

The Shadow Rate of Interest, Macroeconomic Trends, and Time-Varying Uncertainty

Benjamin K. Johannsen*
Federal Reserve Board

Elmar Mertens
Federal Reserve Board

First Draft: February 14, 2015

This Draft: March 27, 2015

Abstract

This paper describes the historical behavior of the shadow rate of interest—the nominal interest rate that would prevail in the absence of its effective lower bound—along with other macroeconomic aggregates in post-war data from the United States. Our flexible time series approach does not impose the no-arbitrage restrictions commonly used in the dynamic term-structure literature, allowing us to include time variation in parameters, such as stochastic volatility. We explicitly incorporate the effective lower bound on nominal interest rates, allowing us to address data before, during, and after the Great Recession. We document that the trend component of the nominal interest rate has fallen almost continuously since the early 1980s. The decline is due to long-standing downward trajectories of the trend component of both inflation and the real short-term interest rate. We also document that the trend level of output growth has declined over our sample. However, we find only modest evidence of correlation between the trend in the real short-term interest rate and the trend in output growth. Finally, we show that term premium estimates constructed by comparing the expected path of future short-rates from our models with data on longer-term interest rates are strikingly similar to the estimates reported in the dynamic term-structure literature. Our results thus offer empirical context to current debates about secular stagnation, movements premiums, and longer-run normal levels of interest rates.

*For correspondence: Benjamin K. Johannsen, Board of Governors of the Federal Reserve System, Washington D.C. 20551. email benjamin.k.johannsen@frb.gov. Tel.: +(202) 530 6221. Future updates of this working paper can be downloaded from www.elmarmertens.com/research/workingpapers/JohannsenMertensShadowrate.pdf. The views in this paper do not necessarily represent the views of the Federal Reserve Board, or any other person in the Federal Reserve System or the Federal Open Market Committee. Any errors or omissions should be regarded as solely those of the authors.

1 Introduction

This paper describes the historical behavior of the shadow rate of interest—the nominal interest rate that would prevail in the absence of its effective lower bound—along with other macroeconomic aggregates in post-war data from the United States. To empirically match the slow-moving changes in the average level the nominal interest rate we use a model that decomposes the data into trend and cycle components. Our flexible time series approach does not impose the no-arbitrage restrictions commonly used in the term-structure literature, allowing us to include time variation in parameters, such as stochastic volatility. We explicitly incorporate the effective lower bound on nominal interest rates, allowing us to address data before, during, and after the Great Recession.

We document that the trend component of the nominal interest rate has fallen almost continuously since the early 1980s. The decline is due to long-standing downward trajectories of the trend component of both inflation and the real short-term interest rate. We also document that the trend level of output growth has declined over our sample. However, similar to results recently reported in Hamilton et al. (2015), we find only modest evidence of correlation between the trend in the real short-term interest rate and the trend in output growth. Our results thus offer empirical context to current debates about secular stagnation, movements in real interest rates, and longer-run normal levels of short-term nominal interest rates.

To study movements in longer-term interest rates, we estimate the expected average path of future short-term interest rates at every date in our sample while explicitly incorporating the effective lower bound on nominal interest rate. We combine these estimates with data on longer-term interest rates to estimate movements in premiums. Our estimates indicate that premiums have declined steadily since the early 1980s, which is also the period over which we estimate that the volatility in trend inflation has declined. Thus, despite using time series models that do not include longer-term interest rates or impose the no-arbitrage restrictions commonly employed in the

dynamic term–structure literature, our results are very similar to those reported in Kim and Wright (2005) and Wright (2011).

Explicitly modeling the effective lower bound on nominal interest rates has profound effects on our model estimates of the shadow rate of interest, which were well below zero during the Great Recession. An alternative approach to ours for modeling short–term nominal interest rate data at the effective lower bound would be to treat those data as missing. In this case, the fact that the short–term nominal interest rate has been at its effective lower bound for more than six years would render no information about changes in trends and volatilities in other macroeconomic variables. When we estimate our models treating the interest rate data as missing when the effective lower bound binds, our estimates of the shadow rate are markedly higher. However, the downward tilt in the trend of the short–term real interest rate is only somewhat attenuated and our estimated path for the trend level of output growth is almost unchanged. We take this as confirmatory evidence that our conclusions are not figments of the way we model the effective lower bound. However, we stress that modeling the effective the lower bound on nominal interest rates incorporates useful information in the estimation and is particularly helpful for understanding movements in longer-term interest rates during the Great Recession.

The way we incorporate the lower bound on nominal interest rates and estimate the model can be extended to a broad class of time series models. With short–term nominal interest rates at or near their effective lower bounds since the onset of the global financial crisis in 2008, time series models that include interest rates but ignore the effective lower bound, like vector autoregressions, have been unable to adequately address the data. Moreover, reduced-form empirical explorations of the relationship between short– and longer–term interest rates—such as Campbell and Shiller (1991)—have often ignored the truncation in the distribution of future short-term interest rates. Our modeling approach overcomes these shortcomings and can accommodate censored data (such as the shadow rate when the short–term interest rate is at its effective lower bound) in wide class of otherwise conditionally–linear Gaussian state–space models.

Examples include the vector autoregressions studied in Sims (1980) and the models with time-varying parameters studied in Primiceri (2005) and Cogley and Sargent (2005).

Following work by Black (1995), several studies have used dynamic factor term-structure models to infer the level of the shadow rate of interest, including Kim and Singleton (2011), Krippner (2013), Wu and Xia (2014), and Bauer and Rudebusch (2014). While imposing the no-arbitrage cross-equation restrictions that define these models has offered interesting insights, the cost is that these authors maintain the strong assumption that the dynamic structure of the model does not vary over time. Our flexible time series approach naturally incorporates time variation in parameters and volatilities.

Iwata and Wu (2006), Nakajima (2011), Chan and Strachan (2014) are the closest papers in the literature to ours. These papers also estimate time series models that incorporate an effective bound on nominal interest rates. In all of these studies lagged observed interest rates (rather than shadow rates) are explanatory variables in the dynamic system. We instead allow lagged shadow rates to be explanatory variables so that our concept of the shadow rate more closely aligns to the shadow rate studied in dynamic term structure models and so that the shadow rate can have persistence from its own lagged values. By allowing the shadow rate to have persistence from its lagged values, we are able to connect the concept of the shadow rate with the desired level of the short term interest rate from the monetary authority because, we allow it to have the persistence observed in historical short-term interest rates. Nevertheless, our approach is flexible enough to include both shadow rates and observed rates as explanatory variables.

The rest of our paper is organized as follows. In section 2, we develop and estimate a simple bivariate model of inflation and the shadow rate of interest to illustrate our estimation procedure. In section 3, we develop and estimates a larger model, which includes real output and unemployment in addition to inflation and the shadow rate. We analyze historical term premiums implied by our model estimates in section 4. We

offer concluding comments in section 5. The details of our general estimation strategy are available in the appendices.

2 A Model of Inflation and the Shadow Rate

In this section we use a bivariate model of inflation and the short-term nominal interest to decompose the historical time series into trend and gap components. This small model allows us to illustrate our technique for estimating the shadow rate of interest in a time series model that accounts for the effective lower bound and lets us make inference about changes in the trend level of the short-term real rate of interest.

2.1 Probability Model

We assume that the observed short-term nominal rate of interest, i_t^* , is the maximum of the shadow rate of interest, i_t , and zero. That is,

$$i_t^* = \max(i_t, 0). \quad (1)$$

The shadow rate follows a process so that

$$i_t = \bar{\pi}_t + \bar{r}_t + \tilde{i}_t, \quad \bar{r}_t = \bar{r}_{t-1} + \sigma_{\bar{r}} \varepsilon_{\bar{r},t}, \quad \rho(L)\tilde{i}_t = \sigma_{\tilde{i}} \varepsilon_{\tilde{i},t} + \beta \tilde{\pi}_t. \quad (2)$$

where $\rho(L)$ is a polynomial in the lag operator, the variables \bar{r}_t and $\bar{\pi}_t$ have the interpretation of the trends in the real short-term interest rate and inflation rate, and \tilde{i} is the interest rate gap, which is meant to capture transitory movements in the level of short-term interest rates. A non-zero value of β allows fluctuations in the gap component of inflation, $\tilde{\pi}_t$, to affect the shadow rate through \tilde{i}_t , presumably through a monetary policy response.

We assume that observed inflation, π_t , follows an unobserved components process

with stochastic volatility, as in the UC-SV model of Stock and Watson (2007), so that

$$\pi_t = \bar{\pi}_t + \tilde{\pi}_t, \quad \bar{\pi}_t = \bar{\pi}_{t-1} + \sigma_{\bar{\pi},t} \varepsilon_{\bar{\pi},t}, \quad \tilde{\pi}_t = \sigma_{\tilde{\pi},t} \varepsilon_{\tilde{\pi},t}. \quad (3)$$

Throughout the paper, $\varepsilon_{\cdot,t}$ and $\eta_{\cdot,t}$ represent standard normal random variables that are independent of all other stochastic processes in the system. Importantly, the volatilities of the innovations to the trend and the gap components of inflation are stochastic, and the volatility processes evolve so that

$$\log(\sigma_{\bar{\pi},t}^2) = \log(\sigma_{\bar{\pi},t-1}^2) + \delta_{\bar{\sigma}} \eta_{\bar{\sigma},t}, \quad (4)$$

$$\log(\sigma_{\tilde{\pi},t}^2) = \log(\sigma_{\tilde{\pi},t-1}^2) + \delta_{\tilde{\sigma}} \eta_{\tilde{\sigma},t}. \quad (5)$$

The defining feature of the gap components of the data, $\tilde{\pi}_t$ and \tilde{i}_t , is that they are assumed to have a zero ergodic mean. The values of $\bar{\pi}_t$ and \bar{r}_t are thus similar in spirit to the Beveridge and Nelson (1981) trend, so that

$$\bar{\pi}_t = \lim_{h \rightarrow \infty} E_t(\pi_{t+h}) \quad \text{and} \quad \bar{r}_t = \lim_{h \rightarrow \infty} E_t(i_{t+h} - \pi_{t+h}), \quad (6)$$

where $E_t(\cdot)$ is the expectation operator assuming all values of the volatilities as well as the trend and gap components of the data up to time t have been observed and that the parameters of the model are known.

2.2 Data

For estimation, we use data on π_t and i_t . For π_t , we use the seasonally adjusted personal consumption expenditures chain-type price index, and for i_t we use the three-month Treasury bill rate from the secondary market.¹ Our sample starts in the first quarter of 1960 and runs through the fourth quarter of 2014. We make the data quarterly by using the quarterly average of the interest rate and the log difference of the quarterly

¹The data can be downloaded from FRED (<http://research.stlouisfed.org/fred2/>) using the mnemonics PCECTPI and TB3MS.

average of the price index. Because of the effective lower bound on nominal interest rates, the shadow rate is assumed to be censored from the first quarter of 2009 through the end of our sample. While the censoring in our data comes only once and at the end of our sample, our estimation approach is applicable to a much wider class of censoring problems, which could accommodate multiple censored variables (like short and long-term interest rates), as well as multiple periods of censoring.

2.3 Estimation

To estimate the parameters and unobserved states of the model, we use Bayesian methods.² The novel modeling contribution of this paper lies in the sampling of the unobserved trend and gap components of the data when the shadow rate data are censored, so we focus on this step of the estimation procedure. Conditional on parameter values and a sequence of volatilities, our model can be put into the form

$$\xi_t = A_t \xi_{t-1} + B_t \varepsilon_t \quad (7)$$

$$y_t = C_t \xi_t \quad (8)$$

$$y_t^* = \tilde{C}(y_t) = \begin{bmatrix} \pi_t \\ \max(i_t, 0) \end{bmatrix} \quad (9)$$

where $y_t = [\pi_t, i_t]'$, $\xi_t = [\tilde{r}_t, \tilde{\pi}_t, \tilde{\pi}_t, \tilde{i}_t, \tilde{i}_{t-1}, \dots, \tilde{i}_{t-p+1}]'$, ε_t is a vector of standard normal random variables, the matrices $\{A_t, B_t, C_t\}_{t=1}^T$ are constructed accordingly from the parameters and volatilities, and the function $\tilde{C}(\cdot)$ encodes the max operator in the equation for the nominal interest rate.

Our approach for drawing from the posterior of $\xi \equiv [\xi_1, \xi_2, \dots, \xi_T]'$ is to first treat the censored data as missing and take draws from the posterior distribution of ξ , which is straightforward using standard filtering and smoothing techniques. Knowing that the data are censored (not simply missing) amounts to knowing that the values of ξ that are consistent with the data imply values of i_t that are less than the effective lower

² We assume that $\rho(L)$ contains four lags ($p = 4$).

bound during the period in which the short-term nominal interest rate has been at its lower bound. Thus, we can draw from the posterior of ξ , treating the censored data as missing and rejecting draws until we find a draw that is consistent with the lower bound. Our estimation procedure is thus a generalization of Park et al. (2007) that applies the methodology of Hopke et al. (2001). Appendix A explains in detail how to construct a draw from the posterior distribution of ξ , conditional on the data and the censoring. With a draw of ξ in hand, the posterior distribution of the parameters can be sampled using standard methods in the literature on conditionally-linear time series models with time-varying parameters and stochastic volatility, such as those used in Primiceri (2005) or Cogley and Sargent (2005).³

To estimate the parameters and unobserved states of the model, we initialize parameters and states from their prior distributions, sample 50,000 draws from the posterior distribution, discard the first 25,000 and check for convergence across multiple sampling chains using the convergence diagnostic statistic of Gelman and Rubin (1992).

The top panel of figure 1 shows features of the posterior distribution of the shadow rate of interest. Note that the posterior is simply a point mass on the observed rate when the nominal rate is above its effective lower bound. When the nominal interest rate is at its effective lower bound, every draw of the shadow rate of interest is required to be below the lower bound. The median of the posterior distribution of the shadow rate reached its lowest level in 2013, at about -2.5 percent and has since stayed near that value. These estimates are not dissimilar from those reported in Wu and Xia (2014). Our estimation procedure also delivers the full posterior distribution of the unobserved shadow rate, allowing us to characterize the uncertainty around it. Notably, the confidence sets are somewhat wide.

[Figure 1 about here.]

