRE and POST residuals are overall very similar and, indeed, statistically indistinguishable.⁵ On the dollar scale, by contrast, the POST distribution is markedly wider, because identical proportional innovations translate into larger dollar swings when the TGA base is larger.

Model Residuals: PRE vs POST (Normal weeks)

Log scale



Dollar scale



One-step-ahead residuals on Normal-regime weeks; dollar residual = observed - fitted TGA. Source: author's calculations.

Figure 5: **Model residuals, PRE vs. POST (Normal weeks)**. One-step-ahead residuals for the PRE (2016–2019) and POST (2021–present) segments, on Normal-regime weeks. Top: log scale. Bottom: dollar scale (billions of USD; observed minus fitted TGA). The log-scale distributions nearly coincide, whereas the dollar-scale POST distribution is wider, reflecting its larger base. Source: author's calculations.

There are some differences in skewness between the PRE and POST distributions that may be worth further research—the implications for rate control are not symmetric. There is also some autocorrelation in the residuals, a bit more marked in the POST segment. The model can be extended to address both, at a substantial cost of parsimony, by allowing a third regime that governs the transition back from EM to Normal. Similarly, segment-specific

⁵Over all weeks (including debt-ceiling episodes) a two-sample Kolmogorov–Smirnov test does reject equality ($D = 0.14$, $p = 0.02$), but the difference lives entirely in the extraordinary-measures tails: on Normal weeks the test is far from significant ($D = 0.11$, $p = 0.33$).

volatility and dynamics would naturally improve the in-sample fit. In particular, the volatile EARLY period somewhat inflates the estimate for volatility of innovations, which is pooled across segments. If anything, the volatility in TGA log-changes has decreased over time.

4 Sterilization Options

As established in Section 2, every dollar of TGA variation drains a dollar of reserves unless the Fed conducts an offsetting operation. *Sterilization* refers to precisely such operations: the Desk injects (or withdraws) reserves to counteract the drain (or injection) caused by a rising (or falling) TGA, so that the supply of reserves is insulated from the Treasury’s cash management.

The Desk could conduct small, high-frequency purchases or sales of Treasury bills to offset anticipated TGA flows, as proposed by Vissing-Jorgensen (2026). Because the purchased assets would have short maturities, such operations would carry negligible effects on duration and could be unwound quickly. As discussed in the introduction, other proposals shift the management of the TGA away from the Fed.

My focus here is on the practical feasibility of sterilization operations by the Desk. As shown in the next subsection, the Desk would face a high-frequency challenge in offsetting TGA fluctuations: much of the variation in the TGA balance occurs and must be addressed on very short notice.

4.1 A High-Frequency Challenge

Sterilization through market operations confronts a basic difficulty: much of the action in the TGA balances occurs week-to-week, with a short half-life, so offsetting operations must be both timely and responsive.

The first piece of evidence is model-free: Figure 6 shows the standard deviation of the h -week change in the log TGA balance, computed on Normal-regime weeks only so that debt-ceiling drawdowns do not contaminate the measure, separately for the PRE and POST segments. Two features stand out. First, dispersion accumulates quickly and then plateaus: in log terms, the standard deviation more than doubles between the one- and six-week horizons, after which it flattens near its unconditional level. Most of the variance the Desk would have to offset therefore builds up within the first month or so, and an operation that reacts with a lag of even a few weeks has already missed the bulk of it. Second, the log-scale profiles for PRE and POST nearly coincide (top panel): the *proportional* volatility of the TGA is essentially unchanged across the two segments. In dollars (bottom panel), however,

the POST profile sits roughly twice as high as PRE at every horizon, because the same proportional swings now apply to a base two to three times larger. This is the dollar-scale volatility increase invoked in the buffer calculation of Section 2, here shown directly and without any model.




Figure 6: **Standard deviation of weekly TGA changes, on Normal regime only.** Standard deviation of the h -week change in the TGA, on Normal-regime (non-episode) weeks only, by horizon h , for the PRE and POST segments. Top panel: log scale; bottom panel: dollars. Source: Federal Reserve H.4.1 via FRED, author's calculations.

