Federal Reserve Bank of Dallas

An Asset-Liability Management Approach to the Federal Reserve Balance Sheet

Hugo De Vere, Srini Ramaswamy and Sam Schulhofer-Wohl

Working Paper 2525 **July 2025**

Research Department

https://doi.org/10.24149/wp2525

Working papers from the Federal Reserve Bank of Dallas are preliminary drafts circulated for professional comment. The views in this paper are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Dallas or the Federal Reserve System. Any errors or omissions are the responsibility of the authors.

An Asset-Liability Management Approach to the Federal Reserve Balance Sheet*

Hugo De Vere†, Srini Ramaswamy‡ and Sam Schulhofer-Wohl§

June 30, 2025

Abstract

The Federal Reserve's liabilities include a mix of floating-rate instruments, such as reserves, and long-duration, non-interest-bearing instruments, such as currency. We investigate the implications of an asset-liability management approach to choosing assets to back these liabilities, with a focus on matching the duration of assets and liabilities. We study the net income volatility and mark-to-market volatility of several different asset maturity ladders using a Monte Carlo simulation of future interest rate paths. Short-duration ladders minimize net income volatility when paired with floating-rate liabilities but maximize it when paired with currency. Long-duration ladders minimize income volatility when combined with currency and also minimize the volatility of the economic value of assets net of liabilities in that case, but at the expense of higher mark-to-market asset volatility. We discuss why barbells that combine long- and short-duration strategies produce much lower income volatility than ladders of similar average duration, when liabilities have a mix of long and short durations. However, a barbell could be challenging to implement at scale. We find that an ”across-the-curve” strategy of buying securities in proportion to outstanding amounts generates somewhat less income volatility than a laddered portfolio, though still more than the barbell portfolio.

JEL Classification: E52, E58

Keywords: Central Bank, Monetary Policy, Federal Reserve Balance Sheet, Asset and Liability Management

*The views expressed in this paper are those of the authors, and do not necessarily reflect those of the Federal Reserve Bank of Dallas or the Federal Reserve System.

†Hugo De Vere, Federal Reserve Bank of Dallas, Hugo.DeVere@dal.frb.org.

‡Srini Ramaswamy, Federal Reserve Bank of Dallas, srini.ramaswamy@dal.frb.org.

§Sam Schulhofer-Wohl, Federal Reserve Bank of Dallas, samuel.schulhofer-wohl@dal.frb.org.

1 Introduction

Popular discussions often apply commercial-bank risk metrics to the Federal Reserve's balance sheet. Such comparisons are imperfect because the Fed differs from commercial banks in two important ways.

- The Fed does not seek to maximize profits but rather uses its balance sheet to advance its public mission, particularly the monetary policy dual mandate of maximum employment and stable prices.
- The Fed's main liabilities, bank reserves and currency, are redeemable only with these same liabilities, which the Fed itself creates. That is, the Fed is a "payment in kind" issuer. Unlike a commercial bank, the Fed therefore faces no risk of being unable to pay its debts.

Notwithstanding these differences, the volatility of income from the Fed's balance sheet can complicate the Fed's communications. Some policymakers and other observers have suggested mitigating this challenge by matching the duration of the Fed's assets to that of its liabilities [17], not unlike approaches used by commercial banks and other financial institutions. In this note, we examine what asset-liability matching would entail for the composition of the Fed's assets and some of its consequences with regards to the volatility of future income as well as the mark-to-market volatility of the Fed's balance sheet assets. Consistent with the Federal Open Market Committee's stated intention to hold primarily Treasury securities in the long run [11], we examine asset portfolios consisting only of Treasuries and do not examine the behavior of mortgage-backed securities, which the Fed also currently holds. While the purpose of our paper is to understand how asset-liability management could be implemented, we emphasize that minimizing volatility of income and/or mark-to-market valuations are likely not the only or the most important objectives for the balance sheet given the Fed's overarching public policy mission.

The outline of this paper is as follows. Section 2 of the paper reviews how commercial banks match assets and liabilities. A key element, and the one most applicable to the Fed, is matching the interest rate or duration risk of assets and liabilities. Section 3 examines the interest rate risks on the liability side of the Fed's balance sheet. We show that the Fed's liabilities can be generally grouped into two categories: floating-rate liabilities with no duration, such as bank reserves, and non-interest-bearing liabilities with very long duration, such as currency.

Section 4 examines how different assets could be matched to these categories of liabilities. Backing floating-rate liabilities with floating-rate assets minimizes the volatility of both net income and mark-to-market valuation for the floating-rate set of liabilities. For long-duration liabilities, there is a trade-off between the volatility of net income and the volatility of mark-to-market asset valuations; backing long-duration liabilities with long-duration assets reduces the volatility of both net income and the net value of assets minus liabilities, but shows higher mark-to-market volatility when considering the asset side of the balance sheet on its own.

These findings suggest intuitively that a barbell portfolio of very-short-duration and very-long-duration assets would most directly match the Fed's mix of short-duration and long-duration liabilities. Indeed, we find that it is not just the average duration of the asset portfolio but the mix of durations in it that matters for the Fed's net income volatility. Assuming the Fed holds securities to maturity rather than actively trading to rehedge its portfolio, a barbell portfolio has only about 10 percent as much net income volatility as a rolling ladder of securities with the same average duration as the barbell. The reason is that the duration of an asset is a measure of how its market value responds to interest rates, and correspondingly duration-matching mitigates mark-to-market volatility. By contrast, mitigating income volatility requires matching the re-couponing of assets to the re-couponing of liabilities. The market value of an asset depends on the present value of its future cash flows, so many coupon (and principal) sequences can produce a given level of duration. But given a barbell of floating-rate liabilities that re-coupon often and long-duration liabilities such as currency that do not re-coupon at all, a barbell that matches re-couponing frequency (by blending short-duration and long-duration assets in the right weights) ought to have lower income volatility.

A barbell could be challenging to implement in practice, as it would require the Fed to concentrate its holdings in certain maturity sectors. We find that an across-the-curve investment portfolio can have somewhat lower income volatility than a duration-equivalent rolling ladder, though not as low as the barbell. Our focus in section 4 is on a long-run or steady state Fed balance sheet, in which the Fed is neither actively conducting large-scale active purchases for monetary policy or financial stability purposes, nor normalizing its balance sheet following an episode of such purchases. Section 5 concludes by discussing how the potential for episodes of large-scale asset purchases might change the analysis, as well as some broader considerations.

2 Key Elements of Asset-Liability Management

The basic idea behind traditional bank asset-liability management is to match the most important risk characteristics of a bank's assets with the corresponding characteristics of its liabilities, such that the bank has sufficient assets to meet its liabilities across potential states of the world. At a high level, two of the most important characteristics are interest rate risk and liquidity.

- **Interest rate risk**: To protect the economic value of its balance sheet from swings in interest rates, a bank may seek to match the effect of interest rate fluctuations on the present value of its assets to the effect of the same fluctuations on the present value of its liabilities. This strategy is sometimes loosely referred to as duration matching or managing a bank's duration of equity. This is an easy notion to state, but difficult to apply in practice because many large balance sheet items often have poorly observed prices. Deposit liabilities are a particular challenge. Because banks typically can retain deposits without raising their interest rates one for one when market rates rise, deposits have an economic value to a bank that increases as market rates rise. However, depositors have the right to withdraw at par at any time; miscalculating the elasticity of deposits with respect to interest rates can lead to errors in estimating both the value of deposits and their duration (see [12]).

- **Liquidity risk**: Commercial banks must have sufficient reserves, or other readily monetizable assets, to meet potential deposit redemptions and other outflows.

In the Fed's current ample-reserves monetary policy implementation framework, the Fed can always create additional reserves without impinging on its other objectives [14] and thus does not need to manage the liquidity risk of its assets. Therefore, we explore the notion of matching duration risk in the Fed's assets to that of its liabilities. We begin by examining the duration of the Fed's liabilities and then turn to the asset side of the balance sheet.

