FEDERAL RESERVE BANK OF SAN FRANCISCO

WORKING PAPER SERIES

Interest Rates Under Falling Stars

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

January 2019

Working Paper 2017-16
http://www.frbsf.org/economic-research/publications/working-papers/2017/16/

Suggested citation:

Bauer, Michael D., Glenn D. Rudebusch. 2019. "Interest Rates Under Falling Stars," Federal Reserve Bank of San Francisco Working Paper 2017-16. https://doi.org/10.24148/wp2017-16

Interest Rates Under Falling Stars

Michael D. Bauer and Glenn D. Rudebusch*

Federal Reserve Bank of San Francisco

December 27, 2018

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 for 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 their long-run trends.[2] 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 that these long-run trends are constant. 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 standard time series methods, we establish some stylized facts about the connections between macro trends and interest rate trends as well as the importance of such trends for predicting excess bond returns. For this largely model-free analysis, we consider a variety of empirical proxies for the macro trends, including both off-the-shelf measures and our own estimates. Second, we introduce a new dynamic term structure model with a stochastic interest rate trend in order to account for and elucidate these stylized facts. With this model, we can estimate the long-run trend in Treasury yields, provide new estimates of the term premium, and predict interest rates out of sample quite accurately. Our paper bridges an important gap between macroeconomics and finance by taking shifting long-run trends seriously.

A simple illustration of the potential importance of macro trends is provided in Figure 1. The secular decline in the 10-year Treasury yield since the early 1980s reflects a gradual downtrend in the general level of U.S. interest rates. The underlying drivers of this decline remain contentious. In finance (no-arbitrage) models of the yield curve, the decline is mostly attributed to the term premium, which is the residual difference between a long-term interest rate and the expectations component (the expected average future short-term rate). This

[1] See, among many others, Ang and Piazzesi (2003), Diebold et al. (2006), Rudebusch and Wu (2008), Bikbov and Chernov (2010), Rudebusch and Swanson (2012), Bansal and Shaliastovich (2013), and Joslin et al. (2014). For a detailed survey, see Gürkaynak and Wright (2012).

[2] Throughout this paper, we use the Beveridge-Nelson concept of a trend, that is, the expectation for an economic series in the (infinitely) distant future.

is because these models assume interest rates are stationary, mean-reverting processes, so expected short-term rates will be close to their mean for horizons more than a few years out in the future. Therefore, the model-implied expectations component of long-term interest rates ends up being quite stable, and low-frequency variation in interest rates must reflect fluctuations in the term premium (e.g., Wright, 2011). Instead, our estimates of the trends underlying interest rates displayed in Figure 1 suggest a different explanation (details of these trend estimates are provided in the next section). First, our measure of U.S. trend inflation, π_t^*, which is based on long-horizon inflation survey forecasts, declined by almost six percentage points from the early 1980s to the late 1990s. Second, as inflation expectations stabilized over the past two decades, the estimated equilibrium real interest rate, r_t^*, which captures expectations of future real short-term interest rates, has fallen. The long-run trend in nominal yields is driven by the equilibrium nominal short-term interest rate, which is the sum of both macro trends, that is, $i_t^* = \pi_t^* + r_t^*$. As evident in Figure 1, this estimate of i_t^* exhibited similar low-frequency movements as the ten-year Treasury yield. This correspondence 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.

While the example illustrated in Figure 1 takes the macro trends as given, the fundamental difficulty for investigating the link between macro trends and interest rates is that the trends are unobserved and, particularly in the case of r_t^*, quite difficult to estimate. To address this issue, we use a variety of approaches and estimates to ensure that our findings are robust. Specifically, for r_t^*, we consider existing estimates from Laubach and Williams (2003), Kiley (2015), Holston et al. (2017), Del Negro et al. (2017), and Johannsen and Mertens (2018), and 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^*. To avoid look-ahead bias in our analysis of predictability, we focus our attention on trend measures that are constructed only with data that would have been available in real time.

Our first stylized fact from a model-free analysis is that variation in both the inflation trend and the equilibrium real interest rate is responsible for the extremely high persistence of interest rates. It has long been recognized that nominal interest rates contain a slow-moving trend component (Nelson and Plosser, 1982; Rose, 1988; King et al., 1991), and Campbell and Shiller (1987) find support for the presence of a single common trend, implying 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. To account for this missing persistence, we show that it is crucial to include movements in r_t^* to fully capture the trend component in interest rates. Cointegration regressions of yields on macro trend proxies uncover a tight long-run relationship, and the associated estimated error-correction models reveal that yields fairly quickly revert back to their trend.

Our second set of stylized facts uses predictive regressions for excess bond returns to assess how important macro trends are for bond risk premia. Relative to the canonical benchmark specification that predicts future excess bond returns with the current yield curve (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.[3] Inclusion of the equilibrium real rate is especially important later in our sample when the inflation trend shows less variation, as evident in Figure 1.[4] The predictive gains are economically large: a decline of one percentage point in i_t^* predicts an increase in the future quarterly excess returns of about 0.5 percentage point. For example, when interest rates adjust to a lower trend, long-term bond holders benefit, just as they did during the adjustment to a lower r_t^* following the 2008 financial crisis. Finally, we show that a parsimonious way to capture the power of i_t^* in the excess return predictive regressions is to use detrended yields. 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^*.

After documenting these two broad sets of stylized facts, we then formulate a new model to explain them. Specifically, we introduce a 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 most existing DTSMs—e.g., Kim and Wright (2005), Joslin et al. (2011) and Adrian et al. (2013) among many others—which impose stationary dynamics. Even though interest rates may be bounded over the very long run, it is often desirable in general to model a highly persistent but stationary time series as integrated (Campbell and Perron,

[3] As discussed further below, our findings provide evidence against the *spanning hypothesis*, which is 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).

[4] The recent reduced variation in the inflation trend explains why the performance of the regressions of Cieslak and Povala (2015), who predict bond returns using a moving average of past inflation, has diminished over time.

1991).[5] For example, a unit root is often assumed when modeling price inflation because of the resulting improvements in forecast performance (Stock and Watson, 2007; Faust and Wright, 2013). In this spirit, we adopt a shifting endpoint (stochastic trend) in a model of nominal interest rates. In our model, the long-run trend is *unspanned*, so that it cannot be backed out from the cross section of yields. This specification allows for a well-behaved, slow-moving trend. Unfortunately, the fact that the trend is unobserved makes estimation somewhat difficult—a common problem for unobserved components models. Therefore, for robustness, we employ two different estimation methodologies, which encouragingly provide similar empirical representations and results. First, we estimate the model using an observed proxy for i_t^* in order to carry out a model-based analysis using the real-time interest rate trend employed in our cointegration and excess return regressions. Second, we use Bayesian estimation of our model to infer i_t^* using only Treasury yield data, imposing a certain degree of smoothness on our trend estimate, similar to Del Negro et al. (2017).

Given these model estimates, we then show that simulated data from the models are able to match the stylized facts that we established about the role of trends for the yield curve. First, the model-based estimates of the trend and cycles components of interest rates are consistent with the cointegration regression results from the actual data. Second, 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, as in the postwar data sample, predictive regressions using simulated data from our model demonstrate notable increases in R^2 when the interest rate trend is added as a predictor. Therefore, both in the model and in the actual data, trends contain important predictive information.

We also use our shifting-endpoint model to investigate two additional issues that are among the most 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 compared with the popular conventional model by Joslin et al. (2011) in which yield factors follow a stationary VAR(1). This standard decomposition, which is representative of a large class of stationary DTSMs, 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). Linking macro

[5] While our model contains a unit root under the real-world probability measure, the largest root under the risk-neutral measure is less than one to avoid violations of no-arbitrage.

trends to the yield curve solves the knife-edge problem of Cochrane (2007), who noted that assuming either a stationary or random walk level of interest rates leads to drastically different implications for the term premium, in each case arguably quite unsatisfactory. Assuming a common, slow-moving underlying trend instead leads to more plausible decompositions of long-term rates.

Out-of-sample forecasting of interest rates is a second important application of DTSMs. In the past, 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 the forecasts implied by our shifting-endpoint DTSM are generally better than this martingale benchmark. By contrast, the standard stationary model (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. Our forecasts also beat long-range projections from the Blue Chip survey of professional forecasters. 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^*.

Our analysis is related to several different strands of the macro and finance literatures. Some studies have found 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 downtrend trend in yields over the past 20 years. Time variation in r_t^* has so far received little attention in finance, despite the growing macroeconomic evidence regarding substantial changes in the equilibrium real interest rate. The underlying economic forces affecting 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 and reduced desired investment and, in turn, pushed down the steady-state real interest rate.[6] 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 to

[6]Discussions and estimates of the decline in r^* 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 the macroeconomics literature, r_t^* is often labeled the neutral or natural rate of interest although, as noted below, there are various definitions with subtle differences.

capture the extreme persistence of interest rates, and a variety of remedies have suggested.[7] Our results suggest that allowing for shifting long-run trends may be a better solution. Some recent 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 a popular example of 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). The model we propose in this paper allows for stochastic macroeconomic trends to affect long-run expectations of interest rates via changes in the equilibrium interest rate i_t^*.

Three recent papers are very 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^*, and we include their full-sample and real-time estimates of r_t^* in the first stage of our analysis. Our DTSM also allows for shifting endpoints, and in one implementation, we use a Bayesian method and smoothness prior for the trend that follows their work. However, in contrast to their time series model, 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. 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—i_t^*, r_t^* and π_t^*—in a common-trend VAR for interest rates and inflation, using surveys of professional forecasters to pin down model-implied expectations. While their model of the yield curve is not arbitrage-free, it takes the important step of allowing for shifting endpoints in interest rate expectations.

The paper is structured as follows: In Section 2 we explain the main concepts, including the long-run Fisher equation, and describe our proxies for each of the macroeconomic trends and our story of falling stars. Section 3 presents stylized facts about the role of macroeconomic trends for 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.

[7] 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

Conceptually, the persistent components in inflation and the real interest rate will be reflected in long-term interest rates, but it remains an open question whether this link between macro trends and the yield curve matters in practice. This section lays the groundwork for answering this question, introducing the concepts and different empirical estimates of π_t^*, r_t^*, and i_t^*.

2.1 Definition of the interest rate trend

The role of trends for interest rates can be illustrated with the usual decomposition of any n-period interest rate into expectations of short-term rates and a term premium residual:

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

where $y_t^{(n)}$ is the yield on an n-period government bond and i_t is the nominal short rate. The yield term premium, $TP_t^{(n)}$, in theory compensates bond investors for the duration risk in longer-term bonds. In practice, it captures the effects of all factors and financial market frictions other than the expectations component. The expectations theory of interest rates assumes that TP_t^n is zero (strong expectations hypothesis) or constant (weak expectations hypothesis), 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) further decomposes the expectations component into short rate trend, i_t^*, and short rate cycle $i_t^c = i_t - i_t^*$. We define this trend in the short-term rate, the *equilibrium nominal short rate*, as the Beveridge-Nelson trend,[8]

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

Yields of all maturities contain i_t^*, and if this trend is time-varying then interest rates are non-stationary and have 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 would be expected to drive the 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 the real rate and inflation

[8]Throughout we assume that there are no deterministic time trends in the variables of interest.

are defined analogously to i_t^*. The main observation is the following: 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.

2.2 Measures of trend inflation: π_t^*

According to the prevailing consensus in empirical macroeconomics, in order to produce accurate forecasts and correctly capture the evolution of expectations, inflation is best modeled as a first-order integrated, $I(1)$, process (Stock and Watson, 2007; Faust and Wright, 2013). That is, π_t^* should be allowed to vary over time. From a macroeconomic perspective, this time-varying inflation endpoint can be viewed as the perceived inflation target of the central bank.

Empirical proxies for trend inflation, π_t^*, 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 survey-based measure, namely, the Federal Reserve's series on the perceived inflation target rate, denoted PTR. It measures long-run expectations of inflation in the price index of personal consumption expenditures (PCE), and is widely used in empirical work—see, for example, Clark and McCracken (2013) and Del Negro et al. (2017).[9] Figure 1 shows that from the beginning of our sample to the late 1990s, this estimate mostly mirrored the increase and decrease in the ten-year yield. Since then, however, it has been essentially flat at two percent, which is the level of the longer-run inflation goal of the Federal Reserve that was first announced in 2012. Other survey expectations of inflation over the longer run, such as the long-range forecasts in the Blue Chip survey, exhibit a similar pattern.