Looking at the raw volatility takes the perspective of a *naïve* Desk that simply matches the TGA level it last observed, making no use of the predictable structure in the series. Next, I deploy the model to illustrate the resulting informational demands on the Desk.

Suppose the Desk observes the TGA only through week $t - L$ —an *information lag* L reflecting publication and settlement delays—but forms the statistically optimal forecast of the current balance from the time-series model of Section 3, and offsets a fraction $a \in [0, 1]$ of the forecast deviation of reserves from their target. Let the *reserve gap* D_t be the residual deviation of reserves from target after this operation. In log terms, its variance decomposes

exactly into two pieces,

$$\text{Var}(D_t) = \underbrace{\text{FEV}(L)}_{\text{information-lag floor}} + \underbrace{(1-a)^2 [V_{\text{unc}} - \text{FEV}(L)]}_{\text{attenuation penalty}}, \quad (4)$$

where V_{unc} is the unconditional variance of the log TGA balance⁶ and $\text{FEV}(L)$ is the L -step-ahead forecast-error variance of the model. The two terms isolate the two constraints on sterilization. The first, the *information-lag floor*, is what a Desk operating L weeks behind cannot offset even at full participation ($a = 1$); it is the variance of the L -step forecast error and does not depend on a . The second, the *attenuation penalty*, is what the Desk forfeits by offsetting only partially ($a < 1$), and it falls on the *forecastable* portion of the variance, $V_{\text{unc}} - \text{FEV}(L)$.

Figure 7 plots the floor as a share of the unconditional variance, $\text{FEV}(L)/V_{\text{unc}}$, against the information lag L . This object is now model-specific. The share rises steeply: if the Desk is just one week behind the data, it would already forfeit about a third of the variance ($\text{FEV}(1)/V_{\text{unc}} \approx 34\%$), two weeks behind about three-fifths, and by three to four weeks behind roughly ninety percent. The model thus inherits this feature of the raw data: the bulk of the variance is concentrated in short-term fluctuations.

The attenuation penalty in (4) is, by contrast, second-order. Cutting the offset fraction a from 1 to 0.25 raises the reserve-gap standard deviation by about 45% when information is fresh ($L = 1$), but by only about 1% once it is stale ($L = 8$). The reason is in the decomposition: the penalty multiplies the forecastable variance $V_{\text{unc}} - \text{FEV}(L)$, which the information lag has already eroded, so its bite is itself contingent on L . In the realistic regime where the Desk is at least a few weeks behind, how aggressively it trades barely matters; the lag dominates.

Now, my model was designed to disentangle the effect of EM, favoring parsimony over fit. As a result, I cannot claim that the model achieves the lowest information floor: it is not optimized to minimize the fraction of unforecastable variance. Moreover, there is no question that the Desk has access to further historical data as well as other sources of information that could improve forecasts beyond what my parsimonious model captures. On the other hand, a closer fit in-sample will call for a larger attenuation factor.

⁶This is the same unconditional variance V_{unc} that sets the segment target balances in Section 3.2: there it fixes the level of the lognormal mean, here the scale of the forecast-error decomposition.




Figure 7: **The information-lag floor.** Share of the unconditional log TGA variance that a Desk operating L weeks behind cannot offset even at full participation, $FEV(L)/V_{unc}$, versus the information lag L (weeks). The dashed line marks the unconditional variance (the no-information ceiling). Source: author's calculations.

4.2 Governance and Communications

The Fed (the Reserve Banks, to be precise) serves as the fiscal agent for the Treasury, managing the TGA and processing daily financial transactions, facilitating the auction, issuance, and redemption of marketable securities and U.S. savings bonds, and much more.