3 Interest Rate Risk in Federal Reserve Liabilities

Figure 1 shows the composition of the Fed's major liabilities through time. The predominant liabilities are reserves, which are held by banks; balances in the Fed's reverse repurchase agreement facilities, which are available to counterparties such as government-sponsored enterprises, money market mutual funds, and foreign central banks; currency in circulation (CIC); and the Treasury General Account (TGA), which the U.S. Treasury Department uses to make and receive federal expenditures and revenues. Together, these four items account for 99 percent of current liabilities and have accounted for about 95 percent of total liabilities on average for more than two decades.

Figure 1: Federal Reserve Liabilities

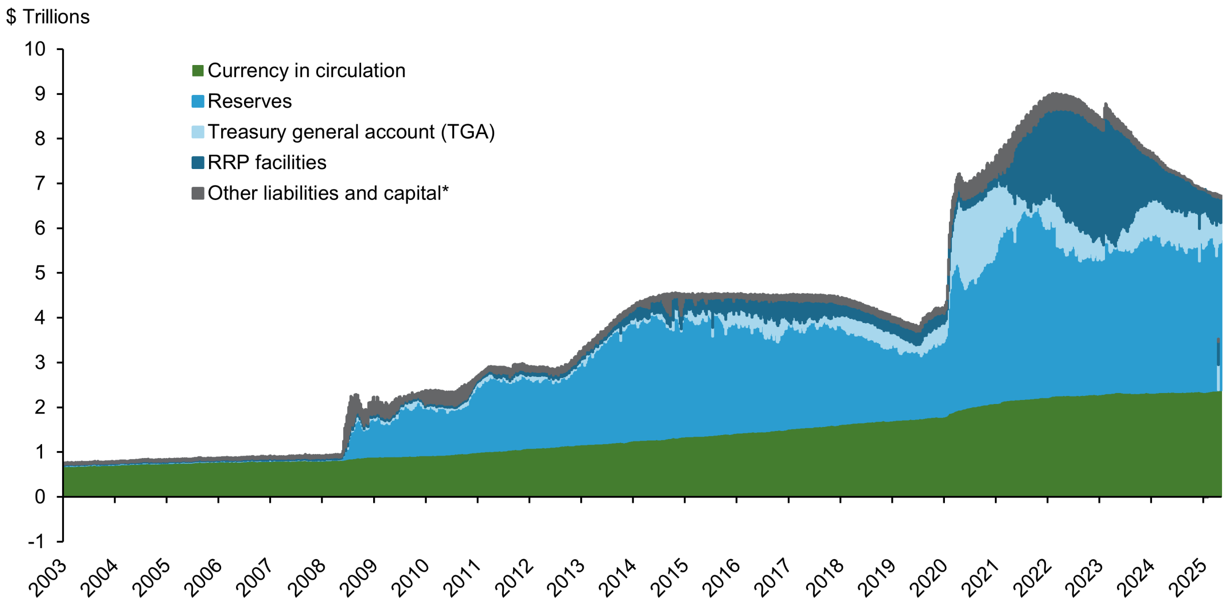

*Includes the liability for earnings remittances due to the U.S. Treasury.
NOTE: RRP facilities include the Foreign Repo Pool and Overnight Reverse Repo Factility (ON RRP).
SOURCES: Board of Governors of the Federal Reserve System, Factors Affecting Reserve Balances (H.4.1).

Federal Reserve Bank of Dallas

3.1 Reserves

Reserves pay an interest rate, the Interest Rate on Reserve Balances (IORB), that the Fed adjusts from time to time as part of the process of implementing the monetary policy decisions of the Federal Open Market Committee (FOMC). Because reserves pay a floating overnight interest rate, reserves have zero interest rate duration risk. (Technically, their economic value does not respond to the interest rate environment because the interest rate they pay, IORB, is also the appropriate discount rate in the Fed's ample-reserves implementation regime.[1])

3.2 Reverse Repo Balances

For our purposes, liabilities that are part of RRP facilities (the Overnight RRP facility for domestic money market participants and the Foreign and International Monetary Authorities facility) are essentially the same as reserves — because they are floating-rate liabilities, they have no interest

[1] This argument abstracts from any effect of the interest rate curve on spreads between IORB and other overnight rates. However, in the ample-reserves regime, such spreads are generally small. Another way to see that reserves have zero interest rate duration risk is to observe that their economic value is par. If reserves were worth more or less than par to private actors, there would be pressure for the supply of reserves to increase or decrease, and money market rates would deviate significantly from IORB. Nearly two decades of experience with the ample reserves regime shows that the Fed has been able to manage reserve supply so as to generally avoid this outcome.

rate duration risk. The small differences in interest rates between RRP liabilities and IORB are not material for our purposes. Therefore, we bucket RRP balances with reserves. We refer to them collectively as floating-rate liabilities.[2]

3.3 Currency in circulation

It may seem strange to consider CIC as contributing to the Fed's interest rate risk. After all, a dollar bill is always worth $1, which suggests that its interest rate sensitivity must be zero. Indeed, in an accounting sense, the present value of CIC is always par and the interest rate risk associated with CIC is zero. However, because the Fed does not pay interest on currency but can use it to fund interest-bearing assets, the economic value of CIC to the Fed rises with short term interest rates, and the associated interest rate risk is not zero. In this sense, there is a close parallel between the interest rate risk associated with a commercial bank's deposit liabilities and the interest rate risk associated with CIC on the Fed's balance sheet.

The duration of CIC depends on how its economic value on the Fed's balance sheet responds to interest rates. In turn, the economic value depends in part on how long currency in circulation can be expected to remain in circulation. Although CIC is a "put-able" liability that can be redeemed at par at any time (for, say, reserves), in practice it has experienced steady growth. Figure 2 shows that, barring a brief period at the turn of the millenium when there was an apparent reversal of pre-Y2K cash hoarding, CIC growth has not dipped below zero for seven decades. Similarly, [15] finds that demand for U.S. banknotes has grown over many years and has continued to grow in the post-COVID period. This experience supports modeling CIC as having a very long life. (Below, we address the implications of factors that could reduce the duration of CIC, including the potential for a shift in demand from currency to remunerated liabilities such as reserves when interest rates increase, or for novel payments technologies to reduce the demand for currency.)

Holding fixed the Fed's assets, each dollar of CIC reduces by one dollar the quantity of interest-bearing liabilities needed to fund those assets. Thus, the economic value of CIC comes from the spread between IORB (the cost of interest-bearing liabilities) and zero (the cost of currency) over the lifetime of the currency. Figure 3 shows this estimated value as a function of the assumed future average of short-term rates, for three representative horizons. The value is non-trivial. If we assume that nominal short-term rates will average 3 percent, which is the median long-run projection for the federal funds rate in the Federal Open Market Committee's June 2025 Summary of Economic Projections (SEP), then each $1 trillion in CIC represents about $610 billion in economic present value over a 30-year horizon, with about half that value coming in the first ten years. For reference purposes, the figure also shows the value of CIC to the Fed if we assume that it is held in perpetuity. This represents a theoretical upper bound on the value of CIC as a function of expected interest rates.

What if rates were to average higher or lower? If nominal short-term rates average 4 percent over the next 30 years, the economic value of currency to the Fed's balance sheet rises to $720 billion per $1 trillion of CIC. Thus, in a sense, each $1 trillion in CIC imparts short-duration risk to the Fed, to the tune of approximately $1.1 billion for each basis point movement in interest rates. This is the equivalent of the duration risk contained in approximately $1.4 trillion of 10-year Treasury notes at current interest rates. Even if we measure economic value of CIC over only a 10-year horizon, the interest rate sensitivity of each $1 trillion of CIC is still about $750 million per basis point (or about $0.9 trillion in 10-year note equivalents). Moreover, as shown in Figure 4, this interest rate risk itself rises as interest rate expectations fall. Using a 30-year horizon for

[2] At times, the Fed has also offered term deposits. These liabilities have been essentially zero since late 2016, with no active usage except for test operations.