In their analysis of inflation trends and yields, Cieslak and Povala (2015) used a moving average of past core inflation.[10] As they noted, this trend estimate has generally co-moved closely with PTR. However, since 2004, low inflation has pushed the moving-average measure of π_t^* below PTR and other survey-based longer-run inflation projections. We will use PTR in our analysis, which is consistent with many model-based estimates (e.g., Christensen et al., 2010; Clark and Doh, 2014; Stock and Watson, 2016) and the official two percent inflation target of the Federal Reserve.

[9]Since 1979 (i.e., for most of our sample), PTR corresponds to long-run inflation expectations from the Survey of Professional Forecasters. Before 1979, PTR 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`.

[10]Specifically, they use the recursion $\pi_t^* = \alpha\pi_{t-1}^* + (1-\alpha)\pi_t$ with $\alpha = 0.987$ and π_t defined as monthly year-over-year inflation in the consumer price index excluding food and energy items.

Given our focus on forecasting, real-time availability is an important issue. For most of our sample, the PTR estimate was available in real time: Since 1979, it has been available in the survey releases at the end of each quarter, which is when our yields are observed.

2.3 Measures of the equilibrium real interest rate: r_t^*

The relevant definition of r_t^* for our analysis of the term structure of interest rates is the long-run trend in the one-period real interest rate, analogous to the definition of i_t^* in equation (2). Accordingly, changes in r_t^* signal the presence of a stochastic trend in the real rate, consistent with evidence that the real rate is $I(1)$ (see, for example, Rose, 1988). Various types of structural economic changes, such as shifts in the trend growth rates of the population or productivity, affect the equilibrium real rate, as investigated in a large recent literature. Two studies that propose appealing new methods to estimate this long-run r_t^* are Del Negro et al. (2017), who estimate a linear state space model with common trends, and Johannsen and Mertens (2016, 2018), who employ a non-linear state space model including a shadow rate. We include both of these estimates in our analysis.[11]

A closely related concept of r_t^*, often referred to as the *neutral* real interest rate, is the level of the real rate at which monetary policy would be neither expansionary nor contractionary.[12] Laubach and Williams (2003, 2016) (LW) implemented this concept empirically using the simple Rudebusch and Svensson (1999) macroeconomic model. While this definition of r_t^* is theoretically somewhat different from the long-run concept we employ, in practice, the LW "longer-run" estimates of r_t^* are effectively comparable to a long-run trend.[13] Accordingly, in some of our analysis, we consider the LW estimate and two closely related estimates from variants of the LW model: the model of Holston et al. (2017) (HLW), which is a slight simplification of the LW model, and the model of Kiley (2015), which accounts for changes in financial conditions.[14]

In addition to these five existing estimates of r_t^*, we also construct three model-based estimates of our own in order to help address real-time estimation issues and ensure the robustness of our results. The first is a simple univariate unobserved components model

[11] 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.

[12] Another definition of r_t^*, which is a short-run concept and therefore less relevant to our purpose here, is the *natural* or *efficient* real interest rate in dynamic stochastic general equilibrium models. This is the real rate that would prevail in the absence of nominal frictions and is often assumed to be stationary (see, for example, Cúrdia et al., 2015).

[13] Appendix A.1 provides a more detailed discussion. For example, LW assume that r_t^* is a martingale.

[14] Specifically, 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 the Kiley model, the IS curve is augmented with credit spreads.

similar to Watson (1986), which decomposes the real short rate into a random walk trend (r_t^*) and a stationary component. Recently, Fiorentini et al. (2018) have argued that such univariate models can infer r_t^* with greater precision than the LW model. The second model augments the first with information about r_t^* from macroeconomic trend proxies that recent work has shown to be important correlates of the real rate trend. Specifically, we include moving averages of past growth of real GDP and of labor force hours, motivated by Lunsford and West (2017) who documented strong and robust low-frequency co-movement of these two macro variables and the real interest rate. Our third model specifies both real rate and inflation processes to obtain an ex-ante real rate using data for the T-bill rate, the quarterly inflation rate, and the long-run inflation expectations measure PTR. We estimate all three state-space models using Bayesian methods, and like Del Negro et al. (2017) we impose a smoothness prior on r_t^*, i.e., a tight prior centered around low values for the trend innovation variance. Finally, we also include a simple, model-free estimate of r_t^*: an exponentially-weighted moving average of past real rates.[15] Appendix A provides an overview of all trend estimates, implementation details for our estimation, and figures with the time series for each individual estimate.

Figure 2 summarizes these various measures of r_t^* by displaying select ranges and averages. The left panel shows the range and average of the available smoothed (two-sided) estimates, for which inference about both the parameters and the state variables uses the full data sample. We have seven smoothed estimates are available, for all models but the HLW model. In retrospect, it appears that r_t^* has remained relatively stable until the late 1990s, but then exhibited a pronounced and sustained decline over the past 20 years—a pattern noted by other studies.

When analyzing the relationship between macro trends and the yield curve, we are interested in what investors could have known about macro trends contemporaneously. Our analysis will therefore focus on filtered (one-sided) and real-time estimates of r_t^*, which are displayed in the right panel of Figure 2. Filtered estimates infer latent state variables using data available up to time t, but condition on full-sample parameter estimates; these include the LW, HLW, and Kiley models. These estimates exhibit large disagreement in the early part of the sample, move mostly sideways from the mid-1980s to the late 1990s, and then decline substantially over the last two decades. A sudden and pronounced decline occured in 2008 during the financial crisis, followed by about a decade of very low levels of r_t^*.

Given our interest in assessing the ability of various models to predict interest rates and

[15] For the smoothing parameter we use a value of 0.98, in line with those used in other work estimating macro trends. For example, Malmendier and Nagel (2011) and Orphanides and Williams (2005) also use $\alpha = 0.98$, based on calibration. The smoothing parameter that Cieslak and Povala (2015) used to estimate their inflation trend translates to about 0.95 in quarterly data, but with this value a large amount of business cycle variation still remains in the r_t^* estimate.

bond returns, the pitfall of look-ahead-bias can be particularly serious. For example, the filtered estimates from the LW-type models are problematic for this purpose because the parameters are estimated based on the full data, i.e., with information that was not available contemporaneously to investors.[16] We therefore consider (pseudo) real-time estimates where inference about both parameters and r_t^* uses data only up until period t. These include the estimates from the Johannsen and Mertens (2016) and Del Negro et al. (2017) models, our own three models, and the moving average trend estimate. Although our models are estimated in real time, most of them are nevertheless recent inventions that investors didn't have access to in real time. To address this issue, we include a simple moving-average estimate of r_t^* that does not require filtering through complicated time series and macro models and was readily available to investors at any point in our sample period.[17]

The average of our six real-time estimates increased from about one to two 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. The moving-average estimate itself displays a very similar pattern. Despite common difficulties with real-time estimation of r_t^*, as noted for example by Clark and Kozicki (2005), the secular decline since the turn of the millenium becomes quite apparent once a variety of models and estimates are considered. These patterns suggest that the long-run value for the real interest rate as perceived by bond investors likely drifted down significantly over the past two decades.

2.4 A story of falling stars

Our macro trend estimates are consistent with a compelling narrative about the evolution of long-term nominal interest rates as illustrated in Figure 1, which shows the ten-year Treasury yield, the PTR estimate of trend inflation, an r_t^* measure that is the average of all nine filtered and real-time estimates described above, and the sum of these two macro trend estimates as a proxy for i_t^*. Interest rates and long-run inflation expectations spiked in the late 1970s, but with the Volcker disinflation of the 1980s, they both trended down together. Around the turn of the millennium, long-run inflation expectations stabilized near 2 percent, but i_t^* and long-term interest rates continued to decline, in part because structural changes in the global economy started pushing down the equilibrium real rate. The analysis in the following

[16] Laubach and Williams (2016) showed, however, that truly real-time estimation of their model delivers an estimated series of r_t^* that is close to their final revised estimate over the period that both are available, suggesting that real-time estimation of their model would likely yield similar results.

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

sections investigates whether the link between macro trends and the yield curve that underlies this narrative is supported by the empirical evidence, and whether accounting for shifts in i_t^* alters our interpretation of interest rate movements and bond risk premia.[18]

Before proceeding with our empirical analysis, which mostly treats these trend estimates as data, a note on the issue of measurement error is in order. Of course, 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. We will show that our measures of the macro trends are closely connected to the yield curve and contain important information for predicting future yields and returns, despite the measurement error that likely works against finding such links. For example, if the trends are subject to classical measurement error, that would make the coefficients in our regressions both less precise and bias them toward zero. Alternatively, because our trend proxies are estimates of the true trends using all available information, the measurement error is more likely to be orthogonal to our trend estimates (instead of being orthogonal to the true trend), which would make our estimates noisy but not necessarily biased (Mankiw and Shapiro, 1986; Hyslop and Imbens, 2001). In either case, because of the presence of measurement error, our results could be viewed as a lower bound for the tightness of the connection between the yield curve and the *true* underlying 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.

3 Macro trends and yields: stylized facts

In this section, we document the empirical significance of macro trends in the dynamics of the Treasury yield curve in several ways. Our aim is to establish stylized facts about the connection between these trends and the yield curve. These facts include the persistence, cointegration, and error-correction behavior of interest rates as well as the incremental predictive power of macro trends for future bond returns.

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

[18] An alternative for estimating the market-perceived equilibrium short rate i_t^* is to use survey measures, 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 starting in 1986, fluctuates around our baseline estimate of i_t^*. We consider survey forecasts in Section 4.6 below.

addition, our proxies for macroeconomic trends are the PTR estimate of π_t^* and the three alternative estimates of r_t^* described and shown above: the average of the three filtered estimates, $r_t^{*,F}$, the average of the six real-time estimates, $r_t^{*,RT}$, and a simple real-time estimate calculated as the (exponential) moving average of past real rates, $r_t^{*,MA}$. As a proxy measure of the equilibrium nominal short rate we use $i_t^* = \pi_t^* + r_t^{*,RT}$ to avoid look-ahead bias; however, we obtained similar results using the filtered and moving-average measures of the real rate trend in our estimate of i_t^*.

3.1 Persistence, unit roots, and cointegration

We first investigate whether the trend components of inflation and the real interest rate play an important role in driving movements in the yield curve, that is, whether these macro trends account for the high persistence of interest rates. A related question is whether changes in r_t^* materially contribute to movements in i_t^* and the persistence in bond yields, or whether accounting for π_t^* alone is sufficient.

The essence of the key result is illustrated in Figure 3, which shows the (demeaned) ten-year yield and two detrended yield series. The first detrended yield series is the difference between the yield and our estimate of i_t^*. In comparison to the ten-year yield, this series exhibits much less of a trend, declining, for example, 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 noticeable trend remains in the difference between the yield and i_t^*, which indicates that the trend component of the yield moves more than one-for-one with i_t^*. Similarly, this behavior suggests that the term premium is to some extent also driven by changes in i_t^*. The second detrended series is the residual from a (Dynamic OLS) cointegration regression of the yield on i_t^*, estimates of which are discussed below and reported in the last column of Table 2. This cointegration residual shows no discernible trend, indicating that changes in i_t^*, properly scaled, can fully account for the trend in the ten-year Treasury yield. The residual can be understood as an interest rate "gap" or "cycle" (Cieslak and Povala (2015) denote the residuals in regressions of yields on π_t^* as "interest rate cycles"). The fitted value of the regression is the trend component of the ten-year yield, for which we also obtain a model-based estimate in Section 4.3. The key point here is that changes in i_t^* account for the persistent component in interest rates, and that we can separate trend and cycle components of nominal interest rates using i_t^*.

