Interest Rates Under Falling Stars

Michael D. Bauer and Glenn D. Rudebusch*
Federal Reserve Bank of San Francisco

July 1, 2019

Abstract

Macro-finance theory implies that trend inflation and the equilibrium real interest rate are fundamental determinants of the yield curve. However, empirical models of the term structure of interest rates generally assume that these fundamentals are constant. We show that accounting for time variation in these underlying long-run trends is crucial for understanding the dynamics of Treasury yields and predicting excess bond returns. We introduce a new arbitrage-free model that captures the key role that long-run trends play in determining interest rates. The model also provides new, more plausible estimates of the term premium and accurate out-of-sample yield forecasts.

Keywords: yield curve, macro-finance, inflation trend, equilibrium real interest rate, shifting endpoints, bond risk premia
JEL Classifications: E43, E44, E47

*Michael D. Bauer (michael.bauer@sf.frb.org), Glenn D. Rudebusch (glenn.rudebusch@sf.frb.org): Economic Research Department, Federal Reserve Bank of San Francisco, 101 Market Street, San Francisco, CA 94105. We thank Anna Cieslak, Todd Clark, John Cochrane, Greg Duffee, Robert Hodrick, Lars Svensson, Jonathan Wright, and seminar participants at various institutions for helpful comments; Elmar Mertens and Mike Kiley for their r^* estimates; and Simon Riddell and Logan Tribull for excellent research assistance. The views in this paper are solely the responsibility of the authors and do not necessarily reflect those of others in the Federal Reserve System.

1 Introduction

Researchers have made numerous attempts to connect macroeconomic variables to the term structure of interest rates using a variety of approaches ranging from reduced-form no-arbitrage models to fully fledged dynamic macro models.[1] Despite both theoretical and empirical progress, there is no clear consensus about how macroeconomic information should be incorporated into yield curve analysis. One important link between the macroeconomy and the yield curve that has been largely overlooked is the connection between the time-varying long-run means of macroeconomic series and interest rates, that is, between their stochastic trends. Specifically, macroeconomic data and models can shed light on estimates of the trend in inflation (π_t^*) and the equilibrium real interest rate (r_t^*), while finance theory—from the Fisher equation to modern no-arbitrage theory—implies that such macroeconomic trends must be reflected in interest rates. Of course, the key outstanding issue is whether there is empirically significant variation in these long-run trends over time. Almost all term structure analyses in finance assume constant long-run means. Instead, in this paper, we document that accounting for changes in π_t^* and r_t^* is essential for modeling the term structure, estimating bond risk premia, and forecasting yields. We quantify the importance of time-varying macroeconomic trends for interest rates in two steps. First, using a variety of empirical proxies for these shifting endpoints, we establish some stylized facts about macroeconomic and interest rate trends. One of these stylized facts is the importance of such trends for predicting excess bond returns. Second, to account for these stylized facts, we introduce a new dynamic term structure model that embeds a stochastic interest rate trend. This no-arbitrage model provides new estimates of the long-run stochastic trend in Treasury yields and the term premium and accurately predicts interest rates out of sample. This analysis bridges an important gap between macroeconomics and finance by incorporating shifting macro-finance trends.

A simple illustration of the potential importance of macro trends for U.S. interest rates is provided in Figure 1. The underlying drivers of the apparent secular decline in longer-term Treasury yields since the early 1980s remain contentious. Standard finance (no-arbitrage) models of the yield curve assume that interest rates are stationary, i.e., mean-reverting—see Kim and Wright (2005), Joslin et al. (2011), and Adrian et al. (2013), among many others. This mean reversion implies that expected short-term rates are invariably close to their historical average at horizons beyond a few years ahead. That is, the expectations component of long-term interest rates varies only modestly. Consequently, low-frequency variation in interest rates must be captured by the term premium, which is the residual difference between a long-

[1] See, among many others, Ang and Piazzesi (2003), Diebold et al. (2006), Bikbov and Chernov (2010), and Rudebusch and Swanson (2012). A detailed survey is Gürkaynak and Wright (2012).

term interest rate and its expectations component (e.g., Wright, 2011). Instead, our analysis stresses the important role of the time-varying trends underlying interest rates. To illustrate this role, Figure 1 displays observable proxies of these trends (described further in Section 2). Trend inflation, π_t^*, measured using long-horizon survey forecasts, declined almost 6 percentage points from the Volcker disinflation of the early 1980s to the late 1990s and then remained stable around 2 percent. In contrast, the trend or equilibrium real interest rate, r_t^*—shown here as the average of several estimates—changed little early on but declined notably over the past two decades, likely reflecting structural changes in the global economy. As a result, the equilibrium nominal short-term interest rate, $i_t^* = \pi_t^* + r_t^*$, which reflects movements in both underlying macro trends, accounts for much of the low-frequency movements in the ten-year Treasury yield through variation in the expectations component. This suggests that the early downward trend in π_t^* and the more recent decline in r_t^*—that is, an environment of falling stars—played a major role in the secular decline in nominal interest rates. This link between macro trends and interest rate trends is at the core of our analysis.

Our key contribution is to establish this link by formulating a no-arbitrage model that simultaneously models the long-run trend and the term premium in yields. To motivate our new empirical model, we first establish two stylized facts about macro-finance trends and yields using both existing and new empirical proxies for these trends.[2] The first stylized fact is that variation in both the inflation trend and the equilibrium real interest rate is responsible for the very high persistence of interest rates. It has long been recognized that nominal interest rates contain a slow-moving trend component, and Campbell and Shiller (1987) uncover a single common trend, which implies that interest rates across maturities are cointegrated. We show that this common trend is i_t^* and that it is driven by quantitatively important fluctuations in both underlying macroeconomic trends, π_t^* and r_t^*. Accounting only for the inflation trend on its own, as in Kozicki and Tinsley (2001) and Cieslak and Povala (2015), leaves a highly persistent component of interest rates unexplained. We show that it is crucial to include movements in r_t^* to fully capture the trend component in interest rates.

Our second stylized fact is that macro stochastic trends are crucial for understanding bond risk premia. Relative to the canonical predictive regressions for excess bond returns that forecast future excess bond returns with current yields (Fama and Bliss, 1987; Cochrane and Piazzesi, 2005), we find that adding π_t^* and r_t^* as regressors provides notable gains in predictive power. Our findings provide evidence against the *spanning hypothesis*, which is

[2]Our estimates of r_t^* include existing estimates from Laubach and Williams (2003, 2016), Kiley (2015), Holston et al. (2017), Del Negro et al. (2017), and Johannsen and Mertens (2016, 2018). In addition, we create four new estimates, ranging from a simple moving-average of the real rate to a multivariate estimate that includes variables commonly viewed as drivers of r_t^*.

implied by most asset pricing models and posits that the current yield curve contains all the relevant information for predicting future interest rates (e.g., Duffee, 2013; Bauer and Rudebusch, 2017; Bauer and Hamilton, 2018). The predictive gains are economically large, so, for example, when interest rates adjust to a lower trend, long-term bond holders benefit substantially, just as they did during the adjustment to a lower r_t^* following the 2008 financial crisis. Indeed, including the equilibrium real rate is especially important later in our sample when the inflation trend shows less variation.[3]

To account for these two stylized facts, we introduce a new no-arbitrage dynamic term structure model (DTSM) with a shifting endpoint for interest rates. This novel specification parsimoniously models the Treasury yield curve using four state variables: three "yield factors" which are simply linear combinations of yields, and one common stochastic trend, i_t^*, which captures movements in the underlying macroeconomic trends. The existence of a common trend in interest rates stands in stark contrast to the usual stationary DTSMs in which yield factors follow a stationary VAR(1). Furthermore, in our model the trend is *unspanned*, so it cannot be backed out from the cross section of yields. This specification is consistent with our evidence that accounting for i_t^* leads to better predictions of bond returns. In addition, it allows for an estimated, unobserved i_t^* that—as in an unobserved components model—does not follow Treasury yields in lockstep. For robustness, we employ two different estimation methodologies, which provide similar empirical results. First, we estimate the model using an observed proxy for i_t^* in order to carry out a model-based analysis using the same real-time interest rate trend we employed in establishing the stylized facts about macro-finance trends. Second, we use Bayesian estimation of our model to infer i_t^* using only Treasury yields, imposing a certain degree of smoothness on our trend estimate, similar to Del Negro et al. (2017).

We then show that simulated data from these estimated models are able to match the stylized facts that we established about the role of trends for the yield curve. The model-based estimates of the trend and cycles components of interest rates are consistent with the cointegration regression results from the actual data. Also, the model accounts for the additional predictive power for future bond returns from the trend underlying interest rates that was evident in our historical data sample. Specifically, predictive regressions using simulated data from our model demonstrate notable increases in R^2 when the interest rate trend is added as a predictor—just as in the actual data.

[3]The improvement we document using trends to predict excess bond returns dovetails with recent research showing predictive gains from including slow-moving averages of past inflation, real output growth, and consumption growth as predictors (e.g., Cieslak and Povala, 2015; Brooks and Moskowitz, 2017; Garg and Mazzoleni, 2017; Jorgensen, 2017)—variables that are closely related to π_t^* and r_t^*.

We also use our shifting-endpoint model to investigate two important applications of DTSMs: (1) decomposing long-term interest rates into an expectations component and a term premium and (2) forecasting interest rates. For the former, our model provides markedly different decompositions of long-term rates than the usual stationary DTSM. A stationary DTSM implies a quite stable expectations component and attributes most of the secular decline in interest rates to the residual term premium, as discussed in critiques by Kim and Orphanides (2012) and Bauer et al. (2014). Instead, our shifting-endpoint model decomposition attributes the majority of the secular decline to the decrease in i_t^*. Consequently, our term premium estimates exhibit only a modestly decreasing trend and more pronounced cyclical swings, similar to other risk premia in asset prices (Fama and French, 1989). As for out-of-sample interest rate forecasting, previous researchers have found it surprisingly difficult to consistently beat the simple random walk forecast, which predicts that a future yield will equal the current yield. But we find that forecasts implied by our shifting-endpoint DTSM are generally more accurate than this martingale benchmark and also better than long-range projections from the Blue Chip survey of professional forecasters. By contrast, the standard stationary DTSM (corresponding to the special case of our model with a constant i^*) forecasts substantially worse than the random walk. The improvements in forecast accuracy from allowing a slow-moving shifting endpoint are both economically and statistically significant, and they are consistent with yields returning over time to their underlying macro trend.[4]

Our analysis is related to several different strands of the macro and finance literatures. Some studies have described a link between the inflation trend and nominal yields (Kozicki and Tinsley, 2001; van Dijk et al., 2014; Cieslak and Povala, 2015), but this leaves unexplained the continuing decline in yields over the past 20 years. Time variation in r_t^* has so far received little attention in finance, despite growing evidence of a sizable decline in the equilibrium real interest rate.[5] Our paper comprehensively documents the empirical importance of both relevant macro trends, π_t^* and r_t^*, for the dynamics of the yield curve.

No-arbitrage models of the yield curve typically rule out time-varying trends by specifying that the dynamics of yields are stationary. However, such models are intrinsically ill-suited

[4]In related work, van Dijk et al. (2014) documented forecast improvements relative to a random walk by including shifting endpoints based on their proxy of π_t^*. We demonstrate how to achieve even greater forecast gains in an arbitrage-free DTSM framework that accounts for time-varying π_t^* and r_t^*.

[5]The economic forces reducing the equilibrium real interest rate likely include lower productivity growth, changing demographics, a decline in the price of capital goods, and strong precautionary saving flows from emerging market economies, which have tended to increase global savings, reduce desired investment, and push down the steady-state real interest rate. Discussions include Summers (2014), Kiley (2015), Rachel and Smith (2015), Carvalho et al. (2016), Hamilton et al. (2016), Laubach and Williams (2016), Johannsen and Mertens (2016, 2018), Christensen and Rudebusch (2017), Del Negro et al. (2017), Holston et al. (2017) and Lunsford and West (2017). In macroeconomics, r_t^* is often labeled the neutral or natural rate of interest although, as noted below, there can be subtle dfferences among various definitions.

to capture the extreme persistence of interest rates, and a variety of remedies have been suggested.[6] Our results suggest that allowing for shifting long-run trends may be a better solution. The idea of a shifting-endpoint term structure model goes back to Kozicki and Tinsley (2001), who specified a restrictive two-factor DTSM that imposed the expectations hypothesis and identified i_t^* with a distant-horizon forward rate (with a unit root under both the risk-neutral and the real-world measure). Some macro-finance models of the yield curve allow for changes in the inflation trend but still assume a constant equilibrium real rate (Hördahl et al., 2006; Rudebusch and Wu, 2008). The model of Campbell et al. (2017), which is designed to capture the changing comovement of stock and bond returns, is an asset pricing model with a time-varying π_t^* and constant r^*. A few no-arbitrage models allow for changes in r_t^*, but they also make strong assumptions about this trend by tightly linking it to π_t^* (Dewachter and Lyrio, 2006) or trend output growth (Dewachter and Iania, 2011). Our model allows both stochastic macroeconomic trends to affect term structure dynamics. We also depart from the existing literature by treating i_t^* as unspanned, which allows this trend to exhibit some variation independent of contemporaneous yields.

Three recent papers propose statistical models of interest rates with time-varying long-run means, and are thus closely related to ours. Del Negro et al. (2017) estimate a Bayesian common-trend VAR for interest rates and inflation that allows for time variation in π_t^* and r^*. We use their estimates of r_t^*, along with others, in our analysis establishing stylized facts. Johannsen and Mertens (2018) also propose a time series model for yields and a Bayesian estimation framework that allows for stochastic volatility, and we also include their real-time r_t^* among our trend proxies. Finally, Crump et al. (2018) also allow for shifting macroeconomic trends in a common-trend VAR for interest rates and inflation with i_t^*, r_t^*, and π_t^* pinned down from surveys of professional forecasters. In contrast to the models in these papers, our model is an arbitrage-free representation that jointly captures the time series and cross section of Treasury yields. Such a specification allows us to investigate the role of trends for the yield curve and the term premium and to estimate expected bond returns.

The paper is structured as follows. Section 2 describes key concepts and our initial proxies for the macroeconomic trends. Section 3 presents stylized facts about the macroeconomic trends and Treasury yields. Section 4 introduces our no-arbitrage model, shows how it captures the stylized facts, and presents novel estimates of the term premium as well as out-of-sample forecast results. Section 5 concludes. An Appendix contains supplementary material.

[6] Proposed remedies include adding survey information (Kim and Wright, 2005; Kim and Orphanides, 2012), restricting risk prices (Joslin et al., 2014; Bauer, 2018), correcting for small-sample bias (Bauer et al., 2012, 2014), and using near-integration or long memory processes (Jardet et al., 2013; Goliński and Zaffaroni, 2016).

2 Interest rate trends: theory and estimates

The role of trends for interest rates can be illustrated with the usual decomposition of any n-period yield, $y_t^{(n)}$, into expectations of short-term rates and a term premium, $TP_t^{(n)}$:

$$y_t^{(n)} = \frac{1}{n}\sum_{j=0}^{n-1} E_t i_{t+j} + TP_t^{(n)} = i_t^* + \frac{1}{n}\sum_{j=0}^{n-1} E_t i_{t+j}^c + TP_t^{(n)}, \tag{1}$$

where i_t is the nominal short rate, $y_t^{(1)}$.[7] The term premium compensates bond investors for the duration risk in longer-term bonds, though in practice, it also captures liquidity and financial market frictions. The expectations theory of interest rates assumes that the term premium is zero (strong form) or constant (weak form), whereas modern finance theories of interest rates explicitly model the time variation in the term premium, as described in Section 4. The second equality in equation (1) introduces the trend or equilibrium short rate, i_t^*, and the short rate cycle $i_t^c = i_t - i_t^*$. We define i_t^* as the Beveridge-Nelson trend in the short-term nominal interest rate,

$$i_t^* \equiv \lim_{j\to\infty} E_t i_{t+j}. \tag{2}$$

We employ this stochastic trend concept throughout this paper and assume no deterministic time trends. All nominal yields are affected by changes in i_t^*, and if this trend is time-varying, then yields are nonstationary with a common trend. In this case, Kozicki and Tinsley (2001) fittingly call i_t^* a "shifting endpoint." The Fisher equation suggests that two macroeconomic trends determine this endpoint:

$$i_t = r_t + E_t\pi_{t+1} \quad \Rightarrow \quad i_t^* = r_t^* + \pi_t^*,$$

where r_t is the real short rate, π_t is inflation, and the trends in r_t and π_t are defined analogously to i_t^*. If inflation or the real interest rate contain long-run trend components, then nominal interest rates will also be driven by these same macroeconomic trends.

The remainder of this section (with further details in Appendix A) describes the empirical proxies of these trends that we employ to establish stylized facts about macro-finance trends in interest rates. For the long-run inflation trend, the prevailing consensus in empirical macroeconomics models inflation as an $I(1)$ process—in order to produce accurate forecasts and correctly capture the evolution of inflation expectations (Stock and Watson, 2007; Faust and Wright, 2013). From a macroeconomic perspective, the resulting time-varying π_t^* can be

[7] This intuitive definition of the term premium is common in the macroeconomic literature but ignores the convexity of bond prices. The model-based term premium estimates in Section 4.5 account for this convexity.

viewed as the perceived inflation target of the central bank. Empirical proxies for trend inflation have been constructed from surveys, statistical models, or a combination of the two—see, for example, Stock and Watson (2016) and the references therein. We employ a well-known, widely used survey-based measure of long-run inflation expectations, the Federal Reserve's series on the perceived target rate (PTR) for inflation (e.g., Clark and McCracken, 2013; Del Negro et al., 2017). Figure 1 shows that, from the beginning of our sample to the late 1990s, this estimate generally mirrored the rise and fall of the ten-year yield.[8] Since then, however, it has remained close to 2 percent, which has been the explicit longer-run inflation goal of the Federal Reserve since 2012. Given our focus on forecasting, real-time availability is an important issue, and for most of our sample (since 1979) the PTR estimate was available in real time from the Survey of Professional Forecasters.

As a proxy for r_t^*, we use a variety of estimates to avoid reliance on one specific model or estimation approach. All of these estimates are consistent with our definition of r_t^* as the stochastic trend in the real short rate, analogous to equation (2). The first two r_t^* estimates are the long-run real rate trends from time series models identified with Bayesian methods: Del Negro, Giannone, Giannoni, and Tambalotti (2017)(DGGT) estimate a linear state space model with common trends, and Johannsen and Mertens (2016, 2018)(JM) employ a nonlinear state space model with a shadow rate. Three additional estimates are based on simple New Keynesian macro models and use the Kalman filter to infer the *neutral* real interest rate, that is, the level of the real rate consistent with real output at potential and inflation at target: Laubach and Williams (2003, 2016) (LW) implemented this concept empirically using the Rudebusch and Svensson (1999) model, Holston, Laubach, and Williams (2017) (HLW) use a slight simplification of the LW model, and Kiley (2015) augments the model to account for changes in financial conditions. The implementation of these macro models assumes that the neutral rate follows a martingale, thus their r_t^* estimates are consistent with the long-run concept we employ. We also construct three model-based estimates of our own. The first is based on a simple univariate unobserved components model, and the second uses the same model augmented with moving averages of real GDP and labor force growth that appear to be correlated with the real rate trend (Lunsford and West, 2017). The third instead uses the observed nominal short rate and inflation in a state-space model with a latent real rate and r_t^*. We estimate these models using Bayesian methods, and like Del Negro et al. (2017) we employ a prior with a low innovation variance for r_t^*. Finally, as a simple, model-free estimate

[8] Other estimates of the inflation trend, such as the long-range forecasts in the Blue Chip survey, the moving-average estimate of Cieslak and Povala (2015), or various model-based estimates (Clark and Doh, 2014; Stock and Watson, 2016) exhibit a similar pattern.

of r_t^*, we include an exponentially-weighted moving average of past real rates.[9]

Figure 2 summarizes these r_t^* estimates. The left panel shows the average and range of the smoothed (two-sided) estimates, which use the full data sample for inference about the parameters and state variables.[10] These show a relatively stable r_t^* until the late 1990s, but then exhibit a pronounced decline. To better capture what bond market investors could have known about macro trends contemporaneously, our analysis will focus on filtered (one-sided) and real-time estimates of r_t^*, which are shown in the right panel of Figure 2. Filtered estimates—from the LW, HLW, and Kiley models—infer latent state variables at time t using data available up to time t. However, they also use model parameter estimates from the full sample, which would not be available to investors in real time.[11] Therefore, to avoid the pitfall of look-ahead bias in our analysis of predictive power, we consider (pseudo) real-time estimates that only use data through period t to construct r_t^*. Such real-time estimates are available for the JM and DGGT models, our own three models, and a moving-average of past real rates.[12] As shown in Figure 2, the average of these six real-time estimates increased from about one to 2 percent early in the sample, moved sideways for about 15 years, and then around 2000 started a substantial, steady decline until the end of the sample. This pattern closely follows the movements of just the moving-average estimate of r_t^*, which has the added benefit of not relying on a model or estimation technique unavailable to investors historically. Finally, the average of all nine filtered and real-time r_t^* estimates, which is displayed in Figure 1, also shows a similar pattern.

Substantial estimation and model uncertainty underlie estimates of any long-run trends, as suggested by the range of estimates in Figure 2. To reduce the uncertainty inherent in individual estimates, we take averages across models. While this approach is in no way a cure-all against estimation uncertainty and measurement error, model-averaging is well-known to improve forecast accuracy (Timmermann, 2006).[13] Our baseline estimate of i_t^* for the analysis in Section 3 will be the sum of the PTR estimate of π_t^* and the average of all six real-time estimates of r_t^*.[14]

[9] Our smoothing parameter of 0.98 follows other research estimating macro trends (Malmendier and Nagel, 2011; Orphanides and Williams, 2005).

[10] We have seven smoothed estimates—for all models but the HLW model; see Appendix A.

[11] Still, Laubach and Williams (2016) argue that, for the past decade, truly real-time r_t^* estimates are similar to their filtered estimates.

[12] These estimated trends all rely on revised inflation data, and the multivariate model also uses revised output data. Although the use of real-time data vintages can be important for near-term forecasting, it is arguably less relevant for estimation of long-run trends.

[13] We discuss the issue of measurement error in the Appendix, arguing that our results are likely a lower bound for the tightness of the connection between the yield curve and the *true* underlying macro trends.

[14] An alternative proxy for i_t^* could be obtained from surveys, such as the long-range forecast for the three-month T-bill rate from the Blue Chip Financial Forecasts. This series, which is only available semiannually

3 Macro trends and yields: stylized facts

3.1 The common trend in nominal yields

We first investigate whether macroeconomic trends account for the high persistence of interest rates. We then assess whether accounting for just the inflation trend is sufficient or if a real rate trend also materially contributes to movements in nominal bond yields. Our proxies for these macroe trends are the PTR survey-based estimate of π_t^* and the three alternative estimates of r_t^*: the average of three filtered estimates, the average of the six real-time estimates, and a simple (exponentially weighted) moving average of past real rates. We proxy for the equilibrium nominal short rate i_t^* with the sum of π_t^* and the real-time average estimate of r_t^* (but obtain similar results with the filtered or moving-average measures).

We find that changes in macro trends account for the persistence of interest rates, and the trend and cycle components of interest rates can be identified using our proxy i_t^*. This is illustrated in Figure 3, which shows the (demeaned) ten-year yield and two detrended versions of that yield.[15] The first detrended yield is the simple difference between the yield and i_t^*, which should be stationary if the yield is indeed driven by this trend and the loading on i_t^* is unity. In comparison to the ten-year yield, this series exhibits much less of a trend, for example, declining by less than half as much as the ten-year yield over the period from the early 1980s to the end of our sample. Still, a notable trend remains in this simple difference, which indicates that the yield trend component moves more than one-for-one with i_t^*. By definition, the expectations component of yields moves one-for-one with i_t^*, so the term premium must contain a trend component positively related to i_t^*. To account for this relationship, the second detrended series is the residual from a (Dynamic OLS) cointegration regression of the yield on i_t^*, discussed below. This cointegration residual shows no discernible trend, indicating that changes in i_t^*, properly scaled, can fully capture the trend in the ten-year Treasury yield. The fitted value of the cointegration regression is the trend component of the ten-year yield, for which we also obtain a model-based estimate in Section 4.3, and the residual—shown in Figure 3—corresponds to the interest rate cycle.

It has long been known that nominal interest rates are highly persistent.[16] Summary statistics and unit root tests reported in Appendix B.1 confirm that Treasury yields are best modeled as having a stochastic trend in our sample. We examine whether this persistence

starting in 1986, fluctuates around our baseline estimate of i_t^*. We consider survey forecasts in Section 4.6.

[15] We use end-of-quarter zero-coupon Treasury yields from Gürkaynak et al. (2007) with maturities from one to 15 years and three-month and six-month Treasury bill rates from the Federal Reserve's H.15 release, over the period from 1971:Q4 to 2018:Q2.

