A Time Series Model of Interest Rates With the Effective Lower Bound*

Benjamin K. Johannsen†
Federal Reserve Board

Elmar Mertens
Bank of International Settlements

December 31, 2017

Abstract

Modeling interest rates over samples that include the Great Recession requires taking stock of the effective lower bound (ELB) on nominal interest rates. We propose a flexible time–series approach which includes a "shadow rate"—a notional rate that is less than the ELB during the period in which the bound is binding—without imposing no–arbitrage assumptions. The shadow rate serves as a latent state variable to characterize the joint dynamics of yields and macro variables. The approach allows us to estimate the behavior of trend real rates as well as expected future interest rates in recent years.

*An earlier draft of this paper has been circulated under the title "The Shadow Rate of Interest, Macroeconomic Trends, and Time-Varying Uncertainty."

†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. We would like to thank seminar and conference participants at the 2017 joint Bank of Canada, Federal Reserve of San Francisco, and Beedie School of Business conference on Advances in Fixed Income and Macro-Finance Research and the 6th Joint BoC-ECB Conference — in particular our discussants Greg Duffee and Feunou Kamkui — the Federal Reserve Banks of Cleveland, and Richmond, the Federal Reserve Board, the Bank for International Settlements, the University of Texas at Austin, the World Congress of the Econometric Society (2015), and the Dynare Conference (2015), for their useful comments and suggestions. The views in this paper do not necessarily represent the views of the Bank for International Settlements, the Federal Reserve Board, 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 models nominal interest rates, along with other macroeconomic data, using a flexible time-series model that explicitly incorporates the effective lower bound (ELB) on nominal interest rates. We employ a modeling device that we refer to as a "shadow rate"—the nominal interest rate that would prevail in the absence of the ELB—which is conceptually similar to the shadow rates studied in the dynamic term–structure literature, as in Kim and Singleton (2011), Krippner (2013), Priebsch (2013), Ichiue and Ueno (2013), Bauer and Rudebusch (2015), Krippner (2015), and Wu and Xia (2016). Our time–series approach allows us to estimate the relationship between interest rates and macroeconomic data in a flexible way and, similar to the approach taken in Diebold and Li (2006), does not impose rigid no–arbitrage restrictions across the term–structure of interest rates.

We use our approach to estimate a trend–cycle model of U.S. data on interest rates, unemployment, and inflation over a sample that includes the recent spell at the ELB. Since the global financial crisis of 2008, real interest rates have been historically low, prompting some — for example, Summers (2014) and Rachel and Smith (2015) — to argue that the long–run normal level of real interest has fallen. Using our estimated model, we find that the trend component of the nominal interest rate has declined 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. While uncertainty bands around our estimate of the trend real rate are wide, we find that any decline since the global financial crisis of 2008 is best characterized as a continuation of a longer term downward trajectory. Similar to Laubach and Williams (2003, 2015) and many others, we find large uncertainty bands notwithstanding some differences in point estimates.[1] However, none of these papers explicitly models the ELB.

Explicitly modeling the ELB has large effects on inference about out–of–sample expected short–term interest rates and term premiums over the past several years. Our estimated shadow rates are less than the ELB by construction, and our shadow–rate model delivers predicted

[1] See, for example, Clark and Kozicki (2005), Hamilton et al. (2015), Kiley (2015), Lubik and Matthes (2015), Hakkio and Smith (2017), Pescatori and Turunen (2015), and Del Negro et al. (2017). Notably, our estimated trend real rate displays less movement than the trend estimates reported by Laubach and Williams (2003, 2015), and does not not dip as low during the recent recession.

paths for future short–term interest rates that include extended periods at the ELB. We compare interest–rate forecasts from our model with forecasts from the Survey of Professional Forecasters (SPF) and find that our model performs better than the SPF at longer horizons. In addition, we compare interest–rate forecasts from our model to forecasts from the shadow-rate term-structure model of Wu and Xia (2016) (WX-SRTSM), and find that our model performs at least as well.

When the economy is away from the ELB, the shadow rate and the policy rate are identical. At the ELB, the shadow rate is an unobserved state variable that matters for forecasting future outcomes in the policy rate and other variables. Unexpected variations in the shadow rate can thus be interpreted as reflecting changes in monetary policy implemented through unconventional tools (such as asset purchases or forward guidance). Using short-run identification restrictions common in the literature, we identify monetary policy shocks from surprises in the shadow rate and report impulse responses of the yield curve and macro variables. In contrast to other related approaches (Wu and Xia, 2016), our impulse responses are estimated jointly with the shadow rate. Our model reveals interesting time variation in those impulse response functions around the recent ELB episode.

The way we incorporate the ELB 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 ELBs in many parts of the world, time series models that include interest rates but ignore the ELB—like a standard vector autoregression—have been unable to adequately address the data. Moreover, reduced-form explorations of the empirical 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 in a 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), the no–arbitrage dynamic term–structure models studied in Kim and Singleton (2011), Krippner (2013), Priebsch (2013), Ichiue and Ueno (2013), Bauer and Rudebusch (2015), Krippner (2015), and Wu and Xia (2016) identify shadow rates

by imposing no–arbitrage cross–equation restrictions. These studies offer interesting insights, yet the no–arbitrage assumptions that these authors maintain may preclude certain model features, like stochastic model parameters. Our time series approach naturally incorporates time variation in parameters, and thus, for some purposes—like including time–varying trends in inflation and interest rate data or modeling stochastic volatility—offers a flexible alternative. In addition, our shadow–rate estimates do not only reflect information embedded in longer–term yield data but also condition on direct readings about business cycle conditions embedded in macro variables such as inflation or the output– and unemployment–rate gaps (as measured by CBO).[2]

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 the ELB. 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. In doing so we are able to more closely align our approach with the no–arbitrage term–structure literature, and, in additional, connect the concept of the shadow rate with the level of the short–term rate that would prevail in the absence of the ELB because we allow it to have the same persistence and co–variance properties as short–term interest rates. Nevertheless, our approach is flexible enough to include both shadow rates and observed rates as lagged explanatory variables.

2 A Model of Interest Rates and the ELB

In this section we describe our time–series model, which explicitly includes the ELB. The model includes inflation (π_t), nominal interest rates of maturity 3–months (i_t), 2–years (y_t^2), 5–years (y_t^5), and 10–years (y_t^{10}), and a cyclical measure of real activity ($\tilde{c}_t$) (henceforth referred

[2]In an alternative time-series approach Lombardi and Zhu (2014) use a dynamic factor model to derive estimates of the stance of monetary policy — labeled "shadow rate" — from interest rates, monetary aggregates and variables characterizing the Federal Reserve's balance sheet. However, as such, their underlying shadow–rate concept is quite different from what is used here or in the dynamic term-structure literature in that their measure needs not be identical to observed interest rates, even when the ELB is not binding, nor is their shadow rate constrained to lie below the ELB when the bound is binding.

to as "business cycle measure"). We use CBO output gap estimates as business cycle measure; similar results are obtained based on CBO unemployment-rate gap estimates.

2.1 The Shadow Rate Approach

Our data set includes a short–term interest rate, which is constrained by the ELB. We model the short–term interest rate as the observation of a censored variable. In particular, we assume that the nominal interest rate is the maximum of the ELB and a shadow rate (s_t) so that

$$i_t = \max\left(s_t, ELB\right) \ . \tag{1}$$

The ELB might arise because of an arbitrage between bonds and cash, though the world has seen negative short–term nominal interest rates in a number of countries. It also might be thought of as a level below which monetary authorities are unwilling to push short–term interest rates. For our purposes, it is taken as an exogenous known constant (which could be made time–varying). In our empirical application, the ELB on nominal rates of all maturities is assumed to bind at 25 basis points. We proceed by modeling the shadow rate, in conjunction with the other variables in the model, using standard time–series methods, and account for the ELB when conditioning the posterior distribution of our model on observed interest rate data.

In principle, the distinction between shadow rates and observed nominal rates could also be extended here to other nominal yields; however, in our application the distinction would be moot since the ELB never binds for the other yields in our data.

2.2 A Time Series Model with Shadow Rates

We assume that all our variables, except for the business cycle measure, can be decomposed into trend and gap components. That is, for any data series x_t we can write

$$x_t = \bar{x}_t + \tilde{x}_t \qquad \text{where} \quad \bar{x}_t = \lim_{h\to\infty} E_t\left(x_{t+h}\right) \qquad \text{and} \quad E(\tilde{x}_t) = 0 \tag{2}$$

Measured as deviation from an infinite-horizon forecast, the defining feature of the gap component, $\tilde{x}_t$, is that it has a zero ergodic mean. Considering our business cycle measures, we assume that $\tilde{c}_t$ satisfies the definition of a gap component, as expressed by $\tilde{x}_t$ in (2).[3]

The trend components $\bar{x}_t$ are similar in spirit to the trend concept of Beveridge and Nelson (1981); however, by treating the trends as unobserved components we allow for the conditional expectations, $E_t(\cdot)$ in (2), to reflect a possibly wider information set than what is known to an econometrician at time t.[4] Defining the trend components as infinite-horizon expectations implies that changes in $\bar{x}_t$ follow martingale-difference processes; and, as a result, the trend components have unit root dynamics.[5]

For example, U.S. inflation dynamics are well captured by such a trend-cycle decomposition when trend shocks have time-varying volatility; see Stock and Watson (2007), Cogley and Sargent (2015) or Mertens (2016). So, for the trend component of inflation, we write:

$$\bar{\pi}_t = \bar{\pi}_{t-1} + \sigma_{\bar{\pi},t}\epsilon_{\bar{\pi},t}. \tag{3}$$

$$\log\left(\sigma^2_{\bar{\pi},t}\right) = (1-\rho_{\bar{\pi}})\mu_{\bar{\pi}} + \rho_{\bar{\pi}}\log\left(\sigma^2_{\bar{\pi},t-1}\right) + \phi_{\bar{\pi}}\eta_{\bar{\pi},t} \tag{4}$$

where $\epsilon_{\bar{\pi},t} \sim N(0,1)$ and $\eta_{\bar{\pi},t} \sim N(0,1)$ and $|\rho_{\bar{\pi}}| < 1$. Assuming a stationary AR(1) specification for stochastic volatility (rather than the also commonly used random walk specification) helps us obtain better behaved predictive densities for multi-step forecasting.

We assume that the trend shadow rate has two components

$$\bar{s}_t = \bar{\pi}_t + \bar{r}_t \tag{5}$$

where $\bar{r}_t$ will be our measure of the trend in real interest rates. We assume that $\bar{r}_t$ evolves so

[3]For the unemployment rate (u_t) we thus assume that the CBO's measure of the long-run natural rate reflects a trend estimate series akin to the $\bar{x}_t$ component in (2). Similarly, for the output gap, we assume that the CBO's measure of potential real GDP reflects a trended series akin to the $\bar{x}_t$ component in (2).

[4]See also the discussion in Mertens (2016).

[5]Similar trend-cycle decompositions have also been used in a variety of structural models, see, for example, Rudebusch and Svensson (1999), Kozicki and Tinsley (2002, 2001, 2012) Ireland (2007), or Cogley et al. (2010).

that

$$\bar{r}_t = \bar{r}_{t-1} + \sigma_{\bar{r}} \epsilon_{\bar{r},t}. \tag{6}$$

To capture a connection between the short–term interest rate and the other interest rates, we assume that the other interest rates in our model (y_t^2, y_t^5, y_t^{10}) share a common trend with the shadow rate, adjusted for average term–premiums.[6] By assuming that yield spreads are stationary, we impose the same cointegrating relationship on nominal yields that has also been used by Campbell and Shiller (1987) and King and Kurmann (2002) or more recently by Bauer and Rudebusch (2017a) and Benati (2017). In sum, there are two stochastic trends in our model: $\bar{\pi}_t$ and $\bar{r}_t$.[7]

While we model shocks to trend inflation with stochastic volatility, our baseline specification assumes homoscedastic shocks to the trend real rate. Prior evidence suggests that time-varying volatility in shocks to trend inflation serves well to capture changes in the anchoring of public perceptions of long–term inflation expectations and the credibility of policymaker's inflation goals in U.S. postwar data (Stock and Watson, 2007). By contrast, the variability of the trend real rate is more likely to reflect slow–moving changes in long–term growth expectations, demographic trends and other secular drivers (Rachel and Smith, 2015; Gagnon et al., 2016). Furthermore, as discussed by Hamilton et al. (2015), trend variability probably accounts for only a small share of the variability in real rates, which cautions us against fitting a stochastic volatility process for changes in this trend. As robustness check, we also estimate an alternative version of our model, where shocks to the real-rate trend are also subject to stochastic volatility. As discussed in Section 3.2, this model variant fits the data worse — as measured by the marginal data density — than our baseline specification.

We assume that the gap components of the series in our model follow a joint autoregressive

[6] Specifically, the trend in the nominal two-year yield is written as $\bar{y}_t^2 = \bar{s}_t + \bar{p}_0^2$ where the constant $\bar{p}_0^2$ represents the average term–premium, and its estimated value reflects the average spread between the short–term nominal interest rate and the two-year yield in our model over our sample. As for all gap variables, the mean of the yield gap, $E(\tilde{y}_t^2)$, has been normalized to zero. Similar relationships hold for y_t^5 and y_t^{10}.

[7] Shocks to both trends are assumed to be mutually uncorrelated.

process. That is, we assume that

$$\boldsymbol{A}(L)\tilde{\boldsymbol{X}}_t = \boldsymbol{B}\tilde{\boldsymbol{\Sigma}}_t^{1/2}\tilde{\boldsymbol{\varepsilon}}_t, \quad \tilde{\boldsymbol{X}}_t \equiv \boldsymbol{X}_t - E_t\boldsymbol{X}_{t+\infty}, \quad \boldsymbol{X}_t \equiv \begin{bmatrix} \pi_t & \tilde{c}_t & s_t & y_t^2 & y_t^5 & y_t^{10} \end{bmatrix}' \tag{7}$$

where $\boldsymbol{A}(L)$ is a polynomial in the lag operator that has all roots outside the unit circle, $\boldsymbol{B}$ is a unit–lower–triangular matrix, $\tilde{\boldsymbol{\varepsilon}}_t$ is multi–variate standard normal, and $\tilde{\boldsymbol{\Sigma}}_t^{1/2}$ is a diagonal matrix of stochastic volatilities; the vector $\boldsymbol{\sigma}_t^2$ collects the diagonal elements of $\tilde{\boldsymbol{\Sigma}}_t = \tilde{\boldsymbol{\Sigma}}_t^{1/2}\left(\tilde{\boldsymbol{\Sigma}}_t^{1/2}\right)'$.[8] In order to give the model the flexibility to capture large changes in the gap components over the business cycle, shocks to the gap–vectorautoregression are assumed to follow stationary stochastic-volatility processes with mutually correlated shocks:

$$\log\left(\tilde{\boldsymbol{\sigma}}_t^2\right) = \left(\boldsymbol{I} - \tilde{\boldsymbol{\rho}}\right)\tilde{\boldsymbol{\mu}} + \tilde{\boldsymbol{\rho}}\log\left(\tilde{\boldsymbol{\sigma}}_{t-1}^2\right) + \tilde{\boldsymbol{\Phi}}\tilde{\boldsymbol{\eta}}_t \qquad \tilde{\boldsymbol{\eta}}_t \sim N(\boldsymbol{0}, \boldsymbol{I}) \tag{8}$$

where $\tilde{\boldsymbol{\mu}}$ is a vector of mean log-variances, $\tilde{\boldsymbol{\rho}}$ is a diagonal matrix of lag coefficients that are all inside the unit circle, and $\tilde{\boldsymbol{\Phi}}$ is the variance-covariance matrix of shocks to the SV processes. Shocks to the stochastic volatilities in the gap shocks can be correlated, which allows the model to pick up on commonalities in time-variation of business-cycle volatility across variables (Jurado et al., 2015; Clark et al., 2016).

2.3 Relationship Between Shadow and Interest Rates

We conceptualize the shadow rate as the nominal interest rate that would prevail in the absence of the ELB. On a period–by–period basis, the interest rate is either equal to the shadow rate or equal to the ELB. The key distinction between shadow rates and interest rates is thus that shadow rates have unbounded support.

In the model we presented in the previous section, we modeled the shadow–rate gap, as well as its lags, as part of a joint dynamic system, which allows the shadow rate to have the same persistence properties when the ELB is binding and when it is not. By contrast, Iwata and Wu (2006) and Nakajima (2011) model the variables in their models as functions of

[8] Due to the stochastic volatilities, estimation is not invariant to the ordering of the variables Negro and Primiceri (2015). Appendix D documents the robustness of estimates to alternative orderings.

lagged observed interest rates. This means that, in those papers, at the ELB the value of the shadow–rate in the previous period has no direct effect on its value today. This approach is in stark contrast to the shadow–rate dynamics from dynamic term–structure literature; see, for example, Wu and Xia (2016).[9]

Similar to the term-structure literature, we embed the shadow rate into a state vector with auto–regressive dynamics, such that the persistence of the shadow rate does not depend on whether the ELB binds. When the ELB is binding on observed interest rates, the shadow rate is intended to capture the hypothetical level of the nominal rate that would prevail in the absence of the ELB constraint; accordingly, we deem it beneficial that the estimated persistence of the shadow rate in our specification reflects to a large degree the persistence of observed interest rates when those are away from the ELB.