Table 1 documents the persistence of the ten-year Treasury yield, denoted $y_t^{(40)}$, and the macro trends.[19] These trends include the inflation trend, π_t^*, the three estimates of the real rate

[19]These results and others throughout this section are robust to using other longer-term interest rates or the overall level of the yield curve (the first principal component of yields).

trend—$r_t^{*,F}$, $r_t^{*,RT}$, and $r_t^{*,MA}$—as well as $i_t^* = \pi_t^* + r_t^{*,RT}$. 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—and the low-frequency stationarity test (LFST) of Müller and Watson (2013), which is reported as a p-value for the null of stationarity.[20]

The ten-year yield is very persistent, with a first-order autocorrelation coefficient of 0.97 and a half-life of around 26 quarters. The macro trends are even more persistent. Our real-time estimate of r_t^* has an autocorrelation coefficient of 0.98 and a half-life of about 37 quarters. The inflation trend and i_t^* have autocorrelation coefficients of 0.99 and half-lives of 85 and 60 quarters, respectively. Neither the ADF or PP tests reject a unit root, while the LFST p-values reject stationarity for each series. In sum, interest rates and macro trends are highly persistent and can be effectively modeled as $I(1)$ processes. The key question is whether this persistence is driven by the same underlying trend, that is, whether there is a cointegration relationship between long-term rates and the macro trend estimates. As a simple first step, we consider the properties of differences between yields and the macro trend proxies. Subtracting π_t^* from the 10-year yield removes a substantial amount of the persistence, as evident from the autocorrelation coefficients and half-lives reported in the second line of Table 1. Still, the resulting series is 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 often favors stationarity, depending on the test and specific real rate trend proxy. 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 below finds an even stronger macro-finance link when the cointegration coefficients are slightly larger than one.

To formally investigate the relationships between the macro trends and interest rates, including the individual contributions of π_t^* and r_t^*, we consider cointegration regressions. This regression is $y_t^{(40)} = \beta_0 + \beta_1' X_t^* + u_t$, where the vector X_t^* contains one or two macro trend proxies, and u_t is a residual that is stationary if there is cointegration. We estimate β_0 and β_1 using the Dynamic OLS estimator of Stock and Watson (1993), by including four

[20] 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.

leads and lags of $\Delta y_t^{(40)}$ and ΔX_t^*.[21] Table 2 reports the coefficient estimates and Newey-West standard errors (using six lags), the R^2 for this regression, as well as persistence statistics for the cointegration residuals and cointegration tests.

The first specification in the first column of Table 2 includes only π_t^*, and the results parallel those obtained by Cieslak and Povala (2015) using simple OLS regressions and their proxy for the inflation trend. The inflation trend coefficient is somewhat higher than one and strongly significant, and the R^2 is very high. While Cieslak and Povala (2015) interpret these results as indicating that trend inflation drives the level of yield curve, it does not fully account for the trend behavior of yields. The same statistical measures of persistence used in Table 1 are applied to the regression residual $\hat{u}_t = y_t^{(40)} - \hat{\beta}_0 + \hat{\beta}_1 \pi_t^*$. This residual remains quite persistent, and there is some evidence from the ADF and LFST statistics that it contains a unit root.[22] Furthermore, a Johansen (1991) trace test does not reject the null of absence of cointegration among the yield and the inflation trend.[23] 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 inflation trend, indicates only a modest amount of error correction ($\hat{\alpha} = -0.11$), suggesting that any equilibrating forces pushing yields back in line with the inflation trend are not very strong.

The missing piece is the real-rate trend. The remaining columns of Table 2 show that once r_t^* is added there is strong evidence for cointegration. Again, for robustness, columns 2, 3, and 4 use the filtered (one-sided), real-time, and moving-average estimates of the real rate trend). In each case, the r_t^* and π_t^* coefficients are generally estimated to be of similar magnitude and strongly significant. Importantly, adding the real rate trend increases the R^2 by 8 to 11 percentage points, and the cointegration residual appears to be stationary. For the real-time and moving-average estimates of r_t^*, the evidence for this cointegration relationship is the strongest: in these cases the cointegration residual (see also Figure 3) is quickly mean-reverting, a unit root is rejected, 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 2 shows that i_t^* appears to be able to fully capture the trend component in the ten-year yield. The evidence for cointegration is as strong as when both macro trends are included individually. The cointegration coefficient is 1.67, which is somewhat high given that expectations of future short rates are affected one-for-one by i_t^*, but Section 4.3 will illuminate this issue further. The error-correction coefficient is estimated as -0.45 and

[21] Standard OLS as well as the reduced-rank maximum likelihood VAR estimation of Johansen (1991) lead to the same conclusions about the presence of cointegration, though we omit those results here for brevity.

[22] For these unit root tests on the 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.

[23] The test uses two lags in the VAR representation, based on information criteria.

is strongly significant. Accordingly, when the ten-year yield is high relative to i_t^*, it tends to fall back toward this trend over time, with almost half of the equilibrium error on average reversed within one quarter. Apparently, knowledge of this trend is quite important when predicting the future evolution of yields, a point that we investigate in more detail below by predicting bond returns.

Overall, these results suggest that the low-frequency trend behavior of Treasury yields is explained by movements in the underlying macroeconomic trends, and both the inflation trend and the real-rate trend play important roles. The evidence supports the presence of exactly one cointegration relationship among the ten-year yield, the inflation trend and the equilibrium real rate. This implies that detrending long-term rates by only π_t^* is insufficient. Furthermore, the coefficients on the two macro trends appear to be of equal magnitude, which implies that changes in i_t^* can fully account for the low-frequency variation in interest rates.

3.2 Predicting excess bond returns

Knowledge of the macroeconomic trends underlying yields, and specifically of i_t^*, appears to be important for understanding the evolution of bond yields. We now examine whether such trends can improve predictions of the excess return of long-term bonds over the risk-free interest rate. Expected excess returns capture bond risk premia and have long been of central interest in financial economics (e.g., Fama and Bliss, 1987).

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. We predict the average excess return for all bonds with maturities from two to 15 years, $\overline{rx}_{t+1}$. Our focus is on a holding period of one quarter to avoid the statistical problems that come with overlapping returns (Bauer and Hamilton, 2018; Cochrane and Piazzesi, 2005), but we also examined a holding period of 4 quarters and obtained qualitatively similar results to those reported here.[24] Since Fama and Bliss (1987) and Campbell and Shiller (1991), it is well-known that the yield curve, and in particular its slope, contains information useful 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 doc-

[24]Since we only have annual maturities for long-term bonds, we calculate one-quarter returns with the usual approximation $y_{t+1}^{(n-1)} \approx y_t^{(n)}$.

umented apparent violations of the spanning hypothesis using various additional predictors. Bauer and Hamilton (2018) demonstrated that inference in these predictive regressions suffers from serious small-sample econometric problems arising from highly persistent predictors, and that accounting for these problems renders most of the proposed predictors insignificant. They found, however, that the proxy for trend inflation of Cieslak and Povala (2015) was a relatively robust predictor. Here we investigate whether including both π_t^* and r_t^* leads to even stronger predictive gains and whether the use of detrended yields can fully uncover the predictive power in the yield curve.

Table 3 reports the 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 (PCs) of yields.[25] The specification in the second column just adds π_t^*, and columns three through five add as well one of our three proxies of r_t^* in order to capture the effects of both trends together. 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 avoid the problematic small-sample distortions noted above.[26] The top panel uses the full sample of data, while the bottom panel also 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; Bauer and Hamilton, 2018).

In the full sample, the inclusion of the inflation trend increases the predictive power substantially compared to using only yield-curve information: Both the inflation trend and the level of yields (PC1) appear highly significant (in addition to the slope, PC2). This parallels the findings of Cieslak and Povala (2015). However, adding r_t^* to the regressions leads to further impressive gains in predictive power, in particular for the moving-average estimate of r_t^*: The R^2 increases notably, the coefficients and significance for π_t^* and PC1 rise, and the coefficient on r_t^* itself is large and significant. Not surprisingly given Figure 1, these results shift in the more recent period as the real-rate trend has gained in importance over time relative to the inflation trend. In the subsample starting in 1985, the inflation trend is not statistically significant when included on its own according to the small-sample p-values. Only with the addition of the equilibrium real rate do both trends matter for bond risk premia; the coeffi-

[25]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 add up to one.

[26]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.

cients 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.[27]

The estimated coefficients have several interesting interpretations. First, the similar magnitude of the coefficients on the two individual macro trends suggests that their sum can capture the relevant information. Indeed, including as a predictor just i_t^* provides the same strong predictive gains. The key to forecasting excess bond returns is the inclusion of a measure of the overall trend in interest rates. Second, the coefficient on i_t^* has a negative sign, because if the trend rises then interest rates also rise in response, producing losses for long-term bond holders. Numerically, an increase in i_t^* by one percentage point predicts a decrease in next quarter's excess bond return by 0.45 percentage point. Third, the coefficient on i_t^* is larger in magnitude then the coefficient on the yield-curve level (PC1), consistent with our finding above that the underlying trend of interest rates is a slightly scaled-up multiple of i_t^*. Changes in the long-run nominal rate i_t^* ultimately have larger effects on bond returns than changes in the level of the yield curve.

In the presence of persistent predictors, it is generally difficult to interpret the magnitude of R^2 as a measure of predictive accuracy, because even predictors that are irrelevant in population can substantially increase R^2 in small samples (Bauer and Hamilton, 2018). We avoid this pitfall by using the bootstrap to generate small-sample distributions of R^2 under the spanning hypothesis (that is, under the null that the three PCs contain all the relevant predictive information) and interpret the statistics obtained in the actual data by comparing them to the quantiles of these distributions. The top panel of Table 5 reports this comparison for four of the specifications that we have considered so far. 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 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.

Our findings so far were based on regressions with highly persistent predictors, which causes econometric problems. While the bootstrap provides one way to address these problems and

[27]In additional, unreported results we have found that the predictive gains from including r_t^* stem mainly from the period since the early 2000s when both r_t^* and long-term interest rates decreased while long-run inflation expectations where anchored close to two percent.

allows us to obtain reliable small-sample inference, a simple alternative is available. Namely, we can use *detrended yields* as predictors, that is, residuals from regressions of yields on macro trends. Besides avoiding the pitfalls of persistent regressors, this approach has the added benefit that it reveals whether it is truly the reversion of interest rates to the underlying macro trends that causes the predictive power, as we have claimed above. We estimate predictive regressions in which the three predictors are PCs of the 17 detrended yields, and compare these to the baseline regression with PCs of the 17 original yields. Each detrended yield is obtained as the (cointegration) residual from a regression of the yield on macro trends, and as before we consider three types of detrending: using only π_t^*, using both π_t^* and $r_T^{*,RT}$, as well as using i_t^*. The regression estimates are shown in Table 4 and the R^2 of these regressions are compared to bootstrap intervals under the spanning hypothesis in the bottom of Table 5. In the full sample, all three detrending methods lead to substantial increases in R^2 over the yields-only baseline regression, and these increases in R^2 comfortably exceed the upper bound of the 95%-bootstrap distribution, indicating their statistical significance. The coefficients on $PC1$ and $PC2$ increase and become (more) statistically significant due to the detrending. However, detrending with i_t^* works particularly well, giving the largest gain of 12 percentage points, which is far outside the bootstrap interval. In the post-1985 sample, R^2 does not noticeably 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 results suggest the following conclusions. First, the regressions using detrended yields provide additional statistical evidence for the importance of macro trends in predictions of future bond returns. Second, the explanation for the observed predictive gains is indeed the trend-reversion of interest rates, i.e., Granger-causation from underlying macro trends to future yields. Third, it is again important to not only account for the inflation trend but for the complete i_t^* trend.

In sum, accounting for the persistent components of yields is important for understanding return predictability and estimating bond risk premia. We find that r_t^* has strong incremental predictive power for bond returns, about on par with the importance of π_t^* as a predictor, suggesting that both macro trends need to be accounted for accurate estimation of bond risk premia. These results are strong evidence against the spanning hypothesis. The predictive power in the yield curve can be fully revealed if yields are detrended, but it is crucial to use i_t^* instead of π_t^* for that detrending. These results document the economic significance of the macro-finance low-frequency linkage and confirm the influence of trends on yields uncovered in Section 3.1.

4 A no-arbitrage model with a stochastic trend

We now introduce and estimate a DTSM with a shifting endpoint for interest rates. While existing DTSMs generally assume stationarity and therefore impose that i_t^* is constant, our new formulation incorporates a common stochastic trend. This novel specification allows us to explain and interpret the stylized facts documented above, estimate the term premium, and accurately forecast interest rates out-of-sample. This demonstrates how allowing for time variation in i_t^* is crucial for understanding important aspects of interest rate dynamics.