[16] See, for example, Nelson and Plosser (1982), Rose (1988), and King et al. (1991).

is driven by a common macro trend. Figure 3 suggests a likely cointegration relationship with a loading on the trend larger than unity. To formally investigate this relationship, including the individual contributions of π_t^* and r_t^*, we use cointegration regressions of the form $y_t^{(40)} = \beta_0 + \beta_1' X_t^* + u_t$, where $y_t^{(40)}$ is the ten-year yield, the vector X_t^* contains one or two macro trend proxies, and u_t is a stationary residual under cointegration. We estimate the coefficients β_0 and β_1 using the Dynamic OLS estimator of Stock and Watson (1993), by including four leads and lags of $\Delta y_t^{(40)}$ and ΔX_t^*.[17]

Table 1 shows the results of our cointegration analysis, including coefficient estimates, Newey-West standard errors (using six lags) and the regression R^2. The first regression specification includes only π_t^*, and the results echo those of Cieslak and Povala (2015), who used simple OLS regressions but a similar inflation trend proxy. The inflation trend coefficient is somewhat higher than one and strongly significant, and the R^2 is very high. However, trend inflation does not fully account for the trend behavior of yields, as evident from the persistence of the cointegration residuals $\hat{u}_t = y_t^{(40)} - \hat{\beta}_0 + \hat{\beta}_1 \pi_t^*$. While $\hat{u}_t$ is much less persistent than the ten-year yield, it is still quite persistent. As shown in the second column, for this series, unit root tests—the Augmented Dickey-Fuller (ADF) and Phillips-Perron (PP) tests—do not reject the null, and the low-frequency stationary test of Müller and Watson (2013) rejects stationarity at the 5 percent level.[18] The inadequacy of detrending with only π_t^* is also supported by a Johansen (1991) trace test, which does not reject the null of absence of cointegration among the yield and the inflation trend.[19] Finally, an error-correction model, where $\Delta y_t^{(40)}$ is regressed on $\hat{u}_{t-1}$ and four lags each of the first-differenced yield and the inflation trend, indicates only modest equilibrating forces pushing yields toward the inflation trend.

The missing piece driving the level of the yield curve is the real rate trend. The remaining columns of Table 1 show that there is strong evidence for cointegration once r_t^* is added. For all three r^* proxies, the estimated coefficients are at or above one and strongly significant. Adding the real rate trend increases the R^2 by 8 to 11 percentage points, and the cointegration residual now appears to be stationary. For the real-time and moving-average estimates of r_t^*, the Johansen test strongly suggests a cointegration rank of one, and there is substantial error-correction behavior of the yield in response to deviations from trend. The last column of Table 1 shows that i_t^* can fully capture the trend component in the ten-year yield, since the

[17] In Appendix B.1, we report results for the level of the yield curve, which are similar to those for the ten-year yield. Also, standard OLS as well as the reduced-rank maximum likelihood VAR estimation of Johansen (1991) lead to the same conclusions as Dynamic OLS (not reported).

[18] For the unit root tests of cointegration residuals, we omit intercepts in the ADF and PP regressions and use the critical values provided by Phillips and Ouliaris (1990), which vary with the number of regressors. Further details are in Appendix B.1.

[19] The test is based on a VAR representation with four lags, based on information criteria.

evidence for cointegration is as strong as when both macro trends are included individually. The cointegration coefficient is 1.67, suggesting a positive loading of the term premium on i_t^*, a result that will be captured by our model in Section 4.3. The error-correction coefficient is estimated as -0.45 and is strongly significant. Accordingly, when the ten-year yield is high relative to i_t^*, it tends to revert quickly toward this trend, with almost half of the equilibrium error on average reversed within one quarter. Therefore, knowledge of the macro trend appears quite useful when predicting future yields. In sum, these results indicate that the low-frequency behavior of Treasury yields is explained by movements in a common underlying macroeconomic trend, i_t^*, in which both π_t^* and r_t^* play important roles.

3.2 Predicting excess bond returns with macro trends

We now consider whether macroeconomic trends can help predict future excess bond returns. The one-quarter return on a bond with maturity n in excess of the risk-free rate is

$$rx_{t+1}^{(n)} \equiv p_{t+1}^{(n-1)} - p_t^{(n)} - y_t^{(1)} = -(n-1)y_{t+1}^{(n-1)} + ny_t^{(n)} - y_t^{(1)},$$

where $p_t^{(n)}$ denotes the log-price of a zero-coupon bond. Since the findings in Fama and Bliss (1987) and Campbell and Shiller (1991), it is well-known that the yield curve, particularly its slope, contains useful information for predicting excess bond returns. The key question is whether the current yield curve contains *all* of the information relevant for predicting future returns, that is, whether the *spanning hypothesis* holds. Several studies (including Ludvigson and Ng, 2009; Joslin et al., 2014; Cieslak and Povala, 2015) have documented apparent violations of the spanning hypothesis using various additional predictors. Bauer and Hamilton (2018) demonstrated that accounting for the small-sample econometric problems associated with these highly persistent predictors renders most of them insignificant. However, they found that the trend inflation proxy of Cieslak and Povala (2015) was a relatively robust predictor. We investigate, in turn, the gains from including π_t^*, r_t^*, and i_t^* in predictive regressions for excess bond returns. We predict the one-quarter excess return averaged across maturities from two to 15 years.[20]

Table 2 reports results for six different predictive regressions. The first column is the usual baseline specification that includes only a constant and the first three principal components

[20] We use one-quarter holding periods to avoid the econometric problems resulting from overlapping observations. The one-quarter returns are calculated with the usual approximation $y_{t+1}^{(n-1)} \approx y_t^{(n)}$. We have obtained similar results with annual returns (not reported). In addition, Appendix B.2 provides estimates for individual bond maturities.

(PCs) of yields.[21] The specification in the second column adds π_t^*, and columns three through five include π_t^* as well as one of our three proxies of r_t^* in order to capture the effects of both trends. The final column adds only i_t^* instead of two separate macro trends. We report White standard errors as well as small-sample p-values for the macro trends using the parametric bootstrap of Bauer and Hamilton (2018) to obtain accurate small-sample tests even though regressors are persistent and correlated with lagged forecast errors.[22] The top panel uses the full sample of data, while the bottom panel considers a subsample from 1985 to 2018 for robustness. There is some evidence of a shift in the conduct of U.S. monetary policy in the mid-1980s, which would likely lead to a change in interest rate dynamics, so such a subsample is often considered for yield curve analysis (e.g., Rudebusch and Wu, 2007; Joslin et al., 2014).

In the full sample, the inclusion of the inflation trend increases the predictive power substantially compared with using only yield curve information. Both the inflation trend and the level of yields (PC1) are highly significant (in addition to the slope, PC2), as in Cieslak and Povala (2015). However, adding r_t^* to the regressions leads to further impressive gains in predictive power. Notably, for the moving-average estimate of r_t^*, the R^2 jumps, the coefficients and significance for π_t^* and PC1 rise, and the coefficient on r_t^* is large and significant.[23] Not surprisingly given Figure 1, the relative predictive power of the two macro trends changes in the more recent period as variation in r_t^* has become more pronounced. In the subsample starting in 1985, the inflation trend by itself is not statistically significant. Only with the addition of the equilibrium real rate do both trends matter for bond risk premia: the coefficients on π_t^* and PC1 more than double, the R^2 increases substantially, and the coefficients on π_t^* and r_t^* are statistically significant. These results imply that the trend in the real interest rate is as important as, and recently more important than, the trend in inflation.

The inflation and real-rate trends have coefficients with the same sign and roughly equal magnitudes. This suggests that they play similar roles for yield dynamics and that their sum, i_t^*, can capture the relevant information. Indeed, as shown in the final column of Table 2, including just i_t^* as a predictor provides the same strong predictive gains. The coefficient on i_t^* is of course negative: if the trend component increases then interest rates also rise in response, producing losses for bond holders.[24] The coefficient is also larger in magnitude than

[21] We scale the PCs such that they correspond to common measures of level, slope, and curvature, as in Joslin et al. (2014). For example, the loadings of yields on PC1 sum to one.

[22] We simulate 5000 bootstrap samples for yields and predictors under the spanning hypothesis, using separate (bias-corrected) VAR(1) models for yield factors and predictors. The small-sample p-values are the fractions of simulated samples in which the t-statistics of the macro trends are at least as large (in absolute value) as in the actual data.

[23] Changes in R^2 can be difficult to interpret in the presence of highly persistent predictors, so Appendix B.2 confirms these results using our bootstrap.

[24] Numerically, an increase in i_t^* by 0.1 percentage point predicts a decrease in the next quarter's excess bond

the coefficient on the yield curve level (PC1), consistent with our findings in Sections 3.1 and 4 that the trend component of yields moves more than one-for-one with changes in i_t^*.

An earlier literature showed that a high yield curve *slope* predicts positive excess bond returns. Cochrane and Piazzesi (2005) also found additional predictive power in other aspects of the shape of the yield curve—captured in mid-curve forward rates or higher-order PCs. Our findings, by contrast, document the importance of the *level* of the yield curve as a powerful predictor. While in yields-only regressions, such as specification (1) in Table 2, the level appears irrelevant, it becomes a strong predictor once i_t^* is accounted for. This result implies that the level of interest rates matters for expected future bond returns, but that it needs to be viewed relative to the underlying yield trend and not a constant historical mean. Furthermore, accounting for i_t^* makes the slope of the yield curve an even more powerful predictor.

In Table 3, we further investigate the importance of viewing interest rates relative to their stochastic trend by using detrended yields as predictors. The use of such predictors with low persistence has the added benefit that conventional, large-sample inference is more reliable. We use the first three PCs of the 17 detrended yields as predictors, where each detrended yield is obtained as the cointegration residual from a regression of the yield on macro trends. As before, we consider three types of detrending: using only π_t^*, using both π_t^* and (real-time) r_t^*, as well as using i_t^*. In the full sample, all three detrending methods lead to substantial predictive gains compared to the baseline regression with just the first three PCs of the 17 original yields. Notably, the detrended level of the yield curve (PC1) becomes strongly significant and the most powerful predictor. Detrending with i_t^* works particularly well, giving the largest gain of 11 percentage points. In the post-1985 sample, the R^2 does not increase when detrending with only π_t^*, but it increases by 9 and 6 percentage points, respectively, when detrending with either both π_t^* and r_t^* or with only i_t^*. These estimates provide strong support for the importance of macro trends in predicting bond returns. Furthermore, the similarity of the original and detrended results confirm that the predictive gains are indeed due to trend-reversion of interest rates (i.e., Granger-causation from underlying macro trends to future yields).

In sum, the key to predicting bond returns and estimating bond risk premia is to account for the long-run trend in interest rates. Both underlying macro trends have strong incremental predictive power for bond returns relative to the observed yield curve. That is, the long-run trend cannot be inferred from the cross-section of yields alone, so our results contradict the spanning hypothesis. An alternative method to uncover their predictive power is to detrend interest rates, but it is crucial to use i_t^* instead of π_t^* for that detrending.

return by 0.45 percentage point. Appendix B.2 shows that this coefficient rises with bond maturity.

4 A no-arbitrage model with a stochastic trend

To account for the stylized facts establishing an important low-frequency macro-finance interplay, we introduce and estimate a new DTSM. While existing DTSMs generally assume stationarity, our new formulation incorporates a shifting endpoint for interest rates, i_t^*, as a common stochastic trend for yields. This novel specification allows us to explain the stylized facts, provide new estimates of the term premium, and more accurately forecast interest rates.

4.1 Model specification

The state variables of our model of yield dynamics include N linear combinations of yields in the vector P_t. We use the loadings for the first three principal components of observed yields to construct three factors P_t ($N = 3$) as is commonly done (e.g., Joslin et al., 2011).[25] The key feature of our model is the presence of a single stochastic trend τ_t common to the three factors. The common trends representation of the dynamic system is:

$$P_t = \bar{P} + \gamma\tau_t + \tilde{P}_t, \quad \tau_t = \tau_{t-1} + \eta_t, \quad \tilde{P}_t = \Phi\tilde{P}_{t-1} + \tilde{u}_t, \tag{3}$$

where γ is an N-vector with the loadings on the common trend, η_t is *iid* normal with variance σ_η^2, $\tilde{P}_t$ are the N cycle components that have *iid* normal innovations $\tilde{u}_t$ with covariance matrix $\widetilde{\Omega}$, and Φ is a mean-reversion matrix that has eigenvalues with modulus less than one so that $\tilde{P}_t$ is stationary. Accordingly, the state variables $Z_t = (\tau, P_t')'$ are cointegrated, and Appendix C.1 gives more details and alternative representations of the dynamic system.[26] We will assume that innovations to trend and cycle are orthogonal, $E(\eta_t\tilde{u}_t) = 0$, an assumption commonly made in unobserved components models (Watson, 1986; Del Negro et al., 2017). Our assumptions imply that the long-run trend components of P_t are

$$P_t^* \equiv \lim_{j\to\infty} E_t P_{t+j} = \bar{P} + \gamma\tau_t. \tag{4}$$

The short-term interest rate is taken to be an affine function of P_t:

$$i_t = \delta_0 + \delta_1' P_t. \tag{5}$$

Finally, we assume absence of arbitrage opportunities so that there exists a "risk-neutral" probability measure, denoted by $\mathbb{Q}$, which prices all financial assets (Harrison and Kreps,

[25] Let the vector Y_t contain J model-implied yields. Then $P_t = WY_t$, where W is an $N \times J$ coefficient matrix with rows containing the first three eigenvectors of the sample covariance matrix of observed yields.

[26] Innovations to P_t are $P_t - E_{t-1}P_t \equiv u_t = \gamma\eta_t + \tilde{u}_t$ with covariance matrix $\Omega = \gamma\gamma'\sigma_\eta^2 + \widetilde{\Omega}$.

1979). The risk-neutral dynamics are specified as

$$P_t = \mu^{\mathbb{Q}} + \Phi^{\mathbb{Q}} P_{t-1} + u_t^{\mathbb{Q}}, \tag{6}$$

where the innovations $u_t^{\mathbb{Q}}$ are *iid* normal (under $\mathbb{Q}$) with covariance matrix Ω. We assume that P_t is stationary under $\mathbb{Q}$ by imposing that $\Phi^{\mathbb{Q}}$ has eigenvalues less than one in absolute value, thus forward rates, yields and bond convexity remain finite in the limit as maturity increases. Appendix C.2 provides additional details, including the stochastic discount factor and prices of risk that are implied by equations (3)–(6), and explains how stationarity under the real-world measure is reconciled with a unit root under the risk-neutral measure. Equations (5) and (6) imply that yields are affine functions of the yield factors P_t, $Y_t = A + BP_t$. The loadings in A and B are determined by the risk-neutral parameters and given in Appendix C.3.

Two types of normalizations are required for the state variables and parameters to be identified. The first set of normalizations, required in any affine term structure model, uniquely pins down the yield factors P_t. In their absence, an affine transformation of P_t would leave implied interest rates Y_t unchanged (Dai and Singleton, 2000). We use the normalization of Joslin et al. (2011), which imposes restrictions only on the $\mathbb{Q}$-dynamics and is ideally suited for the use of linear combinations of yields as yield factors. In this normalization, only four parameters determine the $\mathbb{Q}$-dynamics: a scalar $k^{\mathbb{Q}}$ and an N-vector $\lambda^{\mathbb{Q}}$ containing the eigenvalues of $\Phi^{\mathbb{Q}}$ which we assume to be real, distinct, and less than one. These four parameters and Ω fully determine the yield loadings A and B. Details are in Appendix C.4.

The second normalization is needed to identify the unobserved trend τ_t. Without it, an affine transformation of τ_t would leave the yield factors P_t unchanged. For identification, we require one linear constraint on each of the vectors $\bar{P}$ (to fix the level) and γ (to fix the scale).[27] We choose to identify τ_t as the long-run mean of the short rate, $\tau_t = i_t^* = \delta_0 + \delta_1' P_t^*$, so that in light of (4) our normalizations are $\delta_0 + \delta_1' \bar{P} = 0$ and $\delta_1' \gamma = 1$. In total, the parameters of the model are $k^{\mathbb{Q}}, \lambda^{\mathbb{Q}}, \gamma, \bar{P}, \Phi, \Omega, \sigma_\eta^2$, and σ_e^2, where γ and $\bar{P}$ have two degrees of freedom.

The key innovation of this model relative to existing affine DTSMs is that the equilibrium nominal interest rate i_t^* is allowed to be time-varying. This implies that bond yields have a common stochastic trend and are $I(1)$. Imposing a unit root in this fashion is a convenient way to account for a very persistent process in a dynamic model, and it crystalizes the arguments regarding slow-moving interest rate trends. Taken literally, a unit root specification for interest rates is unlikely because nominal interest rates appear to have remained bounded over the centuries. However, in finite samples, a persistent stationary process can always be

[27] For further discussion of identification requirements in multivariate unobserved components (common trends) models, see, for example, Harvey (1990).

approximated arbitrarily well by a unit root process, and it is well-known that doing so can often be beneficial for forecasting (e.g., Campbell and Perron, 1991).[28] To illustrate the importance of the time-varying trend in interest rates, we will compare our model, which we refer to as the "shifting endpoint" (SE) model, to a conventional DTSM with stationary dynamics, a "fixed endpoint" (FE) model. Such a specification corresponds to a restricted special case of our model where $i_t^* = i^*$, i.e., $\sigma_\eta^2 = 0$, which is equivalent to the popular three-factor affine DTSM of Joslin et al. (2011).

Another interesting aspect of our model is that the long-run trend is *unspanned* by the yield curve, meaning that the cross section of interest rates at time t is not deterministically related to i_t^*. Section 3.2 documented strong evidence against the spanning hypothesis, which is implied by essentially all asset pricing models as agents can generally back out all relevant state variables from the yield curve (Duffee, 2013). Such inference is not possible in our model because the trend does not *directly* affect yields: conditional on P_t, yields do not load on i_t^*. Therefore, the model can match the evidence documented in Section 3.2, as we demonstrate below.[29] Another important rationale for an unspanned trend is to allow it to be unobserved. To incorporate an unobserved components specification within a dynamic term structure model, it is necessary for the trend component to be unspanned, because otherwise i_t^* would be completely determined by the time-t yield curve and thus be observed. Such a trend hardwired to current interest rates would also result in excessively volatile trend estimates.[30]

Our parsimonious specification can accurately fit the cross section of yields, is consistent with absence of arbitrage and time-varying risk premia, and includes a stochastic trend in line with much evidence showing that nominal interest rates are effectively $I(1)$. Importantly, changes in i_t^* may affect expectations of future short rates as well as the term premium, and our estimates below will assess the empirical importance of a trend in the term premium. While our DTSM framework is quite flexible, it is kept deliberately simple by focusing on the nominal yield curve instead of separately modeling inflation and real yields. Such a separation, in the spirit of Christensen et al. (2010), would be a useful future extension that could differentiate r_t^* and π_t^*, but it is unnecessary for our purposes of estimating nominal risk premia and forecasting nominal Treasury yields.

[28] For example, a unit root in inflation may not be literally true, but assuming so often improves forecast performance (Faust and Wright, 2013).

[29] Sufficiently flexible *spanned* DTSM, in which yields depend on all risk factors, can typically match the data equally well as the corresponding *unspanned* DTSM. However, such spanned models require many more parameters and are more difficult to estimate (Bauer and Rudebusch, 2017). Furthermore, this result applies only to models in which the unspanned factors are observed, as in Joslin et al. (2014), whereas we allow for the possibility that the trend is unobserved.

[30] We have estimated term structure models with spanned trend factors and obtained implausibly volatile trend estimates, as movements in long-term yields are tightly linked to movements in the trend factor.

4.2 Model estimation

The shifting trend model can be cast in state-space form for estimation. The measurement equation is $Y_t^o = A + BP_t + e_t$, with $e_t \overset{iid}{\sim} N(0, \sigma_e^2 I_J)$. For our data set, Y_t^o includes $J = 17$ yields. The state equation is the VAR representation for the state variables Z_t, as given in Appendix C.1.

The common trend in our DTSM is both unspanned and unobservable, unlike the observed unspanned factors in models like Joslin et al. (2014). Our specification is therefore effectively an unobserved components model. While this is desirable for matching the stylized facts of trends in interest rates, we face the usual estimation problem that empirically uncovering long-run trends using limited sample sizes is intrinsically difficult (Watson, 1986). For a very persistent time series with relatively infrequent business cycle fluctuations, even several decades of data typically contain limited information that can pin down the underlying trend. Simply estimating an unobserved components model with maximum likelihood and no external additional information is problematic and unreliable.[31] Therefore, in the literature, two broad approaches have been used to facilitate estimation of unobserved trends from time series data: the first includes additional data or proxies for the trend, the second incorporates prior information about properties of the trend such as its smoothness. For robustness, we consider each of these two methods for estimating our stochastic trend DTSM.

The first method adds data that can directly help pin down the trend estimate. For example, long-run survey expectations of inflation are often used to inform estimates of π_t^* (Kozicki and Tinsley, 2012; Del Negro et al., 2017). Similarly, Crump et al. (2018) use surveys of professional forecasters to estimate i_t^* in a factor model for the yield curve. In this spirit, we use the empirical proxy for i_t^* equal to the sum of the PTR estimate of π_t^* and our real-time estimate of r_t^*. Section 3 showed that this trend proxy captures low-frequency movements of yields and predicts excess bond returns. Because we can now treat both i_t^* and P_t as observed, estimation is greatly simplified.[32] We denote the resulting estimates as the "observed shifting endpoint" (*OSE*) model. To guard against the problems that the zero lower bound (ZLB) on nominal interest rates poses for the estimation of affine yield curve models (Bauer and Rudebusch, 2016), we estimate the parameters of the *OSE* model using a sample that ends in 2007:Q4, before the height of the financial crisis and the ZLB period. Once the parameter

[31] The likelihood is essentially flat in some dimensions of the parameter space, and numerical optimization leads to local optima with very different trend estimates, depending on starting values.

[32] Details are in Appendix C.7. Taking i_t^* as observed is justified because maximum likelihood estimation of a state-space model with an additional measurement equation for i_t^* results in a noise variance equal to essentially zero. Taking the yield factors P_t as observed was proposed and justified by Joslin et al. (2011). If estimation with the state-space model is desired, excellent starting values for the model's parameters can be obtained from a first-step estimation with observed state variables.

estimates are obtained, we proceed to use the model over the entire sample period.[33] We find that the *OSE* model fits observed yields well, with a root mean-squared error of 8 basis points over the full sample. The $\mathbb{Q}$-dynamics are highly persistent, with the largest eigenvalue of $\Phi^{\mathbb{Q}}$ equal to 0.997. For the real-world dynamics, the mean reversion to the common trend is relatively quick, with the largest eigenvalue of Φ equal to 0.933. More details on the estimation are given in Appendix C.7.

Our second estimation of the long-run trend model uses only Treasury yield data while making use of prior information in a Bayesian framework. In this regard, an effective approach to estimating unobserved trends imposes a tight prior distribution for the variance of the trend innovations. Such priors are commonly used in a variety of macroeconomic contexts to produce smooth trend estimates, consistent with the macroeconomic view that such trends are determined by gradual variation in structural factors such as demographic change. For interest rates, Del Negro et al. (2017) used such a smoothness prior for estimation of r_t^*. Similarly, we specify a tight inverse-gamma prior for σ_η^2 with a mean of $0.1^2/400$, implying that the standard deviation of the change in i_t^* over a century is 10 percent. In light of the observed variation in interest rates and our trend proxies, we view this as a conservative choice in the sense that it limits the amount of yield variation that can be attributed to the long-run trend. Also, we follow Del Negro et al. (2017) by combining this prior with the stationary version of a Minnessota VAR-prior for the cycle components $\tilde{P}_t$ with the distribution for the diagonal elements of Φ centered around zero. We call this representation the "estimated shifting endpoint" (*ESE*) model. We use a Markov chain Monte Carlo (MCMC) algorithm to simulate draws from the joint posterior distribution of the latent state variables and parameters. The remaining parameters have standard and mostly uninformative priors that are described in Appendix C.8 along with the MCMC algorithm and further details. To account for the ZLB, we treat short-term yields as missing observations when they were constrained.[34]

Figure 4 plots the resulting *ESE* model estimate of i_t^*—specifically, the posterior mean and 95 percent credibility intervals. The *ESE* model i_t^* rose until the early 1980s to over 8 percent and then gradually declined over the following decades to just below 3 percent at the end of our sample. For comparison, the ten-year Treasury yield and the proxy estimate of i_t^* are also plotted. The *ESE* model estimate matches the proxy fairly closely and can account for much

[33] An extension of our model to explicitly account for the zero lower bound, for example with a shadow-rate specification as in Kim and Singleton (2012) and Bauer and Rudebusch (2016), is conceptually straightforward but beyond the scope of this paper.