2.4 Interpretation of $\bar{r}_t$

Because interest rates are truncated shadow rates, the expected interest rate is necessarily weakly larger than the expected shadow rate. In turn, it is also the case that,

$$\lim_{h\to\infty} E_t\left(i_{t+h}\right) \geq \lim_{h\to\infty} E_t\left(s_{t+h}\right) = \bar{s}_t = \bar{\pi}_t + \bar{r}_t. \tag{9}$$

In our model, $\bar{s}_t$ is the median forecast of $\lim_{h\to\infty} i_{t+h}$, offering a direct connection between far-ahead shadow rates and interest rates.[10] Further, assuming that the Fisher hypothesis holds, this connection gives $\bar{r}_t$ the interpretation of the median forecast of the real interest rate in the long run.[11] Notably, the same relationship holds for the other yields in our model, up to a constant offset, because of the co-integrating relationship we have assumed. For the remainder of the paper, we refer to $\bar{r}_t$ as the trend real interest rate.

[9] Our statistical approach is able to accommodate models with both lagged observed interest rates and lagged shadow rates. See the appendix for further discussion.

[10] The value of $\bar{s}_t$ is not a mean forecast because, in our model, at long horizons, there is positive probability that the shadow rate will be less than zero, meaning $\lim_{h\to\infty} E_t\left(i_{t+h}\right) > \lim_{h\to\infty} E_t\left(s_{t+h}\right)$; however, one could conceptualize models in which the inequality needs not to be strict.

[11] So long as $\bar{s}_t \geq 0$, which it is in our estimates, then our interpretation of $\bar{r}_t$ applies. However, if $\bar{s}_t < 0$, then the median forecast of the $\lim_{h\to\infty} i_{t+h} - \pi_{t+h}$ is $ELB - \bar{\pi}_t$. Then $\bar{r}_t$ would help determine the probability that interest rates would be above the ELB over the business cycle and would be less than the long–run median real interest rate.

Alternatively, $\bar{r}_t$ could be interpreted as the point of attraction for the short–term real interest rate, assuming no future shocks. In this way, $\bar{r}_t$ is the expected short–term real interest rate in the longer–run after all shocks have died out. As in the previous interpretation, $\bar{r}_t$ can be though of a the trend real interest rate.

Importantly, the co-integrating relationship between the yields in our model allows the longer–term yields to offer direct evidence on $\bar{r}_t$. This disciplines movements in $\bar{r}_t$ because these yields are above the ELB throughout our sample.

2.5 Estimation Procedure

To estimate the parameters and unobserved states of the model, we use Bayesian methods.[12] The novel modeling contribution of this paper lies in the sampling of the unobserved trend and gap components of the data when the interest rate data are at the ELB, so we focus the text on this step of the estimation procedure. Conditional on parameter values and a sequence of volatilities, our model can be put into the form

$$\boldsymbol{\xi}_t = \boldsymbol{\mathcal{A}}\boldsymbol{\xi}_{t-1} + \boldsymbol{\mathcal{B}}\,\boldsymbol{\Sigma}_t^{1/2}\,\boldsymbol{\varepsilon}_t \qquad \boldsymbol{\varepsilon}_t \sim N(\mathbf{0}, \boldsymbol{I}) \tag{10}$$

$$\boldsymbol{X}_t = \boldsymbol{\mathcal{C}}\boldsymbol{\xi}_t \tag{11}$$

$$i_t = \max\left(s_t, ELB\right) \qquad s_t = \boldsymbol{c}_s\,\boldsymbol{X}_t \tag{12}$$

where $\boldsymbol{\xi}_t$ contains the stochastic trend and gap components of our model, as well as the appropriate number of lags to represent their dynamics in companion form, and $\boldsymbol{\Sigma}^{1/2}$ is a diagonal matrix of stochastic volatilities (as well as the constant volatility of real–rate trend shocks in the baseline model). The matrices $\boldsymbol{\mathcal{A}}$, $\boldsymbol{\mathcal{B}}$, and $\boldsymbol{\mathcal{C}}$ are constructed accordingly from the parameters in our model, [13] $\boldsymbol{c}_s$ is a selection vector, and the $\max$ operator encodes the ELB in the observation equation for the interest rate. We set the value of the ELB to 25 basis points, and assume that the federal funds rate was at the ELB for every quarter in which the average

[12] Our data are quarterly, and we include two lags in $\boldsymbol{A}(L)$.

[13] In Appendix A we show how $\boldsymbol{\mathcal{A}}$, $\boldsymbol{\mathcal{B}}$ and $\boldsymbol{\mathcal{C}}$ can easily be made time–varying. Missing data for the two–year yields prior to 1976:Q2, are, for example, handled by deterministic time–variation in $\mathcal{C}_t$ where rows associated with missing observations are set to zero.

annualized 3–month yield was less than 25 basis points.

We collect the unobserved state variables of our model in the vector $\boldsymbol{\xi}$, which contains each vector $\boldsymbol{\xi}_t$ stacked by time. Our approach for drawing from the posterior of $\boldsymbol{\xi}$ is to first treat the interest rate data at the ELB as missing and take draws from the posterior distribution of $\boldsymbol{\xi}$, which is straightforward using standard filtering and smoothing techniques. Knowing that the interest rate is at the ELB (not simply missing) in period t amounts to knowing that the values of $\boldsymbol{\xi}_t$ that are consistent with the data imply values of s_t that are less than the effective lower bound during that period. Thus, we can draw from the posterior of $\boldsymbol{\xi}$ by first treating interest rate data at the ELB as missing, and then rejecting draws until we find a $\boldsymbol{\xi}$ that is consistent with the ELB. Notice that rejections of $\boldsymbol{\xi}$ are not part of a Metropolis-Hastings algorithm for constructing the posterior of $\boldsymbol{\xi}$. Instead, we reject draws of $\boldsymbol{\xi}$ that are inconsistent with the knowledge that the short–term interest rate was at the ELB over certain periods, and thus use the rejection step to draw directly from the posterior of $\boldsymbol{\xi}$.

Our estimation procedure is a generalization of Park et al. (2007) that applies the methodology of Hopke et al. (2001). Appendix A explains in further detail how to construct a draw from the posterior distribution of $\boldsymbol{\xi}$, conditional on the data, in a conditionally–linear Gaussian state–space model like ours. With a draw of $\boldsymbol{\xi}$ 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). We jointly estimate the parameters and unobserved states of the model using Bayesian MCMC techniques; our priors and details of the MCMC steps are described in Appendix B.

3 Shadow Rate and Trend Estimates

In this section we describe the posterior distribution of our model with regard to estimated shadow rates and trends. Our model is estimated using quarterly data from 1960:Q1 to 2017:Q2, which includes the recent period at the ELB. All data are publicly available from the FRED database maintained by the Federal Reserve Bank of St. Louis. Inflation is mea-

sured by the quarterly rate of change in the PCE headline deflator (expressed as an annualized percentage rate). Readings for the nominal yields are constructed as quarterly averages of the Treasury's constant maturity rates. The unemployment rate gap is computed as the difference between the quarterly average rate of unemployment and the CBO's measure of the natural long–term rate of unemployment for a given quarter. The output gap is computed as the log difference between real GDP and the CBO's measure of potential real GDP for a given quarter.[14] All computations are based on the vintage of FRED data available that has been available at the end of October 2017.[15]

3.1 The Shadow Rate

Our model delivers estimates of the shadow rate of interest that are, by construction, less than the ELB during the period in which the bound is binding. While we consider both the unemployment rate gap and the output gap as cyclical factors in our model, the shadow rate estimates are similar in both cases. Figure 1 shows results from our model when the cyclical factor is measured by the output gap.[16] Panel (a) of Figure 1 shows the posterior mean of our shadow rate, along with uncertainty bands, on the same plot as estimates from Wu and Xia (2016), and Krippner (2013).[17] Two features are worth noting. First, our estimated shadow rate, which also conditions on the business cycle measure and inflation, is lowest during 2009, near the trough of the Great Recession, according to the NBER. By contrast, the other estimates reach low points much later, and those low points are more–negative than our estimate. Second, all three estimates are remarkably similar at the end of 2015, just before the Federal Reserve's departure from the ELB.

[14] In order to be comparable to annualized growth rates, the log difference between actual and potential GDP has been scaled by a factor of 400 when computing the output gap.

[15] The FRED database is available at `https://research.stlouisfed.org/fred2/`. In FRED, the PCE headline deflator has the mnemonic `PCECTPI`. For Treasury yields we used `TB3MS`, `GS2`, `GS5`, and `GS10`. Data for the 5-year yield is available only as of 1976:Q2 and prior observations are treated as missing within our state space model. The unemployment rate and the CBO's estimate of the natural rate of unemployment in the long run are given by `UNRATE` and `NROU`. Real GDP and the CBO's estimate of potential real GDP are given by `GDPC1` and `GDPPOT`.

[16] Figure 8 in Appendix C shows results from our model when the cyclical factor is measured by the unemployment rate gap.

[17] Measures of the shadow rate from Priebsch (2013) and Ichiue and Ueno (2013) are not shown because their sample ends in 2013. Their measures are qualitatively similar in that they reach low points well after the trough of the recent recession.

Panel (b) of Figure 1 shows quasi real–time estimates of the shadow rate, which are conditioned solely on data through period t.

[Figure 1 about here.]

The quasi real–time estimates of the shadow rate are notably similar to the full–sample estimates. This similarity indicates that the data are informative about the shadow rate, even in real time. Additionally, this similarity suggests that the length of the ELB period (7 years) does not heavily influence the particular values of the full–sample estimates of the shadow rate.

3.2 The Real Rate in the Long Run

Figure 2 displays the posterior median of our estimates of the trend real rate, $\bar{r}_t$, along with uncertainty bands from our model when the cyclical factor is measured as the output gap.[18] Panel (a) of Figure 2 shows the smoothed estimates, in that the entire data sample is used to estimate the parameters and $\bar{r}_t$. Panel (b) of Figure 2 shows quasi real–time estimates of $\bar{r}_t$, which are conditioned solely on data through period t. Consistent with results reported by Hamilton et al. (2015), Kiley (2015), and Lubik and Matthes (2015), the uncertainty bands surrounding our estimates of $\bar{r}_t$ are wide.

[Figure 2 about here.]

The median estimate of $\bar{r}_t$ using the entire data sample has about a 1 percentage point decline from the mid 1990s to the end of the sample. The pseudo–real–time estimate of $\bar{r}_t$ is more volatile than the estimate from the entire data sample, and any downward movement in $\bar{r}_t$ seems to start around 2000. Notably, the size of the movement in the pseudo–real–time estimate of $\bar{r}_t$ since 2000 is not unprecedented. To the extent that these estimates could be taken to suggest a downward shift in $\bar{r}_t$, the series peaked well before the Great Recession. Studies like Laubach and Williams (2015) and Lubik and Matthes (2015) also document downward trajectories in the trend real rate; however, our estimates do not dip nearly as much as their

[18] Figure 9 in Appendix C shows results from our model when the cyclical factor is measured by the unemployment rate gap.

estimates. One reason for the differences is our inclusion of stochastic volatility, which allows for large movements in the gap components of our model in 2008 and 2009.

3.3 Stochastic Volatility in Real Rate Trend Shocks

In our benchmark model, the shocks to $\bar{r}_t$ have a constant volatility. To assess if this assumption is appropriate, we estimated a version of our model where $\bar{r}_t$ evolves so that

$$\bar{r}_t = \bar{r}_{t-1} + \sigma_{\bar{r},t}\epsilon_{\bar{r},t}, \tag{13}$$

$$\log\left(\sigma^2_{\bar{r},t}\right) = (1-\rho_{\bar{r}})\mu_{\bar{r}} + \rho_{\bar{r}}\log\left(\sigma^2_{\bar{r},t-1}\right) + \phi_{\bar{r}}\eta_{\bar{r},t}. \tag{14}$$

As with the inflation trend, we assume that $\epsilon_{\bar{r},t} \sim N(0,1)$, $\eta_{\bar{r},t} \sim N(0,1)$, and $|\rho_{\bar{r}}| < 1$. We compare the fit of the model with stochastic volatility in $\bar{r}_t$ to our benchmark specification using the marginal data density (MDD), $p(\boldsymbol{Z})$, where $\boldsymbol{Z}$ is the observable data.

We estimate the MDD using the harmonic mean estimator of Geweke (1999), as presented by Herbst and Schorfheide (2014), who express the MDD as

$$p(\boldsymbol{Z}) \approx \left[\frac{1}{N}\sum_{i=1}^{N}\frac{f(\boldsymbol{\theta}^i)}{p(\boldsymbol{Z}|\boldsymbol{\theta}^i)p(\boldsymbol{\theta}^i)}\right]^{-1}, \tag{15}$$

where N is the number of draws from the posterior distribution, $\boldsymbol{\theta}$ is a vector that collects all of the estimated parameters and $\boldsymbol{\theta}^i$ is a particular draw from the posterior distribution, f is a function of the parameter vector that integrates to one, $p(\boldsymbol{\theta})$ is the prior density of $\boldsymbol{\theta}$, and $p(\boldsymbol{Z}|\boldsymbol{\theta}^i)$ is the likelihood.[19] For our models, computation of the likelihood requires a particle filter because of multiple layers of latent variables — trends and gaps as well as the stochastic volatilities — as well as our handling of the ELB; see also Fuentes-Albero and Melosi (2013)

[19]For f, we use

$$\begin{aligned} f(\boldsymbol{\theta}) =& \tau^{-1}(2\pi)^{-d/2}|\boldsymbol{V}_{\boldsymbol{\theta}}|^{-1/2}\exp\left[-0.5(\boldsymbol{\theta}-\bar{\boldsymbol{\theta}})'\boldsymbol{V}_{\boldsymbol{\theta}}^{-1}(\boldsymbol{\theta}-\bar{\boldsymbol{\theta}})\right] \\ &\times \mathcal{I}\left\{(\boldsymbol{\theta}-\bar{\boldsymbol{\theta}})'\boldsymbol{V}_{\boldsymbol{\theta}}^{-1}(\boldsymbol{\theta}-\bar{\boldsymbol{\theta}}) \leq F^{-1}_{\chi^2_d}(\tau)\right\} \end{aligned}$$

where $\bar{\boldsymbol{\theta}}$ and $\boldsymbol{V}_\theta$ are the mean and variance of the posterior distribution of $\boldsymbol{\theta}$, d is the length of $\boldsymbol{\theta}$, and $F_{\chi^2_d}$ is the cumulative distribution function of the χ^2 distribution with d degrees of freedom, and $\mathcal{I}$ is the indicator function. We set $\tau = 0.9$, and have found that results are robust to other choices of τ.

and Chan and Grant (2015). Our particle filter is an auxiliary particle filter that employs Rao-Blackwellization to handle the linear parts of the state space. Additionally, the particle filter handles the ELB by sampling linear states from the truncated normal distributions implied by the posterior distribution of the shadow rate.[20]

We find substantially higher MDD for the model without stochastic volatility in the real rate trend: When the output gap is used as the business cycle measure, the log-MDD is estimated at 304.6 as opposed to 228.8 for the model with stochastic volatility. Similarly, when the unemployment rate gap is used as the business cycle measure, we find that the MDD is substantially larger for the model without stochastic volatility (570.1 as opposed to 492.6 for the model with stochastic volatility).[21] These results confirm our prior disposition, as informed by results discussed in Hamilton et al. (2015), that $\bar{r}_t$ is best modeled without stochastic volatility.

4 Forecasting Interest Rates

Our shadow–rate approach has significant implications for forecasting interest rates. In this section we offer a number of ways to evaluate our shadow–rate approach by analyzing the model's out–of–sample predictions.

4.1 Predictive Density at Selected Dates

To create out–of–sample forecasts from our model, at each date we use the median of the predictive densities from the posterior distribution (using data only up to that date) to forecast future interest rates. Our forecasting procedure thus captures uncertainty about both the parameter values and the unobserved states in our model. The posterior predictive distribution can be highly asymmetric because of the ELB. As a result, we follow Bauer and Rudebusch (2015) and use the posterior median as the point forecast for forecast evaluation, rather than the posterior mean. We compare these forecasts to forecasts from the SPF, conducted by the Federal Reserve Bank of Philadelphia, and to median forecasts from the WX-SRTSM.

[20] Further details regarding our particle filter are provided in Appendix E.

[21] Differences in log-MDD above one, which correspond to Bayes factors above three, are considered to be substantial; see Kass and Raftery (1995) and Jeffreys (1961).

[Figure 3 about here.]

So as to illustrate its skewness, Figure 3 display statistics from the posterior predictive density of the short–term nominal interest rate from our baseline model (with the output gap as the cyclical factor) at different dates. The forecast horizon extends for five years, and, in addition to mean and median predictions, shaded areas indicate 50 and 90 percent uncertainty bands. The dashed lines that extend below the ELB indicate posterior quantiles of the shadow rate distribution (as opposed to the interest rate distribution). The predictive density of the interest rate is a truncated version of the predictive density of the shadow rate distribution, so the quantiles of the shadow rate distribution become exactly the quantiles of the interest rate distribution if the value is larger than the ELB. The truncation of the shadow–rate distribution causes substantial asymmetry in the interest rate distribution leading to marked differences in the predictive means and medians of our baseline model.

In 2008:Q4, the first period before the ELB (Panel (a) of Figure 3), the predictive density from our model takes the ELB already into account and produces interest rate forecasts based on the truncated distribution of future shadow rates. In doing so, the mean interest–rate forecast rises appreciably above the median for several periods.

In 2009:Q1, the period the ELB begins to bind (Panel (b) of Figure 3), accounting for the ELB produces interest rate forecasts that place substantial probability on remaining exactly at the ELB for several quarters. As in Panel (a), the truncation of the shadow rate distribution in order to produce interest rate forecasts creates a divergence of mean and median estimates of interest rates for several years.

In 2010:Q4, after the ELB had been binding for some time (panel (c) of figure 3), our baseline model still predicts substantial probability of interest rates at the ELB because of the estimated negative shadow rate. Moreover, the median interest rate forecast remains at the ELB for a number of quarters. Toward the end of the ELB period (2015:Q4, shown in panel (d) of Figure 3), the forecasts for the short-term interest rate are similar to the forecasts for the shadow rate, in large part because our estimated shadow rate is only slightly less than the ELB at that point.