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).[28] The key feature of our model is the presence of a single stochastic trend τ_t common to the three factors:

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

where γ is an N-vector with the loadings on the common trend, and Φ is a mean-reversion matrix that has eigenvalues with modulus less than one. Accordingly, the state variables $Z_t = (\tau, P_t')'$ are cointegrated (their VAR and VECM representations are given in Appendix B.1). Shocks to the trend τ_t are $\eta_t \overset{iid}{\sim} N(0, \sigma_\eta^2)$, shocks to the cycles $\tilde{P}_t$ are $\tilde{u}_t \overset{iid}{\sim} N(0, \widetilde{\Omega})$, and we will assume that $E(\eta_t\tilde{u}_t) = 0$, an assumption commonly made in unobserved components models (Watson, 1986; Del Negro et al., 2017). Innovations to P_t are $u_t = \gamma\eta_t + \tilde{u}_t$ with covariance matrix $\Omega = \gamma\gamma'\sigma_\eta^2 + \widetilde{\Omega}$. The long-run trend components of P_t are

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

The short-term interest rate is taken to be an affine function of P_t but not of τ_t:

$$i_t = \delta_0 + \delta_1' P_t. \tag{4}$$

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,

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

1979). The risk-neutral dynamics are specified as

$$P_t = \mu^{\mathbb{Q}} + \Phi^{\mathbb{Q}} P_{t-1} + u_t^{\mathbb{Q}}, \tag{5}$$

where under the $\mathbb{Q}$-measure $u_t^{\mathbb{Q}} \overset{iid}{\sim} N(0, \Omega)$. Equations (4) and (5) 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 B.2. We assume that P_t is stationary under $\mathbb{Q}$ (i.e., that $\Phi^{\mathbb{Q}}$ has eigenvalues less than one in absolute value) because otherwise the model would violate no-arbitrage.[29]

The key innovation of this model relative to existing affine DTSMs[30] is that the equilibrium nominal interest rate is allowed to be time-varying:

$$i_t^* = \delta_0 + \delta_1' P_t^* = \delta_0 + \delta_1' \bar{P} + \delta_1' \gamma \tau_t = \tau_t, \tag{6}$$

where the last equality uses two normalizing assumptions that we make for estimation of the model: $\delta_0 + \delta_1' \bar{P} = 0$ and $\delta_1' \gamma = 1$. In contrast to the standard DTSM, bond yields have a common stochastic trend and are $I(1)$. Imposing a unit root in this fashion is a very 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 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). Similarly, many unit root models for inflation dynamics may not be literally true but nevertheless are very successful in forecasting inflation (Stock and Watson, 2007; Faust and Wright, 2013). To illustrate the importance of the time-varying trend in interest rates, we will compare our stochastic trend model to a conventional DTSM with stationary dynamics. 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). We refer to this specification as the "fixed endpoint" *FE* model.

Another interesting aspect of our model is that the long-run trend is *unspanned* by the

[29] A unit root under $\mathbb{Q}$ would lead to yield convexity terms diverging to minus infinity; see Campbell et al. (1997, p. 433). With a unit root under the real-world measure but not under the risk-neutral measure, our model implies that constant terms for risk-neutral rates and term premia are divergent. This is neither a problem in theory nor in practice: No tradeable securities have payoffs tied to these quantities, and at the maturities we focus on, estimates of the term premium are not noticeably affected by this issue.

[30] Prominent examples of stationary DTSMs include Joslin et al. (2011), Adrian et al. (2013) and Joslin et al. (2014).

yield curve. This specification allows i_t^* to play the role of an unobserved, long-run trend. Section 3.2 documented strong evidence against the spanning hypothesis, which is implied by essentially all asset pricing models (Duffee, 2013). In such models, investors (agents) can generally back out all relevant information (state variables) from the yield curve. This is not possible in our model because it affects expectations of future yields but it does not directly affect current yields. In particular, conditional on P_t, the trend i_t^* does not directly affect the short rate in (4) or the risk-neutral dynamics in (5), so yields do not load on i_t^*. Models with unspanned risk factors are often convenient for matching certain features in the data (Joslin et al., 2014), and we will show that this model matches the stylized facts documented in Section 3. If, by contrast, the common trend was spanned by yields, movements in long-term yields would generally translate directly into movements in i_t^*, generating highly volatile estimates of the trend.[31] While alternative spanned trend representations may be possible, our specification is parsimonious and effective.

Finally, we consider how our model fits into the literature. 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).[32] Cieslak and Povala (2015) proposed a stylized model with a trend component that illustrated important issues related to the behavior of yields and risk premia but was not flexible enough to match both the cross section and time series dynamics of the yield curve. Their model's trend component is a highly persistent but stationary variable, which is precluded from affecting risk premia. Crump et al. (2018) propose a statistical yield-curve model with shifting endpoints that does not include arbitrage-free pricing, which they use to fit survey expectations, similar to Kozicki and Tinsley (2012). In contrast to existing models, our model allows for time-varying risk premia, can accurately fit the cross section of yields, is consistent with absence of arbitrage, and includes a stochastic trend in line with much evidence showing that nominal interest rates are effectively $I(1)$.[33] Importantly, changes in i_t^* may affect not only expectations of future short rates but also the term premium, and our estimates below will show how important a trend in the term premium is empirically. While our DTSM framework is quite flexible, two modeling decisions keep it deliberately simple: First, we focus on the nominal yield curve and do not explicitly model inflation and the real yield curve, which would enable us to consider

[31]In addition, while sufficiently flexible *spanned* DTSM, in which yields depend on all risk factors, can typically match the data equally well as the corresponding *unspanned* DTSM, the spanned models require many more parameters and are more difficult to estimate (Bauer and Rudebusch, 2017).

[32]Extensions of the Kozicki-Tinsley model by Dewachter and Lyrio (2006), Dewachter and Iania (2011) and others allowed for time-varying risk premia, with some limited success.

[33]See Nelson and Plosser (1982), Rose (1988) and King et al. (1991), among many others.

the separate roles of the two underlying trends r_t^* and π_t^*. But this extension is unnecessary for our purpose of estimating nominal risk premia and forecasting Treasury yields, and it would make our DTSM substantially more complex. Second, our model is Gaussian, although both interest rates and their underlying macro trends likely have time-varying volatility (Stock and Watson, 2007; Cieslak and Povala, 2016). How to introduce stochastic volatility in term structure models without sacrificing its ability to match expected bond returns and risk premia remains an important open question in the term structure literature (Andersen and Benzoni, 2010).

4.2 Model estimation

The shifting trend model can be cast in state-space form for the estimation. The measurement equations for yields are $Y_t^o = A + BP_t + e_t$, with $e_t \overset{iid}{\sim} N(0, \sigma_e^2 I_J)$. For our data set Y_t^o includes $J = 17$ yields. Estimation also requires further identifying assumptions; otherwise, it is possible to transform the state variables without changing implied interest rates (Dai and Singleton, 2000). We use the normalization of Joslin et al. (2011), so that four parameters—a scalar $k^{\mathbb{Q}}$ and an N-vector $\lambda^{\mathbb{Q}}$—together with Ω fully determine the yield loadings A and B.[34] 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 the vectors γ and $\bar{P}$ have only two degrees of freedom due to our normalizing assumptions.

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 the purpose of modeling 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.[35] 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

[34] The vector $\lambda^{\mathbb{Q}}$ contains the eigenvalues of $\Phi^{\mathbb{Q}}$, which we assume to be real, distinct, and less than one, and $k^{\mathbb{Q}}$ determines the long-run $\mathbb{Q}$-mean of the short rate. For details see Joslin et al. (2011). 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. Note that this normalization ensures the consistency conditions $WA = 0$ and $WB = I_N$ so that we indeed have $WY_t = P_t$.

[35] 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.

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 our empirical proxy $i_t^* = \pi_t^* + r_t^{*,RT}$, which we showed in Section 3 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.[36] 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 zero-lower-bound-period. Once the parameter estimates are obtained, we proceed to use the model over the entire sample period.[37] 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 a largest eigenvalue of $\Phi^{\mathbb{Q}}$ of 0.997. For the real-world dynamics, the mean reversion to the common trend is relatively quick, with a largest eigenvalue of Φ equal to 0.933. More details on the estimation are given in Appendix B.3.

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 estimate 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 a similar 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%. 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

[36] Details are in Appendix B.3. 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.

[37] 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.

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 are described in Appendix B.4 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.[38]

Figure 4 plots the resulting *ESE* model estimate of i_t^*—specifically, the posterior mean and 95%-credibility intervals. The *ESE* model i_t^* rose until the early 1980s to over eight percent and then gradually declined over the following decades to just below three percent at the end of our sample. For comparison, the 10-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 of the variation in the 10-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 its underlying trend.

4.3 Implications for trend components of yields

Section 3.1 showed that proxies for i_t^* were important drivers of yields, in particular, 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 analogue of the regression-based "interest rate cycles" proposed by Cieslak and Povala (2015) and discussed in Section 3.1—exhibits clear mean reversion around zero and little to no apparent trend. This confirms that our model captures 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

[38]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.

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%-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 2. 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%-Monte Carlo intervals of the coefficient estimates, shown as dashed lines in Figure 6, are quite wide and comfortably contain the coefficients in the actual data.[39]

Our results here and in Section 3.1 show that the trend in Treasury yields does not move 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 that the term premium contains a 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.

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 principle components of yields. Here, we tackle the important task of explaining this result with our shifting endpoint DTSMs. To do so, 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.[40] 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.

[39] 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.

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

Table 6 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 6). The reported ΔR^2 show 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% Monte Carlo interval for ΔR^2 is not centered around zero. However, it remains the case that the upper bound of this interval (4 percentage points) is very much lower than the actual improvement in R^2 from adding trends in the data (12 percentage points), consistent with our 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 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.[41] 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

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

[41] 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 error would make them effectively indistinguishable (Jorgensen, 2017).

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 B.2. 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 B.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 contain a trend component.[42] 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. 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 term premium estimation 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 common modeling approach for DTSMs. 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.[43] 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

[42] 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.

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

1980s. As argued by Kim and Orphanides (2012) and Bauer et al. (2012), 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). The term structure literature has proposed a number of different remedies to avoid such counterfactual decompositions of long-term rates into expectations and term premium. These modifications generally have the goal of making long-run expectations more variable and include restrictions on risk prices (Cochrane and Piazzesi, 2008; Joslin et al., 2014; Bauer, 2018), bias correction of interest rate dynamics (Bauer et al., 2014), and incorporation of survey-based expectations of future interest rates (Kim and Wright, 2005; Kim and Orphanides, 2012).

The remedy we propose here is to allow for a long-run trend in interest rates. 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%-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 expectations are anchored at i_t^*, 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% in the early 1980s to about 0-1% at the end of the sample. The reason is that long-term interest rates exhibited more pronounced secular movements than the estimates of the long-run trend i_t^*, hence 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. Some downward trend in the term premium since the 1980s is plausible from a macro-finance perspective, given that inflation risk declined substantially over this period, and the term premium includes compensation for this risk (Wright, 2011).[44] But the trend in the term premium is much more modest than with a fixed endpoint. Furthermore, the estimated term premium with a shifting endpoint exhibits more pronounced cyclical variation, in line

[44]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).

with the notion that risk premia are countercyclical (Fama and French, 1993; Cochrane and Piazzesi, 2005). Our shifting-endpoint models attribute the majority of the secular decline in interest rates not to the term premium, like the *FE* model and many existing models, 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% of this decline. In contrast, the shifting-endpoint models attribute only 35% (*OSE*) and 37% (*ESE*) to the term premium and the majority to a falling expectations component, in line with the substantial downward shift in the underlying macro trends. 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 walk for the level of interest rates leads to drastically different implications for expectations and term premia, and that results in both cases are at odds with widely held views about interest rates. While imposing a random walk for the level is known to forecast well, it also leads to the implausible implication that the expectations accounts for essentially all variation in yields, as emphasized by Cochrane. By contrast, as shown by Cochrane and evident in the left panel of Figure 7, imposing stationary dynamics leads to the implausible implication that the term premium component accounts for all of the low-frequency variation in yields. Our results suggest a solution to this knife-edge problem, namely to assume that interest rates are I(1) with a slow-moving common trend. Such a formulation avoids both extremes, and delivers term premium estimates that are in line with the common macro-finance priors. In addition, it produces interest rate forecasts that outperform the random walk, as we will show below.

In recent work, Crump et al. (2018) use a shifting-endpoint VAR to fit survey-based interest rate expectations, and calculate the term premium as the difference between yields and survey expectations. Their model does not impose absence of arbitrage, hence estimates of the term premium ignore the role of bond convexity, because risk-neutral yields are approximated using expectations of future short rates, and should therefore be taken with a grain of salt. But more important for their results is the reliance on survey expectations, which leads them to conclude that "term premiums account for the bulk of the [...] variation in yields." The likely reason is that the estimate of i_t^* implied by survey expectations captures less of the low-frequency movements in interest rates, so the residual term premium is instead required to account for most of the secular decline. Their results about the absence of any important role of long-run trends for interest rates are at odds with our stylized facts in Section 3 and our model-based decompositions of long-term interest rates.

