Less Duration, More Durable: The Case for Using Floating-Rate Debt to Reduce Treasury Borrowing Costs*

Anh Le† Haoxiang Zhu‡

This Draft: November 11, 2026
First Draft: October 2026

Abstract

The U.S. government pays over $1 trillion in interest expense in 2026 on a national debt of over $40 trillion. Managing its borrowing costs is critical for the sustainability of U.S. fiscal conditions and the long-term health of the U.S. Treasury market. Building on a dynamic model of interest rates, we show that greater issuance of floating-rate notes (FRNs) during periods of elevated interest rate risk premium can meaningfully reduce the U.S. government's interest expense. Such saving outweighs the illiquidity discount in the market price of existing Treasury FRNs. The saving is also greater if the FRNs are linked to shorter-maturity T-bills.

*For helpful discussions and comments, we thank Mark Cabana, Darrell Duffie, Gil Holmes, Thomas Pluta, Jeremy Stein, and seminar participants at Pennsylvania State University and Carnegie Mellon University.

†Smeal College of Business, Pennsylvania State University. Email: anh.le@psu.edu.

‡MIT Sloan School of Management and NBER. Email: zhuh@mit.edu.

Less Duration, More Durable: The Case for Using Floating-Rate Debt to Reduce Treasury Borrowing Costs

Abstract

The U.S. government pays over $1 trillion in interest expense in 2026 on a national debt of over $40 trillion. Managing its borrowing costs is critical for the sustainability of U.S. fiscal conditions and the long-term health of the U.S. Treasury market. Building on a dynamic model of interest rates, we show that greater issuance of floating-rate notes (FRNs) during periods of elevated interest rate risk premium can meaningfully reduce the U.S. government's interest expense. Such saving outweighs the illiquidity discount in the market price of existing Treasury FRNs. The saving is also greater if the FRNs are linked to shorter-maturity T-bills.

1 Introduction

In late 2026, the U.S. government carries more than $40 trillion of debt,[1] with annual interest expense exceeding $1 trillion.[2] As debt issued during the low-interest rate era matures, the U.S. Treasury must refinance it at prevailing market rates, which are much higher. Long-term bond yields have risen to their highest levels in two decades,[3] and they reflect not only expected short-term interest rates but also a risk premium, namely the extra compensation investors demand for bearing interest rate risk.

This paper argues that broader use of Treasury floating-rate notes (FRNs), particularly during periods of high interest rate risk premium, can meaningfully reduce the government's interest burden, while maintaining a stable maturity structure. The core mechanism is simple: FRNs carry minimal interest rate risk and hence minimal interest rate risk premium. Shifting a marginal dollar of borrowing from fixed-coupon securities to FRNs of the same maturity should therefore yield an expected annual saving approximately equal to the interest rate risk premium, while maintaining the same weighted-average maturity of debt.

Implementing this simple idea, however, is considerably more involved. Three questions arise. First, if FRNs are illiquid, does the Treasury pay a premium to issue them? Second, because the interest rate risk premium varies over time, when should the Treasury issue FRNs? Third, which short-term interest rate should FRNs be indexed to? We address these three questions using a dynamic model of interest rates fitted to market data.

First, is there an illiquidity discount for FRNs, and if so, how large is it? We answer this question by directly examining FRN prices. The U.S. Treasury has issued two-year FRNs indexed to the 13-week Treasury bill since January 2014. New two-year FRNs are auctioned quarterly and reopened in each of the following two months. The key pricing metric is the spread over the 13-week T-bill rate, determined at auction, that investors demand. A spread of 10 bps means that the FRN pays the 13-week T-bill rate plus 10 bps, annualized. Interest accrues daily and is paid quarterly. To establish a benchmark for FRN prices, we begin with a dynamic model of interest rates with two yield curves: one for off-the-run Treasury notes and bonds, and one for T-bills, whose yields include a liquidity premium. Applying this two-curve model to FRNs, we find that both auction prices and secondary-market prices of FRNs lie below the benchmark price. Correspondingly, 45 of the 48 new FRNs auctioned

[1] https://fiscaldata.treasury.gov/datasets/debt-to-the-penny/

[2] https://fiscaldata.treasury.gov/interest-expense-avg-interest-rates/

[3] See https://www.wsj.com/finance/investing/surging-yields-bring-the-bond-market-back-to-the-turn-of-the-century-2b74773f and https://www.reuters.com/business/us-30-year-bond-yield-rises-highest-since-2004-selloff-deepens-2026-09-24/

between January 2014 and December 2025 cleared at a positive spread, and the average is 11.1 bps. The spread at which the benchmark prices a new two-year FRN at par is negative in most weeks, averaging about −3 bps. We interpret this pricing gap as evidence of an illiquidity discount on FRNs.

To account for this illiquidity discount, we extend the benchmark to a three-curve model, adding a third yield curve that discounts FRN cash flows and is fitted to FRN prices. Using this three-curve model, we define the illiquidity discount as the excess of FRN yields over off-the-run yields. Between 2014 and 2025, this discount raised the spread that a new two-year FRN had to pay by about 14 bps per year on average. Figure 1 illustrates the spot interest rates implied by the three curves: generally, the T-bill curve is the lowest, the FRN curve is the highest, and the off-the-run (non-special) Treasury curve is in the middle. The T-bill liquidity premium declines with maturity and vanishes at the one-year maturity, whereas the FRN illiquidity discount increases with maturity.

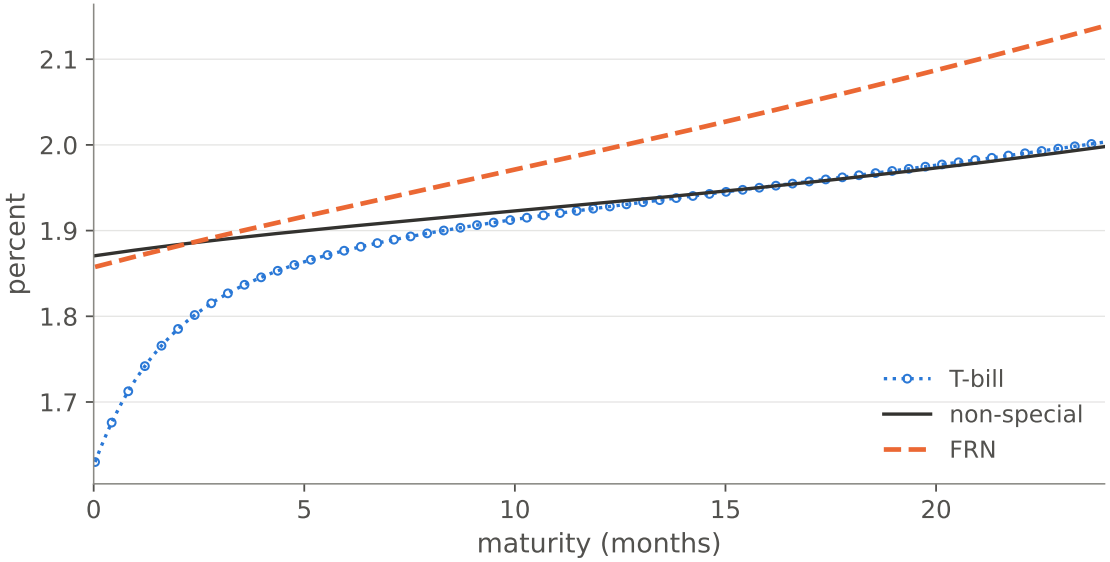


Figure 1: Average spot interest rates for T-bills, off-the-run coupon securities (non-special), and FRNs, in percent, continuously compounded, implied by the three-curve model. Each curve is evaluated at the mean of the filtered states over 2014–2025, the period in which FRNs traded.

The second question concerns timing: when is the best time to issue FRNs to avoid paying the interest rate risk premium? The data suggest that the answer lies in the Federal Reserve interest rate cycle. The gray bars in Figure 2 show the realized annual interest saving of each FRN issued from January 2014 to December 2023, computed as the par yield of two-

year off-the-run coupon securities immediately before the FRN auction, minus the average realized 13-week T-bill rate over the following two years, minus the auction-determined spread. Positive realized savings concentrate at the tail end of the Federal Reserve's hiking cycle from December 2015 to December 2018, where they are around 100 bps per year for FRNs auctioned from mid-2018 to early 2020. This pattern has an intuitive explanation: after a series of rate hikes, bond investors have absorbed losses that reduce their risk-bearing capacity, so they may demand a higher interest rate risk premium. This explanation is consistent with evidence that increases in short-term rates raise term premia (Hanson and Stein, 2015; Hanson et al., 2021; Kekre et al., forthcoming). It is in these periods that FRNs deliver the largest savings.

Without foresight, how can the Treasury identify periods of elevated interest rate risk premium in real time? Our model is particularly useful for this purpose. At each date, we estimate the model using only information available up to that date and then compute the model-implied forecast of the 13-week T-bill rate over the next two years. The model-predicted expected saving is then the two-year par yield from the non-special curve, minus the model-predicted average 13-week T-bill rate, minus the FRN spread in the auction. The blue triangles in Figure 2 show the expected savings implied by the real-time model. The model correctly predicts the positive savings from early 2018 to early 2020, although it mistakenly predicts positive savings in 2023, when interest rates continued to rise. Consider a simple rule: at each quarterly FRN auction between January 2016 and December 2023, the Treasury issues a two-year FRN if the real-time model predicts a positive saving and a standard two-year fixed-coupon note otherwise. Under this rule, the Treasury would have issued FRNs at 19 of the 32 auctions and saved 14.8 bps of interest expense per year on those FRNs, relative to issuing only fixed-coupon notes.

The third question concerns security design. Current FRNs have a two-year maturity and are indexed to the 13-week T-bill; would an alternative design deliver larger savings? Our analysis suggests the answer is yes. For notes with relatively short maturities, such as two years, the embedded interest rate risk premium is approximately the risk-neutral expectation of the average index over the life of the note minus its physical expectation, $E^Q(\overline{r_\tau}) - E^P(\overline{r_\tau})$, plus a convexity adjustment, where r_τ is the index of tenor τ underlying the FRN and $\overline{r_\tau}$ denotes its average. We find that this interest rate risk premium decreases with the tenor of the index, so the 13-week T-bill embeds a smaller interest rate risk premium than the 4-week T-bill does. Figure 1 illustrates this fact: the average T-bill interest rate increases in maturity, i.e., $E^P(\overline{r_{13w}}) > E^P(\overline{r_{4w}})$, but rolling over 4-week or 13-week T-bills for a year or

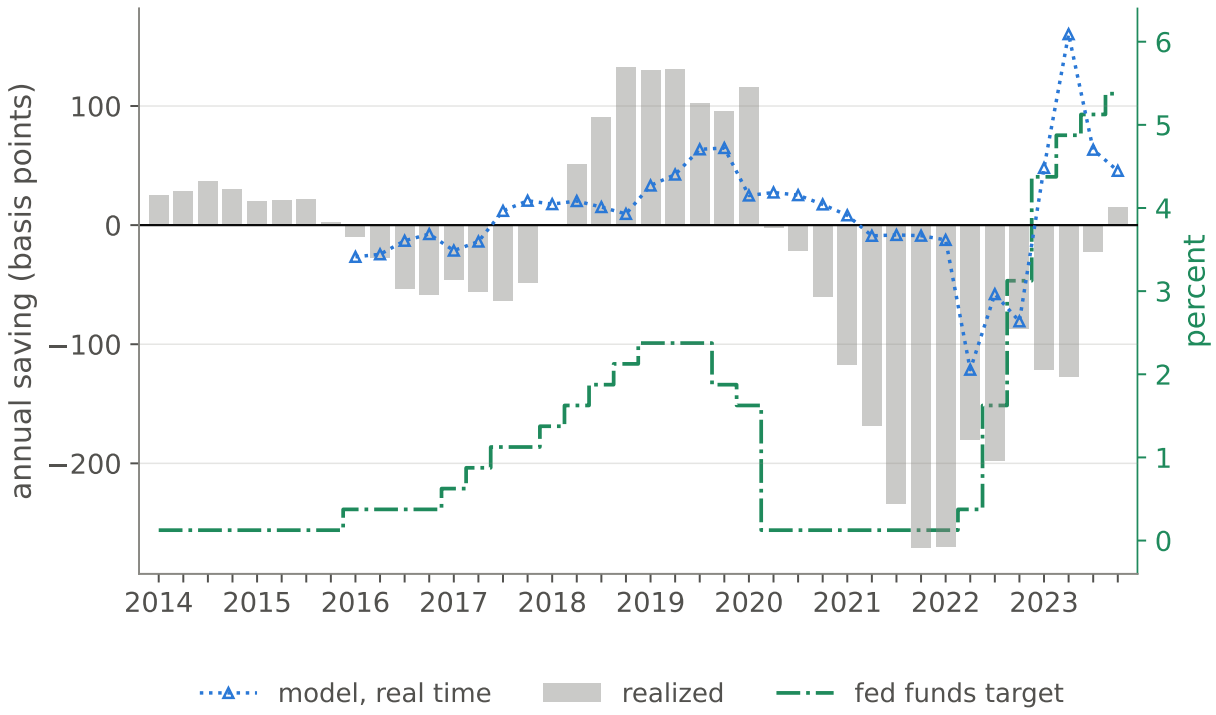


Figure 2: Realized and anticipated annual saving for each two-year FRN auctioned and matured between 2014 and 2025. The realized saving uses the average of the index rate published by the Treasury over the life of each FRN. The anticipated saving replaces that average with the forecast from a real-time model, estimated on data up to the last sampled date before each auction. Real-time forecasts begin with the January 2016 auction, when enough FRN-implied yield data are available. The dash-dotted step line (right axis) is the midpoint of the FOMC's target range for the federal funds rate on each auction date.

longer produces the same spot interest rate under the T-bill curve, i.e., $E^Q(\overline{r_{13w}}) \approx E^Q(\overline{r_{4w}})$. Hence $E^Q(\overline{r_{4w}}) - E^P(\overline{r_{4w}}) > E^Q(\overline{r_{13w}}) - E^P(\overline{r_{13w}})$, and the difference is approximately 7 bps in our sample. Equivalently, indexing FRNs to the 4-week T-bill lets the Treasury capture the larger liquidity premium apparently embedded in shorter-maturity T-bills. Importantly, this gain is not offset by a larger illiquidity discount. In our estimated model, the liquidity component of a two-year FRN differs across index tenors by less than 0.1 bp, so the larger risk premium carries through fully to the net saving.

Along the maturity dimension, the model predicts that the interest rate risk premium increases with the maturity of the security, but so does the illiquidity discount. The expected net saving, defined as the interest rate risk premium minus the illiquidity discount, also

increases with maturity. Because no FRNs have been issued with maturities beyond two years, this result should be interpreted with caution.

What does all of this mean for the U.S. Department of the Treasury? Treasury describes FRN issuance as consistent with its mission "to borrow at the lowest cost over time," alongside managing the maturity profile of its debt, expanding the investor base, and adding flexibility for debt managers (U.S. Department of the Treasury, 2013). With the objective of cost saving in mind, our analysis points to three routes, which are not mutually exclusive.

The first route, and the least risky, is to index FRNs to the 4-week T-bill. Our model predicts that this simple change would reduce interest expense by about 7 bps per year, based on data from the last decade, compared with the current design of indexing FRNs to 13-week T-bills.

The second route is to improve the liquidity of the FRN market and thereby reduce the illiquidity discount on FRNs. In this regard, money market funds are central to this market: they "often represent the single largest class of investors in FRNs and frequently hold more than 50% of FRN issues" (Fleckenstein and Longstaff, 2020). Their holdings, however, are limited by SEC Rule 2a-7, which caps a fund's weighted average life at 120 days and charges an FRN at its final maturity rather than its next reset date (U.S. Securities and Exchange Commission, 2014). Regardless of whether the SEC revises this rule or its interpretation, one-year FRNs seem to be low-hanging fruit, as our analysis suggests that they would carry a smaller illiquidity discount. Once liquidity builds up in the one- and two-year segments, as evidenced by a stable and well-behaved spread, the Treasury can incrementally extend the FRN program to the three- and five-year segments, where the interest rate risk premium avoided is larger.

The third route, which is more challenging, is to actively adjust the mix of fixed-coupon and FRN issuance so that less duration risk is released to the market when the interest rate risk premium is high. This strategy is, admittedly, tantamount to taking an active view on the interest rate risk premium. The premium is estimated rather than observed, and such estimates can be noisy. Moreover, adjustments in the issuance mix need to be consistent with the "regular and predictable schedule" on which Treasury relies to reduce market uncertainty (U.S. Government Accountability Office, 2014). Both risks would be mitigated, however, if the composition adjustments were bounded and well communicated to the market. In addition, if the Treasury auctioned N-year FRNs and N-year fixed-coupon securities simultaneously, it could compare bids across the two securities and choose the composition based on its internal assessment of relative value, including the risk premium.

In this way, higher interest rate risk premia would automatically tilt issuance toward FRNs.

A broader use of FRNs is consistent with the steady shift of Treasury debt toward T-bills in recent years. As of October 2026, about 23.0% of marketable Treasury securities are T-bills, compared to 16.7% in December 2021 and 13.1% in December 2016.[4] By significantly increasing the fraction of debt financed by T-bills rather than fixed coupon securities, the Treasury shortens both the duration and the maturity of debt. Compared to the strategy of rolling over T-bills, FRNs further reduce the rollover risk and the costs of running and participating in Treasury auctions.

Treasury debt management is a vast topic, and we choose to focus on only one aspect of it: for a given maturity, does the Treasury save borrowing costs by moving a marginal dollar of debt from fixed coupon securities to FRNs? How does the saving, if any, vary over time? What design features of FRNs make them less costly to issue? While we have made progress on these questions, our analysis is one piece of a giant puzzle. Complementary questions that are beyond the scope of this paper include the optimal maturity structure of government debt, the modeling of fiscal deficit of the Federal government, the general equilibrium change in risk premium in response to a significant reduction in aggregate duration, and the modeling of demand from various types of Treasury securities investors.

Related Literature

Our paper contributes to four strands of literature: the pricing of Treasury floating rate notes (FRNs), the optimal maturity structure of government debt, the liquidity premium of Treasury securities, and term structure modeling.

Pricing of Treasury FRNs. Fleckenstein and Longstaff (2020) and Hartley and Jermann (2024) reach opposite conclusions about whether FRNs are cheap or expensive for the Treasury. Fleckenstein and Longstaff (2020) compare FRN prices with replicating portfolios built from Treasury bills, notes, and basis swaps, and find that FRNs trade at a persistent premium of roughly 6 to 10 basis points. They attribute this premium to mark-to-market stability, which makes FRNs especially attractive to money market funds, and support this interpretation with fund flows around the 2014 money market fund reform. Hartley and Jermann (2024) instead price the FRN's coupon stream off the zero-coupon Treasury curve and find that the spreads the Treasury pays at auction exceed the frictionless par spread by 14 basis points on average. They interpret this wedge as compensation for FRNs being less

[4]See https://fiscaldata.treasury.gov/datasets/monthly-statement-public-debt/summary-of-treasury-securities-outstanding.

money-like than the bills their coupons reference. Rather than replicating the cash flows of FRNs based on other securities and derivatives, we price FRNs based on a no-arbitrage term structure model. Applied to data from 2014 to 2025, our model reveals that two-year FRNs are issued with an illiquidity discount of about 14 bps per year. This discount, however, is outweighed by the interest rate risk premium in certain stages of the Federal Reserve's interest rate cycle; thus, issuing FRNs at the right time leads to positive expected savings for the Treasury.

Optimal maturity of government debt. Our paper is complementary to the literature on the maturity structure of government debt. In preferred-habitat theories (Vayanos and Vila, 2021), investor clienteles demand specific maturities and risk-averse arbitrageurs absorb the remaining duration, so the supply of duration affects the term premium. Greenwood and Vayanos (2014), Hamilton and Wu (2012), and Li and Wei (2013) provide empirical support for such supply effects. Belton, Dawsey, Greenlaw, Li, Ramaswamy, and Sack (2018) develop a simulation framework in which the Treasury adjusts maturity in response to economic conditions, trading off the lower expected cost of short-term financing against the volatility of interest expense. Our focus on interest rate risk premium is a natural continuation of this line of inquiry. Using a no-arbitrage model of interest rates, we contribute to this strand of literature by estimating expected savings over different time periods and for FRNs of different maturities.