³First, draw ξ . Second, draw ρ and β from their posteriors, which involves only a regression equation in \tilde{i}_t and $\tilde{\pi}_t$. Third, draw $\sigma_{\tilde{\pi}}^2$ and $\sigma_{\tilde{i}}^2$ from their posteriors, conditional on the innovations in each equation. Fourth, draw the stochastic volatilities, $\sigma_{\tilde{\pi},t}^2$ and $\sigma_{\tilde{i},t}^2$, according to Kim et al. (1998). Fifth, draw $\delta_{\tilde{\pi}}^2$ and $\delta_{\tilde{i}}^2$ conditional on the innovations from the stochastic volatility equations, which are observable. Repeat. Note that the ordering is important, as explained by Del Negro and Primiceri (2014).

An important feature of our estimation approach is that the knowledge that the short-term interest rate has been below the effective lower bound informs the estimation of the rest of the model. If we simply assumed that the shadow rate data have been missing since the beginning of 2009, rather than censored by the effective lower bound, our estimates change in interesting ways. The bottom panel of figure 1 shows the shadow rate when we make this alternative assumption. Notably, it quickly returns to its trend level (changes in which are driven by changes in the inflation trend). Without more information about the driving forces behind the interest rate gap, the model does not predict a prolonged period at the effective lower bound from inflation data alone. While the 90 percent confidence set includes zero, only about 5 percent of the draws from the posterior distribution imply a shadow rate below the effective lower bound through the end of our sample.

The top panel of figure 2 shows features of the posterior distribution of the trend for the short-term real interest rate. The estimated trend has a downward tilt, however the confidence sets are wide. It is perhaps not surprising that the model has a difficult time delivering precise estimates of the trend in the short-term real rate of interest given the difficulty surrounding the equilibrium real interest rate (see Clark and Kozicki (2005)). Additionally, in this bivariate model, the only source of persistence in the interest rate gap is from its own lags, which leads to a very persistent estimates of the gap process, making uncertainty about the trend less important for model performance. Nevertheless, the downward trend in the median of the posterior estimates is clear. When we estimate the model assuming that the interest rate data since 2009 are missing, the bottom panel of figure 2 shows that we estimate a somewhat less pronounced decline in the trend for the short-term real interest rate. However, it remains true that at the end of our sample, the median of the posterior is at its lowest point.

[Figure 2 about here.]

Features of the posterior distributions of the trend processes for inflation and the shadow rate are shown in figure 3. Since the early 1980s, the model estimates that the

trend level of inflation has clearly fallen and has since remained relatively stable. The decline in the inflation trend and the decline in the trend of real interest rate combine to have the trend of the shadow rate at its lowest level (about 2 percent) at the end of our sample.

[Figure 3 about here.]

While this small, bivariate model is already able to tell us a great deal, its parsimony makes us unable to comment on movements in unemployment and output, which may also be important determinants of real interest rates. We accommodate these data in the model presented in the next section.

3 Macroeconomic Model

This section presents shadow-rate estimates derived from a richer model, combining the trend-cycle decomposition for nominal rates and inflation considered above with a macroeconomic time-series model that characterizes also the dynamics of real output and the unemployment rate. By considering joint dynamics of trends and cycles in the nominal rate of interest and macroeconomic variables, this modeling effort builds on earlier work by Laubach and Williams (2003) and Clark and Kozicki (2005); as a notable extension our model allows for time-varying parameters, in particular time-varying volatilities of shocks to trends and cycles. Considering only the dynamics for output and unemployment within a similar stochastic volatility model, Mertens (2014) found sizable reductions in the revisions between real-time and revised measures.

3.1 A trend-cycle decomposition for macro variables

As before, our macro model distinguishes between the actual nominal interest rate, i_t^* , which is constrained by the effective lower bound, and a potentially negative shadow rate of interest, i_t . As in Section 2, the shadow rate is decomposed into a gap component, \tilde{i}_t , and a trend component that equals the sum of inflation trend, $\bar{\pi}_t$, and real-rate

trend, \bar{r}_t :

$$i_t^* = \max(i_t, 0) \qquad i_t = \bar{\pi}_t + \bar{r}_t + \tilde{i}_t \qquad (10)$$

While the previous section assumed purely autoregressive dynamics for the shadow-rate gap, \tilde{i}_t , the macro model in this section relates the shadow-rate gap to other macroeconomic gap variables. Specifically, in the spirit of an interest-rate reaction function known from the literature on monetary policy (Taylor, 1999) we assume that the shadow-rate gap depends not only on its own lag but also on current and lagged values of the inflation gap and an output gap, \tilde{y}_t , that is yet to be defined.⁴

$$\tilde{i}_t = \rho_{\tilde{i}} \tilde{i}_{t-1} + d_{\tilde{\pi}}(L) \tilde{\pi}_t + d_{\tilde{y}}(L) \tilde{y}_t + \sigma_{\tilde{i},t} \varepsilon_{\tilde{i},t} \qquad (11)$$

As evident in (11), the shadow-rate gap is also influenced by an idiosyncratic shock $\varepsilon_{\tilde{i},t}$ that is subject to stochastic volatility. Our assumptions about volatilities in the macro model are similar to those made in Section 2 and will also be summarized further below. As before, shocks to different variables x and z , denoted $\varepsilon_{x,t}$ and $\varepsilon_{z,t}$, respectively, are assumed to be mutually independent standard normal random variables.

Similar to the bivariate model, the shadow-rate trend is the sum of inflation trend, $\bar{\pi}_t$, and real-rate trend, \bar{r}_t . Our macro model embeds the processes for inflation and the real rate in a multivariate context that includes also data on real output and the unemployment. While we retain the assumption of an exogenous inflation trend process identical to what has been described in Section 2, the specification for the inflation gap will be augmented with a Phillips-curve relationship. Furthermore, we are also interested in studying the relationship between the real-rate trend and long-term growth. Specifically, denoting long-term growth in GDP by μ_t we assume for the

⁴However, it should be noted that our shadow-rate equation (11) is solely cast in terms of gaps whereas conventional Taylor-type rules describe the reaction of the actual level of the nominal rate, typically assuming constant longer-run values for inflation and the real rate instead of the drifting trends considered here. Furthermore, the reaction function (11) imparts interest rate inertia by smoothing with respect to the lagged value of the *shadow*-rate gap, not the actual interest rate with non-trivial implications for instances when the lower-bound constraint binds for the lagged level of the nominal rate.

following dependence between the trends in output growth and the real rate:⁵

$$\mu_t = \mu_{t-1} + \sigma_\mu \varepsilon_{\mu,t} \quad (12)$$

$$\bar{r}_t = \bar{r}_{t-1} + \beta_{r,\mu} \varepsilon_{\mu,t} + \sigma_{\bar{r}} \varepsilon_{\bar{r},t} \quad (13)$$

$$= \beta_{r,\mu} \mu_t + \bar{r}_t^{\perp\mu} \quad \text{where} \quad \bar{r}_t^{\perp\mu} = \bar{r}_{t-1}^{\perp\mu} + \sigma_{\bar{r}} \varepsilon_{\bar{r},t}$$

The output gap is derived from a well-studied trend-cycle decomposition of real GDP known also as the ‘‘Harvey-Clark’’ model (Harvey, 1985; Clark, 1987). Specifically, the level of real output, y_t , is split into a trend, \bar{y}_t , and a gap component, \tilde{y}_t . The trend is modeled as a random walk with time-varying drift given by the long-term growth process for μ_t specified in (12). Time variation in this expected rate of trend growth allows the model to capture low-frequency changes in real growth experienced by the U.S. economy over the postwar period.⁶ As in the original specification of the Harvey-Clark model, the output gap is characterized by an AR(p) process:⁷

$$y_t = \bar{y}_t + \tilde{y}_t \quad (14)$$

$$\bar{y}_t = \bar{y}_{t-1} + \mu_{t-1} + \sigma_{\bar{y},t} \varepsilon_{\bar{y},t} \quad (15)$$

$$a(L)\tilde{y}_t = \beta_{\bar{y},\tilde{y}} \sigma_{\bar{y},t} \varepsilon_{\bar{y},t} + \sigma_{\tilde{y},t} \varepsilon_{\tilde{y},t}. \quad (16)$$

Trend growth in real output is thus driven by two components: Level shocks $\varepsilon_{\bar{y},t}$ that are subject to stochastic volatility and the constant-volatility shocks to the expected rate of trend growth, μ_t .

We embed the Harvey-Clark assumption of an autoregressive output gap in a small factor model of the macroeconomy. Dynamic factor models relate fluctuations in mul-

⁵When estimating the macro model, real GDP is measured in per-capita terms using the log-difference between the FRED series `GDPC1` and `POP`; the unemployment rate is measured by `UNRATE`.

⁶Average growth rates in real GDP have mostly been trending downwards in U.S. data over the last five decades or so. These low-frequency changes have been documented for example by Gordon (2012) and Stock and Watson (2012) and have already been noted in the textbook of Blanchard and Fischer (1989, Chapter 1).

⁷Throughout we report results for $p = 2$, same as in Watson (1986), Harvey (1985), and Clark (1987). While we are investigating the sensitivity of our results to this choice further, a preliminary analysis suggests the robustness of findings presented here to different choices of lag length.

multiple variables to a smaller set of dynamic factors; in our case, a single factor — the output gap — accounts for the joint persistence in the gap dynamics of output, inflation, and the unemployment rate. The shadow-rate gap inherits persistence from this macro factor which is then propagated further by the own-persistence of the shadow-rate gap embedded in (11). The unemployment rate and inflation are assumed to be linked with the common output-gap factor via an Okun’s Law relationship and an empirical Phillips curve, respectively.

In further contrast to the original Harvey-Clark model, the output gap process in (16) allows for correlation between level shocks to the trend and gap components of output. As noted by Morley et al. (2003) estimates of the trend and cycle in U.S. output can be very sensitive to the originally assumed orthogonality between trend and cycle shocks in the Harvey-Clark model, at least when derived from a univariate model, see also Watson (1986). Furthermore, as documented by Orphanides and van Norden (2002), univariate output gap estimates can be sensitive to the particular data vintage used in the estimation.⁸ To some good extent, this simply reflects the uncertain nature of modeling and measuring these latent variables.⁹ But as demonstrated for example by Laubach (2001) or Fleischman and Roberts (2011), much of this uncertainty can be reduced by conditioning the estimates on multivariate information. At least for U.S. data, the unemployment rate seems to provide a highly useful and timely signal about the output gap.¹⁰ These results have recently been extended by Mertens (2014), who documented not only improvements in the reliability of output gap estimates due the inclusion of unemployment rate data but also from accounting for time-variation in the size of business cycles by including stochastic volatility in the modeling assumptions.

Using a similar specification as the one employed here, Mertens (2014) also found

⁸Orphanides and van Norden (2002) consider a wide range of mostly univariate models for real output as well as a few bivariate models that also condition their output gap estimates on inflation without including the unemployment rate as it is done here.

⁹Clark and Kozicki (2005) report considerable real-time uncertainty in estimates of the trend real rate echoing the results of Orphanides and van Norden (2002) regarding the output gap. We plan to investigate these questions further in the context of our model in a future revision of this draft.

¹⁰Apart from revisions in seasonal adjustments, unemployment data for the U.S. is also not subject to the kind of data revisions due to collection issues and methodological changes known from GDP data for example.

output gap estimates to be more robust to assumptions about correlations between shocks to trend and cycle than suggested by Morley et al. (2003) and Sinclair (2009) in the context of univariate and bivariate models with constant-parameters.

As with the other variables discussed so far, the unemployment rate, u_t , is modeled as the sum of a trend and a gap component; the former is described by a driftless random walk and the latter is linked to the output gap via an Okun’s law relationship.

$$u_t = \bar{u}_t + \tilde{u}_t \quad \bar{u}_t = \bar{u}_{t-1} + \sigma_{\bar{u}}\varepsilon_{\bar{u},t} \quad \tilde{u}_t = b_u(L)\tilde{y}_t + \sigma_{\tilde{u}}\varepsilon_{\tilde{u},t} \quad (17)$$

The unemployment trend identified by our model is conceptually identical to the “NAIRU” estimates of Staiger et al. (1997), Gordon (1997), Laubach (2001), Reifschneider et al. (2013), Stella and Stock (2013) and Watson (2014), to name but a few. However it should be stressed, that there is no presumption here that this long-term forecast for unemployment derived from our purely statistical decomposition is necessarily related to the structural concept of a “non-accelerating inflation rate of unemployment.” As discussed by Watson (2014), estimates of the variance of shocks to the unemployment trend are hard to come by and so we have chosen not to endow this trend with stochastic volatility, see also Stella and Stock (2013).¹¹

Considering the inflation process, we augment the simple UC-SV model for inflation that formed part of our bivariate model presented in Section 2 with a Phillips-curve relationship for the inflation gap.

$$\pi_t = \bar{\pi}_t + \tilde{\pi}_t \quad \bar{\pi}_t = \bar{\pi}_{t-1} + \sigma_{\bar{\pi},t}\varepsilon_{\bar{\pi},t} \quad \tilde{\pi}_t = b_\pi(L)\tilde{y}_t + \sigma_{\tilde{\pi},t}\varepsilon_{\tilde{\pi},t} \quad (18)$$

In constant parameter models, similar specifications have for example been used by Kuttner (1992), Gerlach and Smets (1997), Laubach and Williams (2003) and Fleis-

¹¹Mertens (2014) considers stochastic volatility shocks to trend unemployment and finds no substantial time-variation in the estimated volatilities either. Our prior for $\sigma_{\bar{u}}$ is an inverse Wishart with only three degrees — and thus only mildly informative — and a mean set equal to the lower range of values considered by Watson (2014) ($\sigma_{\bar{u}}^2 = 0.1^2$). Results are insensitive to choosing any of the higher values considered by Watson (2014).

chman and Roberts (2011). However, as recently surveyed by Stock and Watson (2009), empirical linkages between inflation and slack seem to have evolved, and might be less clear than once thought, likely due to non-linearities in joint dynamics of inflation and real activity as reported by Stock and Watson (2010) or Stella and Stock (2013); our approach nests the multivariate UC-SV model for unemployment and inflation of Stella and Stock.