[34] Specifically, three- and six-month yields are assumed to be missing from 2008:Q4 to 2015:Q3, the one-year yield is omitted from 2011:Q4 to 2014:Q3, and the two-year yield is omitted from 2012:Q4 to 2013:Q1. The start dates of these periods are based on results in Swanson and Williams (2014) while the end dates are when the yields lifted off from the ZLB.

of the variation in the ten-year yield. As usual for unobserved trend models, the uncertainty around the estimated trend is quite large. From 2008 to 2015, the estimation intervals widen a bit as uncertainty increases because the yield curve is partially constrained by the ZLB and is less informative about the underlying trend.

4.3 Implications for trend components of yields

Section 3.1 showed that proxies for i_t^* were important drivers of yields, particularly in terms of cointegration. Now we interpret this evidence in the context of our model.

First, note that our shifting-endpoint DTSM can decompose yields into trend and cycle components:

$$Y_t = Y_t^* + \tilde{Y}_t = A + B\bar{P} + B\gamma i_t^* + B\tilde{P}_t, \quad Y_t^* \equiv \lim_{j\to\infty} E_t Y_{t+j}. \tag{7}$$

Figure 5 shows the (actual and model-fitted) ten-year yield and its trend and cycle components. The trend component accounts for most of the low-frequency movements in the yield, while the cycle component—the model-based analog of the regression-based interest rate cycles discussed in Section 3.1—exhibits clear mean reversion around zero and little to no apparent trend. This confirms that our model accounts for the long-run trend in interest rates.

Second, we can compare the model-implied coefficients of yields on i_t^* to the empirical results we obtained from the cointegration regressions of yields on proxies of i_t^* in Section 3.1. In the model, the trend i_t^* is unspanned and only indirectly affects yields through its relationship with P_t. Equation (7) shows that the implied loadings of yields on i_t^* are $B\gamma$. Figure 6 plots these coefficients, which gradually rise from unity at the short end—since i^* is by definition the trend component of the short rate—to around 1.8. In the data, the coefficients (shown with their 95 percent confidence intervals) are relatively flat around 1.6-1.9, with a slight hump around two years and a gradual decline thereafter. Note that the estimated coefficient for the ten-year yield corresponds to the value reported in the last column of the top panel of Table 1. To gauge whether the model is consistent with the estimated cointegration coefficients we need to account for sampling uncertainty, hence we simulate 5,000 artificial samples of the trend and yields with the same size as the actual data, and estimate the cointegration regressions in each simulated sample. The 95 percent Monte Carlo intervals of the coefficient estimates, shown as dashed lines in Figure 6, comfortably contain the values obtained with the actual data.[35]

Our results here and in Section 3.1 show that the trend in Treasury yields does not move

[35] Because innovations to the trend and cycles components of yields are orthogonal in the model, the estimates are unbiased, so that the mean of the Monte Carlo coefficient distributions equals the model-implied loadings $B\gamma$, up to simulation error.

one-for-one with changes in i_t^*, since the relevant coefficients are quite a bit larger than one. Because expectations of future short rates by definition load with a unit coefficient on i_t^*, coefficients above one indicate that the term premium positively responds to changes in i_t^*. For a yield of maturity n, the term premium loading on i_t^* is $B_n'\gamma - 1$. That is, our cointegration regressions and the model-based loadings suggest that the term premium responds to the trend component as well. In Section 4.5, we will show the resulting term premium estimates and quantify the relative importance of the trend in short-rate expectations and in the term premium. Our model has similar implications for the trend behavior of the slope of the yield curve. The slope, if measured as the simple difference between a long-term yield and the short rate, in fact has the exact same trend component and loading on i_t^* as the term premium in that long-term yield. While some macroeconomic time series models, such as DGGT, assume that separate trends drive short rates and the slope, in our no-arbitrage model the slope contains a trend due to the risk in i_t^*.

4.4 Model-based explanation of excess return predictability

A key stylized fact from Section 3 is that various measures of the long-run trends in interest rates have substantial predictive power for excess bond returns, and this additional information is not spanned by the first three principal components of yields. To explain this result with our shifting endpoint DTSMs, we first simulate 5,000 artificial samples from each of our three models: the *FE* model with stationary yields, the *OSE* model with an observed trend proxy, and the *ESE* model with an estimated trend.[36] We then run excess return regressions using the simulated data sets from the three models in order to replicate and understand the regression results from the actual data.

Table 4 presents the R^2 results from the actual and simulated data. The top row reports the R^2 for the predictive regressions of excess bond returns using actual data with and without i_t^* as an additional predictor, as described in Section 3.2 (corresponding to the estimates in the first and last column of the top panel of Table 4). The reported ΔR^2 shows the substantial predictive gains from including the macro trend proxy.

The remaining three rows provide results using data simulated from the *FE*, *OSE*, and *ESE* models. For the *FE* model, there are essentially no predictive gains on average from adding i_t^* to the yield principal components, with a median ΔR^2 of 1 percentage point. This is not surprising because this model imposes the null hypothesis that i_t^* truly has no predictive power. Given the persistent nature of yields and the trend and the econometric issues described by Bauer and Hamilton (2018), the *FE* model's 95 percent Monte Carlo interval for ΔR^2 is not

[36] For the *FE* model, we separately simulate i_t^* from a random walk process using the *OSE* parameters.

centered around zero. However, it remains the case that the upper bound of this interval (4 percentage points) is much lower than the actual improvement in R^2 from adding trends in the data (12 percentage points), consistent with our rejections of the spanning hypothesis in Section 3.2 using the bootstrap, which also simulated data under the null hypothesis.

For the *OSE* model, the predictive gains from adding i_t^* are large: the median gain of 9 percentage points is close to the ΔR^2 of 12 percentage points in the actual data. Importantly, the Monte Carlo interval comfortably contains the value found in the data. For the *ESE* model, the results are similar, with large model-implied predictive gains from including the trend as an additional predictor. Overall, the shifting-endpoint models accurately capture the predictive gains from adding i_t^* in regressions for excess bond returns.

In our shifting-endpoint model, shifts in i_t^* are unspanned by the yield curve by construction. Why are changes in the trend unspanned in the data? The intuitive reason is the following: Shifts in the level of the yield curve can occur because of either changes in the underlying trend (see the loadings in Figure 6) or because of high-frequency movements in interest rates relative to the trend, with opposite implications for expectations of future bond returns, and investors cannot distinguish between the two.[37] Therefore, yields and macro trends contain important separate pieces of information for predicting future interest rates.

To better understand the predictability of excess bond returns through the lens of our model, we consider analytically the model-implied expected excess bond returns. Leaving aside constant terms that are irrelevant for predictability, bond risk premia are

$$E_t rx_{t+1}^{(n)} = \mathcal{B}_{n-1}' \left(E_t P_{t+1} - E_t^{\mathbb{Q}} P_{t+1} \right) = \mathcal{B}_{n-1}' \left((I_N - \Phi^{\mathbb{Q}}) \gamma i_t^* + (\Phi - \Phi^{\mathbb{Q}}) \tilde{P}_t \right), \tag{8}$$

where $\mathcal{B}_n$ contains the affine loadings of the log bond price $p_t^{(n)}$ on the risk factors P_t. The derivation is given in Appendix C.3. The first expression shows the well-known result that time variation in bond risk premia is due to differences in expectations of future yields under the physical and risk-neutral probability measures. In the presence of a trend in yields, as in our model, physical-measure (real-world) expectations depend not only on current yields (P_t) but also on the shifting endpoint i_t^* (see equation (3) or the VAR representation in Appendix C.1). Hence, ignoring the trend misses an important part of the variation in $E_t P_{t+1}$ and understates the predictability of bond returns. The second expression in equation (8) shows the separate roles of trends and cycles for bond risk premia. The first term involves i_t^*, which demonstrates that bond risk premia—expected excess returns and, by extension, term premia—generally

[37] Even in a term structure model with a spanned trend, the cross-sectional loadings of yields on i_t^* would differ only slightly from the loadings of the level factor in detrended yields, and small measurement errors would make them effectively indistinguishable (Jorgensen, 2017).

contain a trend component.[38] The second term demonstrates the important role of the cyclical factors $\tilde{P}_t = W\tilde{Y}_t$, which are just linear combinations of *detrended* yields. Our results in Section 3.2 show that detrended yields capture essentially all of the predictability of excess returns, which indicates that the second term in the final expression of equation (8) is the more relevant one in practice and that the trend component in bond risk premia is small.

4.5 The term premium in long-term yields

Another perspective on the pricing of risk in Treasury bonds is given by the term premium in long-term interest rates.[39] As evident from the definition in equation (1), estimating it requires long-horizon projections of future short-term rates, which will depend crucially on the nature of the long-run trend in interest rates. Furthermore, because the term premium is the average of expected future excess bond returns, $TP_t^{(n)} = n^{-1}\sum_{j=0}^{n-2} E_t rx_{t+j+1}^{(n-j)}$, our results about the importance of trends for predictions of excess bond returns suggest that trends should matter for the estimation of term premia as well.

Our baseline for comparison is the *FE* model in which the state variables follow a stationary VAR. This special case of our no-arbitrage model with a constant endpoint i_t^* corresponds to the model of Joslin et al. (2011) and is representative of the conventional DTSM in the literature. The left panel of Figure 7 shows the five-to-ten-year forward rate as well as its expectations component, i.e., the "risk-neutral rate," and the term premium as estimated from the *FE* model.[40] The expectations component is estimated to be very stable, and it remains near the unconditional mean of the short rate, the fixed endpoint i^*. Therefore, the term premium, the residual, has to account for the trend in the long-term interest rate since the 1980s. As argued by Kim and Orphanides (2012) and Bauer et al. (2012), among others, such behavior by the expectations component and the term premium appears at odds with observed trends in survey-based expectations (Kozicki and Tinsley, 2001) and with the cyclical behavior of risk premia in asset prices (Fama and French, 1989). By allowing for a long-run trend in interest rates, we avoid such counterfactual decompositions of long-term rates into expectations and term premium. The right panel of Figure 7 shows the estimates for the risk-neutral rate and the term premium from our no-arbitrage *OSE* and *ESE* models, i.e., with either an observed or estimated shifting endpoint i_t^*. We also show (Bayesian) 95 percent credibility intervals for the risk-neutral rate obtained from our *ESE* model, which account for the estimation uncertainty for both the parameters and the unobserved trend. Short-rate

[38] Indeed, only if the model also contained a unit root under $\mathbb{Q}$ is it possible for the trend to drop out from risk premia.

[39] Appendix C.6 provides details on the model-implied term premium.

[40] Results for the ten-year yield are qualitatively similar.

expectations are anchored at i_t^*, and the cycle components $\tilde{P}_t$ mean-revert relatively quickly; hence, the expectations component of the five-to-ten-year forward rate closely mirrors this trend. For both shifting-endpoint models, movements in the expectations component account for the majority of the low-frequency variation in the observed interest rate. In particular, the estimates exhibit a pronounced downward trend since the 1980s due to the decline in i_t^*.

The forward term premium estimated from the shifting-endpoint models also exhibits some moderate low-frequency swings, with a decline from around 4-5 percent in the early 1980s to about 0-1 percent at the end of the sample. Because long-term interest rates exhibited more pronounced secular movements than the estimates of the long-run trend i_t^*, the estimated term premium also contains a trend component. As discussed above in Section 4.3, yields load on i_t^* with coefficients larger than unity, and the term premium consequently has a positive loading on this trend. Our estimates show that changes in the trend carry a positive risk premium. From a macro-finance perspective, some downward trend in the term premium since the 1980s is plausible. In large part, inflation risk declined substantially over this period as average inflation declined, and the term premium includes compensation for this risk (Wright, 2011).[41] But the trend in the term premium is much more modest and more plausible with a shifting rather than a fixed endpoint. Our shifting-endpoint models attribute a majority of the secular decline in interest rates not to the term premium but instead to the decline in long-run expectations of future short-term rates. The drop in the forward rate from its average during 1980-1982 to its average during 2015-2017 was close to 10 percentage points. According to the stationary *FE* model, the estimated term premium accounts for 75 percent of this decline. In contrast, the shifting-endpoint models attribute only 35 percent (*OSE*) and 37 percent (*ESE*) to the term premium, while attributing the majority to a falling expectations component, in line with the substantial downward shift in the underlying macro trends. Furthermore, the estimated term premium with a shifting endpoint exhibits more pronounced cyclical variation, in line with the notion that risk premia are countercyclical (Fama and French, 1993; Cochrane and Piazzesi, 2005). This stark difference demonstrates how accounting for the slow-moving trend component in interest rates fundamentally alters our understanding of the driving forces of long-term interest rates, and bridges the gap between the common wisdom of secular macroeconomic trends and no-arbitrage models for the yield curve.

Using a shifting-endpoint model with a slow-moving trend component solves the knife-edge problem of Cochrane (2007), who pointed out that assuming either stationarity or a random

[41] Furthermore, in theory, a decline in the term premium is consistent with the fall in the covariance between nominal bond returns and stock returns since the 1980s, as emphasized by Campbell et al. (2017). A decline in real risk premia would also be implied by a change in business cycle dynamics that increased persistent risk relative to transitory risk, as described by Campbell (1986) and Beeler et al. (2012).

walk for the level of interest rates leads to drastically different implications for expectations and term premia, in both cases at odds with widely held views about interest rates. Imposing a random walk for the level is known to forecast well but leads to the implausible implication that the expectations component accounts for essentially all variation in yields. Imposing stationarity, as shown in the left panel of Figure 7, leads to the implausible implication that the term premium component accounts for all of the low-frequency variation in yields. Our approach of assuming that interest rates contain a slow-moving common trend avoids both extremes. It also delivers term premium estimates that are in line with the common macro-finance priors and produces interest rate forecasts that outperform the random walk, as we will show below.

In recent work, Crump et al. (2018) also estimate the term premium from a model with shifting endpoints but come to a different conclusion, namely that "term premiums account for the bulk of the [...] variation in yields." The fraction of the variance of long-term interest rates accounted for by the term premium is more than one-half based on their estimates, whereas it is only about one-third according to our estimates.[42] While our model specification differs substantially from theirs, including the fact that ours is an arbitrage-free term structure model, the key difference driving the disparate conclusions is how the shifting endpoint is identified: Crump et al. use survey data, and their i_t^* closely mirrors long-run survey expectations of the three-month T-bill rate. This estimate of i_t^* is quite stable, so they find a more limited role for trends and long-run expectations and a more volatile residual term premium.

Finally, we note that when interpreting model-based estimates of the term premium, it is important to keep in mind that they assume rational expectations, which may not represent the expectations of investors at each point in time. As shown by Piazzesi et al. (2015) and Cieslak (2018), statistical estimates of risk premia are affected by persistent forecast errors by investors, i.e., deviations from rational expectations.[43] While this caveat also applies to our estimates, models like ours that allow for shifting endpoints are better suited to capture changes in long-run investor expectations that can result from learning about macroeconomic trends and policies. Estimates of such models are also much less sensitive to the choice of sample period than stationary DTSMs, which require inference about the unconditional mean of the short-term interest rate, assumed to be constant (Orphanides and Wei, 2012). For

[42] Specifically, we calculated the same decomposition of the sample variance of forward rates as Crump et al. report in their Table 2. For the variances of levels, they find that 53 to 56 percent is explained by the forward term premium at maturities of five years and beyond. By contrast, over their sample period, we find that only 35 percent of the variance is explained by the term premium.

[43] While some might conclude that estimates of term premia should therefore rely exclusively on survey expectations, there is no guarantee that expectations from surveys of professional forecasters are any closer to the true subjective expectations of bond traders than statistical models.

these reasons, our shifting-endpoint model is likely to better be able to recover true historical investor expectations and bond risk premia.

4.6 Out-of-sample forecasts of interest rates

We now turn to out-of-sample (OOS) forecasts of long-term interest rates using our new shifting-endpoint DTSM. Despite many advances in yield curve modeling, the random walk model has proven very hard to beat when forecasting bond yields, due to the extreme persistence of interest rates (Duffee, 2013). But our results so far suggest that one might be able to obtain more accurate forecasts by accounting for the interest rate trends.

We construct out-of-sample forecasts using the *OSE* model, which proxies for i_t^* with the sum of the PTR estimate for π_t^* and the real-time estimate of r_t^* described in Section 2. We will compare these forecasts to those from the *FE* model and from a driftless random walk.[44] The models are recursively estimated—that is, using an expanding estimation window using all data available up to each forecast date—starting in 1976:Q1 when five years of data are available.[45] We focus on forecasts of the ten-year yield; we have found results for other maturities to be qualitatively similar. We forecast at horizons of 4, 10, 20, 30, and 40 quarters, and because our data end in 2018:Q2 the last forecast date is 2008:Q2, for a total of 127 (overlapping) forecasts.[46] The top panel of Table 5 reports the root mean-squared errors in percentage points. We also calculate p-values for tests of equal finite-sample forecast accuracy using the approach of Diebold and Mariano (1995) (DM).[47] We calculate the DM p-values using standard normal critical values for one-sided tests of the null hypothesis that the *OSE* model does not improve upon the *RW* or *FE* forecasts. We find that the forecasts from our common-trend model *OSE* are substantially and statistically significantly more accurate than the *FE* model. For example, when forecasting 10 years ahead, model *OSE* lowers the RMSE by about 40 percent relative to the *FE* model. The improvements relative to the

[44] We do not include estimates from the *ESE* model because the estimation of this model using our MCMC sampler takes a significant amount of time, which renders recursive estimation with each additional available observation period as required for out-of-sample forecasting computationally too burdensome.

[45] A rolling scheme, which uses only a fixed number of observations for parameter estimation, allows for an easier asymptotic justification of tests for predictive ability (Giacomini and White, 2006) but requires a specific choice of the window length. We have also obtained forecasts with such a scheme, using a variety of different window lengths, and found qualitatively similar results as when using a recursive scheme.

[46] The last date for which we re-estimate the models is 2007:Q4, consistent with Section 4.2.

[47] Following common use, we construct the DM test with a rectangular window for the long-run variance and the small-sample adjustment of Harvey et al. (1997). Monte Carlo evidence in Clark and McCracken (2013) indicates that this test has good size in finite samples. However, for very long forecast horizons, the long-run variance is estimated with considerable uncertainty because there are only a few such non-overlapping observations in our sample.

RW are smaller but still statistically significant.[48] These results document that, relative to conventional stationary models, more accurate Treasury yields forecasts can be obtained by allowing for an underlying common trend in interest rates. Gains are even possible relative to the random walk benchmark. In contrast to the random walk forecast, which simply assumes all changes are permanent, accounting for i_t^*, the underlying source of the highly persistent changes in interest rates, benefits forecast accuracy.[49]

Finally, we compare the accuracy of our statistical models to that of professional forecasters. Since 1988, the Blue Chip Financial Forecasts (*BC*) survey has asked its respondents for long-range forecasts of interest rates twice a year. The respondents provide their average expectations of the target variable for each of the upcoming five calendar years and for the subsequent five-year period—we will focus on the five annual forecast horizons. We match the available information sets by using only data up to the quarter preceding the survey date for our model-based forecasts, and we exactly match the forecast horizons with the *BC* forecasts by taking averages of model-based forecasts over the relevant calendar years. The sample includes 48 forecast dates from March 1988 to December 2011. The bottom panel of Table 5 shows the RMSEs of the survey forecasts and the three model-based forecasts. Shifting-endpoint forecasts based on i_t^* improve over the *RW* forecasts in this sample as well, which are better than the *BC* forecasts. Because of the smaller sample size and substantial overlap, the gains of the *OSE* model are not as strongly significant, and for some horizons they are insignificant. The gains relative to the *FE* model are again strongly significant. The reason for the poor performance of the survey forecasts is that they consistently overpredict future yields, in particular at long horizons (results not shown). Other studies have documented the poor performance of survey forecasts of interest rates (e.g., van Dijk et al., 2014), which contrasts with the very good performance of survey-based inflation forecasts (Ang et al., 2007; Faust and Wright, 2013). Our results suggest that the underlying reason for this poor performance is that professional forecasters have not sufficiently accounted for the shifting long-run trend component in interest rates.

[48] We have found in additional, unreported analysis—using plots of differences in cumulative sums of forecast errors over time—that the forecast gains of *OSE* are not driven by certain unusual subperiods but instead are a consistent pattern over most of our sample period.

[49] Earlier research by van Dijk et al. (2014) found that, when forecasting interest rates, it is beneficial to link long-run projections of interest rates to long-run expectations of inflation, but their forecasts do not account for time variation in r_t^* and are not based on an arbitrage-free model.

5 Conclusion

In this paper, we have bridged the gap between macroeconomics and finance concerning the role of long-run trends. We have provided compelling new evidence from a variety of perspectives that interest rates and bond risk premia are substantially driven by time variation in the perceived trend in inflation and the equilibrium real rate of interest. Our results demonstrate that the links between macroeconomic trends and the yield curve are quantitatively important, and that accounting for these time-varying trend components is crucial for understanding and forecasting long-term interest rates and bond returns. Importantly, variation in the equilibrium real rate is even more consequential recently than variation in the inflation trend. Our analysis first established the links between macroeconomic trends and yields by taking as data various estimates of the trends from surveys and models. However, to understand these results, we then formulated and estimated a new no-arbitrage dynamic term structure model. In particular, when estimated with just yield curve data, this model provides a new finance-based measure of the trend in nominal interest rates that complements related estimates in the macroeconomics literature. It also provides for the first time an internally consistent formulation of that equilibrium trend and bond risk premia.

Our results provide strong support for using yield curve models that allow for slow-moving changes in the long-run means of nominal and real interest rates instead of the stationary dynamic specifications with constant means that are ubiquitous. We have developed a new common-trend dynamic term structure model, but much future research remains to be done in jointly estimating macro trends and yield curve dynamics. Three future research avenues seem particularly promising: First, a joint model of real yields, nominal yields, and inflation expectations that includes shifting endpoints could provide new separate estimates of π_t^* and r_t^*. Second, combining our shifting-endpoint DTSM specification with the shadow-rate paradigm (Bauer and Rudebusch, 2016) would allow for a sophisticated treatment of the ZLB, which may provide additional insights into the behavior of macro trends over the ZLB period. Third, one could consider changes over time in the variability of the trends by adding stochastic volatility (as in Stock and Watson, 2007) in order to assess how our (mostly unconditional) results about the relative importance of trends are affected by taking a conditional perspective.

References

Adrian, T., R. K. Crump, and E. Moench (2013): "Pricing the Term Structure with Linear Regressions," *Journal of Financial Economics*, 110, 110–138.

Ang, A., G. Bekaert, and M. Wei (2007): "Do Macro Variables, Asset Markets, or Surveys Forecast Inflation Better?" *Journal of Monetary Economics*, 54, 1163–1212.

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

Bauer, M. D. (2018): "Restrictions on Risk Prices in Dynamic Term Structure Models," *Journal of Business & Economic Statistics*, 36, 196–211.

Bauer, M. D. and J. D. Hamilton (2018): "Robust Bond Risk Premia," *Review of Financial Studies*, 31, 399–448.

Bauer, M. D. and G. D. Rudebusch (2016): "Monetary Policy Expectations at the Zero Lower Bound," *Journal of Money, Credit and Banking*, 48, 1439–1465.

——— (2017): "Resolving the Spanning Puzzle in Macro-Finance Term Structure Models," *Review of Finance*, 21, 511–553.

Bauer, M. D., G. D. Rudebusch, and J. C. Wu (2012): "Correcting Estimation Bias in Dynamic Term Structure Models," *Journal of Business & Economic Statistics*, 30, 454–467.

——— (2014): "Term Premia and Inflation Uncertainty: Empirical Evidence from an International Panel Dataset: Comment," *American Economic Review*, 104, 1–16.

Beeler, J., J. Y. Campbell, et al. (2012): "The Long-Run Risks Model and Aggregate Asset Prices: An Empirical Assessment," *Critical Finance Review*, 1, 141–182.

Bikbov, R. and M. Chernov (2010): "No-Arbitrage Macroeconomic Determinants of the Yield Curve," *Journal of Econometrics*, 159, 166–182.

Brayton, F. and P. Tinsley (1996): "A Guide to FRB/US: A Macroeconomic Model of the United States," Finance and Economics Discussion Series 1996-42, Board of Governors of the Federal Reserve System.

Brooks, J. and T. J. Moskowitz (2017): "Yield Curve Premia," Unpublished manuscript.

Campbell, J. Y. (1986): "Bond and Stock Returns in a Simple Exchange Model," *Quarterly Journal of Economics*, 101, 785–803.

Campbell, J. Y., A. W.-C. Lo, and A. C. MacKinlay (1997): *The Econometrics of Financial Markets*, Princeton University Press.

Campbell, J. Y. and P. Perron (1991): "Pitfalls and Opportunities: What Macroeconomists Should Know About Unit Roots," *NBER Macroeconomics Annual*, 6, 141–201.