Georgievska, Pegoraro, and Saunders (2026) show that the U.S. Treasury does not adjust its duration supply in response to swap spreads or term premia, and they argue that more active duration management, including a larger FRN program, could reduce borrowing costs. They quantify the potential savings using auction-level demand elasticities. Our approach, based on a no-arbitrage term structure model of interest rates, is complementary: we estimate the time-varying interest rate risk premium that the Treasury could avoid, as well as the illiquidity discount that FRNs carry, and hence the net savings from FRN issuance. We also investigate how savings vary across reference rates and FRN maturities.

Liquidity premium of Treasury securities. The literature finds meaningful safety and liquidity premia for Treasury securities (Krishnamurthy and Vissing-Jorgensen, 2012). These premia are particularly large for Treasury bills (Greenwood, Hanson, and Stein, 2015; Nagel, 2016), but Fleckenstein and Longstaff (2024) find that long-term Treasury yields have been above "fair value" since 2015. Greenwood, Hanson, and Stein (2015) show that the convenience premium on bills is a reason for the government to issue short-term debt, to be balanced against rollover risk. Consistent with this literature, we find that the liquidity pre-

mium of T-bills is larger at shorter maturities, and FRNs based on shorter-maturity T-bills would produce larger savings.

Term structure modeling. Methodologically, we build on the affine Gaussian term structure model of Joslin, Singleton, and Zhu (2011). Our contribution is to extend this framework to three yield curves (off-the-run, T-bill, and FRN). We show how the FRN zero yield curve can be constructed by "bootstrapping" FRNs of staggered remaining maturities. The bootstrapping procedure is similar in spirit to that of Fama and Bliss (1987), but requires the use of a model to value the floating coupons, which, unlike the fixed coupons of notes and bonds, are not known in advance. The reduced-form nature of our model lets us study the data while imposing minimal structure on the economy. Re-estimating the model each week on data available at that date also allows its parameters, including the long-run mean of the short rate, to potentially change over time, a feature that Bauer and Rudebusch (2020) show matters for estimates of term premia.

2 Institutional Background of Treasury FRNs

2.1 Procedural history and design considerations of Treasury FRNs

Treasury floating rate notes (FRNs), first auctioned in January 2014, were the first new type of Treasury security since the launch of Treasury Inflation-Protected Securities (TIPS) in 1997. The idea was raised briefly at the August 2011 meeting of the Treasury Borrowing Advisory Committee (TBAC) and discussed in detail at the August 2012 TBAC meeting (Treasury Borrowing Advisory Committee, 2011, 2012). The U.S. Treasury issued requests for comments in March and December 2012 on its proposed FRN program, which was adopted in its July 2013 final rule. In the final rule, the Treasury explained: "Treasury views issuance of floating rate notes as consistent with its mission to borrow at the lowest cost over time, manage the maturity profile of our marketable debt outstanding, expand the Treasury investor base, and provide a financing tool that gives debt managers additional flexibility" (U.S. Department of the Treasury, 2013).

On maturity management, a two-year FRN locks in two years of funding while paying a floating coupon, so it carries short-rate exposure similar to a rolling 13-week bill program but requires one refinancing instead of eight. In TBAC's words, floating rate notes "could reduce rollover risk by decoupling the maturity and interest rate risk decisions" (Treasury Borrowing Advisory Committee, 2011). On the investor base, commenters believed that "if appropriately structured, a Treasury floating rate note would be an attractive investment

for a broad base of institutional investors including money market funds, securities lenders, corporations, and foreign central banks" (U.S. Department of the Treasury, 2013).

Central to the analysis of this paper is the first objective, lower cost over time. The mechanism we focus on is term-premium avoidance. By substituting short-rate exposure for duration exposure, the Treasury lowers expected financing cost by the term premium it would otherwise pay. Market participants discussed this channel explicitly. One primary dealer's comment letter in 2012 observed that "forward rates have systematically over-predicted the future evolution of short term rates, primarily because the curve has had positive term premium," so that indexing to a short rate lets the Treasury "reduce its interest expense by an amount equal to the embedded term premium" (Misra, 2012).

Market participants also thoroughly discussed design choices, the most contentious of which was the index. The Treasury considered two candidates: the high rate from the 13-week bill auction, converted to a simple-interest money-market yield on an actual/360 basis, and the General Collateral Finance (GCF) Repo Index published daily by the Depository Trust & Clearing Corporation. Most public commenters favored the GCF Repo Index for its "daily rate changes, the greater nominal value traded in GC agreements, and the potential to attract new investors," while others preferred the 13-week Treasury bill rate, citing disadvantages of the GC rates, "including higher daily rates and volatility than bill markets, the relative newness of the GCF indexes, and the possibility of exposing Treasury's financing costs to private-sector rate shocks" (Copic et al., 2014). In the end, the Treasury adopted the 13-week Treasury bill as the index (U.S. Department of the Treasury, 2013).

Maturity is another design choice. In its December 2012 request for comments, the Treasury stated: "We intend to start our floating rate securities program with a two-year maturity. We anticipate strong demand from money market investors with weighted average portfolio constraints. A two-year maturity might also offer an appealing investment alternative for cash portfolios. We anticipate eventually issuing longer maturity securities and seek comment on the most appropriate maturity for both the initial and future phases of the program" (U.S. Department of the Treasury, 2012). The money market fund constraint is worth emphasizing. Under SEC's Rule 2a-7 of the 1940 Investment Company Act, a government floating rate security is deemed to mature at its next reset for the 397-day and weighted-average-maturity tests, but its final maturity counts in full against the 120-day weighted-average-life limit, so a longer tenor reduces how much money funds can hold (U.S. Securities and Exchange Commission, 2014). Fidelity, a major money-fund manager, preferred "a final maximum maturity of one year (or as close to one year as possible)" (Fi-

delity Investments (FMR Co.), 2013). Maturity extension, however, pulled the other way: in TBAC's assessment, "a 2yr final maturity FRN would slow the rise in outstanding Treasury WAM, a 10yr final would accelerate the rise in WAM and a 5yr final would split the difference" (Treasury Borrowing Advisory Committee, 2012). The Treasury's final rule authorizes issuance of FRNs with maturities of one to ten years, and notes that "[m]ost commenters agreed that this [two years] was a good maturity to start with, and suggested eventual expansion to longer maturities of up to 10 years" (U.S. Department of the Treasury, 2013).

Two developments, shortly after the first Treasury FRN issuance in January 2014, added momentum to the program. The SEC's July 2014 amendments to Rule 2a-7 required institutional prime money market funds to float their net asset values and subjected prime and tax-exempt funds to redemption gates and liquidity fees, while government funds kept a stable NAV (U.S. Securities and Exchange Commission, 2014). In the run-up to the October 2016 compliance date, more than one trillion dollars moved from prime to government funds (Cipriani and La Spada, 2021). As government securities, FRNs meet the need of government money market funds (subject to the cap in weighted average maturity and life of the securities). And with duration of about one week (in terms of interest rate risk), FRNs also suit the price stability goal of prime funds. Separately, the Government Accountability Office concluded in June 2014 that FRNs could help Treasury meet its borrowing objectives but recommended additional tools to monitor exposure to short-rate shocks (U.S. Government Accountability Office, 2014). The terms of the FRN have not changed substantively since 2014.

2.2 Two-year FRNs from 2014 to 2025

Every Treasury FRN has been issued with a two-year original maturity; the contractual terms below follow the official term sheet (U.S. Department of the Treasury, Bureau of the Fiscal Service, 2013) and are encoded directly in the pricing formulas of Section 4. Interest accrues from a Dated Date on the last calendar day of a month, and the note matures on the last calendar day of the month two years later. Original issues settle on the last calendar day of the month and reopenings on the last Friday, with shifts to the next business day for holidays. On each calendar day s, interest accrues per \$100 notional at

$$\text{daily accrual}(s) \;=\; \max\!\left(\frac{b_{s^\star}+\theta}{360},\, 0\right), \tag{1}$$

where $b_{s^\star}$ is the high rate from the most recent 13-week Treasury bill auction in effect on day s, converted to a money-market rate on an actual/360 basis, and θ is the spread, fixed in tenths of a basis point at the original auction for the life of the security. The floor at zero binds only in the rare case $b_{s^\star} + \theta < 0$. Because bill auctions are held weekly, $b_{s^\star}$ is piecewise constant on blocks of about seven days. Under the lockout rule, a bill auction whose result becomes effective within two business days before an Interest Date or Issue Date is ignored and the preceding rate is used; this rule enters the pricing formula of Section 4. Accrued interest is paid quarterly on the last calendar day of January, April, July, and October, and principal is repaid at par at maturity. FRNs are not STRIPS-eligible.

Each auction is a single-price auction in which competitive bidders submit a "discount margin" over the index, positive, zero, or negative, in tenths of a basis point. Noncompetitive bids of up to \$5 million are filled in full at the clearing margin, and competitive bids are filled in ascending order of the discount margin up to the offering amount. At an original-issue auction the spread θ is set at the highest accepted margin. At reopenings the spread carries over, and the margin from the auction should be viewed as a quotation convention that, once applied to a pre-defined formula, recovers the price of the FRN. The FRNs are auctioned on a monthly cadence: a new CUSIP is issued in January, April, July, and October and reopened in each of the following two months. Auctions are announced on the Thursday before auction week and held late in the month, typically on a Wednesday.

From January 2014 to December 2025, the Treasury held 146 FRN auctions: 144 regular monthly auctions and two small off-cycle reopenings.[5] These comprise 48 original issues and 98 reopenings. They are summarized in Table 1 and plotted in Figure 3. For original issues, offering sizes rose from \$15 billion at the inaugural auction to \$30 billion by late 2025. The offered amount at a typical reopening is about \$2 billion lower than the preceding original auction. The average accepted amount per auction is \$20.5 billion. The average bid-to-cover ratio is 3.28, indicating strong demand. Indirect bidders, i.e., those submitting bids through primary dealers, account for just over 50% of the allocations, whereas primary dealers account for 40%-45%. Annual issuance grew from \$164 billion in 2014 to \$354 billion in 2025, and cumulative issuance through December 2025 reached \$2.99 trillion, of which \$1.09 trillion came from original issues and \$1.90 trillion from reopenings. The spread has ranged from -7.5 basis points (April 2022) to $+30.0$ basis points (October 2019), with an average of 11.1 basis points.

[5] On December 6, 2019 and July 14, 2022, Treasury announced and auctioned \$25 million of an outstanding FRN on the same day for settlement three days later, and primary dealers took the entire amount.

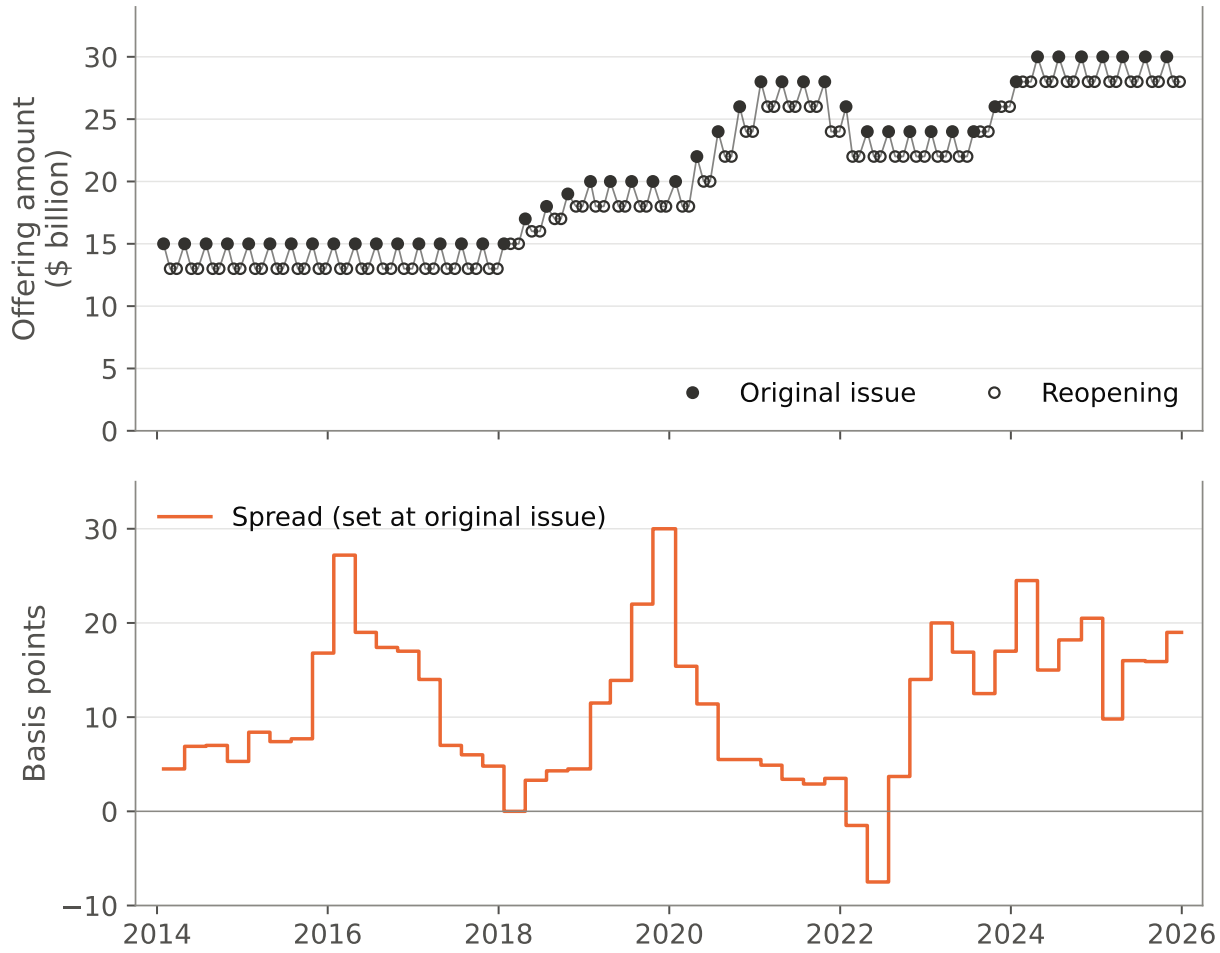


Figure 3: Offering amount (upper panel) and spread (lower panel) of Treasury FRN auctions, from 2014 through 2025, as reported by TreasuryDirect. The two off-cycle $25 million reopenings of December 6, 2019 and July 14, 2022 are omitted.

3 Data

Two datasets underlie the empirical analysis: (i) for FRNs, the initial auction-determined spread and secondary-market prices at daily frequency; (ii) off-the-run Treasury coupon securities, Treasury bills, and general-collateral repo rates, also observed at daily frequency. Section 3.1 describes the first. Section 3.2 describes the second, together with the two zero-coupon yield curves we construct from it, and Section 3.3 states the estimation sample and reports summary statistics.

Table 1: Summary statistics of Treasury FRN auctions as reported by TreasuryDirect, over the 2014–2025 sample, which covers all 146 auctions of two-year FRNs: 48 original issues and 98 reopenings. Each cell reports the mean across auctions with the standard deviation in parentheses. Bidder shares are percentages of total accepted, which also includes noncompetitive and SOMA awards. The high discount margin determines the contractual spread at the original-issue auctions. At reopening auctions, the high discount margin sets the price of the reopened FRN given the contractual spread.

	All (N=146)	Original (N=48)	Reopenings (N=98)
Accepted per auction ($B)	20.5 (6.6)	22.8 (6.4)	19.3 (6.4)
Bid-to-cover ratio	3.28 (0.60)	3.28 (0.56)	3.29 (0.62)
Primary dealer share (%)	43.5 (14.5)	39.6 (11.9)	45.4 (15.2)
Direct bidder share (%)	2.6 (3.8)	2.4 (3.9)	2.6 (3.8)
Indirect bidder share (%)	51.4 (15.4)	52.4 (13.2)	51.0 (16.4)
High discount margin (bp)	11.7 (8.0)	11.1 (7.8)	12.0 (8.1)

3.1 FRN Prices

For every two-year FRN issued between January 2014 and December 2025, we obtain daily bid and ask prices from Bloomberg, organized by maturity cohort. Each cohort is tracked over its full secondary-market life, from the Dated Date to the Maturity Date or, for cohorts still outstanding, to the end of the sample. A cohort issued in January 2014 therefore contributes roughly two years of observations, while one issued in October 2025 is observed only through the end of the sample period. We supplement these quotes with each cohort's fixed spread over the index rate, taken from TreasuryDirect.

Quotes are stated per $100 of face value and exclude accrued interest. For each trading day, we compute the interest accrued since the most recent Interest Date under the security's contractual terms and add it to the bid and ask quotes to obtain dirty prices. The midpoint of these dirty prices serves as the price input to our empirical analysis.

3.2 Off-the-Run and T-Bill Yield Curves

We download the CRSP daily Treasury file from WRDS, which reports bid, ask, and mid quotes together with accrued interest and coupon schedules for every outstanding U.S. Treasury security since June 14, 1961. We retain the three security types that carry no embedded options or other nonlinear features: noncallable bonds (ITYPE=1), noncallable notes (ITYPE=2), and Treasury bills (ITYPE=4). Callable bonds, TIPS, STRIPS, and flower

bonds are excluded. We then drop obviously stale or erroneous quotes and apply the standard Fama-Bliss cross-sectional filters, which impose two criteria. The first removes near-duplicates: when two securities mature within a week of each other, we keep the one with the narrower bid-ask spread. The second removes yield outliers: a security is dropped if its yield lies more than 20 basis points away from the average yield of both its three nearest shorter-maturity neighbors and its three nearest longer-maturity neighbors and does not lie between those two averages, or if it forms part of a run of alternating yield changes of at least that magnitude. A final pass reinstates any dropped security whose yield satisfies the neighbor criterion relative to the securities that survive.

To construct the off-the-run yield curve, we restrict attention on each day to off-the-run securities, excluding all T-bills and the two most recently issued notes or bonds at each standard maturity—the on-the-run and first off-the-run issues. We also drop coupon securities with less than a month to maturity: backing a yield out of a near-maturity price annualizes any pricing error by a factor of roughly $365/\tau$, where τ is the number of days remaining, so that at fifteen days a five-cent error in price—well inside the bid-ask spread of an illiquid near-maturity issue—becomes roughly 120 basis points of yield, against ten basis points at six months. Following Filipović, Pelger, and Ye (2025) (henceforth FPY) and using their publicly available code, we estimate a set of discount factors by kernel ridge (KR) regression, which minimizes overall pricing error subject to a penalty that keeps the fitted discount curve smooth. Specifically, for the cross-section of securities outstanding on day t, with prices P_i and cash flow matrix C (so that P_i equals the sum of C_{ij} times the discount factor at tenor j), the KR discount curve $d^\star(\tau)$, where τ denotes maturity, solves

$$\min_{d \in \mathcal{G}_{\alpha,\delta}} \sum_i \omega_i \big(P_i - C_i \cdot d\big)^2 \;+\; \lambda \, \|d\|^2_{\alpha,\delta},$$

where $\mathcal{G}_{\alpha,\delta}$ is the space of twice weakly differentiable discount curves normalized to $d(0) = 1$, the norm $\|\cdot\|_{\alpha,\delta}$ penalizes the first and second derivatives of d under a maturity weight $e^{\alpha\tau}$, the ω_i are pricing weights, and λ governs the trade-off between fit and smoothness. As shown by FPY, the solution is available in closed form,

$$d^\star(\tau) \;=\; 1 + \sum_j k(\tau, \tau_j)\, \beta_j, \qquad \beta \;=\; C^\top \big(CKC^\top + \Lambda\big)^{-1} \big(P - C\mathbf{1}\big),$$

where k is the reproducing kernel of $\mathcal{G}_{\alpha,\delta}$, given in closed form by FPY, K is the kernel matrix with entries $K_{jl} = k(\tau_j, \tau_l)$, and $\Lambda = \mathrm{diag}(\lambda/\omega_i)$.

We weight each security by the inverse of its squared dollar duration, so that the pricing-error term is approximately a yield-error criterion. Following FPY, we set the maturity weight to $\alpha = 0.05$, the value they interpret as the infinite-maturity yield, and the tension parameter to $\delta = 0$. To choose λ, we run five-fold cross-validation, holding out securities stratified by maturity bucket. This is the design FPY use to select their own hyperparameters.[6] The resulting root-mean-squared pricing error (RMSE) is nearly flat over $\lambda \in [0.5, 2.0]$, so we adopt the FPY default of $\lambda = 1$.