4.2 Forecasting Performance

To assess the forecasting performance of our model, we compare the accuracy of forecasts from our model to the accuracy of forecasts from the SPF and the WX-SRTSM. For the SPF, forecasters submit survey responses in the middle of the quarter. Because forecasts are submitted in the middle of the quarter, we date the forecasts as being made in the previous quarter so as to not give our model an informational advantage. For interest rate forecasts, forecasters submit forecasts for the average value of the 3-month and 10-year Treasury rates in the current and subsequent 4 quarters. These rates are part of our benchmark data set, and as such we compare forecasts from our model directly to forecasts from the SPF. We compare our model's forecast to the average forecast across forecasters in the SPF.

The WX-SRTSM is estimated at a monthly frequency using month-end data on fitted zero-coupon Treasury yields since 1990 from Gurkaynak et al. (2007). So as to be consistent, when comparing our model to the WX-SRTSM, we estimate our model on data since 1990 and use quarter-end interest rate data from the fitted zero-coupon Treasury yields from Gurkaynak et al. (2007) of the same maturity as in our baseline case. We compare our model's quarterly forecast for quarter-end interest rates to the median forecast from the WX-SRTSM in the final month of each quarter.

To assess performance, we focus on two statistics: the root–mean–squared error and the mean absolution deviation. Table 1 compares the forecasts from the SPF with our model for the post 2008 period. The statistics for our model are shown as calculated. The statistics for the SPF are shown on a relative basis to our model.

Except at very short horizons where the SPF has an informational advantage, our model performs better than the SPF at both the 3–month and 10–year horizon. To assess the statistical significance of these differences, we use the statistic proposed by Diebold and Mariano (1995) and West (1996).[22] The differences in forecast performance between our model and the SPF for the 3-month yield are not statistically significant, however our model statistically outperforms

[22]McCracken (2000) discusses issues related to forecast comparison using the mean absolute deviation when parameters of the model producing the forecasts are estimated. Our out-of-sample comparisons begin in 2008, affording us a relatively small sample size. For the purposes of forecast comparison, we construct the statistics as in Diebold and Mariano (1995).

the SPF for forecasting the 10-year yield at horizons greater than 3 quarters.

[Table 1 about here.]

Table 2 compares the forecasts from the WX-SRTSM with our model for the post 2008 period. The statistics for our model are shown as calculated and the statistics for the WX-SRTSM are shown on a relative basis to our model. Our model performs at least as well as the WX-SRTSM, on balance, for both the 3-month and 10-year Treasury yields. Thus, even though our model does not impose the cross-equation restrictions associated with the no-arbitrage assumptions of the WX-SRTSM, our model does no worse over the ELB period in forecasting interest rates.

[Table 2 about here.]

4.3 Forecast Uncertainty

Naturally, the ELB has important effects on the predictive density for nominal interest rates when the predictive density for shadow rates has non-negligible coverage below the ELB. To illustrate the relevance of these effects, Figure 4 compares forecast uncertainty about the shadow rate with forecast uncertainty about the short–term interest rate. For the purpose of this figure, we measure forecast uncertainty by the conditional standard deviation of the predictive densities described above for one and eight-quarters ahead. Overall, near– and medium–term uncertainty about future short–term rates (and the shadow rate) has mostly declined since the mid-1980s. Nevertheless, as the level of nominal rates has been trending down over this period as well, the probability of reaching the ELB has become more and more non-negligible; this is particularly true for longer–horizons forecasts made since 2000, causing forecast uncertainty about the shadow rate to differ from forecast uncertainty about the short–term interest rate.

[Figure 4 about here.]

Not surprisingly, the onset of the last NBER recession in 2007 is reflected in higher estimated levels of stochastic volatility to all shocks in our model, leading to increased shadow–rate uncertainty. By accounting for the ELB, our baseline model recognizes that during the

last recession, increased shadow–rate uncertainty is accompanied by a marked downward shift of the shadow–rate distribution to values below the ELB such that the truncated distribution of actual short–term nominal interest rates almost collapses at values at or slightly above the ELB. Consequently, near–term uncertainty for short–term nominal rates declines during the last recession when properly accounting for the ELB, as shown in Panel (a) of Figure 4. In contrast, as shown in Panel (b) of the figure, medium–term uncertainty about nominal interest rates increases with the increasing shadow–rate uncertainty, though not by quite as much, as nominal rates are projected to return to their estimated non–negative trend level.

4.4 Long–Term Forecasts of Short–Term Rate

Our model's predicted path for future short–term interest rates implies a forecast of the short–term nominal interest rate in 10 years, as does the WX-SRTSM. In addition, once per year, respondents to the SPF offer forecasts of the average short–term interest rate over the next 10 years. These forecasts are shown in Figure 5.

[Figure 5 about here.]

Bauer and Rudebusch (2017b) have emphasized the importance of long–term forecasts for predicting interest rates at shorter horizons. The 10–year ahead forecast of the 3–month Treasury rate from the WX-SRTSM is more volatile than the forecast from our model. While the SPF provides forecasts of the average 3–month Treasury bill rate over the next 10 years, those forecasts are only somewhat less volatile than our model's 10–year ahead forecast of the 3–month rate. While all three lines appear to have some downward movement since 2000, the forecasts from the WX-SRTSM are notable because they, at times, are equal to the ELB. Unlike the WX-SRTSM, our model includes trend components in interest rates, which appears to reduce the volatility of the long–term forecasts of the short–term interest rate.

5 Impulse-Response Analysis

A common question of interest in empirical research is to measure the response of macroeconomic variables to identified economic shocks; for example, shocks due to monetary policy. Absent a binding lower–bound on nominal interest rates, the conventional tool for such impulse-response function (IRF) analysis are linear vector-autoregressions (VARs). In this section, we show how to perform a similar analysis within our shadow-rate approach to time series modeling.

Our application is similar in spirit to the VAR analysis of Wu and Xia (2016) and others who use estimated shadow rates as observables in a (factor-augmented) VAR system and identify monetary policy shocks via short-run restrictions on VAR-implied forecast errors for the shadow rate.[23] However, we also extend their approach in several dimensions: First, instead of a two-step approach that ignores uncertainty about the true shadow-rate values, we estimate impulse responses jointly with our shadow-rate inference. Second, our model implies a time-varying vector moving average (VMA) representation of the data with respect to observable forecast errors, which generates time-varying impulse responses with distinctively different patterns at and away from the ELB. Similar to Wu and Xia (2016), we focus on monetary shocks identified by short-run restrictions on surprise changes in the shadow rate. Specifically, our short-run restrictions are identical to those used by Christiano et al. (1999) for the short-term nominal interest rate. In the same spirit as Christiano et al. (1999), we define identified shocks with respect to the level of the shadow rate, rather than any trend or cyclical component separately.

5.1 Identification of monetary policy with shadow-rate shocks

In a VAR framework, the reduced-form dynamics of a vector of variables ($\boldsymbol{X}_t$) are described by $\boldsymbol{A}(L)\boldsymbol{X}_t = \boldsymbol{e}_t$ where $\boldsymbol{A}(L) = \sum_{i=0}^{\infty} \boldsymbol{A}_i L^i$ is a lag polynominal with $\boldsymbol{A}_0 = \boldsymbol{I}$ and $\boldsymbol{e_t} = \boldsymbol{X}_t - E(\boldsymbol{X}_t|\boldsymbol{X}^{t-1})$ is a vector of mean-zero forecast errors with variance $\boldsymbol{\Omega}$.[24] Economic assumptions then allow to relate the reduced-form forecast errors to structural shocks

[23] See, for example, also Hakkio and Kahn (2014), Doh and Choi (2016), Bundesbank (2017), Francis et al. (2017)

[24] Practical application are typically limited to finite-order VARs where $\boldsymbol{A}_i = \boldsymbol{0} \quad \forall\, i > p$ for some p.

$\boldsymbol{z}_t = \boldsymbol{Q}^{-1}\boldsymbol{e_t}$ where $\boldsymbol{Q}\boldsymbol{Q}' = \boldsymbol{\Omega}$. For example, in their by-now canonical approach, (Christiano et al., 1999, henceforth: CEE) identify monetary policy shocks by placing a short-run restriction on $\boldsymbol{Q}$: Monetary policy shocks are assumed to be given by the residual obtained from regressing forecast error in the policy rate on forecast errors in other macroeconomic (though not financial) variables. This scheme can be implemented by ordering the policy rate in $\boldsymbol{X}_t$ after other macroeconomic variables and before financial variables — similar to our definition of $\boldsymbol{X}_t$ above (ignoring the distinction between shadow rate s_t and actual policy rate i_t for now). $\boldsymbol{Q}$ can then be set equal to the lower–triangular Choleski factorization of the variance-covariance matrix $\boldsymbol{\Omega}$ of forecast errors, and the monetary policy shock is given by the element of $\boldsymbol{z}_t$ corresponding to the location of the policy rate in $\boldsymbol{X}_t$. Impulse responses are then generated by the VMA representation $X_t = \boldsymbol{A}(L)^{-1}Q\boldsymbol{z}_t$.

However, when considering the ELB on nominal interest rates, such a VAR approach runs into complications: First, the linear dynamics implied by a VAR for the actual policy rate are hardly applicable at the ELB. As proposed by our model as well as other shadow-rate approaches, it is, however, conceivable that the dynamics of the shadow rate can be described by a process that is at least conditionally linear.[25] At the ELB, the shadow rate is an unobserved state variable that matters for forecasting future outcomes in the policy rate and other variables. Unexpected variations in the shadow rate can thus be interpreted as reflecting changes in monetary policy implemented through unconventional tools (such as asset purchases or forward guidance). Similar to Wu and Xia (2016), we are thus interested in characterizing the response of macroeconomic and yield-curve variables to monetary policy shocks as measured by surprise changes in the shadow-rate after controlling for endogenous variations in the shadow rate. For simplicity we adopt a short-run restriction that will be identical to the CEE identification for policy-rate shocks when the ELB does not bind. When the ELB binds, we apply the CEE-Choleski identification to surprise changes in the shadow rate, which is jointly estimated with the IRF.

[25] By *conditionally* linear we mean to include cases like our state space model in (10) and (11), which is linear conditional on knowledge of trends and gaps (the components of $\boldsymbol{\xi}_t$) as well as the evolution of time-varying parameters, like the stochastic volatilities that cause variation in $\boldsymbol{\Sigma}_t$ in (10). Of course, this includes also cases of outright linear models, such as the FAVAR considered by Wu and Xia (2016).

Second, economic theory suggests the economy might react differently to policy shocks when the ELB binds – even after accounting for the immediate non-linearity in policy-rate dynamics arising from the ELB.[26] If so, impulse-responses generated from a time-invariant VAR will hardly provide an appropriate characterization.

In principle, a VAR with time-varying parameters (TVP-VAR) could be used to address concerns about time-varying IRFs (Gali and Gambetti, 2009). However, even when considering only a relatively relatively small set of variables as in our case, such an approach can quickly run into curse-of-dimensionality problems; and, as documented already by Cogley and Sargent (2005), there are likely fewer underlying sources of parameter variations than VAR coefficients. Moreover, the many degrees of freedom of a TVP-VAR that exist already in the case of known observables will impair our shadow-rate inference, which is build around a missing observations problem.

In contrast, the unobserved components (UC) model characterized by (10) and (11) implies a time-invariant VMA with respect to innovations $\boldsymbol{e}_t$ that are defined relative to the history of the data vector $\boldsymbol{X}_t$ as well. However, in contrast to the case of a generic TVP-VAR, time-variation in the resulting IRFs is the result of lower–dimensional variations in the stochastic volatilities that affect the UC shocks ε_t in (10). Specifically, assume model parameters are known and a trajectory of stochastic volatilities are given such that $\boldsymbol{\mathcal{A}}$, $\boldsymbol{\mathcal{B}}_t$, $\boldsymbol{\mathcal{C}}$ are known when taking expectations. And denote expectation of $\boldsymbol{X}_t$ and $\boldsymbol{\xi}_t$ relative to the history of X^{t-1} by $\boldsymbol{X}_{t|t-1}$ and $\boldsymbol{\xi}_{t|t-1}$, respectively.[27] Denote innovations of $\boldsymbol{X}_t$ relative to its own history by

$$\boldsymbol{e}_t = \boldsymbol{X}_t - \boldsymbol{X}_{t|t-1}. \tag{16}$$

From (10) and (11) we then obtain a time-varying VMA representation of $\boldsymbol{X}_t$ with respect to

[26] See, for example, the models of Eggertsson and Krugman (2012), Christiano et al. (2011), Johannsen (2014), and Gavin et al. (2015), as well as the DSGE-based estimates of Gust et al. (2017) and Aruoba et al. (2017). Alternatively, estimates by Stock and Watson (2012) suggest that the propagation of shocks did not materially change over the Great Recession (for variables other than the policy rate), at least after conditioning on a broad set of economic factors and when accounting for the unusual depth of the recession. Wu and Zhang (2016) also argue that economic responses are time-invariant once shadow-rate dynamics are included in the analysis.

[27] Formally we have thus $\boldsymbol{X}_{t|t-1} = E\left(\boldsymbol{X}_t|\boldsymbol{X}^{t-1}\right)$ and $\boldsymbol{\xi}_{t|t-1} = E\left(\boldsymbol{\xi}_t|\boldsymbol{X}^{t-1}\right)$ where $\boldsymbol{X}^{t-1}$ denotes the infinite horizon history $\boldsymbol{X}^{t-1} = \{\boldsymbol{X}_{t-1}, \boldsymbol{X}_{t-2}, \ldots\}$.

innovations spanned by its own history – henceforth the "the VMA of $\boldsymbol{X}_t$":

$$\boldsymbol{\xi}_{t|t} = \boldsymbol{\mathcal{A}}\boldsymbol{\xi}_{t-1|t-1} + \boldsymbol{K}_t \boldsymbol{e}_t \tag{17}$$

$$\boldsymbol{X}_t = \left(\boldsymbol{I} + \boldsymbol{\mathcal{C}}\left(\boldsymbol{I} - \boldsymbol{\mathcal{A}}L\right)^{-1} \boldsymbol{K}_t\, L\right) \boldsymbol{e_t} = \boldsymbol{\mathcal{D}}_t(L)\boldsymbol{e_t} \tag{18}$$

Time variation in $\boldsymbol{\mathcal{D}}_t$ stems from $\boldsymbol{K}_t$, which is a time-varying Kalman gain matrix induced by the stochastic volatilities that cause variations in $\boldsymbol{\mathcal{B}}_t$.[28] Time-variation in the VMA of $\boldsymbol{X}_t$ with respect to $\boldsymbol{e}_t$ thus reflects changes in the signal-to-noise ratios involved in filtering out trends and gaps, which leads to time-varying persistence in the effects of $\boldsymbol{e}_t$ on $\boldsymbol{X}_t$.

Applied to the UC model with stochastic volatility, standard Kalman-filtering computations, described in Appendix E, imply a time-varying variance covariance matrix for the innovations $\mathrm{Var}_t\left(\boldsymbol{e}_t\right) = \boldsymbol{\Omega}_t = \boldsymbol{Q}_t\boldsymbol{Q}_t'$ with lower–triangular Choleski factor $\boldsymbol{Q}_t$. The third column of $\boldsymbol{Q}_t$, which is associated with the component of the shadow-rate innovation that is orthogonal to innovations in inflation and the business cycle shall be denoted $\boldsymbol{Q}_{t,m}$. Echoing the CEE-VAR approach, responses of $\boldsymbol{X}_t$ to a monetary policy shock are then given by

$$\boldsymbol{X}_t = \boldsymbol{\mathcal{D}}_t(L)\boldsymbol{Q}_{t,m}\boldsymbol{z}_{t,m} = \boldsymbol{\mathcal{D}^{m,*}}_t(L)z_{t,m}. \tag{19}$$

where $z_{t,m}$ denotes a standard–normal monetary policy shock. To generalize the impulse responses to the case of actual–rate dynamics, which are non–linear due to the censoring constraint (1), note that that the impulse response coefficients in $\boldsymbol{\mathcal{D}^{m,*}}_t(L) = \sum_{k=0}^{\infty} \boldsymbol{\mathcal{D}^{m,*}}_{k,t}L^k$ reflect updates in forecasts of $\boldsymbol{X}_{t+k}$ prompted by the observation of a unit–sized policy shock at time t:

$$\boldsymbol{\mathcal{D}^{m,*}}_{k,t} = E\left(\boldsymbol{X}_{t+k}|\boldsymbol{X}^{t-1}, z_{t,m} = 1\right) - E\left(\boldsymbol{X}_{t+k}|\boldsymbol{X}^{t-1}\right) \tag{20}$$

where the expectations condition also on model parameters. The right-most term in (20) can also be referred to as "baseline" forecast. Of course, once the economy has touched the ELB, $\boldsymbol{X}^{t-1}$ is not observable anymore and we condition instead on the observable data $\boldsymbol{Z}^{t-1}$ via particle filter described further below and in the appendix.

[28] The Kalman filtering computations are standard with details described in Appendix E.

By analogy, we can generate impulse responses for the actual rate by simulation of the predictive density for i_{t+k}, once conditional on $\boldsymbol{X}^{t-1}$ and then conditional on $\left(\boldsymbol{X}^{t-1}, z_{t,m}=1\right)$, compute expectations via numerical integration across the simulated draws of each predictive density and then consider the updates in the forecast trajectories induced by the policy shock.