Turning to the variance-covariance structure of this macro model, the following shocks have been endowed with stochastic volatility: shocks to trend level and gap in output, inflation trend shock and the idiosyncratic shocks to inflation and the shadow rate. The remaining shocks — namely those to expected trend growth, the trends in unemployment rate and real rate as well as the Okun’s law error — seemed not to display substantial time variation in a preliminary analysis and are assumed to be homoscedastic here. All stochastic volatility processes are specified as independent random walks in the log-variances as described in Section 2 above.

3.2 Estimates of the shadow rate and macro trends

As in the bivariate model presented in Section 2, estimates of the shadow rates after 2008 depend quite a bit on whether nominal-rate data is treated as censored or missing when observations are constrained by the effective lower bound. In contrast to the bivariate model, shadow rates sampled from the macro model reflect also information about the state of the macroeconomy, in particular the depth of the last recession, prompting also important differences in the macro model’s shadow-rate estimates, shown in Figure 4, relative to those generated by the bivariate model.

[Figure 4 about here.]

Whether the nominal-rate data has been treated as censored or not, the estimated shadow-rate displays a marked trough in the middle of 2009, dropping below minus 4%. Even when nominal-rate data since 2009 is treated as missing during the model estimation, the simulated posterior places almost no mass on positive values of the

shadow rate during 2009, as shown in the bottom panel of Figure 4. Measured by the posterior median, the shadow rate is estimated to have run below the effective lower bound until about mid-2011 — even as the macro model discarded the censoring information available from nominal-rate data seen since 2009. In contrast, the corresponding posterior generated from the bivariate model places more-than-even odds on positive values for the shadow rate over this entire lower-bound episode, as shown in the bottom panel of Figure 1.

[Figure 5 about here.]

A key factor for the markedly negative shadow-rate estimates generated by the macro model is the depth of the recession reflected in the model’s estimates of the gaps in output and unemployment that are shown in Figure 5. Reflecting a strong Okun-Law relationship estimated from the data, both the output gap — depicted in the top panel of the figure — and the unemployment gap — shown in the bottom panel — display a similar cyclical pattern.¹²

The stochastic-volatility specification of our model endows our model with important non-linearities and allow our “business cycle factor”, measured by the output gap, to flexibly adapt to the well-documented changes in the size of the business cycle over the postwar period (Bernanke, 2004).¹³ Moreover, our estimates of the business cycle reflect the well-known asymmetry that recessions tend to be deeper and shorter than expansions in the U.S. (Morley and Piger, 2012).¹⁴

As shown in the top panel of Figure 5, our output gap estimates accord also remarkably well with what is implied by the most recent estimates for potential real GDP made by the Congressional Budget Office (CBO) with a production function approach (Congressional Budget Office, 2001) that is very different from our purely statistical

¹²The posterior distribution of the sum of Okun’s Law coefficients $b_{\bar{y}}(1)$ from (17) has a posterior median of -0.54 with a 90% credible set ranging from about -0.60 to -0.49 , compared to a normal prior $N(-0.50, 0.2^2)$, corresponding to a prior median of -0.50 and 90% credible set ranging from -0.76 to -0.24 .

¹³Estimates of the stochastic volatility in shocks to the output gap and the level of trend output, not shown, display a marked decrease over the 1980s, while the last recession is mostly mirrored by an increase in gap volatility, similar to the results of Mertens (2014).

¹⁴This result is even more remarkable considering the absence of any direct asymmetries in our non-linear model specification.

trend-cycle decomposition. To the extent that our output gap estimates seem to reflect widely held views, it also makes sense to let our deep output gap estimates for the years 2008 drive the largely negative shadow-rate estimates over that period via our estimated reaction function for the shadow-rate gap defined in (11).

[Figure 6 about here.]

While output gap estimates from the macro model clearly attribute a good portion of weaknesses in real GDP observed since 2008 to an unusually depressed cyclical factor this does not fully account for the low growth rates in real GDP seen over that period. As shown in Figure 6, our estimates of expected trend growth, μ_t , display a steady decline over most of the sample period. Except for temporary increases in the late 1990s, expected trend growth has almost continuously drifted down from an annualized rate of about 3.5% in 1960 to a low near 0.5% at the end of 2014. The figure also reports the posterior median estimate of \bar{r}_t , the real-rate trend, which has also steadily declined from nearly 2% in 1960 to about 1% by the end of 2014. Economically, these are highly significant declines and speak to the concerns raised for example by Fernald (2012), Gordon (2014), and the survey volume edited by Teulings and Baldwin (2014).

While low-frequency changes in growth and the real rate — not to mention evidence about their interdependence — are generally hard to measure, we also find signs of a statistical relationship between longer-term growth and the real-rate trend as measured by the sensitivity $\beta_{\bar{r},\mu}$ in the real-rate equation (13) that we intend to investigate more deeply in future updates of this draft. Currently, our prior for $\beta_{\bar{r},\mu}$ is a relatively vague $N(0, 1)$, leading to a posterior median of 0.53 and a posterior inter-quartile range from -0.02 to 1.08 (compared to -0.76 to 0.76 under the prior).

[Figure 7 about here.]

Figure 7 displays the considerable uncertainty surrounding the aforementioned median estimates of the real-rate trend; the width of the 90% credible set generally covers a couple of percentage points and even the inter-quartile range stretches from almost

zero to about 1.75% for late 2014. As shown in the bottom panel of the figure, these estimates are not much affected by treating the nominal-rate data since 2009 as censored or missing.¹⁵

4 Long-Term Interest Rates

Our models deliver expected paths for future short-term interest rates. In this section, we compute the expectations component of longer-term interest rates using this path. We define the expectations component of longer-term interest rates on bonds with maturity h periods in the future as

$$e_{t,t+h} \equiv E_t \left(\frac{1}{h} \sum_{j=0}^{h-1} i_{t+j}^* \right).$$

The observed longer-term interest rate is given by

$$i_{t,t+h}^* = e_{t,t+h} + p_{t,t+h}$$

where $p_{t,t+h}$ represents premiums. Note that for any j , $E_t(i_{t+j}^*) \geq 0$ because of the effective lower bound on nominal interest rates.

In our models, conditional on parameters, i_{t+j}^* has a non-normal (because of stochastic volatilities) truncated distribution. By explicitly modeling the effective lower bound on nominal interest rates, we take into account this truncation. To construct $E_t(i_{t+j}^*)$ we first simulate $\sigma_{\cdot,t+j}^2$ for all of the stochastic volatilities. The volatilities in our model are assumed to evolve independently of the means of all other variables, so, conditional on parameters, they can be drawn without reference to any other process. Conditional on parameters and a sequence of $\sigma_{\cdot,t+j}^2$, our models are linear Gaussian state-space models, meaning that the distributions of i_{t+j} are normal distributions with known mean and variance. Thus, i_{t+j}^* has a truncated normal distribution, the mean of which

¹⁵Estimates of trend inflation generated by the macro model are similar to those discussed in Section 2 for the bivariate model.

can be calculated directly. For each draw of parameters from the posterior distribution, we simulate draws of $\sigma_{:,t+j}^2$ to create a distribution of $E_t(i_{t+j}^*|\sigma_{:,t+j}^2)$. We then average over those draws to create an estimator of $E_t(i_{t+j}^*)$. With values of $E_t(i_{t+j}^*)$ in hand, our estimate of $e_{t,t+h}$ for each posterior draw of model parameters is constructed directly by averaging over time horizons. We construct posterior distributions of $e_{t,t+h}$ for all time periods in our sample and compare $e_{t,t+h}$ to the observed h -period interest rate.

We focus on the 5-year and 10-year Treasury rates as longer-term interest rates. That is, we set $h = 20$ or $h = 40$. We emphasize that the model estimation does not include the longer-term interest rate in the set of observables. Figures 8 and 11 display quantiles from the posterior distribution of $e_{t,t+h}$ from our models, along with the 5-year and 10-year treasury rates. The figures also show the estimated zero-coupon yields from Kim and Wright (2005), which are very similar to the interest rate data we use. Clearly, the expectation components do not track the nominal interest rates exactly. We call the difference between the data and our estimated expectation component the premium. Figures 9 and 12 show the estimated 5-year and 10-year premiums from each model. While the confidence intervals are wider for the parsimonious model, both models deliver similar histories of premiums in that premiums jumped in the early 1980s and have declined since. As discussed above, we also find that this is the period over which volatility in the trend component of the inflation process has declined. Strikingly, the model indicates that the decline in longer-term interest rates since the onset of the Great Recession appears to be mostly due to changes in trend levels of inflation and the real interest rate, not premiums.

Using a no-arbitrage dynamic term-structure approach Kim and Wright (2005) and Wright (2011) have documented a decline in premiums since 1990. In figures 10 and 13 we compare our term premium estimates (using the same zero-coupon interest rate data) to updated estimates generated from the dynamic term-structure model of Kim and Wright (2005).¹⁶ Our estimates display similar downward trends and are

¹⁶Updated estimates of the Kim and Wright (2005) model are available at <http://www.federalreserve>.

clearly correlated with those from Kim and Wright (2005). Interestingly, our estimates of the 5-year term premium display increased volatility near recession dates, which we are able to capture because our models include stochastic volatility.

[Figure 8 about here.]

[Figure 9 about here.]

[Figure 10 about here.]

[Figure 11 about here.]

[Figure 12 about here.]

[Figure 13 about here.]

Incorporating the effective lower bound on nominal interest rates is crucial when calculating the expectations component of longer-term interest rates. Figures 14 and 15 shows the estimates of the expectation component of the 5-year and 10-year interest rate and compare it with the expected average shadow rate over that horizon. The expected average shadow rate is the analogue of the expectation component of the longer-term interest rate if the short-term interest rate were not constrained by its effective lower bound. Note that toward the end of our sample, the truncation of i_{t+j}^* has large effects on $e_{t,t+h}$, and effectively keeps the longer-term rates from falling. Earlier in our sample, the difference between the expectation component of the longer-term rate and the expected average future shadow rate is small because interest rates were relatively far from the effective lower bound. Our model thus suggests one channel through which the co-movement of short- and longer-term interest rates may decline in an environment in which interest rates are low. A pertinent example of such a decline is the U.S. experience over the recent decade—also known as the period of the “bond conundrum.” Our estimates suggest the subdued increases of longer-term yields during the Federal Reserve’s period of tightening that began in 2004 may at least in part reflect the very low level of the short-term nominal interest rate at that time.

[Figure 14 about here.]

[Figure 15 about here.]

5 Conclusion

In this paper, we developed a methodology to include the effective lower bound on nominal interest rates in and estimate the parameters of a broad class of widely-used time series models. We document a pronounced downward tilt in the trend levels of the short-term nominal interest rate, output growth, and the short-term real interest rate. Because of declines in the trend level of inflation, we also find that the longer-run normal level of the nominal interest rate is now as low as it has been over our entire sample.

When we compute expected average paths of future short-term nominal interest rates and compare them with longer-term bond yields, we find that premiums have declined noticeably since the 1980s. Notably, our estimates indicate that premiums have remained stable and near zero during the Great Recession. By explicitly modeling the effective lower bound on nominal interest rates, our time series model predicts that changes in short-term rates will have a different relationship with the expectation component of longer-term rates when the nominal interest rate is near the effective lower bound as compared to when it is not.

Finally, the posterior estimates presented in this paper take into account all available data. In future versions, we would like to present real-time estimates of important variables like the trend real interest rate, the trend growth rate, and the natural rate of unemployment. These real-time estimates would provide a useful comparison to help understand how more-recent data affect the way the model views the historical time series.

A Sampling States with Censored Data

Our Gibbs sampling procedure is a generalization of Park et al. (2007) that applies the methodology of Hopke et al. (2001). Assume that the vector ξ_t is a random variable that evolves so that

$$\xi_t = A_t \xi_{t-1} + B_t \varepsilon_t$$

where ε_t is a vector of standard normal random variables of appropriate length and the sequence of matrices $\{A_t\}_{t=1}^T$ and $\{B_t\}_{t=1}^T$ are given. Define the vector

$$y_t = C_t \xi_t$$

where the sequence of matrices $\{C_t\}_{t=1}^T$ are known. The vector y_t is equivalent to the observation vector in a linear-Gaussian-state-space model. We generalize the model to incorporate a censoring function, $\tilde{C}_t(\cdot)$ so that the vector

$$y_t^* = \tilde{C}_t(y_t)$$

is actually observed. The censoring functions $\tilde{C}_t(\cdot)$ are defined by

$$\iota_i' \tilde{C}_t(y_t) = \mathbf{1}(\iota_i' y_t \in D_{t,i}) \iota_i' y_t$$

where ι_i is a vector that has a 1 in the i 'th element and otherwise contains zeros, $D_{t,i}$ is a measurable set on the real line, and $\mathbf{1}(\cdot)$ is the indicator function. The function \tilde{C}_t works so that if the i 'th entry of y_t is in the set $D_{t,i}$, $\iota_i' y_t^* = \iota_i' y_t$. Otherwise, $\iota_i' y_t^* = 0$.¹⁷

The standard setup in which all of the data are always observed is the case in which $D_{t,i}$ is the entire real line for all i and t . In this case, given a normal distribution for ξ_0 , the joint distribution of the vectors $y = [y_1', y_2', \dots, y_T']'$ and $\xi = [\xi_1', \xi_2', \dots, \xi_T']'$ is a

¹⁷It is a trivial generalization to make the censored value non-zero. That is, to make $\iota_i' y_t^* = c_i$ when $\iota_i' y_t \notin D_{t,i}$.

multivariate normal distribution

$$\begin{bmatrix} y \\ \xi \end{bmatrix} \sim N \left(\begin{bmatrix} \mu_y \\ \mu_\xi \end{bmatrix}, \begin{bmatrix} V_{yy} & V_{y,\xi} \\ V_{\xi,y} & V_{\xi,\xi} \end{bmatrix} \right)$$

Conditional on $y = y^*$,

$$\xi \sim N \left(\hat{\mu}_\xi, \hat{V}_{\xi,\xi} \right)$$

The Kalman filter gives a convenient way to compute the conditional distribution of ξ and recover

$$\Pr(\xi_T | y = y^*)$$

and the densities

$$\Pr(\xi_t | y = y^*, \xi_{t+1})$$

can be calculated using standard smoothing techniques.