Campbell, J. Y. and R. J. Shiller (1987): "Cointegration and Tests of Present Value Models," *Journal of Political Economy*, 95, 1062–1088.

——— (1991): "Yield Spreads and Interest Rate Movements: A Bird's Eye View," *Review of Economic Studies*, 58, 495–514.

Campbell, J. Y., A. Sunderam, and L. M. Viceira (2017): "Inflation Bets or Ddeflation Hedges? The Changing Risks of Nominal Bonds," *Critical Finance Review*, 6, 263–301.

Carvalho, C., A. Ferrero, and F. Nechio (2016): "Demographics and Real Interest Rates: Inspecting the Mechanism," *European Economic Review*, 88, 208–226.

Chib, S. and B. Ergashev (2009): "Analysis of Multifactor Affine Yield Curve Models," *Journal of the American Statistical Association*, 104, 1324–1337.

Chib, S. and E. Greenberg (1995): "Understanding the Metropolis-Hastings Algorithm," *The American Statistician*, 49, 327–335.

Christensen, J. H. and G. D. Rudebusch (2017): "A New Normal for Interest Rates? Evidence from Inflation-Indexed Debt," Working Paper 2017-07, Federal Reserve Bank of San Francisco.

Christensen, J. H. E., J. A. Lopez, and G. D. Rudebusch (2010): "Inflation Expectations and Risk Premiums in an Arbitrage-Free Model of Nominal and Real Bond Yields," *Journal of Money, Credit and Banking*, 42, 143–178.

Cieslak, A. (2018): "Short-Rate Expectations and Unexpected Returns in Treasury Bonds," *The Review of Financial Studies*, 31, 3265–3306.

Cieslak, A. and P. Povala (2015): "Expected Returns in Treasury Bonds," *Review of Financial Studies*, 28, 2859–2901.

Clark, T. and M. McCracken (2013): "Advances in Forecast Evaluation," in *Handbook of Economic Forecasting*, ed. by G. Elliott and A. Timmermann, Elsevier, vol. 2, Part B, 1107 – 1201.

Clark, T. E. and T. Doh (2014): "Evaluating Alternative Models of Trend Inflation," *International Journal of Forecasting*, 30, 426–448.

Cochrane, J. (2007): "Commentary on 'Macroeconomic Implications of Changes in the Term Premium'," *Federal Reserve Bank of St. Louis Review*, 271–282.

Cochrane, J. H. and M. Piazzesi (2005): "Bond Risk Premia," *American Economic Review*, 95, 138–160.

Crump, R., S. Eusepi, and E. Moench (2018): "The Term Structure of Expectations and Bond Yields," Staff Reports 775, Federal Reserve Bank of New York.

Dai, Q. and K. J. Singleton (2000): "Specification Analysis of Affine Term Structure Models," *Journal of Finance*, 55, 1943–1978.

Del Negro, M., G. Eggertsson, A. Ferrero, and N. Kiyotaki (2010): "The Great Escape? A Quantitative Evaluation of the Fed's Non-Standard Policies," working paper, Federal Reserve Bank of New York and Princeton University.

Del Negro, M., D. Giannone, M. P. Giannoni, and A. Tambalotti (2017): "Safety, Liquidity, and the Natural Rate of Interest," *Brookings Papers on Economic Activity*, 2017, 235–316.

Dewachter, H. and L. Iania (2011): "An Extended Macro-Finance Model with Financial Factors." *Journal of Financial & Quantitative Analysis*, 46.

Dewachter, H. and M. Lyrio (2006): "Macro Factors and the Term Structure of Interest Rates," *Journal of Money, Credit and Banking*, 38, 119–140.

Diebold, F. X. and R. S. Mariano (1995): "Comparing Predictive Accuracy," *Journal of Business & Economic Statistics*, 13, 253–263.

Diebold, F. X., G. D. Rudebusch, and B. S. Aruoba (2006): "The Macroeconomy and the Yield Curve: A Dynamic Latent Factor Approach," *Journal of Econometrics*, 131, 309–338.

Duffee, G. R. (2011): "Information In (and Not In) the Term Structure," *Review of Financial Studies*, 24, 2895–2934.

——— (2013): "Forecasting Interest Rates," in *Handbook of Economic Forecasting*, ed. by G. Elliott and A. Timmermann, Elsevier, vol. 2, Part A, 385–426.

Durbin, J. and S. J. Koopman (2002): "A Simple and Efficient Simulation Smoother for State Space Time Series Analysis," *Biometrika*, 89, 603–616.

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

Fama, E. F. and K. R. French (1989): "Business Conditions and Expected Returns on Stocks and Bonds," *Journal of Financial Economics*, 25, 23–49.

——— (1993): "Common Risk Factors in the Returns on Stocks and Bonds," *Journal of Financial Economics*, 33, 3–56.

Faust, J. and J. H. Wright (2013): "Forecasting Inflation," in *Handbook of Economic Forecasting*, ed. by G. Elliott and A. Timmerman, North-Holland, vol. 2, Part A, 2–56.

Fiorentini, G., A. Galesi, G. Perez-Quiros, and E. Sentana (2018): "The Rise and Fall of the Natural Interest Rate," Documentos de Trabajo 1822, Banco de Espana.

Garg, A. and M. Mazzoleni (2017): "The Natural Rate of Interest and Bond Returns," Unpublished manuscript.

Giacomini, R. and H. White (2006): "Tests of Conditional Predictive Ability," *Econometrica*, 74, 1545–1578.

Goliński, A. and P. Zaffaroni (2016): "Long Memory Affine Term Structure Models," *Journal of Econometrics*, 191, 33–56.

Gürkaynak, R. S., B. Sack, and J. H. Wright (2007): "The U.S. Treasury Yield Curve: 1961 to the Present," *Journal of Monetary Economics*, 54, 2291–2304.

Gürkaynak, R. S. and J. H. Wright (2012): "Macroeconomics and the Term Structure," *Journal of Economic Literature*, 50, 331–367.

Hamilton, J. D., E. S. Harris, J. Hatzius, and K. D. West (2016): "The Equilibrium Real Funds Rate: Past, Present, and Future," *IMF Economic Review*, 64, 660–707.

Harrison, J. M. and D. M. Kreps (1979): "Martingales and Arbitrage in Multiperiod Securities Markets," *Journal of Economic Theory*, 20, 381–408.

Harvey, A. C. (1990): *Forecasting, Structural Time Series Models and the Kalman Filter*, Cambridge University Press.

Harvey, D., S. Leybourne, and P. Newbold (1997): "Testing the equality of prediction mean squared errors," *International Journal of Forecasting*, 13, 281–291.

Holston, K., T. Laubach, and J. C. Williams (2017): "Measuring the Natural Rate of Interest: International Trends and Determinants," *Journal of International Economics*, 108, S59–S75.

Hördahl, P., O. Tristani, and D. Vestin (2006): "A Joint Econometric Model of Macroeconomic and Term-Structure Dynamics," *Journal of Econometrics*, 131, 405–444.

Hyslop, D. R. and G. W. Imbens (2001): "Bias from Classical and Other Forms of Measurement Error," *Journal of Business & Economic Statistics*, 19, 475–481.

Jardet, C., A. Monfort, and F. Pegoraro (2013): "No-Arbitrage Near-Cointegrated VAR (p) Term Structure Models, Term Premia and GDP Growth," *Journal of Banking & Finance*, 37, 389–402.

Johannsen, B. K. and E. Mertens (2016): "The Expected Real Interest Rate in the Long Run: Time Series Evidence with the Effective Lower Bound," Feds notes, Board of Governors of the Federal Reserve System.

——— (2018): "A Time Series Model of Interest Rates with the Effective Lower Bound," BIS Working Papers 715, Bank for International Settlements.

Johansen, S. (1991): "Estimation and Hypothesis Testing of Cointegration Vectors in Gaussian Vector Autoregressive Models," *Econometrica*, 59, 1551–1580.

Jorgensen, K. (2017): "How Learning from Macroeconomic Experiences Shapes the Yield Curve," Unpublished manuscript.

Joslin, S., M. Priebsch, and K. J. Singleton (2014): "Risk Premiums in Dynamic Term Structure Models with Unspanned Macro Risks," *Journal of Finance*, 69, 1197–1233.

Joslin, S., K. J. Singleton, and H. Zhu (2011): "A New Perspective on Gaussian Dynamic Term Structure Models," *Review of Financial Studies*, 24, 926–970.

Kiley, M. T. (2015): "What Can the Data Tell Us about the Equilibrium Real Interest Rate," Finance and Economics Discussion Series 2015-077, Board of Governors of the Federal Reserve System.

Kim, D. H. and A. Orphanides (2012): "Term Structure Estimation with Survey Data on Interest Rate Forecasts," *Journal of Financial and Quantitative Analysis*, 47, 241–272.

Kim, D. H. and K. J. Singleton (2012): "Term Structure Models and the Zero Bound: An Empirical Investigation of Japanese Yields," *Journal of Econometrics*, 170, 32–49.

Kim, D. H. and J. H. Wright (2005): "An Arbitrage-Free Three-Factor Term Structure Model and the Recent Behavior of Long-Term Yields and Distant-Horizon Forward Rates," Finance and Economics Discussion Series 2005-33, Board of Governors of the Federal Reserve System.

King, R. G., C. I. Plosser, J. H. Stock, and M. W. Watson (1991): "Stochastic Trends and Economic Fluctuations," *American Economic Review*, 81, 819–840.

Kozicki, S. and P. Tinsley (2001): "Shifting Endpoints in the Term Structure of Interest Rates," *Journal of Monetary Economics*, 47, 613–652.

——— (2012): "Effective Use of Survey Information in Estimating the Evolution of Expected Inflation," *Journal of Money, Credit and Banking*, 44, 145–169.

Laubach, T. and J. C. Williams (2003): "Measuring the Natural Rate of Interest," *Review of Economics and Statistics*, 85, 1063–1070.

——— (2016): "Measuring the Natural Rate of Interest Redux," *Business Economics*, 51, 57–67.

Le, A., K. J. Singleton, and Q. Dai (2010): "Discrete-Time Affine Term Structure Models with Generalized Market Prices of Risk," *Review of Financial Studies*, 23, 2184–2227.

Lubik, T. A. and C. Matthes (2015): "Calculating the Natural Rate of Interest: A Comparison of Two Alternative Approaches," Economic Brief EB15-10, Federal Reserve Bank of Richmond.

Ludvigson, S. C. and S. Ng (2009): "Macro Factors in Bond Risk Premia," *Review of Financial Studies*, 22, 5027–5067.

Lunsford, K. G. and K. D. West (2017): "Some Evidence on Secular Drivers of US Safe Real Rates," Working Paper 1723, Federal Reserve Bank of Cleveland.

Malmendier, U. and S. Nagel (2011): "Depression Babies: Do Macroeconomic Experiences Affect Risk Taking?" *Quarterly Journal of Economics*, 126, 373–416.

Mankiw, N. G. and M. D. Shapiro (1986): "News or Noise? An Analysis of GNP Revisions," Working Paper 1939, National Bureau of Economic Research.

Müller, U. K. and M. W. Watson (2013): "Low-frequency robust cointegration testing," *Journal of Econometrics*, 174, 66–81.

Nelson, C. R. and C. R. Plosser (1982): "Trends and random walks in macroeconmic time series: some evidence and implications," *Journal of monetary economics*, 10, 139–162.

Ng, S. and P. Perron (1995): "Unit root tests in ARMA models with data-dependent methods for the selection of the truncation lag," *Journal of the American Statistical Association*, 90, 268–281.

Orphanides, A. and M. Wei (2012): "Evolving macroeconomic perceptions and the term structure of interest rates," *Journal of Economic Dynamics and Control*, 36, 239–254.

Orphanides, A. and J. C. Williams (2005): "Imperfect Knowledge, Inflation Expectations, and Monetary Policy," in *The Inflation-Targeting Debate*, University of Chicago Press, 201–246.

Phillips, P. C. and S. Ouliaris (1990): "Asymptotic properties of residual based tests for cointegration," *Econometrica*, 165–193.

Piazzesi, M., J. Salomao, and M. Schneider (2015): "Trend and Cycle in Bond Premia," Unpublished manuscript.

Rachel, L. and T. Smith (2015): "Secular Drivers of the Global Real Interest Rate," Working Paper 571, Bank of England.

Rose, A. K. (1988): "Is the Real Interest Rate Stable?" *Journal of Finance*, 43, 1095–1112.

Rudebusch, G. and L. E. Svensson (1999): "Policy rules for inflation targeting," in *Monetary policy rules*, University of Chicago Press, 203–262.

Rudebusch, G. D. and E. T. Swanson (2012): "The bond premium in a DSGE Model with Long-Run Real and Nominal Risks," *American Economic Journal: Macroeconomics*, 4, 105–143.

Rudebusch, G. D. and T. Wu (2007): "Accounting for a Shift in Term Structure Behavior with No-Arbitrage and Macro-Finance Models," *Journal of Money, Credit and Banking*, 39, 395–422.

——— (2008): "A Macro-Finance Model of the Term Structure, Monetary Policy, and the Economy," *Economic Journal*, 118, 906–926.

Stock, J. H. and M. W. Watson (1988): "Testing for common trends," *Journal of the American statistical Association*, 83, 1097–1107.

——— (1993): "A simple estimator of cointegrating vectors in higher order integrated systems," *Econometrica*, 783–820.

——— (2007): "Why has US inflation become harder to forecast?" *Journal of Money, Credit and banking*, 39, 3–33.

——— (2016): "Core inflation and trend inflation," *Review of Economics and Statistics*, 98, 770–784.

Summers, L. H. (2014): "US economic prospects: Secular stagnation, hysteresis, and the zero lower bound," *Business Economics*, 49, 65–73.

Swanson, E. T. and J. C. Williams (2014): "Measuring the Effect of the Zero Lower Bound On Medium- and Longer-Term Interest Rates," *American Economic Review*, 104, 3154–3185.

Timmermann, A. (2006): "Forecast Combinations," in *Handbook of Economic Forecasting*, ed. by G. Elliott, C. Granger, and A. Timmermann, Elsevier, vol. 1, chap. 4.

van Dijk, D., S. J. Koopman, M. Wel, and J. H. Wright (2014): "Forecasting Interest Rates with Shifting Endpoints," *Journal of Applied Econometrics*, 29, 693–712.

Watson, M. W. (1986): "Univariate detrending methods with stochastic trends," *Journal of monetary economics*, 18, 49–75.

Wright, J. H. (2011): "Term Premia and Inflation Uncertainty: Empirical Evidence from an International Panel Dataset," *American Economic Review*, 101, 1514–1534.

Appendix

A Trend estimates

A.1 Overview of trend estimates

Table A.1 provides an overview of all of the trend estimates used in our analysis. The first five estimates of r_t^* are obtained from models in published studies, and these are described below in Appendix A.2 and shown in Figure A.1. Our own four estimates of r_t^* are described in Appendix A.3 and shown in Figure A.2. In Section 3, we use a "filtered" estimate and a "real-time" estimate of r_t^*, which are averages of the three filtered and six real-time estimates, respectively. The real-time estimate of i_t^* is the sum of the PTR estimate of π_t^* and the real-time estimate of r_t^*. The *ESE* model estimate of i_t^* is described in Section 4.2, with details in Appendix C.8.

Table A.1: Overview of trend estimates

Trend	Source	Real-time	Filtered	Smoothed
r^*	Del Negro et al. (2017) (DGGT)	Yes		Yes
	Johannsen and Mertens (2016) (JM)	Yes		Yes
	Laubach and Williams (2016) (LW)		Yes	Yes
	Holston et al. (2017) (HLW)		Yes	
	Kiley (2015)		Yes	Yes
r^*	UC	Yes		Yes
	Proxies	Yes		Yes
	SSM	Yes		Yes
	MA	Yes		
π_t^*	PTR	Yes		
i^*	Real-time	Yes		
	ESE			Yes

Overview of trend estimates used in the paper: five external sources of r^*, our own four estimates of r^*, the PTR estimate of π^*, and two i^* estimates. "Real-time" indicates (pseudo) real-time estimation of the trend proxy, "filtered" indicates estimates from a (one-sided) Kalman filter using full-sample parameter estimates, and "smoothed" indicates (two-sided) Kalman smoother or full-sample Bayesian estimate of the trend. All trend estimates are quarterly from 1971:Q4 to 2018:Q1.

Our proxy estimate of the inflation trend π_t^* is the Fed's PTR measure, the perceived target rate (PTR) of inflation. PTR measures expectations for inflation in the price index of personal consumption expenditures (PCE). Consistent with this definition, our estimates of r_t^* are based on real interest rates relative to PCE inflation. Since 1979, PTR corresponds to long-run inflation expectations from the Survey of Professional Forecasters (SPF). Before 1979,

it is based on estimates from a learning model for expected inflation. For details see Brayton and Tinsley (1996). Data are available at https://www.federalreserve.gov/econresdata/frbus/us-models-package.htm.

The estimates of r_t^* are described in the following. All of these estimates are based on the assumption that the real rate contains a stochastic trend, i.e., a stochastic endpoint.

A.2 External estimates of r_t^*

Del Negro et al. (2017) (DGGT) propose a number of Bayesian common-trend VARs to estimate r_t^* and its possible drivers. We focus on their baseline model in which three stochastic trends, including r_t^* and π_t^*, are estimated from five data series: (1) observed PCE inflation, (2) long-run inflation expectations from the PTR series, (3) the 3-month T-bill rate, (4) the 20-year Treasury yield and (5) long-run expectations of the 3-month yield. For details about the data and model specification see their section II.A. We estimate their model using data up to 2018:Q1 and replicate their published results. The trends are smoothed (two-sided) estimates, as they are the posterior medians of the (MCMC) sampled trend series conditional on the full data set. In addition, we recursively estimate their model starting in 1971:Q4, expanding the data by adding one quarter at a time. The smoothed estimate of r_t^* is shown in the left panel of Figure A.1, and the recursive, real-time estimate is shown in the right panel.

Johannsen and Mertens (2016, 2018) (JM) propose a time series model for interest rates with explicit treatment of the zero lower bound and stochastic volatility. For the version of their model reported in Johannsen and Mertens (2016), they generously provided us with updated estimates including both full-sample (smoothed/two-sided) and real-time estimated series of r_t^*. These are shown in Figure A.1.

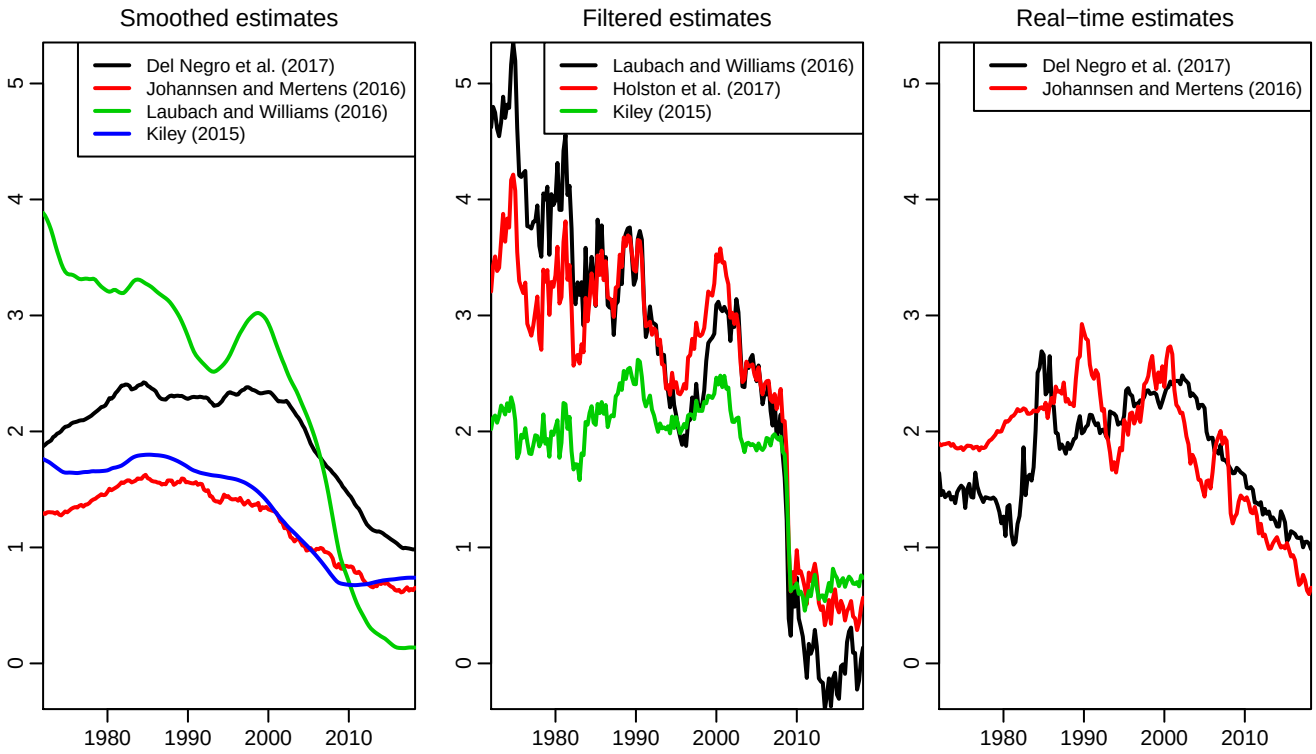
While the two estimation approaches above use the long-run definition of r_t^* that is the relevant one for the trend in nominal interest rates (see Section 2), the following three external estimates use a different definition and estimation approach. The prominent model of Laubach and Williams (2003, 2016)(LM), the slightly modified version of this model by Holston, Laubach, and Williams (2017)(HLW), and the version by Kiley (2015) all use a simple linearized New Keynesian macro model—essentially the Rudebusch and Svensson (1999) model—in which r_t^* is estimated as the *neutral* real interest rate at which monetary policy is neither expansionary nor contractionary.[50] But despite this difference with long-run r_t^*, these estimates are still worth considering in our context for several reasons. First, in practice, their "definition takes a 'longer-run' perspective, in that it refers to the level of real interest rates expected to prevail, say, five to ten years in the future, after the economy has emerged from any cyclical fluctuations and is expanding at its trend rate" (Laubach and Williams, 2016, p. 57), so the definitions are effectively quite close. Second, in all three of these models the neutral rate is a martingale, so at least as implemented within the models, the neutral real rate is also the long-run trend in the real rate. Third, these are among the most widely used estimates of r_t^* both in the academic literature and in practice. Accordingly, these estimates are useful additions to our analysis. A shortcoming of these estimates is that they are not

[50] Compared to the LW model, the HLW model excludes relative price shocks from the Phillips curve, uses a simpler proxy for inflation expectations, and assumes a one-for-one effect of trend output growth on r_t^*. In Kiley's model, the IS curve is augmented with credit spreads.

available in real time. Of course, it would in principle be possible to recursively estimate these models—and for a limited time period towards the end of our sample Laubach and Williams (2016) have done this. But these macro-based estimates are by their nature very sensitive to the macroeconomic data, in particular real GDP, so that real-time data issues become a serious concern. We therefore only have filtered/one-sided estimates based on the full-sample parameter estimates, and smoothed/two-sided estimates, which are shown in the middle and right panels of Figure A.1.

We exclude two other prominent estimates of r_t^*. Lubik and Matthes (2015) estimate r_t^* as the five-year forecast of the real rate from a time-varying parameter VAR model. This horizon is too short to be comparable with our long-run trend, and their implementation leads to a highly volatile estimate of r_t^*. Christensen and Rudebusch (2017) estimate r_t^* from yields on Treasury inflation-protected Securities (TIPS) using a dynamic term structure model that accounts for liquidity premia, but their sample only starts in 1998 with the introduction of TIPS.

Figure A.1: External estimates of r_t^*



Estimates of r_t^* from existing models. Left panel: smoothed/two-sided estimates. Middle panel: filtered/one-sided estimates. Right panel: (pseudo) real-time estimates. The sample period is from 1971:Q4 to 2018:Q1.

A.3 Details on our estimates of r_t^*

Our first estimate is based on a univariate unobserved components (UC) model for the real short-term interest rate, similar to Watson (1986). Recently, Fiorentini et al. (2018) have argued that such univariate models can infer r_t^* with greater precision than the LW model. We use an ex-ante real interest rate calculated as the difference between the three-month Treasury bill rate and core PCE inflation, that is, the four-quarter percent change in the price index for personal consumption expenditures excluding food and energy items. (Inflation over the past year serves as a proxy estimate for expectations of inflation over the next quarter.) The real rate is decomposed into a random walk trend (r_t^*) and a stationary component (the real rate gap, specified as an AR(1) process with zero mean), which are the two state variables of the state-space model. The prior distributions for the parameters are uninformative, with the exception of the variance for the innovations to the random walk component, i.e., for changes in r_t^*. We use a tight prior around a low value for this variance, similar to DGGT. Specifically, the prior distribution for this variance is inverse-gamma, $IG(\alpha/2, \delta/2)$, with $\alpha = 100$ and $\delta = 0.01(\alpha + 2)$. This implies that the mode is 0.01, and the variance of the change in r_t^* over 100 years is 4, i.e., the standard deviation is 2 (percent). This is a slightly higher mode than used by DGGT (their mode implies a standard deviation over 100 years of one percent). In our MCMC sampler, we draw the unobserved state variables using the simulation smoother of Durbin and Koopman (2002) and the parameters using standard Gibbs steps. For this model and for the following two models below, we use random starting values to initialize our MCMC chain, and we carefully monitor convergence of the MCMC sampler. For the full-sample estimation we use 100,000 MCMC draws. For our recursive estimation, we start in 1971:Q4 with 100,000 draws, and every time we add another observation, we obtain 20,000 more draws.

