D'Amico, Stefania; Pancost, N. Aaron

Working Paper

Special repo rates and the cross-section of bond prices: The role of the special collateral risk premium

Working Paper, No. 2018-21

Provided in Cooperation with:
Federal Reserve Bank of Chicago

Suggested Citation: D'Amico, Stefania; Pancost, N. Aaron (2018) : Special repo rates and the cross-section of bond prices: The role of the special collateral risk premium, Working Paper, No. 2018-21, Federal Reserve Bank of Chicago, Chicago, IL, https://doi.org/10.21033/wp-2018-21

This Version is available at:
https://hdl.handle.net/10419/200591

Federal Reserve Bank of Chicago

Special Repo Rates and the Cross-Section of Bond Prices: the Role of the Special Collateral Risk Premium

Stefania D'Amico and N. Aaron Pancost

December 3, 2018

WP 2018-21

https://doi.org/10.21033/wp-2018-21

Special Repo Rates and the Cross-Section of Bond Prices: the Role of the Special Collateral Risk Premium*

Stefania D'Amico† and N. Aaron Pancost‡

This version: December 3, 2018

Abstract

We estimate the joint term-structure of U.S. Treasury cash and repo rates using daily prices of all outstanding Treasury securities and corresponding special collateral (SC) repo rates. This allows us to derive a risk premium associated to the SC value of Treasuries and quantitatively link this premium to various price anomalies, such as the on-the-run premium. We show that a time-varying SC risk premium can explain between 74%–90% of the on-the-run premium, and is highly correlated with a number of other Treasury market anomalies. This suggests a commonality across these price anomalies, explicitly linked to the SC value of the highest-quality securities—recently-issued U.S. nominal Treasuries.

*We thank Gadi Barlevy, Luca Benzoni, Darrell Duffie, Jean-Sébastien Fontaine, John Griffin, Arvind Krishnamurthy, Gregor Matvos, Daniel Neuhann, Nivine Richie (discussant), Sam Schulhofer-Wohl, Krista Schwarz (discussant), Andrea Vedolin, Jan Wrampelmeyer (discussant), and seminar participants at the McCombs Finance Brownbag, the 2017 European Winter Meeting of the Econometric Society, the fifth International Conference on Sovereign Bonds at the Bank of Canada, the Federal Reserve Bank of Chicago, the 2018 Annual Meeting of the Central Bank Research Association, the University of North Carolina at Chapel Hill Kenan-Flagler Finance Brownbag, the 2018 Southern Finance Association Annual Meeting, and the 2018 Term Structure Workshop at the Deutsche Bundesbank for helpful comments and conversations. All remaining errors are our own. Tim Seida provided excellent research assistance. The views expressed here do not reflect official positions of the Federal Reserve.

†Federal Reserve Bank of Chicago. Contact: sdamico@frbchi.org

‡University of Texas at Austin McCombs School of Business. Contact: aaron.pancost@mccombs.utexas.edu

1 Introduction

Previous studies have documented various price anomalies in very large and liquid U.S. fixed-income markets. For example, assets with almost identical payoffs can carry significantly different prices: Just-issued (on-the-run) Treasury bonds trade at higher prices than previously issued (off-the-run) bonds with the same remaining maturity (Krishnamurthy, 2002); prices of nominal Treasury bonds can consistently exceed those of inflation-swapped TIPS with exactly matching cash flows (Fleckenstein, Longstaff and Lustig, 2014); and off-the-run Treasury notes can trade at higher prices than off-the-run Treasury bonds with similar duration (Musto, Nini and Schwarz, 2017 and Pancost, 2017). We conjecture that there could be a commonality across these anomalies: a missing risk premium on Treasuries with special collateral value in the repurchase agreement (repo) market.

Holders of nominal Treasury securities that go "on special" in repo markets can effectively receive a dividend by borrowing at below market rates, using those Treasuries as collateral (Duffie 1996). A rational investor, in pricing a Treasury security, would account for the current and future expected dividends that could be obtained by lending that specific security in the repo market. This stream of dividends is uncertain, as the Treasury collateral value fluctuates unexpectedly depending on demand and supply dynamics in the repo market, and thus commands a compensation for risk. Models that ignore this special collateral (SC) risk premium might interpret a spread between actual and model-implied prices entirely as anomalous, though in fact most of this spread may be justified as compensation for exposure to SC risk.

Our premise is that Treasury market participants price not only expected repo "dividends," but also their associated risk premia. Thus, we develop a dynamic term structure model (DTSM) that, by explicitly accounting for observed SC repo risk factors, can identify and estimate SC risk premia separately from other classical risk premia, (i.e., those related to the level, slope, and curvature risk factors). Subsequent to fitting the model, we study the properties of SC risk premia with the goal of understanding their economic importance, their ability to explain relevant price anomalies, and to serve as an indicator of financial fragility connected to the scarcity of public safe and liquid assets.

In our DTSM, individual Treasury securities with the highest potential to go on special are priced by discounting their cash flows with a specialness-risk-adjusted short-rate process. This approach implies a security-specific discounting that effectively links the pricing in the Treasury cash and repo markets. Seminal work by Duffie (1996) and Vayanos and Weill (2008) as well as related empirical evidence (e.g., Jordan and Jordan, 1997; D'Amico, Fan

and Kitsul, 2017) have suggested that pricing in these two markets are tightly linked.[1] So far, DTSMs in the literature have been modeling and estimating U.S. Treasury cash and repo rates separately, as the common approach implicitly ignores the possibility that investors might be discounting the stream of future cash-flows of certain Treasury securities at a rate lower than the generic short rate.[2]

Most likely, this omission in the term-structure literature is due to both the lack of data on SC repo rates and the complexity of pricing each Treasury security individually within a DTSM. Using daily prices of all outstanding Treasury securities and a proprietary dataset on individual Treasury SC repo rates from 2009 to 2017, we estimate the joint term-structure of U.S. Treasury cash and repo rates, derive a risk premium associated to the SC value of Treasuries, and quantitatively link this premium to various price anomalies. In an economy in which investors often are willing to *pay* up to 3% to *lend* cash overnight to access certain securities, and fails to deliver even seasoned Treasuries are significant, the differences across SC repo rates could have important consequences for cash prices.

An additional advantage of our approach is that, by using a large cross-section of bond prices, we can estimate a SC risk premium without estimating the physical dynamics of the model's latent factors. This is because, with so many more bonds than factors, we can estimate latent factors consistently cross-section by cross-section without worrying about the time-series parameters of the model, as in the SR approach of Andreasen and Christensen (2015). This represents a substantial benefit of using actual Treasury bond data rather than a few smoothed zero-coupon yields, as our sample length is limited by the specific-collateral data.

We find that the SC risk premium is large in magnitude at the beginning of the sample, particularly at the 10-year maturity, as those issues are the "most special" in the repo markets. It fluctuates significantly over time, exhibiting sharp variations following supply shocks, approximated by the FOMC announcements of the Federal Reserve's asset purchase programs, while remaining practically unchanged following demand shocks, approximated by Treasury auction announcements characterized by abnormal bid-to-cover ratios.[3] Importantly, at the 10-year maturity, the SC risk premium account for 83% of the on-the-run premium, consistent with the theory of Vayanos and Weill (2008). In particular, we show that our model matches the size, persistence, and variability of the on-the-run premium by al-

[1] This is nicely illustrated by the formulas in Buraschi and Menini (2002). See also Cherian, Jacquier and Jarrow (2004) and Graveline and McBrady (2011).

[2] The rare exceptions are in the swap spread literature: Grinblatt (2002) and Feldhutter and Lando (2008) approximate the Treasury convenience yield with the Libor-Tbill spread or the Libor-GC spread, respectively, which however is the same across all Treasury securities.

[3] See Cahill et al. (2013) for an in depth discussion of supply shocks, and Gorodnichenko and Ray (2018) for the demand shocks.

lowing for a time-varying SC risk premium that also correlates with the state of the economy. At the same maturity, the SC risk premium also explains 62% of the off-the-run note-bond spread, 68% of the TIPS-Treasury bond puzzle, and about 60% of measures of TIPS illiquidity relative to their nominal counterparts. This indeed suggests a common underlying economic mechanism across these anomalies, linked to the SC value of the highest-quality securities—recently-issued nominal U.S. Treasury bonds.

This paper is related to various studies in the literature on top of those that have motivated our search for a possible common driver of the already mentioned price anomalies. Duffie (1996) shows in a simple static setting how a security's special repo rate that is below the generic short rate of interest (i.e., the general collateral repo rate) implies a higher price for that security in the cash market. However, he does not examine whether observed on-the-run premiums are consistent with special repo rates within a dynamic setting and in the data. Our dynamic model nests Proposition 1 of Duffie (1996) in a setting that allows us to measure the risk premium on special repo rates—one of the contributions of this paper. Krishnamurthy (2002) notes that special spreads on the on-the-run 30-year Treasury bond are high, and that this is consistent with a price premium on the on-the-run bond. He quantifies the 30-year on-the-run premium and estimates the profits from trading on it, and finds that they are small. We go one step further by showing the substantial joint time variation in both the special spread and the on-the-run premium, and by quantifying the premium on the risk that makes the two consistent with each other.

The usual practice in the literature on the term structure of interest rates (see Pancost 2017 for a survey) is to exclude on-the-run bonds from the analysis. Gürkaynak, Sack and Wright (2007) exclude not just on-the-run bonds, but also the first off-the-run (i.e., those bonds that were on-the-run just before the latest bond was issued). Many empirical studies of the term structure of interest rates, including D'Amico, Kim and Wei (2010), Hamilton and Wu (2012), Kim and Orphanides (2012), and Bauer and Rudebusch (2014), use the estimated smooth yield curves from Gürkaynak, Sack and Wright (2007) as if they were data. Other studies, including Ang and Piazzesi (2003), Christensen, Diebold and Rudebusch (2011), Joslin, Singleton and Zhu (2011), and Creal and Wu (2016), use the bootstrap method of Fama and Bliss (1987) applied to only a small number of bonds, which generally do not include on-the-run bonds. Unlike these papers, we also consider the prices of on-the-run bonds directly and seek to explain why they are higher than other bonds with similar cash-flows using CUSIP-level special repo rates.

One of the goals of our paper is to measure the time-varying risk premium that can make prices in the repo and cash markets consistent with one another. This is distinct from the question of why an on-the-run premium exists in the first place. Duffie (1996) conjectures

that it may be difficult to find someone willing to trade an off-the-run bond, making the on-the-run bond more liquid and increasing its price. Vayanos and Weill (2008) formalize the intuition of Duffie (1996) in a model of search frictions in which more-liquid securities can trade at a premium even when their promised cash-flows are identical to less-liquid securities.[4] Pasquariello and Vega (2009) show that the liquidity differential between on- and off-the-run securities varies with the Treasury auction cycle, even after controlling for repo specialness. We document how repo specialness also varies with the auction cycle, and show that the inclusion of the SC risk premium, above and beyond the special repo rate, helps align on-the-run yield differentials with repo special spreads.

The rest of the paper is organized as follows. Section 2 describes the data. Section 3 sets up the model and describes identification and estimation. Section 4 presents our empirical results across alternative specifications, which help assess the contribution of time-varying risk premiums on the special repo factors. Section 5 offers concluding remarks.

2 Data and Motivation

We use daily prices of Treasury bonds from CRSP covering the period from January 2, 2009 to December 29, 2017.[5] We are limited to this sample period by the availability of our SC repo rate data. However, considering that we are interested in understanding how the collateral value of Treasury securities affect their prices in the cash market, this is a very interesting period, which also makes the estimation of the term structure of Treasury yields quite difficult. In particular, in those years, reduced issuance by the Treasury, the sharp increase in Treasury holdings by the Federal Reserve (Fed), and new financial regulation have reportedly shrunk the availability of Treasury securities, making this high-quality collateral quite scarce in the repo market. This, in turn, might have caused special spreads to be positive also for off-the-run securities and increased off-the-run fails to deliver to unusual levels (see for example Fleming and Keane, 2016).

Our data includes 496,420 price observations across 2,252 trading days and 628 unique CUSIPs; we drop all bonds with remaining time to maturity of less than one year. We supplement this data with CUSIP-level SC repo rates from a major electronic broker-dealer trading platform, and compute daily averages of SC repo rates quoted between 7:30 and 10

[4]Although the model of Vayanos and Weill (2008) is dynamic, they do not allow the on-the-run premium or the repo special spread to vary dynamically. Duffie, Gârleanu and Pedersen (2002) derive a dynamic model with search frictions and heterogenous beliefs that generates an endogenous on-the-run premium that declines over time as pessimists borrow the bond and find progressively less-optimistic traders to whom they short-sell it.

[5]Source: CRSP®, the Center for Research in Security Prices, Booth School of Business, The University of Chicago. Used with permission. All rights reserved. crsp.uchicago.edu

a.m. (Eastern time). This time window is chosen because trading in the repo market begins at about 7 a.m., remains active until about 10 a.m., and then becomes light until the market closes at 3 p.m. Repo transactions with SC are executed on a delivery versus payment basis (i.e., same-day settlement). In these transactions, a collateral security is delivered into a cash lender's account in exchange for funds. The exchange occurs via FedWire or a clearing bank. Finally, we use Treasury general collateral (GC) repo rates from the General Collateral Finance (GCF) Repo Index, which is a tri-party repo platform maintained by the Depository Trust & Clearing Corporation (DTCC).[6] This market is characterized as being primarily inter-dealer, although some commercial banks and Fannie Mae also participate. It is a fairly active market although its size is still small compared to that of the overall tri-party repo market.[7] Table 1 provides some descriptive statistics of our data.

[Table 1 about here.]

We highlight four features of our data from Table 1. First, on average there are over 200 bonds per cross-section, so that there is plenty of variation for identifying 3, 4, or 6 pricing factors. The fact that we observe many more bonds in the cross-section than the number of latent factors allows us to identify the latter without implementing a nonlinear Kalman filter; in particular, we need not estimate the physical dynamics of our model's factors in order to estimate our SC risk premium, as described in the next section. Second, although off-the-run bonds on average have much lower special spreads than on-the-run bonds, they are still very likely to go "on special" (i.e., they have a special spread > 0). The third column of Table 1 shows that on average across all maturities, almost 85% of all off-the-run bonds are trading special in our sample. Moreover, the percentage trading special is uniform across maturities at issuance; the maturity least likely to trade special, the 30-year, is still on special over 80% of the time.

Third, although all maturities trade special, the 10-year on-the-run bond is the "most special," in the sense that it has the highest special spread of any maturity. The fifth column of Table 1 reports the average special spread for on-the-run bonds; it is over 35 bps for the 10-year on-the-run bond, which is large given that the average GC repo *rate* in our sample is 30 bps. The difference implies that on average, holders of the 10-year on-the-run bond are *paid* about 5 bps to *borrow* cash using their special collateral. Other maturities trade special at lower spreads, about 20 bps on average. The 30-year on-the-run bond trades at a much lower average special spread of about 8 bps.

[6] DTCC GCF rate data are publicly available at http://www.dtcc.com/charts/dtcc-gcf-repo-index.aspx.

[7] For more detail about the GCF Repo Index see Fleming and Garbade (2003).

Fourth and finally, the higher special spreads on the 10-year bond correspond to higher on-the-run premia in the cash market. The sixth column of Table 1 reports, for on-the-run bonds of the indicated maturity, the average price residual (actual minus fitted) implied by a 3-factor model estimated using only off-the-run bonds and ignoring all special spreads.[8] In our sample, this premium is about 89 bps of par on average for the 10-year on-the-run bond; maturities besides 10 years tend to have much lower cash premia, closer to 17 bps of par. The 30-year bond trades on average only about 7 bps above what it "should" according to a simple 3-factor model. This finding for the 30-year on-the-run bond is unique to our sample period: before 2009 the 30-year on-the-run bond featured a large and time-varying on-the-run premium, often as large as the 10-year note. However, because we do not have special spreads before 2009, we omit this time period from our analysis in this paper.

[Figure 1 about here.]

Because our model is most conveniently formulated in terms of coupon bond prices, rather than yields (see equations 14 and 15 in the next section), in this paper we define the on-the-run premium in prices rather than in yields. This choice is only for convenience and does not materially affect the results: our model produces on-the-run yield spreads that mirror other measures of the on-the-run premium. This is clearly illustrated in the top panel of Figure 1, which plots our 10-year on-the-run premium in yields (red line) together with two other measures of the on-the-run premium commonly used in the literature. Both the blue and the red lines define the 10-year on-the-run premium by subtracting the yield on the 10-year on-the-run note from the yield on a synthetic bond with the exact same coupon structure and maturity. The red line estimates the synthetic bond's yield with our 3-factor model ignoring repo spreads, while the blue line uses the Sevensson yield curve estimated by Gürkaynak, Sack and Wright (2007). Both synthetic yields are constructed using parameters estimated only on off-the-run bonds. The GSW on-the-run premium has a correlation coefficient of 0.98 with the one implied by our 3-factor model.