When interpreting model-based estimates of the term premium, it is important to keep in mind that they are based on an assumption of rational expectations and do not necessarily reflect 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.[45] While this caveat also applies to our estimates, models that allow for shifting endpoints like ours 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 these reasons, our shifting-endpoint model can likely better 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 (for a survey see 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 2.3. We will compare these forecasts to those from the *FE* model and from a driftless random walk.[46] 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.[47] We focus on forecasts of the ten-year yield; we have found results for other

[45] 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.

[46] 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 costly.

[47] 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

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.[48] The top panel of Table 7 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).[49] 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 our common-trend model OSE achieves substantial and statistically significant forecasts gains relative to the FE model. For example, when forecasting ten years ahead, model OSE lowers the RMSE by about 40% relative to the FE model. The improvements relative to the RW are smaller but still statistically significant.[50]

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 "gold standard" 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. 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.

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

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.

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

[49] 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 as in those cases there are only few non-overlapping observations in our sample.

[50] 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 sub-periods, but instead are a consistent pattern over most of our sample period.

includes 48 forecast dates from March 1988 to December 2011. The bottom panel of Table 7 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 gains 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 over-predict 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.

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 introduced and estimated a 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.

Andersen, T. G. and L. Benzoni (2010): "Do bonds span volatility risk in the US Treasury market? A specification test for affine term structure models," *The Journal of Finance*, 65, 603–653.

Ang, A., G. Bekaert, and M. Wei (2007): "Do macro variables, asset markets, or surveys forecast ination 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.

Bansal, R. and I. Shaliastovich (2013): "A Long-Run Risks Explanation of Predictability Puzzles in Bond and Currency Markets," *Review of Financial Studies*, 26, 1–33.

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," *The 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 deflation 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.

——— (2016): "Information in the term structure of yield curve volatility," *The Journal of Finance*, 71, 1393–1436.

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.

Clark, T. E. and S. Kozicki (2005): "Estimating equilibrium real interest rates in real time," *The North American Journal of Economics and Finance*, 16, 395–413.

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.

——— (2008): "Decomposing the Yield Curve," unpublished manuscript, Chicago Booth School of Business.

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.

Cúrdia, V., A. Ferrero, G. C. Ng, and A. Tambalotti (2015): "Has US monetary policy tracked the efficient interest rate?" *Journal of Monetary Economics*, 70, 72–83.

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. Giannoni, and A. Tambalotti (2017): "Safety, liquidity, and the natural rate of interest," Brookings Papers on Economic Activity, Conference Draft.

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. (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, 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 (US).

——— (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*, 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,” FEDS Working Paper 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 Investifation 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, Federal Reserve Board of Governors.

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.

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?" *The Quarterly Journal of Economics*, 126, 373–416.

Mankiw, N. G. and M. D. Shapiro (1986): "News or noise? An analysis of GNP revisions," .

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.

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

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.1 and shown in Figure A.1. Our own four estimates of r_t^* are described in Appendix A.2 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 PTR estimate of π_t^* is described in Section 2.2. 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 B.4.

Table A.1: Overview of trend estimates

Trend	Source	Real-time	Filtered	Smoothed
r^*	Del Negro et al. (2017) (DNGGT)	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^*, and four of our own estimates of r^*, the PTR estimate of π^*, and the two i^* estimates used in the paper. "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.

A.1 External estimates of r_t^*

Del Negro et al. (2017) (DNGGT) 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) 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 2.3), the other three external estimates use a different definition and estimation approach. The prominent model of Laubach and Williams (2003, 2016), the slightly modified version of this model by Holston et al. (2017), 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. 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, 5-10 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. For these reasons, we believe it is worthwhile to include them in our analysis, while being cognizent of the conceptual differences. A shortcoming of these estimates is that they are not available in real time. While 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—they 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.

A.2 Details on our estimates of r_t^*

We include three of our own model-based estimates and one simple moving-average estimate of r_t^* in our analysis. The first estimate is a univariate unobserved components (UC) model, similar to Watson (1986) for the real short-term interest rate, which is taken 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

Figure A.1: External estimates of r_t^*

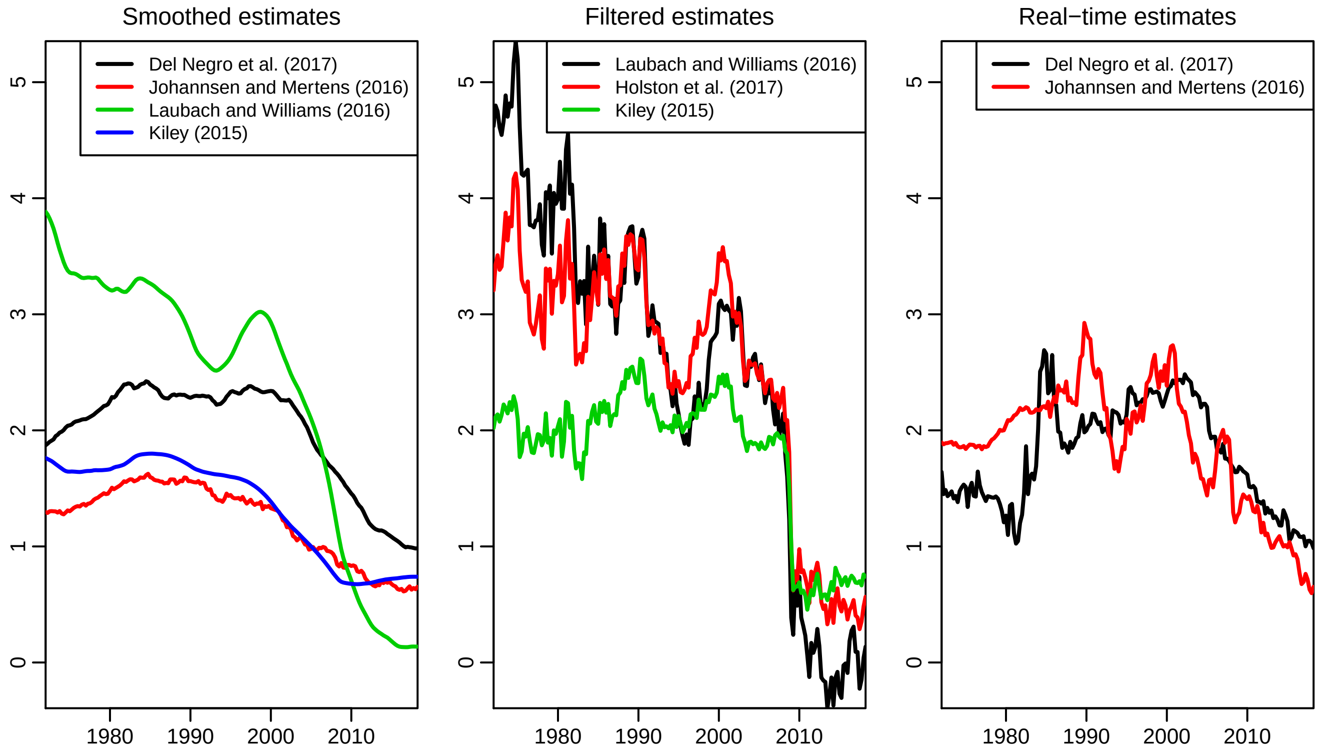

Estimates of r_t^* from existing models. Left panel: smoothed/two-sided estimates. Middle panel: filtered/one-sided estimates. Right panel: (pseudo) real-time estimates. The sample period is from 1971:Q4 to 2018:Q1.

energy items. This rate of inflation over the past year serves as a proxy estimate for expectations of inflation over the next quarter, hence providing a measure of the ex-ante real interest rate. 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^*. Here we use a tight prior around a low value for this variance, similar to DNGGT. 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 DNGGT (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 then 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 directly motivated by the results in Lunsford and West (2017). 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 alleviates any 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 DNGGT 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 Minnesota prior as DNGGT. 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 DNGGT (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 those used in other work estimating macro trends (see footnote 15 in the main text).

Figure A.2: Our estimates of r_t^*

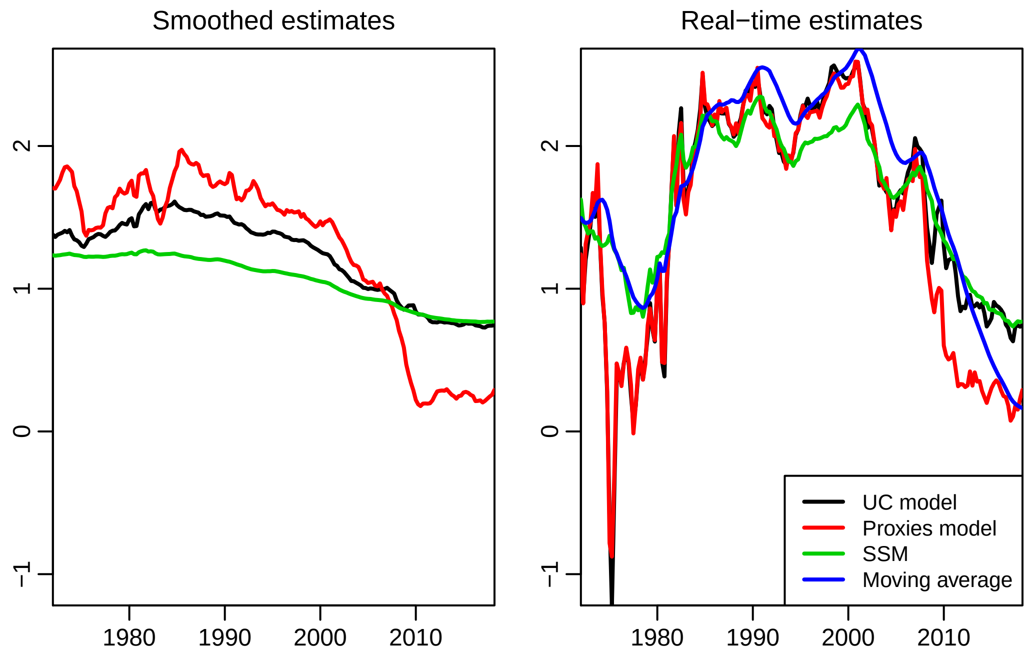

Our own four estimates of r_t^*. Left panel: smoothed/two-sided estimates. Right panel: (pseudo) real-time estimates. The sample period is from 1971:Q4 to 2018:Q1.

B Details on dynamic term structure model

B.1 Dynamic system

Equation (3) gives 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).

Our assumptions imply a cointegrated VAR for Z_t:

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

where Φ_Z has exactly one eigenvalue equal to unity. To understand this result, first note that (3) implies that

$$P_t = (I_N - \Phi)\bar{P} + (I_N - \Phi)\gamma\tau_{t-1} + \Phi P_{t-1} + u_t,$$

with $u_t = \gamma\eta_t + \tilde{u}_t$. Therefore, the parameters for the VAR representation (B.1) for Z_t are

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

and the covariance matrix of the VAR innovations is

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

where $\Omega = E(u_t u_t') = \gamma\gamma'\sigma_\eta^2 + \widetilde{\Omega}$.

The vector-error-correction representation is

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

where $\Phi_Z - I_{N+1} = \alpha\beta'$. Note that, by construction, the intercept does not generate a linear trend in the mean of Z_t because $\mu_Z = -\alpha\beta'(0, \bar{P}')'$. The cointegration residual is $\beta' Z_{t-1} = \tilde{P}_{t-1} + \bar{P}$ and it Granger-causes the yield factors:

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

That is, deviations of yields from their equilibrium $P_t = \bar{P} + \gamma\tau_t$ are eliminated by future changes in yields.