Missing data can be accommodated in a straight forward way by letting \tilde{C}_t be defined so that at time t if $D_{t,i} = \emptyset$ for all i . If we let y^o be the observed values and y^c be the missing (censored) values, then

$$\begin{bmatrix} y^o \\ y^c \\ \xi \end{bmatrix} \sim N \left(\begin{bmatrix} \mu_{y^o} \\ \mu_{y^c} \\ \mu_\xi \end{bmatrix}, \begin{bmatrix} V_{y^o,y^o} & V_{y^o,y^c} & V_{y^o,\xi} \\ V_{y^c,y^o} & V_{y^c,y^c} & V_{y^c,\xi} \\ V_{\xi,y^o} & V_{\xi,y^c} & V_{\xi,\xi} \end{bmatrix} \right)$$

And the distribution of ξ can again be calculated as a conditional distribution from the multivariate normal so that, conditional on y^* ,

$$\begin{bmatrix} y^c \\ \xi \end{bmatrix} \sim N \left(\begin{bmatrix} \tilde{\mu}_{y^c} \\ \tilde{\mu}_\xi \end{bmatrix}, \begin{bmatrix} \tilde{V}_{y^c,y^c} & \tilde{V}_{y^c,\xi} \\ \tilde{V}_{\xi,y^c} & \tilde{V}_{\xi,\xi} \end{bmatrix} \right)$$

In the context of the Kalman filter and smoother, the missing data can easily be accommodated by skipping updating steps when data are not available, allowing ξ to be sampled from its posterior distribution, conditional on y^* .

If $D_{t,i} \neq \emptyset$, but the data are also censored, then we have information that the draws of ξ must satisfy the requirement that the following expression must be equal to 1 for all i and t

$$\mathbf{1}(\iota'_i C_t \xi_t \notin D_{t,i}) \mathbf{1}(\iota'_i y_t^* = 0) + \mathbf{1}(\iota'_i y_t^* \neq 0)$$

The distribution of the censored data, y^c , and ξ conditional on y^* is then

$$\begin{bmatrix} y^c \\ \xi \end{bmatrix} \sim TN \left(\begin{bmatrix} \tilde{\mu}_{y^c} \\ \tilde{\mu}_\xi \end{bmatrix}, \begin{bmatrix} \tilde{V}_{y^c, y^c} & \tilde{V}_{y^c, \xi} \\ \tilde{V}_{\xi, y^c} & \tilde{V}_{\xi, \xi} \end{bmatrix}, D \right)$$

where TN stands for truncated normal so that the density function is

$$\Pr \left(\begin{bmatrix} y^c \\ \xi \end{bmatrix} \right) \propto \phi \left(\begin{bmatrix} \tilde{\mu}_{y^c} \\ \tilde{\mu}_\xi \end{bmatrix}, \begin{bmatrix} \tilde{V}_{y^c, y^c} & \tilde{V}_{y^c, \xi} \\ \tilde{V}_{\xi, y^c} & \tilde{V}_{\xi, \xi} \end{bmatrix} \right) \mathbf{1}(\xi \in D)$$

where ϕ is the multivariate normal density function and D is the set of values of ξ that are consistent with the observed values y^* , given the censoring function $\{\tilde{C}_t\}_{t=1}^T$. Note that for implementation,

$$\mathbf{1}(\xi \in D) = \prod_{t=1}^T \prod_i [\mathbf{1}(\iota'_i C_t \xi_t \notin D_{t,i}) \mathbf{1}(\iota'_i y_t^* = 0) + \mathbf{1}(\iota'_i y_t^* \neq 0)]$$

Notice that, conditional on y^* , draws can be taken from the distribution of ξ using rejection sampling by treating the censored data as though they are missing from the perspective of the Kalman filter and smoother, sampling ξ , and then verifying that $\mathbf{1}(\xi \in D) = 1$. If not, draw ξ again treating the censored data as missing and repeat until $\mathbf{1}(\xi \in D) = 1$.

B Censored Data in State Equation

In our baseline framework, lagged values of ξ_t appear as explanatory variables and are not censored. A straightforward extension is to incorporate lagged values of y_t^* , that is, the censored observations. In this case,

$$\xi_t = A_t \xi_{t-1} + F_t x_{t-1}^* + B_t \varepsilon_t$$

where $\{F_t\}_{t=1}^T$ are conformable matrices that are known, $x_{t-1}^* = [y_{t-1}^{*'}, y_{t-2}^{*'}, \dots, y_{t-p}^{*'}]'$ for some given p , and all other model components are defined as in our baseline model.

We maintain the assumptions that

$$y_t = C_t \xi_t$$

where the sequence of matrices $\{C_t\}_{t=1}^T$ are known and that

$$y_t^* = \tilde{C}_t(y_t).$$

The posterior of ξ can be constructed exactly as in the previous section, treating $\{x_{t-1}^*\}_{t=1}^T$ as exogenous in every period because the rejection step will ensure that the sampled values of ξ are consistent with x_{t-1}^* for all t . For comparison, it is worth noting that the models of Iwata and Wu (2006) and Nakajima (2011) can be cast, conditional on parameter values, as special cases of this setup in which the matrix $A_t = \mathbf{0}$.

References

- Michael D. Bauer and Glenn D. Rudebusch. Monetary policy expectations at the zero lower bound. *Federal Reserve Bank of San Francisco Working Paper Series*, 2014.
- Ben S. Bernanke. The Great Moderation. Remarks by Governor Ben S. Bernanke at the meetings of the Eastern Economic Association, Washington D.C., February 2004.
- Stephen Beveridge and Charles R. Nelson. A new approach to decomposition of economic time series into permanent and transitory components with particular attention to measurement of the business cycle. *Journal of Monetary Economics*, 7(2): 151–174, 1981.
- Fischer Black. Interest rates as options. *The Journal of Finance*, 50(5):1371–1376, 1995.
- Olivier Jean Blanchard and Stanley Fischer. *Lectures on Macroeconomics*. The MIT Press, Cambridge, MA, 1989.
- John Y. Campbell and Robert J. Shiller. Yield spreads and interest rate movements: A bird’s eye view. *The Review of Economic Studies*, 58(3):495–514, 1991.
- Joshua C.C. Chan and Rodney Strachan. The zero lower bound: Implications for modelling the interest rate. *Working Paper*, 2014.
- Peter K Clark. The cyclical component of u.s. economic activity. *The Quarterly Journal of Economics*, 102(4):797–814, November 1987.
- Todd E. Clark and Sharon Kozicki. Estimating equilibrium real interest rates in real time. *The North American Journal of Economics and Finance*, 16(3):395–413, December 2005.
- Timothy Cogley and Thomas J. Sargent. Drifts and volatilities: monetary policies and outcomes in the post WWII US. *Review of Economic Dynamics*, 8(2):262–302, 2005.
- Congressional Budget Office. CBO’s method for estimating potential output: An update, 2001.
- Marco Del Negro and Giorgio Primiceri. Time varying structural vector autoregressions and monetary policy: A corrigendum. *Mimeo*, 2014.
- John Fernald. Productivity and potential output before, during, and after the Great Recession. Working Paper Series 2012-18, Federal Reserve Bank of San Francisco, 2012.
- Charles A. Fleischman and John M. Roberts. From many series, one cycle: Improved estimates of the business cycle from a multivariate unobserved components model. Finance and Economics Discussion Series 2011-46, Board of Governors of the Federal Reserve System (U.S.), 2011.

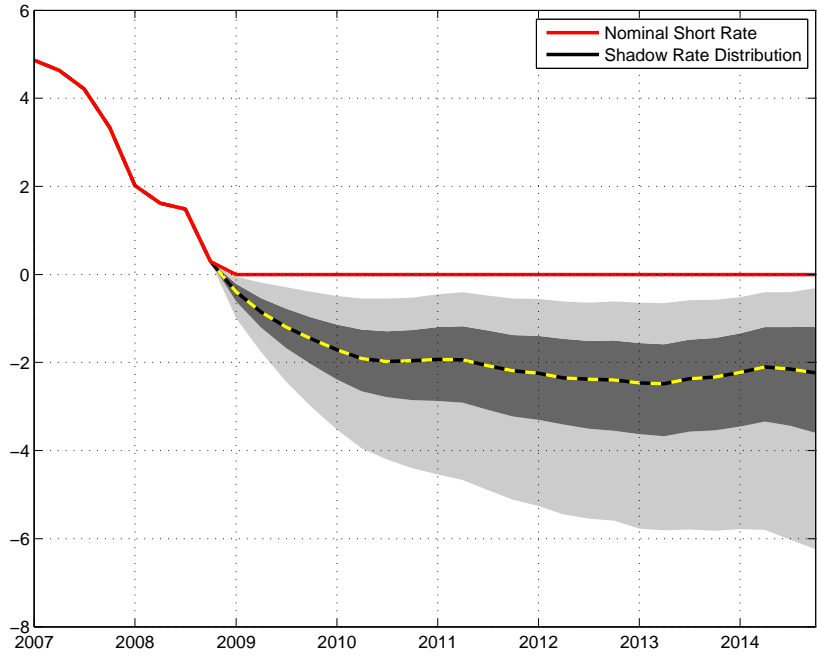
- Andrew Gelman and Donald B. Rubin. Inference from iterative simulation using multiple sequences. *Statistical Science*, 7(4):457–472, 1992.
- Stefan Gerlach and Frank Smets. Output gaps and inflation: Unobservable components estimates for the G-7 countries. *mimeo* Bank for International Settlements, January 1997.
- Robert J. Gordon. The time-varying NAIRU and its implications for economic policy. *Journal of Economic Perspectives*, 11(1):11–32, Winter 1997.
- Robert J. Gordon. Is U.S. economic growth over? Faltering innovation confronts the six headwinds. NBER Working Papers 18315, National Bureau of Economic Research, Inc, August 2012.
- Robert J. Gordon. The demise of U.S. economic growth: Restatement, rebuttal, and reflections. NBER Working Papers 19895, National Bureau of Economic Research, Inc, February 2014.
- James D. Hamilton, Ethan S. Harris, Jan Hatzius, and Kenneth D. West. The equilibrium real funds rate: Past, present and future. *Working Paper*, March 2015.
- A C Harvey. Trends and cycles in macroeconomic time series. *Journal of Business & Economic Statistics*, 3(3):216–27, June 1985.
- Philip K. Hopke, Chuanhai Liu, and Donald B. Rubin. Multiple imputation for multivariate data with missing and below-threshold measurements: Time-series concentrations of pollutants in the arctic. *Biometrics*, 57(1):22–33, 2001. ISSN 0006341X.
- Shigeru Iwata and Shu Wu. Estimating monetary policy effects when interest rates are close to zero. *Journal of Monetary Economics*, 53(7):1395 – 1408, 2006.
- Don Kim and Kenneth J. Singleton. Term structure models and the zero bound: An empirical investigation of Japanese yields. *Mimeo*, 2011.
- Don H. Kim and Jonathan H. Wright. An arbitrage-free three-factor term structure model and the recent behavior of long-term yields and distant-horizon forward rates. *FEDS Working Paper Series*, 2005.
- Sangjoon Kim, Neil Shephard, and Siddhartha Chib. Stochastic volatility: Likelihood inference and comparison with arch models. *The Review of Economic Studies*, 65(3):361–393, 1998.
- Leo Krippner. Measuring the stance of monetary policy in zero lower bound environments. *Economics Letters*, 118(1):135–138, 2013.
- Kenneth Kuttner. Monetary policy with uncertain estimates of potential output. *Economic Perspectives*, 16:2–15, January 1992.
- Thomas Laubach. Measuring the NAIRU: Evidence from seven economies. *The Review of Economics and Statistics*, 83(2):218–231, May 2001.

- Thomas Laubach and John C. Williams. Measuring the natural rate of interest. *The Review of Economics and Statistics*, 85(4):1063–1070, November 2003.
- Elmar Mertens. On the reliability of output-gap estimates in real time. *mimeo* Federal Reserve Board, November 2014.
- James Morley and Jeremy Piger. The asymmetric business cycle. *The Review of Economics and Statistics*, 94(1):208–221, February 2012.
- James C. Morley, Charles R. Nelson, and Eric Zivot. Why are the Beveridge-Nelson and unobserved-components decompositions of GDP so different? *The Review of Economics and Statistics*, 85(2):235–243, May 2003.
- Jouchi Nakajima. Monetary policy transmission under zero interest rates: An extended time-varying parameter vector autoregression approach. *Bank of Japan Working Paper Series*, 2011.
- Athanasios Orphanides and Simon van Norden. The unreliability of output-gap estimates in real time. *The Review of Economics and Statistics*, 84(4):569–583, November 2002.
- Jung Wook Park, Marc G. Genton, and Sujit K. Ghosh. Censored time series analysis with autoregressive moving average models. *The Canadian Journal of Statistics / La Revue Canadienne de Statistique*, 35(1):151–168, 2007.
- Giorgio E. Primiceri. Time varying structural vector autoregressions and monetary policy. *The Review of Economic Studies*, 72(3):821–852, 2005.
- Dave Reifschneider, William Wascher, and David Wilcox. Aggregate supply in the United States: Recent developments and implications for the conduct of monetary policy. Finance and Economics Discussion Series 2013-77, Board of Governors of the Federal Reserve System (U.S.), 2013. URL <http://ideas.repec.org/p/fip/fedgfe/2013-77.html>.
- Christopher A. Sims. Macroeconomics and reality. *Econometrica*, 48(1):1–48, 1980.
- Tara M. Sinclair. The relationships between permanent and transitory movements in U.S. output and the unemployment rate. *Journal of Money, Credit and Banking*, 41(2-3):529–542, 03 2009.
- Douglas O. Staiger, James H. Stock, and Mark W. Watson. How precise are estimates of the natural rate of unemployment? In *Reducing Inflation: Motivation and Strategy*, NBER Chapters, chapter 5, pages 195–246. National Bureau of Economic Research, Inc, 1997.
- Andrea Stella and James H. Stock. A state-dependent model for inflation forecasting. *forthcoming* in *Unobserved Components and Time Series Econometrics*, Oxford University Press, 2013.
- James H. Stock and Mark W. Watson. Why has U.S. inflation become harder to forecast? *Journal of Money, Credit and Banking*, 39:3–33, 2007.