Figure 2: Historical growth in Currency in Circulation (CIC)

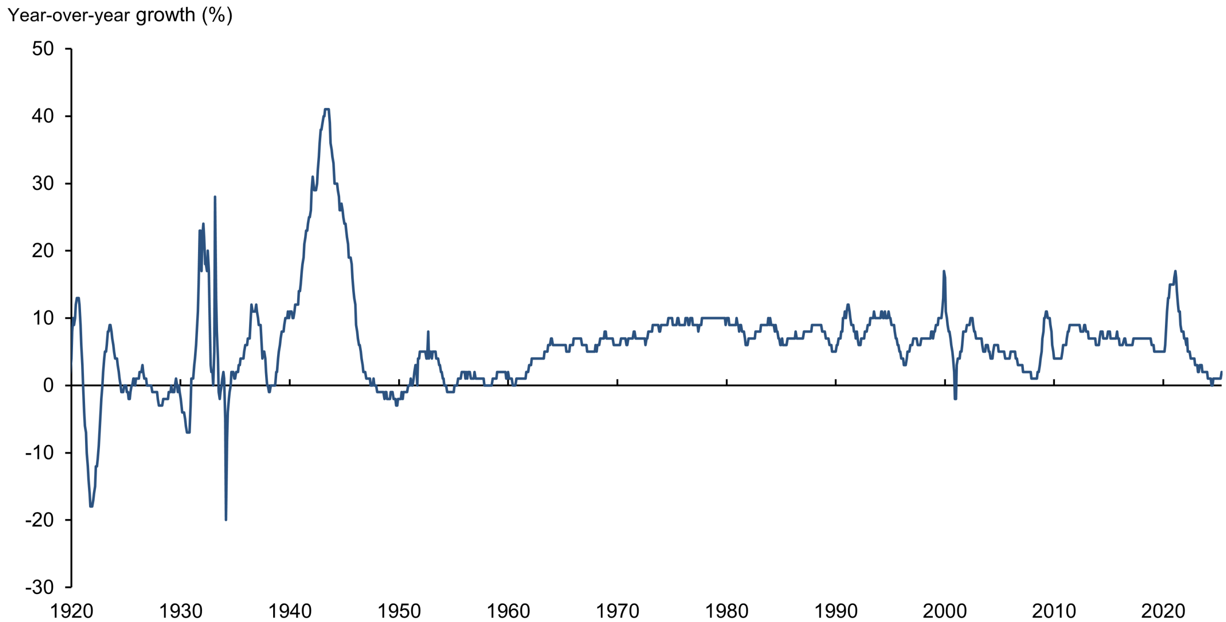

SOURCES: Board of Governors of the Federal Reserve System, Factors Affecting Reserve Balances (H.4.1), FRED. Author's calculations.
Federal Reserve Bank of Dallas

Figure 3: Estimated value to the Fed from $1 trillion in CIC, versus interest rates

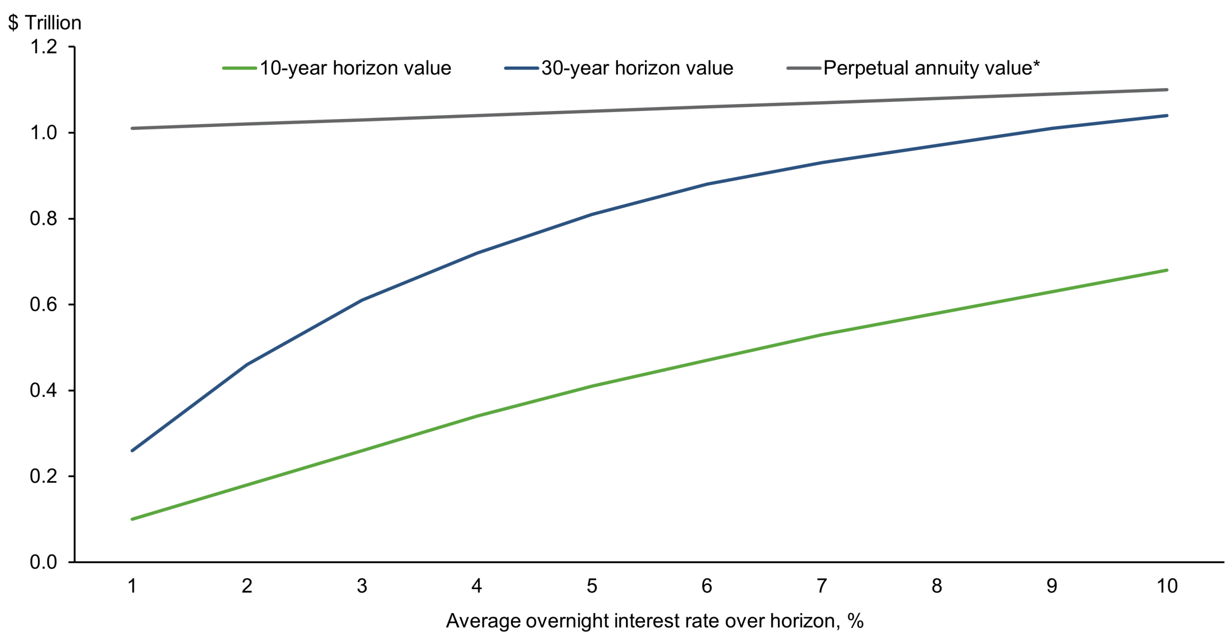

*Perpetual annuity value calculated is calculated under the assumption that CIC is held in perpetuity, and represents a theoretical upper bound on the value of issuing currency to the Fed.
SOURCES: Author's calculations.
Federal Reserve Bank of Dallas

valuation, the duration risk in $1 trillion of currency rises to $2.5 trillion 10-year note equivalents if interest rates average only 1 to 2 percent.

Figure 4: Interest rate risk in $1 trillion of CIC, for various ranges of interest rate expectations

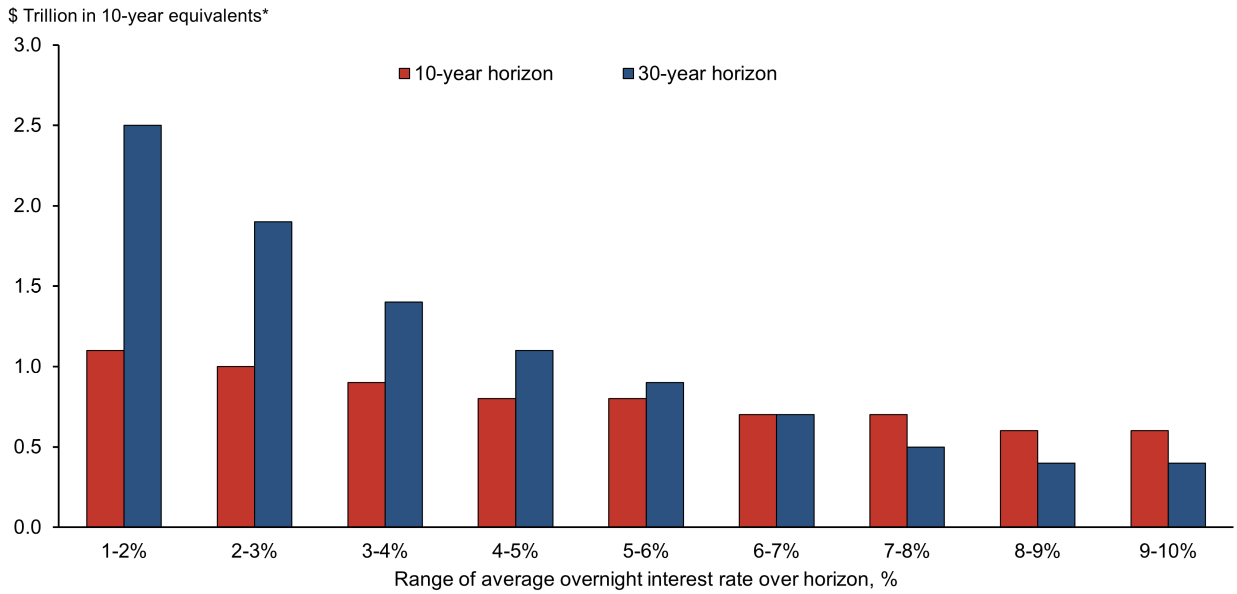

*Average duration risk within each range is calculated by taking the change in the market value of $1 trillion of CIC between the end points of each range, dividing by 100 to get the average present value of a basis point (PVBP), and dividing by $800 million (which is approximately the PVBP of $1 trillion 10-year notes).
SOURCES: Author's calculations.
Federal Reserve Bank of Dallas

Currently, the Fed's liabilities include about $2.4 trillion in currency in circulation. Assuming a horizon of 30 years, this translates into a duration risk that is the equivalent of about $3.9 trillion in 10-year note equivalents if discounted at the long-run 3 percent interest rate in the SEP, or $3.4 trillion 10-year note equivalents if discounted at the roughly 4 percent interest rates priced in the forward swap curve.