While the two synthetic-bond on-the-run measures are very similar, a simpler "2-bond" measure looks completely different. The black line in the top panel of Figure 1 reports the on-the-run premium calculated as the difference between the yield on the first off-the-run 10-year bond and the yield on the on-the-run 10-year bond (Krishnamurthy, 2002). This 2-bond measure is model-free, but as noted by Gürkaynak, Sack and Wright (2007), it will tend to understate the on-the-run premium when the yield curve is upward sloping, because the first off-the-run bond usually has a lower duration than the on-the-run bond. This bias

[8] Gürkaynak, Sack and Wright (2007) omit on-the-run and first-off-the-run bonds from their yield curve estimation because these bonds usually trade at a premium.

can be so severe that the 2-bond measure even returns an on-the-run *discount*, albeit a small one, as in the top panel of Figure 1. In other time periods (not reported), the on-the-run bond usually has a lower yield than the first off-the-run (i.e., a positive on-the-run premium), but the premium is generally an order of magnitude smaller than the premium implied by the other two measures. The reason is that, as shown in Table 1, not only on-the-run securities but also off-the-run securities can have special collateral value in the repo market, and thus, at times, the first off-the-run, being the on-the-run closest substitute, can be at least as valuable as collateral. In addition, security-level data from the DTCC show that fails to deliver seasoned Treasury securities—defined as securities issued more than 180 days prior—were increasing in early 2011 through late 2012 from previously-negligible levels, the same period in which we observe the 2-bond premium fall into negative territory.[9]

[Figure 2 about here.]

The top panel of Figure 1 shows that the 10-year on-the-run premium is large; the bottom panel of Figure 1 shows that the 10-year bond also has a large and time-varying special spread. The figure plots the SC repo rates on 10-year on-the-run bonds over time (blue line), as well as the overnight GC repo rate (red line). Consistent with its large and positive cash premium, the 10-year SC repo rate is often substantially lower than the GC rate, often dropping as low as -3% (annualized).

Often, the SC repo rates in the bottom panel of Figure 1 tend to spike downwards at regular intervals. This "seasonality" is a feature of the Treasury auction cycle, which has three important periodic dates: the auction announcement date, the auction date, and the issuance date. There is usually about one week from the announcement to the auction. During a typical auction cycle, the supply of Treasury collateral available to the repo market is at its highest level when the security is issued, therefore the repo specialness spread should be close to zero. As time passes, more and more of the security may be purchased by holders who are not very active in the repo market, consequently the security's availability may decline over time and the repo specialness spread may increase. When forward trading in the next security begins on the auction announcement date, holders of short positions will usually roll out of the outstanding issue, implying that demand for that specific collateral should decrease and that the repo specialness spread will rapidly decline (Fisher 2002).

The top panel of Figure 2 illustrates the auction-cycle dynamics of repo special spreads by showing their averages as function of days since issuance for bonds maturing in 2, 3, 5, 7, 10, and 30 years. As in Keane (1995), specialness spreads on the securities that are auctioned monthly (2-, 3-, 5-, and 7-year notes) climb with the time since the last auction

[9]See Fleming and Keane (2014) and Fleming and Keane (2016).

until around the announcement of the next auction, after which they decline sharply. In contrast, the quarterly auction cycle of the 10-year note looks quite different, mainly because since 2009 the Treasury has introduced two regular re-openings following each 10-year note auction. Therefore, it is possible to observe three separate auction sub-cycles: the most dramatic run-up in specialness spread takes place before the first reopening, approximately 30 days after the auction; a second run-up, similar in shape but smaller in magnitude, immediately follows and peaks just before the second reopening, approximately 60 days after the auction; and finally, during the third sub-cycle the specialness spread is practically flat. This would suggest that the increased availability of the on-the-run security after each reopening strongly diminishes the impact of the seasonal demand for short positions around these dates (Sundaresan 1994).

The evidence of a recurrent auction cycle suggests that some time variation in special spreads is predictable and therefore is not itself a source of risk. Thus, when modeling the special repo risk factor we should carefully account for the deterministic component in special spreads. However, there is a good deal of unpredictable variation around the averages plotted in the top panel of Figure 2. The bottom panel of Figure 2 shows a scatterplot of the square root of the observed special spreads around the auction cycle component, estimated using a smoothing spline.[10] It is precisely the risk of the special spread varying significantly around the deterministic auction-cycle component that market participants ought to be pricing and that should generate a SC risk premium. This implies that the stochastic component of our special repo risk factor should be focused on capturing those deviations from the deterministic component and should have a market price of risk associated to it.

The goal of our model is to quantitatively link both the level and fluctuations of the risky special repo spread to the cash premium plotted in the top panel of Figure 1. Because the 10-year on-the-run note displays the largest premium in the cash market, and also has the largest special spread on average, we begin with a model specification that considers only the 10-year on-the-run and first off-the-run bonds as "special," that is, exposed to repo risk factors. Then, we extend the baseline model to incorporate special spreads and cash price premia for other maturities as well.

[10]We model the square root of the special spread over the auction cycle, rather than the level, to ensure that the special spread is weakly positive.

3 Model

We assume that the prices of Treasury bonds depend on a $k \times 1$ vector X_t that consists of both observable and unobservable (latent) factors, which evolve according to

$$X_{t+1} = \mu + \Phi X_t + \Sigma \varepsilon_{t+1} \tag{1}$$

where the vector of shocks ε_{t+1} are independent of X_t and of each other across time, and are normally distributed:

$$\varepsilon_{t+1} \sim \mathbb{N}\left(\vec{0}, I\right).$$

The generic short rate of interest is assumed to be an affine function of the factors:

$$\log\left(1 + R_t\right) = \delta_0 + \delta_1' X_t. \tag{2}$$

The stochastic discount factor (SDF) is given by

$$\log \frac{M_{t+1}}{M_t} = -\delta_0 - \delta_1' X_t - \frac{1}{2}\lambda_t'\lambda_t - \lambda_t'\varepsilon_{t+1}, \tag{3}$$

where

$$\lambda_t \equiv \lambda + \Lambda X_t.$$

3.1 Special Repo Spreads

A repo contract can be thought of as a collateralized loan, where the repo seller borrows at the repo rate in exchange for a Treasury bond, and regains the bond when she repays the loan plus interest at maturity. Duffie (1996) shows that SC repo rates can be below the GC repo rate, without creating an arbitrage opportunity, because the supply of special collateral is fixed.

A SC repo rate that is below the GC rate is a dividend that is proportional to the collateral's current price. Let $1 + R_t$ denote the current (gross) GC rate, and $1 + r_t^i$ the SC rate on bond i with current price P_t. Assume no haircut for simplicity. At time t, the owner of the special collateral borrows P_t against the collateral, at r_t^i, and simultaneously lends an amount Δ at the GC rate R_t. At time $t + 1$, she earns

$$\left(1 + R_t\right)\Delta - \left(1 + r_t^i\right) P_t$$

so that if $\Delta = \frac{1+r_t^i}{1+R_t} P_t$, she has no gain or loss at $t+1$, and earns

$$P_t - \Delta P_t = \left(1 - \frac{1+r_t^i}{1+R_t}\right) P_t$$

at time t.

In what follows, it will be convenient to parameterize the log gross special spread on bond i as

$$y_t^i \equiv \log \frac{1+R_t}{1+r_t^i} \geq 0, \tag{4}$$

which implies that the price of a zero-coupon bond that is on special and with special spread equal to y_t^i must have a price given by

$$\begin{aligned} P_t &= \left(1 - e^{-y_t^i}\right) P_t + E_t^* P_{t+1} \\ &= e^{y_t^i} E_t^* P_{t+1} \end{aligned}$$

where E_t^* is the risk-neutral expectation. Because SC rates are always weakly less than the GC rate, the "special dividend" $e^{y_t^i} \geq 1$ or, equivalently, $y_t^i \geq 0$.

In order to ensure that y_t^i is nonnegative for all values of the state vector X_t, we parameterize it as a quadratic form in the factors:

$$y_t^i = X_t' \Gamma_i X_t, \tag{5}$$

where Γ_i is a symmetric positive semidefinite matrix that picks out some i-specific factors from the state vector X_t. This modeling device is commonly used in DTSMs accounting for the zero lower bound on the short rate (e.g., Ahn, Dittmar and Gallant, 2002; Kim and Singleton, 2012), and leads to the following proposition for pricing zero-coupon bonds that are on special for their entire life. We show later how we allow bonds to be on special for only a limited (deterministic) time, for example while they are on-the-run and first off-the-run.

Proposition 1. *Consider a zero-coupon bond on special with τ periods to maturity, where the repo spread is given by equation (5). Then the zero-coupon bond price satisfies*

$$\log P_{t,\tau}^Z = A_\tau + B_\tau' X_t + X_t' C_\tau X_t \tag{6}$$

where the A_τ, B_τ, and C_τ loadings are given by

$$\begin{aligned} C_\tau &= \Gamma + \Phi^{*\prime} C_{\tau-1} D_{\tau-1} \Phi^* \\ B_\tau' &= -\delta_1' + \left(2\mu^{*\prime} C_{\tau-1} + B_{\tau-1}'\right) D_{\tau-1} \Phi^* \\ A_\tau &= -\delta_0 + A_{\tau-1} + \frac{1}{2} B_{\tau-1}' \Sigma G_{\tau-1} \Sigma' B_{\tau-1} \\ &\quad + \frac{1}{2} \log |G_{\tau-1}| + \left(\mu^{*\prime} C_{\tau-1} + B_{\tau-1}'\right) D_{\tau-1} \mu^* \end{aligned} \tag{7}$$

where

$$\begin{aligned} G_{\tau-1} &= \left[I - 2\Sigma' C_{\tau-1} \Sigma\right]^{-1} \\ D_{\tau-1} &= \Sigma G_{\tau-1} \Sigma^{-1}, \end{aligned}$$

$C_0 = \vec{0}_{k\times k}$, $B_0 = \vec{0}_{k\times 1}$, and $A_0 = 0$, and the risk-neutral parameters μ^* and Φ^* are given by

$$\begin{aligned} \mu^* &\equiv \mu - \Sigma\lambda \\ \Phi^* &\equiv \Phi - \Sigma\Lambda. \end{aligned} \tag{8}$$

Proof. See Appendix A. ∎

The loadings in equation (7) include the loadings of an affine term structure model as a special case when $\Gamma = \vec{0}$, since in this case $C_\tau = \vec{0}$ for all τ and therefore $G_\tau = D_\tau = I$ for all τ. Further, these loadings are usually obtained in quadratic-Gaussian term-structure models (e.g., Kim 2004; Breach, D'Amico and Orphanides 2016).

In order to incorporate the observed features of repo special spreads described in Section 2, we specify the individual repo special dividend y_t^i as the sum of three observable factors that apply only to bonds on special. In particular, we assume that the special spread on bond i at time t is given by

$$y_t^i = \left[y_{\tau(t,i)}^D + y_t^S + x_t^i\right]^2 \tag{9}$$

where y_τ^D is a deterministic auction-cycle process that depends on calendar time t only through $\tau(t,i)$, which is the number of days since issuance of bond i at time t, y_t^S is a stochastic aggregate repo factor applying to all special bonds, and x_t^i is a bond-specific residual. The square in equation (9) ensures that special spreads are always weakly positive.

It is important that we do not attribute all of the common movement in special spreads to our aggregate repo risk factor y_t^S, because some of its variation over time is predictable,

as shown in the top panel of Figure 2. For example, market participants are surely aware that special spreads on the 10-year note tend to be high on average just before the first reopening. Thus, this predictable variation in special spreads is not itself a source of risk and is captured by y_τ^D. Further, including this deterministic component in equation (9) also allows us to include more bonds, at different points in the auction cycle, into the aggregate repo factor y_t^S. For example, the first off-the-run bond (days 90 to 180 in the bottom panel of Figure 2) does not display much of an auction cycle. Thus including y_τ^D in the analysis helps us strip out the auction-cycle component and compare apples to apples.

However, as shown in the bottom panel of Figure 2, there are substantial deviations from the deterministic auction-cycle component, whose risk market participants ought to be pricing. Thus, we identify the aggregate repo risk factor y_t^S as the average deviation across all special securities:

$$y_t^S = \frac{1}{n_t} \sum_{i=1}^{n_t} \left[\sqrt{y_t^i} - y_{\tau(t,i)}^D \right]. \tag{10}$$

We then estimate x_t^i as a residual:

$$x_t^i = \sqrt{y_t^i} - y_t^S - y_{\tau(t,i)}^D. \tag{11}$$

Importantly, both equations (10) and (11) can be estimated directly on the special-spread data, independently of the term structure model. In this sense they are observable pricing factors, compared to the first three latent factors of the model.

We add the repo factors $y_{\tau(t,i)}^D$, y_t^S, and x_t^i to the state vector X_t and assume, along with the three latent factors, that the full vector X_t follows a VAR process as in equation (1). Then, in the basic case of a bond that is on special for its entire life, the matrix Γ from equation (7) is pinned down by

$$\Gamma = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & 0 & 1 & 1 & 1 \end{bmatrix} \tag{12}$$

where the three repo factors, two of them specific to bond i at time t, enter as the fourth, fifth, and sixth components of the state. Importantly, we allow the aggregate repo risk factor y_t^S to have unrestricted VAR dynamics and prices of risk. We assume that the idiosyncratic

special spread residuals x_t^i follow the process

$$x_{t+1}^i = \rho x_t^i + \sigma_x \varepsilon_{t+1}^i \tag{13}$$

where i indexes individual CUSIPs and ε_{t+1}^i is a standard normal random variable that is iid over time and independent of the aggregate VAR shocks ε_{t+1}. We assume x_t^i is unconditionally mean-zero in order to allow the average repo spreads to be governed by the auction cycle and the aggregate repo factor. Because the x_t^i factors are idiosyncratic, we assume that the dynamics of each x_t^i are identical under the physical and risk-neutral measures, so that the shocks ε_{t+1}^i do not enter the SDF or carry any risk prices.

Because the auction-cycle component y_τ^D does not depend directly on calendar time and is not random, we model its effect on bond prices by inserting it directly into the appropriate element of μ, and setting the corresponding rows of Φ and Σ to zero.[11]

Overall, this setup implies that for "non-special" bonds (i.e., second off-the-runs and older) the price loadings are given by equation (7) with $\Gamma = \vec{0}$ and initial conditions $A_0 = B_0 = C_0 = 0$. In other words, we price these bonds with a standard 3-factor Gaussian DTSM. In contrast, a zero-coupon bond that is on special until maturity would have price loadings again given by equation (7), with the same initial condition, but now with Γ given by equation (12). Thus the price of a special bond is driven by six factors rather than three, although the additional three factors are not latent.

However, as shown in Figure 2 and the fourth and fifth columns of Table 1, bonds do not remain on special for their entire life. Bonds are "more special," in the sense of a higher special spread, when they are on-the-run. We model this by noting that at issuance, the date at which the current on-the-run bond goes off-the-run is known with certainty. At that point, it will have a known remaining maturity m. Thus, from that time $\tau = m$ all the way to expiration ($\tau = 0$), it must have price loadings corresponding to an off-the-run bond with $\Gamma = \vec{0}$. For maturities from $\tau = m$ up to its issuance maturity, it has loadings given by A_τ^m, B_τ^m, and C_τ^m, which we compute using equations (7) and (12). In other words, instead of maturing and paying \$1 per unit of face-value, the initial condition for the on-the-run loadings is that at a known maturity m they become an off-the-run bond and inherit the off-the-run loadings.

[11] We could in addition allow σ_x^i to vary deterministically over the auction cycle as well, which we conjecture might further eliminate auction-cycle seasonality from our estimates.

3.2 Coupon-bearing bonds

So far, we have derived the price of zero-coupon bonds that can be on special for their entire life. However, special spreads typically accrue to coupon-bearing bonds that only trade special for a limited time. Coupon-bearing bonds are linear combinations of the zero-coupon bonds priced in Proposition 1, weighted by their coupon payments: the price of bond i at time t is given by

$$\begin{aligned} P_t^i &= \sum_j c_j \exp\left\{A_{\tau_j} + B'_{\tau_j} X_t + X_t' C_{\tau_j} X_t\right\} \\ &\equiv \sum_j c_j P^Z\left(\tau_j, X_t\right) \\ &\equiv \vec{P}^Z\left\{\vec{c}\left(i\right), \vec{\tau}\left(i\right), X_t\right\} \end{aligned} \tag{14}$$

where c_j denotes the size of the jth coupon payment, τ_j its time to maturity, and the last two lines define notation. Equation (14) includes the repo special spread at time t through the C_{τ_j} loadings (which all contain a Γ term). Details on how we identify the term structure of special spread loadings on actual coupon bonds are given in section 3.3 below.