Because the off-the-run cross-section begins only about a month out, we anchor the short end with an overnight general-collateral repo rate, entered as a synthetic one-day zero-coupon bond. From 2005 onward we use the DTCC GCF Treasury repo rate; before then, when that index is unavailable, we use an overnight Treasury general-collateral quote from Bloomberg (ticker `RPGT01D`). Both rates are secured against a general basket of Treasury collateral and are therefore largely insulated from supply and demand for any individual issue. To filter out extreme single-day swings in the overnight market, we apply a weekly rolling median to the series before use.

The T-bill curve is estimated by the same procedure, with T-bill prices as the input. No short-end anchor is required here: bills run down to within days of maturity, so the cross-section constrains the front of the curve directly. At the long end we supplement the bill prices with the on-the-run and first off-the-run two-year Treasury notes. These are the two securities the non-special curve excludes, so the bill curve, as constructed, captures the potential specialness embedded in the short government curve out to the two-year horizon. We again set $\alpha = 0.05$ and, after a cross-validation exercise analogous to that in FPY, $\lambda = 0.1$.

The two curves serve different purposes. Spot rates from the off-the-run curve can be read as default-free rates, free of the convenience components embedded in the prices of on-the-run coupon securities and of T-bills. We therefore refer to it as the *non-special* curve, reserving "off-the-run" for the securities it is fitted to. T-bills, by contrast, carry a money-like liquidity premium (Greenwood, Hanson, and Stein, 2015; Nagel, 2016), so the spot rates they imply typically lie below their non-special counterparts. Because FRN coupons are indexed to the auction rate on newly issued 13-week bills, we use the T-bill curve to compute FRN cash flows.

[6] Our implementation differs from FPY's in granularity rather than design: we use five folds at each month-end over five maturity buckets (0–2y, 2–5y, 5–10y, 10–20y, 20–30y), where FPY use ten folds at each quarter-end over ten buckets.

3.3 Estimation Sample and Summary Statistics

For estimation we use yields sampled weekly from the beginning of 2000 through the end of 2025, a total of 1,350 dates, together with the FRN-implied yields of Section 5.1 on the same weekly grid from 3 August 2015, when they begin. Both curves fitted in Section 3.2 are defined at every maturity, but only a fixed subset of maturities from each is used in estimation: the non-special curve at 3 and 6 months and 1, 2, 3, 4 and 5 years, and the T-bill curve at 30, 60, 91, 180, 270, 365 and 730 days. To the non-special yields we append a seven-day trailing median of the overnight repo series of Section 3.2 as a one-day yield. The repo rate is quoted as simple interest on an actual/360 basis, so we convert each quote q to the continuously compounded one-day yield $365\ln(1+q/360)$, the basis of every other yield in the estimation.

Two of the T-bill maturities, 365 and 730 days, exceed the maturity of any outstanding bill; at those points the fitted curve is primarily identified by the on-the-run and first off-the-run two-year notes. We include them because pricing a two-year FRN requires risk-neutral forecasts of the 13-week bill rate at horizons out to two years, and without an observation at the long end of the segment those forecasts would rest largely on extrapolation from the factor structure.

Table 2 reports summary statistics of the two yield curves. At every maturity where both curves are observed, the average T-bill yield lies below the average non-special yield: by 3.7 bps at 91 days, 2.7 bps at one year and 0.9 bps at two years. At 91 days and one year, this gap is the T-bill convenience premium, visible before any model is imposed and declining with maturity. At two years, where the bill curve is identified by the on-the-run and first off-the-run notes rather than by bills, the 0.9 bps gap is instead the on-the-run premium of those notes, which Section 4.3 measures directly.

4 A Three-Curve Term Structure Model

4.1 Risk-Neutral Dynamics

We adopt the Joslin-Singleton-Zhu affine term structure framework (Joslin, Singleton, and Zhu, 2011). Under the risk-neutral measure Q, the state vector Z_t follows

$$dZ_t = -\kappa\, Z_t\, dt + \Sigma\, dB_t, \tag{2}$$

Table 2: Summary statistics of yields used in estimation, in percent, weekly, 2000–2025. Panel A: T-bill zero yields implied by T-bill prices and the on-the-run and first off-the-run two-year notes. Panel B: non-special zero yields implied by off-the-run coupon securities. Both sets of constant-maturity zero-yields are estimated by kernel ridge regression as described in the text. The overnight rate is a seven-day trailing median of the overnight general-collateral repo series, converted to continuous compounding.

	Mean	**Std. dev.**	**P25**	**P75**
Panel A: T-bill curve				
30-day	1.818	1.921	0.074	3.493
60-day	1.862	1.949	0.086	3.594
91-day	1.900	1.960	0.097	3.721
180-day	1.980	1.947	0.149	3.734
270-day	2.023	1.909	0.198	3.768
365-day	2.078	1.873	0.274	3.796
730-day	2.232	1.750	0.625	3.817
Panel B: non-special curve				
Overnight (repo)	2.024	2.047	0.178	4.018
3-month	1.937	1.943	0.144	3.758
6-month	2.010	1.940	0.181	3.787
1-year	2.105	1.888	0.326	3.853
2-year	2.241	1.746	0.652	3.809
3-year	2.419	1.638	0.974	3.864
4-year	2.608	1.546	1.310	3.928
5-year	2.780	1.471	1.571	3.978

where B_t is a standard Brownian motion and Z_t is a 5×1 vector partitioned as $Z_t = (X_t', S_t, F_t)'$. The three factors X_t drive the non-special Treasury curve;[7] an additional factor S_t captures the T-bill yield curve, while a fifth factor F_t captures the market pricing of Treasury FRNs.

We impose a block-diagonal structure on the mean-reversion matrix,

$$\kappa = \text{blkdiag}\big(\kappa_X, \kappa_S, \kappa_F\big), \tag{3}$$

where κ_X is a 3×3 block in Jordan form with real eigenvalues, and κ_S and κ_F are scalars. Together with the zero risk-neutral long-run mean of Z_t implied by (2) and the unit loadings in (4)–(6) below, these conditions follow the standard normalizations of Joslin, Singleton,

[7] Litterman and Scheinkman (1991) show that three common factors account for nearly all of the variation in Treasury bond yields.

and Zhu (2011); we take Σ to be lower triangular for the same reason, since Σ enters the model only through $\Sigma\Sigma'$, which remains an unrestricted covariance matrix. Because the eigenvalues of κ_X are real and distinct in our estimates, the Jordan block is similar to a diagonal matrix; equation (3) therefore entails five free mean-reversion parameters in total.

We specify three different short rates, all as affine functions of the same state vector:

$$r_t = r_\infty + \mathbf{1}_3' X_t \qquad \text{(non-special)}, \tag{4}$$

$$r_t^b = r_\infty^b + \delta_X^{b\prime} X_t + S_t \qquad \text{(T-bill)}, \tag{5}$$

$$r_t^f = r_\infty^f + \delta_X^{f\prime} X_t + F_t \qquad \text{(FRN)}. \tag{6}$$

The unit loadings on X_t in equation (4) and on S_t and F_t in equations (5) and (6) are normalizations that fix the scale of the latent factors. The zero restrictions are separate, and they give S_t and F_t their interpretation as curve-specific spread factors: S_t does not enter the non-special or FRN short rates, while F_t does not enter the non-special or T-bill short rates.

We use the separate short rate r_t^b to price T-bills, in particular the 13-week bills that determine the cash flows of Treasury FRNs. The spread between the non-special and T-bill short rates is

$$c_t \equiv r_t - r_t^b. \tag{7}$$

We leave c_t unrestricted in sign. Thus, the model can accommodate both a T-bill convenience premium and an inconvenience yield. Modeling r_t and c_t jointly is equivalent to modeling r_t and r_t^b jointly. Importantly, the spread can depend both on the non-special curve factors X_t and on the independent T-bill factor S_t. We therefore do not hard-wire that T-bills always trade at a yield below the non-special Treasury curve. This flexibility allows the model to capture the T-bill convenience premium documented in Greenwood et al. (2015); Nagel (2016), as well as the possibility of an inconvenience yield, which He, Nagel, and Song (2022) document for longer-maturity Treasuries during the COVID-19 crisis.

We similarly introduce a separate factor for Treasury FRNs. FRNs may have liquidity and pricing characteristics that differ substantially from those of off-the-run coupon Treasuries or T-bills. In particular, an FRN may command a mark-to-market stability premium (Fleckenstein and Longstaff, 2020) or instead trade at an illiquidity discount (Hartley and Jermann, 2024).

The short rate used to discount FRN cash flows is r_t^f in equation (6). The corresponding

spread relative to the non-special curve is

$$\lambda_t \equiv r_t^f - r_t. \tag{8}$$

We likewise leave λ_t unrestricted in sign, allowing the model to capture either an FRN illiquidity discount or a liquidity premium. The factor F_t captures the component of this spread that is specific to the FRN market, while δ_X^f allows the FRN spread to co-move with the non-special Treasury curve.

All three short rates therefore share the same state dynamics and variance parameters (κ, Σ) but differ in their intercepts and factor loadings. The two spread curves can nonetheless respond to the common non-special factors through the unrestricted loadings δ_X^b and δ_X^f. This structure allows us to model three related but distinct Treasury curves within a single affine term structure framework.

4.2 FRN Pricing

Throughout this subsection we state cash flows and prices per \$1 of face value, for both the T-bill and the FRN.

T-bill price. The 13-week T-bill auction price takes an exponentially affine form in Z_t:

$$P_t = \exp(-\alpha^b - \beta^{b\prime} Z_t), \tag{9}$$

where the loadings α^b and β^b are derived from the T-bill short rate (5) and account for the settlement lag between the auction and issue dates (Appendix A).

FRN daily interest. A Treasury FRN with the auction-determined spread θ accrues daily interest, so that on each accrual day s the interest is

$$\frac{b_s + \theta}{360} = \frac{1 - P_s}{91\, P_s} + \frac{\theta}{360} = \frac{e^{\alpha^b + \beta^{b\prime} Z_s} - 1}{91} + \frac{\theta}{360}, \tag{10}$$

where b_s is the simple interest rate from the most recent 13-week T-bill auction before day s, subject to the 2-business-day lockout before each coupon and issue date described in Section 2.2. Interest accrues daily and is paid quarterly on dates $T_1, T_2, \ldots, T_M$, where $M = 4T$ is the number of quarterly periods in the life of a T-year security. We omit the zero floor of (1) here and in the estimation. It never binds over our sample: the smallest

value of $b_s + \theta$ during the life of any outstanding issue is +3.3 bp, on 23 March 2020, and the 13-week auction High Rate has never been negative.

FRN price. The model-implied price of a T-year FRN as of date t is

$$V_t = E_t^Q \left\{ \sum_{k:\, T_k > t} e^{-\int_t^{T_k} r_u^f \, du} \left(\sum_{s=T_{k-1}+1}^{T_k} \frac{b_s + \theta}{360} \right) + e^{-\int_t^{T_M} r_u^f \, du} \right\}. \tag{11}$$

Three features of this expression are worth emphasizing. First, the discount factor $e^{-\int_t^{T_k} r_u^f du}$ uses the FRN discount rate r_u^f, embedding the liquidity discount or premium λ. Second, the coupon rate b_s is determined by the T-bill auction price P_s, which depends on the T-bill short rate process r_u^b and therefore reflects the convenience premium process c. Third, for future floating coupon payments the expectation involves three distinct time indices: the valuation date t, the T-bill auction date s, and the coupon payment date T_k, with $t < s < T_k$. This requires evaluating

$$E_t^Q \left[e^{-\int_t^{T_k} r_u^f \, du} \, e^{\beta^{b\prime} Z_s} \right], \tag{12}$$

which mixes the FRN discount curve (for determining the discount rate) and the T-bill curve (for determining the cash flows). Mixing the two curves is what the security's structure requires. A convenience premium compensates an investor for holding a bill, and an FRN is not a bill: its coupons are indexed to bill auction rates, and so track the premium embedded in them, but the security itself confers none of the money-like services that give rise to it. Applying the law of iterated expectations at date s collapses the discounting from s to T_k into a term exponentially affine in Z_s. What remains amounts to the standard problem of pricing an exponentially affine payoff on the FRN discount curve, and its solution is again exponentially affine in Z_t; the derivation is in Appendix A.

When c_t and λ_t are identically zero—so that bills and FRNs are priced on the non-special curve—the model reduces to standard no-arbitrage FRN pricing under a single risk-free curve. Deviations of the market price from this no-arbitrage benchmark are absorbed by the three-curve structure: c shifts the FRN coupon, a T-bill convenience premium suppressing it and an inconvenience raising it, while λ shifts the rate at which the FRN's cash flows are discounted.

Par spread. The daily accrual (10) is affine in θ and enters (11) linearly, so the model price is exactly affine in the spread,

$$V_t = A_t(T) + \theta\, B_t(T), \tag{13}$$

where $A_t(T)$ is the price of the same security with $\theta = 0$ and $B_t(T)$ is the present value of a unit spread paid over its remaining life. The *par spread* of a newly issued T-year FRN is the value of θ at which the security prices at par, $V_t = 1$:

$$\theta_t(T) = \frac{1 - A_t(T)}{B_t(T)}. \tag{14}$$

The par spread is the model's counterpart of the discount margin that clears an FRN auction.

4.3 Physical Dynamics

Under the physical measure P the state follows

$$dZ_t = \kappa^P\big(\mu^P - Z_t\big)\, dt + \Sigma\, dB_t^P, \tag{15}$$

with the same diffusion matrix Σ as in (2). The absence of an intercept in (2) is a normalization; under P both the long-run mean μ^P and the mean-reversion matrix κ^P are free, and κ^P in particular is not restricted to be block diagonal. The two drifts differ by $\Sigma\Lambda_t$, where Λ_t is the market price of risk, so leaving κ^P and μ^P unrestricted is the same as letting Λ_t be an unrestricted affine function of the state.

Both the P and Q measures are needed for subsequent analysis in the paper. Consider issuing a T-year FRN in place of a T-year fixed-coupon note of the same face value. The note commits the Treasury to a known annual cost, its par yield $y_t^{\text{par}}(T)$. The FRN commits it to the reference rate as that rate turns out, plus the contractual spread, so its expected annual cost is $E_t^P[\overline{r^{\text{ref}}}] + \frac{365}{360}\,\theta_t(T)$, where $r_t^{\text{ref}} = y_t^b(\tau)$ is the τ-maturity bill yield to which the coupon is indexed—τ is thirteen weeks for the two-year FRN currently issued by the Treasury—and $\overline{r^{\text{ref}}}$ is its average over $[t,\, t+T]$. The expected annual saving/dissaving is the

difference,[8]

$$\text{Saving}_t(T) \;=\; \underbrace{y_t^{\text{par}}(T)}_{\text{fixed-rate cost}} \;-\; \underbrace{\left(E_t^P\left[\overline{r^{\text{ref}}}\right] + \tfrac{365}{360}\,\theta_t(T)\right)}_{\text{expected floating cost}}. \tag{16}$$

The quantity in (16) is closely related to the term premium, but is not the same object. The term premium at horizon T is the gap between a T-year yield and the average short rate expected to prevail over the same horizon, and is the compensation investors require for bearing duration risk. The saving/dissaving departs from it in two ways: the fixed leg enters as the par yield of a coupon-bearing security rather than as a zero-coupon yield, and the calculation further deducts the par spread, which is absent from the term premium calculation.

The three terms also do not come from the same side of the model. The par yield $y_t^{\text{par}}(T)$ and the par spread $\theta_t(T)$ are both risk-neutral quantities. The par spread is given by (14); the par yield is the quarterly coupon rate at which a non-special fixed-coupon security of maturity T prices at par, obtained from the zero-coupon bond prices that (2) implies. On the other hand, the expected average reference rate $E_t^P[\overline{r^{\text{ref}}}]$ is a time-series forecast and requires the physical dynamics (15). The saving/dissaving thus sets the market's required compensation for duration risk against what the Treasury should expect to pay without it.

Equation (16) compares an FRN with off-the-run fixed-rate debt rather than the on-the-run note Treasury would issue. To the extent that on-the-run yields are lower than off-the-run yields, the implied saving is overestimated. To gauge this bias, we measure the average on-the-run premium over 2014–2025: it is 1.16, 0.55, 1.19, 0.34 and 0.76 bps at two, three, five, seven and ten years, each with a Newey-West t statistic of at least 3.8. The two-year figure agrees with Table 2: the T-bill curve, anchored by the on-the-run and first off-the-run two-year notes, averages 2.232% against 2.241% for the non-special curve, a gap of about one basis point.

[8] The 365/360 factor on the spread puts the two legs on a common day count. The par yield is an annual rate on a 365-day year, while the FRN accrues $(b_s + \theta)/360$ per calendar day by (10), so over a year the spread costs $\frac{365}{360}\theta$. The reference-rate leg needs no such factor. The 365-day annualization of the index leg of (10), $365 \times \frac{b_s}{360} = (1 - P_s)/(\tau P_s)$, corresponds, up to convexity, to the τ-maturity yield itself, which is what r^{ref} denotes.

5 Empirical Implementation

The model in Section 4 can be used to price any Treasury security. Estimating the model, however, is considerably easier when using spot rates than prices. A spot rate is affine in the state, whereas the price of a coupon-paying bond, or of an FRN, is generally a sum of exponentially affine terms rather than an affine function of the state. The standard approach is therefore to first strip zero curves from observed prices, as we do for the off-the-run and T-bill curves in Section 3.2. FRNs, however, present a different challenge because their cash flows are not fixed, so the same stripping procedure does not apply.

In Section 5.1, we develop a two-step bootstrap procedure for FRNs. In the first step, we estimate a two-curve model using only T-bill and non-special yields. In the second step, we augment this model with a constant spread, maturity by maturity, so that it exactly fits the observed prices of FRNs. The resulting two-curve model, together with these maturity-specific spreads, produces a set of spot rates that incorporates the liquidity characteristics of the FRN market. Section 5.2 then describes our Kalman filtering approach, including the measurement and transition equations and the restrictions imposed. Section 5.3 reports the model's fit to the three curves used in estimation. Finally, Section 5.4 provides a more demanding test: whether the model can match the observed market prices of Treasury FRNs, which are not themselves direct inputs to the estimation.

5.1 Bootstrapping the FRN-Implied Spot Rates

We use a two-step procedure to construct a term structure of interest rates that shares the liquidity characteristics of the FRNs themselves.

The first step is to estimate, over the full sample, a *two-curve* term structure model: the model of Section 4 with the FRN block removed. We keep the first four elements of the state, $(X_t', S_t)'$—the three non-special factors and the T-bill factor—together with the blocks of (2) and (3) that govern them, and the two short rates they price: the non-special rate (4) and the T-bill rate (5). The FRN factor F_t and the FRN short rate (6) play no part at this stage. We call this model two-curve because it prices exactly two: the non-special curve, fitted to off-the-run coupon securities and the overnight repo rate, and the T-bill curve.[9]

Estimation of this two-curve model is standard. Every bill yield and every non-special yield is affine in (X_t, S_t), and the state follows a Gaussian vector autoregression, so the

[9]The physical dynamics matter little at this stage. Pricing FRN cash flows requires only the risk-neutral parameters; the physical parameters enter the bootstrapped spot rates solely through the filtered state.

system is linear and the Kalman filter delivers the likelihood exactly. As is standard, we assign one measurement-error standard deviation per curve.

The second step applies the estimated two-curve model to FRN prices. The gap between model-implied and observed FRN prices is what we use to back out spot rates that embed the risk characteristics of the FRN market. For simplicity, consider first, on date t, an FRN with a single remaining cash flow, due at $T_1 > t$, and market price $P_t^{(1)}$. That cash flow is the principal together with interest accrued at the daily 13-week bill rate $\{b_s\}$ plus the auction-determined spread $\theta^{(1)}$. Discounting it at the non-special short rates $\{r_u\}$ would in general not reproduce $P_t^{(1)}$. We therefore add a constant spread η_1 to the short rate over $[t, T_1]$:

$$P_t^{(1)} = E_t^Q \left\{ e^{-\int_t^{T_1} (r_u + \eta_1)\, du} \left(\sum_{s=t}^{T_1} \frac{b_s + \theta^{(1)}}{360} \; + \; 1 \right) \right\}. \tag{17}$$

The η_1 that solves this equation measures the illiquidity discount or premium on an FRN maturing at T_1. With η_1 in hand we can price a zero-coupon FRN that matures at T_1 and inherits the liquidity characteristics of the FRN market,

$$E_t^Q \left\{ e^{-\int_t^{T_1} (r_u + \eta_1)\, du} \right\}, \tag{18}$$

whose yield is the FRN-implied spot rate at maturity T_1.