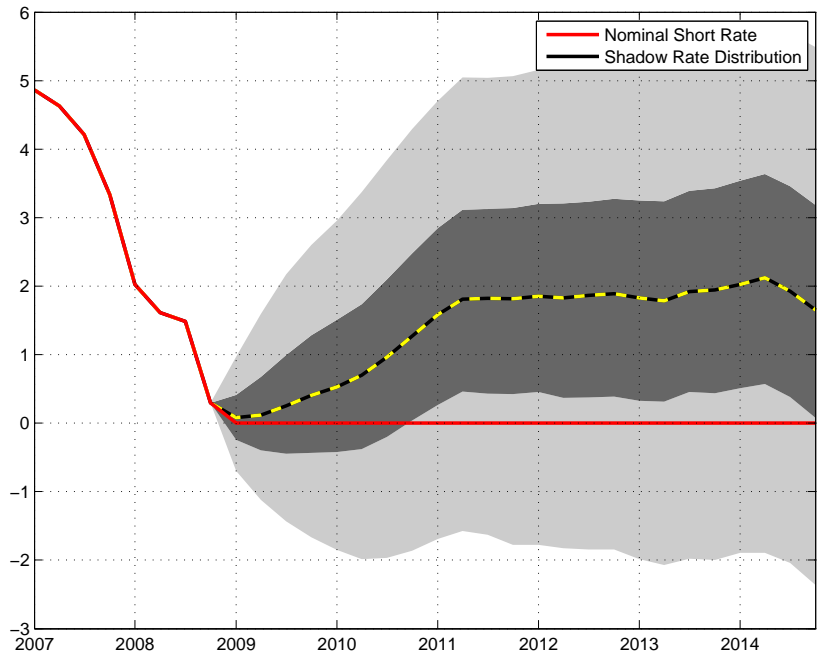
- James H. Stock and Mark W. Watson. Phillips curve inflation forecasts. In Jeff Fuhrer, Yolanda K. Kodrzycki, Jane Sneddon Little, and Giovanni P. Olivei, editors, *Understanding Inflation and the Implications for Monetary Policy — A Phillips Curve Retrospective*, chapter 3, pages 99–184. MIT Press, Cambridge, MA, September 2009.
- James H. Stock and Mark W. Watson. Modeling inflation after the crisis. NBER Working Papers 16488, National Bureau of Economic Research, Inc, October 2010.
- James H. Stock and Mark W. Watson. Disentangling the channels of the 2007-2009 recession. NBER Working Papers 18094, National Bureau of Economic Research, Inc, May 2012.
- John B. Taylor, editor. *Monetary Policy Rules*. The University of Chicago Press, Chicago, 1999. (Papers based on the National Bureau of Economic Research Conference on Research in Business Cycles, held at Stanford University 1998).
- Coen Teulings and Richard Baldwin, editors. *Secular Stagnation: Facts, Causes and Cures*. VOXeu.org ebook. CEPR Press, 2014.
- Mark W. Watson. Univariate detrending methods with stochastic trends. *Journal of Monetary Economics*, 18(1):49–75, July 1986.
- Mark W. Watson. Inflation Persistence, the NAIRU, and the Great Recession. *American Economic Review*, 104(5):31–36, May 2014.
- Jonathan H. Wright. Term premia and inflation uncertainty: Empirical evidence from an international panel dataset. *American Economic Review*, 101(4):1514–34, 2011.
- Cynthia Wu and Fan Dora Xia. Measuring the macroeconomic impact of monetary policy at the zero lower bound. *Mimeo*, 2014.

Figure 1: Shadow rate of interest from the bivariate model.

(a) Nominal rate data treated as censored



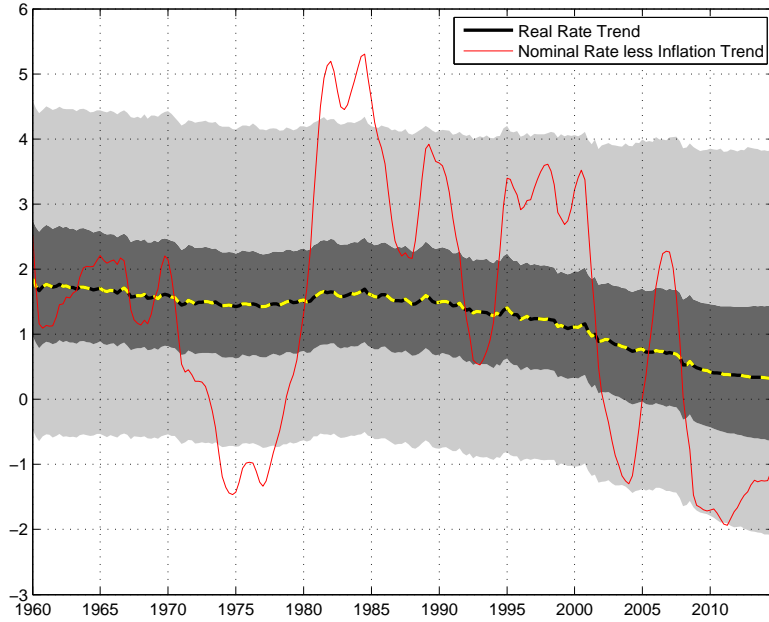
(b) Nominal rate data treated as missing since 2009:Q1



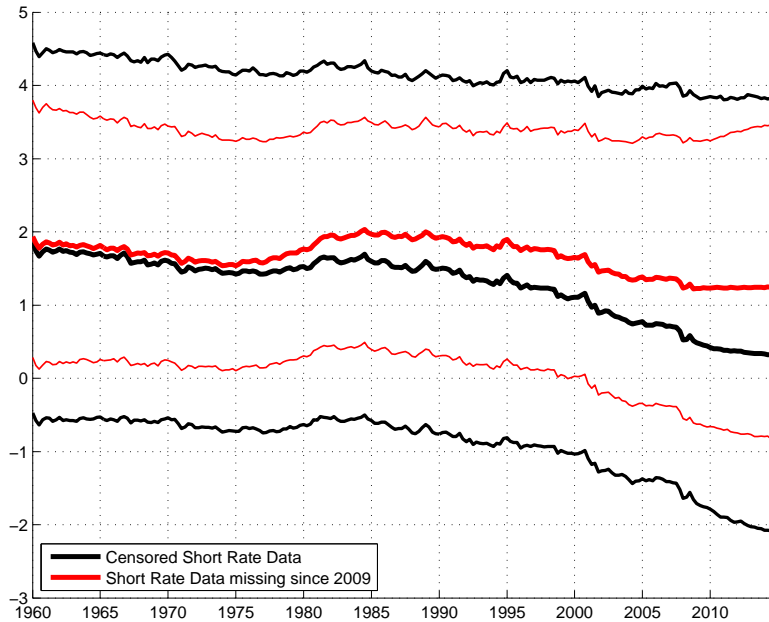
Note: The black-yellow dashed line is the median of the posterior of the distribution. The darkly-shaded area is the 50 percent confidence set. The lightly-shaded area is the 90 percent confidence set. All estimates reflect full-sample information from 1960:Q1 through 2014:Q4.

Figure 2: Real interest rate trend from the bivariate model.

(a) Nominal rate data treated as censored

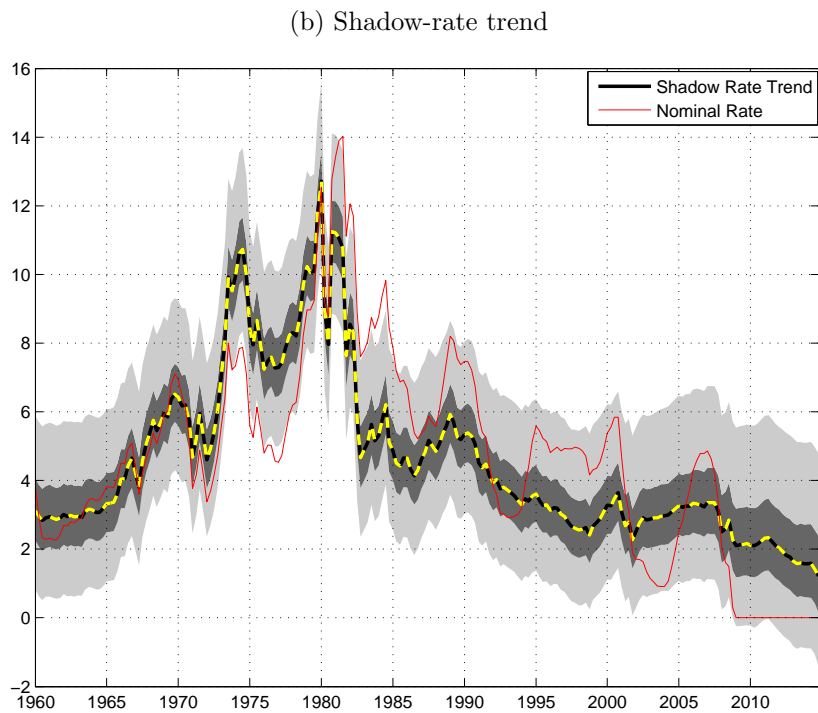
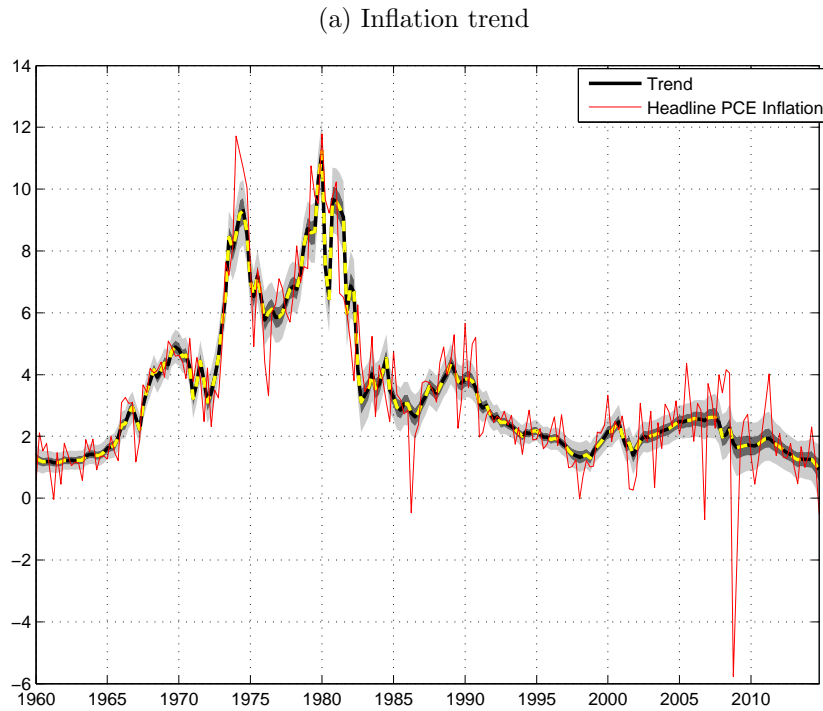


(b) Comparison with data treated as missing



Note: Top Panel: The black-yellow dashed line is the median of the posterior of the distribution. The darkly-shaded area is the 50 percent confidence set. The lightly-shaded area is the 90 percent confidence set. All estimates reflect full-sample information from 1960:Q1 through 2014:Q4. Bottom Panel: The bold lines are the median of the posterior distribution, the thin lines represent 90 percent confidence sets.

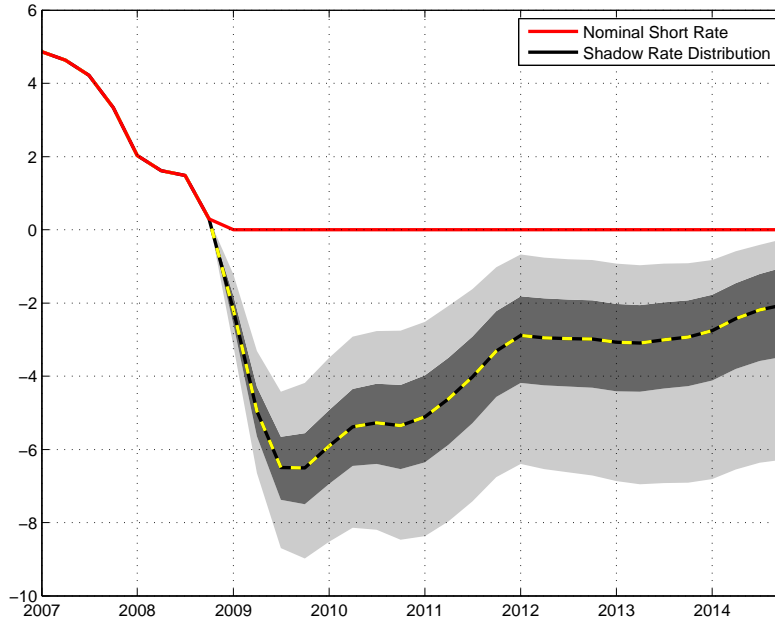
Figure 3: Inflation and shadow-rate trend from the bivariate model.



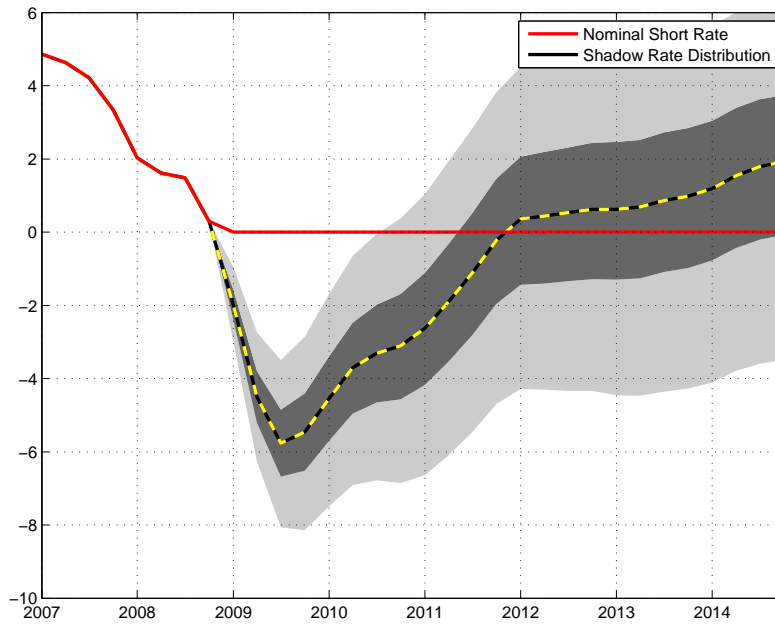
Note: The black-yellow dashed line is the median of the posterior of the distribution. The darkly-shaded area is the 50 percent confidence set. The lightly-shaded area is the 90 percent confidence set. All estimates reflect full-sample information from 1960:Q1 through 2014:Q4.