Treasury bonds in our sample pay the same coupon amount every six months; accounting for these coupons, and pricing accrued interest, implies that the price of bond i is given by

$$P_t^i = P^Z\left(\tau_i, X_t\right) + \frac{c_i}{2 \times 100}\left[\sum_{j=1}^{N_{it}} P^Z\left(\tau_{ij}, X_t\right) - \frac{\tau_{i1}}{\xi_{it}} P^Z\left(\tau_{i1}, X_t\right)\right] \tag{15}$$

where N_{it} is the number of remaining coupon payments for bond i at time t, and ξ_{it} is the time between the next and previous coupon payment for bond i at time t.[12] The last term in equation (15) accounts for the accrued interest on bond i at time t, which is shared pro rata between the buyer and seller depending on the time remaining to the next coupon payment. Equation (15) describes how the coupon rate c and time to maturity τ of a given bond i in the data translate into the cash-flows $\vec{c}\left(i\right)$ and their maturities $\vec{\tau}\left(i\right)$ in equation (14).

[12]For the bonds that mature on February 29th, we assign their coupon payments in non-leap-years to February 28th.

Stacking all n_t bonds at time t yields the measurement equation

$$\vec{P}_t = \begin{bmatrix} \vec{P}^Z\left\{\vec{c}\,(1), \vec{\tau}\,(1), X_t\right\} \\ \vec{P}^Z\left\{\vec{c}\,(2), \vec{\tau}\,(2), X_t\right\} \\ \cdots \\ \vec{P}^Z\left\{\vec{c}\,(n_t), \vec{\tau}\,(n_t), X_t\right\} \end{bmatrix} + \vec{\eta}_t \tag{16}$$

which, along with equation (1), constitute the state-space system to be estimated. In practice, the number of bonds n_t in each cross-section is so large relative to the number of factors X_t that the latter can be estimated on each cross-section individually without regard to the state equation (1) (see Andreasen and Christensen 2015 for a proof); that is, our method is in practice very similar to a formal non-linear Kalman filter. However, unlike estimating a non-linear Kalman filter, our method does not require estimates of the physical measure parameters. This is important because our sample of repo special spreads is only about nine years long, and does not include enough business cycles to reliably estimate physical parameters.

Coupon payments complicate the computation of prices for special securities somewhat, in particular because those securities do not remain on special for their entire life. However, we are helped by the fact that securities go off-the-run at a fixed time since issuance, at which point their special spreads are close to zero (see Figure 2). Thus, it is reasonable to assume that second off-the-run and older securities have zero special spreads, which guarantees that their remaining maturity when they go "off special" is known in advance. We model these complications as follows: for each special bond i in our sample, we compute the price loadings of each coupon j and principal payment using the off-the-run loadings with $\Gamma = \vec{0}$ for maturities from 0 to $m\,(j)$, where $m\,(j)$ is the maturity of coupon payment j when the bond goes off special, and with Γ as in equation (12) from $m\,(j)$ up to the initial maturity τ_j *of that coupon (or principal) payment*. Thus, every coupon payment of a special bond is discounted using the special spread and the repo risk factor, but the sensitivity to that special repo factor varies across coupons according to their time to maturity.

An example will help illustrate the issue. Consider a 10-year bond that trades special while it is on-the-run and first off-the-run. This bond has 3650 days to maturity and is special for 180 days; assume, as is true in the data, that its first coupon payment is made just before it goes off special, at (say) 175 days after issuance.[13] For simplicity, assume that

[13]Because we include both on-the-runs and first off-the-runs as "special" in our benchmark model, bonds are special for roughly six months. In practice the first semi-annual coupon payment of each 10-year note arrives a few days before the next note is issued and the bond becomes a second off-the-run. Thus the first coupon payment of each on-the-run bond is indeed "special" its entire life.

this bond makes only three payments: the first coupon payment at 175 days, another coupon payment at 365 days, and a principal plus coupon payment at maturity. Denote the price loadings of each coupon payment j of this bond by $A_{\tau_j}^{m(j)}$, $B_{\tau_j}^{m(j)}$, and $C_{\tau_j}^{m(j)}$, where again $m(j)$ is the maturity of the j coupon when it goes off special, while A_τ, B_τ, and C_τ denote price loadings of non-special bonds. The latter triplet of loadings follow equation (7) but with $\Gamma = \vec{0}$. If the bond is issued at time t, its price is given by:

$$P_t = e^{y_t}\left[ce^{A_{175}^0+B_{175}^{0\prime}X_t+X_t^\prime C_{175}^0 X_t} + ce^{A_{365}^{185}+B_{365}^{185\prime}X_t+X_t^\prime C_{365}^{185}X_t} + (1+c)\,e^{A_{3650}^{3470}+B_{3650}^{3470\prime}X_t+X_t^\prime C_{3650}^{3470}X_t}\right]$$

where y_t is the bond's current special spread. The idea is that although all three sets of loadings satisfy equation (7), they each have a different dependence on the matrix Γ coming through their initial conditions, which depend on the maturity of each payment at the time the bond goes off special. No matter when the bond goes off special, the current level of the special spread needs to be taken into account, and crucially, this special spread is a dividend proportional to the value of the entire bond (all coupon payments). This is because, for the first 180 days, the principal and all coupon payments are exposed to the repo risk factor shocks.

The complication arises because each coupon has a different remaining time to maturity when the bond goes off special. Consider the first coupon payment, which occurs in 175 days. On this date, the bond is still special (and will be for an additional 5 days), so that the A^0, B^0, and C^0 loadings follow equation (7) with Γ from equation (12) and initial condition $A^0 = B^0 = C^0 = 0$. These loadings are valid for dates after t (issuance) up to the maturity of that coupon payment at $t+175$. After $t+175$ this coupon no longer appears in the pricing equation for this bond.

The second coupon payment in 365 days uses a different set of loadings, because when the bond goes off special in 180 days, that coupon still has $365-180 = 185$ days to maturity. Thus the A^{185}, B^{185}, and C^{185} loadings follow equation (7) with Γ from equation (12) but with initial condition $A_{185}^{185} = A_{185}$, $B_{185}^{185} = B_{185}$, and $C_{185}^{185} = C_{185}$, because at $t+180$ this bond is no longer exposed to the repo risk factors. Likewise, the final coupon and principal payment follows yet another set of loadings, with Γ from equation (12) but with initial condition $A_{3470}^{3470} = A_{3470}$, $B_{3470}^{3470} = B_{3470}$, and $C_{3470}^{3470} = C_{3470}$.

In all cases, price loadings follow equation (7); the issue is which intial condition to use. In order to price this bond for the 3650 days that it is observed in the data, we must compute loadings using the recursion in equation (7) four times: once for each coupon payment, and once for non-special bonds.[14] Given these loadings, the price of the bond depends only on

[14]To price the actual Treasury bonds in the data, we must compute the recursion (7) a total of 1,040

calendar time through the factors X_t. Thus, for example, the price 10 days after issuance would be

$$P_{t+10} = e^{y_{t+10}} \Big[c e^{A^0_{165} + B^{0\prime}_{165} X_{t+10} + X'_{t+10} C^0_{165} X_{t+10}} + c e^{A^{185}_{355} + B^{185\prime}_{355} X_{t+10} + X'_{t+10} C^{185}_{355} X_{t+10}} + (1+c)\, e^{A^{3470}_{3640} + B^{3470\prime}_{3640} X_{t+10} + X'_{t+10} C^{3470}_{3640} X_{t+10}} \Big]$$

while the prices at $t + 177$ (when the first coupon payment has been paid, but the bond is still special) and $t + 200$ (when the bond is off special) are given by

$$P_{t+177} = e^{y_{t+177}} \Big[c e^{A^{185}_{188} + B^{185\prime}_{188} X_{t+177} + X'_{t+177} C^{185}_{188} X_{t+177}} + (1+c)\, e^{A^{3470}_{3473} + B^{3470\prime}_{3473} X_{t+177} + X'_{t+177} C^{3470}_{3473} X_{t+177}} \Big]$$

$$\begin{aligned} P_{t+200} &= c e^{A^{185}_{165} + B^{185\prime}_{165} X_{t+200} + X'_{t+200} C^{185}_{165} X_{t+200}} + (1+c)\, e^{A^{3470}_{3450} + B^{3470\prime}_{3450} X_{t+200} + X'_{t+177} C^{3470}_{3450} X_{t+200}} \\ &= c e^{A_{165} + B'_{165} X_{t+200} + X'_{t+200} C_{165} X_{t+200}} + (1+c)\, e^{A_{3450} + B'_{3450} X_{t+200} + X'_{t+177} C_{3450} X_{t+200}} \end{aligned}$$

where the last line follows from the chosen initial conditions of the "special" loadings, which now coincide with the loadings of non-special bonds.

Notice that there is nothing crucial about the assumption that the bond is no longer special at the switching date $t + 180$. For example, we could extend the model to allow for a second, "off-the-run" repo factor $\widetilde{y}^S_t$ that applies to bonds that are second off-the-run or older. The only way this would change the computations would be that now the off-special loadings A_τ, B_τ, and C_τ would have a Γ different from zero, loading on that new factor, and equation (12) would have to be modified accordingly.

3.3 Identification

By Proposition 1, in order to price the cross-section of bonds we need only the "risk-neutral" parameters of the model: μ^*, Φ^*, Σ, δ_0, δ_1, and Γ. The remaining time-series parameters μ and Φ (or, equivalently, the price of risk parameters λ and Λ) can be identified by the dynamics of the factors X_t.[15] However, because some of the elements of X_t are latent, they must be invariant to translation and rotation; this means that not all elements of μ^*, Φ^*, δ_0, δ_1, and Δ are identifiable.

times: because we price bonds with time to maturity as long as 30 years, and Treasury bonds pay semi-annual coupons, in the data we have as many as sixty different maturities at off-special, per bond, for which to compute loadings. However, there are many "shared" coupon maturities, so that the calculation does not need to be redone for every single coupon of every single bond. For example, the 38 different 10-year notes in our sample have 8 maturities when they go off special (between 3,465 and 3,472 days), but there are only 5 different maturities at off-special for their penultimate coupons (between 3,284 and 3,288 days). In all 38 cases the first coupon payment is made just before the bond goes off special (i.e., becomes a second off-the-run bond).

[15] Although for the analysis in this paper, we need not do so, as described earlier.

Proposition 2. *Multiplying the factors X_t by any invertible matrix R^{-1}, and adjusting the risk-neutral parameters according to*

$$\begin{aligned}
\Gamma &\to R'\Gamma R \\
\Phi^* &\to R^{-1}\Phi^* R \\
\mu^* &\to R^{-1}\mu^* \\
\Sigma &\to R^{-1}\Sigma \\
\delta_1' &\to \delta_1' R
\end{aligned}$$

will lead to an identical fit to any set of bond prices. Likewise, adding the vector r to the factors, and adjusting the risk neutral parameters according to

$$\begin{aligned}
\mu^* &\to \mu^* + (I - \Phi^*)\, r \\
\delta_1' &\to \delta_1' + 2r'\Gamma \\
\delta_0 &\to \delta_0 - \left(\delta_1' + r'\Gamma\right) r
\end{aligned}$$

will also lead to an identical fit to the data. Finally, any $\widetilde{\Sigma}$ such that $\widetilde{\Sigma}\widetilde{\Sigma}' = \Sigma\Sigma'$ will also lead to an identical fit to the data.

Proof. The result can be verified by induction, comparing equations (1) and (6) using equation (7) for the two sets of factors and parameters. ∎

Our benchmark model without repo specials has three latent factors. We use Proposition 2 to identify the parameters of the model that pertain to these factors by assuming that $\Sigma = I/365$ (so that is equal to the identity matrix at an annual frequency), $\mu^* = \vec{0}$, δ_0 and δ_1 are free parameters, and the top-left corner of Φ^* has the form

$$\Phi^*_{\text{TL}} = \begin{bmatrix} \phi_1^* & \phi_2^* & 0 \\ \phi_3^* & \phi_4^* & 0 \\ \phi_5^* & \phi_6^* & \phi_7^* \end{bmatrix}. \tag{17}$$

Allowing the parameter ϕ_2^* to be non-zero allows Φ^*_{TL} to have two complex eigenvalues, which we find provides a better fit to the data than a lower-triangular Φ^*_{TL}. These identification assumptions leave us with 11 free parameters that affect the model's fit to the cross-section of only off-special bond prices.

Because the first three factors are latent, once we add the observed repo factors to the VAR, any non-zero values in the top-right corner of Φ^* (i.e., the dependence of future latent

factors on the repo factors under the risk-neutral measure) can be rotated away by re-defining the latent factors appropriately. Notice that according to Proposition 2, equation (12) allows us complete freedom to rotate and translate the first three latent factors in X_t without affecting Γ. It is the fact that $\Sigma \propto I$ and the form of Φ^* in equation (17) that prohibit us from such rotations of the latent factors, and the fact that $\mu^* = 0$ that prevents translations of the latent factors.[16]

Given these assumptions, the total set of risk-neutral parameters to be estimated is given by

parameter	size	number of free parameters
δ_0	1×1	1
$\delta_1 = [\delta_{11}, \delta_{12}, \delta_{13}, 0, 0, 0]'$	6×1	3
$\mu^* = \left[0, 0, 0, \mu_4^*, 0, y_\tau^D\right]'$	6×1	1
$\Phi^* = \begin{bmatrix} \Phi^*_{\text{TL}} & \vec{0} & \vec{0} & \vec{0} \\ \Phi^*_{\text{BL}} & \Phi^*_{\text{BR}} & \vec{0} & 0 \\ \vec{0} & \vec{0} & \rho & 0 \\ \vec{0} & \vec{0} & 0 & 0 \end{bmatrix}$	6×6	$7 + 3 + 1 + 1 = 12$
$\Sigma = \begin{bmatrix} \frac{1}{365} I & \vec{0} & \vec{0} & \vec{0} \\ \Sigma_{21} & \Sigma_{22} & 0 & 0 \\ \vec{0} & 0 & \sigma_x & 0 \\ \vec{0} & 0 & 0 & 0 \end{bmatrix}$	6×6	$3 + 1 + 1 = 5$

for a total of 22 free risk-neutral parameters in the benchmark model with a single aggregate repo factor y_t^S. Of course simpler models, for example a model without special spreads at all, have fewer parameters, as described in the next section. In addition, some parameters, including ρ and σ_x, can be estimated directly on the repo special spread data without needing to estimate a DTSM. In particular, we estimate x_t^i for each CUSIP i using equation (11). We then estimate ρ and σ_x from equation (13) via OLS, pooling across CUSIPs.

To estimate the remaining parameters we proceed in steps. First we minimize the sum of squared price residuals using only second off-the-run and older bonds and ignoring repo special spreads, as follows. Given values for the eleven parameters in δ_0, δ_1, and Φ^*_{TL}, we

[16] Notice however that this identification will not guarantee that the latent factor loadings have the usual level, slope, and curvature shapes. However, after estimation we can re-rotate parameters, and thus loadings, to reproduce these typical DTSM loading shapes.

estimate the factors X_t on each date that minimize the pricing residuals. These estimates are consistent for X_t as the number of bonds in each cross-section goes to infinity. We search over parameters values to maximize the combined log-likelihood from the residuals in equation (16) pooled across all dates t.[17]

To estimate the repo parameters of the model, we hold the latent factors X_t fixed and search over the remaining parameters in μ^*, Φ^*_{BL}, and Φ^*_{BR} to minimize the sum of squared pricing residuals on on-the-run and first off-the-run bonds. We estimate Σ_{21} and Σ_{22} directly from an estimated 4-VAR of the three factors and y^S_t.

4 Results

In this section, we estimate models of increasing complexity to illustrate the importance of including a time-varying risk premium on special repo factors for pricing Treasury securities with special collateral value in the repo market; we label this premium the SC risk premium.

In light of the fact that the 10-year bond has the largest price residuals and the largest special spreads in our sample (Table 1, Figure 2), in our benchmark specification we include only the 10-year on-the-run and first off-the-run notes in equation (10). We then proceed in stages.