CIC could have less duration risk if the private sector's relative demand for currency versus reserves decreases as interest rates rise. Conceptually, such an influence is possible because the opportunity cost of holding CIC to a private actor is something close to IORB (if we assume that holders have access to interest-bearing bank deposits or money market funds that pay a relatively high fraction of IORB).

Our goal here is not to make definitive statements regarding models of currency demand. Rather, we seek to adapt models that have been studied in the literature to assess the approximate degree to which CIC substitutes for Reserves as a function of IORB. To this end, we estimate a regression model (motivated by the models explored in [16]) that fits year-over-year CIC growth as a function of its own lags, lagged year-over-year nominal GDP growth and lagged levels of IORB. The inclusion of these variables reflects the idea that the amount of currency in use likely increases with the nominal size of the economy, may respond to interest rate incentives, and can be persistent over time. We estimate the model on quarterly data and include two lags of each variable. This is similar to the approach taken in [16]. We also considered variations such as using quarter-on-quarter changes, which produced a poorer fit.

Results are shown in Table 1. The sum of the coefficients on the two lags of IORB measures the net response of CIC growth to IORB. Directionally, the sign of response is as expected: higher levels of IORB translate into lower CIC growth. But the magnitude of this response is relatively small. Each 100-basis-point move in IORB is associated with a 0.3% decrease in CIC growth. Even accounting for the confidence interval surrounding the coefficient estimate, it appears reasonable to conclude that the current stock of CIC is very sticky.

Table 1: A model for CIC growth

	Lag	
	t-1	***t-2***
CIC growth	1.22	-0.48
	(0.11)	(0.12)
NGDP growth	-0.04	0.06
	(0.06)	(0.06)
IOR level	-0.95	0.64
	(0.39)	(0.4)
R^2	0.89	
Observations	64	
Start	6/30/2009	
End	3/31/2025	

Memo: Standard errors are presented below each coefficient in parentheses.

Our estimation approach does not distinguish between domestic and foreign demand for U.S. currency. However, studies that have separately assessed domestic and foreign demand and allowed for separate drivers of each, such as [15], have also found that demand for U.S. currency is highly persistent.

3.4 Treasury General Account (TGA)

Conceptually, the TGA can be viewed as a liability with either short or long duration. Over short horizons, the TGA balance can be highly volatile as a result of seasonal factors such as tax payments as well as debt-ceiling management. These fluctuations in the TGA result in equal and opposite fluctuations in the Fed's floating-rate liabilities, all else equal, which could motivate considering the TGA as having similar duration to the floating-rate liabilities. Over the longer term, however, the size of the TGA is a policy choice of the Treasury Department and therefore potentially sticky on average. Additionally, like currency, the TGA does not pay interest.

4 Implications for asset-liability matching on the Federal Reserve Balance Sheet

The analysis above shows that the Fed has both floating-rate liabilities with zero duration and long-lived, non-interest-bearing liabilities. Conceptually, each type of liability has different implications for the Fed's income and balance sheet over time.

From a net income perspective, floating-rate liabilities are a source of variability while long-lived, non-interest-bearing liabilities are not:

- Floating-rate liabilities will cause the Fed's interest expense to vary over time. If the variation in interest expense is not offset by variation in the interest income on the assets backing these liabilities, the Fed's net income — which it remits to the Treasury Department — will also vary over time as a result.

- Long-lived, non-interest-bearing liabilities do not create any variability in interest expense. The Fed can minimize the volatility of the net income associated with these liabilities by minimizing the volatility of the income on the assets backing them.

From a balance sheet perspective, the implications are the opposite. The Fed's accounting rules do not call for marking assets or liabilities to market at current prices. Rather, liabilities are shown on the Fed's financial statements at par and assets at amortized cost, similar to the accounting framework used by commercial banks. Neither floating-rate nor long-duration liabilities affect the accounting value of the balance sheet under this convention. Nevertheless, the mark-to-market value of assets can easily be computed and is shown in the notes to the Fed's statements, and the mark-to-market value of liabilities can similarly be computed. When marked to market, floating-rate liabilities do not generate volatility, while long-lived liabilities do:

- Floating-rate liabilities' economic value is par regardless of interest rates. Backing these liabilities with assets whose mark-to-market value also does not fluctuate will minimize the volatility of the net mark-to-market value of this part of the Fed's balance sheet.

- Long-lived liabilities' economic value varies with interest rates, as discussed above. Unless these liabilities are backed with assets that offset this exposure, the net mark-to-market value of the balance sheet will vary similarly.

As a result, the Fed can minimize both the volatility of its net income and the volatility of the net mark-to-market value of its balance sheet by backing floating-rate liabilities with floating rate assets and long-duration liabilities with long-duration, fixed rate assets.

While long-duration securities help to offset the economic interest rate risk intrinsic to long-duration liabilities such as CIC, such duration matching still has meaningful optical manifestations. A long duration asset portfolio shows unrealized mark-to-market losses when rates rise. While there is also an offsetting increase in the economic value of CIC in a higher rate regime (see Figure 3), CIC is always marked at par in the Fed's quarterly financial statements and such fluctuations in the economic value of CIC are therefore not part of the reporting. As a result, with long-duration liabilities, there is a trade-off between income volatility and mark-to-market volatility on the asset side of the balance sheet. Neutralizing interest rate risk from an economic perspective can result in reported mark-to-market (unrealized) losses on the asset side in a higher-rate environment, as has been the case recently. Instead, the benefits of defeasing the interest rate risk in CIC materialize over time, in the form of reduced volatility of net income.

In addition, the Fed historically has considered many factors other than income and mark-to-market volatility in designing its asset holdings. These considerations can include, though may not be limited to, (i) minimizing the impact of the Fed's holdings on valuations along the curve, (ii) not materially altering the maturity structure of Treasury debt available to the public relative to what Treasury issues, (iii) avoiding cliff-effects in maturities that can lead to lumpy reinvestments, and (iv) not owning too large a share of any one outstanding security (see, for example, [18, 17]). Such considerations could motivate an "across-the-curve" investment portfolio, rather than purely holding very long and very short duration assets. These considerations motivate our analysis of maturity ladders as the elemental building blocks of realistic investment strategies.

4.1 Simulating the performance of maturity ladders

To illustrate the trade-off between income volatility and mark-to-market volatility, we examine the simulated performance of five different "maturity ladder" strategies. A ladder of a given maturity involves distributing the amount to be invested equally among all maturities up to that point.

A 5-year ladder, for instance, involves buying (at time $t = 0$) one-fifth of the total amount each in Treasury securities with one, two, three, four and five years to maturity. In each subsequent time step (we use one year increments), one-fifth of the portfolio would mature, while all the other positions would have aged down by a year. The maturing amount is assumed to be invested in a new 5-year Treasury note, at the par yield prevailing at that time in that path.

Such a ladder has desirable properties:

- It avoids large and concentrated maturations, which in turn would necessitate large and concentrated reinvestments.
- Its interest rate duration risk evolves smoothly. The 5-year ladder, for instance, would only lose about a year of duration before reinvestments of maturing amounts take the portfolio's duration back up.

For these reasons, a ladder strategy is somewhat more realistic in comparison to active duration matching strategies that produce choppy reinvestment needs. But it does have an important downside: a 30-year ladder (the longest available choice in the Treasury market) would have a weighted average maturity of about 15 years, and an effective duration of only a little over 10 years. This means it is possible to duration-match the CIC liability only if that liability is assumed to have a duration on the order of 10 years. Even with that assumption, duration-matching CIC (currently at around $2.4 trillion) would require a similar amount invested fully in 30-year maturity ladders, which could be difficult to conduct at scale since issuance of Treasuries skews towards shorter maturities. In practice, this could make a pure barbell strategy (i.e., backing the Fed's floating rate liabilities with short duration investments and its fixed rate liabilities with a duration-matched strategy) less feasible on the scale needed. Nevertheless, ladders provide an illustrative, simple strategy to study the trade-off between mark-to-market volatility and net income volatility.