A similar process extends to an FRN with two remaining cash flows, due at T_1 and T_2. Its first cash flow can be priced as

$$E_t^Q \left\{ e^{-\int_t^{T_1} (r_u + \eta_1)\, du} \left(\sum_{s=t}^{T_1} \frac{b_s + \theta^{(2)}}{360} \right) \right\}, \tag{19}$$

where $\theta^{(2)}$ is the auction-determined spread on this FRN. Both η_1 and the dynamics of $\{r_u\}$ and $\{b_s\}$ are already known, so this expression is readily computable; call it $V_t^{(1)}$. Subtracting it from the market price leaves the present value of the cash flow due at T_2, namely the interest accrued between T_1 and T_2 together with the principal:

$$P_t^{(2)} - V_t^{(1)} \;=\; E_t^Q \left\{ e^{-\int_t^{T_2} (r_u + \eta_2)\, du} \left(\sum_{s=T_1+1}^{T_2} \frac{b_s + \theta^{(2)}}{360} \; + \; 1 \right) \right\}. \tag{20}$$

As with η_1, the spread η_2 applies to the short rate over the whole interval $[t, T_2]$. Solving (20) for η_2 gives the price of a zero-coupon FRN maturing at T_2,

$$E_t^Q \left\{ e^{-\int_t^{T_2} (r_u + \eta_2)\, du} \right\}, \tag{21}$$

whose yield is the FRN-implied spot rate at maturity T_2. Iterating in this way delivers zero-coupon FRN prices out to T_M.

The procedure produces a panel of zero-coupon prices on a curve that carries the FRNs' own liquidity characteristics. The corresponding yields are affine in the state vector, with loadings given by the same formulas as for the other two curves (Appendix A), and it is this affinity that keeps estimation tractable. An FRN price itself is not affine in Z_t: by (11) it is a sum of exponentially affine functions, one per coupon date, and the coupon terms involve the three-time-index expectation (12). Handling that nonlinearity directly in estimation would be substantially more involved.

Treasury issues a new two-year FRN each quarter, so outside the early part of the sample eight distinct securities are outstanding on almost every date. Each of them ages from one date to the next, so the bootstrap delivers zero yields at maturities that shift daily. To obtain a panel at fixed maturities, we interpolate the bootstrapped yields on each date onto a grid of seven tenors—one through seven quarters, or 0.25 to 1.75 years—using a monotone piecewise cubic scheme. The grid stops at seven quarters because a full two years of remaining maturity is available only on an issue date.

Our procedure is similar in spirit to the constant-forward bootstrap of Fama and Bliss (1987), which recovers a cross-section of zero yields on each date by assuming that the forward rate between successive observed maturities is constant.[10] As with Fama-Bliss, what comes out is an approximation rather than an observed quantity. The real test is whether a model estimated on these approximate inputs can price actual FRNs, whose prices are nonlinear in the state. Section 5.4 takes up that question: there we price every outstanding FRN on every date using the exact formula (11) rather than any linearization, and compare the result with observed market quotes. The pricing errors turn out to be small at every maturity.

[10] Fama-Bliss recover the curve one maturity at a time by holding the forward rate constant between successive observed maturities. Our η_n is instead a constant spread on the short rate over the whole interval $[t, T_n]$: the FRN-implied zero yield at T_n is the non-special zero yield plus η_n. What the two algorithms share is the recursive logic—each new maturity is identified by netting out the present value of the cash flows already priced.

5.2 Estimation by Kalman Filtering

The three-curve model is estimated by the Kalman filter, with three sets of inputs: the T-bill curve, the non-special curve, and the FRN-implied spot rates derived in Section 5.1.

Measurement equation. Let Y_t stack the yields observed at date t: the seven T-bill maturities and the eight non-special maturities of Section 3.3, together with the seven FRN-implied tenors of Section 5.1. Each row is affine in the state,

$$Y_t = A_Z + B_Z Z_t + \varepsilon_t, \tag{22}$$

where A_Z and B_Z stack the loadings corresponding to the three curves. These loadings are computed from the risk-neutral dynamics (2)–(6), using the recursions given in Appendix A. Following standard practice, measurement errors are assumed independent across rows and dates, with one standard deviation per curve—σ_e^b on the T-bill rows, σ_e on the non-special rows and σ_e^f on the FRN rows—so the specification adds three error parameters rather than one per maturity.

State rotation. For estimation we work with a "rotation," or a change of state variables, rather than with the latent state Z_t itself. Specifically, we use the yield portfolios $P_t = WY_t$ as factors; the dimension of P_t is 5×1. The loading matrix W is specified such that the first three entries of P_t are the leading principal components of the non-special yields, the usual level, slope, and curvature factors. The fourth is the average bill yield less the average non-special yield at maturities of one year or less, which captures the gap between the two curves at the short end. The fifth is an equally weighted average of the seven FRN-implied yields less an equally weighted average of all eight non-special yields, which captures the difference in level between those two curves. The rotation leaves the model observationally equivalent. Its benefit, as JSZ show, is that it makes the physical parameters more interpretable and better identified, which substantially improves the conditioning of the numerical optimization of the Kalman objective.

Transition equation. Under the physical measure the portfolios follow a first-order Gaussian vector autoregression at the weekly observation interval,

$$P_{t+1} = K_{0P} + K_{1P} P_t + u_{t+1}, \qquad u_{t+1} \sim \mathcal{N}(0, \Sigma_P). \tag{23}$$

This is the discrete-time counterpart of (15), written in the rotated states P_t rather than in Z_t.

The FRN block before August 2015. Treasury FRNs were first issued in January 2014, and the bootstrap of Section 5.1 requires a cross-section of at least seven outstanding issues spanning the full tenor range. The FRN rows of Y_t are therefore populated on 523 of the 1,350 weekly dates, the first being 3 August 2015. Rather than restrict the sample to those 523 dates, we let the dimension of the measurement equation vary with the date: where the FRN rows are missing, the Kalman objective uses the fifteen T-bill and non-special rows alone. The full twenty-six years of yield data thus identify the non-special and T-bill blocks, while the FRN factor is identified from the 523 dates on which the entire maturity spectrum of FRNs trades and the bootstrap output is complete.

Restrictions. One restriction is imposed beyond the specification of Section 4: the mean reversion κ_F of the FRN factor is restricted to be non-negative, so that the FRN spread is stationary under Q. This is not a matter of numerical convenience. Were κ_F zero or negative, the FRN curve would have no well-defined long end: the convexity term in a Gaussian zero yield grows faster than the maturity divides it, so the asymptotic zero-coupon rate diverges rather than settling down. Q-stationarity is what makes the asymptotic rate finite, and the restriction imposes no more than that. The mean reversions of the remaining four factors are left unrestricted in sign, and all are estimated to be positive, so those factors are Q-stationary without any restriction being imposed.

Concentrated parameters. Four groups of parameters admit a closed-form optimizer conditional on the rest and are analytically concentrated out rather than numerically searched over. The three short-rate intercepts r_∞, r^b_∞ and r^f_∞ enter the measurement equation linearly, so the Kalman objective is exactly quadratic in them and each implies a one-dimensional linear first-order condition; the intercepts are separable because each enters only its own curve's rows. The physical intercept K_{0P} concentrates out of the transition density in the same way. This reduces the dimension of the numerical search without approximation.

5.3 Fit to the Three Yield Curves

Table 3 reports in-sample pricing errors, defined as observed minus model-implied yield and evaluated at the filtered state. Overall, all three curves fit reasonably well. The T-bill

curve has an overall RMSE of 4.43 bps, with no maturity worse than 5.62 bps and mean errors no larger than 1.4 bps in absolute value; at the 91-day maturity, the tenor that sets FRN coupons, the RMSE is 4.59 bps. The non-special curve fits slightly less tightly, with an overall RMSE of 6.64 bps, though this is on par with pricing errors typically reported in the literature.[11] The largest errors are at the very short end: at the overnight and 91-day maturities, the RMSEs are 13.44 bps and 10.32 bps, and the mean errors 3.77 bps and -5.90 bps, respectively. The remaining six non-special maturities have mean errors no larger than 2.9 bps in absolute value.

The overnight repo rate, which resides on the non-special curve as a one-day yield (Section 3.3), averages 8.7 bps *above* the 91-day non-special yield. This results in an inverted 0–90 day segment of the curve in 65% of weeks. The model must fit both sets of observations, so at 91 days the model-predicted yield is above the observed one, and this inversion accounts for essentially the whole of the mean errors.

The inverted yield curve at the very short maturity is informative. An overnight general-collateral repo rate is secured against a broad basket of Treasury collateral rather than against any individual security, so it is largely free of the specialness attaching to particular issues. That it exceeds the fitted 91-day yield over a large part of the sample suggests that even off-the-run short-maturity coupon securities retain some specialness of their own.[12]

The FRN-implied curve is fit more tightly than either of the other two, with an overall RMSE at 1.18 bps with no tenor worse than 1.42 bps. We note that the small RMSE is not mechanical since the model has a single FRN-specific factor, F_t, with which to match seven tenors. Section 5.4 further illustrates the goodness of fit of this model by directly comparing the model-implied FRN prices to observed prices.

5.4 Pricing Observed FRNs

We further assess the model's performance by comparing model-implied and actual FRN prices. We take the estimated parameters of Section 5.3 as given, extract the filtered state on each weekly date of the estimation sample, and compute the model price of each outstanding

[11] For example, Joslin, Le, and Singleton (2013) report root-mean-squared fitting errors of 8 to 18 bps across maturities from six months to ten years.

[12] An alternative reading of the inversion is that the repo rate is high rather than the coupon yield low: GCF repo is interdealer borrowing and can carry a funding premium of its own, particularly at quarter- and year-ends. Such episodes are short-lived, however, and the seven-day trailing median we apply to the overnight series (Section 3.3) removes them before estimation. The inversion that remains is a persistent feature, present in 65% of weeks in the sample, rather than a residue of dated funding spikes, which is why we read it as specialness in short-maturity coupon securities rather than as a premium in the repo rate.

Table 3: In-sample pricing errors by curve and maturity, in basis points. Errors are the yields used in estimation minus model-implied yields evaluated at the filtered state, weekly, 2000–2025 (FRN-implied yields from August 2015). P5 and P95 are the 5th and 95th percentiles of the error distribution.

Maturity	**Mean**	**RMSE**	**P5**	**P95**
Panel A: T-bill curve				
30-day	0.88	4.44	−6.5	6.8
60-day	−0.87	3.40	−5.2	3.9
91-day	−1.06	4.59	−7.2	5.9
180-day	0.42	4.25	−5.7	7.3
270-day	0.26	4.19	−6.0	6.4
365-day	1.40	4.25	−5.0	7.0
730-day	−1.00	5.62	−9.3	7.3
All	**0.01**	**4.43**	−6.7	6.6
Panel B: non-special curve				
overnight	3.77	13.44	−9.0	26.0
3-month	−5.90	10.32	−20.8	3.3
6-month	−0.61	4.95	−8.5	7.2
1-year	2.83	6.04	−6.0	11.9
2-year	0.06	6.03	−8.7	9.7
3-year	−0.37	5.21	−9.1	7.9
4-year	0.34	4.00	−6.2	6.6
5-year	0.04	7.85	−13.0	11.8
All (excl. overnight)	**−0.52**	**6.64**	−11.3	8.9
Panel C: FRN-implied curve				
1 quarter	−0.10	1.31	−2.0	1.1
2 quarters	0.40	1.42	−0.8	1.7
3 quarters	−0.10	1.19	−1.7	1.4
4 quarters	−0.37	1.16	−2.3	1.1
5 quarters	−0.01	0.84	−1.4	1.4
6 quarters	−0.01	0.86	−1.3	1.4
7 quarters	0.18	1.35	−1.7	2.6
All	**−0.00**	**1.18**	−1.7	1.6

FRN from the present-value expression (11), projecting coupon payments off the model's T-bill curve and discounting them on the model's FRN curve. The calculation uses the full contractual structure of the security, including weekly index resets, the two-business-day lockout, actual/360 daily accrual, and the three-time-index expectation (12). We report two errors. The *price error* is the observed bid-ask midpoint less the model price, in cents per $100 of face value. The *discount-margin error* expresses the same discrepancy as a spread:

the fixed spread over the reference rate that reconciles the model price to the observed midpoint, less the security's contractual spread. By (13) the model price is exactly affine in that spread, so the inversion is closed-form. The contractual spread is the discount margin at which the issue's original auction cleared.

Table 4: Root-mean-squared FRN pricing errors by remaining maturity. Price errors are the observed bid-ask midpoint minus the model price, in cents per $100, sampled weekly over the 2014–2025 sample. Margin errors are the same errors expressed as discount margins (model minus observed), in basis points. The two-curve benchmark model, estimated separately on the T-bill and non-special curves, prices the FRNs off the non-special curve.

		Price RMSE (cents)		**Margin RMSE (bps)**	
Remaining maturity	N	Two-curve	Three-curve	Two-curve	Three-curve
0–3 months	521	1.13	0.38	9.38	3.85
3–6 months	531	2.37	1.18	6.21	2.81
6–9 months	541	3.55	1.80	5.58	2.74
9–12 months	559	5.40	2.60	6.10	2.96
12–15 months	578	8.70	3.16	7.64	2.79
15–18 months	577	13.48	3.91	9.76	2.84
18–21 months	595	18.92	5.31	11.65	3.28
21–24 months	638	25.90	7.58	13.86	4.06
All	**4,540**	**13.41**	**4.05**	**9.34**	**3.22**

Table 4 reports both errors by remaining-maturity bucket, alongside the same statistics from a separately estimated two-curve model, fitted to the T-bill and non-special curves alone, as a benchmark. Across all maturities, the mean price error is +0.03 cents and the mean discount-margin error +0.06 bps, with bucket-level means (unreported) never exceeding 0.55 bps in absolute value. The average magnitude of the pricing errors is small as well: the overall RMSE is 4.05 cents in price terms and 3.22 bps in discount-margin terms. The three-curve model therefore prices observed FRNs accurately on average and without systematic bias.

The price-error RMSE varies sharply across the buckets, from 0.38 cents at zero to three months remaining to 7.58 cents at twenty-one to twenty-four. This maturity pattern, to some extent, is to be expected. An FRN carries almost no interest-rate duration, since the coupon resets weekly, but its spread is fixed at auction for the whole of its life, so the fixed spread introduces a duration effect proportional to the remaining maturity. Expressing the same discrepancy as a discount margin largely removes the duration effect. The margin RMSE

then ranges only from 2.74 to 4.06 bps, and it shows no tendency to widen as maturity lengthens. By this metric, the estimated model prices FRNs about equally well at every remaining maturity.

Figure 4 shows the same contrast: the interquartile range of price errors narrows from [−4.59, 4.38] cents at twenty-one to twenty-four months remaining to [−0.15, 0.20] at zero to three, while in discount-margin terms it moves only from [−2.47, 2.28] to [−1.22, 1.81] bps. Across the sample as a whole the interquartile range of price errors is [−0.99, 1.68] cents, against an average quoted bid-ask half-spread of 1.60 cents, the cost of crossing the market on these securities.

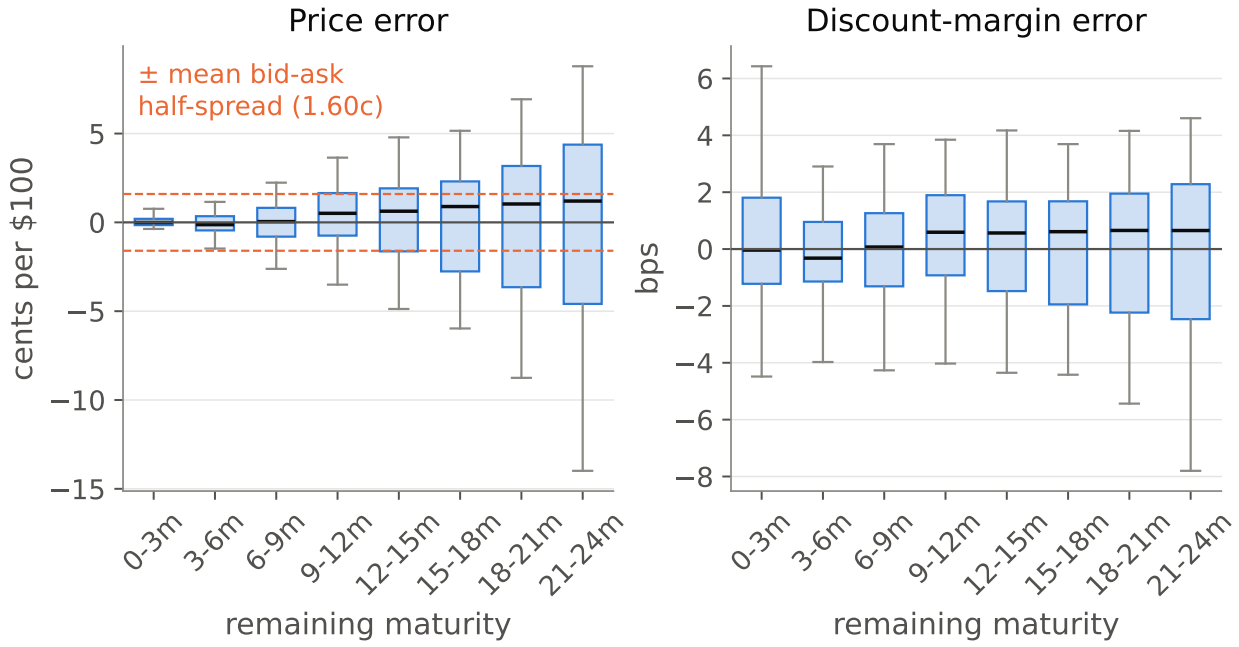


Figure 4: Three-curve pricing errors by remaining maturity over 2014–2025. Boxes show the interquartile range and the median; whiskers show the 5th and 95th percentiles. Left: price errors (observed minus model) in cents per $100, with dashed lines at plus/minus the mean bid-ask half-spread. Right: the same errors as discount-margin errors (model minus observed), in basis points.

Table 4 also illustrates the cost of omitting a dedicated FRN curve. The two-curve model overprices FRNs by 8.53 cents on average, or 5.90 bps of discount margin, and its price RMSE rises monotonically with remaining maturity, reaching 25.90 cents for securities with twenty-one to twenty-four months remaining. Figure 5 shows the same pattern for six

individual securities: the two-curve model-implied price lies above the observed midpoint for most of each security's life, whereas the three-curve model-implied price tracks the midpoint much more closely. This is the pattern anticipated in Section 4: without an FRN-specific discount curve, a model fitted to off-the-run coupon securities and Treasury bills systematically overvalues FRNs relative to the prices investors actually pay. Adding the FRN curve reduces the price RMSE by about 70% and the discount-margin RMSE by about 66%. On average the estimated model discounts FRN cash flows more heavily than those of off-the-run coupon securities, implying that an illiquidity discount, rather than a mark-to-market premium, dominates in our sample.

5.5 T-Bill Premium and FRN Illiquidity Discount

Figure 6 plots the filtered path of the model's factors and the spreads it implies between the three curves over the full sample from 2000 through 2025, and Figure 1 shows the model-implied average T-bill, non-special, and FRN-implied yield curves over the FRN-era sample, 2014–2025.

The upper panel of Figure 6 shows the first three principal components of non-special bond yields, which play the usual roles of level, slope, and curvature.

The middle panel plots the difference between non-special and T-bill yields at one month and at one year. This spread is the convenience premium: investors accept a lower yield on bills because bills deliver cash-like services beyond their cash flows (Greenwood, Hanson, and Stein, 2015; Nagel, 2016; Krishnamurthy and Vissing-Jorgensen, 2012). The premium is concentrated at the very short end, where bills are treated as near-cash: it averages 15.2 basis points at one month over 2014–2025 against 0.5 basis points at one year, and the same ordering holds over the full sample, at 17.9 against 1.3 basis points. Accordingly, in Figure 1 the T-bill curve lies below the non-special curve, markedly so at short maturities. Unsurprisingly, the one-month series peaks during the 2007–2009 financial crisis, reaching 116.6 basis points in the week of 16 September 2008 and 108.6 basis points in August 2007.