First, we estimate a standard 3-factor model that ignores special spreads completely and thus includes only non-special bonds in the estimation (i.e., all securities with maturity at issuance other than 10 years, and 10-year bonds that are second off-the-run or older); and, we use this simple model to price the 10-year on-the-run and first off-the-run bonds (the "special" bonds), continuing to ignore the special repo factors. We report these parameters in Panel A of Table 2.

Second, we use the 3-factor model parameters, the estimated latent factors X_t, and the estimated dynamics of the special spreads to price the special repo factors in a risk-neutral fashion. That is, we force the parameters pertaining to the repo risk factor y^S_t to be the same under the physical and risk-neutral measures.[18] To do so, we set μ^* and Φ^* such that the fourth element of λ is zero, and the fourth row and column of Λ are all zeroes.

[Table 2 about here.]

[17] We ignore the resduals that come from equation (1) in this exercise. In practice, the more than 200 bonds in each cross-section contain much more information than the three factos X_{t-1} from the previous cross-section, so that our method is in practice very similar to a formal non-linear Kalman filter. In addition, nothing in our main results requires estimates of the physical measure parameters μ and Φ.

[18] Notice that although we need to estimate the physical dynamics of X_t in order to label these estimates "risk-neutral," we do not need to know the physical dynamics to estimate the full model with time-varying risk premia, since we estimate μ^* and Φ^* from equation (8) directly.

Third, we incorporate a constant risk premium on special bonds linked to the repo factor y_t^S by allowing the elements of μ^* beyond the first three to differ from the corresponding elements of μ. Finally, we allow for a time-varying risk premium associated to the repo factor y_t^S by estimating new values for the elements of μ^*, as well as for the bottom rows of Φ^* (that is, Φ^*_{BL} and Φ^*_{BR}).

In addition, for robustness, we also explore two alternative specifications that differ from the benchmark model in the bonds that they include as "special," and how they price specialness. We estimate a model that includes all on-the-run and first-off-the-run maturities (and not just the 10-year) in equation (10); we label this estimation the "pooled" model. Then, we move away from a single repo risk factor and extend y_t^S to be a 3×1 vector that averages over 30-year special spreads, 10-year special spreads, and 2-, 3-, 5-, and 7-year special spreads, separately; we label this estimation the "split" model. This model, differently from the benchmark and pooled models, is a 6-factor model that includes 3 latent factors and 3 observed repo risk factors. Similarly to the benchmark model, also for the pooled and split models we will estimate the risk-neutral, constant risk premium, and time-varying risk premium versions.

[Table 3 about here.]

Panel B of Table 2 reports the repo risk factor parameters for the 10-year only model and for the pooled model. Table 3 reports the repo risk factor parameters for the split model in which there are three separate repo risk factors, and thus six pricing factors in total. For each model, the tables report the values for the risk-neutral, constant risk premium, and time-varying risk premium versions as separate columns.

4.1 A missing risk factor

Before describing the results, we briefly analyze the individual residuals of the simple 3-factor model, which ignores special spreads, in order to stress the importance of this missing component. Figure 3 plots the standard deviation of the implied fitting error at each date across all the off-the-run bonds. The fit is generally good, although it varies over time: in particular, the model fits relatively less well at the beginning of the sample, during the height of the financial crisis after the failure of Lehman Brothers. Pancost (2017) examines the pricing residuals of a similar model in detail during this period and in late 2008, and finds that these errors are related to large price differences between bonds older and younger than fifteen years that cannot be explained by their different coupon rates and maturities.

[Figure 3 about here.]

Next, we explore in a reduced-form fashion whether the price residuals are related to special spreads, which would indicate an omitted factor. Duffie (1996) shows in a static setting that a security on special should have a higher cash price than an equivalent security that is not on special, and the price differential should be increasing in the specialness spread. In Table 4, we regress price residuals, $\eta_{i,t}$, from the simple 3-factor model on special spreads and their own past pricing residuals:

$$\eta_{i,t} = \alpha_i + \beta_1 y_t^i + \beta_2 \eta_{i,t-1} + \xi_{i,t}, \tag{18}$$

Although the parameters of this model are estimated only on non-special bonds, we estimate equation (18) on all bonds in our sample with more than one year left to maturity.[19]

The first column of Table 4 confirms Duffie's prediction by showing a positive and significant relation between the model price residuals and special spreads. In other words, the higher the special spread, the larger the amount of underpricing implied by the 3-factor model, which ignores that risk factor.

[Table 4 about here.]

The second column of Table 4 shows that higher price residuals continue to correlate with higher special spreads even after including the lagged price residual. Price residuals on individual CUSIPs are persistent: the R^2 jumps to almost 98% after adding yesterday's price residual to the regression. Controlling for previous price residuals greatly reduces the magnitude and statistical significance of the coefficient on special spreads, but does not eliminate it: bonds with high special spreads tend to have higher prices than they "should" according to a standard DTSM, even after controlling for yesterday's pricing error. The third and fourth columns show that these effects are strongest across, rather than within, bonds. All the coefficients on special spreads become smaller, though they remain economically and statistically significant.

The reduced-form analysis in this section shows that standard 3-factor DTSMs that ignore special spreads underprice bonds that are on special. In the next section we show that incorporating special spreads into a DTSM can explain a substantial fraction of special-bond prices, but only after allowing for time-varying risk premia linked to special repo spreads.

4.2 A missing risk premium

Although the results in Table 4 suggest that standard DTSMs are missing something by ignoring special spreads, we can do much more to quantify their importance. The top panel

[19] Results are largely similar if we exclude the 10-year on-the-run and first off-the-run notes.

of Figure 4 plots the price residuals on the 10-year on-the-run bond implied by our benchmark model, in which we treat 10-year on-the-run and first off-the-run bonds as "special." The black line (barely visible beneath the red line) is the price residual ignoring special spreads completely, which we report for comparison.

The red line is the price residual assuming that the risky special spreads are priced as risk-neutral dividends. In theory, if the profits from short-selling the on-the-run bond and going long a bond with similar cash-flows, for example the first off-the-run bond, were about zero on average, then the red line should hover near zero. That is, the profits from selling the more-expensive bond would be roughly offset by the cost of borrowing that bond (earning a negative interest rate) in the repo market. This is what Krishnamurthy (2002) finds for the 30-year on-the-run bond over his 1995–1999 sample period: the special spread and cash price are roughly consistent with risk-neutral repo special spreads. In contrast, in our sample period, the 10-year on-the-run premium is a good deal higher (almost 1% of par on average according to Table 1, and as high as 60 bps at times as shown in the top panel of Figure 1); while average special spreads, which accrue to the bond only for the first six months, are not high enough to wipe out the short-selling profits.[20] This is not consistent with a risk-neutral valuation of special spreads.

[Figure 4 about here.]

The green and blue lines in the top panel of Figure 4 plot price residuals for the models with constant and time-varying risk premia on the repo factor y_t^S, respectively. The model with constant risk premia lowers residuals at all points in the sample, which improves the fit until about 2012, after which the new residuals hover near zero (or go slightly negative). Any further increase in the risk premium would fit the financial crisis period even better, but at the cost of "over-predicting" prices after 2012, so that they would be actually higher than observed prices, pushing residuals into negative territory. The model with time-varying risk premia, on the other hand, is able to make the pricing errors almost mean-zero during the crisis without sacrificing much of the fit in the later part of the sample; although, in the second half of 2013, it overestimates the price by as much as 50–75 bps of par. Interestingly, this time period coincides with the immediate aftermath of the "Taper Tantrum" episode that we will discuss shortly.

The bottom panel of Figure 4 plots our yield measure of the SC risk premium: the yield difference between two 10-year *zero-coupon* bonds, both receiving the repo factor dividend y_t^S for 180 days, but one estimated using a risk-neutral valuation of this stream of repo

[20] Allowing bonds to be on special for their entire life and re-estimating the model does not solve the problem, because special spreads on far off-the-run bonds are so low on average (fourth column of Table 1).

dividends, and the other evaluated using our estimated time-varying prices of risk. Using such a measure of the SC risk premium, rather than just the difference between the red and blue lines plotted in the top panel of Figure 1, eliminates some of the auction-cycle wiggles evident in those series, which come in part because the actual coupon-bearing 10-year note has a time-varying exposure to the repo factor y_t^S that depends on the time since issuance. The measure plotted in the bottom panel of Figure 4 eliminates this seasonality by pricing a "constant maturity" synthetic 10-year bond that is consistently exposed to the repo factor for the next 180 days.[21]

The 10-year SC risk premium is large from 2009–2011, with a peak of 100 basis points, and then hovers around 25 basis points. Thus, its magnitude is similar to that one of the other classical risk premia, such as the real and inflation risk premium, estimated in the literature over this sample period (e.g., Breach, D'Amico and Orphanides, 2016). To try to understand what drives our SC risk premium, the vertical lines in the bottom panel of Figure 4 mark the times of major macroeconomic events particularly relevant for U.S. Treasury bonds. The blue lines indicate supply shocks approximated by the FOMC announcements of the Federal Reserve's asset purchase programs (Cahill et al. 2013). By far, the largest drop downward in the SC risk premium occurred near the beginning of the sample, on March 18, 2009 when the Federal Reserve announced the first round of quantitative easing (QE1). It is possible that the prospect of a large new investor (the Fed) willing to buy seasoned Treasuries, at that time regarded as considerably less liquid than recently-issued securities, increased their trading opportunities, making the 10-year on- and first off-the-run relatively less special. Similarly, the FOMC announcement on August 10, 2010, in which the Fed promised to reinvest agency-MBS principal payments into Treasury securities also seems to have induced a decline in the SC risk premium.

In contrast, the second round of quantitative easing (QE2), which was announced on November 3, 2010, and the September 21, 2011 announcement of the Maturity Extension Program (MEP), in which the Fed promised to sell shorter-term Treasuries in order to buy longer-term securities to extend the duration of its portfolio, produce a spike upwards in the SC risk premium. One possibility is that those additional rounds of purchases, mostly focused on longer-term Treasuries, were perceived as potentially increasing the relative scarcity of 10-year bonds, making them more special. However, the QE3 announcement on December 12, 2012 has essentially no effect on the repo risk premium.

Another event that seems to have impacted the SC risk premium is Chairman Bernanke's Congressional testimony on May 22, 2013, in which his remarks about the possibility of mod-

[21]The synthetic bond also has no coupon payments; some of the seasonality could be due in part to the changing coupon rate on subsequent 10-year notes.

erating the pace of asset purchases later that year led to the subsequent "Taper Tantrum." That is, a sharp and large increase in longer-term yields that lasted for about 3 months, also accompanied by increased volatility. This, is turn, might have induced a sudden reversal in hedging and speculative positions involving short-selling of the 10-year bond, pushing the premium up.

The red lines, in contrast, plot the dates at which the bid-to-cover ratio in the 10-year note auction exceeded two standard deviations of its mean over our sample period. We use these dates to proxy for unexpected heightened demand for 10-year Treasuries (Gorodnichenko and Ray 2018). A glance at the figure shows that there is very little change in the SC risk premium at these times; thus, the evidence from the bottom panel of Figure 4 suggests that the SC risk premium primarily varies with shocks to Treasury supply rather than demand, at least over our sample period and given our proxy for demand shocks.

[Figure 5 about here.]

We also consider an additional proxy of demand shocks for U.S. nominal Treasuries: flight-to-safety episodes. These are events during which higher uncertainty leads investors to reduce holdings of riskier and less liquid securities in favor of safer and more liquid assets. In the spirit of Baele, Bekaert, Inghelbrecht and Wei (2018), we define flight-to-safety as an episode that satisfies three criteria at a daily frequency: high market stress measured by an increase in equity market volatility (identified by at least a 5-basis-point increase in VIX), a decrease in equity prices (i.e., the S&P500 index loses more than 1% of value), and an increase in Treasury bond prices (identified by a decline of at least 5 basis points in the 10-year Treasury yield), which in turn implies a negative correlation between bond and stock returns over these episodes.

The vertical blue lines in Figure 5 mark the flight-to-safety episodes identified by those criteria in our sample and, as can be seen, only the first 3, which are magnified in the bottom panel, trigger an increase in the 10-year SC risk premium. It is possible that in periods of higher risk aversion, as during the financial crisis and its immediate aftermath, the SC risk premium is more reactive to these types of events; while, during periods of normal market conditions they do not matter much.

4.2.1 Robustness

Up to this point we have defined only the 10-year on-the-run and first off-the-run notes as "special" for calculating the value of y_t^S in equation (10). To examine the robustness of our results to alternative definitions of the repo factor, we also estimate a "pooled" model incorporating special spreads on bonds of maturities of 2, 3, 5, 7, and 30 years in addition to

10 years. The estimated parameters are reported in the last three columns Panel B of Table 2. We then estimate a "split" model in which we allow for three separate y_t^S factors: one for 30-year bonds, one for 10-year bonds, and one for 2-, 3-, 5-, and 7-year bonds together. The parameters of this third model are reported in Table 3.

[Figure 6 about here.]

Figure 6 plots the time series of our repo factor y_t^S for the three estimated models. y_t^S as defined in equation (10) is an average across bonds of special spread deviations from the deterministic auction-cycle component; the top panel plots y_t^S for the benchmark and pooled models, with the blue line including only 10-year special bonds and the red line including special bonds of all maturities. The two series are qualitatively similar, with a correlation coefficient of 0.76, although it is evident that a couple of spikes, in late 2011 and summer of 2013, are specific to the 10-year securities. Those episodes coincide with the MEP and the "Taper Tantrum," and seem to align with the increases in the SC risk premium.

The bottom panel of Figure 6 plots the three time-series of y_t^S for the split model, where the black line includes 30-year special bonds, the blue line includes 10-year special bonds, and the red line includes 2-, 3-, 5-, and 7-year bonds. In the first half of the sample the three series behave quite differently, and towards the end of the sample they tend to move together, though the 10-year repo risk factor still displays the sharpest fluctuations. Interestingly, all three series exhibit the same upward time trend as the benchmark y_t^S in the top panel of Figure 6.

[Figure 7 about here.]

For each model we estimate risk-neutral, constant risk premium, and time-varying risk premium versions and use each to price the 10-year on the run note, as in the top panel of Figure 4; the results are plotted in Figure 7. They are qualitatively similar to the 10-year only model, although the pooled model, being forced to explain more variation with the same number of parameters, tends to capture substantially less of the 10-year price premium and generates a smaller repo risk premium than the other two models (the difference between the red and blue lines).

To more easily compare results across models, Table 5 quantifies the amount of variation in on-the-run price residuals explained by our estimated models, by reporting the sum of squared residuals for each model and its alternative versions: no special spreads, risk-neutral evaluation, a constant risk premium, and a time-varying risk premium.

[Table 5 about here.]

In all three model specifications, a time-varying risk premium is necessary to price 10-year on-the-run and first off-the-run bonds, as shown by the drastic reduction in the sum of squared residuals in Table 5. The same is true for the 30-year bond. In contrast, this is not the case for short-term securities, as evidenced by the lack of improvement across the first 4 columns.

The last three columns report the share of variation explained, for each model and maturity, in the form of an R^2, defined as follows:

$$R_i^2 \equiv 1 - \frac{\eta_i' \eta_i}{\eta_0' \eta_0}, \tag{19}$$

where η_0 is the vector of price residuals under the risk-neutral evaluation, and η_i is the vector of price residuals of the ith alternative model. Thus, for example, from the first row of Table 5, the 10-year only model delivers, for 10-year special bonds, a total sum of squared residuals of 0.379; with a time-varying risk premium this value drops to 0.0423, which is an R^2 of 88.9%. The model achieves this 89% fit on 2,252 price observations with only 5 free parameters (as the other repo parameters reported in Panel B of Table 2, ρ, σ_x, Σ_{21}, and Σ_{22}, are the same across the risk-neutral, constant, and time-varying risk premium models).

While the explanatory power of the pooled and split models is roughly similar to the benchmark model at the the 10-year maturity, they perform differently at other maturities. In particular, the pooled model actually has a negative R^2 for maturities other than 10 years. This is because the price residuals on the 10-year note are so much larger than the other maturities, that to minimize the total sum of squared residuals the pooled model sacrifices the fit on other maturities to better fit the 10-year. However, with only a single y_t^S that averages across special spreads of all maturities (not just the 10-year) and only five free parameters, it is still not able to fit the 10-year prices as well as the 10-year only model. Nevertheless, the SC risk premium that comes out of this model has a correlation coefficient of 0.97 with the one from our benchmark model.