Figure 4: Shadow rates estimated from the macro model

(a) Nominal-rate data treated as censored



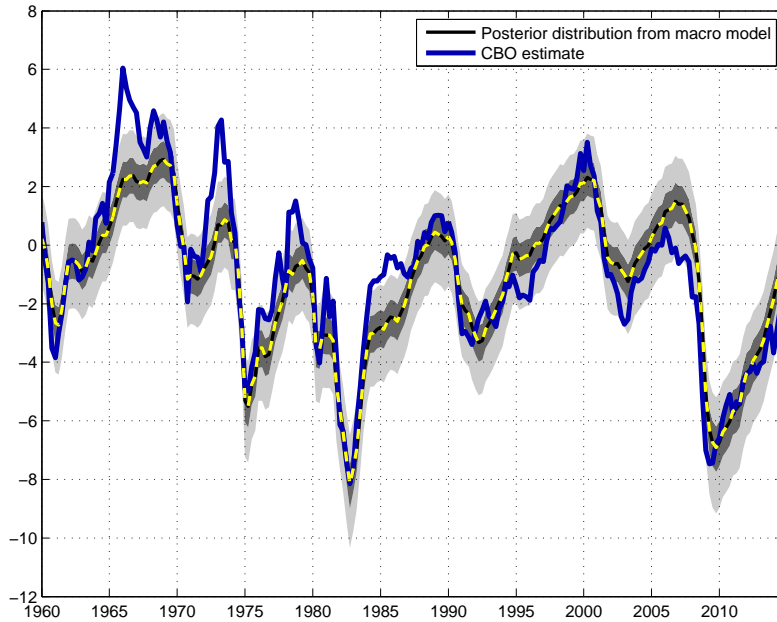
(b) Nominal-rate data treated as missing since 2009:Q1



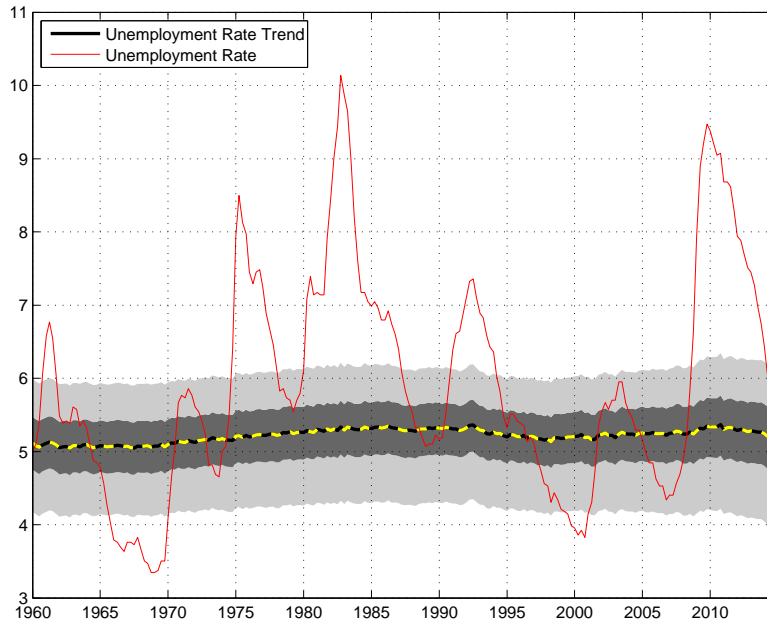
Note: The black-yellow dashed line is the median of the posterior of the distribution. The darkly-shaded area is the 50 percent confidence set. The lightly-shaded area is the 90 percent confidence set. All estimates reflect full-sample information from 1960:Q1 through 2014:Q4.

Figure 5: Output gap and unemployment

(a) Output gap

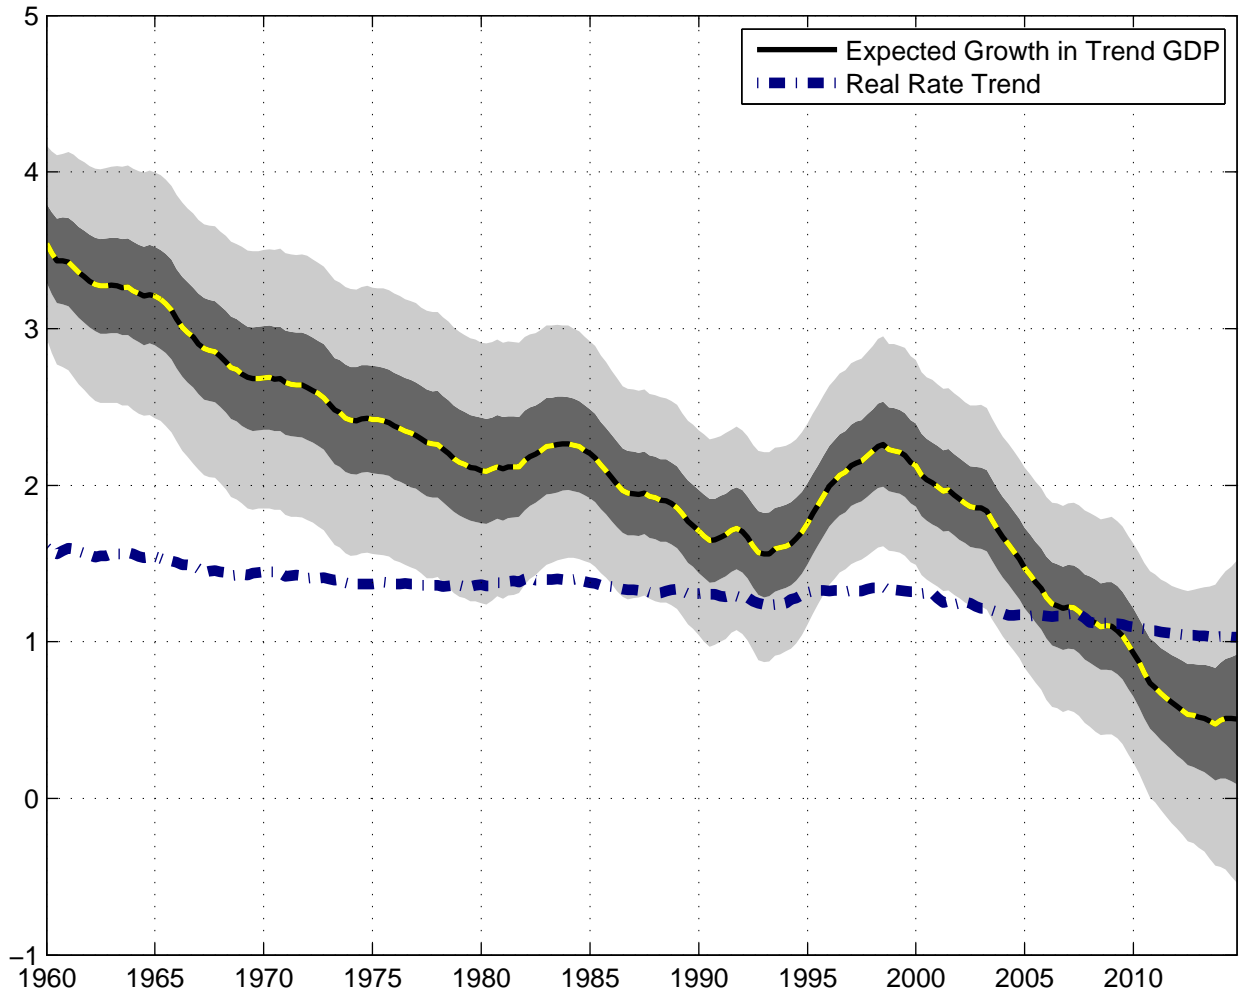


(b) Unemployment rate



Note: Estimates generated from a sampler that treats the nominal rate data as censored at the effective lower bound. The black-yellow dashed line is the median of the posterior of the distribution. The darkly-shaded area is the 50 percent confidence set. The lightly-shaded area is the 90 percent confidence set. All estimates reflect full-sample information from 1960:Q1 through 2014:Q4.

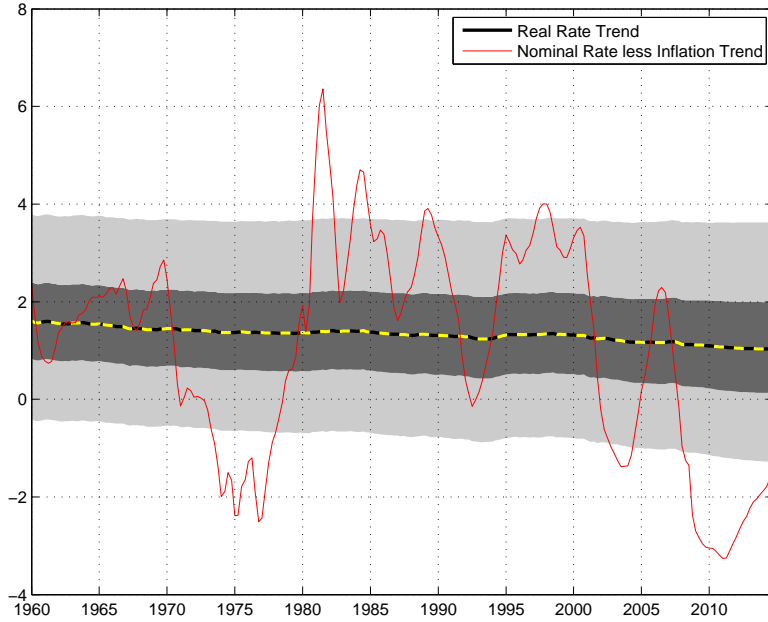
Figure 6: Expected growth in trend output and the real-rate trend



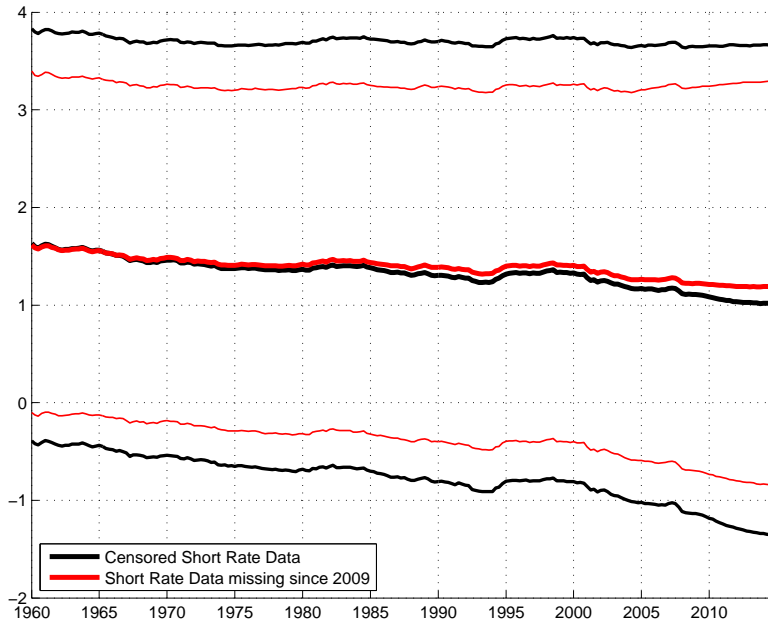
Note: Quantiles of the posterior distribution of the expected growth rate of trend output, denoted μ_t in (12), and posterior median for the real-rate trend \bar{r}_t . The black-yellow dashed line is the median of the posterior of the distribution. The darkly-shaded area is the 50 percent confidence set. The lightly-shaded area is the 90 percent confidence set. All estimates reflect full-sample information from 1960:Q1 through 2014:Q4. Estimates generated from a sampler that treats the nominal rate data as censored at the effective lower bound.

Figure 7: Real-rate trend in the macro model

(a) Nominal-rate data treated as censored



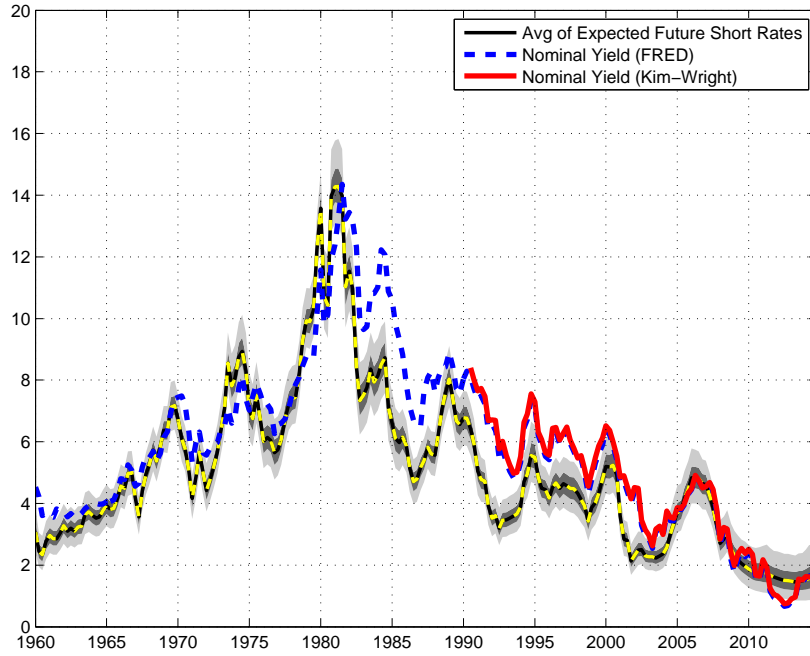
(b) Comparison with data treated as missing



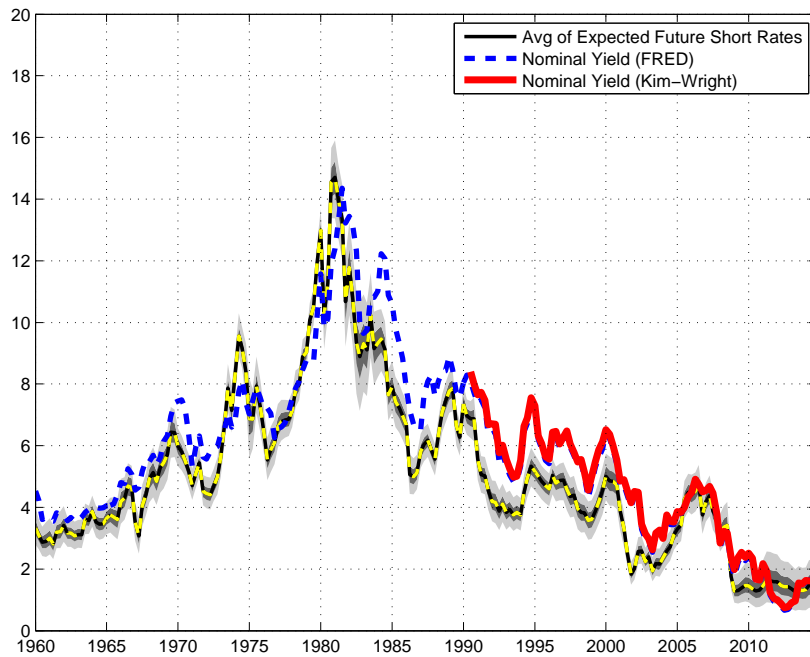
Note: Top Panel: The black-yellow dashed line is the median of the posterior of the distribution. The darkly-shaded area is the 50 percent confidence set. The lightly-shaded area is the 90 percent confidence set. All estimates reflect full-sample information from 1960:Q1 through 2014:Q4. Bottom Panel: The bold lines are the median of the posterior distribution, the thin lines represent 90 percent confidence sets.