The lower panel plots the difference between FRN-implied and non-special yields at one month and at two years. Unlike the T-bill convenience premium, the term structure of FRN spreads slopes strongly upward: the one-month difference averages -0.8 basis points over 2014–2025 and is negative in 59% of weeks, while the two-year difference averages 14.1 basis points and stays between 7.1 and 26.3. Correspondingly, in Figure 1 the FRN curve lies above the other two beyond the shortest maturities, with the gap widening as maturity lengthens. Fleckenstein and Longstaff (2020) argue that FRNs may command a

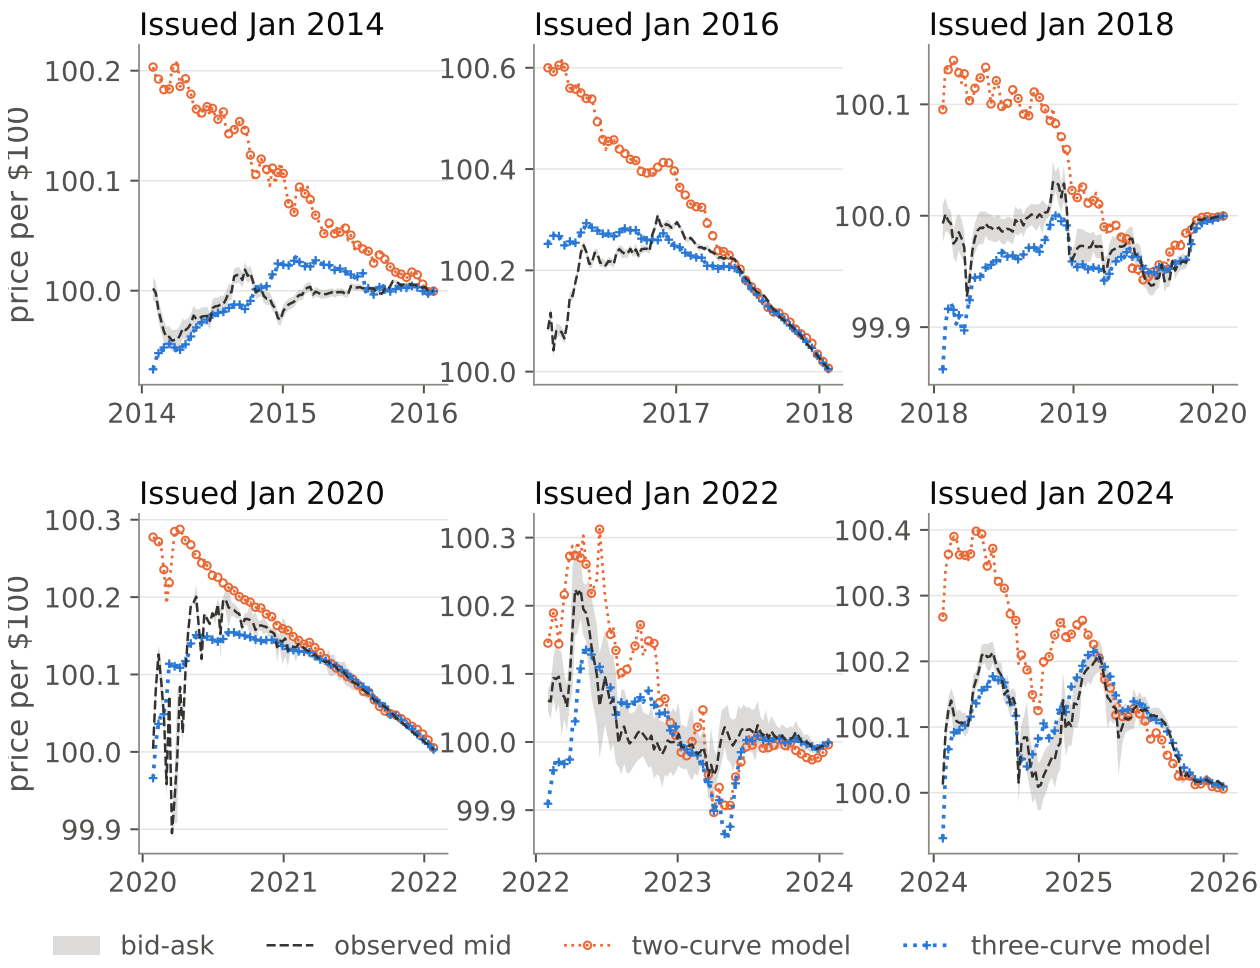


Figure 5: Observed and model-implied prices of the January FRN issue in each even year in the sample. The shaded band is the bid-ask range from Bloomberg and the dashed line its midpoint. The two-curve and three-curve models, estimated separately, price the FRNs off the non-special curve and the FRN curve, respectively.

pricing premium, and hence a lower yield, for their inherent mark-to-market stability. Our estimates suggest long-dated FRNs carry an illiquidity discount that more than offsets any value-stability premium. This is possible since the FRN market is relatively small and new and investors typically buy and hold in this market. We therefore refer to the excess yield of FRNs over the corresponding off-the-run yield as their (net) illiquidity discount.

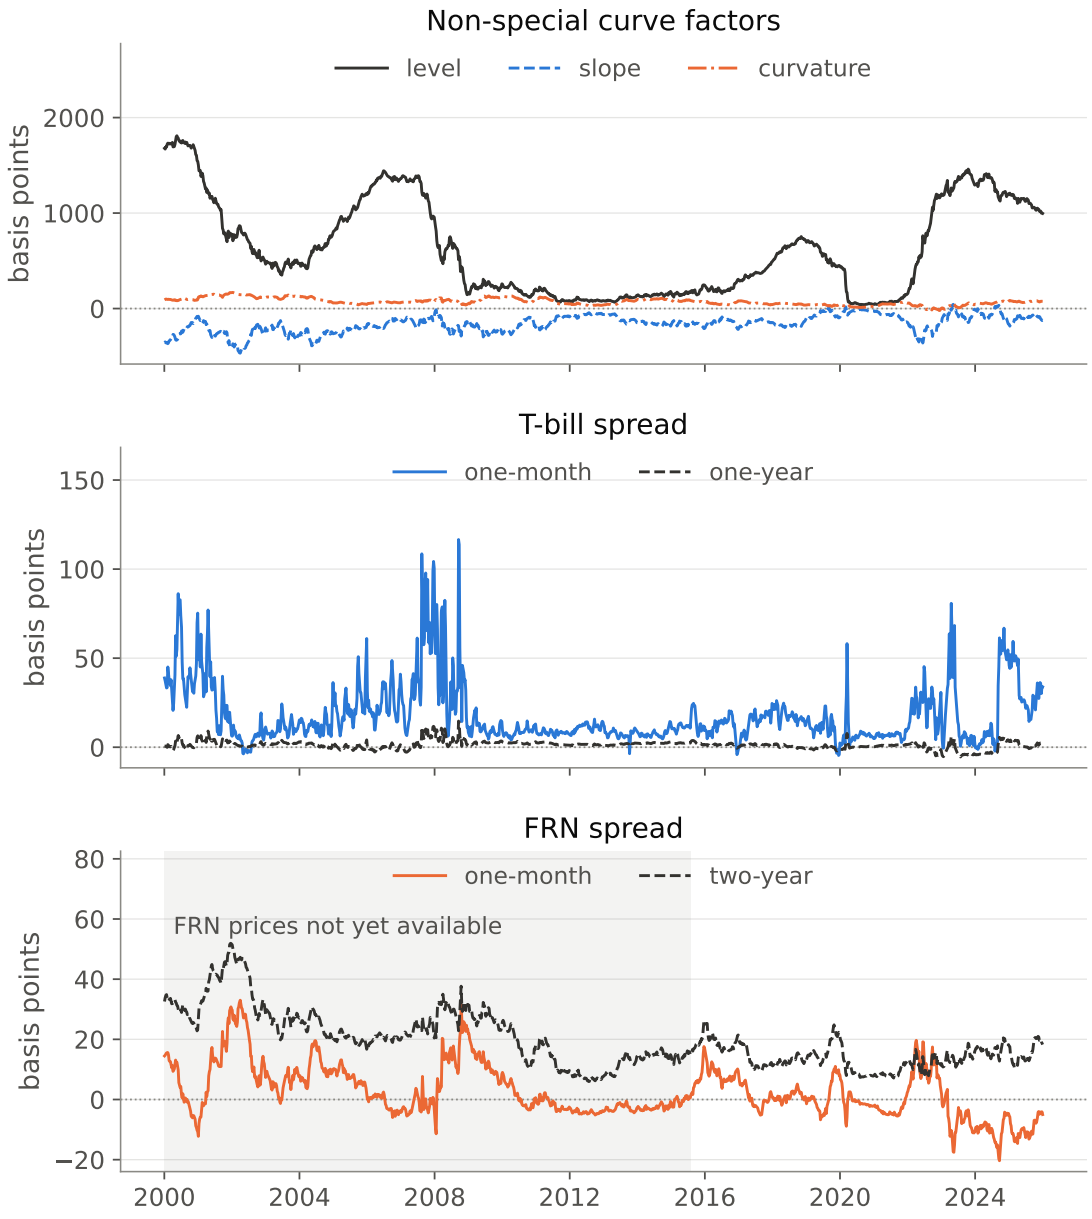


Figure 6: Filtered factors and spreads, weekly, 2000–2025. Upper panel: the first three principal components of non-special yields. Middle panel: non-special minus T-bill yields at one month and one year. Lower panel: FRN-implied minus non-special yields at one month and two years. FRN-implied yields begin in August 2015, so before then (shaded) the FRN spread is inferred only from its estimated covariance with the other factors.

6 Implications for Treasury Interest Cost

Section 5 established that the three-curve model gives an adequate description of the data. In this section we use the estimated model to study the interest costs faced by the Treasury. Section 6.1 measures what investors charge the Treasury for the illiquidity of floating-rate debt. Section 6.2 sets that cost against the term premium the Treasury avoids by issuing floating. Section 6.3 recomputes the expected saving under two alternative forecasts of the reference rate, from surveys and from the model re-estimated in real time, and compares all three with the savings realized on matured issues.

6.1 The FRN Liquidity Component

Recall that equation (14) gives the model-implied constant spread θ over the reference rate at which a newly issued T-year FRN prices at par. For a T-year FRN issued at t we define

$$\text{liquidity component}_t(T) \;\equiv\; \tfrac{365}{360}\Big(\theta_t^{\text{full}}(T) - \theta_t^{\text{ns}}(T)\Big), \tag{24}$$

where $\theta_t^{\text{full}}(T)$ is the par spread under the full three-curve model and $\theta_t^{\text{ns}}(T)$ the par spread under the counterfactual that the FRN's cash flows are discounted on the non-special curve and the security is therefore assumed to trade with the liquidity of an ordinary off-the-run fixed-rate issue.[13]

To illustrate, consider the FRN auction on October 25, 2023. The auction-determined spread was 17 bps. The estimated three-curve model prices the FRN par spread at 17.14 bps, nearly identical to the auction spread. If, however, we discounted the FRN cash flows using the non-special curve that day, the hypothetical spread would be -0.99 bps, essentially zero. In this example, the difference between the two spreads is $17.14 - (-0.99)$, or 18.13 bps. Adjusting for the actual/360 convention, we obtain $18.13 \times \frac{365}{360} = 18.38$ bps, which is our measure of the liquidity component of the FRN. On this date, virtually the entire spread the Treasury committed to pay on the new FRN was compensation for the security's illiquidity.

The top part of Figure 7 plots the total par spread θ^{full} in solid black and the non-special spread θ^{ns} in dotted blue for the two-year FRN from 2014 to 2025. Plotted in the lower part of Figure 7 is the liquidity component as defined in (24), which stays positive in every week of the sample. Investors have required a higher spread to hold the FRN than to hold an otherwise identical security that traded with the liquidity of the off-the-run market. The

[13] The $\frac{365}{360}$ factor annualizes the difference in spreads on an actual/365 basis, since θ is quoted on an actual/360 convention.

average liquidity component on the two-year FRN indexed to the 13-week bill is 14.04 bps from 2014 to 2025.

The hypothetical spread θ^{ns} is usually negative: it averages -3.07 bps since 2014 and lies below zero in 72.0% of weeks. The reason is the choice of index: the security pays the 13-week bill rate but resets weekly, so its holder earns a thirteen-week yield while bearing only a week of interest-rate risk. Such a mismatch in tenors tends to overcompensate the FRN holder when the short end of the yield curve slopes upward. Absent any illiquidity discount, par pricing of the FRN is typically restored by a negative spread.

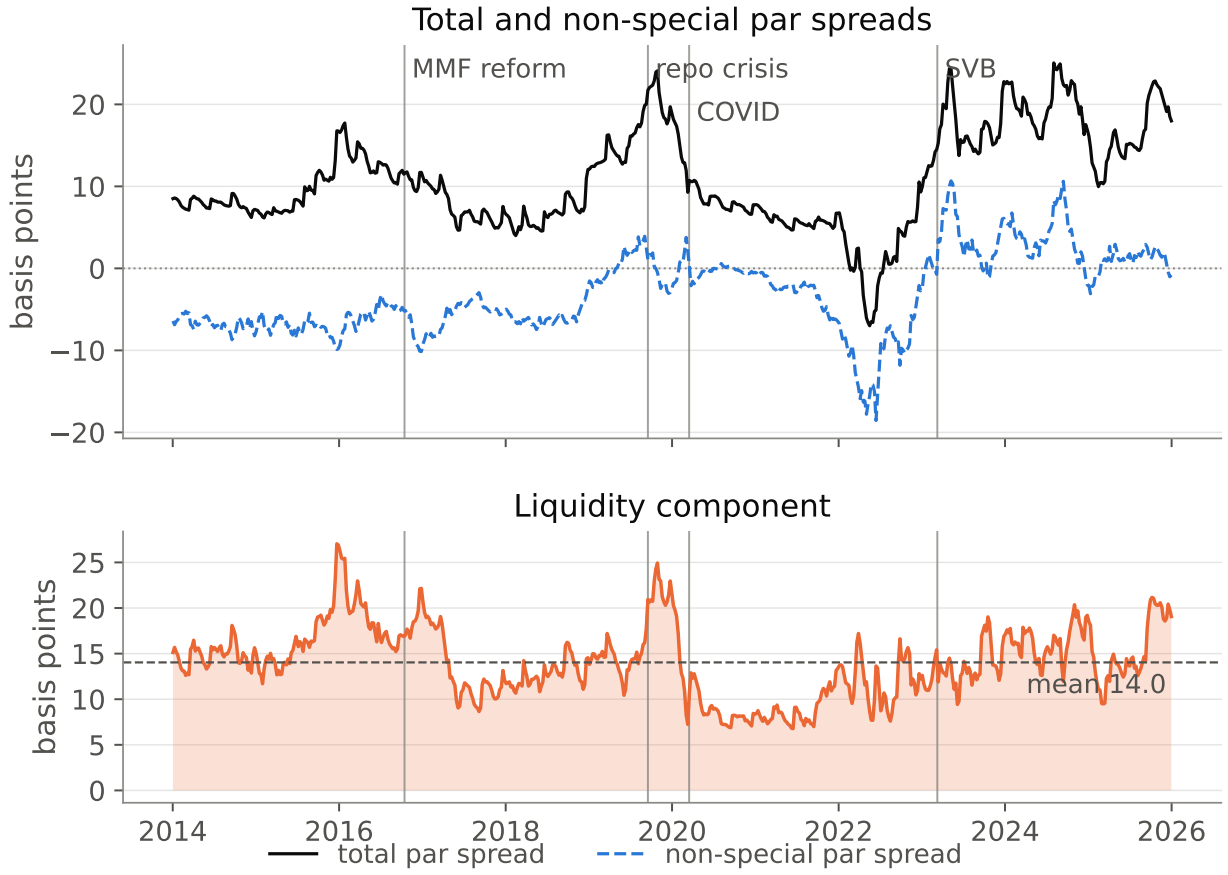


Figure 7: Model-implied par spreads on a newly issued two-year FRN of the current design, indexed to the 13-week bill. The total (non-special) par spread is obtained by discounting the FRN's cash flows off the model's FRN (non-special) curve. The lower panel plots the liquidity component—the difference between the two model-implied par spreads, adjusting for the actual/360 quoting convention (as in (24)). Vertical lines mark the October 2016 money market fund reform, the September 2019 repo market disruption, the March 2020 onset of COVID-19 and the March 2023 failure of Silicon Valley Bank.

6.2 Expected Savings against the Liquidity Component

Among the four objectives Treasury gave for issuing floating-rate notes was to "borrow at the lowest cost over time" (U.S. Department of the Treasury, 2013, 78 FR 46426). In response to Treasury's 2012 request for information, one primary dealer pointed out that FRNs indexed to a floating short rate would let Treasury "reduce its interest expense by an amount equal to the embedded term premium" (Misra, 2012).

The possibility that these savings could be offset by a liquidity effect was also acknowledged, although never quantified. Responding to the same request, Nomura wrote that "the savings in term premium may be offset, at least in the early stages, by a liquidity discount (similar to TIPS where debate still continues over whether the program has been cost effective)" (Nomura Securities International, Inc., 2012). Armed with twelve years of FRN market data, we can quantify both the term premium and illiquidity discount while maintaining cross-security consistency through the no-arbitrage restrictions.

We therefore compare the model-implied expected saving from pure term-premium avoidance, abstracting from any liquidity effect,

$$\underbrace{y_t^{\text{par}}(T)}_{\text{fixed-rate cost}} \; - \; \underbrace{\left(E_t^P\left[\overline{r^{\text{ref}}}\right] + \tfrac{365}{360}\,\theta_t^{\text{ns}}(T) \right)}_{\text{expected floating cost}} \tag{25}$$

against the liquidity component defined in (24). Equation (25) differs from the net cost-saving formula (16) of Section 4.3 in one respect only: the spread θ_t^{full} is replaced by the counterfactual θ_t^{ns}. Subtracting the liquidity component (24) from (25) returns (16) exactly.

Figure 8 plots the two quantities against each other over the FRN period. The liquidity component is relatively stable, confined to a band between 7 and 27 bps over the entire period. By contrast, the expected saving before liquidity swings far more widely: from about +130 bps in mid-2023 and mid-2024 to about −110 bps in mid-2022. Averaged over January 2014 through December 2025, the expected saving is 15.7 bps against a liquidity component of 14.0 bps, leaving a net saving of 1.7 bps. Over the FRN era as a whole, the term premium the Treasury avoids and the liquidity it gives up very nearly offset.

The sign of the expected saving follows a slow-moving cycle. The peaks occur at or shortly after the end of Federal Reserve hiking cycles, in 2019 and in 2023–24, whereas the trough in mid-2022 comes shortly after the start of the 2022 hiking cycle. This is not a coincidence. The level of short-term interest rates tends to mean-revert, a key fact captured by term structure models of interest rate including ours. Thus, periods with higher than usual short-

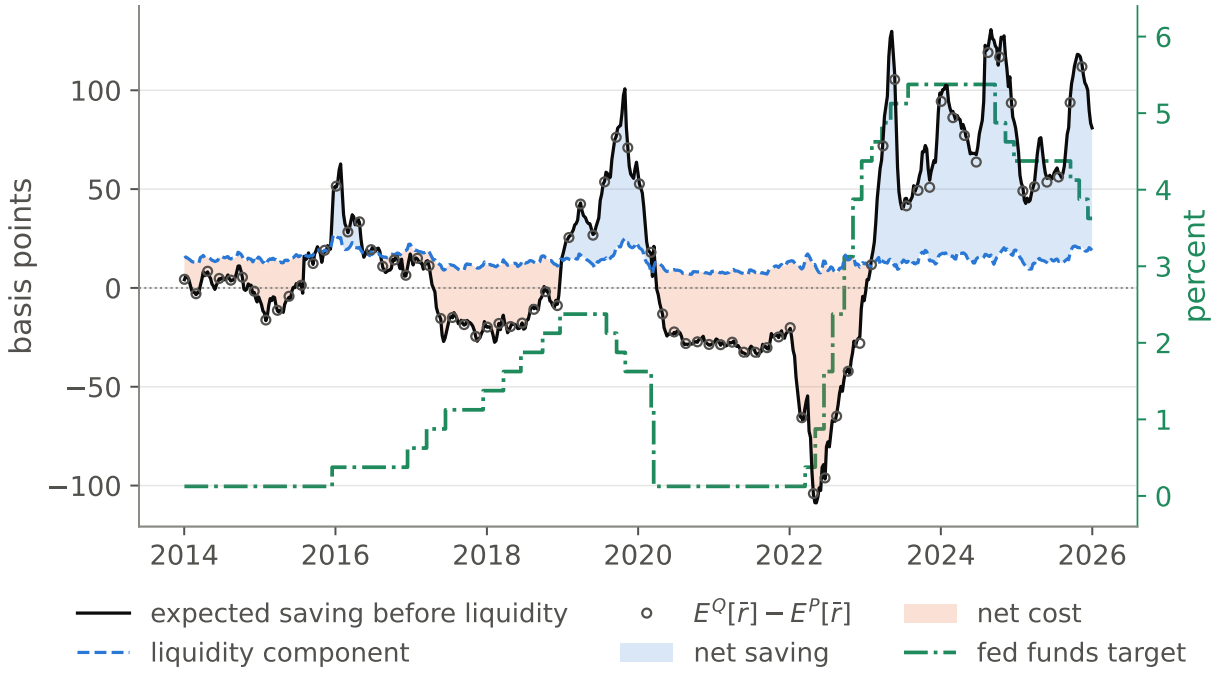


Figure 8: Expected saving and liquidity component for a newly issued two-year FRN of the current design, indexed to the 13-week bill, over 2014–2025. Solid line: expected saving before liquidity, (25). Dashed line: liquidity component, (24). The shading between them is the net saving or net cost. Circles: $E_t^Q[\overline{r^{\text{ref}}}] - E_t^P[\overline{r^{\text{ref}}}]$, sampled every eighth week. Dash-dotted line (right axis): midpoint of the FOMC's target range for the federal funds rate.

term interest rate are expected to be followed by falling short-term interest rates, but that is only half the puzzle. The other half is whether the two-year par yield sufficiently reflects the reduced short rate expectation: it does not. Thus, these periods also feature a high term premium, as measured by the difference between the two-year par yield and expected short rate in the next two years. Yet, they are not obvious to the naked eye because the overnight–2 year segment of the yield curve tends to be flat or even inverted during these periods. This particular combination—a material expected decline in the short-term interest rate, coupled with an insufficient drop in the two-year yield—makes the tail end of hiking cycles appealing for issuing FRNs on an *ex ante* basis.