In contrast, the split model is able to explain the 10-year prices as well as (actually slightly better than) the 10-year only model, in addition to explaining over 40% of the 30-year on-the-run premium and over 13% of the 2-year on-the-run premium. It does so at the cost of a very poor fit to the 5- and 7-year on-the-run premia. But such a poor fit at these maturities is "optimal" in the sense that those on-the-run price residuals are already quite small, as can be seen in the first column of Table 5, relative to the 10- and 30-year sum of squared residuals. Thus, the split model makes the fit at these maturities just a bit worse in order to better explain the larger deviations at the 10- and 30-year maturities. We conjecture that the performance of the split model can be further improved by separating

the repo risk factor for the 2- and 3-year maturities from the repo risk factor for the 5- and 7-year maturities.

4.3 A commonality across price anomalies

To analyze the ability of the SC risk premium to explain relative price anomalies in large and liquid fixed-income markets, Table 6 presents correlations and regression R^2 between our measures of the SC risk premium and various measures of price anomalies documented in the literature. The first row shows that the SC risk premium has a 91% correlation with, and explains 83% of, the GSW on-the-run premium, plotted over our sample period in the top panel of Figure 1. This series is a standard measure of the on-the-run premium (see for example Adrian, Fleming and Vogt 2017) that is independent of our model. However, because the GSW model has six free parameters on each day to price Treasury bonds, it does a slightly better job than our 4-factor benchmark model, which has only 3 degrees of freedom on each date, in addition to the 10-year note's special spread. This clarifies why our model explains 89% of the 10-year price residuals plotted in the top panel of Figure 4, but 83% of the GSW on-the-run premium. This finding is in line with the calibration exercise of Vayanos and Weill (2008) that shows that the majority of the on-the-run premium is due to specialness rather than liquidity.

[Table 6 about here.]

The 10-year SC risk premium is also highly correlated with four other Treasury price anomalies documented in the literature: the 10-year TIPS-Treasury bond puzzle of Fleckenstein, Longstaff and Lustig (2014); the off-the-run note-bond spread of Musto, Nini and Schwarz (2017), which we derive as the difference between the average prices of off-the-run Treasury notes and bonds with duration between 7 and 9 years to make their maturity as close as possible to 10 years; the 10-year TIPS liquidity premium of D'Amico, Kim and Wei (2018); and the average absolute nominal fitting errors, which can be interpreted as a measure of limits to arbitrage (see for example Hu, Pan and Wang 2013). Except for the latter, which is derived from all outstanding Treasuries, the other measures are all roughly 10-year spreads to recently-issued nominal Treasury securities, and as such may appear anomalous when those high-quality securities are particularly valuable as collateral and that value is not priced in as a risk factor.

Table 6 shows that whether we derive the 10-year SC risk premium using all special maturities or just the 10-year, and whether we "split" y_t^S into three maturity groups or pool all maturities together, the SC risk premium can explain a substantial fraction of these price

anomalies. In particular, it accounts for 68% of the variation in the TIPS-Treasury bond puzzle, 62% of the variation in the off-the-run note-bond spread, and 58% of the variation in the TIPS liquidity premium. It is not surprising that, relative to the 10-year on-the-run premium, a smaller share of those spreads is explained by the 10-year SC risk premium. This is because, as the lower-quality asset becomes increasingly dissimilar from the first off-the-run nominal security (either because it is more seasoned or in a different asset category), factors other than collateral value should start playing a bigger role. For example, trading liquidity, as opposed to funding liquidity, might explain part of the remaining variation.

In addition, as shown in the last row, the 10-year SC risk premium accounts also for 36% of the variation in average absolute nominal fitting errors, indicating that, despite the mismatch in maturities (i.e., averaging across all maturities versus an exact 10-year maturity), relative differences in the observed SC value of Treasury securities play a significant role in explaining daily deviations from the yield curve. In other words, it seems reasonable that the cost of short-selling Treasury securities might affect investors' ability to take advantage of arbitrage opportunities, proxied by yield curve fitting errors.

The second two columns of Table 6 report correlations and R^2 for the "pooled" model, while the remaining columns report the same quantities for two versions of the SC risk premium from the "split" model: one constructed as the yield spread on a synthetic 10-year zero-coupon bond, and the other as a synthetic 30-year zero-coupon bond. In all cases, the R^2 are at least 43% and often higher, suggesting that these price anomalies are largely justified by time variation in the risk associated with the value of special collateral.

5 Conclusion

This paper shows that in traditional DTSMs of U.S. nominal Treasury securities there seems to be a missing risk factor: the special-collateral repo spread. This, in turn, implies that there is a missing risk premium: the special-collateral repo risk premium. This omission has led some to interpret price differences relative to the highest quality collateral—recently-issued nominal Treasury securities—entirely as price anomalies. In contrast, if observed special-collateral repo risk factors are explicitly priced in, these anomalies are mostly justified by compensation for exposure to special-collateral risk, and thus are not that anomalous. That is, rational investors seem to account for the uncertain stream of expected repo dividends, in addition to the classical risk factors, when pricing nominal Treasuries.

This paper has only begun to scratch the surface of what is possible with the special repo rate data. For example, although special spreads on far off-the-run bonds tend to be low on average (Table 1), they do correlate with price residuals on these bonds (Table 4), and

it may be the case that an additional, off-the-run repo risk factor can improve the model fit to off-the-run bonds plotted in Figure 3. Such a model can in principle even improve the fit to *on*-the-run bonds, since it would change the initial condition used for the on-the-run loadings, as a newly-issued ten-year bond would be exposed to the on-the-run repo risk factor for six months (while on-the-run and first off-the-run), and then to the "off-the-run" factor for an additional 9.5 years. 30-year bonds would be exposed to the off-the-run factor for 29.5 years. However, because there are so many more off-the-run than on-the-run bonds, maximum-likelihood estimation of this model could sacrifice the fit to the on-the-run bonds, in order to marginally improve the fit on all the off-the-run bonds. Notwithstanding such complications, we view this as an exciting area of future research.

Further, and ideally, it would be great if, on top of pricing jointly Treasury cash and repo rates, we could also bring into the model Treasury future rates. This would complicate further the security-level pricing within a DTSM and we leave these challenges to future research.

A Proofs

A.1 Proof of Proposition 1

The result follows by induction. First note that a zero-coupon bond pays \$1 at maturity, so that $A_0 = 0$, $B_0' = \vec{0}$, and $C_0 = \vec{0}$ as in equation (7) prices bonds at maturity. Next, fix n and assume that at any time t, the price of an $n-1$ period bond satisfies

$$\log P_t^{(n-1)} = A_{n-1} + B_{n-1}' X_t + X_t' C_{n-1} X_t. \tag{20}$$

It then suffices to show that equation (20) implies equation (7) for bonds with maturity n.

The log price of an n-period zero-coupon bond at time t with special spread y_t is given by

$$\begin{aligned}
\log P_t^{(n)} &= y_t + \log E_t \left\{ \frac{M_{t+1}}{M_t} P_{t+1}^{(n-1)} \right\} \\
&= X_t' \Gamma X_t + \log E_t \left\{ \frac{M_{t+1}}{M_t} \exp\left\{ A_{n-1} + B_{n-1}' X_{t+1} + X_{t+1}' C_{n-1} X_{t+1} \right\} \right\} \qquad (21) \\
&= X_t' \Gamma X_t + \log E_t \exp\Bigg\{ -\delta_0 - \delta_1' X_t - \frac{1}{2}\lambda_t'\lambda_t - \lambda_t' \varepsilon_{t+1} \\
&\quad + A_{n-1} + B_{n-1}' \left(\mu + \Phi X_t + \Sigma \varepsilon_{t+1}\right) \\
&\quad + \left(\mu + \Phi X_t + \Sigma \varepsilon_{t+1}\right)' C_{n-1} \left(\mu + \Phi X_t + \Sigma \varepsilon_{t+1}\right) \Bigg\} \\
&= X_t' \Gamma X_t - \delta_0 - \delta_1' X_t - \frac{1}{2}\lambda_t'\lambda_t + A_{n-1} + B_{n-1}' \left(\mu + \Phi X_t\right) + \left(\mu + \Phi X_t\right)' C_{n-1} \left(\mu + \Phi X_t\right) \\
&\quad + \log E_t \exp\left\{ m' \varepsilon_{t+1} + \varepsilon_{t+1}' \Sigma' C_{n-1} \Sigma \varepsilon_{t+1} \right\}
\end{aligned}$$

where the second line uses equations (5) and (20), the next line plugs in equations (1) and (3), and m in the last line is given by

$$\begin{aligned}
m &\equiv -\lambda_t + \Sigma' B_{n-1} + 2\Sigma' C_{n-1} \left(\mu + \Phi X_t\right) \\
&= -\lambda + \Sigma' B_{n-1} + 2\Sigma' C_{n-1} \mu + \left(2\Sigma' C_{n-1} \Phi - \Lambda\right) X_t.
\end{aligned}$$

Because ε_{t+1} is a standard multivariate normal random variable, we have that

$$\log E_t \exp\left\{ m' \varepsilon_{t+1} + \varepsilon_{t+1}' \Sigma' C_{n-1} \Sigma \varepsilon_{t+1} \right\} = \frac{1}{2} m' G_{n-1} m + \frac{1}{2} \log \left| G_{n-1} \right|, \tag{22}$$

where $G_{n-1} = \left[I - 2\Sigma' C_{n-1} \Sigma \right]^{-1}$ and $|G_{n-1}|$ denotes the determinant of G_{n-1}. Equation (22) holds provided G_{n-1} is positive semi-definite, and can be derived by completing the square.

Plugging equation (22) into equation (21) and combining quadratic, linear, and scalar terms yields

$$\begin{aligned}
\log P_t^{(n)} &= \begin{pmatrix} 1 & X_t' \end{pmatrix} M_{n-1} \begin{pmatrix} 1 \\ X_t \end{pmatrix} \\
M_{n-1} &\equiv -\widetilde{\Delta} - \frac{1}{2}\widetilde{\Lambda}'\widetilde{\Lambda} + \frac{1}{2}\widetilde{\Lambda}'G_{n-1}\widetilde{\Lambda} + \widetilde{\Phi}'C_{n-1}\widetilde{\Phi} + 2\widetilde{\Phi}'C_{n-1}\Sigma G_{n-1}\Sigma' C_{n-1}\widetilde{\Phi} \\
&\quad - \left(\widetilde{\Lambda}'G_{n-1}\Sigma' C_{n-1}\widetilde{\Phi} + \widetilde{\Phi}'C_{n-1}\Sigma G_{n-1}\widetilde{\Lambda}\right) \\
&\quad + \widetilde{B}_{n-1}' \left(\widetilde{\Phi} - \Sigma G_{n-1}\widetilde{\Lambda} + 2\Sigma G_{n-1}\Sigma' C_{n-1}\widetilde{\Phi}\right) + \widetilde{A}_{n-1}
\end{aligned} \tag{23}$$

where

$$\begin{aligned}
\widetilde{\Delta} &\equiv \begin{bmatrix} \delta_0 & \frac{1}{2}\delta_1' \\ \frac{1}{2}\delta_1 & -\Gamma \end{bmatrix} \\
\widetilde{\Lambda} &\equiv \begin{bmatrix} \lambda & \Lambda \end{bmatrix} \\
\widetilde{\Phi} &\equiv \begin{bmatrix} \mu & \Phi \end{bmatrix} \\
\widetilde{B}_{n-1} &\equiv \begin{bmatrix} B_{n-1} & \vec{0}_{n\times n} \end{bmatrix} \\
\widetilde{A}_{n-1} &\equiv \begin{bmatrix} A_{n-1} + \log|G_{n-1}| + \frac{1}{2}B_{n-1}'\Sigma G_{n-1}\Sigma' B_{n-1} & \vec{0}_{1\times n} \\ \vec{0}_{n\times 1} & \vec{0}_{n\times n} \end{bmatrix}.
\end{aligned}$$

The remainder of the proof consists of showing that equation (23) is equivalent to equation (7). To do so, we use the fact (proven below in Lemma 1) that the matrix $C_{n-1}D_{n-1} = C_{n-1}\Sigma G_{n-1}\Sigma^{-1}$ is symmetric. For notational simplicity, in what follows we drop all the $n-1$ subscripts.

Equation (23) can be written as

$$\begin{aligned}
M &= -\widetilde{\Delta} - \frac{1}{2}\widetilde{\Lambda}'G\left(G^{-1} - I\right)\widetilde{\Lambda} + \widetilde{\Phi}'C\Sigma G\Sigma'\left(2C + J^{-1}\right)\widetilde{\Phi} \\
&\quad - \left(\widetilde{\Lambda}'G\Sigma' C\widetilde{\Phi} + \widetilde{\Phi}'C\Sigma G\widetilde{\Lambda}\right) \\
&\quad + \widetilde{B}'\Sigma G\left(G^{-1}\Sigma^{-1}\widetilde{\Phi} - \widetilde{\Lambda} + 2\Sigma' C\widetilde{\Phi}\right) + \widetilde{A},
\end{aligned}$$

where $J \equiv \Sigma G \Sigma'$ so that

$$\begin{aligned} J^{-1} &= \Sigma'^{-1} G^{-1} \Sigma^{-1} \\ &= \Sigma'^{-1} \left(I - 2\Sigma' C \Sigma \right) \Sigma^{-1} \\ &= \Sigma'^{-1} \Sigma^{-1} - 2C. \end{aligned} \tag{24}$$

Plugging in equation (24) and the fact that $G^{-1} = I - 2\Sigma' C \Sigma$ and rearranging yields

$$\begin{aligned} M = &-\widetilde{\Delta} + \widetilde{\Lambda}' G \Sigma' C \Sigma \widetilde{\Lambda} + \widetilde{\Phi}' C \Sigma G \Sigma^{-1} \widetilde{\Phi} \\ &- \left(\widetilde{\Lambda}' G \Sigma' C \widetilde{\Phi} + \widetilde{\Phi}' C \Sigma G \widetilde{\Lambda} \right) \\ &+ \widetilde{B}' \Sigma G \left(\Sigma^{-1} \widetilde{\Phi} - \widetilde{\Lambda} - 2\Sigma' C \widetilde{\Phi} + 2\Sigma' C \widetilde{\Phi} \right) + \widetilde{A} \\ = &-\widetilde{\Delta} + \widetilde{\Lambda}' G \Sigma' C \left(\Sigma \widetilde{\Lambda} - \widetilde{\Phi} \right) + \widetilde{\Phi}' C \Sigma G \Sigma^{-1} \left(\widetilde{\Phi} - \Sigma \widetilde{\Lambda} \right) \\ &+ \widetilde{B}' \Sigma G \Sigma^{-1} \left(\widetilde{\Phi} - \Sigma \widetilde{\Lambda} \right) + \widetilde{A} \\ = &-\widetilde{\Delta} + \left(\widetilde{\Phi}' C \Sigma G \Sigma^{-1} - \widetilde{\Lambda}' \Sigma' \Sigma^{-1\prime} G \Sigma' C \right) \left(\widetilde{\Phi} - \Sigma \widetilde{\Lambda} \right) \\ &+ \widetilde{B}' \Sigma G \Sigma^{-1} \left(\widetilde{\Phi} - \Sigma \widetilde{\Lambda} \right) + \widetilde{A} \\ = &-\widetilde{\Delta} + \left(\widetilde{\Phi} - \Sigma \widetilde{\Lambda} \right)' C \Sigma G \Sigma^{-1} \left(\widetilde{\Phi} - \Sigma \widetilde{\Lambda} \right) + \widetilde{B}' \Sigma G \Sigma^{-1} \left(\widetilde{\Phi} - \Sigma \widetilde{\Lambda} \right) + \widetilde{A}, \end{aligned}$$

where the last line applies Lemma 1. Separating M into quadratic, linear, and scalar terms gives the loadings in equation (7).

Lemma 1. *The matrix $C_{n-1} D_{n-1} = C_{n-1} \Sigma G_{n-1} \Sigma^{-1}$ is symmetric for all n.*

Proof. Using equation (23), so long as Γ is symmetric, then C_{n-1} and G_{n-1} are both symmetric for all n. In what follows we drop the $n-1$ subscripts. Let $H \equiv C \Sigma G \Sigma^{-1}$, so that we need to show that $H = H' = \Sigma'^{-1} G \Sigma' C$.