Figure 8: Expectation component of 5-year interest rate.

(a) Bivariate model



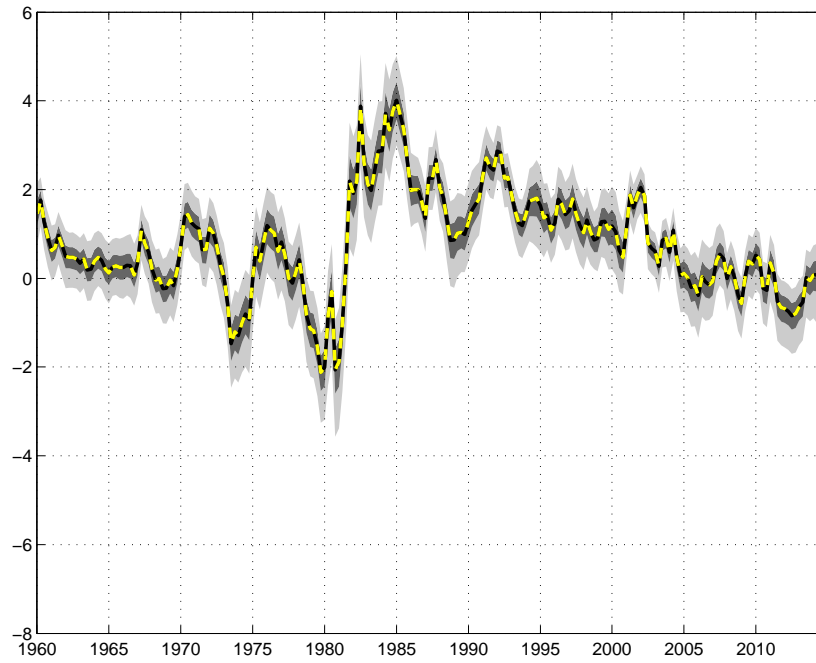
(b) Macro model



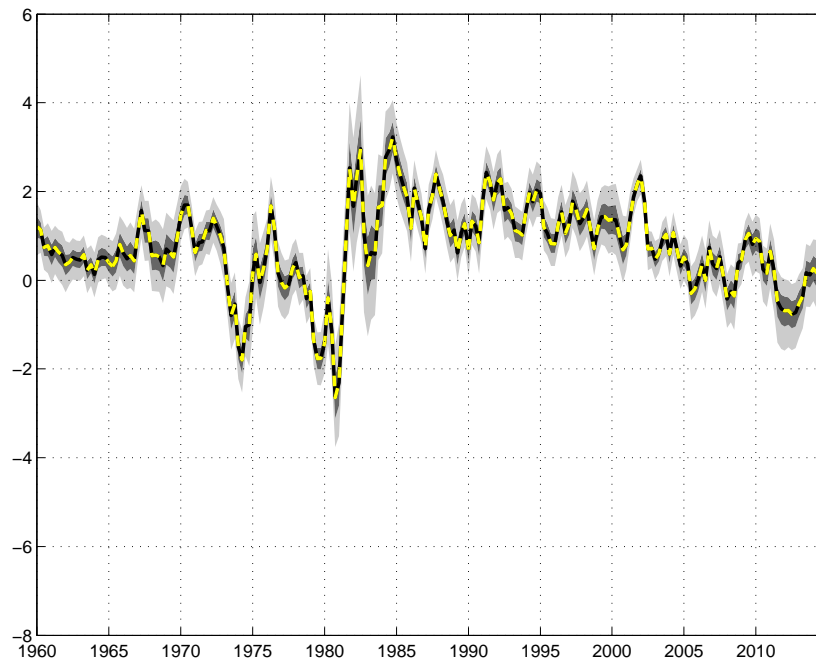
Note: The black-yellow dashed line is the median of the posterior of the distribution. The darkly-shaded area is the 50 percent confidence set. The lightly-shaded area is the 90 percent confidence set. All estimates reflect full-sample information from 1960:Q1 through 2014:Q4.

Figure 9: Premium component of 5-year interest rate.

(a) Bivariate model



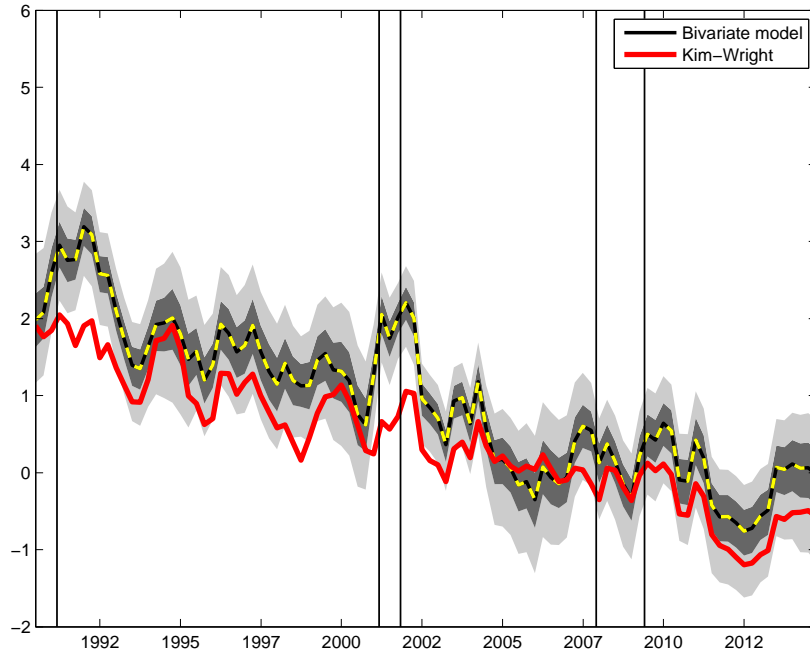
(b) Macro model



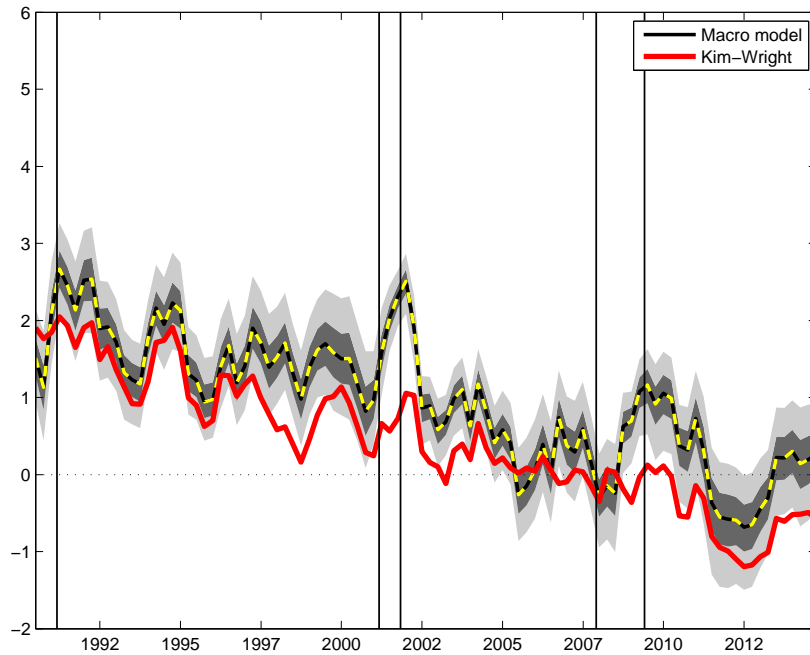
Note: The black-yellow dashed line is the median of the posterior of the distribution. The darkly-shaded area is the 50 percent confidence set. The lightly-shaded area is the 90 percent confidence set. All estimates reflect full-sample information from 1960:Q1 through 2014:Q4.

Figure 10: Premium component of 5-year interest rate.

(a) Bivariate model



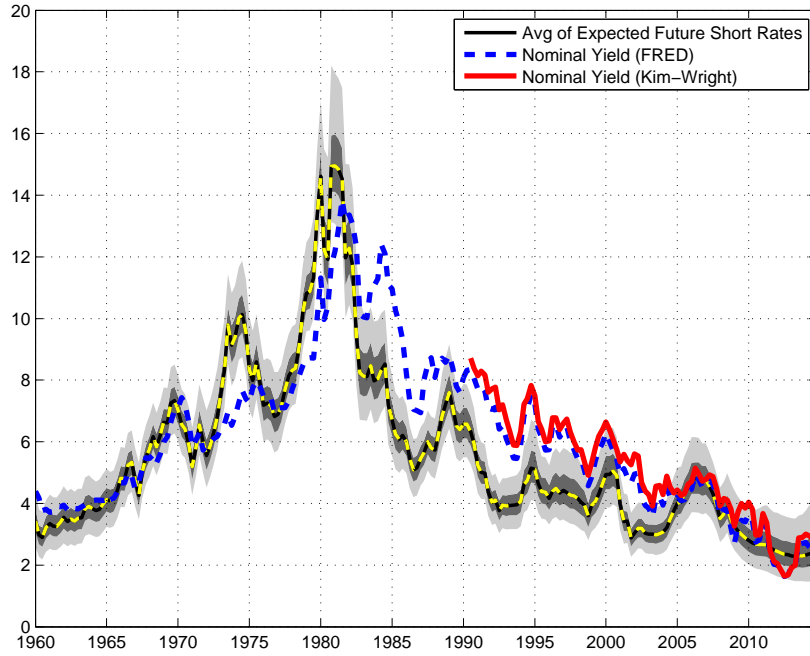
(b) Macro model



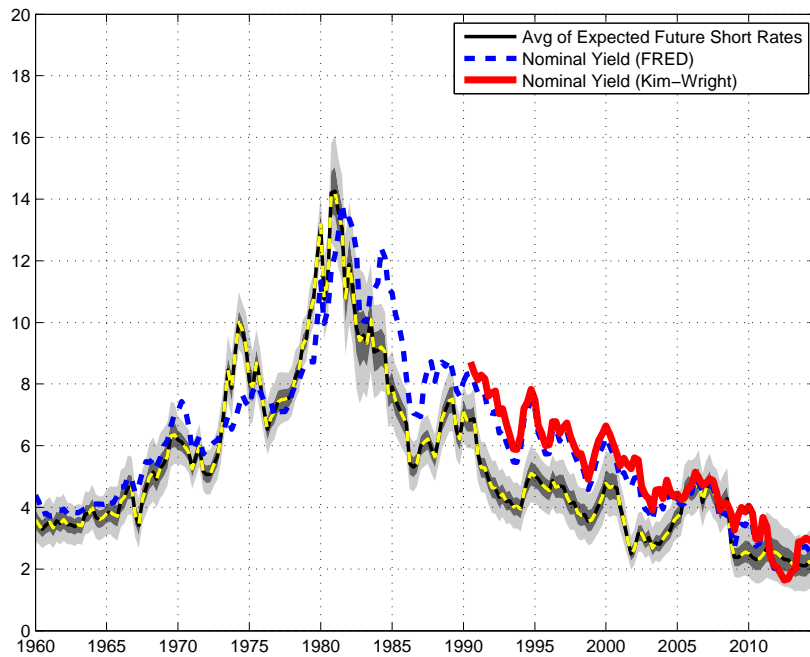
Note: The black-yellow dashed line is the median of the posterior of the distribution. The darkly-shaded area is the 50 percent confidence set. The lightly-shaded area is the 90 percent confidence set. All estimates reflect full-sample information from 1960:Q1 through 2014:Q4.

Figure 11: Expectation component of 10-year interest rate.