We study the behavior of these ladder strategies using a Monte Carlo simulation, generating 1,000 paths spanning the next 30 years using 1-year time steps. At each time step, we simulate interest rates at different maturity points. Appendix A presents details of the simulation. We then evaluate the average performance of these ladder strategies over time across these paths, focusing on three metrics: (i) the volatility of net income, (ii) the volatility of the mark-to-market value of assets and (iii) the volatility of the net mark-to-market (or economic) value of the balance sheet (assets as well as liabilities). As noted earlier, the first two measures are visible in the Fed's financial statements. The third illustrates the potential drawbacks of focusing on the mark-to-market volatility of assets alone.

Given significant differences between CIC (and the TGA, to the extent that one assumes a long duration for it) and floating rate liabilities such as reserves, it is instructive to study these two extremes first, before considering a balance sheet that includes both types of liabilities.

4.2 Portfolio behavior when funded with currency

We begin by considering a balance sheet where the only liability is currency (or, equivalently, a sub-part of the Fed's balance sheet consisting of currency and the assets that back currency).

Let I denote interest income on the Fed's assets, and let E denote interest expense on the Fed's liabilities, so net income is $I - E$. For a fully currency-funded portfolio, $E = 0$ and net income is the same as gross income. The assets that minimize net income volatility for this portfolio are those that have the least volatile gross income — i.e., those with the longest duration, which make fixed coupon payments for a longer period of time. However, there is a trade-off. Such assets have greater volatility in their mark-to-market values.

Figure 5 shows the trade-off between income volatility and the mark-to-market volatility in various maturity ladders that are funded by CIC. A short-duration strategy such as a one-year ladder produces low asset mark-to-market volatility, but it does so at the expense of much higher year-over-year income volatility. Similarly, the 30-year ladder strategy significantly reduces income volatility, but at the expense of higher asset mark-to-market volatility.

Figure 5: Income volatility versus mark-to-market volatility for a currency-funded portfolio

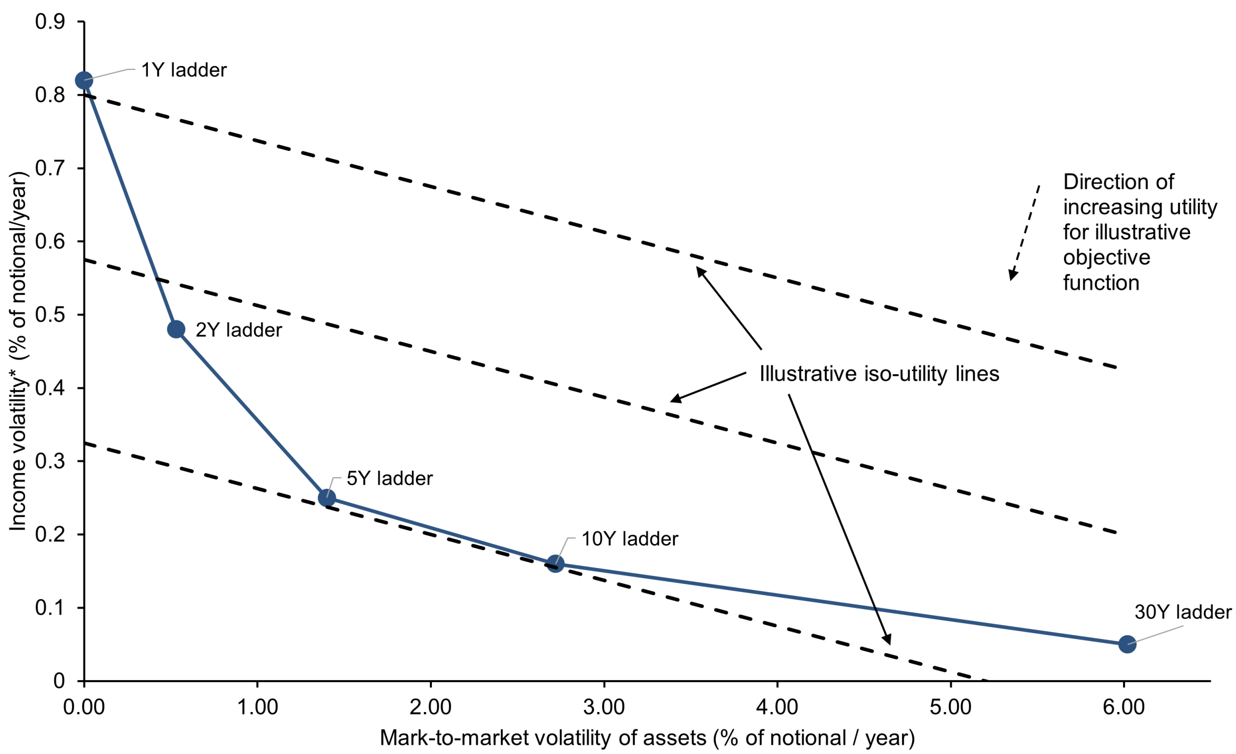

*Income volatility is the standard devation of changes in coupon income.
NOTE: Simulated results are based on Monte Carlo simulations of 1,000 interest rate paths over a 30-year horizon.
SOURCES: Author's calculations.
Federal Reserve Bank of Dallas

Given CIC's long duration, a 30-year ladder strategy goes the farthest towards offsetting duration risk exposure and minimizing income volatility. In contrast, owning a one-year ladder against CIC results in considerably higher income volatility. Moreover, if one were to take a purely economic view of the Fed's balance sheet, considering the mark-to-market volatility of liabilities as well as assets, then longer-duration ladders fare better on both measures of volatility (Figure 6).

That said, long-duration ladders are not an unambiguous choice for backing CIC. If, for example, the Fed employed an objective function that penalized both income volatility and the mark-to-market volatility of assets, then the optimal choice would depend on the relative weight placed on these two forms of volatility. Figure 5 illustrates this by including iso-utility lines for an illustrative linear utility function. For this particular (and purely illustrative) choice of weights, the five-year and 10-year ladders are both about equally preferred and significantly preferred to the 30-year ladder, even though the duration of the five-year and 10-year ladders is much shorter than that of CIC. In addition, the duration of CIC depends on the assumption that it is long-lived. Should policymakers wish to hedge the risk that CIC's duration may be shorter (whether because of a more sensitive response to interest rates or because of novel payment technologies), that could motivate the consideration of shorter-duration portfolios.

Figure 6: Income volatility versus mark-to-market volatility of the balance sheet for a currency-funded portfolio

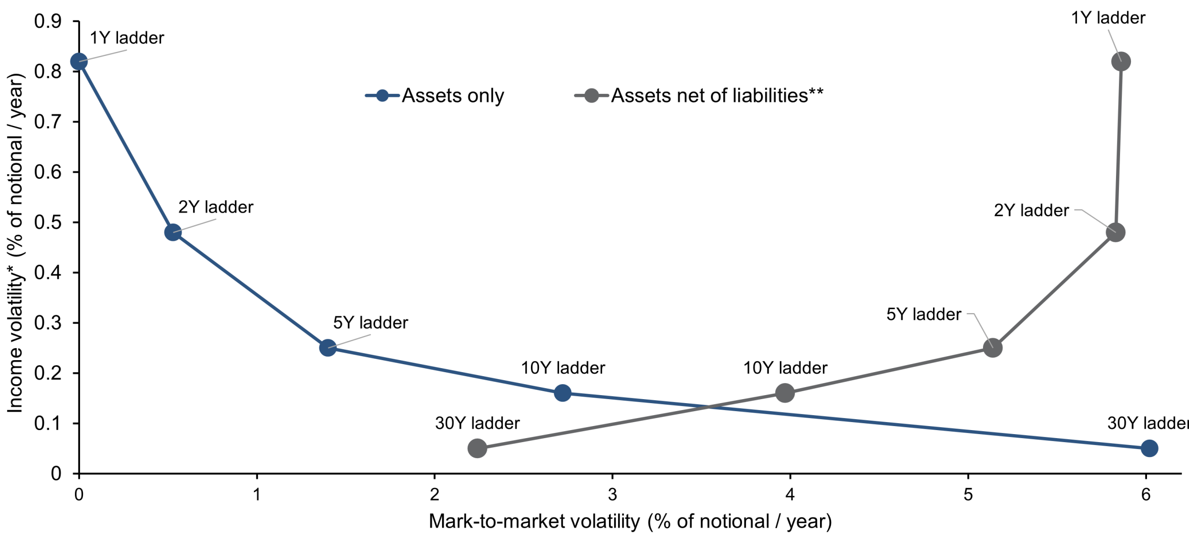

*Income volatility is the standard devation of changes in coupon income.
**Calculated as the standard devation of the changes in the mark-to-market value of both assets and liabilities.
NOTE: Simulated results are based on Monte Carlo simulations of 1,000 interest rate paths over a 30-year horizon.
SOURCES: Author's calculations.