By definition, $G^{-1} = 1 - 2\Sigma' C \Sigma$, so that

$$\begin{aligned} \Sigma G^{-1} \Sigma^{-1} &= \Sigma \left(1 - 2\Sigma' C \Sigma \right) \Sigma^{-1} \\ &= I - 2\Sigma \Sigma' C \end{aligned} \tag{25}$$

$$\begin{aligned} \Sigma'^{-1} G^{-1} \Sigma' &= \Sigma'^{-1} \left(I - 2\Sigma' C \Sigma \right) \Sigma' \\ &= I - 2C \Sigma \Sigma'. \end{aligned} \tag{26}$$

Then we have that

$$
\begin{aligned}
H &= C\Sigma G\Sigma^{-1} \\
&= \underbrace{\left(\Sigma'^{-1}G\Sigma'\right)\left(\Sigma'^{-1}G^{-1}\Sigma'\right)}_{=I} C\Sigma G\Sigma^{-1} \\
&= \left(\Sigma'^{-1}G\Sigma'\right) \underbrace{\left(I - 2C\Sigma\Sigma'\right)}_{\text{by equation (26)}} C\Sigma G\Sigma^{-1} \\
&= \left(\Sigma'^{-1}G\Sigma'\right)\left(C - 2C\Sigma\Sigma'C\right)\Sigma G\Sigma^{-1} \\
&= \left(\Sigma'^{-1}G\Sigma'C\right)\left(I - 2\Sigma\Sigma'C\right)\Sigma G\Sigma^{-1} \\
&= H' \underbrace{\left(\Sigma G^{-1}\Sigma^{-1}\right)}_{\text{by equation (25)}} \Sigma G\Sigma^{-1} \\
&= H'.
\end{aligned}
$$

∎

References

Adrian, Tobias, Michael Fleming, and Erik Vogt. 2017. "An Index of Treasury Market Liquidity: 1991–2017." *Federal Reserve Bank of New York Staff Report 827.* 29

Ahn, Dong-Hyun, Robert F. Dittmar, and A. Ronald Gallant. 2002. "Quadratic Term Structure Models: Theory And Evidence." *The Review of Financial Studies*, 15(1): 243–288. 11

Andreasen, Martin M., and Bent Jesper Christensen. 2015. "The SR Approach: A New Estimation Procedure For Non-Linear And Non-Gaussian Dynamic Term Structure Models." *Journal of Econometrics*, 184(2): 420–451. 3, 16

Ang, Andrew, and Monika Piazzesi. 2003. "A No-Arbitrage Vector Autoregression of Term Structure Dynamics with Macroeconomic and Latent Variables." *Journal of Monetary Economics*, 50(4): 745–787. 4

Bauer, Michael D., and Glenn D. Rudebusch. 2014. "The Signaling Channel For Federal Reserve Bond Purchases." *International Journal of Central Banking*, 10(3): 233–289. 4

Breach, Tomas, Stefania D'Amico, and Athanasios Orphanides. 2016. "The Term Structure And Inflation Uncertainty." Federal Reserve Bank of Chicago Working Paper Series WP-2016-22. 12, 25

Buraschi, Andrea, and Davide Menini. 2002. "Liquidity Risk And Specialness." *Journal of Financial Economics*, 64(2): 243–284. 3

Cahill, Michael E., Stefania D'Amico, Canlin Li, and John S. Sears. 2013. "Duration Risk Versus Local Supply Channel In Treasury Yields: Evidence From The Federal Reserve's Asset Purchase Announcements." Board of Governors of the Federal Reserve System (U.S.) Finance and Economics Discussion Series 2013-35. 3, 25

Cherian, Joseph A., Eric Jacquier, and Robert A. Jarrow. 2004. "A Model Of The Convenience Yields In On-The-Run Treasuries." *Review of Derivatives Research*, 7: 79–97. 3

Christensen, Jens H.E., Francis X. Diebold, and Glenn D. Rudebusch. 2011. "The Affine Arbitrage-Free Class Of Nelson?OOsiegel Term Structure Models." *Journal of Econometrics*, 164(1): 4–20. 4

Creal, Drew D, and Jing Cynthia Wu. 2016. "Bond Risk Premia in Consumption-based Models." Working paper. 4

D'Amico, Stefania, Don H. Kim, and Min Wei. 2010. "Tips from TIPS: the Informational Content of Treasury Inflation-Protected Security Prices." *Federal Reserve Board Finance and Economics Discussion Series*, 2010–19. 4

D'Amico, Stefania, Don H. Kim, and Min Wei. 2018. "Tips From TIPS: The Informational Content Of Treasury Inflation-Protected Security Prices." *Journal of Financial and Quantitative Analysis*, 53(1): 395–436. 29, 52

D'Amico, Stefania, Roger Fan, and Yuriy Kitsul. 2017. "The Scarcity Value Of Treasury Collateral: Repo Market Effects Of Security-Specific Supply And Demand Factors." Forthcoming, *Journal of Financial and Quantitative Analysis*. 2

Diebold, Francis X., Glenn D. Rudebusch, and S. Borağan Aruoba. 2006. "The Macroeconomy and the Yield Curve: a Dynamic Latent Factor Approach." *Journal of Econometrics*, 131(1-2): 309–338.

Duffie, Darrell. 1996. "Special Repo Rates." *The Journal of Finance*, 51(2): 493–526. 2, 4, 5, 10, 23

Duffie, Darrell, Nicolae Gârleanu, and Lasse Heje Pedersen. 2002. "Securities Lending, Shorting, And Pricing." *Journal of Financial Economics*, 66(2): 307–339. Limits on Arbitrage. 5

Fama, Eugene F., and Robert R. Bliss. 1987. "The Information in Long-Maturity Forward Rates." *The American Economic Review*, 77(4): 680–692. 4

Feldhutter, Peter, and David Lando. 2008. "Decomposing Swap Spreads." *Journal of Financial Economics*, 88(2): 375–405. 3

Fisher, Mark. 2002. "Special Repo Rates: An Introduction." *Economic Review*, , (Q2): 27–43. 8

Fleckenstein, Matthias, Francis A. Longstaff, and Hanno Lustig. 2014. "The TIPS-Treasury Bond Puzzle." *The Journal of Finance*, 69(5): 2151–2197. 2, 29

Fleming, Michael, and Frank Keane. 2014. "What Explains the June Spike in Treasury Settlement Fails?" *Liberty Street Economics*, Accessed 5/14/2018. 8

Fleming, Michael, and Frank Keane. 2016. "Characterizing the Rising Settlement Fails in Seasoned Treasury Securities." *Liberty Street Economics*, Accessed 10/16/2018. 5, 8

Fleming, Michael J., and Kenneth D. Garbade. 2003. "The Repurchase Agreement Refined: GCF Repo." *Current Issues in Economics and Finance*, 9(6). 6

Gorodnichenko, Yuriy, and Walker Ray. 2018. "Unbundling Quantitative Easing: Taking a Cue from Treasury Auctions." Working paper. 3, 26

Graveline, Jeremy J., and Matthew R. McBrady. 2011. "Who Makes On-The-Run Treasuries Special?" *Journal of Financial Intermediation*, 20(4): 620–632. 3

Grinblatt, Mark. 2002. "An Analytic Solution For Interest Rate Swap Spreads." *International Review of Finance*, 2(3): 113–149. 3

Gürkaynak, Refet S., Brian Sack, and Jonathan H. Wright. 2007. "The U.S. Treasury Yield Curve: 1961 to the Present." *Journal of Monetary Economics*, 54(8): 2291–2304. 4, 7, 40, 52

Gürkaynak, Refet S, Brian Sack, and Jonathan H Wright. 2010. "The TIPS Yield Curve and Inflation Compensation." *American Economic Journal: Macroeconomics*, 2(1): 70–92.

Hamilton, James D., and Jing Cynthia Wu. 2012. "The Effectiveness Of Alternative Monetary Policy Tools In A Zero Lower Bound Environment." *Journal of Money, Credit and Banking*, 44: 3–46. 4

Hu, Grace Xing, Jun Pan, and Jiang Wang. 2013. "Noise As Information For Illiquidity." *The Journal of Finance*, 68(6): 2341–2382. 29

Jordan, Bradford D., and Susan D. Jordan. 1997. "Special Repo Rates: An Empirical Analysis." *The Journal of Finance*, 52(5): 2051–2072. 2

Joslin, Scott, Kenneth J. Singleton, and Haoxiang Zhu. 2011. "A New Perspective On Gaussian Dynamic Term Structure Models." *Review of Financial Studies*, 24(3): 926–970. 4

Keane, Frank. 1995. "Expected Repo Specialness Costs And The Treasury Auction Cycle." Federal Reserve Bank of New York Research Paper No. 9504. 8

Kim, Don H. 2004. "Time-varying risk and return in the quadratic-gaussian model of the term-structure." Ph.D. Dissertation, Stanford University. 12

Kim, Don H., and Athanasios Orphanides. 2012. "Term Structure Estimation With Survey Data On Interest Rate Forecasts." *Journal of Financial and Quantitative Analysis*, 47(1): 241–272. 4

Kim, Don H., and Kenneth J. Singleton. 2012. "Term Structure Models And The Zero Bound: An Empirical Investigation Of Japanese Yields." *Journal of Econometrics*, 170(1): 32–49. 11

Klingler, Sven, and Suresh Sundaresan. 2018. "An Explanation of Negative Swap Spreads: Demand for Duration from Underfunded Pension Plans." Forthcoming, *Journal of Finance*.

Krishnamurthy, Arvind. 2002. "The Bond/Old-Bond Spread." *Journal of Financial Economics*, 66(2–3): 463–506. 2, 4, 7, 24

Musto, David, Greg Nini, and Krista Schwarz. 2017. "Notes on Bonds: Illiquidity Feedback During the Financial Crisis." Working Paper. 2, 29

Pancost, N. Aaron. 2017. "Zero-Coupon Yields and the Cross-section of Bond Prices." Working Paper, University of Chicago. 2, 4, 22

Pasquariello, Paolo, and Clara Vega. 2009. "The On-The-Run Liquidity Phenomenon." *Journal of Financial Economics*, 92(1): 1–24. 5

Rudebusch, Glenn D., and Tao Wu. 2008. "A Macro-Finance Model of the Term Structure, Monetary Policy and the Economy." *The Economic Journal*, 118(530): 906–926.

Sundaresan, Suresh. 1994. "An Empirical Analysis Of U.S. Treasury Auctions." *The Journal of Fixed Income*, 4(2): 35–50. 9

Vayanos, Dimitri, and Pierre-Olivier Weill. 2008. "A Search-Based Theory Of The On-The-Run Phenomenon." *The Journal of Finance*, 63(3): 1361–1398. 2, 3, 5, 29

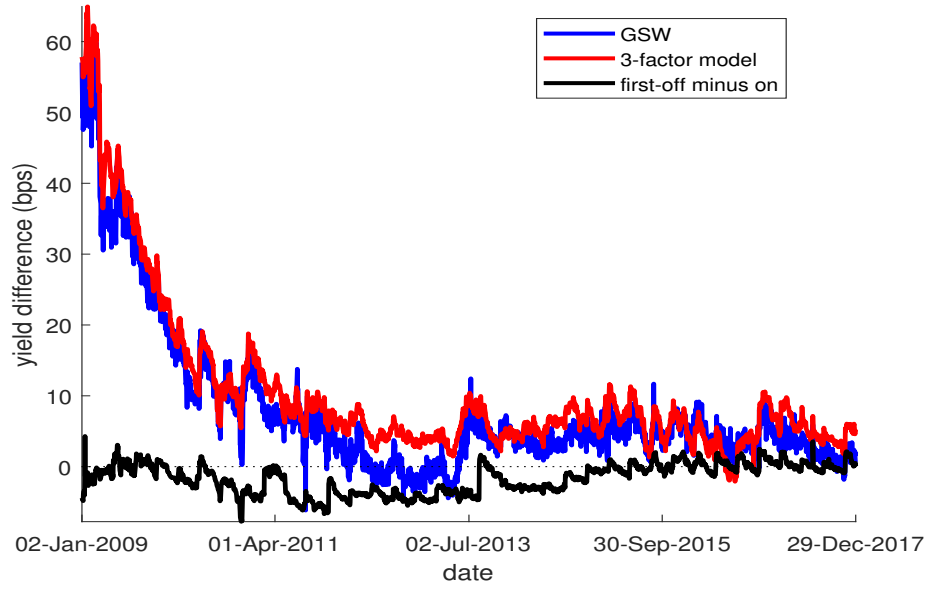

(a) On-the-Run Yield Premia, 10-year

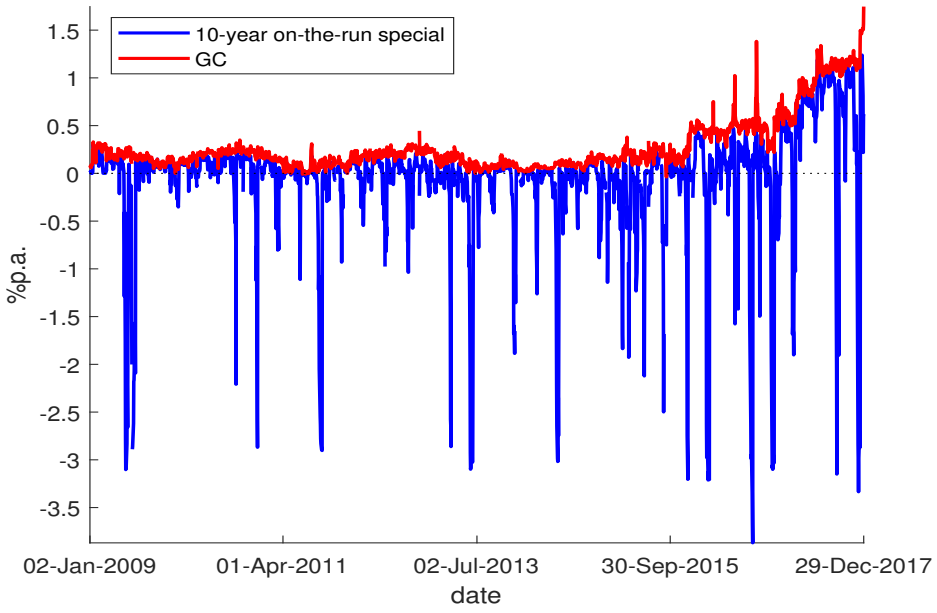

Figure 1. On-the-Run Premia and Special Rates
The top panel plots the difference between the yields to maturity of the 10-year on-the-run bond and the 10-year first off-the-run bond (black line), a synthetic bond with the same cash-flows as the 10-year on-the-run bond, according to the Gürkaynak, Sack and Wright (2007) estimated yield curve (blue line), and according to our estimated 3-factor model (red line), whose parameters are reported in Panel A of Table 2. Each line subtracts the on-the-run yield from the alternative yield. The bottom panel plots the GC repo (red line) rate along with the special rates for 10-year on-the-run bonds (blue line) over time, in annualized percent.

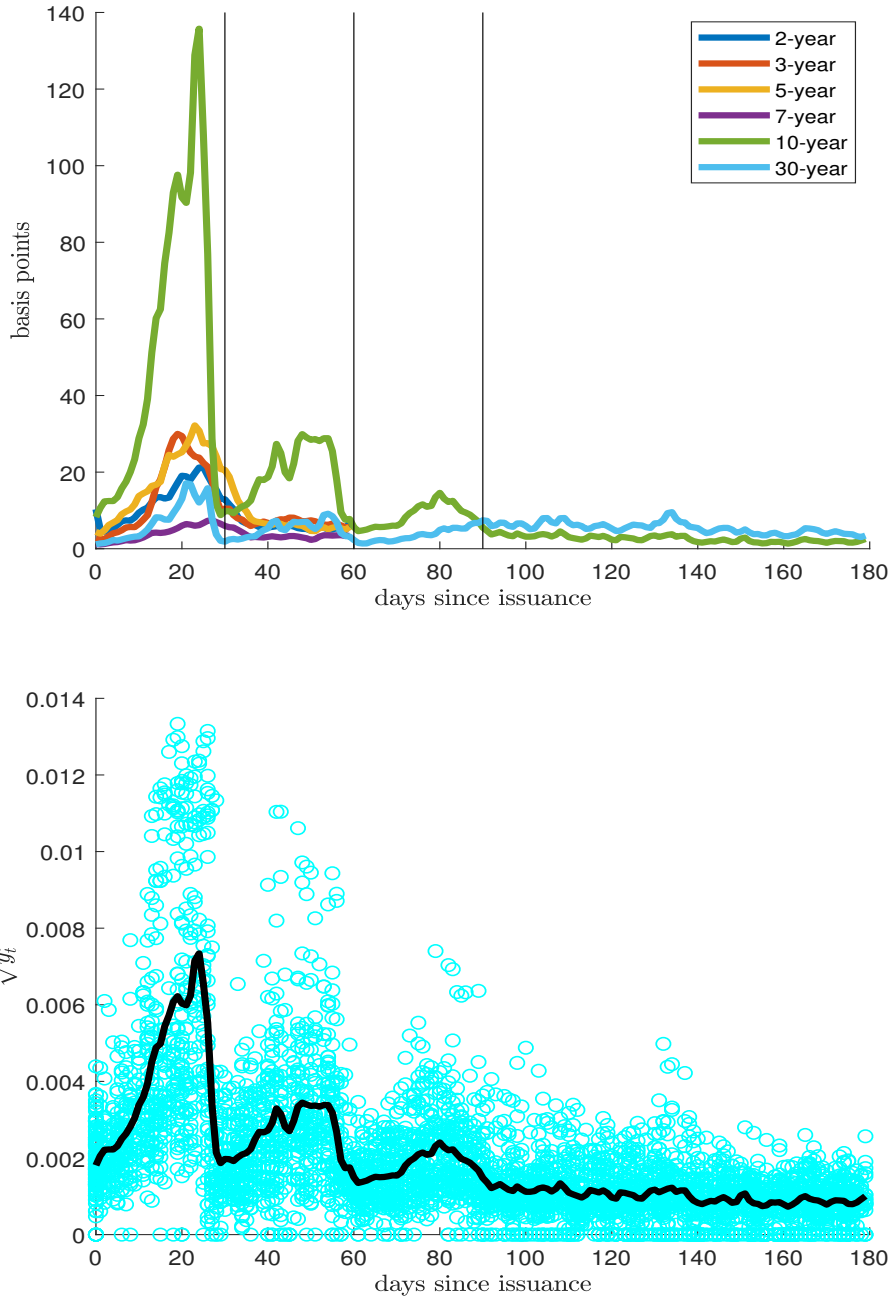

Figure 2. Special Spreads over the Auction Cycle
The top panel plots average special spreads, in annualized basis points, as a function of days since issuance for Treasury notes and bonds of all maturities. The vertical lines at 30-, 60-, and 90-days indicate the timings of the next auction (for 2-, 3-, 5-, and 7-year notes) and of the first and second re-openings, and the next auction, for 10-year notes and 30-year bonds. The bottom panel plots the observed square root of the special spread (not annualized) $\sqrt{y_t^i}$ for the 10-year note as blue circles against its estimated deterministic auction-cycle component y_τ^D as a function of time since issuance τ.