B.2 Affine loadings and excess returns

Prices of zero-coupon bonds with maturity n, denoted by $P_t^{(n)}$, are easily verified to be exponentially affine, i.e., $\log(P_t^{(n)}) = \mathcal{A}_n + \mathcal{B}_n' P_t$, using the pricing equation $P_t^{(n+1)} = e^{-i_t} E_t^{\mathbb{Q}}(P_{t+1}^{(n)})$. (In a slight abuse of notation, $P_t^{(n)}$ with superscript is a bond price while P_t without superscipt is a vector with the N yield factors.) The coefficients follow the usual recursions (e.g., Ang and Piazzesi, 2003):

$$\mathcal{A}_{n+1} = \mathcal{A}_n + \mathcal{B}_n' \mu^{\mathbb{Q}} + \frac{1}{2}\mathcal{B}_n' \Omega \mathcal{B}_n - \delta_0$$
$$\mathcal{B}_{n+1} = (\Phi^{\mathbb{Q}})' \mathcal{B}_n - \delta_1$$

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

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

Model-implied excess bond returns are

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

so that expected excess returns are

$$
\begin{aligned}
E_t rx_{t+1}^{(n)} &= -\frac{1}{2}\mathcal{B}_{n-1}'\Omega\mathcal{B}_{n-1} + \mathcal{B}_{n-1}'\left(E_t P_{t+1} - E_t^{\mathbb{Q}} P_{t+1}\right) \\
&= -\frac{1}{2}\mathcal{B}_{n-1}'\Omega\mathcal{B}_{n-1} + \mathcal{B}_{n-1}'\left(\bar{P} + \gamma\tau_t + \Phi\tilde{P}_t - \mu^{\mathbb{Q}} - \Phi^{\mathbb{Q}} P_t\right) \\
&= c_n + \mathcal{B}_{n-1}'\left(\gamma\tau_t + \Phi\tilde{P}_t - \Phi^{\mathbb{Q}}(\gamma\tau_t + \tilde{P}_t)\right) \\
&= c_n + \mathcal{B}_{n-1}'\left((I_N - \Phi^{\mathbb{Q}})\gamma\tau_t + (\Phi - \Phi^{\mathbb{Q}})\tilde{P}_t\right),
\end{aligned}
$$

where c_n is a maturity-dependent constant.

B.3 Estimation with observed shifting endpoint

The loadings of yields Y_t on the factors P_t, A, and B, are obtained from the fundamental parameters $k^{\mathbb{Q}}$, $\lambda^{\mathbb{Q}}$ and Ω, as described in Joslin et al. (2011). This formulation also provides the "rotated" parameters $\mu^{\mathbb{Q}}$, $\Phi^{\mathbb{Q}}$, δ_0 and δ_1. To ensure that the common trend is $\tau_t = i_t^*$, the constraints $\delta_0 + \delta_1'\bar{P} = 0$ and $\delta_1'\gamma = 1$ must be satisfied (see equation 6). Therefore, the vectors γ and $\bar{P}$ are parameterized in terms of vectors g and p with length $N-1$, using the mappings

$$\gamma = (g', (1 - g'\delta_{1,-N})/\delta_{1,N})', \bar{P} = (p', (-\delta_0 - p'\delta_{1,-N})/\delta_{1,N})',$$

where $\delta_{1,-N}$ are the first $N-1$ elements of δ_1 and $\delta_{1,N}$ is the last element.

In this implementation of our shifting-endpoint DTSM, we use an empirical proxy for i_t^*. Because of the normalizing restrictions just described, we have $\tau_t = i_t^*$ and 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.)

B.4 Estimation with Bayesian priors and Markov chain Monte Carlo

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 DNGGT, which also uses their model's prior to limit "the amount of variation that it attributes to the trends" (DNGGT, p. 249).

- For Φ and $\widetilde{\Omega}$, we specify a Minnesota-type normal/inverse-Wishart conjugate prior similar to DNNGT. 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 DNGGT 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 from; therefore, we must use M-H steps. Note that 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 independence proposal distributions, more specifically, tailored independence proposals similar to Chib and Ergashev (2009): Conditional on all other parameters and the state variables, we use numerical optimization to find the mode and Hessian of the conditional posterior distribution. Then our proposal distribution is 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. It's is not necessary to do this numerical optimization every iteration. We find that doing it only every tenth iteration leads to sufficiently high M-H acceptance probabilities, which generally should be in a range between 20% and 50% (Chib and Greenberg, 1995). This way of constructing M-H proposal distributions is still computationally somewhat burdensome, but it has the great benefit that no fine-tuning of the proposal distributions is necessary, and the efficiency of the sampler is generally 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: Persistence of interest rates, macroeconomic trends, and differences

Series	SD	$\hat{\rho}$	Half-life	ADF	PP	LFST
$y_t^{(40)}$	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
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, calculated as $\ln(0.5)/\ln(\hat{\rho})$; 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 ten-year yield, the detrended yield, and macro trends. The r^*-estimates are the filtered ("F"), real-time ("RT") and moving-average ("MA") estimates shown in Figure 2 and described in Section 2.3. The data are quarterly from 1971:Q4 to 2018:Q1.

Table 2: Cointegration regressions and tests

	(1)	(2)	(3)	(4)	(5)
constant	0.65	-0.22	-1.82	-0.62	-2.17
	(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	1.31	1.09	0.70	0.87	0.70
$\hat{\rho}$	0.88	0.85	0.65	0.76	0.64
Half-life	5.6	4.3	1.6	2.5	1.5
ADF	-2.60	-3.94**	-5.32***	-4.32**	-5.37***
PP	-18.32*	-26.73**	-68.45***	-46.05***	-70.30***
LFST	0.03	0.16	0.71	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 on macroeconomic trends, including four leads and lags of first-differenced trend variables. Newey-West standard errors using six lags in parentheses. The r^* estimates are shown in Figure 2 and described in Section 2.3. The long-run nominal short rate i_t^* is the sum of π_t^* and the real-time estimate of r_t^*. Persistence statistics and unit root tests for cointegration residuals are described in the notes for Table 1 and in the text. Johansen trace statistics for testing whether the cointegration rank (r) among $y_t^{(40)}$ and the macro trends is zero or one against the alternative that it exceeds zero or one, using four lags in the VAR. 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 3: 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 r^* estimates are shown in Figure 2 and described in Section 2.3. The long-run nominal short rate i_t^* is the sum of π_t^* and the real-time estimate of r_t^*. 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 4: 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 r^* estimates are shown in Figure 2 and described in Section 2.3. The long-run nominal short rate i_t^* is the sum of π_t^* and the real-time estimate of r_t^*. Numbers in parentheses are White standard errors with 6 lags. The data are quarterly from 1971:Q4 to 2018:Q1.

Table 5: 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).

Table 6: 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]	[0.04, 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	0.06	0.16	0.10
	[0.02, 0.12]	[0.08, 0.24]	[0.03, 0.18]

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 3 (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 7: 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

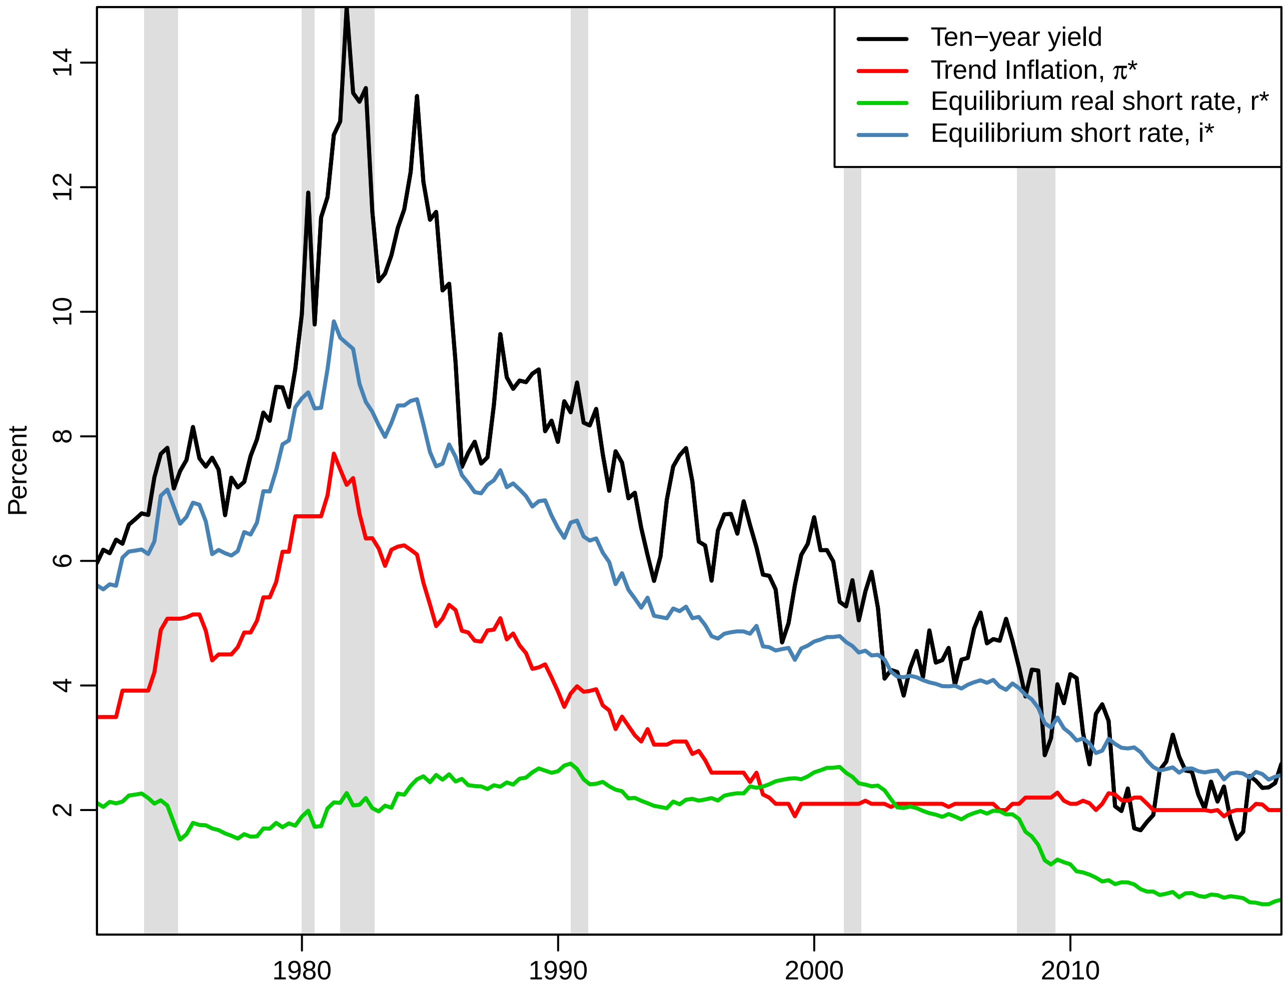

Ten-year Treasury yield and estimates of trend inflation, π_t^* (the mostly survey-based PTR measure from FRB/US), the equilibrium real rate, r_t^* (an average of all filtered and real-time estimates, see Section 2.3), and the equilibrium short rate, $i_t^* = \pi_t^* + r_t^*$. Shaded areas are NBER recessions. The data are quarterly from 1971:Q4 to 2018:Q1.