In our empirical implementation, we compute impulse response conditional on parameter values that are set equal to their posterior median estimates from a full-sample MCMC estimation of the model. However, in order to avoid look-ahead bias imparted by the smoothed SV estimates generated by the MCMC estimation, we have embedded the IRF computations in a particle filter. The IRF in (20) thus condition on a fixed set of model parameters estimated over the full sample, but continuously update underlying *filtered* estimates of the stochastic volatilities. Specifically, we exploit the conditionally linear structure of our model with a Rao-Blackwellized particle filter that simulates a large range of SV trajectories ("the particles") and associated Kalman filters based on the "mixture Kalman filters" of Chen and Liu (2000); see, for example, Creal (2012) and Lopes and Tsay (2011) and the applications in Mertens and Nason (2017) and (Mertens, 2016). Output for each particle is then weighted by their respective likelihood. We perform the computations described above separately for each particle and then integrate the resulting IRF estimates across particles. Similar to our MCMC algorithm, shadow rate estimates are simulated from treating shadow rates as unobserved while imposing that $s_t < ELB$ when the policy rate is constrained by the ELB; details are provided in Appendix E.

5.2 Economic responses to shadow-rate shocks

Figure 6 reports impulse response to monetary policy shocks identified as surprise innovations in the shadow rate that are uncorrelated with contemporaneous forecast errors in the business cycle and inflation. Shocks are identified as part of a particle filtering loop, which simultaneously estimates the linear states of our model ($\boldsymbol{\xi}$) as well as the stochastic volatilities. Variations in the latter cause variations in the filters assessment of the permanent effects of shocks defined relative to the history of $\boldsymbol{X}^t$. We consider impulse-responses at four different dates that represent different points in time before, during and after the ELB had been binding in recent

U.S. history: 2007, 2009, 2011, and 2016.[29] For the sake of comparability, monetary policy shocks have been scaled to generate a percentage point drop in the shadow rate on impact.

[Figure 6 about here.]

Monetary policy shocks are estimated to have transitory and permanent effects on the shadow rate. As shown in Panel 6a, it takes roughly a year for the shadow rate to settle on a new permanent level after a shock. Moreover, shadow-rate shocks are estimated to have had much less permanent effects on the level of the policy rate during the ELB years compared to non-ELB times. A similar pattern is also visible in the responses of longer-term yields, see Panels 6c and 6d.

At least part of the permanent effect of a policy shock on the level of interest rates is attributed to a shift in the inflation, with the remainder accounted for by the real rate trend. As shown in Panel 6e, the response of inflation to a monetary policy shock is characterized by a fairly swift adjustment to the new trend level without much transitory dynamics. If anything, at the ELB, inflation slightly overreacts to its new trend level (i.e. drops initially by more) whereas the adjustment is more monotonic when away from the ELB. Similar to the nominal-rate responses, shadow-rate shocks have much less permanent effects on inflation when the ELB binds.

The responses of the output gap (our business cycle measure) also differ markedly at times when the ELB binds rather than not. In particular, same-sized monetary policy shocks are estimated to generate a more sizable pick-up in real activity — that peaks after about a year and then wanes only gradually over the next few years — when the ELB was binding rather than when not.[30]

The actual policy rate depends on the shadow rate via the censoring function (1). Because of this censoring, responses of the actual rate to shocks depend on the current level of the actual rate as well as the size and the sign of the impulse.[31]

[29] Results reported for the years 2007 and 2016 are also similar to what we obtained during earlier decades.

[30] Since the trend level of the business cycle measure is fixed at zero, the effects of any shock on the business cycle are ultimately transitory.

[31] All other responses shown in Figure 6 are symmetric and correspond up to their sign to the effects of a contractionary shock as well. (While the ELB could in principle also bind for longer-term yields, their levels remained sufficiently elevated for the ELB not to matter for the impulses responses shown here.)

The effects of the ELB on the actual rate impulses visible by comparing the responses of actual rate and shadow rate shown in Panels 6b and 6a, respectively. Not surprisingly, during the ELB years, the actual-rate responses to a *negative* shadow-rate shock are close to zero. However, given the estimated low levels of the shadow rate — reported in Figure 1 — for the years considered here, actual-rate responses to a percentage-point *increase* in the shadow rate (not shown) also flatline close to zero. Considering responses estimated for 2007:Q4 — when the 3-month rate was still above 3 percent and thus well away from the ELB — actual-rate and shadow-rate response are fairly similar. In contrast, the actual-rate response to a one-percentage point drop in the shadow rate estimated for in 2016:Q4 is still quite muted, as the 3-month rate still just hovered below 0.5 percent at that time.

[Figure 7 about here.]

Given the duration and tightness of the ELB period over the last decade, it might seem surprising to see any decline at all in the estimated actual-rate responses to a drop in the shadow-rate reported in Panel 6b. However, impulse response paths merely reflect changes in forecasted trajectories, not changes relative to the initial level of the actual rate. During the ELB years, the actual rate is estimated to have been below its estimated longer-run level to which it is forecasted to rise back again. Negative values for the actual-rate responses after a drop in the shadow rate merely indicated a slower return to trend levels than before, not a cut below the level of the actual rate that prevailed on impact. The resulting asymmetries are also illustrated in Figure 7, which contrasts the hypothetical effects of a positive and a negative shadow rate shock in 2015:Q4 (the quarter of the first rate hike since the economy had reached the ELB in 2009).

6 Conclusion

In this paper, we develop a methodology to account for the ELB in time series models. Our method applies not only to models that — apart from the ELB constraint – are Gaussian and linear, but also to models that are only conditionally linear and Gaussian, for example, due to

time-varying parameters, like stochastic volatilities. Further, we demonstrate how to estimate the parameters and latent states of such a model with an otherwise standard Bayesian MCMC sampler.

We document that including the ELB can have drastic effects for interest rate forecasts, as well as the expectations component of longer–term yields and thus also the computation of term–premiums. Further, accounting for the ELB using our shadow–rate approach appears to improve forecast performance for interest rates. We also estimate changes in the trend real rate, defined as a long–term forecast of the real interest rate, and find that any decline in the trend real rate since the onset of the Great Recession is best characterized as a continuation of a downward trajectory that began well before. In addition, our model estimation empirically reveals interesting time variation in impulse response functions around the recent ELB episode.

APPENDIX

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 $\boldsymbol{\xi}_t$ is a random variable that evolves so that

$$\boldsymbol{\xi}_t = \boldsymbol{\mathcal{A}}_t \boldsymbol{\xi}_{t-1} + \boldsymbol{\mathcal{B}}_t \boldsymbol{\varepsilon}_t \tag{21}$$

where $\boldsymbol{\varepsilon}_t$ is a vector of standard normal random variables of appropriate length and the sequence of matrices $\{\boldsymbol{\mathcal{A}}_t\}_{t=1}^T$ and $\{\boldsymbol{\mathcal{B}}_t\}_{t=1}^T$ are given.[32]

Define the vector

$$\boldsymbol{X}_t = \boldsymbol{\mathcal{C}}_t \boldsymbol{\xi}_t \tag{22}$$

[32]In our application, described in the main body of the paper note that we have a constant $\boldsymbol{\mathcal{A}}_t = \boldsymbol{\mathcal{A}}$ while $\boldsymbol{\mathcal{B}}_t$ is given by $\boldsymbol{\mathcal{B}}_t = \boldsymbol{\mathcal{B}}\,\boldsymbol{\Sigma}_t^{1/2}$ where $\boldsymbol{\mathcal{B}}$ is unit–lower–triangular.

where the sequence of matrices $\{\boldsymbol{\mathcal{C}}_t\}_{t=1}^T$ are known.[33] We assume that $\boldsymbol{X}_t$ has a partition made up of a vector shadow rates, ($\boldsymbol{S}_t$), and a partition of variables that are unconstrained by the ELB, ($\boldsymbol{M}_t$).[34] That is,

$$\boldsymbol{X}_t = \begin{bmatrix} \boldsymbol{S}_t \\ \boldsymbol{M}_t \end{bmatrix} \tag{23}$$

The observed data are

$$\boldsymbol{Z}_t = \begin{bmatrix} \max\left(\boldsymbol{S}_t, ELB\right) \\ \boldsymbol{M}_t \end{bmatrix} \tag{24}$$

where the $\max$ operator is applied element by element. The ELB acts as a censoring function in the model through the $\max$ operator, though more general censoring functions could be used.

Define $\boldsymbol{X} \equiv [\boldsymbol{X}_1', \boldsymbol{X}_2', \ldots, \boldsymbol{X}_T']'$, and $\boldsymbol{Z} \equiv [\boldsymbol{Z}_1', \boldsymbol{Z}_2', \ldots, \boldsymbol{Z}_T']'$. We split $\boldsymbol{Z}$ into two parts, one part containing all non-interest rate data and all observations for interest rates that are not constrained by the ELB, $\boldsymbol{Z}^{NC}$, and another part with the interest rate data constrained at the ELB, $\boldsymbol{Z}^C$.[35] The corresponding elements of $\boldsymbol{X}$ are $\boldsymbol{X}^{NC}$ and $\boldsymbol{X}^C$. Note that, the elements of $\boldsymbol{X}^C$ are all shadow rates that are less than the ELB.

Given a normal distribution for $\boldsymbol{\xi}_0$, it follows that the vectors $\boldsymbol{X}^{NC}$ and $\boldsymbol{\xi} = [\boldsymbol{\xi}_1', \boldsymbol{\xi}_2', \ldots, \boldsymbol{\xi}_T']'$ have a multivariate normal (prior) distribution

$$\begin{bmatrix} \boldsymbol{X}^{NC} \\ \boldsymbol{\xi} \end{bmatrix} \sim N\left(\begin{bmatrix} \boldsymbol{\mu}_X \\ \boldsymbol{\mu}_\xi \end{bmatrix}, \begin{bmatrix} \boldsymbol{V}_{XX} & \boldsymbol{V}_{X,\xi} \\ \boldsymbol{V}_{\xi,X} & \boldsymbol{V}_{\xi,\xi} \end{bmatrix} \right) \tag{25}$$

and we can derive the posterior distribution for $\boldsymbol{\xi}$ conditional on observed interest rates, when the ELB is not binding, as well as all observations for macroeconomic, non-interest-rate vari-

[33] In our application, described in the main body of the paper note that we have a constant $\boldsymbol{\mathcal{C}}_t = \boldsymbol{\mathcal{C}}$.

[34] In the application described above, there are in principle two shadow rates, one associated with the short–term interest rate and one associated with the medium–term yield described; in practice, the ELB constraint has been binding only for the former, however.

[35] Accordingly, $\boldsymbol{Z}^C$ consists solely of observations for interest rates that are equal to ELB.

ables ($\boldsymbol{M}_t$):

$$\boldsymbol{\xi}\big|\left(\boldsymbol{X}^{NC}=\boldsymbol{Z}^{NC}\right)\sim N\left(\hat{\boldsymbol{\mu}}_{\xi},\hat{\boldsymbol{V}}_{\xi,\xi}\right). \tag{26}$$

In general, the posterior moments in (26) are given by

$$\hat{\boldsymbol{\mu}}_{\xi}=\boldsymbol{\mu}_{\xi}+\boldsymbol{K}\left(\boldsymbol{Z}^{NC}-\boldsymbol{\mu}_{X}\right) \qquad \text{with} \quad \boldsymbol{K}=\boldsymbol{V}_{\xi,X}\boldsymbol{V}_{XX}^{-1} \tag{27}$$

$$\text{and} \quad \hat{\boldsymbol{V}}_{\xi,\xi}=\boldsymbol{V}_{\xi,\xi}-\boldsymbol{V}_{\xi,X}\boldsymbol{V}_{XX}^{-1}\boldsymbol{V}_{X,\xi}\,. \tag{28}$$

Typically, these posterior moment matrices will by quite large; $\boldsymbol{\mu}_{\xi}$ is, for example, a vector of length $T\times N_{\xi}=224\times 12=2,688$ in our application, However, the Kalman smoother, adapted for handling missing observations for interest rates when the ELB binds, provides a convenient way to recursively compute the moments in (27) and (28) while recovering the distribution of $\boldsymbol{\xi}$ conditional on observations for $\boldsymbol{Z}^{NC}$. To this point, our procedure amounts to treating the observations for $\boldsymbol{Z}^{C}$ as missing data.[36]

We then note that the information contained in the interest rate data at the ELB is that $\boldsymbol{X}^{C}\leq ELB$ (for every element of $\boldsymbol{X}^{C}$). The posterior distribution of $\boldsymbol{\xi}$, conditional on $\boldsymbol{Z}$, is then

$$\boldsymbol{\xi}|(\boldsymbol{X}=\boldsymbol{Z})\sim TN\left(\hat{\boldsymbol{\mu}}_{\xi},\hat{\boldsymbol{V}}_{\xi,\xi};\boldsymbol{X}^{C}\leq ELB\right) \tag{29}$$

where TN stands for a truncated normal such that its density function is

$$\Pr\left(\boldsymbol{\xi}|\boldsymbol{Z}\right)\propto\phi\left(\hat{\boldsymbol{\mu}}_{\xi},\hat{\boldsymbol{V}}_{\xi,\xi}\right)\mathbf{1}(\boldsymbol{X}^{C}\leq ELB) \tag{30}$$

where ϕ is the multivariate normal density function and $\boldsymbol{X}^{C}$ is a shadow rate draw where every element is below the ELB. To sample from the posterior distribution of $\boldsymbol{\xi}$ conditional on all observations in $\boldsymbol{Z}$, we first draw $\boldsymbol{\xi}$ from $\Pr\left(\boldsymbol{\xi}\big|\boldsymbol{X}^{NC}=\boldsymbol{Z}^{NC}\right)$. We then reject draws until we find a draw that satisfies the requirement that $\boldsymbol{X}^{C}\leq ELB$ for every element. Rejection

[36] Alternatively, Chan and Jeliazkov (2009) describe ways to efficiently compute the moments in (27) and (28) based on sparse matrices that exploit the state space structure inherent in (25).

sampling is thus done on an entire draw of $\boldsymbol{\xi}$, which corresponds to an entire draw of the time series for $\boldsymbol{\xi}_t$.

In our baseline framework, lagged values of $\boldsymbol{\xi}_t$ appear as explanatory variables and are not censored. A straightforward extension is to incorporate a given number of p lags of $\boldsymbol{Z}_t$, which includes interest rate data that can be constrained the ELB. In this case, we change (21) to be

$$\boldsymbol{\xi}_t = \boldsymbol{\mathcal{A}}_t \boldsymbol{\xi}_{t-1} + \boldsymbol{\mathcal{F}}_t \boldsymbol{\zeta}_{t-1} + \boldsymbol{\mathcal{B}}_t \boldsymbol{\varepsilon}_t \tag{31}$$

$$\boldsymbol{\zeta}_{t-1} \equiv [\boldsymbol{Z}'_{t-1}, \boldsymbol{Z}'_{t-2}, \ldots, \boldsymbol{Z}'_{t-p}]' \tag{32}$$

where $\{\boldsymbol{\mathcal{F}}_t\}_{t=1}^T$ are conformable matrices that are known. The posterior of $\boldsymbol{\xi}$ can be constructed exactly as in our baseline model, treating $\{\boldsymbol{\zeta}_{t-1}\}_{t=1}^T$ as exogenous in every period because the rejection step will ensure that the sampled values of $\boldsymbol{\xi}$ are consistent with $\boldsymbol{\zeta}_{t-1}$ for all t. For comparison, 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 $\boldsymbol{A}_t = \boldsymbol{0}$. A notable difference in the posterior simulation of the model is that the truncated distributions in Iwata and Wu (2006) and Nakajima (2011) can be cast as period–by–period truncated normals. By contrast, our posterior estimates require rejection sampling on an entire time series draw of $\boldsymbol{\xi}$.

B Priors and Posterior Sampling

MCMC estimates of the model are obtained from a Gibbs sampler. The sampler is run multiple times with different starting values and convergence is assessed with the scale reduction test of Gelman et al. (2003).[37] For each run, $10,000$ draws are stored after a burn-in period of $10,000$ draws; the post-burnin draws from each run are then merged.

Our models consist of two layers of latent states as well as various parameters; the latent states are:

$$\boldsymbol{\xi}_t = \begin{bmatrix} \bar{\boldsymbol{X}}'_t & \tilde{\boldsymbol{X}}'_t & \tilde{\boldsymbol{X}}'_{t-1} & \ldots & \tilde{\boldsymbol{X}}'_{t-p+1} \end{bmatrix}$$

[37] Specifically, for every model, 10 independent runs for the Gibbs sampler were evaluated; each run initialized with different starting values drawn from the model's prior distribution. Convergence is deemed satisfactory when the scale reduction statistics for every parameter and latent variable are below 1.2; (values close to 1 indicate good convergence).

where $p = 2$ is the lag length of the gap VAR (7), as well as the vector of stochastic log-variances of the gap shocks, $\tilde{\boldsymbol{h}}_t \equiv \log\left(\boldsymbol{\sigma}_t^2\right)$ in (8) and the stochastic log-variance process of shocks to the inflation trend, $\bar{h}_t \equiv \log\left(\sigma_{\bar{\pi},t}^2\right)$ in (4).[38]

The parameter vector of the model comprises the means, persistence and variance-covariances of shocks to $\bar{h}_t$ and $\tilde{\boldsymbol{h}}_t$, see (4) and (8). as well as the variance of shocks to the trend real rate, denoted $\sigma_{\bar{r}}^2$, the transition coefficients of the gap VAR (7), stacked in a vector $\boldsymbol{a}$, and the lower diagonal elements of gap shock loadings $\boldsymbol{B}$ in (7) that can be stacked in a vector denoted $\boldsymbol{b}$. For ease of reference, all parameters are collected in the vector $\boldsymbol{\theta}$. Furthermore, since MCMC estimates of $\bar{h}^T$ and $\tilde{\boldsymbol{h}}^T$ will be obtained from the multi–move filter of Kim et al. (1998), the use of a set of discrete indicator variables, $\boldsymbol{s^T}$, is required to to approximate $\log \boldsymbol{\eta}_{\bar{\pi},t}^2$ and $\log \tilde{\boldsymbol{\eta}}_t^2$ in (4) and (8), respectively, with a mixture of normals. For ease of exposition, we stack the log of the stochastic variances of trend and gap shocks into the vector $\boldsymbol{h}_t$; combining (4) and (8) yields the following vector system:[39]

$$\boldsymbol{h}_t = (\boldsymbol{I} - \boldsymbol{\rho})\,\boldsymbol{\mu} + \boldsymbol{\rho}\boldsymbol{h}_{t-1} + \boldsymbol{\Phi}\boldsymbol{\eta}_t \qquad\qquad \boldsymbol{\eta}_t \sim N(\boldsymbol{0}, \boldsymbol{I}) \tag{33}$$

Conditional on draws for the various parameters, ($\boldsymbol{\theta}$), and log-volatilities $\boldsymbol{h}^T$, we can construct matrices $\boldsymbol{\mathcal{A}}$, $\boldsymbol{\mathcal{B}}$, $\boldsymbol{\mathcal{C}}$. and $\{\boldsymbol{\Sigma}_t\}_{t=1}^T$ and obtain the linear, Gaussian state space system described by equations (21) and (22) in Appendix A.