The second estimate (labeled "proxies") is from a multivariate model, which augments the UC model by two additional measurement equations relating r_t^* to two proxies that recent work has shown to be important correlates of the real rate trend. The first proxy is a ten-year moving average of quarterly real GDP growth, and the second proxy is a ten-year moving average of the quarterly growth rate in the total number of hours worked in the business sector, i.e., labor force hours. The choice of these proxies is motivated by the results in Lunsford and West (2017), who document strong low-frequency correlations between these series and the real rate over a post-war sample. While these macroeconomic data are subject to data revisions, in particular real GDP, the long moving averages and the use of these series in extracting the long-run trend in the real rate lessens concerns that data revisions would materially affect our real-time trend estimate. In the additional measurement equations, the proxies are scaled by a parameter to be estimated and the measurement error is allowed to be serially correlated. The model has four state variables: r_t^*, the real-rate gap, and the two measurement errors (AR(1) processes with non-zero means). The prior distributions are uninformative, except again for the trend innovation variance, where the prior is the same as in the UC model. We design a hybrid MCMC sampler: The scaling parameters are drawn using random-walk Metropolis-Hastings steps with the state variables integrated out (i.e., using the Kalman filter to obtain the likelihood for calculation of the acceptance probability). Then, the state variables are sampled with the simulation smoother and the remaining parameters

are drawn using Gibbs steps.

Our third estimate is from a state-space model (SSM) that is similar to the specification by DGGT in that it includes both inflation and the nominal short rate, and it estimates both r_t^* and π_t^*. The main differences are that we include neither survey expectations nor long-run yields, and that our measurement equations are somewhat more standard. Our three observation series are quarterly PCE inflation, the 3-month T-bill rate, and the PTR series for long-run inflation expectations. In addition to the trends, the two other state variables are the real-rate gap, r_t^g and the inflation gap, π_t^g, which follow a bivariate VAR with four lags. The measurement equations are

$$\begin{aligned} \pi_t &= \pi_t^* + \pi_t^g + e_t^\pi, \\ PTR_t &= \pi_t^* + e_t^{PTR}, \\ y_t^{(3m)} &= \pi_t^* + E_t\pi_{t+1}^g + r_t^* + r_t^g + e_t^y, \end{aligned}$$

where e_t^π, e_t^{PTR}, and e_t^y are *iid* measurement errors, and $E_t\pi_{t+1}^g$ is implied by the VAR. The priors are generally uninformative, and for the VAR we use the same Minnessota prior as DGGT. For the variance of the innovations to both r_t^* and π_t^* we use smoothness prior similar to the previous two models and the same prior modes for the variances as in DGGT (equivalent to one percent and two percent standard deviation, respectively, for changes in trends over 100 years). The MCMC sampler simply combines the simulation smoother for the state variables and Gibbs steps for the other parameters.

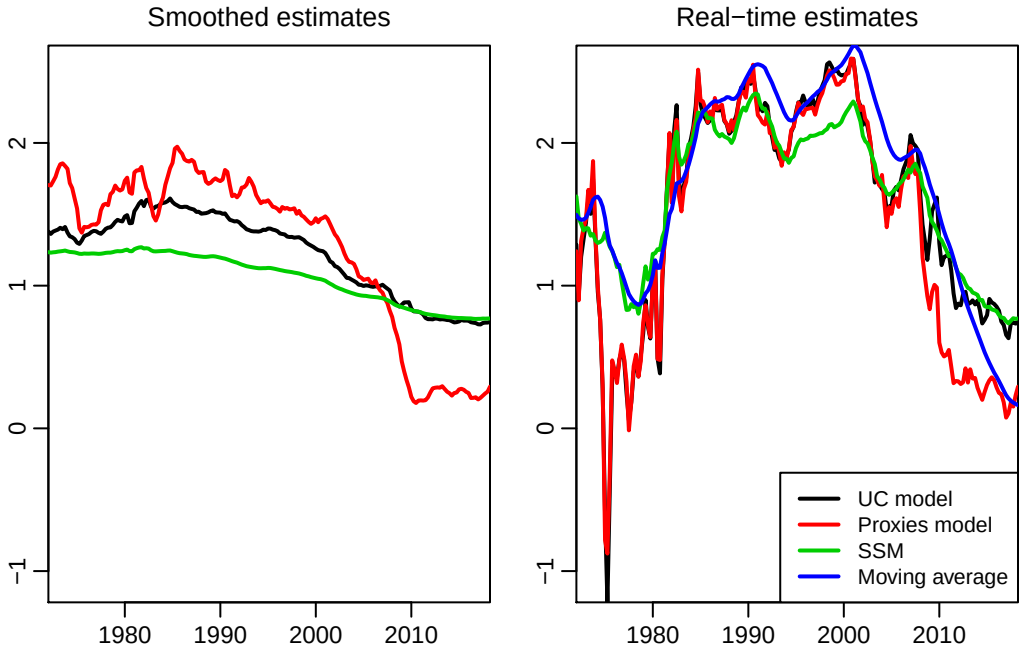
Finally, we also calculate a simple moving-average (MA) estimate of r_t^*. The observed real-rate series is the same as in the first two models above. Denoting this series by r_t, we calculate an exponentially-weighted moving average using the recursion $r_t^* = \alpha r_{t-1}^* + (1-\alpha)r_t$, which we start ten years before the beginning of our sample, in 1961:Q4, with $r_t^* = r_t$. We use $\alpha = 0.98$, a value in line with earlier work estimating macro trends.

A.4 Measurement error in macro trends

In establishing some stylized facts about a low-frequency macro-finance connection in Section 3, we treat the estimated macro trends as data. However, there is substantial model and estimation uncertainty attached to the various point estimates of r_t^*. Similarly, our survey-based measure of the long-run inflation trend, π_t^*, is also imprecise. How concerned should we be about measurement error?

Our results show that measures of macro trends are closely connected to the yield curve and that they contain important information for predicting future yields and returns. The effect of measurement error in the trend estimates on our results depends on the structure of that error. One possibility is that the trends are subject to classical measurement error, an error that is orthogonal to the unobserved trend. In this case, it is well known that the coefficients in our regressions would be both less precise and biased toward zero. An alternative possibility is that the measurement error is orthogonal to the estimate of the trend, which may be more plausible since our trend proxies are estimates of the underlying trend based on the available information set (Mankiw and Shapiro, 1986). This type of measurement error would make our estimates noisy though not necessarily biased (Hyslop and Imbens, 2001). For either type of

Figure A.2: Our estimates of r_t^*



Our own four estimates of r_t^*. Left panel: smoothed/two-sided estimates. Right panel: (pseudo) real-time estimates. The sample period is from 1971:Q4 to 2018:Q1.

measurement error, the use of a noisy trend would likely lead to weaker estimated relationships than the use of the true underlying trend. Hence, accounting for measurement error suggests that our results should be viewed as a lower bound for the tightness of the connection between the yield curve and the *true* macro trends.

Finally, we note that our trend measures were not created on purpose to match the evolution of Treasury yields. Therefore, the connections we find could well be stronger if the trend estimates were optimized to exhibit a tight connection with or predictive power for long-term yields.

B Additional results for Section 3

B.1 Persistence and common trends (Section 3.1)

Table B.1 documents the extremely high persistence of various measures of nominal Treasury yields. It reports summary statistics for the two-year yield, $y_t^{(8)}$, the ten-year yield, $y_t^{(40)}$, and the level of the yield curve, L_t, measured as the first principal component of yields (scaled so that the loadings add up to one). Summary statistics are also reported for our macro trends, including the inflation trend, our three estimates of the real rate trend—the filtered estimate, $r_t^{*,F}$, the real-time estimate, $r_t^{*,RT}$, and the moving-average estimate, $r_t^{*,MA}$—and our real-time proxy for the equilibrium nominal short rate, $i_t^* = \pi_t^* + r_t^{*,RT}$. To assess whether

a simple difference between yields and trends can remove the extreme persistence and make yields stationary we also report summary statistics for yields detrended in this way. Along with standard deviations, the table reports two measures of persistence: the estimated first-order autocorrelation coefficient, $\hat{\rho}$, and the half-life, which indicates the number of quarters until half of a given shock has died out and is calculated as $\ln(0.5)/\ln(\hat{\rho})$. It also includes two tests for a unit root, the Augmented Dickey-Fuller (ADF) t-statistic and the non-parametric Phillips-Perron (PP) Z_α statistic. For the ADF test, we include a constant and k lagged difference in the test regression, where k is determined using the general-to-specific procedure suggested by Ng and Perron (1995), starting from four lags. For the PP test, we use a Newey-West estimator of the long-run variance with four lags. In the last column, we report the p-value for the null hypothesis of stationarity using the low-frequency stationarity test (LFST) of Müller and Watson (2013).

The yields are highly persistent, with a first-order autocorrelation coefficient of 0.97 and a half-life of around 22-27 quarters. Neither the ADF or PP tests reject a unit root, while the LFST p-values reject stationarity for each series. This evidence suggests that nominal yields can be effectively modeled as $I(1)$ processes.

The macro trends are even more persistent. For example, our real-time estimate of r_t^* has an autocorrelation coefficient of 0.98 and a half-life of about 36 quarters. The inflation trend and i_t^* have autocorrelation coefficients of 0.99 and half-lives of 85 and 60 quarters, respectively. Unsurprisingly, these macro trend proxies exhibit pronounced trend behavior, that is, they also behave like I(1) processes.

The naive detrending method of subtracting trend proxies from yields leads to time series with much lower persistence. If only the inflation trend is subtracted, the resulting series is less persistent but still sufficiently trending that the evidence from all three tests favors a unit root. When a real rate trend—$r_t^{*,F}$, $r_t^{*,RT}$, or $r_t^{*,MA}$—is also subtracted off together with π_t^*, the persistence drops further, and the evidence sometimes favors stationarity, depending on the specific interest rate, real rate trend proxy, and unit root test. This extremely simple detrending method, which makes the strong assumption that coefficients in the cointegration vector of yields and macro trends are all one (in absolute value), has some success, but only if both macro trends are accounted for. The evidence in Section 3.1 finds an even stronger macro-finance link when the cointegration coefficients are slightly larger than one.

Table B.2 shows the results for a cointegration analysis similar to the one carried out in Section 3.1, here applied to the level of the yield curve. The results are qualitatively similar to those we reported in Table 1 for the ten-year yield.

B.2 Predicting excess returns (Section 3.2)

For assessing the predictive power of macro trends for excess bond returns, the main text focuses on returns that are averaged across bond maturities from two to 15 years. Table B.3 shows (full-sample) estimates for individual maturites. It compares the usual yields-only specification to the specification that also includes our estimate of the long-run nominal short rate i_t^*. For all maturities, the addition of i_t^* (i) substantially raises R^2, (ii) makes both the level (PC1) and slope (PC2) strongly significant predictors, and (iii) leads to a large and strongly significant coefficient on i_t^* itself. This trend coefficient has the opposite sign and

Table B.1: Persistence of interest rates, macroeconomic trends, and differences

Series	SD	$\hat{\rho}$	Half-life	ADF	PP	LFST
$y_t^{(8)}$	3.54	0.97	21.9	-1.25	-4.59	0.00
$y_t^{(8)} - \pi_t^*$	2.26	0.93	9.1	-2.07	-11.09	0.00
$y_t^{(8)} - \pi_t^* - r_t^{*,F}$	1.61	0.87	4.9	-3.21**	-22.58***	0.11
$y_t^{(8)} - \pi_t^* - r_t^{*,RT}$	2.00	0.92	7.8	-2.24	-13.23*	0.00
$y_t^{(8)} - \pi_t^* - r_t^{*,MA}$	2.05	0.91	7.8	-2.35	-13.83*	0.01
$y_t^{(40)}$	2.94	0.97	26.4	-1.13	-3.11	0.00
$y_t^{(40)} - \pi_t^*$	1.67	0.93	9.4	-2.32	-9.72	0.01
$y_t^{(40)} - \pi_t^* - r_t^{*,F}$	1.17	0.87	5.0	-3.61***	-22.87***	0.23
$y_t^{(40)} - \pi_t^* - r_t^{*,RT}$	1.33	0.90	6.7	-2.51	-15.16**	0.00
$y_t^{(40)} - \pi_t^* - r_t^{*,MA}$	1.36	0.90	6.8	-3.04**	-16.01**	0.01
L_t	3.13	0.97	26.6	-1.20	-3.42	0.00
$L_t - \pi_t^*$	1.84	0.93	9.8	-1.91	-9.69	0.00
$L_t - \pi_t^* - r_t^{*,F}$	1.24	0.87	4.8	-2.99**	-23.03***	0.20
$L_t - \pi_t^* - r_t^{*,RT}$	1.54	0.91	7.7	-2.19	-13.05*	0.00
$L_t - \pi_t^* - r_t^{*,MA}$	1.58	0.91	7.6	-2.36	-13.90*	0.01
π_t^*	1.60	0.99	85.4	-0.62	-1.16	0.00
$r_t^{*,F}$	1.03	0.97	27.2	-0.67	-1.16	0.00
$r_t^{*,RT}$	0.59	0.98	36.2	-0.77	-1.04	0.03
$r_t^{*,MA}$	0.70	0.98	43.4	-0.75	1.31	0.03
$i_t^* = \pi_t^* + r_t^{*,RT}$	1.72	0.99	59.7	-0.30	-0.28	0.00

Standard deviation (SD); first-order autocorrelation coefficient ($\hat{\rho}$); half-life ($\ln(0.5)/\ln(\hat{\rho})$); Augmented Dickey-Fuller (ADF) and Phillips-Perron (PP) unit root test statistics (with *,** and *** indicating significance at 10%, 5%, and 1% level) and p-values for Mueller-Watson low-frequency stationary test (LFST), for the two-year yield, $y_t^{(8)}$, the ten-year yield, $y_t^{(40)}$, the level of the yield curve (the first principal component of yields), L_t, the detrended yields, and macro trends. The r^*-estimates are the filtered ("F"), real-time ("RT") and moving-average ("MA") estimates described in Section 2. The data are quarterly from 1971:Q4 to 2018:Q1.

Table B.2: Cointegration regressions and tests for the level of the yield curve

	Level	(1)	(2)	(3)	(4)	(5)
constant	6.48	-0.23	-1.17	-2.59	-0.84	-3.06
	(0.55)	(0.62)	(0.35)	(0.44)	(0.42)	(0.37)
π_t^*		1.76	1.31	1.56	1.45	
		(0.13)	(0.11)	(0.08)	(0.09)	
r_t^*			1.13	1.90	1.11	
			(0.15)	(0.16)	(0.12)	
i_t^*						1.75
						(0.07)
R^2		0.83	0.93	0.96	0.97	0.94
Memo: r^*			filtered	real-time	mov. avg.	real-time
SD	2.94	1.39	1.10	0.85	1.10	0.81
$\hat{\rho}$	0.97	0.87	0.82	0.72	0.81	0.67
Half-life	26.4	5.1	3.5	2.1	3.3	1.8
ADF	-1.13	-2.68	-3.44	-3.92**	-3.44	-4.26**
PP	-3.11	-20.07*	-31.16**	-51.73***	-32.57**	-61.30***
LFST	0.00	0.04	0.28	0.37	0.10	0.84
Johansen $r = 0$		10.10	29.31	43.26***	38.11**	21.19**
Johansen $r = 1$		0.90	5.96	11.13	10.22	0.71
ECM $\hat{\alpha}$		-0.10	-0.16	-0.29	-0.33	-0.32
		(0.03)	(0.05)	(0.07)	(0.09)	(0.07)

Dynamic OLS regressions for the level of the yield curve, L_t, (the first principal component of yields) on macroeconomic trends, including four leads and lags of ΔL_t and differenced trend variables. Newey-West standard errors using six lags are in parentheses. The r^* estimates are described in Section 2, and the long-run nominal short rate i_t^* is the sum of π_t^* and the real-time estimate of r_t^*. For the cointegration residuals (and, in the first column, for the yield level itself), the second panel reports standard deviations (SD), first-order autocorrelation coefficients ($\hat{\rho}$), half-lives ($\ln(0.5)/\ln(\hat{\rho})$), Augmented Dickey-Fuller (ADF) and Phillips-Perron (PP) unit root test statistics, and p-values for Mueller-Watson low-frequency stationary test (LFST). The table also reports the Johansen trace statistic which tests whether the cointegration rank (r) among the L_t and the macro trends is zero/one against the alternative that it exceeds zero/one, using four lags in the VAR. For ADF, PP and Johansen tests *,** and *** indicate significance at 10%, 5%, and 1% level. Estimates of the coefficient α (with White standard errors) on the cointegration residual in the error-correction model (ECM) for ΔL_t that also includes an intercept, four lags of ΔL_t, and four lags of differenced macro trends. The data are quarterly from 1971:Q4 to 2018:Q1.

Table B.3: Excess return regressions for individual maturities

	Yields only				Yields and i^*				
	PC1	PC2	PC3	R^2	PC1	PC2	PC3	i^*	R^2
2y	0.06	0.05	-0.65	0.05	0.70	0.12	-0.26	-1.19	0.15
	(0.06)	(0.06)	(0.50)		(0.22)	(0.05)	(0.52)	(0.37)	[0.02]
5y	0.08	0.22	-1.72	0.08	1.53	0.38	-0.80	-2.74	0.18
	(0.11)	(0.12)	(0.87)		(0.41)	(0.10)	(0.89)	(0.70)	[0.00]
7y	0.08	0.34	-2.19	0.08	2.06	0.56	-0.94	-3.74	0.19
	(0.14)	(0.15)	(1.11)		(0.53)	(0.13)	(1.14)	(0.90)	[0.00]
10y	0.09	0.51	-2.74	0.09	2.90	0.83	-0.96	-5.31	0.20
	(0.20)	(0.20)	(1.56)		(0.71)	(0.18)	(1.57)	(1.23)	[0.00]
15y	0.08	0.83	-3.70	0.09	4.13	1.28	-1.15	-7.64	0.21
	(0.28)	(0.29)	(2.20)		(0.99)	(0.26)	(2.17)	(1.72)	[0.00]

Predictive regressions for quarterly excess returns on bonds with maturities 2, 5, 7, 10 and 15 years. The predictors are the first three principal components of yields (PC1, PC2, PC3) and, in the second specification, the long-run nominal short rate i_t^*, the sum of π_t^* and the real-time estimate of r_t^* (see Section 2). Numbers in parentheses are White standard errors and in squared brackets are small-sample p-values for the spanning hypothesis (that the coefficient on i_t^* is zero) obtained with the bootstrap method of Bauer and Hamilton (2018). The data are quarterly from 1971:Q4 to 2018:Q1.

a somewhat larger magnitude than the coefficient on the level. These findings all closely parallel those in reported in Table 2. In addition, the results show that the coefficients on both yield predictors and i_t^* rise with maturity, due to the fact that return volatility scales with bond maturity. Finally, it is noteworthy that for the two-year maturity, excess returns appear unpredictable with yield information alone, but become strongly predictable once our trend proxy is added.

In the presence of persistent predictors, it is generally difficult to interpret the magnitude of R^2 as a measure of predictive accuracy, because persistent predictors that are truly irrelevant in population can substantially increase R^2 in small samples (Bauer and Hamilton, 2018). To be able to better interpret the R^2 we use our bootstrap to generate small-sample distributions of R^2 under the spanning hypothesis (that is, under the null that the three PCs of yields contain all the relevant information for predicting returns). We then compare the statistics obtained in the actual data to the quantiles of these bootstrap distributions. If the level or change of a regression R^2 is outside of the 95%-bootstrap interval then the regression results are inconsistent with the spanning hypothesis and suggest that additional predictors beyond yields alone are statistically significant.

The top panel of Table B.4 reports this comparison for the specifications with (i) only yields, (ii) yields and π_t^*, (iii) yields, π_t^* and the real-time estimate of r_t^*, and (iv) yields and i_t^*. Adding π_t^* to the regression increases R^2 by 7 percentage points, while the 95%-bootstrap interval indicates that under the null hypothesis it would be uncommon to observe an increase in R^2 of more than 5 percentage points. Adding the real-time r_t^* estimate increases R^2 to 21%, and the increase relative to the yields-only specification is 12 percentage points, while

Table B.4: Predictive power: R^2 with small-sample bootstrap intervals

	Full sample: 1971:Q4–2018:Q1		Subsample: 1985:Q1–2018:Q1	
Predictors	R^2	ΔR^2	R^2	ΔR^2
Yields only	0.09		0.08	
	[0.03, 0.19]		[0.03, 0.18]	
Yields and π_t^*	0.16	0.07	0.10	0.02
	[0.04, 0.20]	[0.00, 0.05]	[0.04, 0.19]	[0.00, 0.05]
Yields, π_t^* and r_t^*	0.21	0.12	0.19	0.10
	[0.05, 0.21]	[0.00, 0.07]	[0.05, 0.20]	[0.00, 0.07]
Yields and i_t^*	0.21	0.12	0.16	0.07
	[0.04, 0.20]	[0.00, 0.04]	[0.04, 0.19]	[0.00, 0.04]
Yields detrended by π_t^*	0.15	0.07	0.08	-0.00
	[0.03, 0.19]	[-0.03, 0.04]	[0.03, 0.18]	[-0.03, 0.04]
Yields detrended by π_t^* and r_t^*	0.18	0.09	0.17	0.09
	[0.04, 0.19]	[-0.04, 0.05]	[0.03, 0.18]	[-0.04, 0.05]
Yields detrended by i_t^*	0.20	0.12	0.14	0.06
	[0.03, 0.19]	[-0.03, 0.03]	[0.03, 0.18]	[-0.03, 0.03]

Predictive power of regressions for quarterly excess bond returns, averaged across two- to 15-year maturities. The predictors are three principal components (PCs) of yields, the PTR estimate of the inflation trend π_t^*, our real-time estimate of the equilibrium real rate r_t^*, and the equilibrium nominal short rate i_t^* taken as the sum of these inflation and real-rate trend estimates. The last three specifications use detrended yields, that is, three PCs of the yield residuals in regressions on (i) π_t^*, (ii) π_t^* and r_t^*, or (iii) i_t^*. Increase in R^2 (ΔR^2) is reported relative to the first specification with only PCs of yields. Numbers in square brackets are 95%-bootstrap intervals obtained by calculating the same regressions statistics in 5,000 bootstrap data sets generated under the (spanning) null hypothesis that only yields have predictive power for bond returns, using the bootstrap method of Bauer and Hamilton (2018).

the bootstrap suggests an increase of at most 7 percentage points would be plausible under the null. Adding just i_t^* also increases R^2 by 12 percentage points—much more than is plausible under the null. In the post-1985 subsample, the increase in R^2 from only adding π_t^* is not statistically significant, whereas the increases from adding either both macro trends or only i_t^* are significant.

C Details on dynamic term structure model

C.1 Dynamic system

The evolution of the state variables under the real-world (or physical) probability measure, denoted as the $\mathbb{P}$-measure, is given by equation (3), the common trends representation of the dynamic system (Stock and Watson, 1988). The state variables $Z_t = (\tau, P_t')'$ are cointegrated: $\beta' Z_t \sim I(0)$, with $\beta = (-\gamma, I_N)'$ an obvious choice for the N cointegration vectors. Yields are also cointegrated with a single common trend, as evident from equation (7).