Federal Reserve Bank of Dallas

Additionally, factors other than duration matching could come into play. For example, one consideration could be the seigniorage revenue the Fed generates by providing non-interest-bearing currency backed by interest-bearing assets. Long-term interest rates are typically higher on average than short-term interest rates, so backing CIC with longer-duration assets would tend to increase seigniorage revenue. And, as we noted above, a range of policy considerations beyond the effects on the level and volatility of income could also arise.

4.3 Portfolio behavior when funded with floating-rate liabilities

In a fully reserve-funded portfolio, $E = IORB$, so the variance of net income is $var(I - E) = var(I - IORB)$. An asset that pays IORB will set this variance to zero, which is the minimum. Rolling overnight assets that pay coupons that are highly correlated to IORB can achieve this lower bound, and short-maturity ladders will likely come close. Moreover, short-duration assets will also have less volatile market values.

Our simulation illustrates that backing floating-rate liabilities with short-duration assets will indeed minimize the volatility of both income and market value. As shown in Figure 7, for a portfolio funded with reserves (or other floating-rate liabilities), the 1-year ladder produces the lowest volatility on both dimensions. As a result, regardless of the relative weights on income volatility and mark-to-market volatility, an objective function that considered only these two forms of volatility would call for backing reserves with short-duration assets.

Figure 7: Income volatility versus mark-to-market volatility for a portfolio funded with reserves

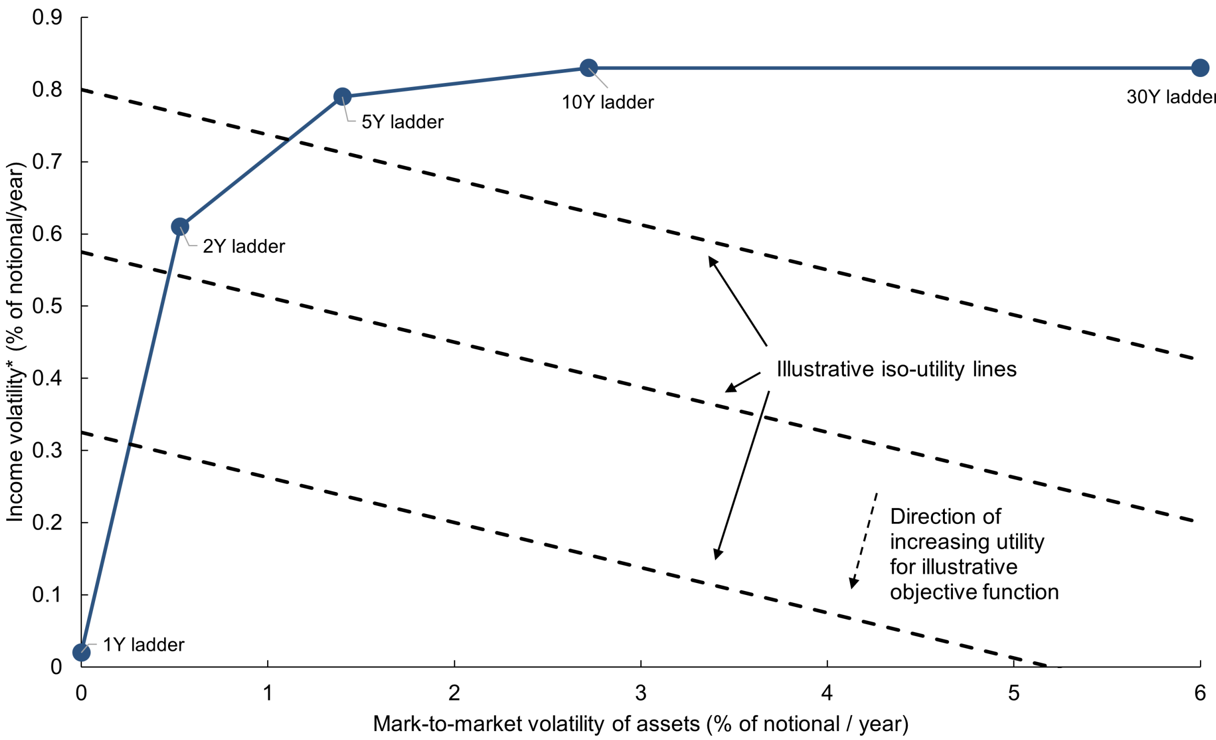

*Income volatility is the standard devation of changes in coupon income minus funding cost (i.e. interest paid on reserves/IORB).
NOTE: Simulated results are based on Monte Carlo simulations of 1,000 interest rate paths over a 30-year horizon.
SOURCES: Author's calculations.
Federal Reserve Bank of Dallas

As with currency, considerations other than duration matching could also be relevant here. Historically, the Fed and other central banks have also taken maturity liquidity to be an important consideration in portfolio design [13]. Maturity liquidity is the percentage of a portfolio that will mature over a given time horizon, allowing the central bank to shrink its balance sheet over that time without actively selling assets, which could have transactions or other costs. Assets with shorter maturity dates have more maturity liquidity; however, maturity liquidity is distinct from duration because a floating-rate asset with a maturity date in the far future has low duration but also low maturity liquidity. Maturity liquidity was considered particularly important in the scarce reserves monetary policy implementation regime that the Fed used before the Global Financial Crisis (GFC), because in some circumstances it could be necessary to rapidly drain reserves to maintain interest rate control. In the post-GFC ample reserves regime, maturity liquidity has received less focus because the primary tools of rate control are administered interest rates.

4.4 Portfolio behavior when funded with a blend of CIC and Reserves

As of mid-June 2025, reserves and RRP balances make up about 60 percent of the Fed's liabilities, while CIC and the TGA make up about 40 percent. The FOMC is continuing to normalize its balance sheet by allowing securities holdings to mature, which reduces the sum of reserves and RRP balances. The FOMC has indicated it will continue this process until reserves are somewhat above the level consistent with ample reserves and then allow reserves to continue to decline to the ample level as a result of growth in other liabilities (see [10]). This plan implies reserves and RRP balances will ultimately be a smaller share of liabilities than they are at present.

The ultimate level of reserves remains uncertain, but as a rough benchmark, we examine the behavior of a balance sheet on which 50 percent of liabilities are floating rate while 50 percent are long lived and not interest bearing. If the TGA is treated as a floating-rate or long-duration liability, this benchmark is similar to the projection in the New York Fed's most recent annual report on open market operations that reserves would reach a minimum in 2026 and equal 44 percent of the Fed's securities holdings at that time (see [18], charts 23 and 25).

Figure 8 shows the income and mark-to-market volatility of various asset ladders in the case where they are funded using an equal-weighted blend of reserves (or other floating rate liabilities) and CIC (or other long-duration liabilities). In addition, we show the performance of two other strategies: a "barbell" (i.e., a 50:50 weighted mix) of the 1-year and 30-year ladders, and an "across-the-curve" strategy that involves combining ladders with weights that would result in the resulting Treasury purchases being proportional to the current stock of Treasuries.