For the initial values of the latent states, the following priors were used:

$$\boldsymbol{\xi_0} \sim N\left(E(\boldsymbol{\xi}_0), \boldsymbol{\Omega}\right) \qquad \text{with} \quad E(\boldsymbol{\xi}_0) = \begin{bmatrix} \bar{\boldsymbol{\xi}} \\ \boldsymbol{0} \end{bmatrix} \qquad \text{and} \quad \boldsymbol{\Omega} = \begin{bmatrix} \bar{\boldsymbol{\Omega}} & \boldsymbol{0} \\ \boldsymbol{0} & \tilde{\boldsymbol{\Omega}} \end{bmatrix} \tag{34}$$

An uninformative prior for the initial gap levels is obtained by setting $\tilde{\boldsymbol{\Omega}}$ equal to the ergodic variance-covariance matrix of the gaps implied by the VAR (7), evaluated at the time zero draws for the stochastic volatilities, encoded in $\boldsymbol{\Sigma}_0$, for every MCMC draw.[40] The prior for the

[38] In the extended model with stochastic volatility in shocks to the real rate trend, the latter can also be wrapped into $\bar{h}_t$.

[39] Since trend SV in (4) are independent of the gap SV in (8), $\boldsymbol{\Phi}$ is block-diagonal, and both SV blocks can actually be estimated in separate Gibbs steps.

[40] In the case of a VAR(1), the ergodic variance-covariance matrix solves $\tilde{\boldsymbol{\Omega}} = \boldsymbol{A}\tilde{\boldsymbol{\Omega}}\boldsymbol{A}' + \boldsymbol{B}\,\boldsymbol{\Sigma}_0\,\boldsymbol{B}'$ for given values

initial trend levels are set to be consistent with

$$\begin{bmatrix} \bar{\pi}_0 \\ \bar{r}_0 \\ \bar{r}_0^2 \\ \bar{r}_0^5 \\ \bar{r}_0^{10} \end{bmatrix} \sim N \left(\begin{bmatrix} 2.0 \\ 2.0 \\ 2.5 \\ 3.0 \\ 3.5 \end{bmatrix}, 100 \cdot \boldsymbol{I} \right) \tag{35}$$

which implies generally vague prior levels for the various trend components.

The prior for the average levels of the log-variances is normal, with a mean value of $\log(0.1)$ and variance corresponding to the ergodic distribution implied by draws for the shock variance and AR(1) lag coefficients associated with the corresponding log-variance process. For each AR(1) lag coefficient, the prior is $N(0.8, 0.2^2)$, as in Clark and Ravazzolo (2014). For the variance of shocks to SV in the inflation trend, the prior is inverse gamma with 6 degrees of freedom and centered around a mean of 0.2^2, which coincides with the fixed coefficient-value of 0.2 used by Stock and Watson (2007) in their univariate model for inflation. For the vector of shocks to the gap SV vector, $\boldsymbol{\eta}_t$ in (8), the prior is inverse Wishart, centered around a mean of $0.2^2 \cdot \boldsymbol{I}$ and $N + 11$ degrees of freedom where N is the number of gap variables ($N = 6$ in our baseline model).

The parameter governing the variability of real-rate trend shocks, $\sigma_{\bar{r}}^2$, has a univariate inverse–Wishart distribution with 3 degrees of freedom (corresponding to an inverse–gamma distribution with a shape parameter equal to 1.5 degrees of freedom) and is centered around a prior mean of 0.2^2. While this vaguely informative prior embeds the belief that the trend shocks explain only a small share of variations in real rates, it also helps to avoid the pile–up problem — known, for example, from Stock and Watson (1998) and considered in the context of estimating $\sigma_{\bar{r}}^2$ also by Laubach and Williams (2003) as well as Clark and Kozicki (2005) — by keeping posterior draws for the parameter away from zero.

A Minnesota-style prior (centered around a mean of zero) is used for the VAR coefficients $\boldsymbol{a}$, with hyperparameters $\lambda_1 = 0.5$ (own lags) and $\lambda_2 = 0.2$ (cross lags). The prior $\boldsymbol{b}$ is

of $\boldsymbol{A}$, $\boldsymbol{B}$, and $\boldsymbol{\Sigma}_0$.

multivariate normal, $\boldsymbol{b} \sim N(\mathbf{0}, \boldsymbol{I})$. And

The Gibbs sampler is initialized with values drawn from the prior for $\boldsymbol{h}^T$, and $\boldsymbol{\theta}$ and then generates draws from the joint posterior distribution

$$p\left(\boldsymbol{\xi}^T, \boldsymbol{h}^T, \boldsymbol{a}, \boldsymbol{b}, \boldsymbol{\rho}, \boldsymbol{\mu}, \sigma_{\bar{r}}^2, \boldsymbol{\phi}_{\bar{\boldsymbol{\pi}}}^{\mathbf{2}}, \boldsymbol{\Phi}, \boldsymbol{s}^T \middle| \boldsymbol{Z}^{\boldsymbol{T}}\right)$$

by iterating over draws from the following conditional distributions:[41]

1. Draw from $p\left(\boldsymbol{\xi}^T \middle| \boldsymbol{h}^T, \boldsymbol{a}, \boldsymbol{b}, \boldsymbol{\rho}, \boldsymbol{\mu}, \sigma_{\bar{r}}^2, \boldsymbol{\phi}_{\bar{\boldsymbol{\pi}}}^{\mathbf{2}}, \boldsymbol{\Phi}, \boldsymbol{s}^T, \boldsymbol{Z}^T\right)$ with the disturbance smoothing sampler of Durbin and Koopman (2002) and rejection sampling for the shadow rate when the observed nominal short–term rate is at the ELB as described in Appendix A.

2. Draw from $p\left(\boldsymbol{a} \middle| \boldsymbol{\xi}^T, \boldsymbol{h}^T, \boldsymbol{b}, \boldsymbol{\rho}, \boldsymbol{\mu}, \sigma_{\bar{r}}^2, \boldsymbol{\phi}_{\bar{\boldsymbol{\pi}}}^{\mathbf{2}}, \boldsymbol{\Phi}, \boldsymbol{s}^T, \boldsymbol{Z}^T\right) = p\left(\boldsymbol{a} \middle| \boldsymbol{\xi}^T, \boldsymbol{h}^T, \boldsymbol{b}\right)$, a normal conjugate posterior for a VAR with known heteroscedasticity, with rejection sampling to ensure a stationary VAR (Cogley and Sargent, 2005; Clark, 2011)

3. Draw from $p\left(\boldsymbol{b} \middle| \boldsymbol{\xi}^T, \boldsymbol{h}^T, \boldsymbol{a}, \boldsymbol{\rho}, \boldsymbol{\mu}, \sigma_{\bar{r}}^2, \boldsymbol{\phi}_{\bar{\boldsymbol{\pi}}}^{\mathbf{2}}, \boldsymbol{\Phi}, \boldsymbol{s}^T, \boldsymbol{Z}^T\right)$ via recursive Bayesian regressions with known heteroscedasticity to orthogonalize the gap shocks of the VAR in (7).

4. Draw from the inverse-gamma conjugate posteriors for $\sigma_{\bar{r}}^2$ and $\boldsymbol{\Phi}$:

$$p\left(\sigma_{\bar{r}}^2 \middle| \boldsymbol{\xi}^T, \boldsymbol{h}^T, \boldsymbol{a}, \boldsymbol{b}, \boldsymbol{\rho}, \boldsymbol{\mu}, \boldsymbol{\phi}_{\bar{\boldsymbol{\pi}}}^{\mathbf{2}}, \boldsymbol{\Phi}, \boldsymbol{s}^T, \boldsymbol{Z}^T\right) = p\left(\sigma_{\bar{r}}^2 \middle| \boldsymbol{\xi}^T\right)$$

5. Draw from the inverse-gamma conjugate posterior for $\boldsymbol{\phi}_{\bar{\boldsymbol{\pi}}}^{\mathbf{2}}$:

$$p\left(\boldsymbol{\phi}^{\mathbf{2}} \middle| \boldsymbol{\xi}^T, \boldsymbol{h}^T, \boldsymbol{a}, \boldsymbol{b}, \boldsymbol{\rho}, \boldsymbol{\mu}, \sigma_{\bar{r}}^2, \boldsymbol{\Phi}, \boldsymbol{s}^T, \boldsymbol{Z}^T\right) = p\left(\bar{\phi}^2 \middle| \bar{h}^T\right)$$

6. Draw from the inverse-Wishart conjugate posterior for $\boldsymbol{\Phi}$:

$$p\left(\boldsymbol{\Phi} \middle| \boldsymbol{\xi}^T, \boldsymbol{h}^T, \boldsymbol{a}, \boldsymbol{b}, \boldsymbol{\rho}, \boldsymbol{\mu}, \sigma_{\bar{r}}^2, \boldsymbol{\phi}_{\bar{\boldsymbol{\pi}}}^{\mathbf{2}}, \boldsymbol{s}^T, \boldsymbol{Z}^T\right) = p\left(\boldsymbol{\Phi} \middle| \tilde{\boldsymbol{h}}^T\right)$$

[41] For ease of notation, $\bar{h}^T$ and $\tilde{\boldsymbol{h}}^T$ are stacked into $\boldsymbol{h}^T$ unless when the distinction becomes material.

7. Draw from the normal conjugate posterior for $\boldsymbol{\rho}$:

$$p\left(\boldsymbol{\rho}\middle|\boldsymbol{\xi}^T, \boldsymbol{h}^T, \boldsymbol{a}, \boldsymbol{b}, \sigma_{\bar{r}}^2, \boldsymbol{\mu}, \boldsymbol{\phi}_{\bar{\boldsymbol{\pi}}}^{\mathbf{2}}, \boldsymbol{\Phi}, \boldsymbol{s}^T, \boldsymbol{Z}^T\right) = p\left(\boldsymbol{\rho}\middle|\tilde{\boldsymbol{h}}^T\right)$$

(Not that the presence of correlated shocks in (8) necessitates a SUR regression to construct the posterior for $\boldsymbol{\rho}$.)

8. Draw the mixture indicators $\boldsymbol{s}^T$ from:

$$p\left(\boldsymbol{s}^T\middle|\boldsymbol{\xi}^T, \boldsymbol{h}^T, \boldsymbol{a}, \boldsymbol{b}, \sigma_{\bar{r}}^2, \boldsymbol{\phi}_{\bar{\boldsymbol{\pi}}}^{\mathbf{2}}, \boldsymbol{\Phi}, \boldsymbol{\mu}, \boldsymbol{Z}^T\right)$$

9. Draw from $p\left(\boldsymbol{h}^T, \boldsymbol{\mu}\middle|s^T, \boldsymbol{\xi}^T, \boldsymbol{a}, \boldsymbol{b}, \boldsymbol{\rho}, \sigma_{\bar{r}}^2, \boldsymbol{\phi}_{\bar{\boldsymbol{\pi}}}^{\mathbf{2}}, \boldsymbol{\Phi}, \boldsymbol{s}^T, \boldsymbol{Z}^T\right) = p\left(\boldsymbol{h}^T\middle|s^T, \boldsymbol{\xi}^{\boldsymbol{T}}, \boldsymbol{\phi}_{\bar{\boldsymbol{\pi}}}^{\mathbf{2}}, \boldsymbol{\Phi}, \boldsymbol{\rho},\right)$ embedding the disturbance smoothing sampler of Durbin and Koopman (2002) in a linear state space for $\left(\boldsymbol{h}^T, \boldsymbol{\mu}\right)$ as in Kim et al. (1998).[42]

Strictly speaking, this is not a simple Gibbs sampler consisting of steps 1 – 9, but rather a Gibbs-within-Gibbs sampler with the outer Gibbs sampler iterating between

$$p\left(s^T, \boldsymbol{\xi}^T, \boldsymbol{\theta}\middle|\boldsymbol{h}^T, \boldsymbol{Z}^T\right) \quad \text{(thus, a block consisting of steps 1 through 8)}$$

and $$p\left(\boldsymbol{h}^T, \boldsymbol{\mu}\middle|s^T, \boldsymbol{\xi}^T, \boldsymbol{a}, \boldsymbol{b}, \sigma_{\bar{r}}^2, \boldsymbol{\phi}^{\mathbf{2}}, \boldsymbol{Z}^T\right) \quad \text{(step 9),}$$

similar to the discussion by Negro and Primiceri (2015).

C Results with Unemployment Rate Gap

[Figure 8 about here.]

[Figure 9 about here.]

[42] The constant $\boldsymbol{\mu}$ is embedded in the state space as a unit root without shocks, which improves the efficiency of the Gibbs sampler by jointly sampling $\boldsymbol{h}^T$ and $\boldsymbol{\mu}$.

D Results with Alternative Orderings of Gap Variables

Our model embeds a VAR for the gap components with stochastic-volatility in its orthogonalized shocks, see (7). In a constant–variance case, estimation of the VAR would be invariant of the ordering of shocks in the Choleski decomposition implied by the unit–lower–triangular structure of $\boldsymbol{B}$ in (7) However, due to the stochastic volatilities, estimation of the VAR coefficients is not invariant to the ordering of variables in the stochastic—volatility case Negro and Primiceri (2015). This appendix documents the robustness of trend and gap estimates to different variable orderings.

D.1 Using Output Gap as Business Cycle Measure

[Figure 10 about here.]

[Figure 11 about here.]

[Figure 12 about here.]

D.2 Using Unemployment Gap as Business Cycle Measure

[Figure 13 about here.]

[Figure 14 about here.]

[Figure 15 about here.]

E Computation of Likelihood and IRF Using Particle Filter

Computation of the marginal data densities (MDD) and the impulse responses in Sections 3.3 and 5, respectively, are based on particle filter estimates of our model.

For the MDD computations, the particle filter provides an approximation to the likelihood of the data ($\boldsymbol{Z}$) conditional on values for the model's constant parameters (stacked into $\boldsymbol{\theta}$);

see equation (15) in Section 3.3. Using notation introduced in Appendix B, our model has two layers of latent state variables, trends and gaps stacked in $\boldsymbol{\xi}_t$ as well as the stochastic volatilities captured by $\boldsymbol{h}_t$. Conditional on the stochastic volatilities, the dynamics of ξ_t are linear, and we will refer to ξ_t also as "linear state variables." Both layers of latent variables need to be integrated out for the likelihood computation, which is not provided by the MCMC sampler. However, for given parameter values $\boldsymbol{\theta}$, estimation is possible via particle filter. In fact, apart from the ELB constraint, estimation is straightforward with a Rao-Blackwellized particle filter that exploits the conditionally linear structure of the model; see, for example, Creal (2012) and Lopes and Tsay (2011) and the applications in Mertens and Nason (2017) and (Mertens, 2016). For the impulse-responses, we identify shocks from the space of forecast errors to the vector of observables ($\boldsymbol{e}_t$) as described in Section 5 which are again not provided by the MCMC sampler but a particle filter. This appendix first describes the by now fairly standard application of an auxiliary particle filter with Rao-Blackwellization and then turns to our handling of the ELB constraint.

E.1 Rao-Blackwellized Particle Filter (Ignoring the ELB)

The particle filter approximates the posterior density of the latent state vector $\boldsymbol{S_t} = \begin{bmatrix} h_t' & \boldsymbol{\xi}_t' \end{bmatrix}$ for given parameters values $\boldsymbol{\theta}$.[43] The priors for the initial values $\boldsymbol{S_0}$ are as described in Appendix B.[44] For now, we are ignoring the ELB constraint at this point and treat X^t as synonymous with observed data Z^t.

At each point in time, indexed by t, the filter tracks a swarm of M "particles", indexed by i, that consist of the stochastic volatilities $\boldsymbol{h}_t^{(i)}$ and Kalman filtered estimates of the linear state that condition on the particles history of $\boldsymbol{h}_t^{(i)}$; the particle's Kalman filtering distribution of the linear states and conditional on data X^t is characterized by mean vector $\boldsymbol{\xi}_{\boldsymbol{t|t}}^{\boldsymbol{(i)}}$ and variance-covariance matrix $\boldsymbol{\Psi}_{\boldsymbol{t|t}}^{\boldsymbol{(i)}}$.

[43] As described in the main text, the MDD computation involves separate evaluation of the particle–filter generated likelihood conditional on each draw $\boldsymbol{\theta}$ obtained from the MCMC sampler applied to the full data sample. For the IRF, we keep parameters fixed at their posterior medians estimated by MCMC over the full data sample.