To derive the VAR representation of the dynamic system, substitute for $\tilde{P}_t$ and $\tilde{P}_{t-1}$ in the VAR equation, using $\tilde{P}_t = P_t - \bar{P} - \gamma\tau_t$. This yields

$$\begin{aligned} P_t &= (I_N - \Phi)\bar{P} + \Phi P_{t-1} + \gamma\tau_t - \Phi\gamma\tau_{t-1} + \tilde{u}_t \\ &= (I_N - \Phi)\bar{P} + (I_N - \Phi)\gamma\tau_{t-1} + \Phi P_{t-1} + u_t, \end{aligned}$$

where the last equation defines the innovations to P_t, $u_t = \gamma\eta_t + \tilde{u}_t$, which have covariance matrix $\Omega = E(u_t u_t') = \gamma\gamma'\sigma_\eta^2 + \widetilde{\Omega}$. Thus, the VAR representation for Z_t is

$$Z_t = \mu_Z + \Phi_Z Z_{t-1} + v_t, \quad v_t = (\eta_t, u_t')', \tag{C.1}$$

with

$$\mu_Z = \begin{pmatrix} 0 \\ (I_n - \Phi)\bar{P} \end{pmatrix},$$

$$\Phi_Z = \begin{pmatrix} 1 & 0_{1\times N} \\ (I_N - \Phi)\gamma & \Phi \end{pmatrix},$$

and innovation covariance matrix

$$\Omega_v = E(v_t v_t') = \begin{pmatrix} \sigma_\eta^2 & \gamma'\sigma_\eta^2 \\ \gamma\sigma_\eta^2 & \Omega \end{pmatrix}.$$

This is obviously a cointegrated VAR with one common trend and N cointegration vectors since Φ_Z has exactly one eigenvalue equal to unity.

The vector-error-correction (VEC) representation is

$$\Delta Z_t = \mu_Z + \alpha\beta' Z_{t-1} + v_t, \quad \alpha = \begin{pmatrix} 0 \\ \Phi - I_N \end{pmatrix}, \ \beta = \begin{pmatrix} -\gamma' \\ I_N \end{pmatrix},$$

where $\Phi_Z - I_{N+1} = \alpha\beta'$. Since the intercept can be written as $\mu_Z = -\alpha\beta'(0, \bar{P}')'$, the cointegration residual is $\beta'\left(Z_{t-1} - (0, \bar{P}')'\right) = \tilde{P}_{t-1}$. It Granger-causes the yield factors according to the last N equations of the VEC representation, which can be written as

$$\Delta P_t = (\Phi - I_N)\tilde{P}_{t-1} + u_t.$$

That is, deviations of the time-t yield factors from their equilibrium, $P_t = \bar{P} + \gamma\tau_t$, are reduced over time by future changes in yields.

C.2 Risk-neutral dynamics and stochastic discount factor

The dynamic system under the risk-neutral measure $\mathbb{Q}$ is given in equation (6). We assume that under $\mathbb{Q}$ the yield factors follow a *stationary* VAR. This specification avoids the counterfactual implication of a unit root under $\mathbb{Q}$ that yields and forward rates diverge to minus infinity with maturity—in that case the conditional variance of future short rates, and hence the convexity in yields, would be unbounded, as discussed below in C.3.

From the short rate equation (5) and the $\mathbb{Q}$-dynamics (6) it follows that the long-run mean

of the short rate under $\mathbb{Q}$ is constant and given by

$$E^{\mathbb{Q}}(i_t) = \delta_0 + \delta_1' E^{\mathbb{Q}}(P_t) = \delta_0 + \delta_1'(I_N - \Phi^{\mathbb{Q}})^{-1}\mu^{\mathbb{Q}}.$$

It further follows from our specification that the trend τ_t is unspanned by yields, meaning that there is no deterministic mapping from P_t (or model-implied yields) to τ_t. Formally, the requirement for τ_t to be an unspanned risk factor is

$$E^{\mathbb{Q}}(i_{t+h}|P_t, \tau_t) = E^{\mathbb{Q}}(i_{t+h}|P_t), \quad \text{for all } h,$$

see Joslin et al. (2014). Since neither the short rate i_t (see equation 5) nor expectations of future yield factors P_{t+h} (see equation 6) depend on τ_t once we condition on P_t, this requirement is trivially satisfied. Note that the $\mathbb{Q}$-dynamics of τ_t are not identified because yields and bond prices are not sensitive to movements in τ_t. If we had allowed for unit eigenvalue of $\Phi^{\mathbb{Q}}$, so that the long-run risk-neutral mean of the short rate would be time-varying, this $\mathbb{Q}$-endpoint would be spanned by P_t and not deterministically related to τ_t.

Two key features of our risk-neutral distribution for P_t are its stationarity and its independence of τ_t. We now explain how these risk-neutral dynamics are reconciled with the non-stationary unit root dynamics under the real-world probability measure that are central to our paper. As in Joslin et al. (2014), the stochastic discount factor (SDF) for the bond market is the projection of the economy-wide SDF on the risk factors driving bond prices, P_t. Our specification implies that the SDF is exponentially affine:

$$M_{t+1} = \exp(-i_t - \frac{1}{2}\lambda_t'\lambda_t - \lambda_t'\varepsilon_{t+1}) \tag{C.2}$$

with $\varepsilon_t = \Omega^{-1/2}u_t$, where $\Omega^{1/2}$ is the Cholesky decomposition of Ω.[51] Shocks to τ_t do not affect the SDF because these risks are unspanned. The risk prices λ_t depend on all risk factors—including τ_t—according to the affine function

$$\lambda_t = \Omega^{-1/2}(\lambda_0 + \lambda_1 Z_t), \tag{C.3}$$

where λ_0 is a N-vector and λ_1 is a $N \times (N+1)$ matrix of risk sensitivities. To see how this SDF reconciles our real-world and risk-neutral dynamic specifications, and to find λ_0 and λ_1, we derive the risk-neutral distribution of P_{t+1} conditional on P_t and τ_t from the real-world distribution and the SDF, using standard tools (Ang and Piazzesi, 2003; Le et al., 2010; Bauer, 2018). The conditional Laplace-transform (that is, the conditional moment-generating

[51] That is, $\Omega^{1/2}$ is lower triangular and satisfies $\Omega = \Omega^{1/2}(\Omega^{1/2})'$.

function) under $\mathbb{Q}$ is

$$\begin{aligned} E^{\mathbb{Q}}(\exp(u'P_{t+1})|Z_t) &= E(\exp(u'P_{t+1} - \frac{1}{2}\lambda_t'\lambda_t - \lambda_t'\varepsilon_{t+1})|Z_t) \\ &= E(\exp(u'(\mu_P + \Phi_{PZ}Z_t + \Omega^{1/2}\varepsilon_{t+1}) - \frac{1}{2}\lambda_t'\lambda_t - \lambda_t'\varepsilon_{t+1})|Z_t) \\ &= \exp(u'(\mu_P + \Phi_{PZ}Z_t) + \frac{1}{2}u'\Omega u - u'\Omega^{1/2}\lambda_t) \\ &= \exp(u'(\mu_P - \lambda_0 + (\Phi_{PZ} - \lambda_1)Z_t) + \frac{1}{2}u'\Omega u). \end{aligned}$$

The first equality changes the probability measure using our SDF and the fact that for any $t+1$-measurable random variable X_{t+1} we have $E_t^{\mathbb{Q}}(X_{t+1}) = E_t(M_{t+1}X_{t+1})/E_t(M_{t+1})$.[52] The second equality plugs in the dynamic process for P_{t+1} from equation (C.1), with μ_P denoting the last N elements of μ_Z and Φ_{PZ} denoting the last N rows of Φ_Z. The third equality uses the mean of a log-normal distribution. The fourth and last equality substitutes for $\Omega^{1/2}\lambda_t$ from equation (C.3). Since the moment-generating function for a multivariate normal distribution with mean μ and covariance matrix Σ is $\exp(u'\mu + 0.5u'\Sigma'u)$ we see that under $\mathbb{Q}$ the conditional distribution of P_{t+1} is Gaussian with mean $\mu_p - \lambda_0 + (\Phi_{PZ} - \lambda_1)Z_t$ and covariance matrix Ω. That is, $E_t^{\mathbb{Q}}P_{t+1} = E_tP_{t+1} - \lambda_0 - \lambda_1 Z_t = E_tP_{t+1} - \Omega^{1/2}\lambda_t$. In light of equations (6) and (C.1) we conclude that

$$\lambda_0 = \mu_P - \mu^{\mathbb{Q}} = (I_n - \Phi)\bar{P} - \mu^{\mathbb{Q}} \tag{C.4}$$

and

$$\lambda_1 = \Phi_{PZ} - \left[0_{N\times 1}, \Phi^{\mathbb{Q}}\right] = \left[(I_N - \Phi)\gamma, \Phi - \Phi^{\mathbb{Q}}\right]. \tag{C.5}$$

By construction, the trend τ_t does not affect the risk-neutral conditional expectations of future P_{t+1}. Thus, the risk sensitivities for τ_t (in the first column of λ_1) have to exactly equal the sensitivities of real-world conditional expectations of P_{t+1} to τ_t (in the first column of Φ_{PZ}). Some authors (e.g. Duffee, 2011; Bauer and Rudebusch, 2017) speak of "knife-edge" restrictions when constraints lead to such cancellations and render certain factors unspanned. These restrictions imply that movements in a factor that is unspanned (τ_t in our case) lead to changes in term premia that exactly offset the changes in expectations, leaving yields unchanged. The risk sensitivities for P_t (in the last three columns of λ_1) are unconstrained, except that the implied $\Phi^{\mathbb{Q}}$ must have eigenvalues less than unity in absolute value.

C.3 Affine loadings

Prices of zero-coupon bonds are exponentially affine. Specifically, if $p_t^{(n)}$ denotes the log-price of an n-period bond, $p_t^{(n)} = \mathcal{A}_n + \mathcal{B}_n'P_t$, and the coefficients can be found using the pricing equation $\exp\left(p_t^{(n+1)}\right) = E_t^{\mathbb{Q}}\exp\left(-i_t p_{t+1}^{(n)}\right)$. They follow the usual recursions (e.g., Ang and

[52] The density of the risk-neutral measure with respect to the real-world measure, also known as the Radon-Nikodym derivative, is

$$\frac{d\mathbb{Q}}{d\mathbb{P}} = \frac{M_{t+1}}{E_t(M_{t+1})} = \exp\left(-\frac{1}{2}\lambda_t'\lambda_t - \lambda_t'\varepsilon_{t+1}\right).$$

Piazzesi, 2003):

$$\mathcal{A}_{n+1} = \mathcal{A}_n + \mathcal{B}_n'\mu^{\mathbb{Q}} + \frac{1}{2}\mathcal{B}_n'\Omega\mathcal{B}_n - \delta_0 \tag{C.6}$$

$$\mathcal{B}_{n+1} = (\Phi^{\mathbb{Q}})'\mathcal{B}_n - \delta_1 \tag{C.7}$$

with initial conditions $\mathcal{A}_0 = 0$ and $\mathcal{B}_0 = 0$.[53] Yields are affine functions of the factors:

$$y_t^{(n)} = A_n + B_n'P_t \quad A_n = -\frac{1}{n}\mathcal{A}_n, \quad B_n = -\frac{1}{n}\mathcal{B}_n.$$

If we denote the vector of J yields that are included in our model as Y_t, we can write $Y_t = A + BP_t$ for J-vector $A = (A_{n_1}, \ldots, A_{n_J})'$, and $J \times N$ matrix $B = (B_{n_1}, \ldots, B_{n_J})'$.

It is often useful to work with forward rates instead of yields, and we use them here to develop intuition about the role of the trend τ_t. We consider one-period forward rates $f_t^{(n)} = p_t^{(n)} - p_t^{(n+1)}$. Yields are averages of forward rates, $y_t^{(n)} = n^{-1}\sum_{j=0}^{n-1} f_t^{(n)}$. Forward rates depend on the risk factors as follows:

$$\begin{aligned} f_t^{(n)} &= \mathcal{A}_n - \mathcal{A}_{n+1} + (\mathcal{B}_n - \mathcal{B}_{n+1})'P_t \\ &= \underbrace{-\frac{1}{2}\mathcal{B}_n'\Omega\mathcal{B}_n}_{\text{convexity}} + \underbrace{\delta_0 - \mathcal{B}_n'\mu^{\mathbb{Q}} + \delta_1'(\Phi^{\mathbb{Q}})^n P_t}_{\text{expected short rate under } \mathbb{Q}} \\ &= \text{constant} + \delta_1'(\Phi^{\mathbb{Q}})^n\gamma\tau_t + \delta_1'(\Phi^{\mathbb{Q}})^n\tilde{P}_t \end{aligned} \tag{C.8}$$

The first equality simply uses the definition of the forward rate and the affine form for log bond price. The second equality follows from the recursions (C.6)-(C.7). It shows that forward rates are equal to a (negative) convexity term plus risk-neutral expected future short rates, $E_t^{\mathbb{Q}} i_{t+n} = \delta_0 + \delta_1' E_t^{\mathbb{Q}} P_{t+n}$. The last equality substitutes for P_t from (3) to show the loadings of forward rates on the trend and cycle factors.

The convexity term in the expression for forward rates is $-\frac{1}{2}\mathcal{B}_n'\Omega\mathcal{B}_n$. Yield convexity is the average of forward rate convexity. Since the $\mathbb{Q}$-dynamics are stationary, $\mathcal{B}_n$ converges to a finite limit as n tends to infinity, and so does the convexity in yields and forward rates. If, by contrast, $\Phi^{\mathbb{Q}}$ had an eigenvalue on the unit circle, convexity and therefore yields and forward rates would diverge to minus infinity. This fact, discussed in Campbell et al. (1997, p. 433), is an important reason why term structure models are usually specified with stationary risk-neutral dynamics, and we follow this tradition here.

Of particular interest are the loadings of forward rates on the trend, which are $\delta_1'(\Phi^{\mathbb{Q}})^n\gamma$. For the short rate, $i_t = f_t^{(0)}$, this loading is one, since we have normalized $\delta_1'\gamma = 1$ in order

[53] The closed-form solution for $\mathcal{B}_n$ is

$$\mathcal{B}_n = -\left[\sum_{j=0}^{n-1}\left(\Phi^{\mathbb{Q}}\right)^j\right]'\delta_1 = -\left[\left(I_N - \Phi^{Q}\right)^{-1}\left(I_N - (\Phi^{\mathbb{Q}})^n\right)\right]'\delta_1, \quad \lim_{j\to\infty}\mathcal{B}_n = -\left[\left(I_N - \Phi^{Q}\right)^{-1}\right]'\delta_1,$$

where the second equality and the existence of the limit rely on our assumption that $\Phi^{\mathbb{Q}}$ has no eigenvalues on the unit circle.

to identify τ_t as the equilibrium short rate i_t^*. Because the $\mathbb{Q}$-dynamics are stationary, these loadings converge to zero with increasing maturity n, as do all yield and forward rate loadings—the limiting forward rate/bond yield is a constant. Importantly, for the yield maturities included in our data, our model estimates imply yield and forward rate loadings on τ_t that are *above one*, due to the effects of the term premium to be discussed below. That is, loadings on the trend show a pronounced hump shape, increasing from one (for the short rate) to above one (for the horizons of practical interest), and then gradually decreasing thereafter. Of course, in practice the largest eigenvalue of $\Phi^{\mathbb{Q}}$ is below but very close to one, which makes it possible that even the interest rates with the longest observed maturities can exhibit significant time variation as well as larger than one-for-one responses to changes in the long-run trend.

C.4 JSZ normalization

We follow Joslin et al. (2014), who also estimate a model with unspanned risk factors, in imposing the normalization of Joslin et al. (2011)(JSZ) on the yield factors P_t. This normalization is convenient because it imposes all identifying restrictions on the $\mathbb{Q}$-dynamics, so that the real-world dynamics can be estimated without restrictions.

The JSZ normalization works as follows: First, start with N latent factors X_t and impose normalizing restrictions on their $\mathbb{Q}$-dynamics. Second, use the affine mapping

$$P_t = WY_t = WA_X + WB_X X_t, \tag{C.9}$$

where A_X and B_X are the affine yield loadings on X_t, to obtain the $\mathbb{Q}$-parameters and yield loadings for P_t in terms of the parameters/loadings of X_t.

The normalizations are (i) $i_t = \iota' X_t$, (ii) $\mu_X^{\mathbb{Q}} = (k^{\mathbb{Q}}, 0, 0)'$ and (iii) $\Phi_X^{\mathbb{Q}} = diag(\lambda^{\mathbb{Q}})$, with real, distinct, descending-ordered diagonal elements $\lambda^{\mathbb{Q}}$ that are assumed to be less than one in absolute value.[54] The loadings of yields on P_t, A and B, are easily calculated from A_X and B_X using the mapping in equation (C.9). Alternatively, one can first calculate the "rotated" parameters $\mu^{\mathbb{Q}}$, $\Phi^{\mathbb{Q}}$, δ_0 and δ_1, and then use the recursions (C.6)-(C.7) to calculate the loadings A and B. Importantly, the yield loadings only depend on the parameters $k^{\mathbb{Q}}$, $\lambda^{\mathbb{Q}}$ and Ω. The normalization ensures the consistency conditions $WA = 0$ and $WB = I_N$ so that we indeed have $WY_t = P_t$.

C.5 Excess bond returns

Model-implied excess bond returns are

$$\begin{aligned} rx_{t+1}^{(n)} &= p_{t+1}^{(n-1)} - p_t^{(n)} - i_t = \mathcal{A}_{n-1} + \mathcal{B}_{n-1}' P_{t+1} - \mathcal{A}_n - \mathcal{B}_n' P_t - \delta_0 - \delta_1' P_t \\ &= -\frac{1}{2}\mathcal{B}_{n-1}'\Omega\mathcal{B}_{n-1} + \mathcal{B}_{n-1}'\left(P_{t+1} - \mu^{\mathbb{Q}} - \Phi^{\mathbb{Q}} P_t\right) \\ &= -\frac{1}{2}\mathcal{B}_{n-1}'\Omega\mathcal{B}_{n-1} + \mathcal{B}_{n-1}'\left(E_t P_{t+1} - E_t^{\mathbb{Q}} P_{t+1}\right) + \mathcal{B}_{n-1}' u_{t+1}. \end{aligned}$$

[54] JSZ show that it is possible to allow for repeated and complex eigenvalues.

where the second line uses the coefficient recursions (C.6)-(C.7) and the third line uses conditional expecations of P_{t+1} under the real-world and risk-neutral mesures (from equations C.1 and 6). Expected excess returns, adjusted for the convexity term, are thus

$$\begin{aligned} E_t rx_{t+1}^{(n)} + \frac{1}{2} Var_t rx_{t+1}^{(n)} &= \mathcal{B}_{n-1}' \Omega^{1/2} \lambda_t = \mathcal{B}_{n-1}'(\lambda_0 + \lambda_1 Z_t) \\ &= \mathcal{B}_{n-1}' \left[\lambda_0 + (I_N - \Phi)\gamma\tau_t + (\Phi - \Phi^{\mathbb{Q}})P_t \right] \qquad \text{(C.10)} \\ &= \mathcal{B}_{n-1}' \left[\lambda_0 + (\Phi - \Phi^{\mathbb{Q}})\bar{P} + (I_N - \Phi^{\mathbb{Q}})\gamma\tau_t + (\Phi - \Phi^{\mathbb{Q}})\tilde{P}_t \right], \qquad \text{(C.11)} \end{aligned}$$

where the first equality uses the fact that the prices of risk capture the difference between real-world and risk-neutral conditional expectations, the second line uses the risk sensitivities given in (C.5), and the third line substitutes for P_t to show the loadings of expected excess returns on the ternd and cycle factors.

Again, the loadings on the trend τ_t are most interesting. Equation (C.10) shows that when holding the yield factors P_t constant, the loadings on τ_t are $\mathcal{B}_{n-1}'(I_N - \Phi)\gamma$. Our model estimates imply substantially negative values for these loadings, across all n, in line with the regression estimates in Section 3.2. An increase in the trend *ceteris paribus* means that yields will adjust upward causing negative excess returns for bond holders. If instead we hold the yield cycles $\tilde{P}_t$ constant, the loadings on τ_t are $\mathcal{B}_{n-1}'(I_N - \Phi^{\mathbb{Q}})\gamma$, as shown by equation (C.11). These loadings tend to be an order of magnitude smaller, consistent with the findings in Section 3.2 that predictive regressions including only detrended yields capture most of the predictability of excess bond returns, and they are positive because risk premia are positively related to movements in the trend.

C.6 Term premium

The term premium is the difference of a long-term interest rate and that rate's expectations component. In equation (1) this expectations component was for simplicity given as the average of expected future short rates. But in the context of no-arbitrage term structure models, it is advantageous to instead use "risk-neutral rates" which, like yields, account for the convexity of bond prices. Specifically, the risk-neutral yield on an n-period bond is

$$\tilde{y}_t^{(n)} = -\log E_t \exp\left(-\sum_{j=0}^{n-1} i_{t+j} \right) / n,$$

which differs from the average expected future short rate $\sum_{j=0}^{n-1} E_t i_{t+j}/n$ by a Jensen's inequality term. The risk-neutral yields are affine in Z_t, and the loadings are calculated using recursions similar to those in equations (C.6)- (C.7), with μ_Z and Φ_Z replacing $\mu^{\mathbb{Q}}$ and $\Phi^{\mathbb{Q}}$, Ω_v replacing Ω, and an $N+1$ vector $(0, \delta_1')'$ replacing δ_1. Crucially, risk-neutral yields depend on both τ_t and P_t, since τ_t affects real-world expectations of future short rates. The term premium on an n-period bond is then defined as

$$TP_t^{(n)} = y_t^{(n)} - \tilde{y}_t^{(n)}.$$

Because the real-world dynamics feature a unit root, risk-neutral rates diverge to minus infinity (and the term premium to infinity) as maturity increases. This issue is neither of theoretical nor practical concern: First, no tradeable securities have payoffs tied to these quantities. Second, at the maturities we focus on, due to the high persistence of the risk-neutral dynamics, the convexity in actual and risk-neutral yields is quite similar, and estimates of the term premium are not noticeably affected by convexity.

We again turn to forward rates for intuition. Using the notation $\tilde{p}_t^{(n)} = \tilde{\mathcal{A}}_n + \tilde{\mathcal{B}}_n Z_t$ for log risk-neutral bond prices and their loadings, the risk-neutral forward rate is

$$\begin{aligned}
\tilde{f}_t^{(n)} &= \tilde{p}_t^{(n)} - \tilde{p}_t^{(n+1)} = \tilde{\mathcal{A}}_n - \tilde{\mathcal{A}}_{n+1} + (\tilde{\mathcal{B}}_n - \tilde{\mathcal{B}}_{n+1})' Z_t \\
&= -\frac{1}{2}\tilde{\mathcal{B}}_n' \Omega_v \tilde{\mathcal{B}}_n + \delta_0 - \tilde{\mathcal{B}}_n' \mu_Z + (0, \delta_1') \Phi_Z^n Z_t \\
&= -\frac{1}{2}\tilde{\mathcal{B}}_n' \Omega_v \tilde{\mathcal{B}}_n + \delta_0 + (0, \delta_1') E_t Z_{t+n} \\
&= -\frac{1}{2}\tilde{\mathcal{B}}_n' \Omega_v \tilde{\mathcal{B}}_n + \delta_0 + \delta_1' (\bar{P} + \gamma\tau_t + E_t \tilde{P}_{t+n}) \\
&= \text{constant} + \tau_t + \delta_1' \Phi^n \tilde{P}_t, \qquad \text{(C.12)}
\end{aligned}$$

where the second line follows from the recursions for the risk-neutral loadings, the third line recognizes and replaces the conditional expectations for Z_{t+n}, the fourth line uses the fact that $(0, \delta_1') E_t Z_{t+n} = \delta_1' E_t P_{t+n}$ and substitutes for $E_t P_{t+n}$, and the last line uses the facts that $\delta_1' \gamma = 1$ and $E_t \tilde{P}_{t+n} = \Phi^n \tilde{P}_t$. These derivations clearly show that $\tilde{f}_t^{(n)} = E_t i_{t+n} + \text{convexity}$, and that these expectations are of course affected one-for-one by the trend $\tau_t = i_t^*$. Equations (C.8) and (C.12) imply that the forward term premium is

$$ftp_t^{(n)} = f_t^{(n)} - \tilde{f}_t^{(n)} = \text{constant} + \delta_1' \left[(\Phi^{\mathbb{Q}})^n - I_N \right] \gamma\tau_t + \delta_1' \left[(\Phi^{\mathbb{Q}})^n - \Phi^n \right] \tilde{P}_t.$$

The loadings on τ_t have a limit of -1 as n tends to infinity, but for the maturities relevant in practice our model estimates imply that these loadings are positive. That is, for the interest rate maturities we focus on (up to 15 years), forward rates have loadings on τ_t that are above one, as discussed above in C.3, because forward term premia have loadings on τ_t that are above zero. Since yields are just averages of forward rates, the same holds for yields and the yield term premium: yield loadings on the trend are larger than one, consistent with the cointegration regressions reported in Section 3.1, and term premia have loadings on the trend that are positive.