To better understand the time variation in the expected saving, it is helpful to note that

the expected saving is very well approximated by:

$$E_t^Q\left[\overline{r^{\text{ref}}}\right] - E_t^P\left[\overline{r^{\text{ref}}}\right], \tag{26}$$

a premium on the average reference rate itself, which is closely related to the bond risk premia investors demand in the Treasury markets. Bond risk premia are well documented to vary with the business cycle and can be positive or negative (Dai and Singleton, 2002; Cochrane and Piazzesi, 2005). More specifically, whenever a rise in the average 13-week Treasury bill rate over the two-year horizon is perceived as bad news, the risk-neutral expectation exceeds the physical one, resulting in a positive saving (before taking into account the illiquidity discount of FRNs).

The approximation can be formalized as follows. To first order,[14]

$$y_t^{\text{par}}(T) \approx E_t^Q\left[\overline{r^{\text{ref}}}\right] + \tfrac{365}{360}\,\theta_t^{\text{ns}}(T). \tag{27}$$

Substituting (27) into (25) leaves (26). The markers in Figure 8 plot the approximation in (26) against the saving (25) itself. They sit almost on top of one another: over January 2014 through December 2025 the two differ by -0.61 basis points on average, with a standard deviation of 1.36 and a largest single-week gap of 3.90.

6.3 Alternative Measures of Expected Savings: Survey Data and Real Time Estimation

Of the three terms in (25), the expected average reference rate $E_t^P\left[\overline{r^{\text{ref}}}\right]$ comes with a certain amount of statistical uncertainty. Whereas the other two terms, y_t^{par} and θ^{ns}, are risk-neutral

[14] To see this, note that a bond with par yield y_t^{par} is priced at par, and so is an FRN with reference rate r_t^{ref}, spread θ_t^{ns}, and the same cash flow dates and liquidity characteristics as the par coupon bond. The coupon cash flows of the two securities are discounted by the same curve. Denote by d_j the stochastic discount factor from t to the jth cash flow date, for $j \in \{1, 2, ..., J\}$, so that $E_t^Q[d_j]$ is the price of a zero-coupon bond maturing on that date. Because both securities are priced at par, we have

$$y_t^{\text{par}}(T)\left(\sum_j E_t^Q[d_j]\right) = \tfrac{365}{360}\,\theta_t^{\text{ns}}(T)\left(\sum_j E_t^Q[d_j]\right) + \sum_j E_t^Q(r_j^{\text{ref}} d_j).$$

The last term can be further decomposed: $\sum_j E_t^Q(r_j^{\text{ref}} d_j) = E_t^Q\left[\overline{r^{\text{ref}}}\right]\left(\sum_j E_t^Q[d_j]\right) + \sum_j E_t^Q\left[\left(r_j^{\text{ref}} - E_t^Q\left[\overline{r^{\text{ref}}}\right]\right) d_j\right]$. Dropping the second order term, $\sum_j E_t^Q\left[\left(r_j^{\text{ref}} - E_t^Q\left[\overline{r^{\text{ref}}}\right]\right) d_j\right]$, and then dividing both sides by $\sum_j E_t^Q[d_j]$, we obtain (27).

quantities, strongly identified from the cross section of prices on each date, $E_t^P\left[\overline{r^{\text{ref}}}\right]$ is a time-series forecast of where short rates go over the life of the security. We measure $E_t^P\left[\overline{r^{\text{ref}}}\right]$ based on the estimated model so far. This subsection recomputes the saving under two alternative measures. The first replaces the model's forecast with a Blue Chip survey forecast of the three-month Treasury yield in the two to five quarters ahead. While the Blue Chip forecast window is shorter than the two-year FRN's life,[15] it nonetheless provides a useful comparison. The second alternative measure keeps the model but re-estimates it at each date on data available only up to that date, so that the forecast uses no information the Treasury would not have had at the time.

Table 5 compares the savings under the three measures: model-based measure estimated in the full sample (panel A), survey-based measure (panel B), and model-based measure estimated in real time (panel C). Because the first FRN was issued in January 2014, the real-time model estimation is only available two years later after we observe the complete life cycle of the first FRN. Thus, we also report the results for 2016–2025, in addition to 2014–2025.

The survey-based expected saving is the lowest of the three. In panel A and panel B, the par yield, the frictionless par spread θ_t^{ns} and the liquidity component are all tightly identified from the cross section of securities prices and are held at their estimated values. Thus, the only difference is in $E_t^P\left[\overline{r^{\text{ref}}}\right]$. Here, the survey-based forecast of short-term interest rate exceeds the model-based forecast by about 12 bps between 2014 and 2025, cutting the estimated saving from 15.9 bps to 3.9 bps before accounting for the illiquidity discount. Net of the liquidity component, the Blue Chip surveys suggest that FRNs impose an additional cost of 10.4 bps per year. Again, the primary reason of this difference is that the surveys predicted a higher short-term interest rate than the model does.

So, which forecast is closer to the realized short-term interest rates in the future? This question can be answered directly in the data. Across the 289 monthly surveys from December 1999 to December 2023, the model-implied two-year average of the 13-week bill rate was 1.74%, against 2.09% by the survey. The realized two-year average of the 13-week bill rate over the following two years was 1.67%. The model therefore has little bias, +6 bps, whereas the surveys over-predict the short-term interest rate by 42 bps. The model's forecast RMSE of 0.83 percentage point is also lower than the survey's 1.21 percentage points. The larger

[15] Blue Chip reaches five quarters ahead in most months and six in March, June, September and December. We take horizons two through five throughout, so that the averaging window is the same in every month; on the dates where the sixth horizon is available, including it shifts the sample mean by 1.5 bps, and individual observations by 8 bps in absolute value on average.

errors in survey-based interest rate forecast are consistent with the literature. Cieslak (2018) documents large and persistent errors in investors' short-rate expectations over the business cycle. In particular, investors overestimate future short rates during economic downturns and Federal Reserve easing. In our sample, the survey forecast is 79 bps too high on average before 2014 and 11 bps too low on average after 2014, consistent with the cyclical pattern Cieslak (2018) documents.

The model estimated in real time produces virtually identical par yield and non-special spread because, as mentioned before, these quantities are strongly identified in the cross section. The primary difference is in $E_t^P\left[\overline{r^{\text{ref}}}\right]$. Between 2016 and 2025, the model estimated in real time expects the short rate to be about 15 bps lower than the model estimated based on the full sample (199.4 bps versus 214.8 bps). This difference is intuitive: the short-term interest rate rose sharply in the last five years of the sample, so the model estimated using the full sample has the benefit of "hindsight." As a result, the model estimated in real time predicts a higher expected saving.

Table 5: Decomposition of average expected savings on a newly-issued two-year FRN of current design, indexed to the 13-week bill, under three expectation measures of the average reference rate: from the full-sample model (Panel A), the Blue Chip consensus forecast of the three-month bill yield averaged over two to five quarters ahead (Panel B), and the real-time model (Panel C). The real-time model starts in January 2016. Averages are over the monthly Blue Chip survey dates in each sample period. All numbers are in basis points, except the two $\% > 0$ columns. To facilitate comparison, we use the same par yield, θ^{ns}, and liquidity component, implied by the full-sample model, across all three panels.

					Before liquidity			**Net of liquidity**	
Sample	N	Par yield	$E[\overline{r^{\text{ref}}}]$	θ^{ns}	Saving	$\% > 0$	Liquidity component	Net saving	$\% > 0$
Panel A: full sample model expectation									
2014–2025	144	202.0	189.3	−3.2	15.9	55.6	14.3	1.6	40.3
2016–2025	120	230.6	214.8	−2.4	18.2	55.0	14.0	4.3	45.8
Panel B: Blue Chip survey expectation									
2014–2025	144	202.0	201.3	−3.2	3.9	50.7	14.3	−10.4	31.2
2016–2025	120	230.6	228.5	−2.4	4.5	51.7	14.0	−9.4	30.8
Panel C: real-time model expectation									
2016–2025	120	230.6	199.4	−2.4	33.7	80.0	14.0	19.7	65.0

Further insight is gained by examining the time series behavior of expected savings. Figure 9 plots both measures over time, against the liquidity component (left axis), and

the Fed Funds target (right axis), as in Figure 8. The full-sample estimated savings is also plotted as comparison (dashed line). The real-time estimated saving follows the monetary policy cycle as the full-sample estimated saving does, leading to expected savings ranging from about −130 bps in May 2022 (shortly after the beginning of the hiking cycle) to about +200 bps in May 2023 (the tail of the hiking cycle). By contrast, the survey-based saving moves in a narrow range around zero and is unrelated to the savings later realized: across the 40 matured issues its correlation with the realized saving is −0.17, against +0.44 for the full-sample model; across the 32 issues for which real-time forecasts are available (from January 2016), the real-time model's correlation is +0.37. A closer inspection reveals that the real-time estimated saving responds faster to the rise and fall of the Fed Funds rate, and reaches higher peaks and lower troughs, compared to the full-sample estimated saving. That pattern, too, is expected, because the real-time model revises its estimate of the long-term mean of short rates as new data points are observed.

Perhaps the ultimate test is how the estimated savings under the three measures compare to the realized savings. Figure 10 illustrates exactly that comparison. For every original FRN issue that matured inside the sample, it shows the saving that actually materialized in the following two years (gray bars), together with estimated savings under the full sample model, real time model, and survey, each taken from the last model date, real-time vintage, or survey before the auction. Because we need two years of post-issuance data to measure the realized savings, this comparison runs from January 2014 to December 2023, with real-time model estimate available starting in January 2016. All three expected saving measures hold the auction spread and the par yield from the non-special curve fixed, so the only thing that differs across them is the forecast of the average index.

In addition, for each FRN, the realized savings and the various components, including the auction spread, actual average short rate over the life of the FRN, and the par yield at the time of the auction, are reported in Table 6.

Figure 10 and Table 6 reveal the low-frequency swing of realized savings and model-implied expected savings across the Federal Reserve's interest rate cycle. Across the 40 matured issues, the realized saving is positive on 18. The first eight FRN issues from January 2014 to October 2015, when the federal funds rate was held steady, all incurred interest expenses below those of two-year fixed coupon Treasury notes, but the real-time model did not have enough data to make a prediction during this period, and the survey-based estimated saving is mostly negative for them. The next eight issues, from January 2016 to October 2017, all incurred interest expenses exceeding those of two-year fixed coupon

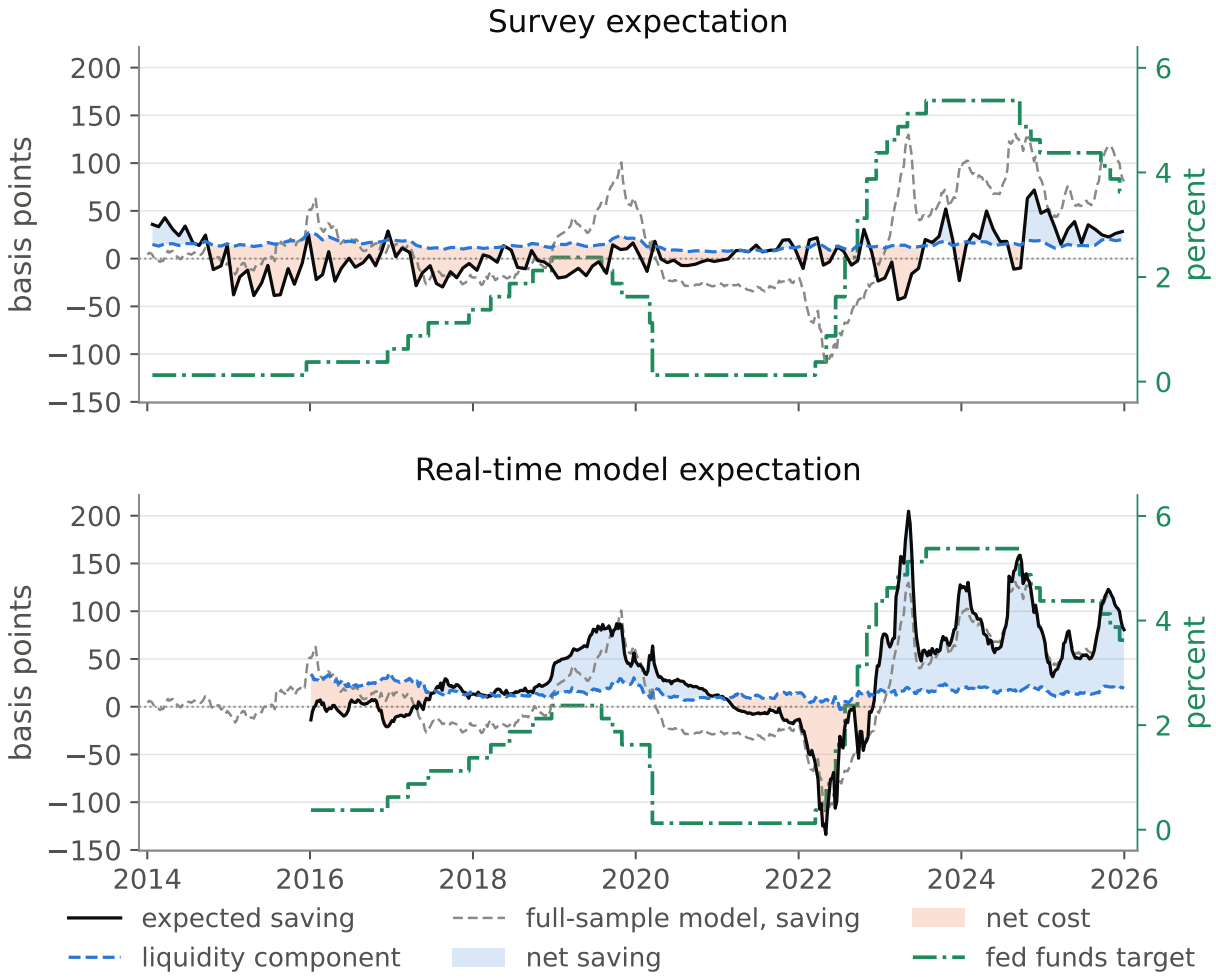


Figure 9: Expected saving and liquidity component for a newly issued two-year FRN of the current design, under two alternative expectations of the reference rate. Upper panel: the model's $E_t^P[\overline{r^{\text{ref}}}]$ is replaced by the Blue Chip consensus forecast of the three-month Treasury yield, averaged over two to five quarters ahead, on monthly survey dates, 2014–2025. Lower panel: saving and liquidity component from the three-curve model re-estimated each week on data available at that date, 2016–2025. Solid line: expected saving before liquidity. Blue dashed line: liquidity component. The shading between them is the net saving or net cost. The faint dashed line is the full-sample saving before liquidity. Dash-dotted line (right axis): midpoint of the FOMC's target range for the federal funds rate.

Treasury notes. It is unsurprising, given that the Fed started the hiking cycle in December 2015. During this period, the survey is negative throughout; the real-time model is negative except for the July and October 2017 issues, and the full-sample model except for the January and April 2016 issues.

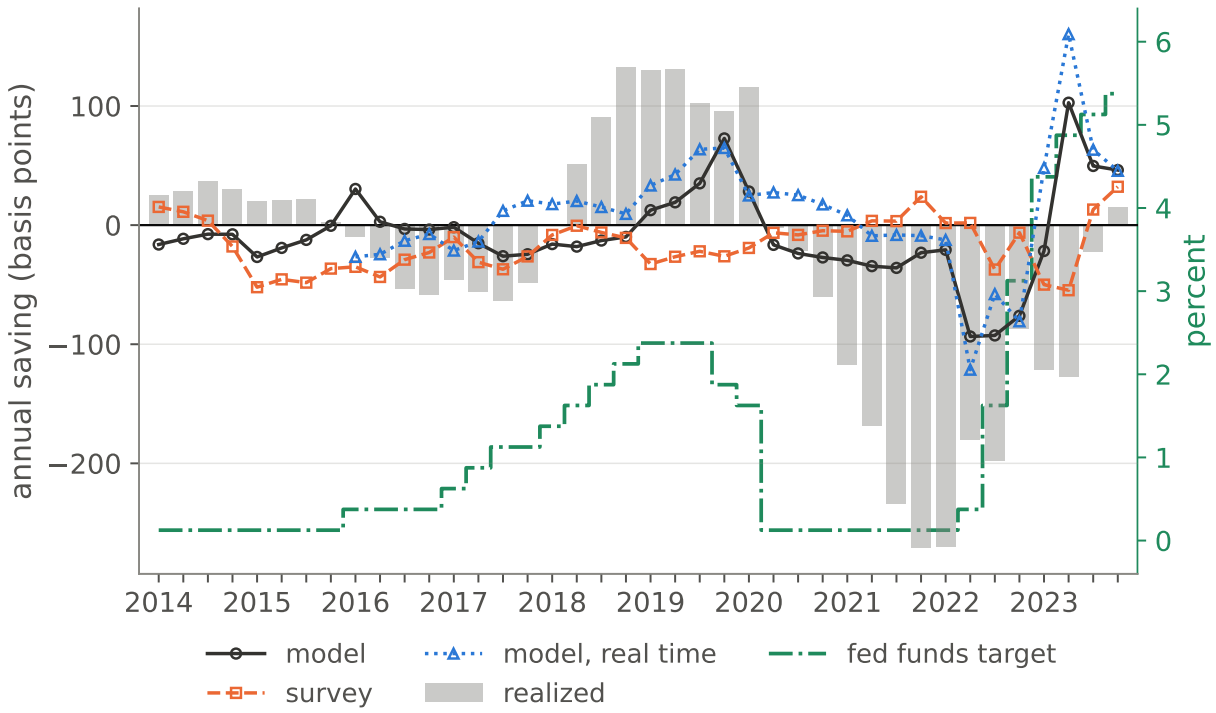


Figure 10: Realized and anticipated annual saving for each two-year FRN auctioned and matured between 2014 and 2025. The realized saving uses the average of the index Treasury published over the life of each FRN. The anticipated savings replace that average with three forecasts made on the last date before each auction: $E_t^P[\overline{r^{\text{ref}}}]$ from the full-sample model, the Blue Chip consensus, and $E_t^P[\overline{r^{\text{ref}}}]$ from real-time models. All four use the auction spread and the two-year par yield from the KR non-special curve. Real-time forecasts begin with the January 2016 auction. The dash-dotted line (right axis) is the midpoint of the FOMC's target range for the federal funds rate on each auction date.

The next nine issues, from January 2018 to January 2020, all have positive realized savings. In this time window, the Fed hiked the Fed Funds target rate from 1.5% to 2.5% and then cut it to 1.75%. The estimated saving from the real-time model stays positive throughout this period, but it climbed up sharply only after the Fed paused hiking. The survey-based estimate is negative in this period throughout. This episode again confirms the intuition that the tail end of hiking cycles coincides with high risk premium and hence high expected savings.