Figure 8: Income volatility versus mark-to-market volatility for a portfolio funded with a 50/50 blend of CIC and Reserves

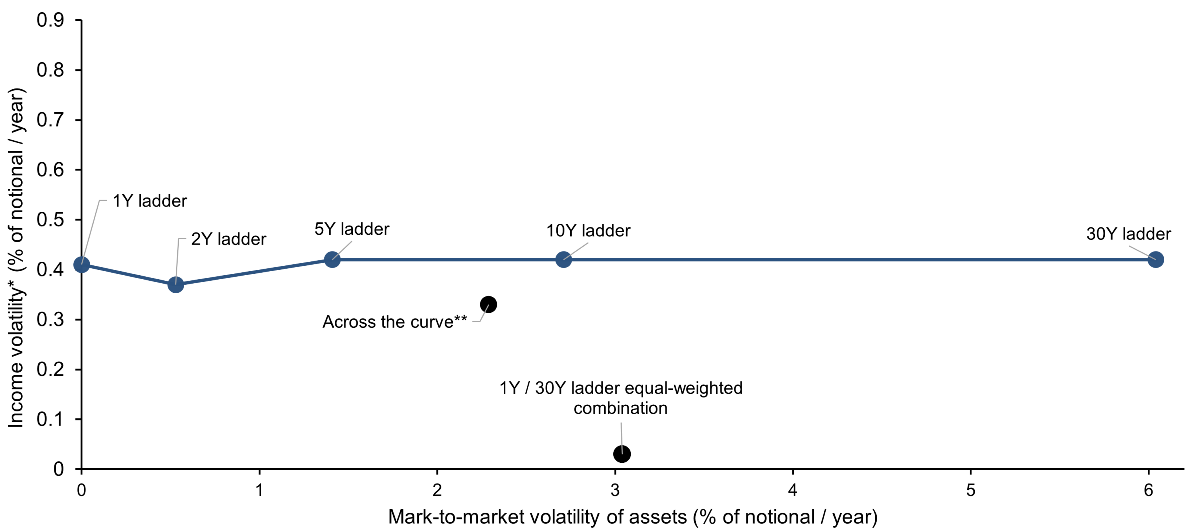

*Income volatility is the standard devation of changes in coupon income minus funding cost (IORB on floating-rate liabilities).
**Across the curve strategy combines individual ladders with weights chosen so that the resulting maturity profile of assets matches the maturity profile of the current stock of Treasuries.The resulting ladders weights are 0%, 37%, 30%, 3%, 30% on 1-year, 2-year, 5-year, 10-year and 30-year ladders.
NOTE: Simulated results are based on Monte Carlo simulations of 1,000 interest rate paths over a 30-year horizon.
SOURCES: Author's calculations.

Federal Reserve Bank of Dallas

We draw two insights from this exhibit. First, income volatility is relatively constant across the different ladders. While the nearly-flat profile is somewhat of a coincidence, it should be expected that it should lie somewhat between the profiles shown in Figures 5 and 7. Indeed, when funded with a mix of currency and floating-rate liabilities, the 1-year ladder has a net income volatility of around 0.4 percent of notional per year, which lies between its currency-funded level (around 0.8 percent of notional per year) and reserve-funded level (zero). This is also true in the 30-year ladder, where the income volatility when funded with a mix of liabilities is around 0.4 percent of notional per year, which lies between the corresponding extremes of (approximately) 0.05 and 0.8

percent of notional per year, respectively.

Second, not all asset strategies with the same duration are created equal. Using a 50:50 mix of the 1-year and 30-year ladder, for instance, results in a duration that is similar to that of a 10-year ladder, but with significantly lower income volatility. To understand the intuition behind this, we go back to our stylized equation - income variance is simply $var(I - E)$. If liabilities are a 50:50 mix of non-interest-bearing currency and floating rate liabilities, this is just $var(I - 0.5 \times IORB)$. Now, for any portfolio, we can always write: $I(t) = 0.5 \times IORB(t) + C(t)$ for some choice of $C(t)$. But the closer $C(t)$ is to a constant, the lower the income volatility of the portfolio. Now the intuition becomes clear: a portfolio that is comprised of 50 percent floating rate assets (which generates interest income that is approximately one half IORB) and 50 percent long duration assets (whose coupon is fixed for a considerable time, causing $C(t)$ to have low volatility) would do the trick. Therefore, among the many ways in which one can construct a portfolio of a given duration, barbells of short and long duration sub-portfolios are likely to produce the lowest income volatility when suitably weighted.

As we noted above, a pure barbell portfolio could be challenging to implement at scale. It turns out, though, that an easily implementable across-the-curve portfolio can still produce lower income volatility than a similar duration ladder. For example, Figure 8 also shows the income and mark-to-market volatility of an illustrative portfolio invested 37 percent in the 2-year ladder, 30 percent in the 5-year ladder, 3 percent in the 10-year ladder, and 30 percent in the 30-year ladder. Those weights on the individual ladders results in Treasury purchases across all different maturities that are in proportion to the maturity profile of outstanding Treasuries and makes it more practical to implement. As can be seen, such a strategy also produces somewhat lower income volatility than individual ladder strategies, but not as low as what is achievable with barbells.

5 Conclusion: large-scale asset purchase programs and broader implications

Our discussion so far has considered a steady-state Fed balance sheet, in which the size and mix of both assets and liabilities are constant. However, over the past two decades, the Fed has repeatedly engaged in temporary large-scale purchases of long-duration assets. The purposes of large-scale asset purchase programs (LSAPs) have included providing economic stimulus by lowering term premiums when the overnight rate was at its effective lower bound and supporting the smooth functioning of Treasury markets and the resulting flow of credit to households, businesses and governments (see [4, 5, 3, 6] for LSAP announcements in the aftermath of the Global Financial Crisis, and see [8, 9] for LSAP announcements at the onset of the Covid-19 pandemic).

LSAPs also increase the stock of reserves (or, potentially, RRP balances). Thus, LSAPs result in a temporary expansion of the Fed's balance sheet consisting of floating-rate liabilities backed by long-duration assets. In other words, an LSAP entails a temporary addition to the balance sheet that, to achieve its policy goals, must not be duration-matched. Moreover, once the asset purchases cease, this addition remains on the balance sheet until either the assets mature, the central bank sells them, or steady-state demand for central bank liabilities increases. In this section, we examine ways to think about duration matching over time when the Fed may occasionally engage in LSAPs, as well as some broader policy considerations.

One possibility is to focus on duration matching only for the steady-state balance sheet. Under this approach, the asset portfolio could be designed to match the steady-state mix of liabilities, and the duration (and other) effects of the assets and liabilities associated with an LSAP could be considered as unrelated. Such a separation would be reminiscent of the Bank of England's decision, when it intervened to support the smooth function of gilt securities markets in late September 2022, to hold the resulting assets in a distinct portfolio (see [1]). LSAPs tend to occur in environments of relatively low long-term rates, heightening the probability that rates will be higher after the programs end. This implies that the Fed is likely to experience mark-to-market losses on assets purchased in an LSAP (though these losses are never crystalized if the Fed holds assets to maturity) as well as relatively low net income during the post-LSAP period.

Alternatively, one could seek to match the average duration of liabilities over time. The duration of assets could be reduced below what would match liabilities during normal times to offset the increase in duration during an LSAP. Reducing duration during normal times would prevent using an exactly matched barbell portfolio and so could increase the Fed's net income volatility during normal times, though as shown above, the effect would depend on the exact construction of the asset portfolio and on whether a barbell portfolio is considered feasible. This approach would also have other benefits and costs.

During an LSAP, for any given increase in the Fed's duration holdings, the level of the Fed's duration exposure would not rise as high because the starting point would be lower. Thus, the Fed would be less exposed to mark-to-market losses from LSAP purchases. However, as we discuss above, reducing duration during normal times would tend to reduce seigniorage revenue during normal times. Moreover, to effectively match the average duration of liabilities over time, one would need to model the probability and size of LSAPs. Precise modeling would be challenging because LSAPs have historically been rare events driven by the materialization of tail risks to the economy and financial system.

Our analysis has focused on the income implications of holding different assets to back the Fed's liabilities. As we discussed at the outset, income implications are at minimum an incomplete description of the relevant policy considerations, given that the Fed does not seek to maximize profits. Other policy considerations that policymakers and analysts have suggested are relevant for the Fed's asset composition include the effect on the Treasury's ability to control the maturity composition of debt held by the private sector [17], the effect on the risks on the consolidated government balance sheet [19], and effects on credit availability and the economy [2].