[44] As before, the ergodic variance of the gaps components in $\boldsymbol{\xi}_0$, depends on the parameter vector $\boldsymbol{\theta}$; the gap variance is recomputed based on the parameter vector used when evaluating the particle filter.

We utilize an auxiliary particle filter as described by Lopes and Tsay (2011). First introduced byPitt and Shephard (1999), the auxiliary particle filter (APF) is a refinement of the bootstrap filter. While the bootstrap filter propagates particles from one period to the next based on their prior distribution, the APF seeks to adapt new particles based on the likelihood they imply for the data. Before turning to the unique steps of the auxiliary particle filtering steps, we describe the standard bootstrap algorithm.

The filter begins by generating M initial particles $h_0^{(i)}$, $\boldsymbol{\xi}_{\boldsymbol{t|t}}^{\boldsymbol{(i)}}$ and $\boldsymbol{\Psi}_{\boldsymbol{t|t}}^{\boldsymbol{(i)}}$ from the their respective priors described in Appendix B. For $t = 1, \ldots, T$, the filter repeats the following steps:

1. For $i = 1, \ldots, M$ draw new particles $\boldsymbol{h}_t^{(i)}$ based on its prior conditional on $\boldsymbol{h}_{t-1}^{(i)}$ implied by (33) and construct the corresponding diagonal matrix of log-volatilities $\boldsymbol{\Sigma}_t^{1/2,(i)}$ and $\boldsymbol{\mathcal{B}}_t^{(i)} = \boldsymbol{\mathcal{B}}\, \boldsymbol{\Sigma}_t^{1/2,(i)}$.

2. For each particle i, engage the Kalman filter for (21) and (22) to compute:

$$\boldsymbol{\Psi}_{\boldsymbol{t|t-1}}^{\boldsymbol{(i)}} = \boldsymbol{\mathcal{A}}_t\, \boldsymbol{\Psi}_{\boldsymbol{t-1|t-1}}^{\boldsymbol{(i)}}\, \boldsymbol{\mathcal{A}}_t' + \boldsymbol{\mathcal{B}}_{\boldsymbol{t}}^{\boldsymbol{(i)}} \boldsymbol{\mathcal{B}}_{\boldsymbol{t}}^{\boldsymbol{(i)}\prime} \tag{36}$$

$$\boldsymbol{\Omega}_{\boldsymbol{t|t-1}}^{\boldsymbol{(i)}} = \boldsymbol{\mathcal{C}}_{\boldsymbol{t}}\, \boldsymbol{\Psi}_{\boldsymbol{t|t-1}}^{\boldsymbol{(i)}}\, \boldsymbol{\mathcal{C}}_{\boldsymbol{t}}{}' \tag{37}$$

$$\boldsymbol{e}_{\boldsymbol{t}}^{\boldsymbol{(i)}} = \boldsymbol{X}_{\boldsymbol{t}} - \boldsymbol{\mathcal{C}}_{\boldsymbol{t}}\, \boldsymbol{\mathcal{A}}_t\, \boldsymbol{\xi}_{\boldsymbol{t-1|t-1}}^{\boldsymbol{(i)}} \tag{38}$$

$$l_t^{(i)} = -\frac{1}{2}\left\{\log\left(2\cdot\pi\right)\cdot N_x^* + \log\left|\boldsymbol{\Omega}_{\boldsymbol{t|t-1}}^{\boldsymbol{(i)}}\right|_+ + \boldsymbol{e}_{\boldsymbol{t}}^{\boldsymbol{(i)}\prime}\left(\boldsymbol{\Omega}_{\boldsymbol{t|t-1}}^{\boldsymbol{(i)}}\right)^+ \boldsymbol{e}_{\boldsymbol{t|t-1}}^{\boldsymbol{(i)}}\right\} \tag{39}$$

$$\boldsymbol{K}_{\boldsymbol{t}}^{\boldsymbol{(i)}} = \boldsymbol{\Psi}_{\boldsymbol{t|t-1}}^{\boldsymbol{(i)}}\, \boldsymbol{\mathcal{C}}_{\boldsymbol{t}}{}'\, \left(\boldsymbol{\Omega}_{\boldsymbol{t|t-1}}^{\boldsymbol{(i)}}\right)^+ \tag{40}$$

$$\boldsymbol{\xi}_{\boldsymbol{t|t}}^{\boldsymbol{(i)}} = \boldsymbol{\xi}_{\boldsymbol{t|t-1}}^{\boldsymbol{(i)}} + \boldsymbol{K}_{\boldsymbol{t}}^{\boldsymbol{(i)}} \boldsymbol{e}_{\boldsymbol{t}}^{\boldsymbol{(i)}} \tag{41}$$

$$\boldsymbol{\Psi}_{\boldsymbol{t|t}}^{\boldsymbol{(i)}} = \boldsymbol{\Psi}_{\boldsymbol{t|t-1}}^{\boldsymbol{(i)}} - \boldsymbol{\Psi}_{\boldsymbol{t|t-1}}^{\boldsymbol{(i)}}\, \boldsymbol{\mathcal{C}}_{\boldsymbol{t}}{}'\, \left(\boldsymbol{\Omega}_{\boldsymbol{t|t-1}}^{\boldsymbol{(i)}}\right)^+ \boldsymbol{\mathcal{C}}_{\boldsymbol{t}}\, \boldsymbol{\Psi}_{\boldsymbol{t|t-1}}^{\boldsymbol{(i)}}{}' \tag{42}$$

In light of the possibility of missing data—encoded as elements of $\boldsymbol{X_t}$ fixed at zero and a rank deficient $\boldsymbol{\mathcal{C}_t}$—note that N_x^* is the number of actual observations in $\boldsymbol{X_t}$, corresponding to the number of non-zero rows of $\boldsymbol{\mathcal{C}_t}$, and $|\cdot|_+$ and $\cdot^+$ denote the pseudo-determinant and pseudo-inverse operators, respectively.

3. Compute the particle weights

$$w_t^{(i)} = \frac{\exp\left(l_t^{(i)}\right)}{\sum_i^M \exp\left(l_t^{(i)}\right)}.$$

The filtered distribution of $\boldsymbol{h}_t$ is approximated by the discrete distribution of particle draws $\boldsymbol{h}_t^{(i)}$ using the pdf described by $w_t^{(i)}$. The associated filtered distribution of $\boldsymbol{\xi_t}$ is approximated by a mixture of normals $N\left(\boldsymbol{\xi}_{\boldsymbol{t|t}}^{\boldsymbol{(i)}}, \boldsymbol{\Psi}_{\boldsymbol{t|t}}^{\boldsymbol{(i)}}\right)$ with weights $w_t^{(i)}$. (For this purpose, the values for $\boldsymbol{h}_t^{(i)}$, $\boldsymbol{\xi}_{\boldsymbol{t|t}}^{\boldsymbol{(i)}}$ and $\boldsymbol{\Psi}_{\boldsymbol{t|t}}^{\boldsymbol{(i)}}$ are stored before the resampling described in the next step.)

4. For $t < T$ prepare the next iteration by applying systematic resampling to the particles $\boldsymbol{h}_t^{(i)}$, $\boldsymbol{\xi}_{\boldsymbol{t|t}}^{\boldsymbol{(i)}}$ and $\boldsymbol{\Psi}_{\boldsymbol{t|t}}^{\boldsymbol{(i)}}$ based on the particle weights $w_t^{(i)} = \frac{\exp\left(l_t^{(i)}\right)}{\sum_i^M \exp\left(l_t^{(i)}\right)}$.

The likelihood of the date t observation, conditional on parameters and the previous history of observations is estimated by averaging over the likelihoods of each particle:[45]

$$p\left(\boldsymbol{X_t}|\boldsymbol{X^{t-1}}, \boldsymbol{\theta^*}\right) \propto \frac{1}{M}\sum_{i=1}^{M} \exp\left(l_t^{(i)}\right) \tag{43}$$

The log-likelihood, which corresponds also to the log-predictive score for given parameter values $\boldsymbol{\theta^*}$ (Geweke and Amisano, 2010; Creal et al., 2010), is then given by

$$\mathcal{L}\left(\boldsymbol{X^T}|\boldsymbol{\theta^*}\right) = \sum_{t=1}^{T} \log\left\{p\left(\boldsymbol{X_t}|\boldsymbol{X^{t-1}}, \boldsymbol{\theta^*}\right)\right\} \tag{44}$$

The bootstrap filter described above generates new particles for item t merely by propagating stochastic volatilities based on their prior values, $\boldsymbol{h}_{t-1}^{(i)}$ and the law of motion (33) but without regard for the likelihood they will attract at t. The APF seeks to adapt new particle draws $\boldsymbol{h}_t^{(i)}$ to data at t. To do so, we employ an algorithm described byLopes and Tsay (2011) for Rao-Blackwellized particle filters that performs two resampling steps. First, before simulating new particles $\boldsymbol{h}_{t-1}^{(i)}$ from (33), $t-1$ particles are resampled based on the particle weights implied at t when using the auxiliary particles $\boldsymbol{h}_t^{(i),*} = \boldsymbol{\mu} + \boldsymbol{\rho}\left(\boldsymbol{h}_{t-1}^{(i)} - \boldsymbol{\mu}\right)$.[46] De-

[45] Since particles get reweighted at every step, the simple average is appropriate, see, for example, Creal (2012).
[46] The auxiliary particles are thus generated at the means implied by the prior particles.

noting the weights implied by the auxiliary particles $\tilde{w}_t^{(i)}$, (non-normalized) particle weights are then given by $\hat{w}_t^{(i)} = l_t^{(i)}/\tilde{w}_t^{(i)}$ where $l_t^{(i)}$ continue to describe the likelihood contributions in (39) generated by the particles obtained from propagating $\boldsymbol{h}_{t-1}^{(i)}$ after the auxiliary reweighting. Further details are provided by Lopes and Tsay (2011).

E.2 Particle Filter with ELB constraint

The Rao-Blackwellized (RB) particle filter exploits the conditionally linear structure of our unobserved components, which enables an analytic description of posterior for the linear states $\boldsymbol{\xi}_t$ conditional on particle draws for the stochastic volatilities. For data histories where the ELB has not been binding, we have $X^t = Z^t$ and can directly employ the RB-APF described above. However, once the ELB binds, the shadow rate becomes a latent variable, and thus $X_t \neq Z_t$, and the posterior for $\boldsymbol{\xi}_t$ is characterized by a truncated normal as described in Appendix A. When the ELB binds, we need to (partially) abandon the Rao-Blackwellization and proceed as described below. Specifically, for the MDD computation, we add draws for $\boldsymbol{\xi}_t^{(i)}$ to the particle vector that satisfies the truncation constraint $s_t = \boldsymbol{c}_s \boldsymbol{\xi}_t^{(i)} \leq ELB$ exactly. For the IRF, we employ a slightly different modification that approximates the truncated distribution for $\boldsymbol{\xi}_{t|t}^{(i)}$ while adhering to the filtering framework that underpins the Wold decomposition in the space of $\boldsymbol{X}_t$ described in Section 5.

For the MDD computations, the ELB is enforced exactly as follows: Consider first the case where the ELB has not been binding for Z_{t-1} but binds for Z_t. In this case, we inherit particles $\boldsymbol{\xi}_{t-1|t-1}^{(i)}$ and $\boldsymbol{\Psi}_{t-1|t-1}^{(i)}$. Analogously to the MCMC algorithm described in Appendix A, these priors for $\boldsymbol{\xi}_t$ can be updated to $\boldsymbol{\xi}_{t|t}^{NC,(i)}$ and $\boldsymbol{\Psi}_{t|t}^{NC,(i)}$ by conditioning only on the elements of Z_t for which the ELB did not bind. To account for the ELB that binds at t, we then sample particles $\boldsymbol{\xi}_t^{(i)}$ from $N\left(\boldsymbol{\xi}_{t|t}^{NC,(i)}, \boldsymbol{\Psi}_{t|t}^{NC,(i)}\right)$ that are consistent with $s_t \leq ELB$. The particle-filtering algorithm described above can then simply be amended by replacing the prior particles $\left(\boldsymbol{\xi}_{t|t-1}^{(i)}, \boldsymbol{\Psi}_{t|t-1}^{(i)}\right)$ by the draw $\boldsymbol{\xi}_t^{(i))}$ as well as a matrix of zeros before evaluating the time t Kalman filtering steps described above.

Furthermore, we amend the likelihood contributions $l_t^{(i)}$ in (39) by adding the probability

of the shadow rate being below the ELB conditional on all other elements of Z_t, which is straightforward to compute based on the normal CDF for s_t implied by $\boldsymbol{\xi}_{t|t}^{NC,(i)}$, and $\boldsymbol{\Psi}_{t|t}^{NC,(i)}$.

This add factor accounts for the information contained in observing $i_t = ELB$ by tilting particle weights towards particles that place higher likelihood on the ELB binding.

Specifically, (39) is replaced by

$$l_t^{(i)} = -\frac{1}{2}\left\{\log\left(2\cdot\pi\right)\cdot N_x^* + \log\left|\boldsymbol{\Omega}_{\boldsymbol{t|t-1}}^{\boldsymbol{(i)}}\right|_+ + \boldsymbol{e}_{\boldsymbol{t}}^{\boldsymbol{(i)}\prime}\left(\boldsymbol{\Omega}_{\boldsymbol{t|t-1}}^{\boldsymbol{(i)}}\right)^+\boldsymbol{e}_{\boldsymbol{t|t-1}}^{\boldsymbol{(i)}}\right\} + \log\left(\text{Prob}\left(s_t \leq ELB|\boldsymbol{\xi}_{t|t}^{NC,(i)}, \boldsymbol{\Psi}_{t|t}^{NC,(i)}\right)\right) \quad (45)$$

For the IRF computation, we compute the time-varying VMA $\boldsymbol{\mathcal{D}}_t$ defined in (18) as well as the innovation variance $\boldsymbol{\Omega}_{t|t-1}$ and the associated structural VMA $\boldsymbol{D}_t^{m,*}$ separately for each particle. Impulse responses are then reported as particle–weighted averages.

As discussed in Section 5, the key driver for time–variation in the VMA is time–varying uncertainty about trend and gap shocks. When abandoning the Rao-Blackwellization at the ELB, as described above, no uncertainty remains about trend and gap levels *per particle* (except for current shocks) since each particle conditions on a specific draw $\boldsymbol{\xi}_{t-1}^{(i)}$ before propagating. In order to maintain uncertainty about trend and gap levels per particle for the purpose of the IRF computations, we approximate the ELB constraint as follows: When the ELB binds at t, we generate $N\left(\boldsymbol{\xi}_{t|t}^{NC,(i)}, \boldsymbol{\Psi}_{t|t}^{NC,(i)}\right)$ as before. Instead of drawing a specific realization of $\boldsymbol{\xi}_t$ that obeys the ELB constraint, we approximate the truncated normal for $\boldsymbol{\xi}_t^{(i)}$ implied by $TN\left(\boldsymbol{\xi}_{t|t}^{NC,(i)}, \boldsymbol{\Psi}_{t|t}^{NC,(i)}|s_t \leq ELB\right)$ by a normal distribution $N\left(\boldsymbol{\xi}_{t|t}^{*,(i)}, \boldsymbol{\Psi}_{t|t}^{*,(i)}\right)$ where $\boldsymbol{\xi}_{t|t}^{*,(i)}$ and $\boldsymbol{\Psi}_{t|t}^{*,(i)}$ are mean and variance of the truncated normal. Estimates of the model's latent variables obtained with this approximation are close to this obtained from the previously described particle filter where the ELB is enforced exactly.

F Computation of Predictive Densities

In order to derive interest-rate forecasts that conform to the ELB (and other data in $\boldsymbol{Z}_t$), we first proceed by characterizing the predictive density for the shadow rate. Forecasts for ac-

tual rates can then be computed by integrating over the censored shadow–rate density. The shadow rate is included in the non-censored vector of variables $\boldsymbol{X}_t$ described in Appendix A. Apart from handling the truncation issues related to the ELB, our approach is fairly standard, building, for example, on the work by Geweke and Amisano (2010), Christoffel et al. (2010), and Warne et al. (2015). Given the truncation issues for interest rates and the fat tails introduced into the predictive density by the stochastic volatility specification, we have chosen to compute the predictive density based on the mixture of normals that is implied by the draws from our MCMC sampler, instead of approximating the predictive density solely based on its first two moments, treating the predictive density as a normal distribution, as has been done, for example, by Adolfson et al. (2007) in the case of linearized, constant-parameter DSGE models.

In order to compute the predictive density for $\boldsymbol{Z}_{t+h}$ jumping off data at time t, we first employ the MCMC sampler described in Appendix B to re-estimate all model parameters and latent variables ($\boldsymbol{\theta}$, $\boldsymbol{\xi}^t$ and $\boldsymbol{h}^t$) conditional on data available through time t. Draws from this MCMC sampler will henceforth be indexed by k.

Conditional on draws $(\boldsymbol{\xi}_t^k, \boldsymbol{h}_t^k, \boldsymbol{\theta}^k)$, it is straightforward to compute the predictive mean for uncensored variables:

$$E\left(\boldsymbol{X}_{t+h}|\boldsymbol{\xi}_t^k, \boldsymbol{h}_t^k, \boldsymbol{\theta}^k\right) = \boldsymbol{C}^k\left(\boldsymbol{\mathcal{A}}^k\right)^h \boldsymbol{\xi}_t^k \tag{46}$$

and the predictive mean, conditional solely on data through t, can then be approximated by averaging over the means derived from each MCMC draw:

$$E\left(\boldsymbol{X}_{t+h}|\boldsymbol{Z}^t\right) \approx \sum_k E\left(\boldsymbol{X}_{t+h}|\boldsymbol{\xi}_t^k, \boldsymbol{h}_t^k, \boldsymbol{\theta}^k\right) \tag{47}$$

However, in order to characterize the entire predictive density for uncensored variables or even the predictive density for interest rates, which are subject to censoring due to the ELB constraint, we need to account for non-linearities in the distribution for future $\boldsymbol{\xi}_{t:t+h}$ arising from the stochastic volatility shocks in our model. We simulate $J = 100$ trajectories, each indexed by j, of $\boldsymbol{h}_{t:t+h}^{k,j}$ as well as shocks to $\boldsymbol{\xi}_{t:t+h}$ for each draw k from the MCMC sampler.