Figure 2: Measures of the equilibrium real interest rate

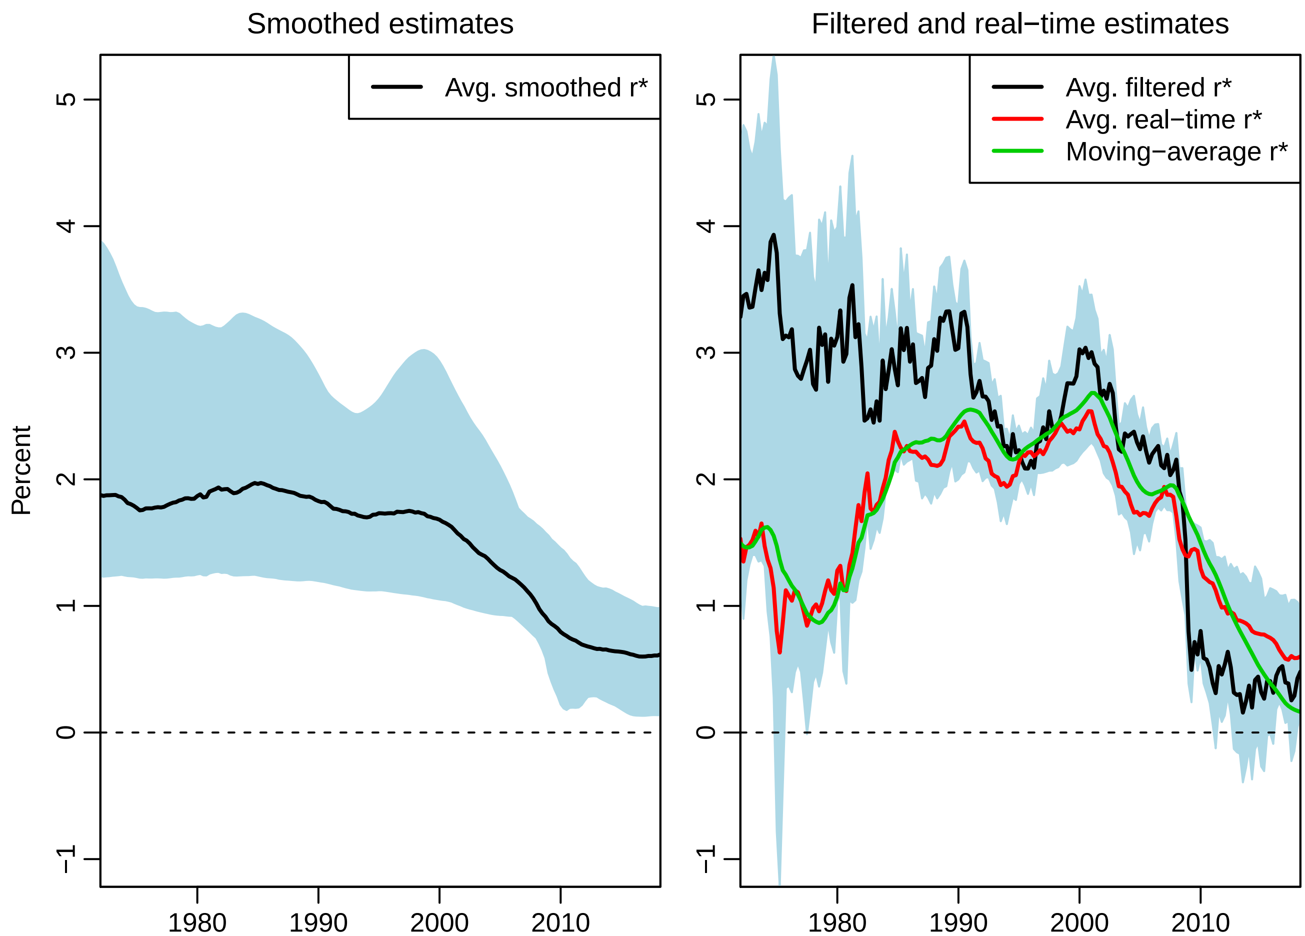

Smoothed (two-sided), filtered (one-sided) and real-time estimates of the equilibrium real interest rate, or long-run r_t^*. Left panel: range and average of seven smoothed estimates, including the four estimates by Del Negro et al. (2010) (DNGGT), Johannsen and Mertens (2018) (JM), Laubach and Williams (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 DNGGT, 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 2.3 and Appendix A. The data are quarterly from 1971:Q4 to 2018:Q1.

Figure 3: Detrending the ten-year yield

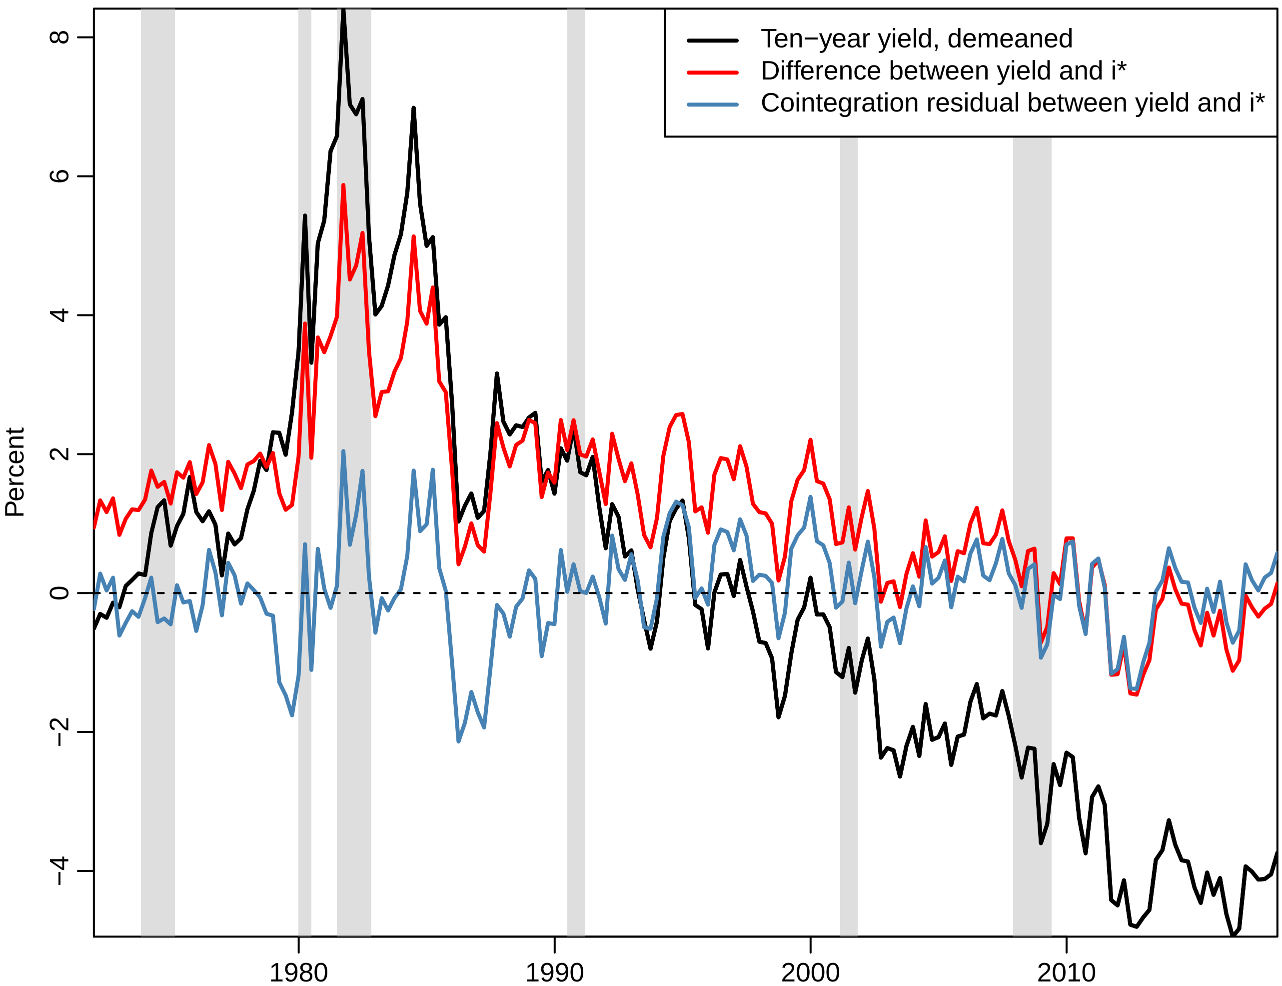

Ten-year Treasury yield and two detrended yield series using as the proxy for the equilibrium interest rate i_t^* the sum of the PTR estimate of π_t^* and the real-time estimate of r_t^* shown in Figure 2. The first detrended yield series is the difference between the yield and i_t^*. The second detrended 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 2. Shaded areas are NBER recessions. The data are quarterly from 1971:Q4 to 2018:Q1.

Figure 4: Model-based estimate of equilibrium interest rate

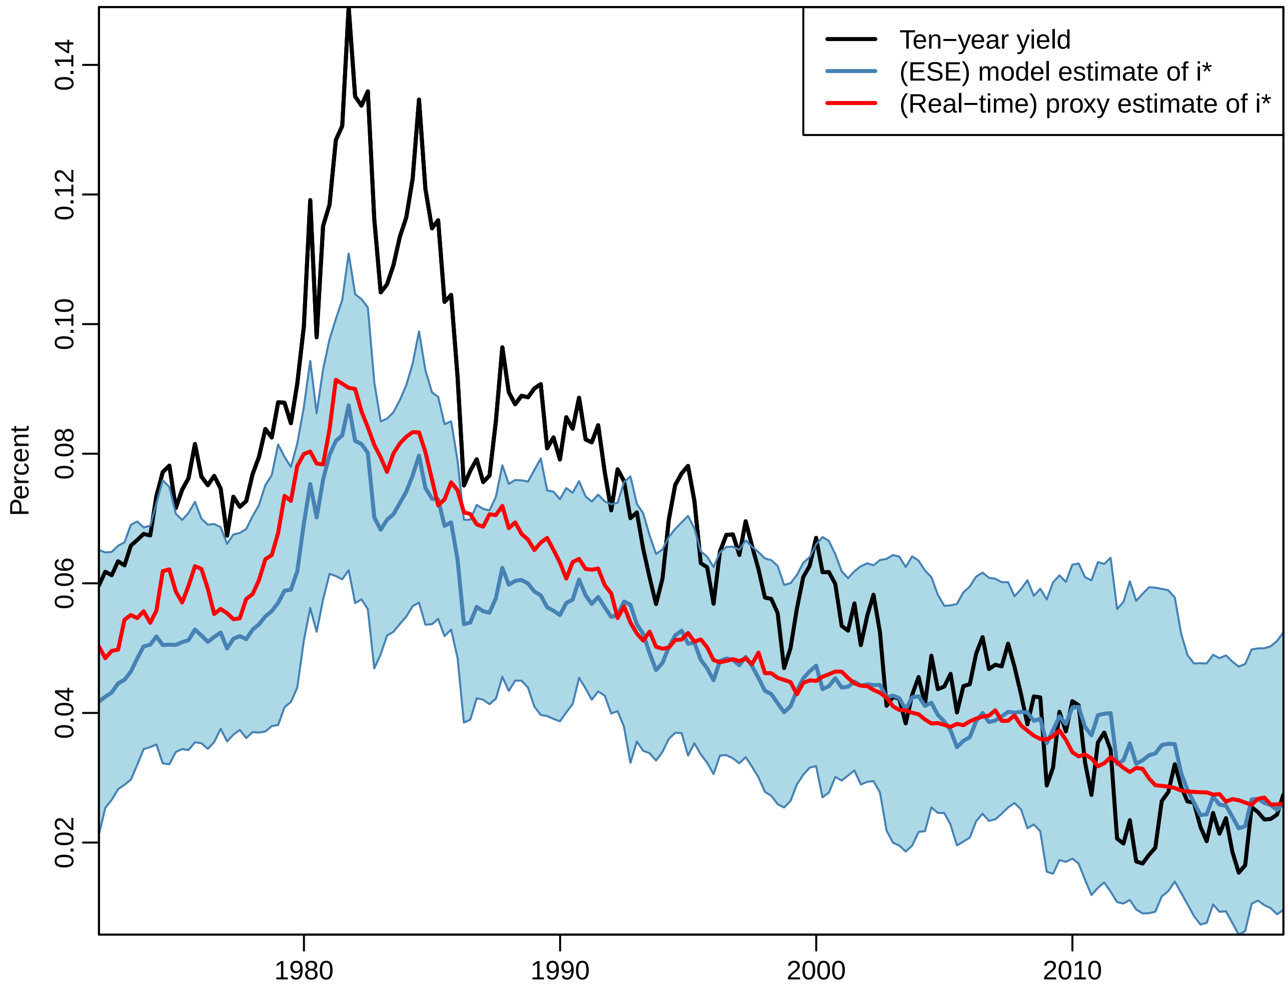

Bayesian (MCMC) estimate—posterior mean and 95%-posterior credibility intervals—of the equilibrium nominal short rate, i^*, from the "estimated shifting-endpoint" (ESE) dynamic term structure model. Also shown is a proxy estimate of i^*, the sum of the PTR estimate of π^* and the real-time macro estimate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