Finally, we also have not investigated how the steady-state composition of the Fed's asset holdings might affect the shape of the yield curve and other asset prices. Such effects depend importantly on how the Treasury Department and other issuers adjust their debt issuance, if at all, in response to any changes in the Fed's desired holdings (see page 8, [7]). These relationships are an important subject for further research.

A Appendix - Details of the simulation approach

We simulate the behavior of asset ladders by first simulating a distribution of potential paths for short-term and long-term interest rates and calculating how these rate paths would affect the income and mark-to-market value through time of each asset ladder. Our simulation is not a "risk-neutral" simulation that is calibrated to recover the prices of instruments along the current yield curve, such as the kind one might use in the pricing of path-dependent derivatives under current market conditions. Rather, it is designed to capture a wide range of plausible paths that might be encountered in the future. Therefore, the equations guiding our generation of rate paths are grounded in assumptions that are reasonably reflective of historical experience.

First, we assume that there are two stochastic factors, given that a significant portion of the variation in yields is typically explained by two principal components. This means that once we simulate shocks to two different maturity points on the curve, we may use linear combinations of those two to calculate changes to yields at other maturity points.

We treat the short rate (IORB) as the first independent point. We simulate this using the equation:

$$y_{t+1}^{IORB} = \max(y_t^{IORB} + \epsilon_{t+1}^{IORB}, 0.10\%) \tag{1}$$

where ϵ_{t+1}^{IORB} is a Gaussian shock with a standard deviation of σ^{IORB}, and y_t^{IORB} denotes the IORB rate at time t. We floor simulated short rates at 0.10%, corresponding to the lowest setting of interest on reserves in 2020, the most recent episode when the FOMC moved its policy target to the effective lower bound.

Next, we simulate changes in the 10-year rate. We allow this rate to respond both to the short rate and to a second (independent) random shock. Specifically, we use the equation:

$$y_{t+1}^{10Y} = y_t^{10Y} + \beta_{IORB}^{10Y} \times (y_{t+1}^{IORB} - y_t^{IORB}) + \epsilon_{t+1}^{10Y} \tag{2}$$

where ϵ_{t+1}^{10Y} is an independent Gaussian shock with a standard deviation of σ^{10Y}, y_t^{10Y} denotes the 10-year rate at time t, and β_{IORB}^{10Y} denotes the linear combination weight on the change in IORB.

We then calculate changes to 1-year, 2-year, 5-year, and 30-year yields as linear combinations of the changes in IORB and 10-year rates. The general form of this equation is:

$$y_{t+1}^{M} = y_t^{M} + \beta_{IORB}^{M}(y_{t+1}^{IORB} - y_t^{IORB}) + \beta_{10Y}^{M}(y_{t+1}^{10Y} - y_t^{10Y}) \quad \forall \quad M \in \{1Y, 2Y, 5Y, 30Y\} \tag{3}$$

The chosen values for all parameters reflect the authors' judgment, but are based on observed volatilities and relationships between yields at different maturity points over the past fifteen years. Table 2 shows the assumed values. We initialize the simulation using recently observed levels: at $t = 0$, we set IORB at 4.3 percent, the 1-year yield at 4.0 percent, the 2-year at 4.0 percent, the 5-year at 4.0 percent, the 10-year at 4.4 percent and the 30-year yield at 4.9 percent.

Finally, for "off-benchmark" maturity points that have not been specifically described above, we read off simulated yields from an interpolated yield curve constructed from the previously described benchmark maturity points.

Table 2: Simulation parameter values

		Linear Comb. Weights		
Simulated Rate	**Sigma (Indep. Shock)**	**Change IORB**	**Change 10Y rate**	**Comment**
IORB	0.875			Random shock
1Y		1.0		Chg. same as IORB
2Y		0.7	0.3	Blend of chg. in IORB and 10Y
5Y		0.25	0.75	Blend of chg. in IORB and 10Y
10Y	0.6	0.4		40% of chg. in IORB + indep. shock
30Y		-0.05	0.95	Blend of chg. in IORB and 10Y

References

[1] Bank of England. Press release on 28th Sep, 2022. URL: https://www.bankofengland.co.uk/news/2022/september/bank-of-england-announces-gilt-market-operation.

[2] Ben S. Bernanke. The crisis and the policy response. Speech at the Stamp Lecture, London School of Economics, London, England, 2009. URL: https://www.federalreserve.gov/newsevents/speech/bernanke20090113a.htm.

[3] Federal Reserve Bank of New York. Statement Regarding Purchases of Treasury Securities, on 3rd Nov, 2010. URL: https://www.newyorkfed.org/markets/opolicy/operating_policy_101103.html.

[4] Federal Reserve Board. Press Release on 25th Nov, 2008. URL: https://www.federalreserve.gov/newsevents/pressreleases/monetary20081125b.htm.

[5] Federal Reserve Board. FOMC Statement on 18th Mar, 2009. URL: https://www.federalreserve.gov/newsevents/pressreleases/monetary20090318a.htm.

[6] Federal Reserve Board. FOMC Statement on 13th Sep, 2012. URL: https://www.federalreserve.gov/newsevents/pressreleases/monetary20120913a.htm.

[7] Federal Reserve Board. Meeting of the Federal Open Market Committee. Transcript, 2019. URL: https://www.federalreserve.gov/monetarypolicy/files/FOMC20190501meeting.pdf.

[8] Federal Reserve Board. FOMC Statement on 15th Mar, 2020. URL: https://www.federalreserve.gov/newsevents/pressreleases/monetary20200315a.htm.

[9] Federal Reserve Board. FOMC Statement on 23rd Mar, 2020. URL: https://www.federalreserve.gov/newsevents/pressreleases/monetary20200323a.htm.

[10] Federal Reserve Board. Plans for reducing the size of the federal reserve's balance sheet. Press Release on 4th May, 2022. URL: https://www.federalreserve.gov/newsevents/pressreleases/monetary20220504b.htm.

[11] Federal Reserve Board. Principles for reducing the size of the federal reserve's balance sheet. Press Release on 26th Jan, 2022. URL: https://www.federalreserve.gov/newsevents/pressreleases/monetary20220126c.htm.

[12] Emily Greenwald, Sam Schulhofer-Wohl, and Joshua Younger. Deposit Convexity, Monetary Policy and Financial Stability. *Working Paper No. 2315, Research Department, Federal Reserve Bank of Dallas*, 2023. doi:10.24149/wp2315.

[13] Jeffrey Huther, Jane Ihrig, and Elizabeth Klee. The Federal Reserve's Portfolio and its Effect on Interest Rates. *Finance and Economics Discussion Series, Board of Governors of the Federal Reserve System*, (2017-075), 2017. doi:10.17016/feds.2017.075.

[14] Jane Ihrig, Zeynep Senyuz, and Gretchen C. Weinbach. The Fed's "Ample-Reserves" Approach to Implementing Monetary Policy. *Working Paper 2020-022, Finance and Economics Discussion Series, Federal Reserve Board*, 2020. doi:10.17016/FEDS.2020.022.

[15] R. Judson. Demand for U.S. Banknotes at Home and Abroad: A Post-Covid Update. *International Finance Discussion Paper No. 1387, Federal Reserve Board*, 2024. `doi:10.17016/ifdp.2024.1387`.

[16] Mariam El Hamiani Khatat. Monetary Policy and Models of Currency Demand. *IMF Working Paper WP/18/28*, 2018. `doi:10.5089/9781484341926.001`.

[17] Lorie K. Logan. Efficient and Effective Central Bank Balance Sheets. *Remarks delivered at the Bank of England Agenda for Research Conference*, 2025. URL: `https://www.dallasfed.org/news/speeches/logan/2025/lkl250225`.

[18] Markets Group of the Federal Reserve Bank of New York. Open market operations during 2024. Report prepared for the FOMC, 2024. URL: `https://www.newyorkfed.org/medialibrary/media/markets/omo/omo2024-pdf.pdf`.

[19] Sam Schulhofer-Wohl. Monetary policy implementation and the consolidated government balance sheet. Dallas Fed Economics Blog, 2025. URL: `https://www.dallasfed.org/research/economics/2025/0415`.