References

Adolfson, M., J. Linde, and M. Villani (2007). Forecasting Performance of an Open Economy DSGE Model. *Econometric Reviews 26*(2-4), 289–328.

Aruoba, S. B., P. Cuba-Borda, and F. Schorfheide (2017). Macroeconomic dynamics near the zlb: A tale of two countries. *The Review of Economic Studies in press*.

Bauer, M. D. and G. D. Rudebusch (2015). Monetary policy expectations at the zero lower bound. *Federal Reserve Bank of San Francisco Working Paper Series*.

Bauer, M. D. and G. D. Rudebusch (2017a, July). Interest Rates Under Falling Stars. Working Paper Series 2017-16, Federal Reserve Bank of San Francisco.

Bauer, M. D. and G. D. Rudebusch (2017b). Interest rates under falling stars. *Federal Reserve Bank of San Francisco Working Paper*.

Benati, L. (2017). A new approach to estimating the natural rate of interest. *Mimeo, University of Bern*.

Beveridge, S. and C. R. Nelson (1981). 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.

Black, F. (1995). Interest rates as options. *The Journal of Finance 50*(5), 1371–1376.

Bundesbank, D. (2017, September). Monetary policy indicators at the lower bound based on term structure models. In *Monthly Report*, pp. 13–34. Deutsche Bundesbank.

Campbell, J. Y. and R. J. Shiller (1987, October). Cointegration and tests of present value models. *The Journal of Political Economy 95*(5), 1062–1088.

Campbell, J. Y. and R. J. Shiller (1991). Yield spreads and interest rate movements: A bird's eye view. *The Review of Economic Studies 58*(3), 495–514.

Chan, J. C. and A. L. Grant (2015). Pitfalls of estimating the marginal likelihood using the modified harmonic mean. *Economics Letters 131*(C), 29–33.

Chan, J. C. and I. Jeliazkov (2009). Efficient simulation and integrated likelihood estimation in state space models. *International Journal of Mathematical Modelling and Numerical Optimization 1*(1/2), 101–120.

Chan, J. C. and R. Strachan (2014). The zero lower bound: Implications for modelling the interest rate. *Working Paper*.

Chen, R. and J. S. Liu (2000). Mixture Kalman filters. *Journal of the Royal Statistical Society Series B 62*(3), 493–508.

Christiano, L., M. Eichenbaum, and S. Rebelo (2011). When Is the Government Spending Multiplier Large? *Journal of Political Economy 119*(1), 78–121.

Christiano, L. J., M. Eichenbaum, and C. L. Evans (1999). Monetary policy shocks: What have we learned and to what end? In J. B. Taylor and M. Woodford (Eds.), *Handbook of Macroeconomics*, Number 15 in Handbooks in Economics, Chapter 2. Amsterdam: Elsevier. Volume 1A.

Christoffel, K., A. Warne, and G. Coenen (2010, May). Forecasting with DSGE models. Working Paper Series 1185, European Central Bank.

Clark, T. E. (2011). Real-time density forecasts from bayesian vector autoregressions with stochastic volatility. *Journal of Business & Economic Statistics 29*(3), 327–341.

Clark, T. E., A. Carriero, and M. Massimiliano (2016, October). Measuring Uncertainty and Its Impact on the Economy. Working Paper 1622, Federal Reserve Bank of Cleveland.

Clark, T. E. and S. Kozicki (2005, December). Estimating equilibrium real interest rates in real time. *The North American Journal of Economics and Finance 16*(3), 395–413.

Clark, T. E. and F. Ravazzolo (2014). Macroeconomic forecasting performance under alternative specifications of time-varying volatility. *Journal of Applied Econometrics*. Published online, doi: 10.1002/jae.2379.

Cogley, T., G. E. Primiceri, and T. J. Sargent (2010, January). Inflation-gap persistence in the U.S. *American Economic Journal: Macroeconomics 2*(1), 43–69.

Cogley, T. and T. J. Sargent (2005). Drifts and volatilities: monetary policies and outcomes in the post WWII US. *Review of Economic Dynamics 8*(2), 262–302.

Cogley, T. and T. J. Sargent (2015, October). Measuring Price-Level Uncertainty and Instability in the United States, 18502012. *The Review of Economics and Statistics 97*(4), 827–838.

Creal, D. (2012). A survey of sequential monte carlo methods for economics and finance. *Econometric Reviews 31*(3), 245–296.

Creal, D., S. J. Koopman, and E. Zivot (2010, June/July). Extracting a robust us business cycle using a time-varying multivariate model-based bandpass filter. *Journal of Applied Econometrics 25*(4), 695–719.

Del Negro, M., D. Giannone, M. P. Giannoni, and A. Tambalotti (2017). Safety, liquidity, and the natural rate of interest. *Brookings Papers on Economic Activity*.

Diebold, F. X. and C. Li (2006). Forecasting the term sturcutre of government bond yields. *Journal of Econometrics* (130), 337–364.

Diebold, F. X. and R. S. Mariano (1995, July). Comparing predictive accuracy. *Journal of Business & Economic Statistics 13*(3), 253–63.

Doh, T. and J. Choi (2016). Measuring the Stance of Monetary Policy on and off the Zero Lower Bound. *Economic Review* (Q III), 5–24.

Durbin, J. and S. J. Koopman (2002, September). A simple and efficient simulation smoother for state space time series analysis. *Biometrika 89*(3), 603–615.

Eggertsson, G. B. and P. Krugman (2012). Debt, Deleveraging, and the Liquidity Trap: A Fisher-Minsky-Koo Approach. *The Quarterly Journal of Economics 127*(3), 1469–1513.

Francis, N., L. E. Jackson, and M. T. Owyang (2017, October). How Has Empirical Monetary Policy Analysis Changed After the Financial Crisis? Working Papers 2014-19 (revised), Federal Reserve Bank of St. Louis.

Fuentes-Albero, C. and L. Melosi (2013). Methods for computing marginal data densities from the Gibbs output. *Journal of Econometrics 175*(2), 132–141.

Gagnon, E., B. K. Johannsen, and J. D. Lopez-Salido (2016, September). Understanding the New Normal : The Role of Demographics. Finance and Economics Discussion Series 2016-080, Board of Governors of the Federal Reserve System (U.S.).

Gali, J. and L. Gambetti (2009, January). On the sources of the Great Moderation. *American Economic Journal: Macroeconomics 1*(1), 26–57.

Gavin, W. T., B. D. Keen, A. W. Richter, and N. A. Throckmorton (2015). The zero lower bound, the dual mandate, and unconventional dynamics. *Journal of Economic Dynamics and Control 55*(C), 14–38.

Gelman, A., J. B. Carlin, H. S. Stern, and D. B. Rubin (2003). *Bayesian Data Analysis* (2nd ed.). Chapman and Hall.

Geweke, J. (1999). Using simulation methods for bayesian econometric models: inference, development, and communication. *Econometric reviews 18*(1), 1–73.

Geweke, J. and G. Amisano (2010). Comparing and evaluating bayesian predictive distributions of asset returns. *International Journal of Forecasting 26*(2), 216 – 230. Special Issue: Bayesian Forecasting in Economics.

Gurkaynak, R. S., B. Sack, and J. H. Wright (2007, November). The U.S. Treasury yield curve: 1961 to the present. *Journal of Monetary Economics, Elsevier 54*(8), 2291–2304.

Gust, C., E. Herbst, D. Lpez-Salido, and M. E. Smith (2017, July). The Empirical Implications of the Interest-Rate Lower Bound. *American Economic Review 107*(7), 1971–2006.

Hakkio, C. S. and G. A. Kahn (2014). Evaluating monetary policy at the zero lower bound. *Economic Review* (Q II), 5–32.

Hakkio, C. S. and A. L. Smith (2017). Bond Premiums and the Natural Real Rate of Interest. *Economic Review* (Q I), 5–39.

Hamilton, J. D., E. S. Harris, J. Hatzius, and K. D. West (2015, March). The equilibrium real funds rate: Past, present and future. *Working Paper*.

Herbst, E. and F. Schorfheide (2014, September). Bayesian inference for DSGE models. *mimeo* Board of Governors of the Federal Reserve System.

Hopke, P. K., C. Liu, and D. B. Rubin (2001). Multiple imputation for multivariate data with missing and below-threshold measurements: Time-series concentrations of pollutants in the arctic. *Biometrics 57*(1), 22–33.

Ichiue, H. and Y. Ueno (2013). Estimating term premia at the zero bound: An analysis of Japanese, US, and UK yields. Working Paper Series 2013–E–8, Bank of Japan.

Ireland, P. N. (2007, December). Changes in the Federal Reserve's inflation target: Causes and consequences. *Journal of Money, Credit and Banking 39*(8), 1851–1882.

Iwata, S. and S. Wu (2006). Estimating monetary policy effects when interest rates are close to zero. *Journal of Monetary Economics 53*(7), 1395 – 1408.

Jeffreys, H. (1961). *The Theory of Probability* (3rd ed.). Oxford University Press.

Johannsen, B. K. (2014, May). When are the Effects of Fiscal Policy Uncertainty Large? Finance and Economics Discussion Series 2014-40, Board of Governors of the Federal Reserve System (U.S.).

Jurado, K., S. C. Ludvigson, and S. Ng (2015, March). Measuring Uncertainty. *American Economic Review 105*(3), 1177–1216.

Kass, R. E. and A. E. Raftery (1995). Bayes factors. *Journal of the American Statistical Association 90*(430), 773–795.

Kiley, M. T. (2015). What can the data tell us about the equilibrium real interest rate? *Finance and Economics Disucssion Series* (2015-77).

Kim, D. and K. J. Singleton (2011). Term structure models and the zero bound: An emprirical investigation of Japanese yields. *Mimeo*.

Kim, S., N. Shephard, and S. Chib (1998, July). Stochastic volatility: Likelihood inference and comparison with arch models. *The Review of Economic Studies 65*(3), 361–393.

King, R. G. and A. Kurmann (2002, August). Expectations and the term structure of interest rates: Evidence and implications. *Federal Reserve Bank of Richmond Economic Quarterly (Fall)*, 49–95.

Kozicki, S. and P. Tinsley (2012, 02). Effective use of survey information in estimating the evolution of expected inflation. *Journal of Money, Credit and Banking 44*(1), 145–169.

Kozicki, S. and P. A. Tinsley (2001, June). Shifting endpoints in the term structure of interest rates. *Journal of Monetary Economics 47*(3), 613–652.

Kozicki, S. and P. A. Tinsley (2002). Alternative sources of the lag dynamics of inflation. Research Working Paper RWP 02-12, Federal Reserve Bank of Kansas City.

Krippner, L. (2013). Measuring the stance of monetary policy in zero lower bound environments. *Economics Letters 118*(1), 135–138.

Krippner, L. (2015). *Zero Lower Bound Term Structure Modeling: A practitioner's Guide*. Palgrave Macmillan.

Laubach, T. and J. C. Williams (2003, November). Measuring the natural rate of interest. *The Review of Economics and Statistics 85*(4), 1063–1070.

Laubach, T. A. and J. C. Williams (2015). Measuring the natural rate of interest redux. *Hutchins Center on Fiscal & Monetary Policy at Brookings* (15).

Lombardi, M. J. and F. Zhu (2014, June). A shadow policy rate to calibrate US monetary policy at the zero lower bound. BIS Working Papers 452, Bank for International Settlements.

Lopes, H. F. and R. S. Tsay (2011). Particle filters and bayesian inference in financial econometrics. *Journal of Forecasting 30*, 168–209.

Lubik, T. A. and C. Matthes (2015). Calculating the natural rate of interest: A comparison of two alternative approaches. *Economic Brief* (15-10).

McCracken, M. W. (2000). Robust out-of-sample inference. *Journal of Econometrics 99*(2), 195 – 223.

Mertens, E. (2016, December). Measuring the Level and Uncertainty of Trend Inflation. *The Review of Economics and Statistics 98*(5), 950–967.

Mertens, E. and J. M. Nason (2017, September). Inflation and professional forecast dynamics: An evaluation of stickiness, persistence, and volatility. CAMA Working Papers 2017-60, Centre for Applied Macroeconomic Analysis, Crawford School of Public Policy, The Australian National University.

Nakajima, J. (2011). Monetary policy transmission under zero interest rates: An extended time-varying parameter vector autoregression approach. *Bank of Japan Working Paper Series*.

Negro, M. D. and G. E. Primiceri (2015). Time Varying Structural Vector Autoregressions and Monetary Policy: A Corrigendum. *Review of Economic Studies 82*(4), 1342–1345.

Park, J. W., M. G. Genton, and S. K. Ghosh (2007). Censored time series analysis with autoregressive moving average models. *The Canadian Journal of Statistics / La Revue Canadienne de Statistique 35*(1), 151–168.

Pescatori, A. and J. Turunen (2015, June). Lower for Longer; Neutral Rates in the United States. IMF Working Papers 15/135, International Monetary Fund.

Pitt, M. K. and N. Shephard (1999). Filtering via simulation: Auxiliary particle filters. *Journal of the American Statistical Association 94*(446), 590–599.

Priebsch, M. A. (2013). Computing arbitrage-free yields in multi-factor Gaussian shadow-rate term structure models. Finance and Economics Discussion Series 2013-63, Board of Governors of the Federal Reserve System (U.S.).

Primiceri, G. E. (2005). Time varying structural vector autoregressions and monetary policy. *The Review of Economic Studies 72*(3), 821–852.

Rachel, L. and T. D. Smith (2015). Secular drivers of the global real interest rate. *Bank of England Staff Working Paper* (571).

Rudebusch, G. and L. E. Svensson (1999, December). Policy Rules for Inflation Targeting. In *Monetary Policy Rules*, NBER Chapters, pp. 203–262. National Bureau of Economic Research, Inc.

Sims, C. A. (1980). Macroeconomics and reality. *Econometrica 48*(1), 1–48.

Stock, J. H. and M. W. Watson (1998). Median unbiased estimation of coefficient variance in a time-varying parameter model. *Journal of the American Statistical Association 93*(441), 349–358.

Stock, J. H. and M. W. Watson (2007, February). Why has U.S. inflation become harder to forecast? *Journal of Money, Credit and Banking 39*(S1), 3–33.

Stock, J. H. and M. W. Watson (2012, May). Disentangling the channels of the 2007-2009 recession. NBER Working Papers 18094, National Bureau of Economic Research, Inc.

Summers, L. H. (2014). Us economic prospects: Secular stagnation, hysteresis and the zero lower bound. *Buiness Economics 49*(2), 65–73.

Warne, A., G. Coenen, and K. Christoffel (2015, December). Marginalized predictive likelihood comparisons of linear gaussian state-space models with applications to dsge, dsge-var, and var models. *mimeo* European Central Bank.

West, K. D. (1996). Asymptotic inference about predictive ability. *Econometrica: Journal of the Econometric Society*, 1067–1084.

Wu, J. C. and F. D. Xia (2016). Measuring the macroeconomic impact of monetary policy at the zero lower bound. *Journal of Money, Credit and Banking 48*(2-3), 253–291.

Wu, J. C. and J. Zhang (2016, November). A Shadow Rate New Keynesian Model. NBER Working Papers 22856, National Bureau of Economic Research, Inc.

Figure 1: Shadow Rate Estimates

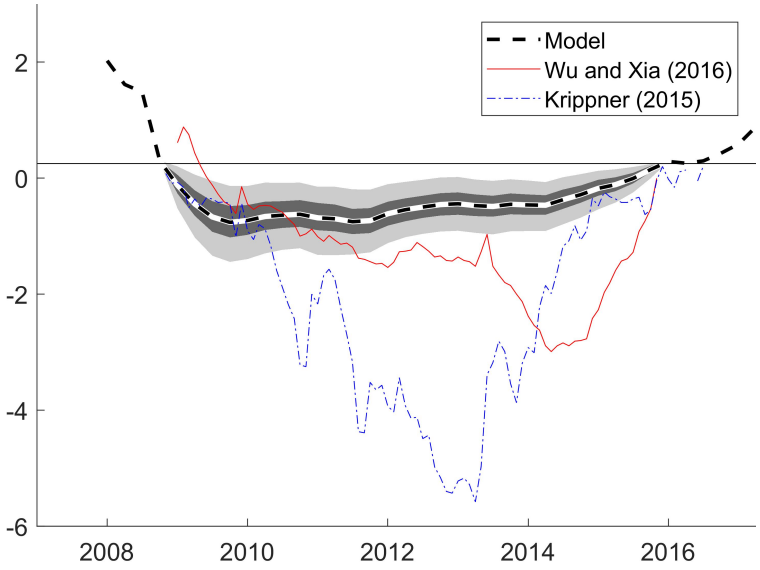

(a) Full-Sample Estimates and Other Estimates

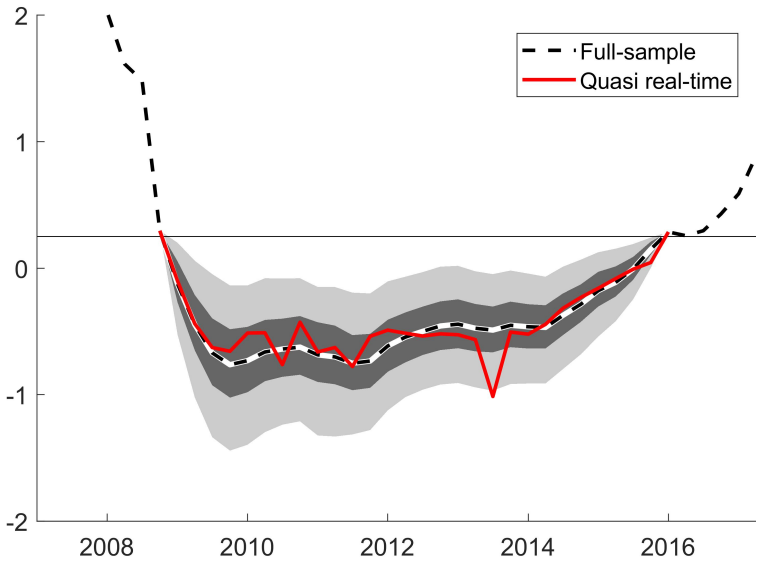

(b) Full-Sample and Quasi Real–Time Estimate

Note: Shaded areas indicate 50 and 90 percent uncertainty bands, dashed lines are posterior means. Results shown in Panel (a) reflect the endpoints of sequentially re-estimating the entire model over growing samples of quarterly observations starting in 1960:Q1, thus reflecting "filtered" estimates of the model's latent variables. Results shown in Panel (b) reflect "smoothed" estimates using all available observations from 1960:Q1 through 2017:Q2. Uncertainty bands reflect the joint uncertainty about model parameters and states.