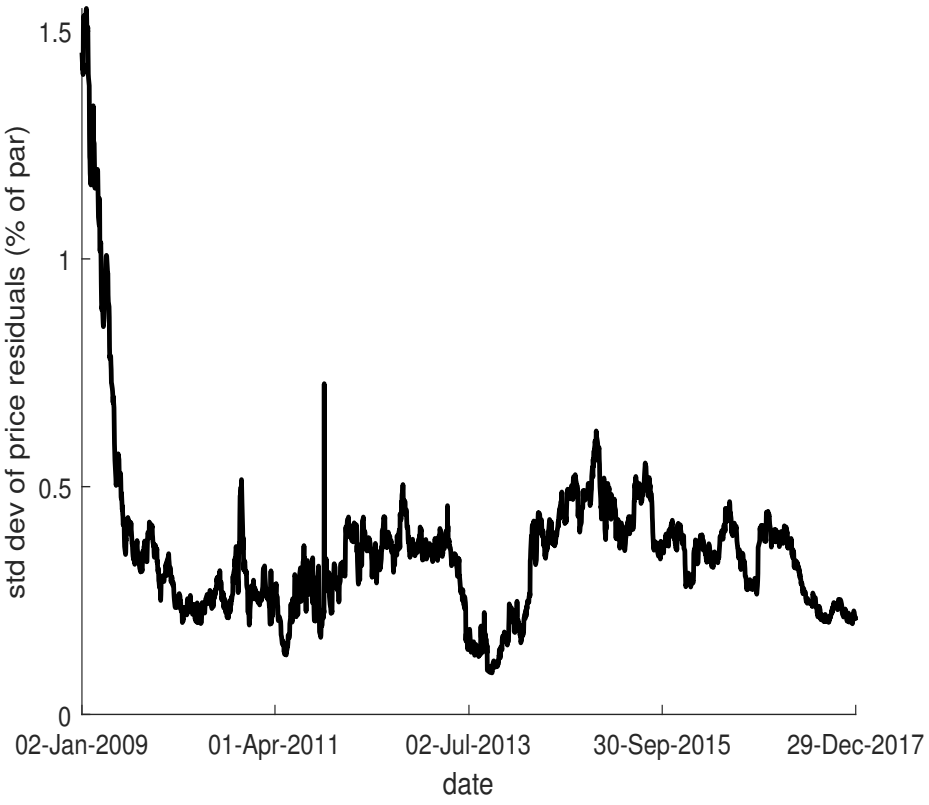

Figure 3. The figure plots the standard deviation of the residuals from equation (16) for each date in the sample, estimated only on off-the-run bonds and setting all special spreads to zero. The parameters of this 3-factor model are reported in Panel A of Table 2. The values are reported in percent of par.

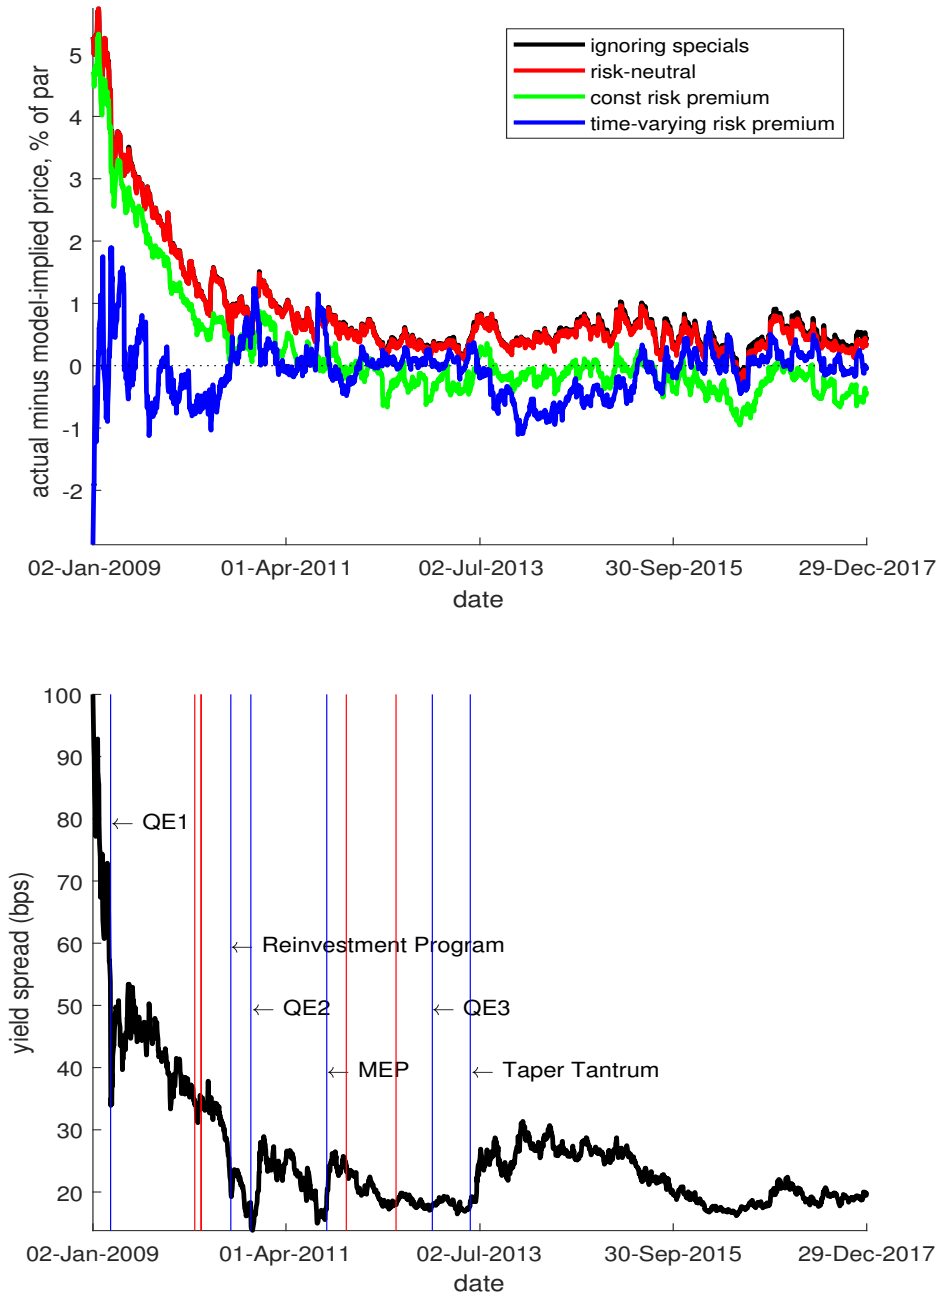

Figure 4. The top panel plots the price residuals from equation (16) on the 10-year on-the-run bond for four different estimates of the 10-year only model: a model ignoring special spreads (black), a model with risk-neutral special spreads (red); a model with a constant risk premium on the repo factor (green); and a model with time-varying risk premia tied to the special factor (blue). The bottom panel plots the repo risk premium, i.e. the estimated yield spread between two 10-year zero-coupon bonds that are special at the aggregate special spread y_t^S for 180 days, one that prices the special spread in a risk-neutral way and the other using our time-varying repo risk premium. The blue vertical lines mark several major economic news announcements, and the red vertical lines denote 10-year Treasury auctions with bid/cover ratios greater than two standard deviations away from the average.

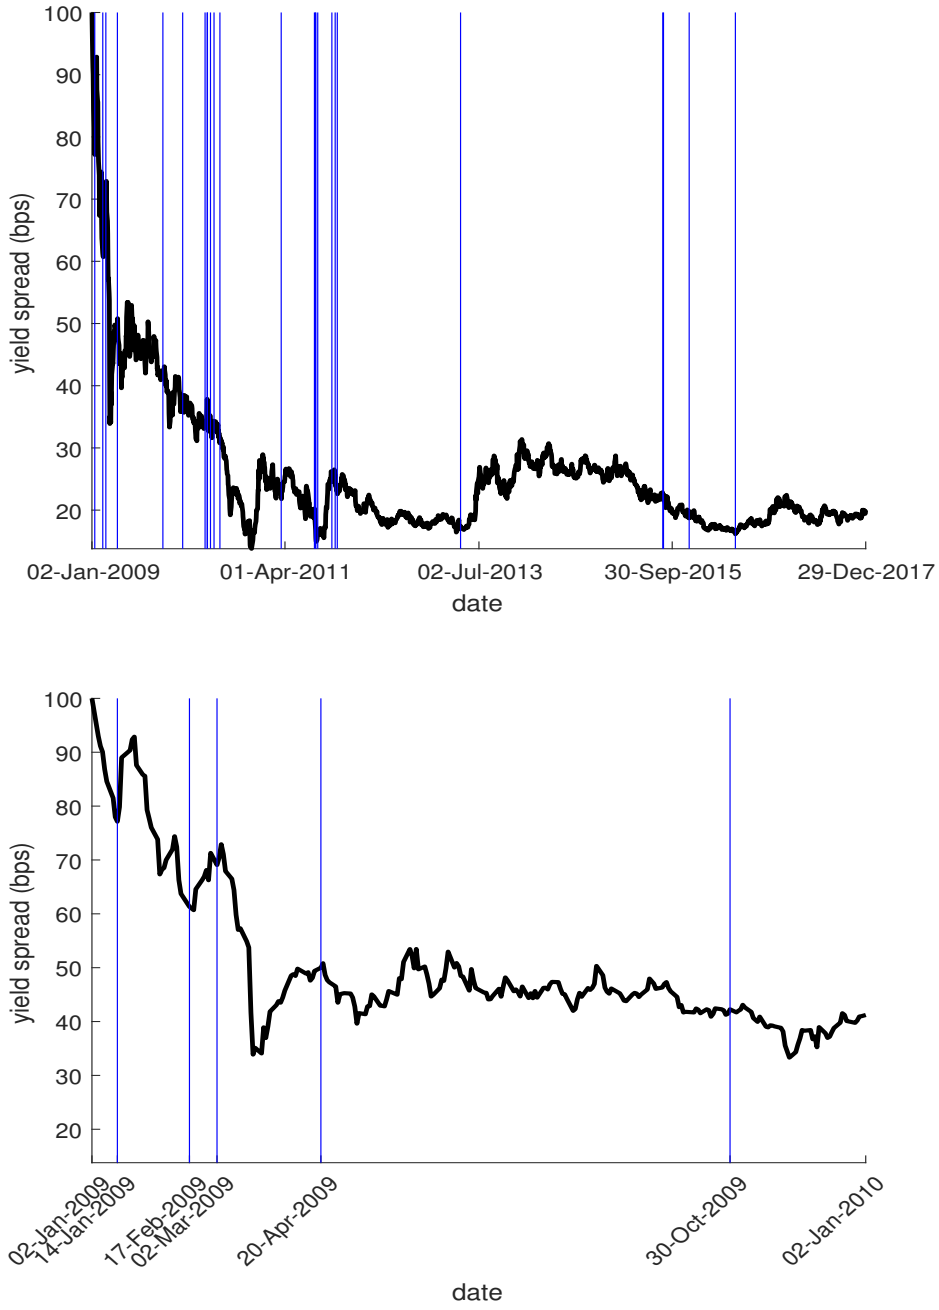

Figure 5. The top panel plots the repo risk premium, i.e. the estimated yield spread between two 10-year zero-coupon bonds that are special at the aggregate special spread y_t^S for 180 days, one that prices the special spread in a risk-neutral way and the other using our time-varying repo risk premium. The blue vertical lines mark flight to safety episodes, i.e. dates with at least a 5-basis-point increase in VIX, S&P500 index loses more than 1% of value, and a decline of at least 5 basis points in the 10-year Treasury yield. The bottom panel zooms in on theyear 2009.

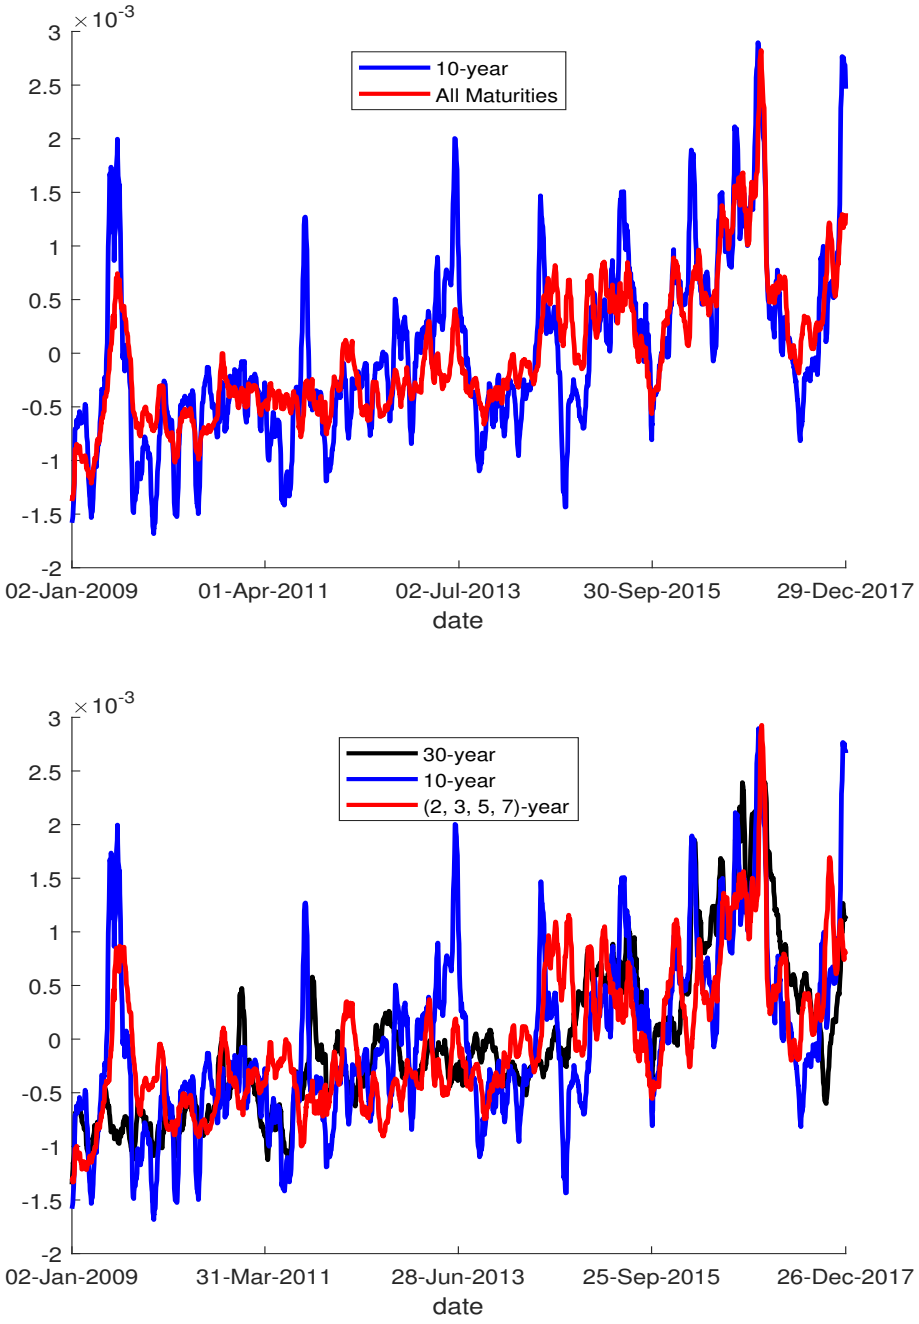

Figure 6. Both panels plot the estimated repo factors y_t^S, defined in equation (10), over time. To make the lines easier to distinguish each figure plots 15-day moving averages of y_t^S. The top panel plots y_t^S for the two 4-factor models: the red line plots y_t^S including special spreads on the 10-year on-the-run and first off-the-run bonds, while the blue line plots y_t^S including special spreads on on-the-run and first off-the-run bonds of maturity at issuance of 2, 3, 5, 7, 10, and 30 years. The bottom panel plots the 3×1 vector y_t^S over time for the model in which the first element of y_t^S averages over 30-year bonds (blue line), the second element averages over 10-year on-the-run bonds (red line), and the third element averages over 2-, 3-, 5-, and 7-year bonds (black line).