There is neither contradiction nor confusion between the Fed's roles as the independent authority for monetary policy, and as fiscal agent for the Treasury. Several proposals discussed above call for the Treasury to manage at least some of its liquidity needs *outside* the Fed, either returning to the use of commercial bank accounts or investing in reverse repos with the private sector. Diminishing the Fed's role as fiscal agent poses no governance or communication challenges.

A sterilization program would see the Desk, and hence the Fed, conducting market operations that are explicitly a counteraction to the flux of revenues and expenses of the Treasury. Moreover, an effective sterilization program would require coordinating and frequently sharing information with the Treasury. These two features would be shared with deficit monetization: It is imperative that neither contradiction nor confusion seep in, nor the appearance of either.

In my view, governance is the foundation of effective communications—both internal and

external. The Treasury department must be actively and publicly involved as the leading stakeholder in any sterilization program, undertaken exclusively for the optimal management of public funds. Similarly, Treasury should have discretion regarding the design and parameters of the program, akin to how a host of Treasury services already offered by the Fed are carefully specified and managed under Treasury guidance.⁷

With the Treasury actively and publicly in charge, all communications will have a head start. The Desk can receive instructions from the Treasury, along with the information needed to deliver such instructions. Whether the Treasury is crossing the t's and dotting the i's, or chooses instead to delegate operational details, the flow of information will remain tied to the service provided. The amount of transparency—itself a question of design but also subject to the law—will also not interfere with the needs elsewhere.

Public communications can then be clear on the role of the Treasury as the principal and the Fed as an agent; and the nature of the arrangement as a fiscal service is unlikely to attract ongoing attention. Any reduction to the Fed's balance sheet will simply accrue over time as the Desk continues to “maintain ample reserves.” There would be no need to tie the management of the TGA to the size of the Fed's balance sheet.

5 Conclusion

This note has documented and quantified the contribution of the Treasury General Account to the size and volatility of the Federal Reserve's balance sheet. A simple exercise suggests the potential for a non-negligible reduction in the supply of reserves from a return to the pre-COVID TGA level and volatility. A parsimonious model allowed disentangling the effects of extraordinary measures and pinpointing the changes in TGA target and volatility. Data and model suggest that the Desk would face stringent information demands to effectively sterilize fluctuations in the TGA balances.

The obvious problem with a sterilization program is its optics, as it may be perceived as deficit monetization. Perhaps this alone favors proposals for the Treasury to manage its liquidity outside the Fed. Alternatively, the sterilization would be best structured as a fiscal service provided by the Fed, as agent, to the Treasury, as principal. Doing so would enable solid public communications that minimize misperceptions regarding monetary and fiscal policy.

⁷The Treasury Financial Manual clearly states that “The Secretary of the Treasury under Section 15 of the Federal Reserve Act, as amended (12 U.S.C. 391), is authorized to require the FRBs to act as depositories and fiscal agents of the United States and may prescribe regulations in connection with these responsibilities.” Treasury Financial Manual.

A Buffer Derivation

This appendix derives the buffer formula (1) used in Section 2.

Time is continuous and indexed by $t \geq 0$. At regular intervals of length Δ , the central bank supplies “ex-ante” reserves $M(t)$, which are fixed over the interval $[t, t + \Delta)$. The problem is identical across intervals, so we can focus on the first interval $[0, \Delta)$ without loss of generality.

At all times, autonomous factors determine the effective supply according to

$$\hat{M}(t) = M(t) + \sigma X(t), \quad (5)$$

where $X(t)$ is a standard Wiener process, with zero mean (without loss of generality) and unit variance.

There exists a fixed threshold M^* such that if, at any time t , $\hat{M}(t) < M^*$, an undesirable failure event arises. A key advantage of the formulation here is that we can obtain a guess-estimate for the buffer without needing to elaborate on the value of M^* or the nature of the failure event.