(a) Bivariate model



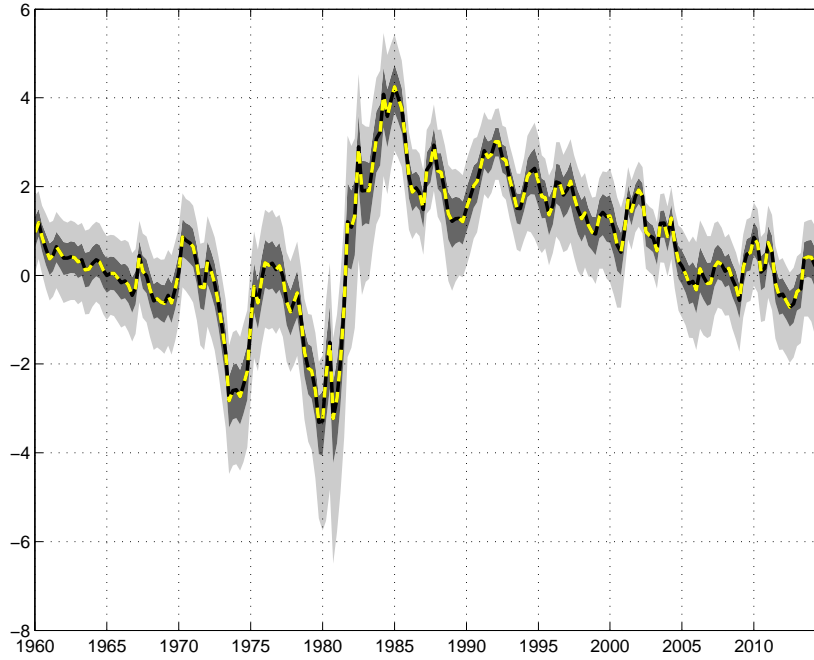
(b) Macro model



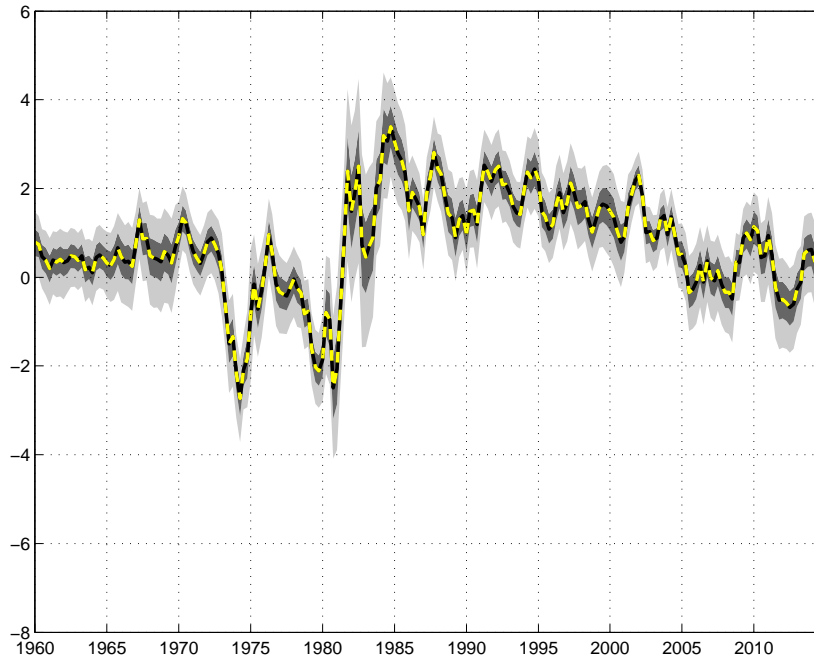
Note: The black-yellow dashed line is the median of the posterior of the distribution. The darkly-shaded area is the 50 percent confidence set. The lightly-shaded area is the 90 percent confidence set. All estimates reflect full-sample information from 1960:Q1 through 2014:Q4.

Figure 12: Premium component of 10-year interest rate.

(a) Bivariate model



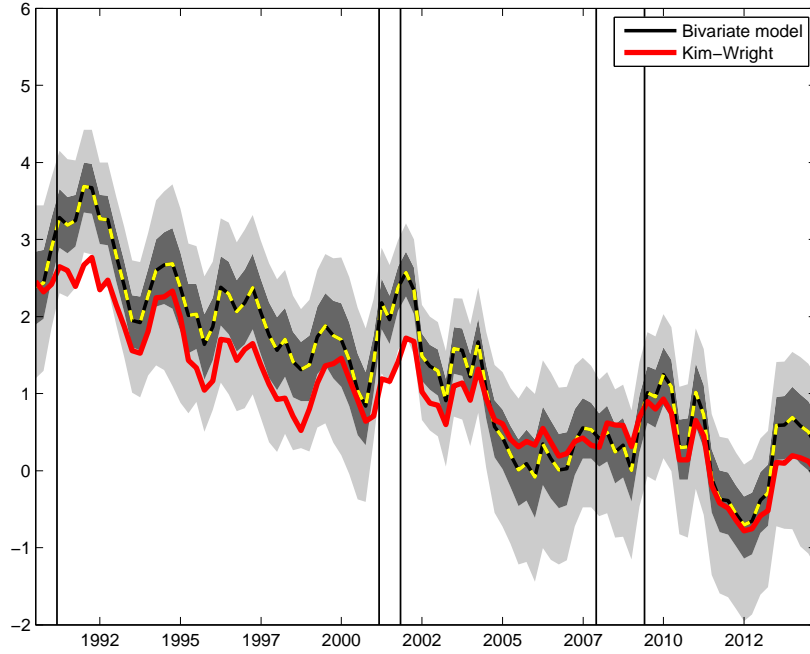
(b) Macro model



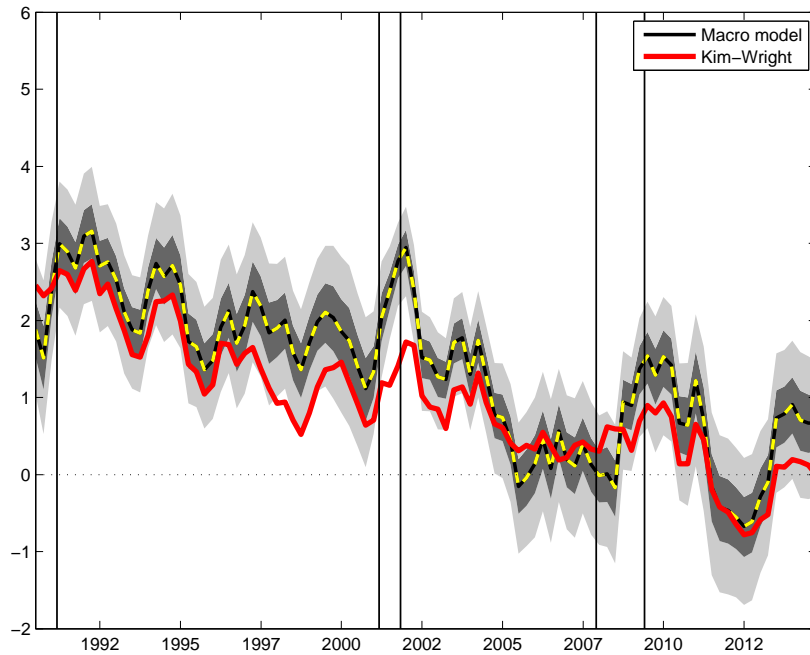
Note: The black-yellow dashed line is the median of the posterior of the distribution. The darkly-shaded area is the 50 percent confidence set. The lightly-shaded area is the 90 percent confidence set. All estimates reflect full-sample information from 1960:Q1 through 2014:Q4.

Figure 13: Premium component of 10-year interest rate.

(a) Bivariate model

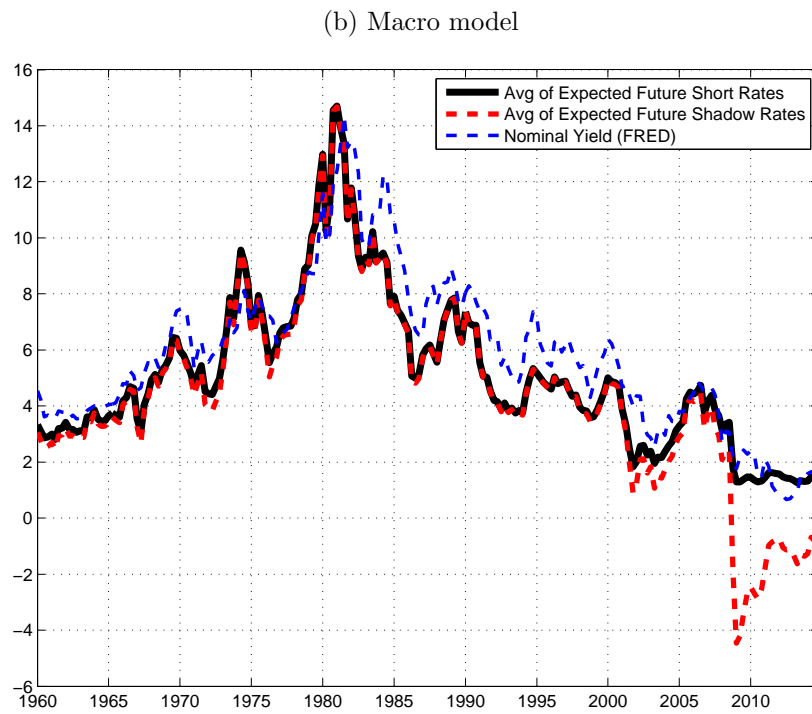
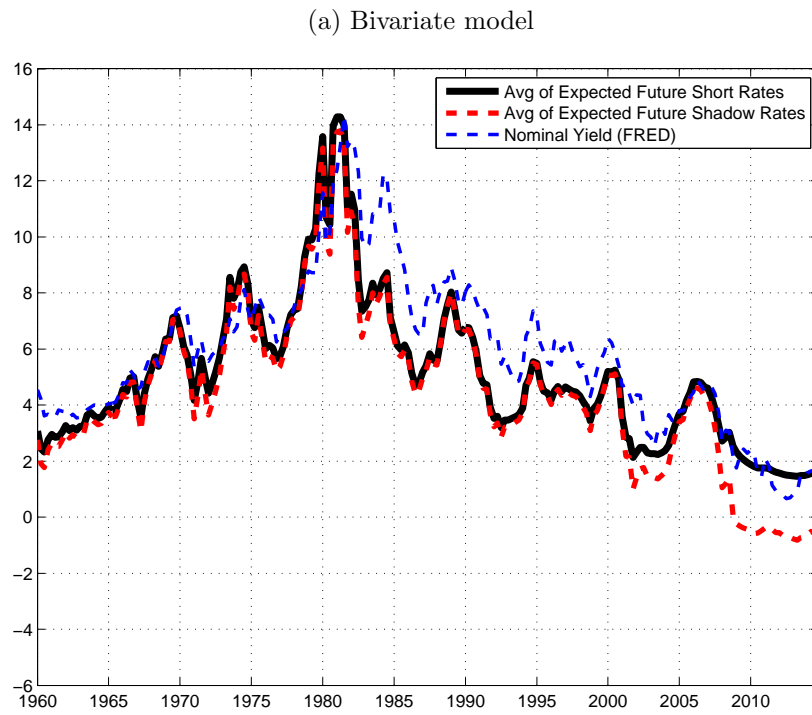


(b) Macro model



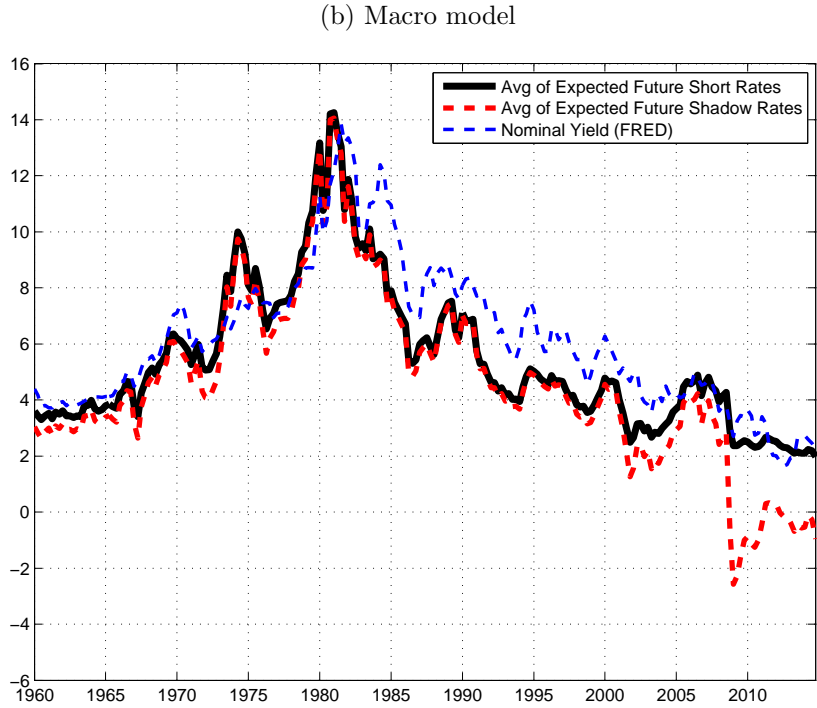
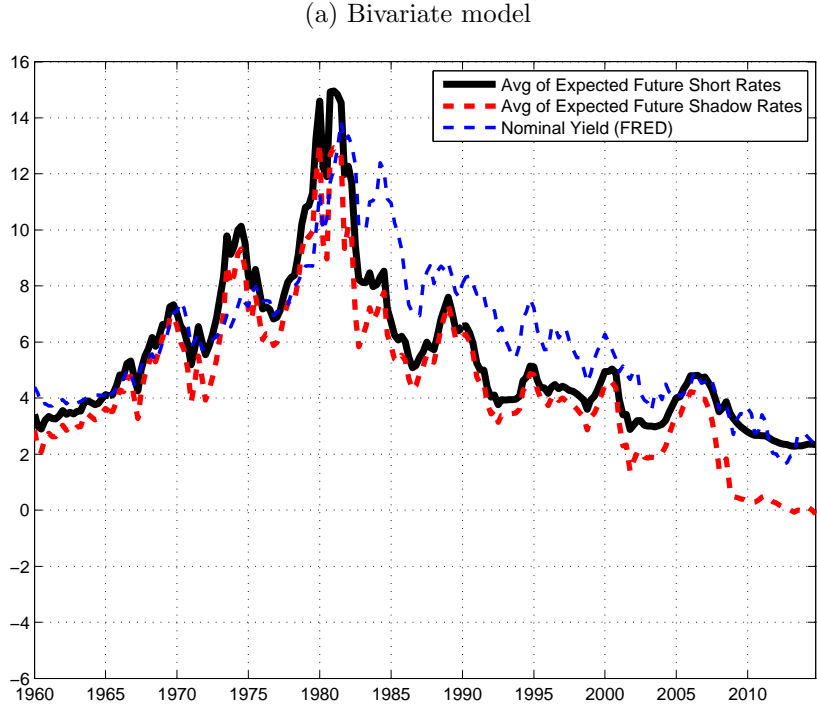
Note: The black-yellow dashed line is the median of the posterior of the distribution. The darkly-shaded area is the 50 percent confidence set. The lightly-shaded area is the 90 percent confidence set. All estimates reflect full-sample information from 1960:Q1 through 2014:Q4.

Figure 14: The effective lower bound and the expectation component of the 5-year rate.



Note: Medians of the posterior distributions which reflect full-sample information from 1960:Q1 through 2014:Q4.

Figure 15: The effective lower bound and the expectation component of the 10-year rate.



Note: Medians of the posterior distributions which reflect full-sample information from 1960:Q1 through 2014:Q4.