C.7 Estimation with observed shifting endpoint

In this implementation of our shifting-endpoint DTSM, we use an empirical proxy for i_t^*. Since $\tau_t = i_t^*$ we can treat this state variable as observable. Furthermore, P_t is assumed to be observable, in which case, $WY_t^o = WY_t$ or $We_t = 0$. Thus, there are effectively only $J - N$ independent measurement errors (see Joslin et al., 2011). Because all of the state variables in Z_t are observable, estimation is quite simple: Not only can the model's likelihood function be evaluated without the Kalman filter, we can also concentrate out the parameters Φ and σ_e^2

from the likelihood. That is, for given values of $k^{\mathbb{Q}}$, $\lambda^{\mathbb{Q}}$, γ, $\bar{P}$, Ω, and σ_η^2, we can analytically obtain the values of Φ and σ_e^2 that maximize the likelihood function. Computationally then, we maximize the log-likelihood over 15 parameters ($k^{\mathbb{Q}}$, 3 in $\lambda^{\mathbb{Q}}$, 2 in γ, 2 in $\bar{P}$, 6 in Ω, and σ_η^2) instead of over 25 parameters.

Alternatively, it could be assumed that the state variables are unobservable, which would require estimation of the state-space form of the model using the Kalman filter. We used this alternative method to estimate our model and obtained results essentially identical to those reported for the version with observable state variables. Still, in some applications the Kalman filter may necessary. For example, it can accommodate missing observations or include multiple measurement equations to pin down i_t^* using several different proxies. Using the Kalman filter is straightforward but computationally intensive (although excellent starting values for the model parameters can be obtained from a first-step estimation with observable state variables.)

C.8 Bayesian estimation with unobserved shifting endpoint

We first describe the prior distributions:

- For the trend innovation variance, σ_η^2, we assume an inverse-gamma distribution, specifically, $IG(\alpha_\eta/2, \delta_\eta/2)$ with $\alpha_\eta = 100$ and $\delta_\eta = 0.1^2/400(\alpha_\eta + 2)$. This parametrization implies a tight prior distribution around a mode of $0.1^2/400$, which corresponds to a standard deviation of 10% for change in τ_t over 100 years. We view this prior as a conservative choice that is justified in large part by consideration of the slow-moving macroeconomic drivers underlying π_t^* and r_t^* and by the circumscribed evolution of the various available estimates of those macro trends. Our prior for σ_η^2 is very similar to DGGT, which also uses their model's prior to limit "the amount of variation that it attributes to the trends" (DGGT, p. 249).

- For Φ and $\widetilde{\Omega}$, we specify a Minnesota-type normal/inverse-Wishart conjugate prior similar to DGGT. The prior for $\widetilde{\Omega}$ is inverse-Wishart, $IW(\kappa, \Psi)$ with $\kappa = N + 2$ and Ψ diagonal with elements $2 \cdot 10^{-3}$, $2 \cdot 10^{-4}$, and $2 \cdot 10^{-5}$ on the diagonal, values that are based on maximum likelihood estimation of the *OSE* model. The prior mean for $\widetilde{\Omega}$ is $\Psi/(\kappa - N - 1) = \Psi$, and the prior mode is $\Psi/(\kappa + N + 1)$. That is, we use the lowest value of κ for the mean to exist, so that the distribution is very disperse. The prior for Φ is a Minnessota prior, but centered around a zero matrix. For the hyperparameter controlling the overall tightness, we use the value of 0.2 like DGGT and others. Furthermore, we also restrict the eigenvalues of Φ to be inside the unit circle.

- For the measurement error variance σ_e^2, we use an inverse-gamma distribution with $\alpha_e = 4$ and $\delta_e = 0.001^2(\alpha_e + 2)$, implying a prior mode for the variance of 0.001^2 which corresponds to a standard deviation of 10 basis points. This prior mode is motivated by the fact that sufficiently flexible DTSMs can generally achieve a good fit to observed yields with only about 5-10 basis points root-mean-squared error.

- The priors for the remaining parameters are completely uninformative.

We use a state-space formulation in terms of the state variables $\tilde{Z}_t = (\tau_t, \tilde{P}_t')'$, which is of course equivalent to using τ_t and P_t. The measurement equation is

$$Y_t^o = A + B\bar{P} + B\gamma\tau_t + B\tilde{P}_t,$$

and the state equation is

$$\tilde{Z}_t = \begin{pmatrix} 1 & 0 \\ 0 & \Phi \end{pmatrix} \tilde{Z}_{t-1} + \begin{pmatrix} \eta_t \\ \tilde{u}_t \end{pmatrix}.$$

Our MCMC algorithm is a block-wise Metropolis-Hastings (M-H) sampler (Chib and Greenberg, 1995). First, note that the log-likelihood is the sum of the cross-sectional log-likelihood (from the measurement equation) and the "dynamic" log-likelihood (from the transition equation). The former, called the (log of the) $\mathbb{Q}$-likelihood by Joslin et al. (2011), for observation t is

$$\log f(Y_t^o | P_t, k^{\mathbb{Q}}, \lambda^{\mathbb{Q}}, \Omega, \sigma_e^2) \propto -\frac{1}{2}|Y_t^o - A - BP_t|/\sigma_e^2,$$

where A and B depend on $(k^{\mathbb{Q}}, \lambda^{\mathbb{Q}}, \Omega)$ and $|\cdot|$ denotes the $\mathcal{L}^2$ (Euclidean) norm. The latter, the $\mathbb{P}$-likelihood, for observation t is

$$\log f(Z_t | Z_{t-1}, \bar{P}, \gamma, \Phi, \Omega, \sigma_\eta^2) = -\frac{1}{2}\left(\log|\widetilde{\Omega}| + \tilde{u}_t'\widetilde{\Omega}^{-1}\tilde{u}_t + \log(\sigma_\eta^2) + \eta_t^2/\sigma_\eta^2\right),$$

where $\widetilde{\Omega} = \Omega - \gamma\gamma'\sigma_\eta^2$, $\tilde{u}_t = \tilde{P}_t - \Phi\tilde{P}_{t-1}$, $\tilde{P}_t = P_t - \gamma\tau_t - \bar{P}$, and $\eta_t = \Delta\tau_t$.

The first block is the sampling of the state variables using the simulation smoother of Durbin and Koopman (2002), and we condition on Z_t when drawing the other blocks. The parameter blocks are (1) $(k^{\mathbb{Q}}, \lambda^{\mathbb{Q}})$, (2) Ω, (3) g (which determines γ), (4) p (which determines $\bar{P}$), (5) σ_η^2, (6) Φ, and (7) σ_e^2. The sampling of Φ and σ_e^2 is straightforward: Gibbs steps can be employed since the full conditional posterior is available. For Φ, the draw is only accepted if the matrix has all eigenvalues inside the unit circle, to ensure stationarity of $\tilde{P}_t$. For blocks (1)-(5), the conditional posterior distributions cannot be sampled; therefore, we must use M-H steps. Only the parameters in blocks (1) and (2) affect the loadings A and B and hence the $\mathbb{Q}$-likelihood. The parameters in blocks (3)-(5) only affect the $\mathbb{P}$-likelihood, which saves on computational costs. For our M-H steps, we use tailored independence proposal distributions similar to Chib and Ergashev (2009). Specifically, conditional on all other parameters and the state variables, we use numerical optimization to find the mode and Hessian of the conditional posterior distribution. Our proposal distribution is then a multivariate t-distribution (with five degrees of freedom) that is centered around this mode and has a covariance matrix equal to the inverse of this Hessian. This way of constructing M-H proposal distributions has the great benefit that no fine-tuning of the proposal distributions is necessary, and the efficiency of the sampler is generally much better than using random walk proposals (Chib and Ergashev, 2009).

We initialize our sampler with the parameter values that maximize the posterior mode. We run the sampler for 500,000 iterations, and discard the first half as a burn-in sample. Of the remaining iterations we use, for computational efficiency, every tenth iteration, so that we have an MCMC sample of 25,000 iterations for our posterior analysis.

Table 1: Cointegration regressions and tests

	Yield	(1)	(2)	(3)	(4)	(5)
constant	6.48	0.65	-0.22	-1.82	-0.62	-2.17
	(0.55)	(0.54)	(0.33)	(0.43)	(0.39)	(0.30)
π_t^*		1.65	1.26	1.53	1.47	
		(0.11)	(0.08)	(0.07)	(0.07)	
r_t^*			0.99	1.76	1.17	
			(0.15)	(0.16)	(0.13)	
i_t^*						1.67
						(0.06)
R^2		0.85	0.93	0.96	0.96	0.95
Memo: r^*			filtered	real-time	mov. avg.	real-time
SD	2.94	1.31	1.09	0.70	0.87	0.70
$\hat{\rho}$	0.97	0.88	0.85	0.65	0.76	0.64
Half-life	26.4	5.6	4.3	1.6	2.5	1.5
ADF	-1.13	-2.60	-3.94**	-5.32***	-4.32**	-5.37***
PP	-3.11	-18.32*	-26.73**	-68.45***	-46.05***	-70.30***
LFST	0.00	0.03	0.16	0.72	0.23	0.70
Johansen $r = 0$		13.34	33.08*	46.83***	45.46***	30.69***
Johansen $r = 1$		1.29	5.92	11.57	9.11	0.73
ECM $\hat{\alpha}$		-0.11	-0.18	-0.44	-0.49	-0.45
		(0.03)	(0.05)	(0.08)	(0.09)	(0.08)

Dynamic OLS regressions of the ten-year yield, $y_t^{(40)}$, on macroeconomic trends, including four leads and lags of $\Delta y_t^{(40)}$ and differenced trend variables. Newey-West standard errors using six lags are in parentheses. The r^* estimates are described in Section 2, and the long-run nominal short rate i_t^* is the sum of π_t^* and the real-time estimate of r_t^*. For the cointegration residuals (and, in the first column, for the ten-year yield itself), the second panel reports standard deviations (SD), first-order autocorrelation coefficients ($\hat{\rho}$), half-lives ($\ln(0.5)/\ln(\hat{\rho})$), Augmented Dickey-Fuller (ADF) and Phillips-Perron (PP) unit root test statistics, and p-values for Mueller-Watson low-frequency stationary test (LFST). The table also reports the Johansen trace statistic which tests whether the cointegration rank (r) among the yield and the macro trends is zero/one against the alternative that it exceeds zero/one, using four lags in the VAR. For ADF, PP and Johansen tests *,** and *** indicate significance at 10%, 5%, and 1% level. Estimates of the coefficient α (with White standard errors) on the cointegration residual in the error-correction model (ECM) for $\Delta y_t^{(40)}$ that also includes an intercept, four lags of $\Delta y_t^{(40)}$, and four lags of differenced macro trends. The data are quarterly from 1971:Q4 to 2018:Q1.

Table 2: Predictive regressions: yields and macro trends

	(1)	(2)	(3)	(4)	(5)	(6)
Full sample: 1971:Q4–2018:Q1						
PC1	0.08	0.98	1.39	2.38	2.04	2.47
	(0.17)	(0.26)	(0.39)	(0.67)	(0.56)	(0.61)
PC2	0.43	0.47	0.43	0.67	0.68	0.70
	(0.17)	(0.17)	(0.17)	(0.15)	(0.15)	(0.15)
PC3	-2.37	-1.79	-1.92	-0.92	-0.90	-0.86
	(1.34)	(1.27)	(1.22)	(1.39)	(1.43)	(1.35)
π_t^*		-1.95	-2.21	-4.40	-3.88	
		(0.44)	(0.47)	(1.10)	(0.92)	
		[0.00]	[0.00]	[0.00]	[0.00]	
r_t^*			-1.19	-3.89	-2.70	
			(0.59)	(1.47)	(1.04)	
			[0.14]	[0.07]	[0.04]	
i_t^*						-4.50
						(1.05)
						[0.00]
R^2	0.09	0.16	0.18	0.21	0.20	0.21
Memo: r^*			filtered	real-time	mov. avg.	real-time
Subsample: 1985:Q1–2018:Q1						
PC1	0.25	0.59	1.67	2.65	2.38	1.93
	(0.16)	(0.22)	(0.47)	(0.57)	(0.51)	(0.47)
PC2	0.41	0.50	0.49	0.53	0.65	0.58
	(0.15)	(0.16)	(0.16)	(0.15)	(0.15)	(0.15)
PC3	-1.09	-0.97	0.14	1.74	2.12	0.56
	(1.14)	(1.12)	(1.30)	(1.48)	(1.55)	(1.19)
π_t^*		-1.05	-1.95	-3.44	-3.34	
		(0.73)	(0.75)	(0.87)	(0.83)	
		[0.38]	[0.10]	[0.01]	[0.01]	
r_t^*			-2.03	-5.80	-4.10	
			(0.82)	(1.54)	(1.08)	
			[0.07]	[0.01]	[0.01]	
i_t^*						-3.08
						(0.91)
						[0.02]
R^2	0.08	0.10	0.14	0.19	0.18	0.16
Memo: r^*			filtered	real-time	mov. avg.	real-time

Predictive regressions for quarterly excess bond returns, averaged across two- to 15-year maturities. The predictors are the first three principal components of yields (PC1, PC2, PC3) and estimates of the inflation trend π_t^*, the real-rate trend r_t^*, and the long-run nominal short rate i_t^*. The long-run nominal short rate i_t^* is the sum of π_t^* and the real-time estimate of r_t^*. The macro trend estimates are described in Section 2. Numbers in parentheses are White standard errors and in squared brackets are small-sample p-values obtained with the bootstrap method of Bauer and Hamilton (2018). The data are quarterly from 1971:Q4 to 2018:Q1.

Table 3: Predictive regressions: detrended yields

	Full sample: 1971–2018				Subsample: 1985–2018			
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
PC1	0.08	0.98	1.25	1.36	0.33	0.41	1.14	0.81
	(0.17)	(0.25)	(0.51)	(0.50)	(0.16)	(0.22)	(0.39)	(0.38)
PC2	0.43	0.48	0.76	0.78	0.34	0.44	0.62	0.62
	(0.17)	(0.17)	(0.16)	(0.16)	(0.14)	(0.14)	(0.14)	(0.14)
PC3	-2.37	-1.77	-0.79	-0.73	-1.04	-0.96	1.53	0.45
	(1.34)	(1.26)	(1.38)	(1.33)	(1.01)	(1.01)	(1.32)	(1.05)
R^2	0.09	0.15	0.18	0.20	0.08	0.08	0.17	0.14

Predictive regressions for quarterly excess bond returns, averaged across two- to 15-year maturities. In specification (1) the predictors are the first three principal components of observed yields (PC1, PC2, PC3). In specifications (2), (3), and (4) the predictors are the first three principal components of detrended yields, that is, of residuals in regressions for yields on π_t^*, π_t^* and r_t^*, or i_t^*, respectively. The long-run nominal short rate i_t^* is the sum of π_t^* and the real-time estimate of r_t^*. The macro trend estimates are described in Section 2. Numbers in parentheses are White standard errors with six lags. The data are quarterly from 1971:Q4 to 2018:Q1.

Table 4: Model-implied predictability of excess bond returns

	R^2 PCs only	R^2 with i_t^*	ΔR^2
Data	0.09	0.21	0.12
FE model	0.09	0.10	0.01
	[0.04, 0.17]	[0.05, 0.17]	[0.00, 0.04]
OSE model	0.10	0.19	0.09
	[0.04, 0.18]	[0.13, 0.26]	[0.02, 0.18]
ESE model	0.07	0.15	0.08
	[0.02, 0.13]	[0.05, 0.27]	[0.00, 0.20]

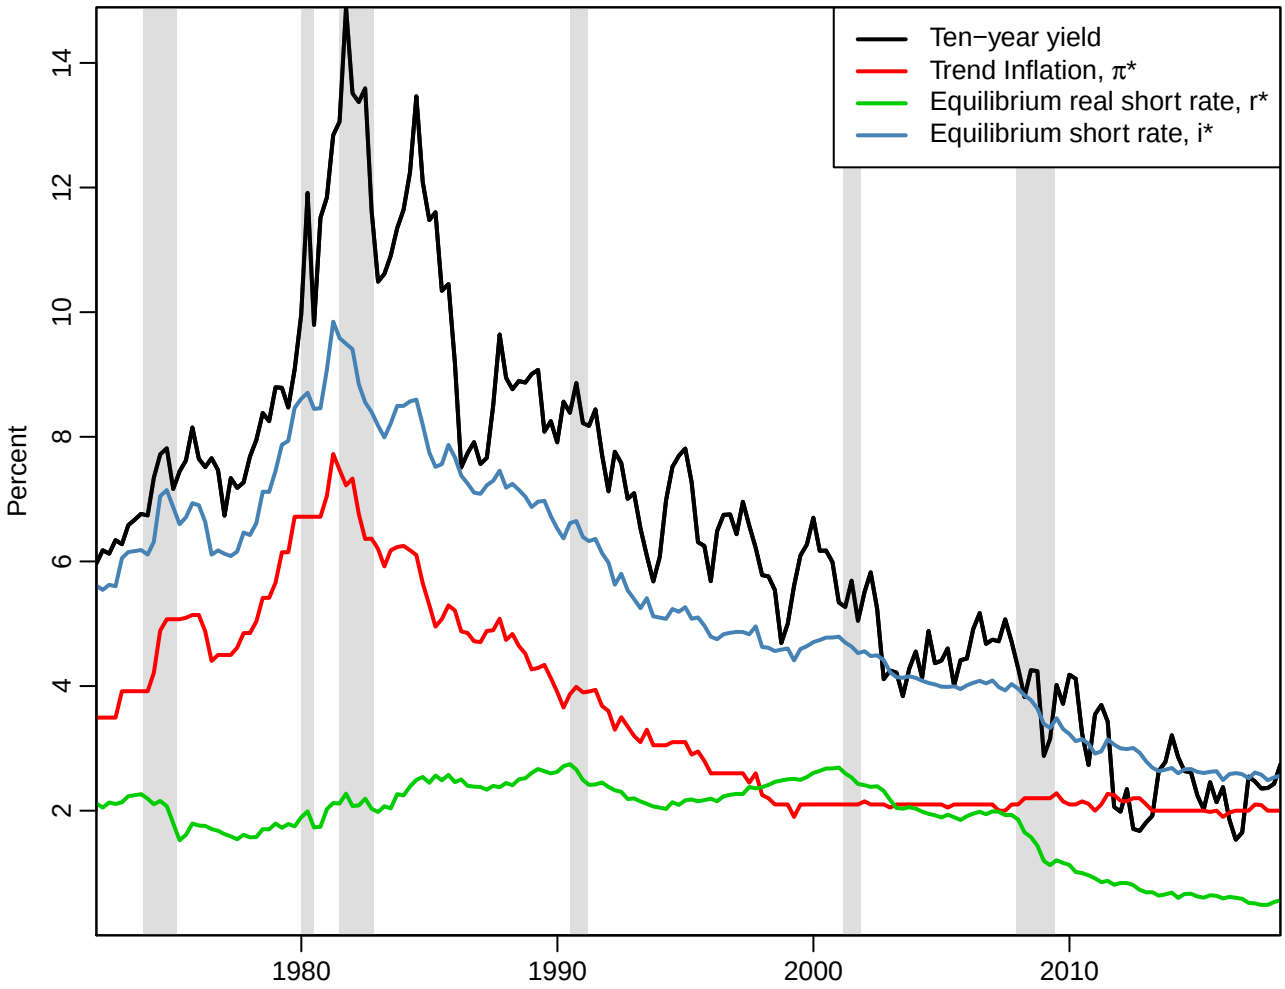
The R^2 of predictive regressions for quarterly excess bond returns, averaged across maturities of 2 to 15 years. The R^2 in the data correspond to the full-sample estimates in Table 2 (first and last columns of top panel). The model-implied R^2 are based on 5,000 simulations of artificial data sets of the same size as the full sample. The table reports means and 95%-Monte Carlo intervals (in square brackets) of the R^2 of predictive regressions estimated in these simulated data.

Table 5: Accuracy of out-of-sample forecasts for the ten-year yield

Quarterly sample: 1976:Q3-2008:Q1	Horizon in quarters				
(127 quarters)	4	10	20	30	40
Random walk (RW)	1.33	1.85	2.52	2.60	2.88
Fixed endpoint (FE)	1.42	2.25	3.28	3.72	4.19
Observed shifting endpoint (OSE)	1.17	1.76	2.37	2.39	2.60
p-value: $OSE \geq RW$	0.05	0.00	0.00	0.03	0.04
p-value: $OSE \geq FE$	0.00	0.00	0.01	0.03	0.05
Blue Chip sample: 1988:Q1-2011:Q4	Horizon in years				
(48 Blue Chip surveys)	1	2	3	4	5
Blue Chip (BC)	1.06	1.39	1.59	1.79	1.99
Random walk (RW)	0.85	1.08	1.21	1.37	1.56
Fixed endpoint (FE)	1.53	2.08	2.52	2.96	3.34
Observed shifting endpoint (OSE)	0.87	0.95	1.04	1.18	1.37
p-value: $OSE \geq BC$	0.10	0.08	0.15	0.18	0.20
p-value: $OSE \geq RW$	0.58	0.05	0.01	0.04	0.08
p-value: $OSE \geq FE$	0.00	0.00	0.00	0.00	0.00

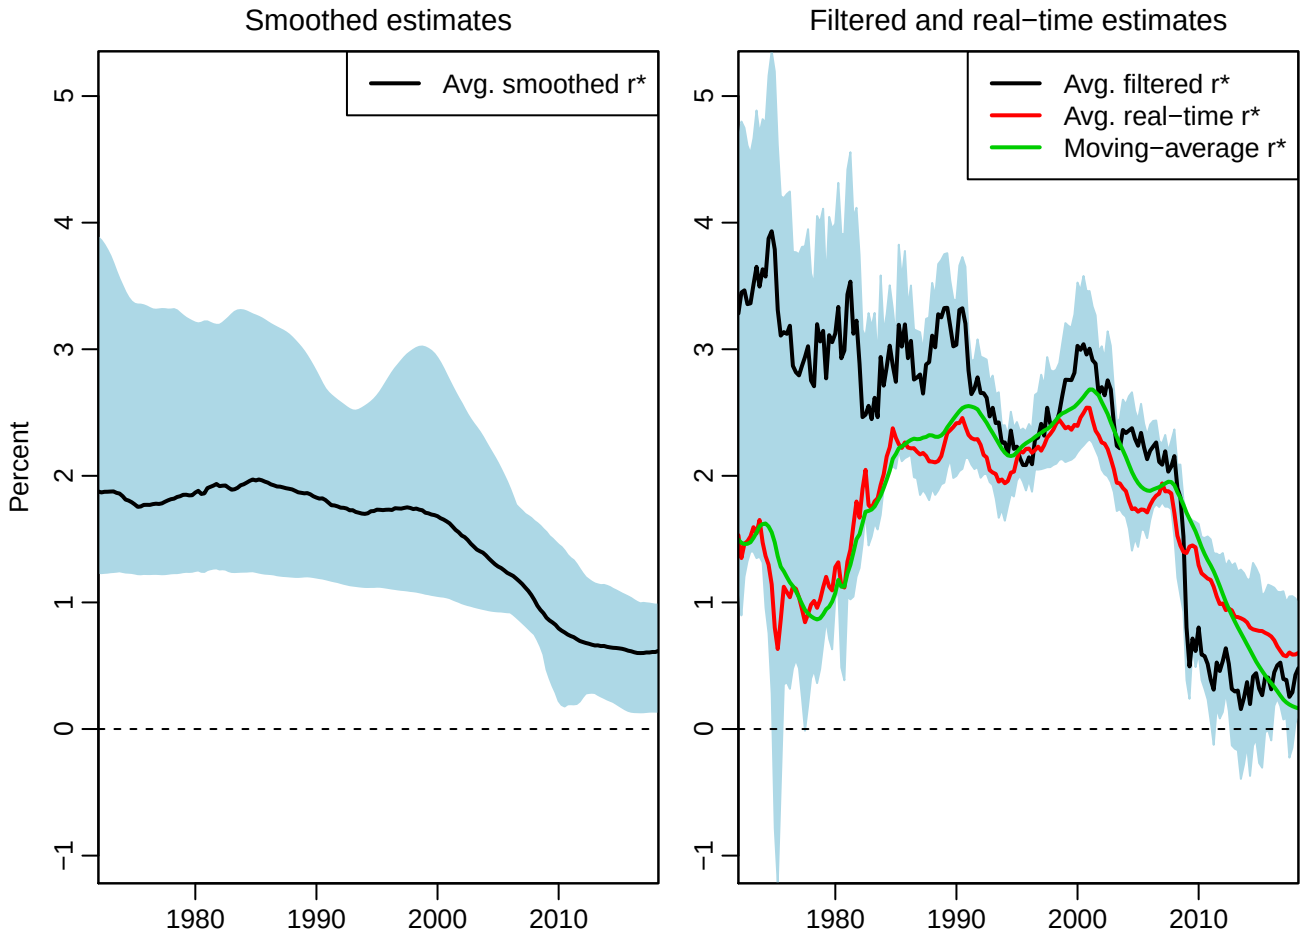
Root mean-squared errors for forecasts for the ten-year Treasury yield. Method RW is a driftless random walk, FE is a common stationary dynamic term structure model (DTSM), which has a fixed endpoint, and OSE is a model with an observed shifting endpoint that uses the proxy $i_t^* = \pi_t^* + r_t^{*,RT}$. The data are quarterly from 1971:Q4 to 2018:Q1. Top panel: forecasts for horizons from 4 to 40 quarters made each quarter from 1976:Q3 (once five years of data are available) to 2008:Q1 for a total of 127 (overlapping) observations. Bottom panel: forecasts for horizons from 1 to 5 years (yearly averages of quarterly forecasts) made each quarter preceding a Blue Chip survey (long-range forecasts in Blue Chip Financial Forecasts) from 1988:Q1 to 2011:Q4 (48 surveys). Models are re-estimated each quarter up to 2007:Q4 (expanding window). Last rows in each panel report one-sided p-values for testing the null hypothesis of equal forecast accuracy against the alternative that method OSE is more accurate, using the method of Diebold and Mariano (1995) with a small-sample correction.