We assume the central bank will tolerate a small failure rate ϵ , scaling linearly with the interval Δ .⁸ That is, the central bank aims to minimize $M(0)$ subject to

$$\Pr\left(\min_{t \leq \Delta} \hat{M}(t) \leq M^*\right) \leq \epsilon \Delta. \quad (6)$$

Note that failure is one-sided, unlike Haubrich (2023), which considers a similar “inventory management” problem but with two-sided costs.

Substituting (5) and defining the normalized buffer $b = (M(0) - M^*)/\sigma > 0$, the constraint becomes

$$\Pr\left(\min_{t \leq \Delta} X(t) \leq -b\right) \leq \epsilon \Delta. \quad (7)$$

By the reflection principle for a standard Wiener process starting at zero,

$$\Pr\left(\min_{t \leq \Delta} X(t) \leq -b\right) = 2 \Phi\left(\frac{-b}{\sqrt{\Delta}}\right), \quad (8)$$

where Φ is the standard normal CDF. Rather trivially, the central bank will set initial

⁸Alternatively, we could specify the fault probability directly. As long as the interval Δ is commensurably short, the two approaches are equivalent.

reserves such that (6) is binding. We can now solve for b :

$$b = \sqrt{\Delta} \left(-\Phi^{-1}(\epsilon\Delta/2) \right). \quad (9)$$

Multiplying by σ to return to dollar units gives

$$B = \sigma \sqrt{\Delta} \left(-\Phi^{-1}(\epsilon\Delta/2) \right), \quad (10)$$

which is equation (1) in the main text.

There are several manifestly heroic assumptions in this rough back-of-the-envelope calculation. That said, the standard Wiener process appears to fit the TGA changes reasonably well up to horizons of a month or two. Figure 8 shows the variance of TGA changes by horizon, which grows approximately linearly up to 8 weeks, consistent with the Wiener process assumption. Past that horizon, there seems to be some mean-reversion (or possibly seasonality), and the derived formula may be substantially off for longer horizons without further adjustments.




Figure 8: **Variance of TGA changes by horizon.** Variance of the h -week change in the TGA against the horizon h ; the near-linear growth up to about eight weeks is consistent with the Wiener-process approximation. Source: Federal Reserve H.4.1 via FRED, author's calculations.

References

- Anderson, Alyssa G. et al. (2026). *A User’s Guide to Reducing the Federal Reserve’s Balance Sheet*. Finance and Economics Discussion Series 2026–019. Board of Governors of the Federal Reserve System. DOI: 10.17016/FEDS.2026.019.
- Federal Open Market Committee (2019). *Statement Regarding Monetary Policy Implementation and Balance Sheet Normalization*. Press Release. Board of Governors of the Federal Reserve System.
- (2025). *FOMC Statement*. Press Release. Board of Governors of the Federal Reserve System.
- Haubrich, Joseph G (2023). *Federal Reserve Balance-Sheet Policy in an Ample Reserves Framework: An Inventory Approach*. WP 23-25. Federal Reserve Bank of Cleveland.
- Logan, Lorie (2023). “Ample Reserves and the Friedman Rule”. In: ECB Conference on Money Markets. European Central Bank.
- Logan, Lorie and Sam Schulhofer-Wohl (2026). *Options for Reducing the Size of the Fed’s Balance Sheet*. In Depth. Federal Reserve Bank of Dallas.
- Santoro, Paul J. (2012). “The Evolution of Treasury Cash Management During the Financial Crisis”. In: *Current Issues in Economics and Finance* 18.3. DOI: 10.2139/ssrn.2047088.
- U.S. Department of the Treasury (2015). *Quarterly Refunding Statement of Acting Assistant Secretary for Financial Markets Seth B. Carpenter*. Press Release JL-10045. U.S. Department of the Treasury.
- (2026). *Treasury Announces Marketable Borrowing Estimates*. Press Release SB-0485. U.S. Department of the Treasury.
- Vissing-Jorgensen, Annette (2026). “Fluctuations in the Treasury General Account and Their Effect on the Fed’s Balance Sheet”. In: *FEDS Notes*. DOI: 10.17016/2380-7172.3873.