The rest of the FRN issues from April 2020 to October 2023 all happened after COVID. Because of the spike in post-COVID inflation and the resulting rate hikes by the Fed, the realized savings for these FRNs are overwhelmingly negative, reaching a trough of about

Table 6: Realized annual savings, in basis points, on the forty original two-year FRN issues auctioned from January 2014 to October 2023, all of which matured within the 2014–2025 sample. The auction spread θ is the spread set at each issue's original auction, as reported by TreasuryDirect. The average 13-week T-bill rate $\bar{b}$ is the realized average of each security's index over its life, from Treasury's published index determination rates. The two-year par yield, assuming quarterly coupon payments, is implied by the non-special curve of Section 3.2. The saving is computed from $y^{\text{par}} - \frac{365}{360}(\bar{b} + \theta)$.

Auction	Auction spread	2-year par yield	Avg. 13-wk T-bill	Annual saving	Auction	Auction spread	2-year par yield	Avg. 13-wk T-bill	Annual saving
Jan 2014	4.5	35.3	5.3	25.3	Jan 2019	11.5	256.8	113.9	129.7
Apr 2014	6.9	43.4	8.4	28.0	Apr 2019	13.9	231.3	85.0	131.1
Jul 2014	7.0	55.8	11.6	37.0	Jul 2019	22.0	182.6	57.1	102.4
Oct 2014	5.3	50.5	15.3	29.6	Oct 2019	30.0	160.9	34.2	95.8
Jan 2015	8.4	49.7	21.2	19.7	Jan 2020	15.4	147.5	15.7	115.9
Apr 2015	7.4	58.3	29.3	21.1	Apr 2020	11.4	22.6	12.9	−2.0
Jul 2015	7.7	71.4	41.5	21.6	Jul 2020	5.5	15.0	30.5	−21.5
Oct 2015	16.8	74.0	54.1	2.1	Oct 2020	5.5	15.0	68.7	−60.3
Jan 2016	27.2	87.0	68.5	−10.0	Jan 2021	4.9	12.1	122.6	−117.1
Apr 2016	19.0	78.0	85.2	−27.6	Apr 2021	3.4	18.8	181.0	−168.2
Jul 2016	17.4	71.6	105.8	−53.2	Jul 2021	2.9	19.7	246.9	−233.6
Oct 2016	17.0	89.4	128.7	−58.3	Oct 2021	3.5	50.8	314.1	−271.2
Jan 2017	14.0	122.9	152.6	−46.0	Jan 2022	−1.5	113.7	380.2	−270.2
Apr 2017	7.0	127.1	173.6	−56.0	Apr 2022	−7.5	257.3	438.7	−179.9
Jul 2017	6.0	134.9	189.4	−63.2	Jul 2022	3.7	298.7	486.3	−198.1
Oct 2017	4.8	159.2	199.7	−48.2	Oct 2022	14.0	442.3	508.1	−87.1
Jan 2018	0.0	206.1	202.4	0.9	Jan 2023	20.0	414.6	508.2	−121.0
Apr 2018	3.3	248.1	190.7	51.3	Apr 2023	16.9	397.8	501.2	−127.5
Jul 2018	4.3	265.9	168.4	90.8	Jul 2023	12.5	486.0	489.0	−22.5
Oct 2018	4.5	281.6	142.8	132.2	Oct 2023	17.0	510.8	471.7	15.3

−270 bps. The real-time model correctly predicts the sign of realized savings for about half of these FRNs. The misses are for issues shortly after COVID (for not foreseeing rate hikes) and in 2023 (for failing to foresee how sustained the hikes are). The full-sample model, with the benefit of seeing all the data, correctly predicts negative savings for issuing FRNs shortly after COVID, but it also misses the negative savings for FRNs issued in 2023. Interestingly, survey-based estimates hover around zero for about eight quarters before drifting into negative territory.

A forecast is the most useful if it informs issuance decisions in real time. Table 7 reports the average realized savings for the real-time model rule and the (real-time) survey rule, under which the issuer chooses a floating-rate note before the auction when its forecast says

the expected saving is positive and a fixed-coupon note otherwise. The simple rule of always issuing floating and (the impossible) perfect foresight are also provided for comparison.

Table 7: Issuance rules scored on realized savings relative to always issuing a two-year fixed-rate par note. The survey and real-time model rules issue an FRN if and only if the forecast saving is positive; they differ only in the expected average reference rate, from the Blue Chip survey or the real-time model. Both use the spread estimated by the real-time model. Forecasts and spreads are taken from the last date strictly before the auction: one to seven days earlier for the model, two to thirty-three days for the survey. Realized savings use the actual auction spread. Perfect foresight issues an FRN exactly when the realized saving is positive. The hit rate is the share of auctions on which the forecast and realized savings have the same sign. The sample is the 32 auctions from January 2016, when the real-time model starts, to October 2023, the last issue to mature by the end of 2025.

	Mean realized saving (bp)	Auctions floated	Hit rate (%)
Always floating	−43.0	32	
Real-time survey rule	−17.0	5	59
Real-time model rule	8.8	19	72
Perfect foresight	27.0	10	100

The most salient result from this comparison is that the real-time model rule saves interest expenses, but the real-time survey rule does not. (The full-sample model also produced positive savings, but it would not be available in real time and hence not reported.) If the Treasury used the real-time model to choose between FRNs and two-year fixed rate coupon securities between 2016 and 2023, it would have saved 8.8 bps per year on average across all 32 FRN issues (or 14.8 bps per year on the 19 FRNs that the real-time model predicts to have positive savings). This saving is about a third of the hypothetical savings under perfect foresight (27 bps, averaged over 32 FRN issues), which is again a strong indication of the slow-moving and predictable nature of interest rate risk premium. In our sample, always issuing floating regardless of market conditions ended up costing the Treasury 43 bps annually, largely due to the significant rise of interest rates post-COVID.

7 Alternative FRN Designs

Before launching the FRN program, Treasury requested comments from the market about how the security should be designed to minimize the cost of financing. The two particular design parameters are the choice of the index and the maturity.

On maturity, commenters split. Mizuho recommended that "initial offerings of Treasury FRNs include maturities of one to five years," predicting that "final maturities inside three years are expected to see higher investor demand resulting in a lower liquidity premium" (Rizzieri, 2012). HSBC took the other side: "maturities longer than two years would be suboptimal for Treasury floaters" (HSBC Global Banking and Markets, 2012).

On the index, multiple commenters point out that an index that is higher on average should imply a narrower spread, so the choice of index may not matter for Treasury interest cost. For instance, HSBC commented: "Investors price FRNs with different indices to achieve an expected yield target, so an index with a lower yield would result in a wider spread" (HSBC Global Banking and Markets, 2012). RBS expected that "investors may require a higher spread to account for the lower nominal yields versus other reference rates" (O'Donnell and Briggs, 2012). According to BMO, the fact that the general collateral rate had exceeded the 13-week bill rate on average "should not matter" (Graham, 2013). Nomura said the same of a repo index, that a compressing auction spread meant "not necessarily increasing the Treasury's borrowing cost at all" (Wolfthal, 2013).

To the best of our knowledge, Goldman Sachs's response provided the most detailed quantitative assessment. In particular, it estimated the spreads at which a Treasury FRN would trade against each candidate index and at each of three maturities: over the 13-week bill, the FRN spread at two-, three-, and five-year maturities would be 10.5, 11.5 and 14.05 bps; over the GCF repo index, the respective spreads would be 3.5, 5.25 and 10.25 bps (Hammack, 2013).[16] Goldman Sachs added: "the US Treasury must determine which of these floating rate issuances will command the largest premium, and which rate is most likely to stay low or trend lower over time."

Now, equipped with twelve years of FRN data and a model that jointly characterizes FRN pricing and Treasury yield curves, we are ready to revisit the FRN design question. We ask: Are there alternative designs of FRN that could reduce the interest expenses of the Treasury? We explore five choices of maturities, one, two, three, four and five years, as well as five candidate indices, the 13-week T-bill, the 4-week T-bill, model-implied hypothetical 1-week and overnight T-bill rates, and the overnight GC repo rate.

Combining equations (16) and (24), we can write the expected savings as

$$\text{Saving}_t(T) = y_t^{\text{par}}(T) - E_t^P\left(\overline{r^{\text{ref}}}\right) - \frac{365}{360}\theta_t^{\text{ns}} - \text{LiquidityComponent}_t(T), \tag{28}$$

[16] From January 2014 to December 2025, the FRN spread set at issuance auctions averaged 11.1 bps (see Table 1), very close to Goldman's estimate of 10.5 in 2013.

Here, the T-year par yield $y_t^{\text{par}}(T)$ is determined by the non-special curve at t and is invariant to the choice of index, $E_t^P\left(\overline{r^{\text{ref}}}\right)$ obviously depends on the index, θ_t^{ns} is derived from the T-bill curve and the non-special curve, and the liquidity component is derived from the three-curve model.

Table 8 reports the savings, as well as the various components, across the design grid. Panels A and B cover the FRN era under the full-sample model, and Panels C and D the models estimated in real time over 2016–2025. Within each block, the first panel (A and C) varies the FRN's maturity (fixing the 13-week T-bill index) and the second panel (B and D) varies the index (fixing two-year maturity). Figure 11 plots the net expected savings for the five maturities and four T-bill indices.

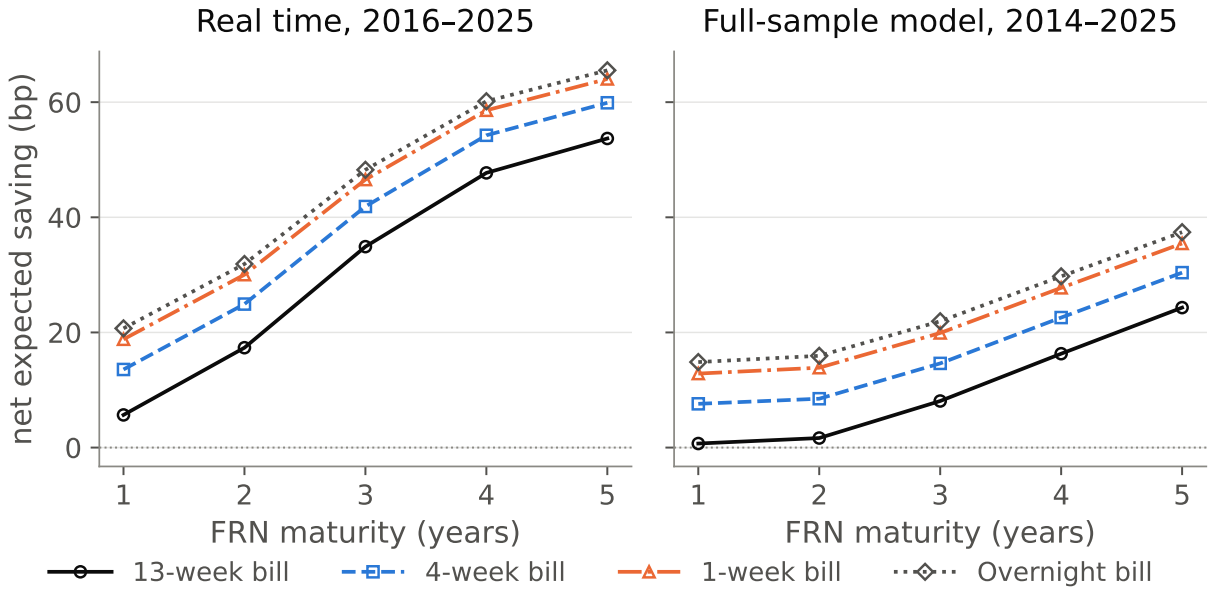


Figure 11: Expected annual saving net of the liquidity component, by FRN maturity and T-bill index tenor. Left: real-time models, re-estimated each week on data available at that date, averaged over 2016–2025 (Panels C and D of Table 8). Right: full-sample model, averaged over 2014–2025 (Panels A and B).

7.1 Choice of FRN maturity

How does the choice of FRN maturity affect the expected savings in using Treasury FRNs? Four observations emerge from Panels A and C. First, regardless of whether the model is estimated using the full sample or real time, the non-special spread is stable across maturity

Table 8: Average expected annual savings, in basis points, from issuing an FRN in place of a fixed-coupon par note of the same maturity. Averages are taken over weekly dates. The saving before liquidity is the par yield less the expected average reference rate, $E^P[\overline{r^{\text{ref}}}]$, and the frictionless par spread, θ^{ns}. The liquidity component is the full three-curve model's par spread less θ^{ns}. The net saving is the saving before liquidity less the liquidity component. Both spreads are scaled by 365/360 to match the par yield's 365-day year. Panels A and B use the full-sample model over 2014–2025. Panels C and D use the real-time models over 2016–2025. Panels A and C vary the FRN maturity with the 13-week bill as the reference rate; Panels B and D vary the reference rate for a two-year FRN. The overnight GC repo rate is treated as the overnight yield on the non-special curve.

	Par yield	$E^P[\overline{r^{\text{ref}}}]$	θ^{ns}	Saving before liquidity	Liquidity component	Net saving
Full-sample model: 2014–2025						
Panel A: by FRN maturity, 13-week bill reference						
1-year	194.1	189.6	−2.3	6.8	6.1	0.7
2-year	200.7	188.1	−3.1	15.7	14.0	1.7
3-year	210.0	183.6	−3.7	30.1	22.0	8.1
4-year	221.4	179.3	−4.2	46.3	30.0	16.3
5-year	234.0	176.3	−4.6	62.4	38.0	24.3
Panel B: by reference rate, two-year FRN						
13-week bill	200.7	188.1	−3.1	15.7	14.0	1.7
4-week bill	200.7	178.9	−0.8	22.5	14.0	8.5
1-week bill	200.7	172.6	0.1	27.9	14.0	13.9
Overnight bill	200.7	170.3	0.4	30.0	14.0	15.9
Overnight GC repo	200.7	195.3	0.9	4.4	14.0	−9.6
Real-time models: 2016–2025						
Panel C: by FRN maturity, 13-week bill reference						
1-year	228.3	219.3	−2.0	11.1	5.4	5.7
2-year	227.9	197.3	−2.7	33.3	16.0	17.3
3-year	230.8	171.8	−3.5	62.6	27.7	34.9
4-year	237.7	154.1	−4.2	87.8	40.1	47.7
5-year	247.3	145.5	−4.7	106.7	53.0	53.7
Panel D: by reference rate, two-year FRN						
13-week bill	227.9	197.3	−2.7	33.3	16.0	17.3
4-week bill	227.9	188.2	−1.2	40.9	16.0	24.9
1-week bill	227.9	182.4	−0.6	46.0	16.0	30.0
Overnight bill	227.9	180.4	−0.4	47.9	16.0	31.9
Overnight GC repo	227.9	209.0	1.0	17.9	16.0	1.9

choice and is between -1 and -5 bps. This is unsurprising, as the non-special spread is determined by the slope of the yield curve at very short maturities. Specifically, interest on FRNs accrues daily, but the index is the 13-week T-bill. Since the yield curve is upward sloping most of the time, the 13-week T-bill rate tends to "overcompensate" investors, implying a negative spread, without accounting for the liquidity component. Second, the liquidity component of FRNs increases monotonically in maturity. In the full-sample model, the liquidity component increases from 6.1 bps at one-year maturity to 38.0 bps at five-year maturity. In the real-time model, the corresponding figures are 5.4 bps at one year and 53.0 bps at five years. Here, the model plays a critical role in inferring the liquidity component at maturities beyond two years, which is not observed directly in the data. This fact can explain why the two model-implied liquidity components diverge by about 15 bps at the five-year maturity.

The third observation is that the expected saving before liquidity—which is equal to par yield of maturity T minus expected short rate over a T-year horizon, minus the non-special spread, as in equation (25)—also increases in maturity. This saving amount is approximated by $E^Q\left[\overline{r^{\text{ref}}}\right] - E^P\left[\overline{r^{\text{ref}}}\right]$, which is a measure of interest rate risk premium over the relevant horizon. On average, interest rate risk premium is positive and increases in maturity; thus, the expected saving, before deducting the liquidity component, also increases in maturity. Finally, interest rate risk premium increases faster in maturity than the liquidity component does, leading to expected savings that increase in the maturity of FRNs. In the real-time model, the model-implied net saving per year is 17.3 bps at two-year maturity, 34.9 bps at three-year maturity, and 53.7 bps at five-year maturity.

7.2 Choice of index

Now turning to the choice of the index, shown in Panels B and D of Table 8 and Figure 11. Here, the FRN maturity is fixed at two years. As a result, the par yield is invariant across the choices of the index, and so is the liquidity component. As discussed before, the non-special spread θ^{ns} is small in magnitude, between -3 and $+1$ bps. The most salient result here is that a shorter-maturity T-bill implies larger savings. In the real-time model, savings before liquidity are 33.3 bps for the 13-week T-bill, 40.9 bps for the 4-week T-bill, 46.0 bps for a hypothetical 1-week T-bill, and 47.9 bps for a hypothetical overnight T-bill. Because this saving is approximated by the difference $E^Q\left[\overline{r^{\text{ref}}}\right] - E^P\left[\overline{r^{\text{ref}}}\right]$, the result reveals that the interest rate risk premium is higher for shorter-maturity T-bills, even within a horizon as short as three months. In other words, while bond investors are generally averse to higher

interest rates, they are particularly averse to higher shorter-term interest rates. The model predicts that the Treasury can save an additional 7.6 bps by moving from the 13-week T-bill to the 4-week T-bill as the index underlying FRNs.

The results on GC repo rate are somewhat surprising: model-implied savings under the GC repo rate as the index are significantly lower than those under T-bills as indices. Compared to FRNs whose cash flows are based on 13-week T-bills, FRNs based on overnight GC repo rate are 11.3 bps more expensive to issue in the full-sample model and 15.5 bps more expensive in the real-time model. If the index is the (hypothetical) overnight T-bill, the overnight GC repo rate implies 25.5 bps and 30.0 bps lower savings in the full-sample model and real-time model, respectively, even though both indices are overnight. Why?

To see the intuition, recall from Figure 1 that very short term T-bill rates settle below the GC repo rate, which anchors the very short end of the non-special curve. Moreover, the difference between the T-bill curve and the non-special curve is the most pronounced at overnight maturity, but the two curves converge at one-year maturity. Thus, to a first-order approximation, $E^Q\left[\overline{r^{\text{T-bill}}}\right] \approx E^Q\left[\overline{r^{\text{GCF}}}\right]$, i.e., the risk-neutral expectations of the average non-special interest rate and the average T-bill interest rate at the one-year horizon coincide. But because the actual expected T-bill rate is lower due to money-like liquidity premium, the Treasury can harvest the T-bill liquidity premium by issuing FRNs linked to T-bills.

7.3 SOFR-linked FRNs?

The lackluster performance of a hypothetical FRN indexed to GCF repo rate should not be interpreted as negative verdict on FRNs linked to alternative repo rates, in particular the secured overnight financing rate, or SOFR. SOFR is a trimmed volume-weighted median of overnight repo rates, including triparty repos (some of which are GCF repos), centrally cleared bilateral repos, and (whenever data are available) noncentrally cleared bilateral repos. Starting in April 2018, SOFR has been published by the Federal Reserve Bank of New York. Since then, SOFR gradually replaced LIBOR as the primary floating rate index in various segments of fixed income markets in the United States, ending with an automatic transition of all "tough legacy" LIBOR-based contracts to SOFR-based contracts in June 2023.

We believe SOFR-linked Treasury FRNs merit serious consideration for a different reason: the deep liquidity of SOFR markets and the ease of hedging and "arbitrage." An FRN that pays SOFR cash flows, say daily compounded SOFR in arrears, can be hedged almost perfectly with a maturity-matching interest rate swap whose floating leg is daily compounded SOFR in arrears. If this hedging transaction faces minimal friction, issuing an N-year SOFR-

linked FRN is equivalent to issuing a fixed-coupon bond whose par yield is the N-year SOFR swap rate. The fact that SOFR swap rates currently trade below the par yield of Treasury coupon securities across maturities—known as negative swap spreads—suggests that the Treasury can harvest part, if not all, of such mispricing by issuing SOFR-based FRNs.[17] Between 2018 and 2025, the two-year SOFR swap rate typically trades below the two-year non-special par yield, and the average difference is about 10 bps. In the best case scenario, the Treasury would be able to issue two-year SOFR FRNs with zero spread and immediately enter payer interest rate swaps (receive SOFR, pay fixed); this pair of transactions would reduce the cost of issuing fixed-coupon securities by the size of the negative swap spread, which is 10 bps for the two-year maturity in the last few years. The more negative is the swap spread, the larger is the saving by issuing SOFR-based FRNs. To the extent that arbitrage is imperfect, the Treasury would harvest a fraction of the negative swap spread. This arbitrage channel, which does not involve market timing, should be distinguished from the channel based on (selectively) avoiding interest rate risk premium. Put another way, there is no reason why the Treasury should not take advantage of both channels.