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

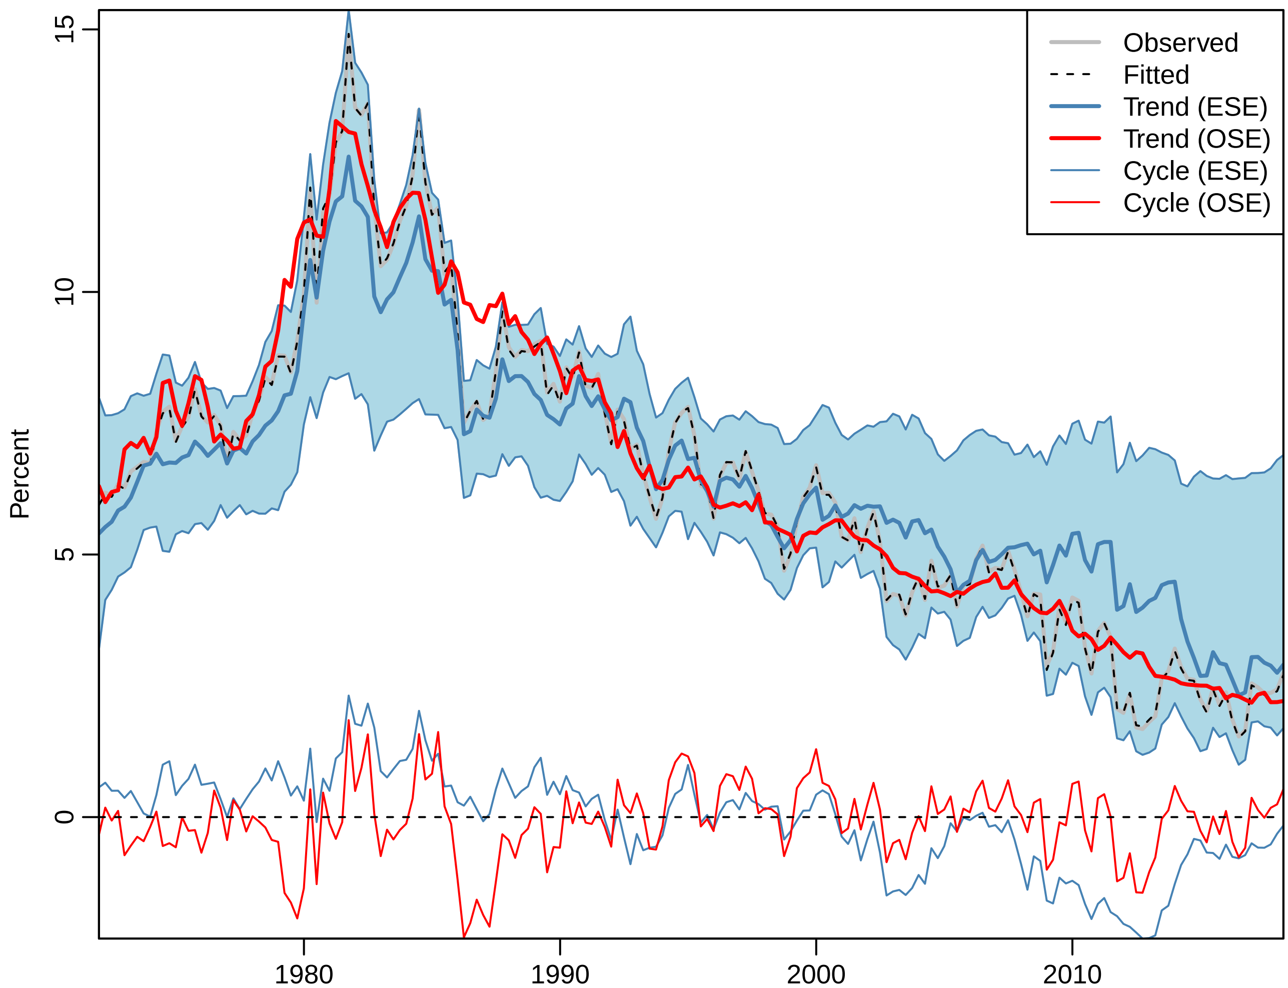

Model-based decomposition of ten-year yield into trend and cycle components. Blue lines show trend and cycle (posterior means) for the model with estimated shifting endpoint (*ESE*), shaded area shows 95%-posterior credibility intervals for trend component; red lines show trend and cycle for model with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

Figure 6: Loadings of yields on the equilibrium interest rate

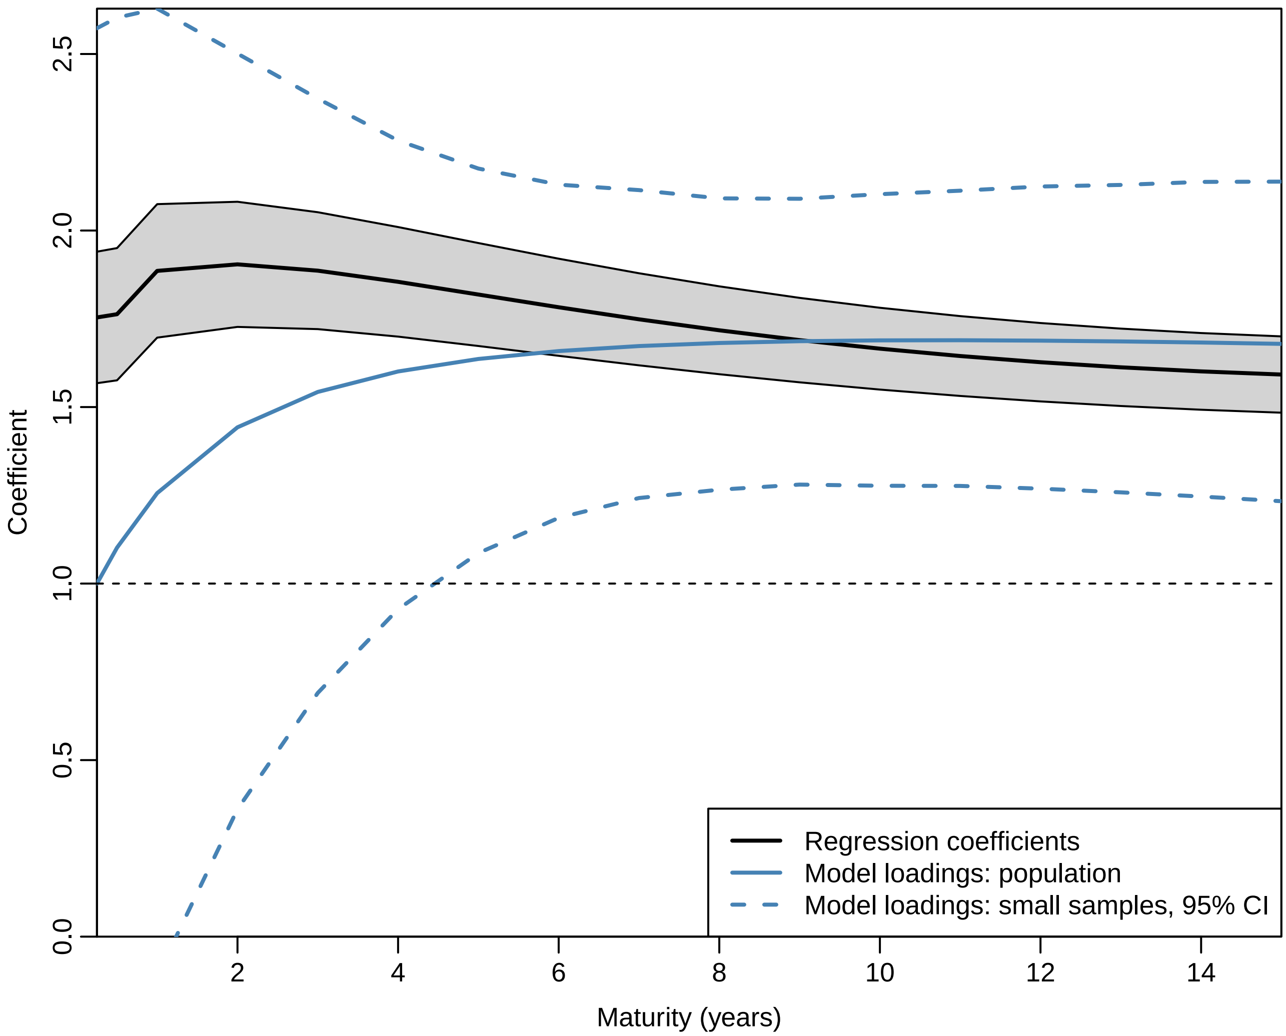

Comparison of loadings of Treasury yields on the equilibrium nominal short rate i_t^* in the data and in the shifting-endpoint dynamic term structure model. Data: regression coefficients (with 95%-confidence intervals) for cointegration regressions of yields on the proxy-estimate of i_t^* (the PTR estimate 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

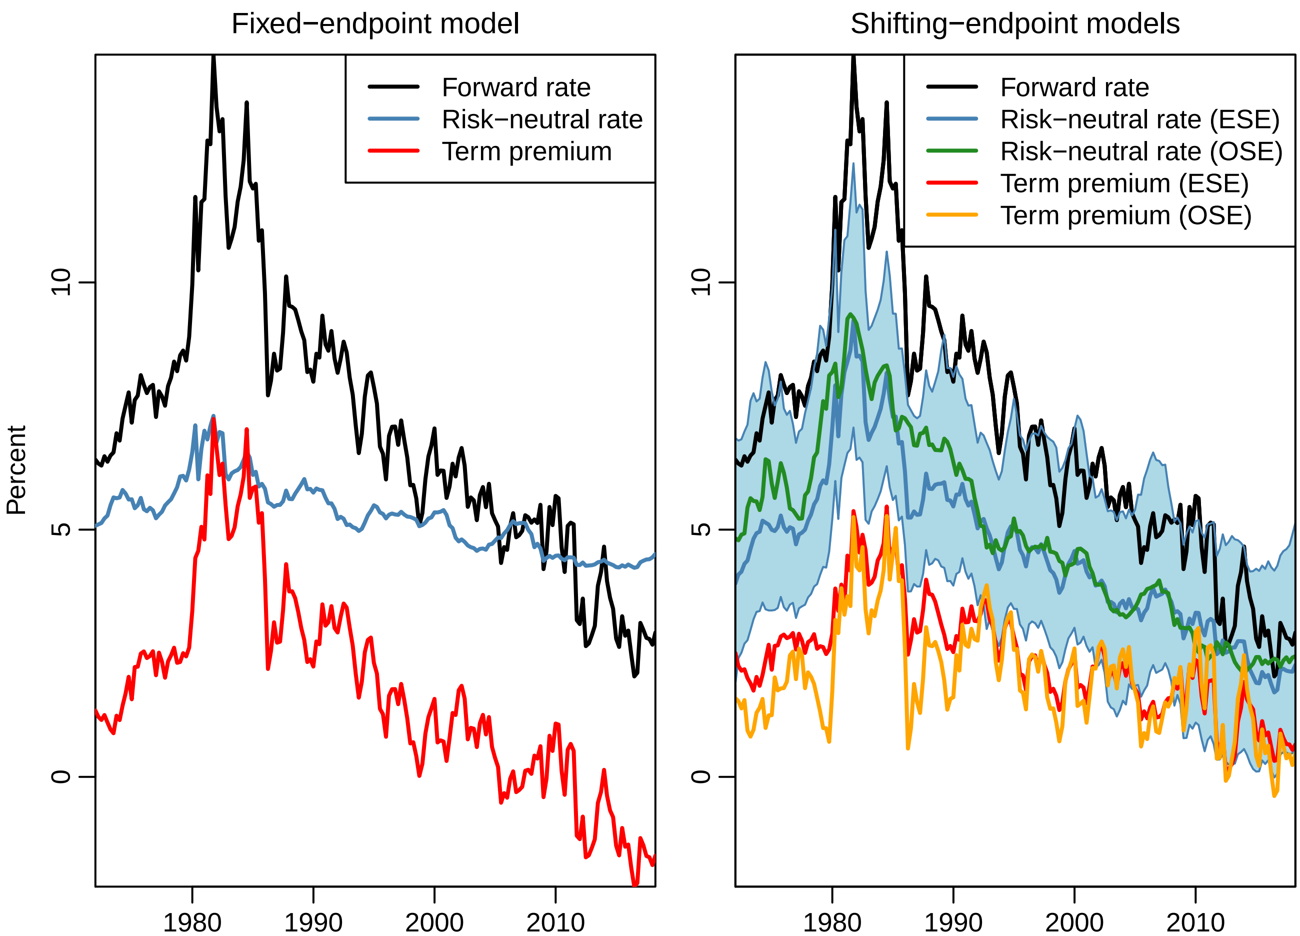

Five-to-ten-year forward rate with estimated expectations component (risk-neutral rate) and term premium. Left panel: conventional dynamic term structure model (DTSM) with a fixed endpoint i^* (Joslin et al., 2011). Right panel: DTSM with a common stochastic trend, where the shifting endpoint is either 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.