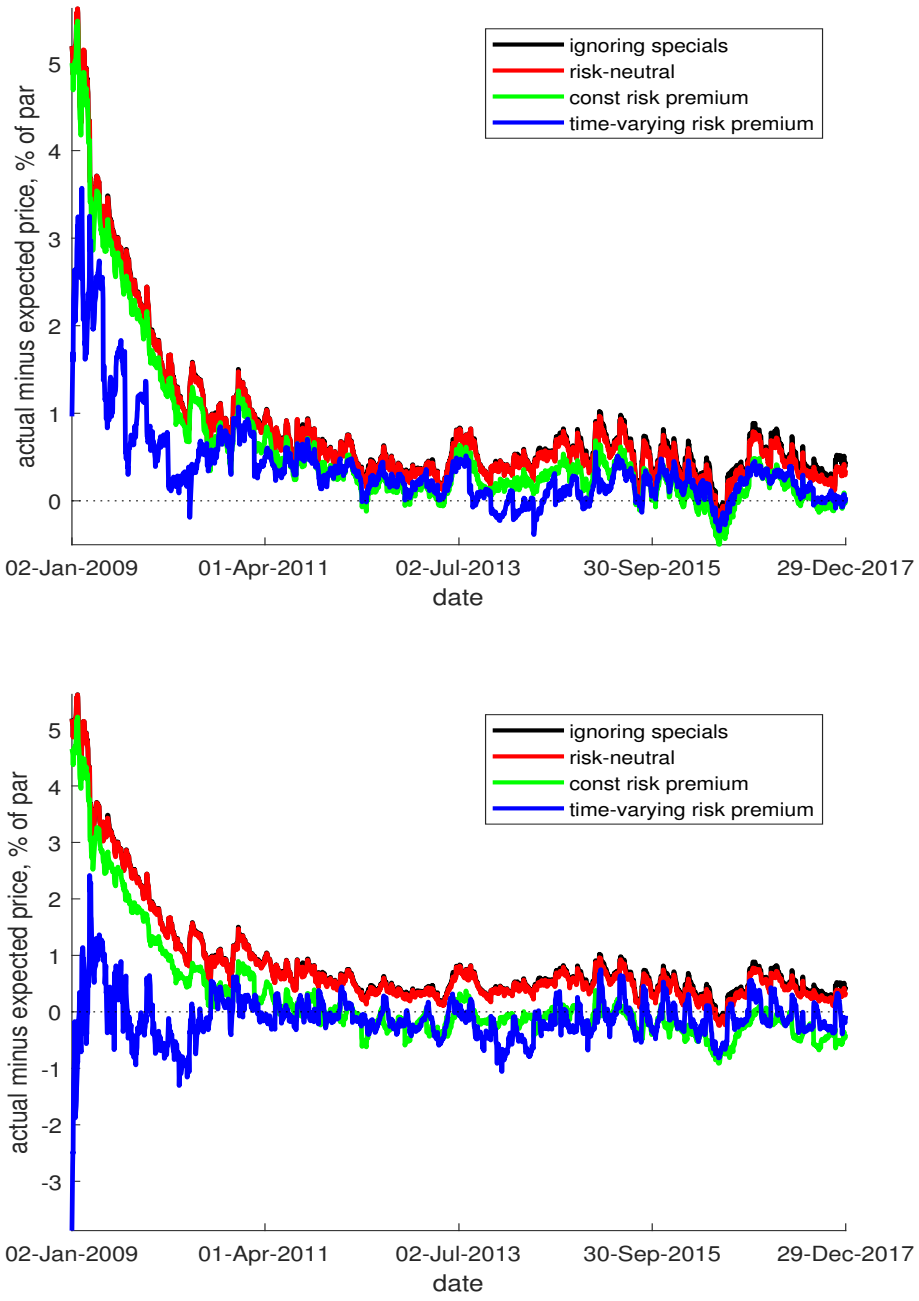

Figure 7. Both plot the price residuals from equation (16) on the 10-year on-the-run bond for four different estimates: a model ignoring special spreads (black), a model with risk-neutral special spreads (red); a model with a constant risk premium on the repo factor (green); and a model with time-varying risk premia tied to the special factor (blue). The top panel uses the "pooled" model, whose parameters are reported in the nlast three columns of Panel B of Table 2, while the bottom panel uses the "split" model whose parameters are reported in Table 3.

Table 1. Sample Summary Statistics
The table reports summary statistics for our sample from January 2, 2009 to December 29, 2017, by maturity at issuance in years (first column). The second column reports the average number of bonds in each cross-section with the indicated maturity at issuance. The third column reports the number of off-the-run bonds that are on special, i.e. have a positive y according to equation (4). The fourth and fifth columns report the average specialness spread, in bps, for off- and on-the-run bonds, respectively. The last column reports the average price residual (actual minus expected), in percent of par, of on-the-run bonds from the 3-factor model without repo specials. Numbers in parenthese are standard deviations.

Maturity (years)	Avg # Bonds	% On Special (off-the-run)	Avg Spread (off-the-run)	Avg Spread (on-the-run)	Avg Price Residual (on-the-run)
All	220	84.7	4.79 (6.16)	19.7 (41.6)	0.167 (0.608)
2	11.2	87.3	5.37 (8.8)	20.5 (38.7)	0.0546 (0.361)
3	21.7	88.9	5.72 (7.6)	21.1 (36)	0.073 (0.247)
5	46.9	84.5	3.84 (6.09)	24.5 (41.7)	-0.022 (0.239)
7	47.8	86.4	4.85 (5.47)	6.18 (11)	-0.0625 (0.396)
10	35.2	83.2	3.44 (3.69)	35.4 (66.5)	0.885 (0.95)
30	58.6	82.4	5.86 (6.43)	8.97 (23.2)	0.0681 (0.522)

Table 2. The table reports the estimated parameters of the 3- and 4-factor models discussed in the text. Panel A reports the 3-factor model parameters. The first three elements of μ^* are normalized to zero, and the top-left corner of Σ is normalized to be $I/365$ so that it is identity at an anual frequency. Panel B reports the repo-factor parameters for two versions of the model with non-zero repo-specials: in the first three, the 4th factor incorporates special spreads on 10-year special bonds (on-the-run and first off-the-run) only, while in the second three, all on-the-run and first off-the-run maturities are incorporated into the 4th factor. For each model, the table reports a risk-neutral version (columns 1 and 4), a model where the fourth element of μ^* is allowed to differ from μ (columns 2 and 5), and a model where the fourth row of Φ^* is allowed to differ from Φ (columns 3 and 6).

Panel A: 3-Factor Model Parameters				
δ_0	δ_1	Φ^*_{TL}		
-6.6335e-06	-0.00098895	0.99974	-0.00012053	0
	0.0006596	0.0021676	0.99982	0
	0.00030939	-0.00067494	3.4403e-05	1.0001

Panel B: 4-Factor Model Parameters						
Parameter	10-Year Only			Pooled		
	risk-neutral	constant	time-varying	risk-neutral	constant	time-varying
ρ	0.795	0.795	0.795	0.77	0.77	0.77
σ_x	0.000382	0.000382	0.000382	0.000377	0.000377	0.000377
μ^*	0.000192	0.00214	0.0589	9.71e-05	0.00165	0.0176
Φ^*_{BL}	-0.00809	-0.00809	0.534	-0.011	-0.011	0.189
	0.00289	0.00289	0.0121	0.00545	0.00545	0.0307
	0.000849	0.000849	-0.181	0.00216	0.00216	-0.0559
Φ^*_{BR}	0.713	0.713	-0.98	0.62	0.62	0.27
Σ_{21}	-1.56e-05	-1.56e-05	-1.56e-05	9.74e-06	9.74e-06	9.74e-06
	1.45e-05	1.45e-05	1.45e-05	1.81e-06	1.81e-06	1.81e-06
	-1.07e-07	-1.07e-07	-1.07e-07	-4.28e-06	-4.28e-06	-4.28e-06
Σ_{22}	0.000666	0.000666	0.000666	0.000451	0.000451	0.000451

Table 3. The table reports estimated repo parameters for the 6-factor model, in which y_t^S is a 3×1 vector of average special spreads: an average over 30-year special bonds, 10-year special notes, and 2-, 3-, 5-, and 7-year special notes, separately. Panel A of Table 2 reports the values of δ_0, δ_1, and Φ^*_{TL} for this model. The first column reports the risk-neutral parameters, the second column reports the parameters of the model in which the 4th, 5th, and 6th elements of μ^* are allowed to differ from μ, and the third column reports the parameters of a model in which the 4th, 5th, and 6th rows of Φ^* are allowed to differ from Φ.

Parameter	risk-neutral	constant	time-varying
ρ	0.768	0.768	0.768
σ_x	0.000373	0.000373	0.000373
μ^*	0.000423	-0.000848	0.0894
	0.000297	0.00115	0.0932
	0.00011	-4.78e-05	0.0193
Φ^*_{BL}	-0.0133	-0.0133	1.06
	-0.0133	-0.0133	1.1
	-0.0116	-0.0116	0.229
	0.00427	0.00427	0.172
	0.00494	0.00494	0.171
	0.00574	0.00574	0.0381
	0.000885	0.000885	-0.252
	0.00151	0.00151	-0.268
	0.00225	0.00225	-0.0536
Φ^*_{BR}	0.671	0.671	3
	-0.144	-0.144	2.15
	-0.112	-0.112	0.436
	-0.015	-0.015	-1.96
	0.794	0.794	-0.955
	-0.00754	-0.00754	-0.424
	-0.0771	-0.0771	-3.54
	-0.0984	-0.0984	-3.56
	0.725	0.725	0.209
Σ_{21}	3.29e-05	3.29e-05	3.29e-05
	-4.12e-06	-4.12e-06	-4.12e-06
	1.91e-05	1.91e-05	1.91e-05
	1.71e-05	1.71e-05	1.71e-05
	1.09e-05	1.09e-05	1.09e-05
	-4.31e-06	-4.31e-06	-4.31e-06
	-1.18e-05	-1.18e-05	-1.18e-05
	-2.55e-06	-2.55e-06	-2.55e-06
	-3.87e-06	-3.87e-06	-3.87e-06
Σ_{22}	0.000504	0.000792	0.000792
	0.000475	0.000905	0.000905
	0.000384	0.000727	0.000727
	0.000657	0.000152	0.000152
	0.000429	3.96e-05	3.96e-05
	0.000464	0.000131	0.000131

Table 4. The table reports regression results from estimating $\eta_{i,t} = \alpha_i + \beta_1 y_t^i + \beta_2 \eta_{i,t-1} + \xi_{i,t}$, where $\eta_{i,t}$ are estimated from equation (16) with the 3-factor model ignoring all special spreads, and the special spread y_t^i is defined in equation (4). Special spreads in the regression are annualized, and the price residuals are in percent of par value. The table reports t-statistics, clustered at the CUSIP level, in parentheses.

y_t^i	0.611***	0.00535***	0.442***	0.00491***
	(8.302)	(4.765)	(6.523)	(4.139)
$\eta_{i,t-1}$		0.989***		0.986***
		(844.2)		(628.3)
R^2	0.019	0.984	0.012	0.978
Observations	496,420	495,792	496,420	495,792
CUSIP FE	NO	NO	YES	YES
Number of CUSIP			628	628

Table 5. The table reports the share of the variation of on-the-run price residuals explained by the three estimated models, the first ("10-Year") using a factor that incorporates special spreads from 10-year on-the-run and first off-the-run bonds; the second ("Pooled") using special spreads from on-the-run and first off-the-run bonds of all maturities in a single average y_t^S factor, and the third ("Split") using three separate y_t^S factors: one for 30-year bonds, one for 10-year bonds, and one for 2-, 3-, 5-, and 7-year bonds. The first four columns report the sum of squared residuals for four versions of each model: one ignoring repos completely (i.e. a 3-factor model), one in which special spreads are priced risk-neutral, one with a constant risk premium, and one with time-varying risk premia. The last three columns convert the sum of squares into R^2 using equation (19).

Maturity	Model	Sum of Squares				R^2		
		no repo	risk -neutral	constant	time -varying	risk -neutral	constant	time -varying
10	10-year	0.379	0.37	0.238	0.0423	0.024	0.374	0.889
	pooled	0.37	0.363	0.293	0.0977	0.02	0.21	0.736
	split	0.37	0.36	0.232	0.039	0.0272	0.375	0.895
30	pooled	0.106	0.106	0.113	0.13	0.000257	-0.0621	-0.227
	split	0.106	0.107	0.107	0.0597	-0.00813	-0.00611	0.437
2	pooled	0.0294	0.028	0.0257	0.0406	0.0491	0.127	-0.377
	split	0.0294	0.0278	0.0282	0.0256	0.0542	0.0421	0.13
3	pooled	0.0144	0.0136	0.0121	0.0304	0.0595	0.16	-1.11
	split	0.0144	0.0135	0.0137	0.0139	0.0658	0.0496	0.0349
5	pooled	0.0131	0.0137	0.0166	0.0207	-0.0461	-0.264	-0.576
	split	0.0131	0.0138	0.0136	0.0157	-0.0558	-0.0367	-0.199
7	pooled	0.0359	0.0367	0.0415	0.0373	-0.0231	-0.155	-0.0394
	split	0.0359	0.0368	0.0367	0.0389	-0.0265	-0.0229	-0.0852

Table 6. The table reports correlation coefficients and regression R^2 between the 10-year repo risk premium for the three models and other Treasury anomalies. The GSW on-the-run spread is defined as the difference in yield to maturity between the 10-year on-the-run note and the yield on a synthetic bond with the exact same cash flows, priced using the daily parameters in Gürkaynak, Sack and Wright (2007) and available on the Federal Reserve's website. The Off Note-Bond Spread is the average yield difference between notes and bonds of duration between 7 and 9 years. The TIPS-Treasury Puzzle is 10-year inflation swap rate minus the 10-year TIPS break-even rate. The TIPS liquidity premium is the difference between 10-year TIPS yields and the model-implied 10-year risk-free rate as estimated by D'Amico, Kim and Wei (2018). The Nominal Fitting Errors are the average absolute pricing residual from equation (16), applied to all bonds, using the parameters in Panel A of Table 2.

Anomaly	10-Year		Pooled		10-Year (split)		30-Year (split)	
	Corr	R^2	Corr	R^2	Corr	R^2	Corr	R^2
GSW 10-Year On-the-Run	0.91	0.83	0.92	0.84	0.85	0.73	0.88	0.78
Off Note-Bond Spread	0.79	0.62	0.78	0.61	0.8	0.64	0.79	0.62
TIPS-Treasury Puzzle	0.82	0.68	0.82	0.68	0.81	0.65	0.82	0.68
TIPS Liq Premium	0.76	0.58	0.74	0.55	0.66	0.43	0.73	0.54
Nominal Fitting Errors	0.6	0.36	0.6	0.37	0.63	0.4	0.66	0.43