Figure 1: Ten-year yield and macroeconomic trends



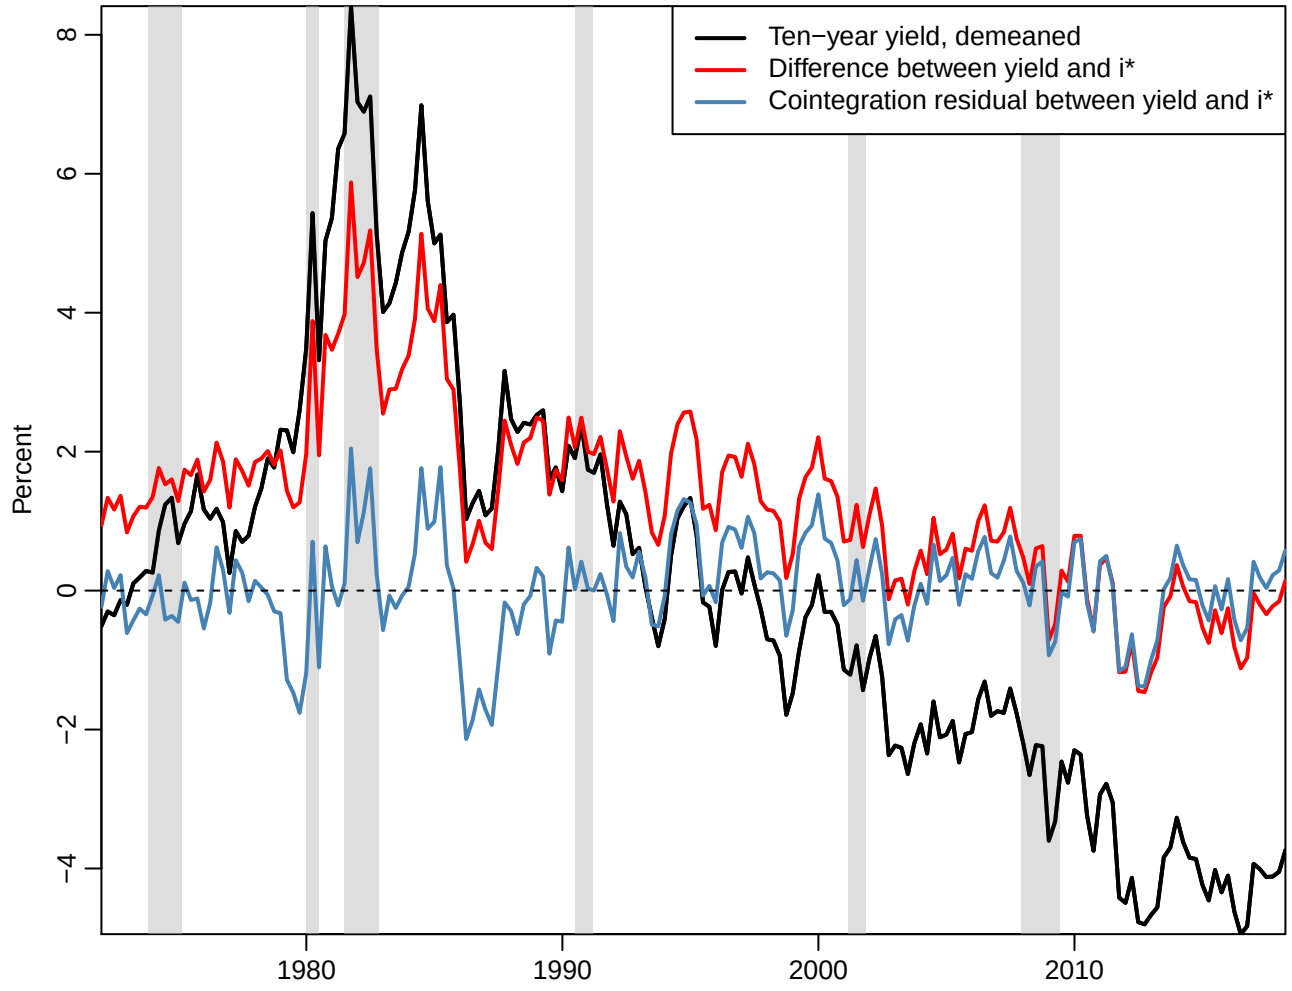
Ten-year Treasury yield and estimates of trend inflation, π_t^* (the mostly survey-based PTR measure from FRB/US), the equilibrium real rate, r_t^* (an average of all filtered and real-time estimates, see Section 2), and the equilibrium short rate, $i_t^* = \pi_t^* + r_t^*$. Shaded areas are NBER recessions. The data are quarterly from 1971:Q4 to 2018:Q1.

Figure 2: Measures of the equilibrium real interest rate



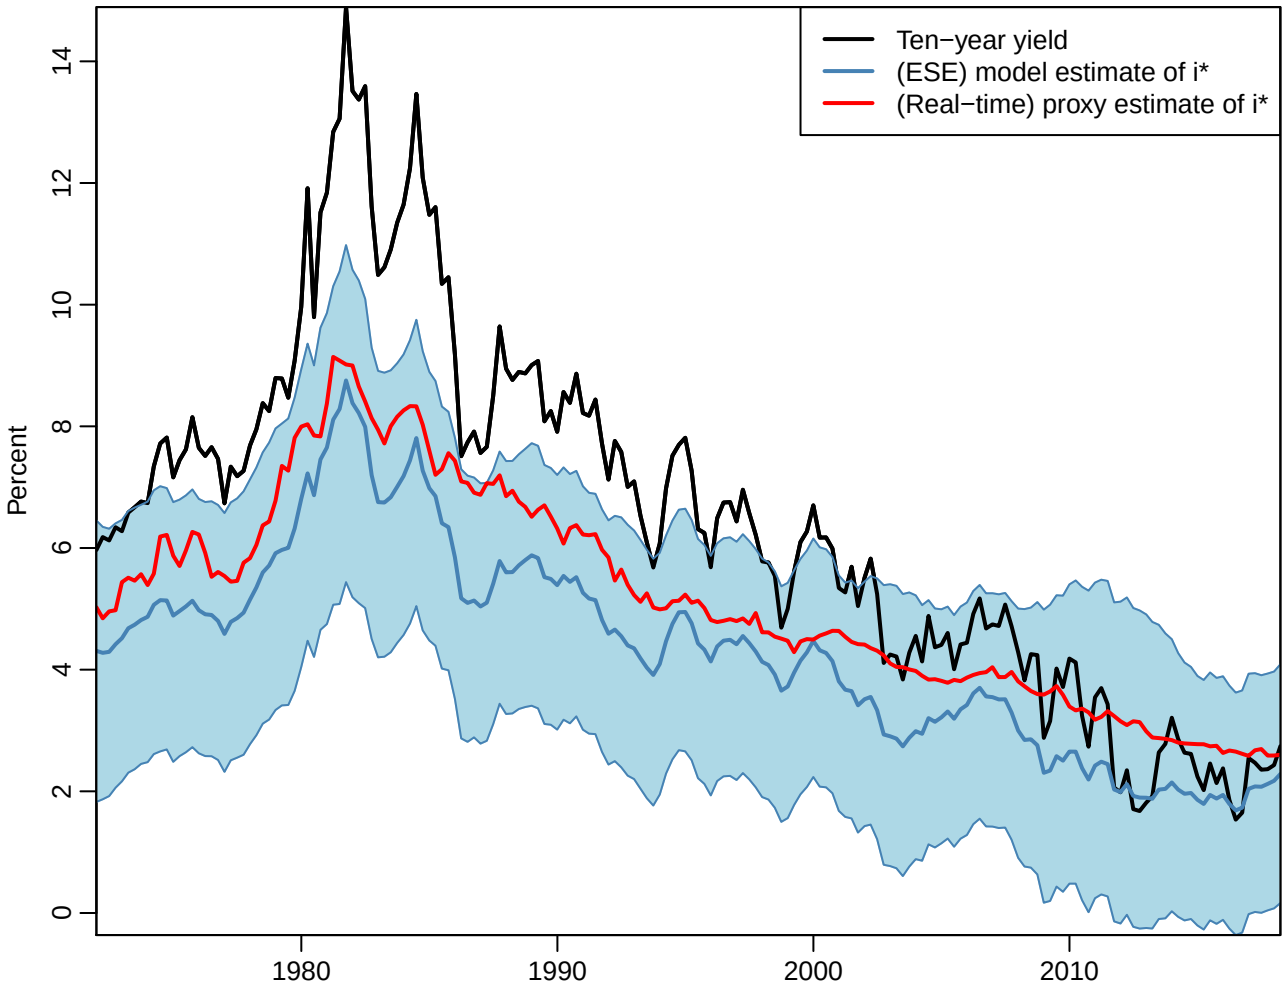
Smoothed (two-sided), filtered (one-sided), and real-time estimates of the equilibrium real interest rate, or long-run r_t^*. Left panel: range and average of seven smoothed estimates, including the four estimates by Del Negro et al. (2010) (DGGT), Johannsen and Mertens (2018) (JM), Laubach and Williams (2016) (LW), and Kiley (2015), as well as our own three model-based estimates. Right panel: range and average of filtered and real-time estimates. Filtered estimates are the estimates of LW, Holston et al. (2017), and Kiley. Real-time estimates are those of DGGT, JM, our own three real-time model-based estimates, and the model-free moving-average estimate, which is also shown separately. For details, see Section 2 and Appendix A. The data are quarterly from 1971:Q4 to 2018:Q1.

Figure 3: Detrending the ten-year yield



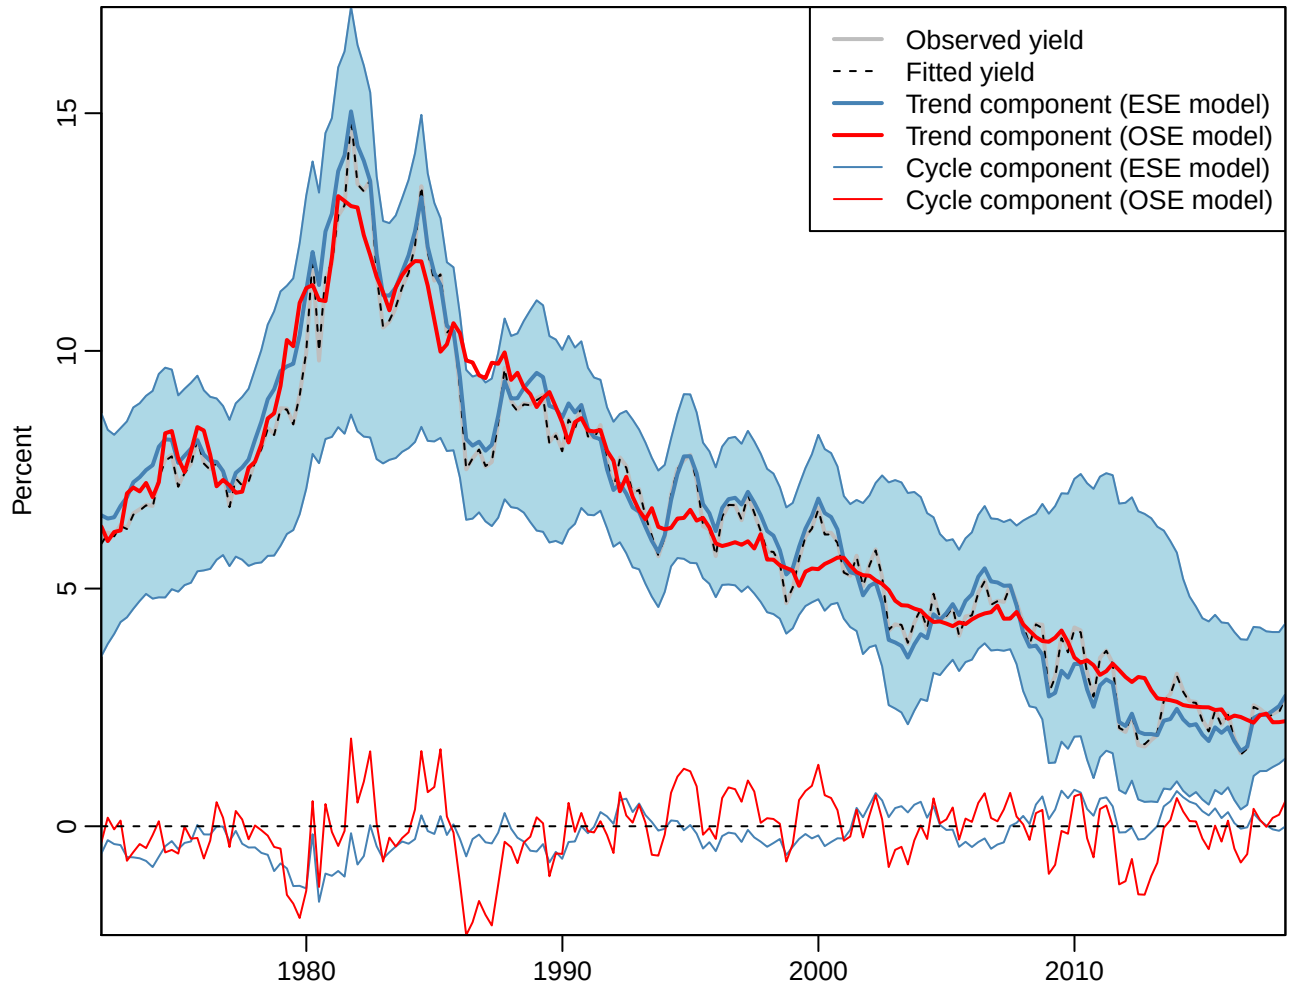
Ten-year Treasury yield and two detrended yield series using as the proxy for the equilibrium interest rate i_t^* the sum of the PTR estimate of π_t^* and the real-time estimate of r_t^* shown in Figure 2. The first detrended yield series is the difference between the yield and i_t^*. The second detrended series is the residual from the (Dynamic OLS) cointegration regression of the yield on i_t^*, estimates of which are reported in the last column of Table 1. Shaded areas are NBER recessions. The data are quarterly from 1971:Q4 to 2018:Q1.

Figure 4: Model-based estimate of equilibrium interest rate



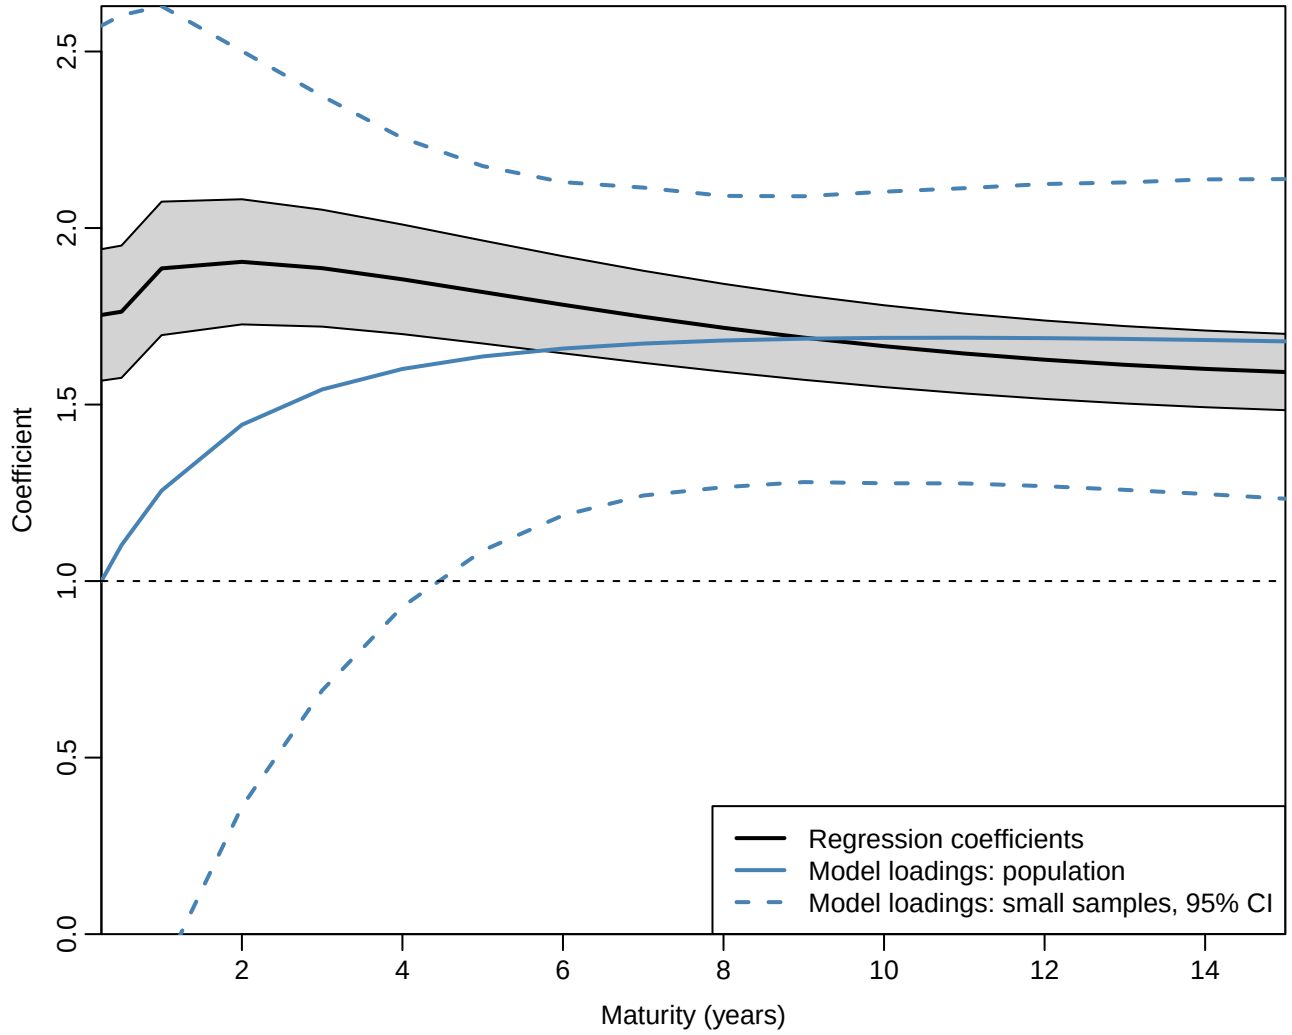
Bayesian (MCMC) estimate—posterior mean and 95%-posterior credibility intervals—of the equilibrium nominal short rate, i^*, from the "estimated shifting-endpoint" (*ESE*) dynamic term structure model. Also shown is a proxy estimate of i^*, the sum of the PTR estimate of π^* and the real-time macro estimate of r^*. The data are quarterly from 1971:Q4 to 2018:Q1.

Figure 5: Trend and cycle components of ten-year yield



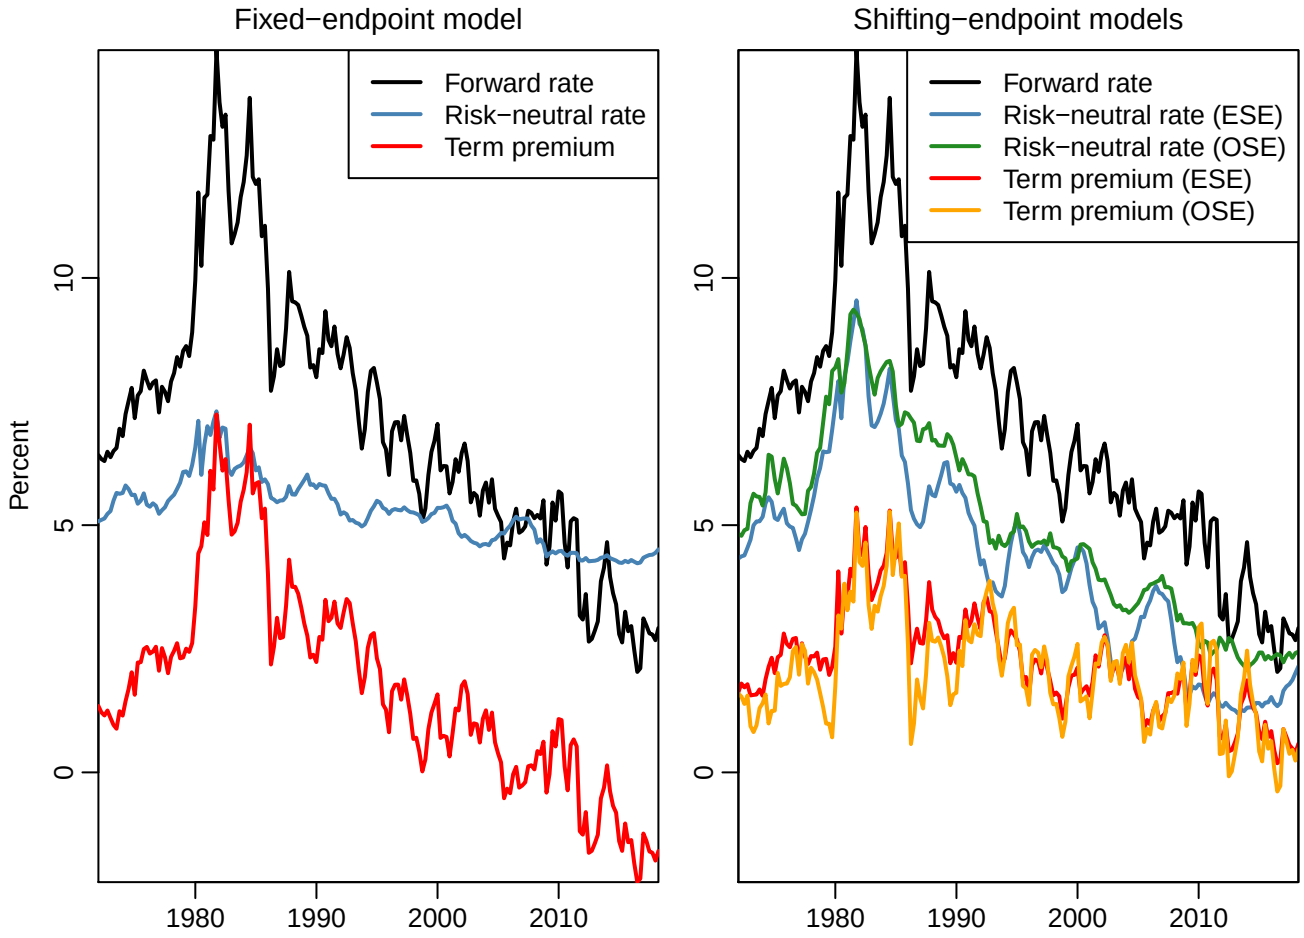
Model-based decomposition of ten-year yield into trend and cycle components. Blue lines show trend and cycle (posterior means) for the model with estimated shifting endpoint (*ESE*), shaded area shows 95%-posterior credibility intervals for trend component; red lines show trend and cycle for model with observed shifting endpoint (*OSE*). The data are quarterly from 1971:Q4 to 2018:Q1.

Figure 6: Loadings of yields on the equilibrium interest rate



Comparison of loadings of Treasury yields on the equilibrium nominal short rate i_t^* in the data and in the shifting-endpoint dynamic term structure model. Data: regression coefficients (with 95%-confidence intervals) for cointegration regressions of yields on the proxy-estimate of i_t^* (the PTR estimate of π^* plus the real-time macro estimate of r^*). Model: loadings of yields on i_t^* in population and 95%-Monte Carlo intervals for regression coefficients in small samples simulated from the model with an observed shifting endpoint (OSE).

Figure 7: Expectations and term premium components in long-term interest rates



Five-to-ten-year forward rate with estimated expectations component (risk-neutral rate) and term premium. Left panel: conventional dynamic term structure model (DTSM) with a fixed endpoint i^* (Joslin et al., 2011). Right panel: DTSM with a common stochastic trend, where the shifting endpoint is either estimated (*ESE*, estimated shifting endpoint) or taken as observed using a proxy of i_t^* (*OSE*, observed shifting endpoint). The data are quarterly from 1971:Q4 to 2018:Q1.