SOFR-linked FRNs already exist in the market, as analyzed by Klingler and Syrstad (2025). Since Fannie Mae's first issue in July 2018, government-sponsored enterprises, large banks, and multinational corporations have issued SOFR-linked FRNs. Between July 2018 and December 2021, SOFR-linked FRNs issued by three GSEs in their sample—Federal Home Loan Banks, Federal Farm Credit Banks, and Farmer Mac—have average spread over SOFR of 6.9 bps, 8.5 bps, and 12.4 bps, with average maturities of 1.27, 1.75, and 2.86 years, respectively (Klingler and Syrstad, 2025). More recent examples of SOFR-linked FRNs include a one-year issue by Federal Farm Credit Banks in September 2024 with a spread of 6.5 bps and an 11-month issue by Farmer Mac in March 2025 with a spread of 3.5 bps.[18]

[17] The negative swap spread is commonly attributed to demand for receiving fixed in swaps from asset managers, pension funds, and other institutional investors seeking duration. Hedge funds and other arbitrageurs take the other side by paying fixed on SOFR swaps, buying fixed-coupon Treasury securities, and financing the cash Treasury positions in the repo market (Klingler and Sundaresan, 2019; Aquilina et al., 2024).

[18] See https://www.farmcreditfunding.com/ffcb_live/termsheet/3133ERSA0.pdf and https://www.farmermac.com/hubfs/FarmerMac-u4m/Pdfs/pricing-supplement-search/31424WWK4_Jefferies-LLC.pdf.

8 Conclusion

Can the U.S. Treasury reduce its borrowing costs by relying more on floating-rate notes? From the perspective of avoiding an elevated interest rate risk premium, the answer is a qualified yes. The interest rate risk premium varies over time, but slowly, and our dynamic model of interest rates provides signals early enough to make the idea actionable. Based on the experience with two-year FRNs issued since 2014, the benefit of avoiding a heightened interest rate risk premium outweighs the cost of the illiquidity discount on these securities. Admittedly, taking an active view on the level of the interest rate risk premium carries inherent risks, but, in our view, the rapid growth of Treasury debt raises the weight on the level of borrowing costs relative to their variability.

Our results also suggest that shorter-maturity T-bills carry larger liquidity premia and therefore make better reference rates for FRNs from the perspective of borrowing costs. Reducing the illiquidity discount of FRNs would be another productive direction. For example, adding a one-year maturity to the FRN program would better meet the demand of money market mutual funds and further improve the liquidity of this market. Longer-maturity FRNs could be added at the same time or sequentially to avoid a larger interest rate risk premium.

We believe SOFR-linked FRNs deserve a serious separate study. In addition to market-value stability similar to that of T-bill-linked FRNs, SOFR-linked FRNs would allow the Treasury to tap the deep liquidity of SOFR derivatives, such as interest rate swaps. The Treasury's borrowing costs would be further reduced if SOFR-linked FRNs allowed it to harvest part of the negative swap spread.

Appendices

A Pricing Derivations

This appendix collects the computational details underlying the pricing formulas in Section 4. All derivations apply to a generic affine short rate $\tilde{r}_t = \tilde{r}_\infty + \tilde{\delta}_Z' Z_t$ with state dynamics (2). It is convenient here to write the short rates (4)–(6) against the full state rather than against its blocks, so that the three stacked loading vectors are

$$\delta_Z = \begin{bmatrix} \mathbf{1}_3 \\ 0 \\ 0 \end{bmatrix}, \qquad \delta_Z^b = \begin{bmatrix} \delta_X^b \\ 1 \\ 0 \end{bmatrix}, \qquad \delta_Z^f = \begin{bmatrix} \delta_X^f \\ 0 \\ 1 \end{bmatrix}. \tag{29}$$

The generic pair $(\tilde{r}_\infty, \tilde{\delta}_Z)$ is then specialized to one of the three curves: non-special (r_∞, δ_Z), T-bill (r_∞^b, δ_Z^b), or FRN discount (r_∞^f, δ_Z^f). To distinguish solutions under different curves, we use superscripts a^j, b^j when needed.

A.1 Exponential-Affine Expectation

For any vector γ,

$$E_t^Q\left[e^{-\int_t^T \tilde{r}_u\, du}\, e^{\gamma' Z_T}\right] = e^{-a(T-t,\gamma) - b(T-t,\gamma)' Z_t}, \tag{30}$$

where a and b satisfy the Riccati equations

$$\dot{a}(\tau,\gamma) = \tilde{r}_\infty - \tfrac{1}{2}\, b(\tau,\gamma)' \Sigma\Sigma' b(\tau,\gamma), \tag{31}$$

$$\dot{b}(\tau,\gamma) = \tilde{\delta}_Z - \kappa' b(\tau,\gamma), \tag{32}$$

with boundary conditions $a(0,\gamma) = 0$ and $b(0,\gamma) = -\gamma$.

A.2 Discretization and Precomputed Loadings

Discretizing (31)–(32) at interval Δ and letting $a_n \equiv a(n\Delta,\gamma)$, $b_n \equiv b(n\Delta,\gamma)$, $\rho \equiv I - \kappa'\Delta$:

$$a_n - a_{n-1} = \tilde{r}_\infty \Delta - \tfrac{1}{2}\, b_{n-1}' \Sigma\Sigma' b_{n-1}\, \Delta, \tag{33}$$

$$b_n = \tilde{\delta}_Z \Delta + \rho\, b_{n-1}. \tag{34}$$

Expanding the recursion for b_n and defining $\phi_n \equiv (I + \rho + \rho^2 + \cdots + \rho^{n-1})\Delta$:

$$b_n = \phi_n \tilde{\delta}_Z - \rho^n \gamma. \tag{35}$$

Substituting (35) into (33) and summing:

$$a_n = n\Delta\,\tilde{r}_\infty - \tfrac{1}{2}\,\gamma' C_n \gamma + \tilde{\delta}_Z' D_n \gamma - \tfrac{1}{2}\,\tilde{\delta}_Z' E_n \tilde{\delta}_Z, \tag{36}$$

where

$$C_n = \sum_{i=0}^{n-1} (\rho^i)' \Sigma\Sigma' \rho^i\,\Delta, \tag{37}$$

$$D_n = \sum_{i=0}^{n-1} \phi_i' \Sigma\Sigma' \rho^i\,\Delta, \tag{38}$$

$$E_n = \sum_{i=0}^{n-1} \phi_i' \Sigma\Sigma' \phi_i\,\Delta. \tag{39}$$

The matrices ρ^n, ϕ_n, C_n, D_n, E_n depend only on (κ, Σ, Δ)—not on $(\tilde{r}_\infty, \tilde{\delta}_Z, \gamma)$—and are therefore computed once and shared across all three pricing curves.

A.3 Yield Loadings

The price of a zero-coupon bond with $\tau = n\Delta$ years to maturity is $e^{-\tau(A_n + B_n' Z_t)}$, where setting $\gamma = 0$ in (36)–(35) gives:

$$A_n = \tilde{r}_\infty - \frac{1}{2n\Delta}\,\tilde{\delta}_Z' E_n \tilde{\delta}_Z, \tag{40}$$

$$B_n = \frac{1}{n\Delta}\,\phi_n \tilde{\delta}_Z. \tag{41}$$

A.4 T-bill Price Loadings

A 13-week bill is auctioned on a Monday and issued on the Thursday of the same week, a three-day lag, and matures 91 days after issue. We apply this convention on every date: the buyer pays P_t on the issue date (day $t+3\Delta$) and receives \$1 at maturity (day $t+94\Delta$, where $\Delta = 1/365$). Federal holidays displace the cycle on a minority of dates, moving the auction to a Tuesday or the issue to a Friday. Across the 626 auctions held between January 2014 and December 2025 the lag from auction to issue is three days in 84.7% of cases, two days in 12.1% and four days in 3.2%, so the convention is exact for most auctions and misaligned by one calendar day for the remainder. Equating discounted cash flows under the T-bill curve:

$$P_t = \frac{\exp\!\big(-a^b(94\Delta,\,0) - b^b(94\Delta,\,0)' Z_t\big)}{\exp\!\big(-a^b(3\Delta,\,0) - b^b(3\Delta,\,0)' Z_t\big)}. \tag{42}$$

Comparing with $P_t = \exp(-\alpha^b - \beta^{b\prime} Z_t)$:

$$\alpha^b = a^b(94\Delta,\ 0) - a^b(3\Delta,\ 0), \tag{43}$$

$$\beta^b = b^b(94\Delta,\ 0) - b^b(3\Delta,\ 0). \tag{44}$$

A.5 Three-Time-Index Expectation

The key pricing object (12) involves the FRN curve in the discount rate and the T-bill curve in the coupon. With coupon date $T = t + k\Delta$ and auction date $s = t + n\Delta$, applying the law of iterated expectations:

$$\begin{aligned}
&E_t^Q\left[e^{-\int_t^{t+k\Delta} r_u^f\, du}\, e^{\beta^{b\prime} Z_{t+n\Delta}}\right] \\
&\quad = E_t^Q\left[e^{-\int_t^{t+n\Delta} r_u^f\, du}\, e^{\beta^{b\prime} Z_{t+n\Delta}}\, E_{t+n\Delta}^Q\left[e^{-\int_{t+n\Delta}^{t+k\Delta} r_u^f\, du}\right]\right] \\
&\quad = e^{-a^f((k-n)\Delta,\,0)}\, E_t^Q\left[e^{-\int_t^{t+n\Delta} r_u^f\, du}\, e^{(\beta^b - b^f((k-n)\Delta,0))' Z_{t+n\Delta}}\right] \\
&\quad = e^{-a^f((k-n)\Delta,\,0) - a^f(n\Delta,\,\gamma^*) - b^f(n\Delta,\,\gamma^*)' Z_t},
\end{aligned} \tag{45}$$

where

$$\gamma^* \equiv \beta^b - b^f((k-n)\Delta,\,0) = \beta^b - \phi_{k-n}\delta_Z^f, \tag{46}$$

and all functions a^f, b^f are evaluated using the FRN discount curve parameters (r_∞^f, δ_Z^f). All quantities on the right-hand side of (45) are computable from the precomputed loadings ϕ_n, ρ^n, C_n, D_n, E_n via (36)–(35).

References

Matteo Aquilina, Andreas Schrimpf, Vladyslav Sushko, and Dora Xia. Negative interest rate swap spreads signal pressure in government debt absorption. *BIS Quarterly Review*, December 2024. Box B.

Michael D. Bauer and Glenn D. Rudebusch. Interest rates under falling stars. *American Economic Review*, 110(5):1316–1354, 2020.

Terry Belton, Kristopher Dawsey, David Greenlaw, Huachen Li, Srini Ramaswamy, and Brian Sack. Optimizing the maturity structure of U.S. Treasury debt: A model-based framework. Hutchins Center Working Paper No. 46, Brookings Institution, 2018.

Anna Cieslak. Short-rate expectations and unexpected returns in Treasury bonds. *Review of Financial Studies*, 31(9):3265–3306, 2018. doi: 10.1093/rfs/hhy051.

Marco Cipriani and Gabriele La Spada. Investors' appetite for money-like assets: The MMF industry after the 2014 regulatory reform. *Journal of Financial Economics*, 140(1):250–269, 2021.

John H. Cochrane and Monika Piazzesi. Bond risk premia. *American Economic Review*, 95 (1):138–160, 2005. doi: 10.1257/0002828053828581.

Ezechiel Copic, Luis Gonzalez, Caitlin Gorback, Blake Gwinn, and Ernst Schaumburg. Introduction to the Floating-Rate Note Treasury security. Liberty Street Economics, Federal Reserve Bank of New York, April 21, 2014, April 2014. URL https://libertystreeteconomics.newyorkfed.org/2014/04/introduction-to-the-floating-rate-note-treasury-security/.

Qiang Dai and Kenneth J. Singleton. Expectation puzzles, time-varying risk premia, and affine models of the term structure. *Journal of Financial Economics*, 63(3):415–441, 2002. doi: 10.1016/S0304-405X(02)00067-3.

Eugene F. Fama and Robert R. Bliss. The information in long-maturity forward rates. *American Economic Review*, 77(4):680–692, 1987.

Fidelity Investments (FMR Co.). Comment letter on advance notice of proposed rulemaking, docket no. BPD-2012-0002. Letter from Scott C. Goebel to Matthew S. Rutherford, January 23, 2013, available at TreasuryDirect, January 2013.

Damir Filipović, Markus Pelger, and Ye Ye. Stripping the discount curve — a robust machine learning approach. *Management Science*, 2025. doi: 10.1287/mnsc.2023.01401. Articles in Advance, published online December 22, 2025.

Matthias Fleckenstein and Francis A. Longstaff. The US Treasury floating rate note puzzle: Is there a premium for mark-to-market stability? *Journal of Financial Economics*, 137 (3):637–658, 2020. doi: 10.1016/j.jfineco.2020.04.006.

Matthias Fleckenstein and Francis A. Longstaff. Treasury richness. *Journal of Finance*, 79 (4):2797–2844, 2024. doi: 10.1111/jofi.13371.

Ljubica Georgievska, Stefano Pegoraro, and Anthony Saunders. The U.S. Treasury's biggest short: Duration in the shadows. Working paper, 2026.

Scott J. Graham. Comment letter on advance notice of proposed rulemaking, docket no. BPD-2012-0002. BMO Capital Markets Corp., letter to Matthew S. Rutherford, January 10, 2013, available at TreasuryDirect, January 2013.

Robin Greenwood and Dimitri Vayanos. Bond supply and excess bond returns. *Review of Financial Studies*, 27(3):663–713, 2014. doi: 10.1093/rfs/hht133.

Robin Greenwood, Samuel G. Hanson, and Jeremy C. Stein. A comparative-advantage approach to government debt maturity. *Journal of Finance*, 70(4):1683–1722, 2015. doi: 10.1111/jofi.12253.

James D. Hamilton and Jing Cynthia Wu. The effectiveness of alternative monetary policy tools in a zero lower bound environment. *Journal of Money, Credit and Banking*, 44(s1): 3–46, 2012. doi: 10.1111/j.1538-4616.2011.00477.x.

Beth Hammack. Comment letter on advance notice of proposed rulemaking, docket no. BPD-2012-0002. Goldman, Sachs & Co., letter to the Department of the Treasury, January 22, 2013, available at TreasuryDirect, January 2013.

Samuel G. Hanson and Jeremy C. Stein. Monetary policy and long-term real rates. *Journal of Financial Economics*, 115:429–448, 2015.

Samuel G. Hanson, David O. Lucca, and Jonathan H. Wright. Rate-amplifying demand and the excess sensitivity of long-term rates. *Quarterly Journal of Economics*, 136(3): 1719–1781, 2021.

Jonathan S. Hartley and Urban J. Jermann. The pricing of U.S. Treasury floating rate notes. *Journal of Financial Economics*, 155:103833, 2024. doi: 10.1016/j.jfineco.2024.103833.

Zhiguo He, Stefan Nagel, and Zhaogang Song. Treasury inconvenience yields during the COVID-19 crisis. *Journal of Financial Economics*, 143(1):57–79, 2022. doi: 10.1016/j.jfineco.2021.05.044.

HSBC Global Banking and Markets. Treasury floating rate notes. Comment letter responding to Treasury's Request for Information on Treasury Floating Rate Notes (Docket BPD-2012-001), April 2012, available at TreasuryDirect, 2012.

Scott Joslin, Kenneth J. Singleton, and Haoxiang Zhu. A new perspective on Gaussian dynamic term structure models. *Review of Financial Studies*, 24(3):926–970, 2011. doi: 10.1093/rfs/hhq128.

Scott Joslin, Anh Le, and Kenneth J. Singleton. Gaussian macro-finance term structure models with lags. *Journal of Financial Econometrics*, 11(4):581–609, 2013.

Rohan Kekre, Moritz Lenel, and Federico Mainardi. Monetary policy, segmentation, and the term structure. *American Economic Review*, forthcoming. doi: 10.1257/aer.20240491.

Sven Klingler and Suresh Sundaresan. An explanation of negative swap spreads: Demand for duration from underfunded pension plans. *Journal of Finance*, 74(2):675–710, 2019.

Sven Klingler and Olav Syrstad. The SOFR discount. *Journal of Financial Economics*, 164: 103989, 2025. doi: 10.1016/j.jfineco.2024.103989.

Arvind Krishnamurthy and Annette Vissing-Jorgensen. The aggregate demand for treasury debt. *Journal of Political Economy*, 120(2):233–267, 2012. doi: 10.1086/666526.

Canlin Li and Min Wei. Term structure modeling with supply factors and the Federal Reserve's large-scale asset purchase programs. *International Journal of Central Banking*, 9(1):3–39, 2013.

Robert B. Litterman and José Scheinkman. Common factors affecting bond returns. *The Journal of Fixed Income*, 1(1):54–61, 1991. doi: 10.3905/jfi.1991.692347.

Priya Misra. Public comment on Treasury Floating Rate Notes. Bank of America Merrill Lynch, comment letter on Docket BPD-2012-001, April 18, 2012, available at TreasuryDirect, April 2012.

Stefan Nagel. The liquidity premium of near-money assets. *Quarterly Journal of Economics*, 131(4):1927–1971, 2016. doi: 10.1093/qje/qjw028.

Nomura Securities International, Inc. Response to request for information on treasury floating rate notes. Comment letter signed by Brad Wolfthal responding to Treasury's Request for Information on Treasury Floating Rate Notes (Docket BPD-2012-001), April 18, 2012, available at TreasuryDirect, 2012.

William O'Donnell and John Briggs. Public comment on Treasury Floating Rate Notes. RBS Securities Inc., comment letter on Docket BPD-2012-001, April 18, 2012, available at TreasuryDirect, April 2012.

Gerald Rizzieri. Public comment on Treasury Floating Rate Notes. Mizuho Securities USA Inc., comment letter on Docket BPD-2012-001, April 11, 2012, available at TreasuryDirect, April 2012.

Treasury Borrowing Advisory Committee. TBAC charge presentation to the U.S. Treasury: What adjustments to debt issuance, if any, should Treasury make in consideration of its financing needs in the short-, medium-, and long-term? Discussion charts, August 2, 2011, slide "Cost and Benefits of Extending Average Maturity", August 2011. URL https://home.treasury.gov/system/files/276/TBAC-discussion-charts-Aug-2011.pdf.

Treasury Borrowing Advisory Committee. TBAC charge #2 presentation to the U.S. Treasury. Discussion charts, August 2012, p. 12, August 2012. URL https://home.treasury.gov/system/files/276/TBAC-Discussion-Charts-August-2012.pdf.

U.S. Department of the Treasury. Sale and issue of marketable book-entry Treasury bills, notes, and bonds. Advance Notice of Proposed Rulemaking, Federal Register 77 FR 72278 (31 C.F.R. Part 356, Doc. 2012-29307, Docket BPD-2012-0002), December 5, 2012, December 2012.

U.S. Department of the Treasury. Sale and issue of marketable book-entry Treasury bills, notes, and bonds. Final Rule, Federal Register 78 FR 46426 (31 C.F.R. Part 356, Doc. 2013-18178, Docket Fiscal-BPD-2013-0001), July 31, 2013, July 2013.

U.S. Department of the Treasury, Bureau of the Fiscal Service. Treasury Floating Rate Note Term Sheet. TreasuryDirect, Auction Regulations (UOC), 2013.

U.S. Government Accountability Office. Debt management: Floating Rate Notes can help Treasury meet borrowing goals, but additional actions are needed to help manage risk. Technical report, GAO-14-535, June 2014. URL https://www.gao.gov/products/gao-14-535.

U.S. Securities and Exchange Commission. Money market fund reform; Amendments to Form PF. Release No. 33-9616; Federal Register 79 FR 47736, August 14, 2014, July 2014.

Dimitri Vayanos and Jean-Luc Vila. A preferred-habitat model of the term structure of interest rates. *Econometrica*, 89(1):77–112, 2021. doi: 10.3982/ECTA17440.

Brad Wolfthal. Comment letter on advance notice of proposed rulemaking, docket no. BPD-2012-0002. Nomura Securities International, Inc., letter to the Department of the Treasury, January 22, 2013, available at TreasuryDirect, January 2013.