Figure 2: The Real Rate in the Long Run

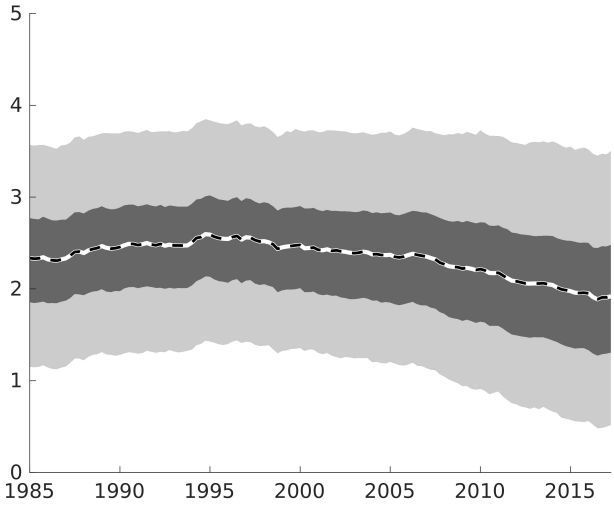

(a) Smoothed Estimates

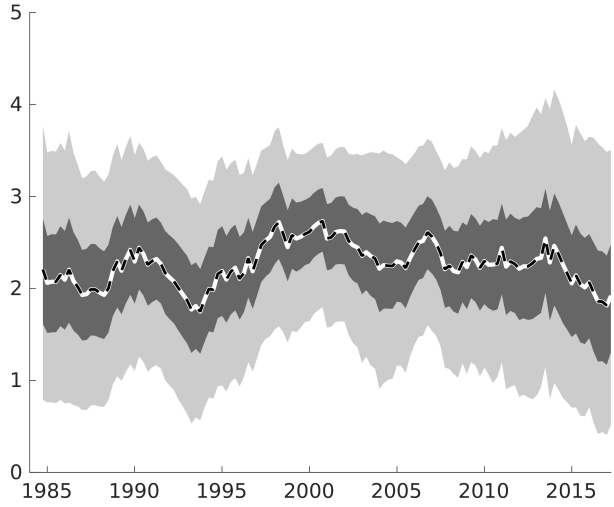

(b) Quasi Real–Time Estimate

Note: Shaded areas indicate 50 and 90 percent uncertainty bands, dashed lines are posterior means. Results shown in Panel (a) reflect the endpoints of sequentially re-estimating the entire model over growing samples of quarterly observations starting in 1960:Q1, thus reflecting "filtered" estimates of the model's latent variables. Results shown in Panel (b) reflect "smoothed" estimates using all available observations from 1960:Q1 through 2017:Q2. Uncertainty bands reflect the joint uncertainty about model parameters and states.

Figure 3: Short–Term Interest Rate Forecasts During the ELB Period

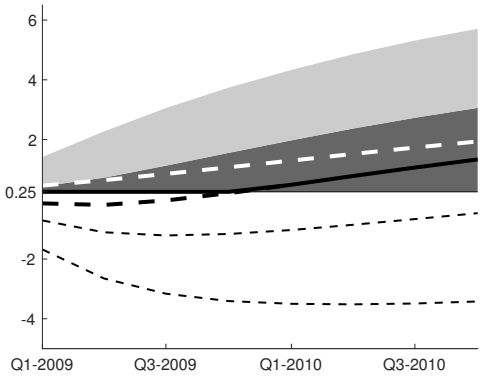

(a) Forecasts from 2008:Q4, Baseline

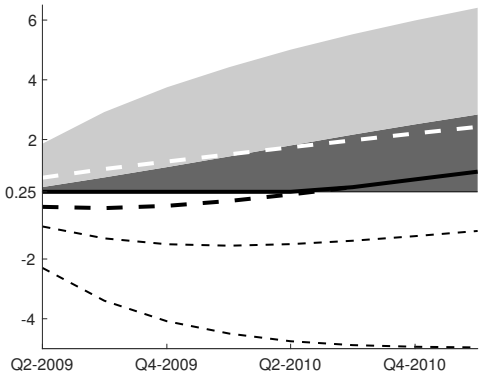

(b) Forecasts from 2009:Q1, Baseline

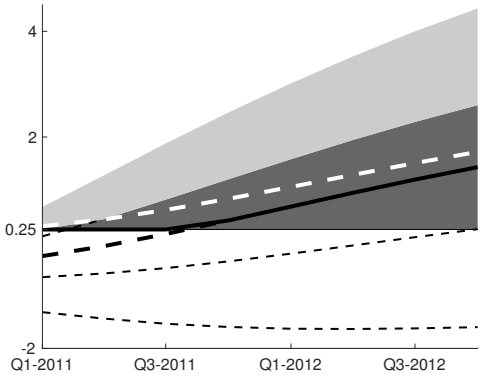

(c) Forecasts from 2010:Q4

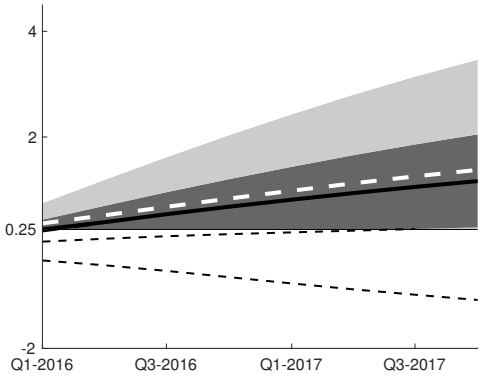

(d) Forecasts from 2015:Q4

Note: Shaded areas indicate 50 and 90 percent credible sets, solid lines are posterior medians, wide dashed lines are posterior means of the projected interest rate. In the left panels, dashed lines less than the ELB are the posterior median and 50 and 90 percent credible sets of the shadow rate.

Figure 4: Time-varying Uncertainty about Future Short–Term Interest Rates

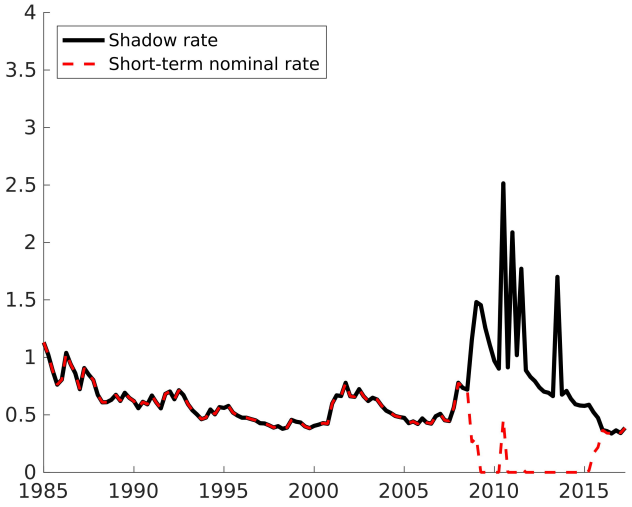

(a) One-quarter ahead uncertainty

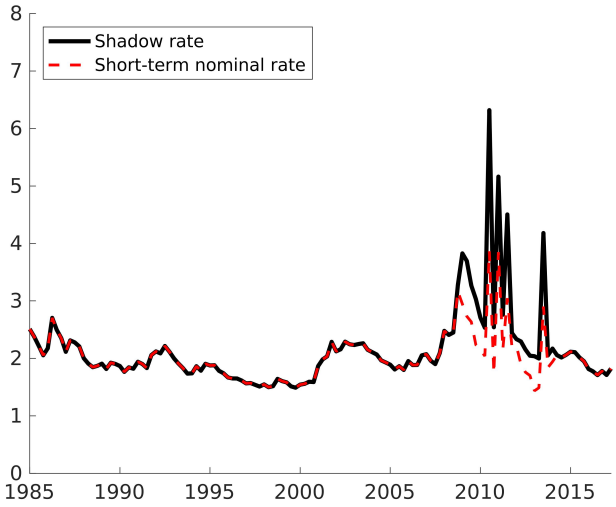

(b) Eight-quarter ahead uncertainty

Note: Forecast uncertainty about future short–term interest rates as measured by the interquartile range of the model's predictive densities for i_t. The predictive densities are re-estimated over growing samples that all start in 1960:Q1. In the baseline model (black solid lines), the predictive density is truncated at the ELB, whereas no constraint is imposed on the predictive density in the alternative model that ignores the ELB (red dashed line).

Figure 5: Long–Term Forecasts of 3-Month Treasury Rate

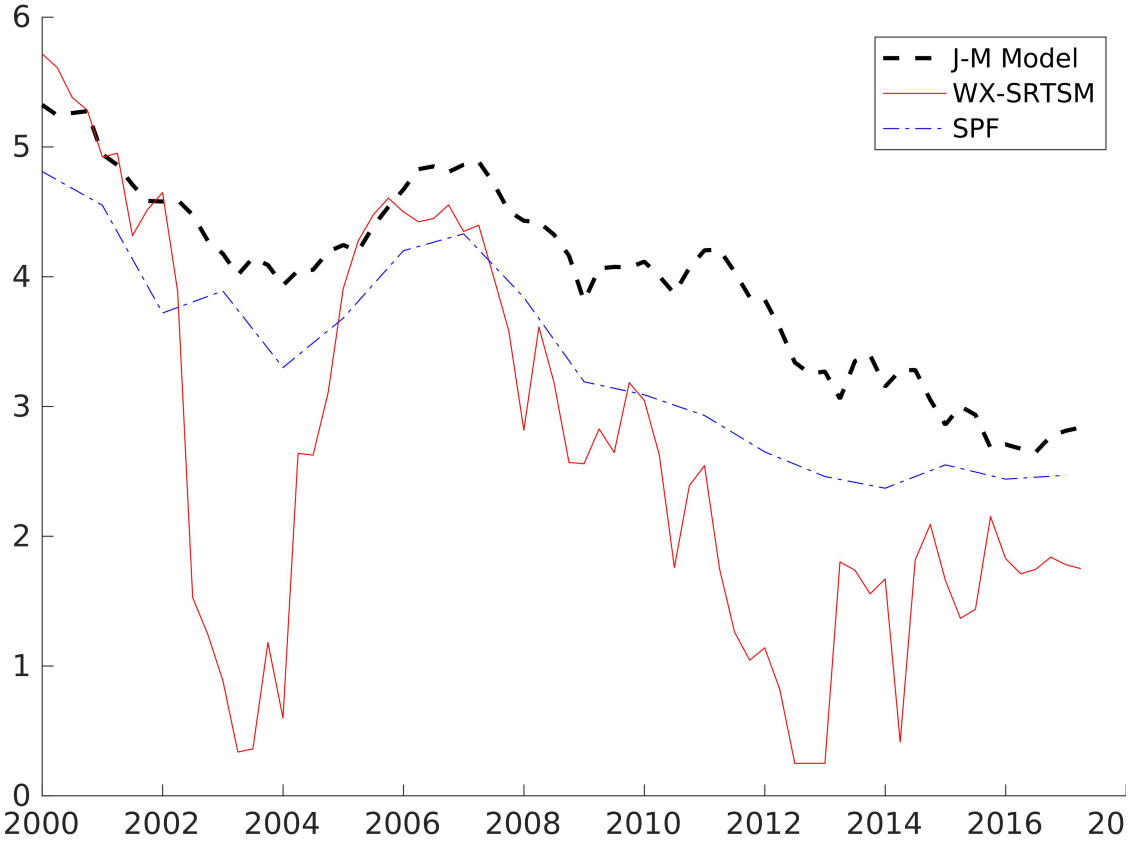

Note: The series labeled "J-M Model" is the median forecast of the 3-month Treasury rate 10-years ahead from the benchmark version of our model with the output gap as the business cycle component. The series labeled "WX-SRTSM" is the median forecast of the 3-month Treasury rate 10-years from the model of Wu and Xia (2016). The series labeled "SPF" is the mean forecast in the Survey of Professional Forecasters of the average value of the 3-month treasury bill over the next 10 years.

Figure 6: Responses to Monetary Policy Shock

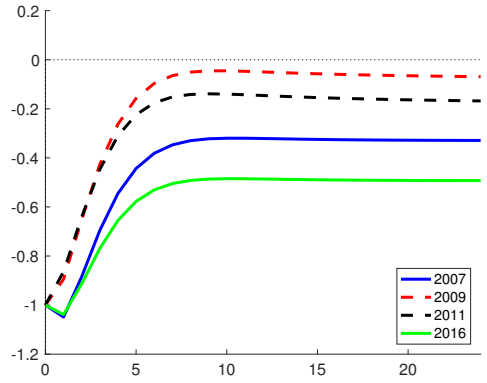

(a) Shadow Rate

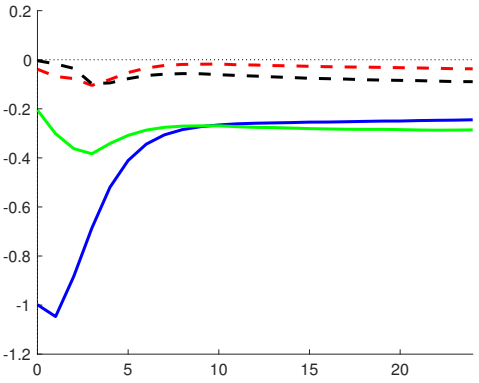

(b) Policy Rate

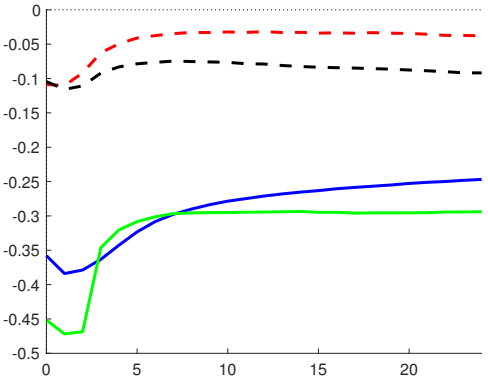

(c) 2-year Yield

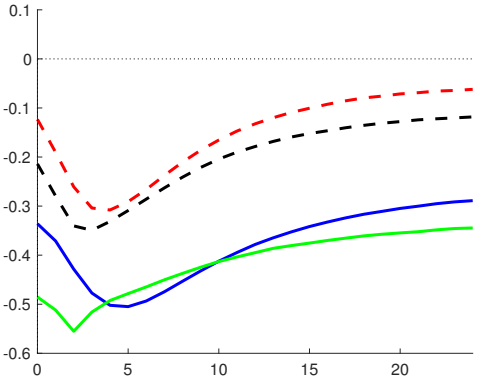

(d) 10-year Yield

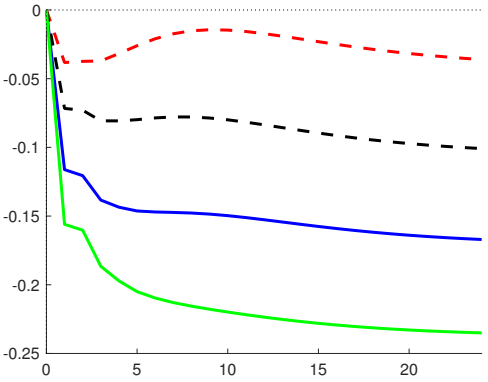

(e) Inflation

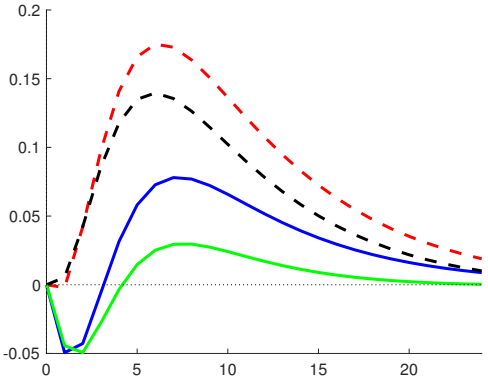

(f) Business Cycle

Note: Responses to monetary policy shocks estimated for 2007:Q4, 2009:Q4, 2011:Q4 as well as 2016:Q4; dashed lines indicate responses at times when the ELB was binding for actual data. Shocks are scaled to generated a one percentage point drop in the shadow rate on impact. All y-axis units are in annualized percentage points for yields and inflation, as well as fluctuations in the business cycle (as measured by the output gap). The x-axis measures time after impact in quarters.

Figure 7: Projected Policy Path Before And After Monetary Policy Shocks

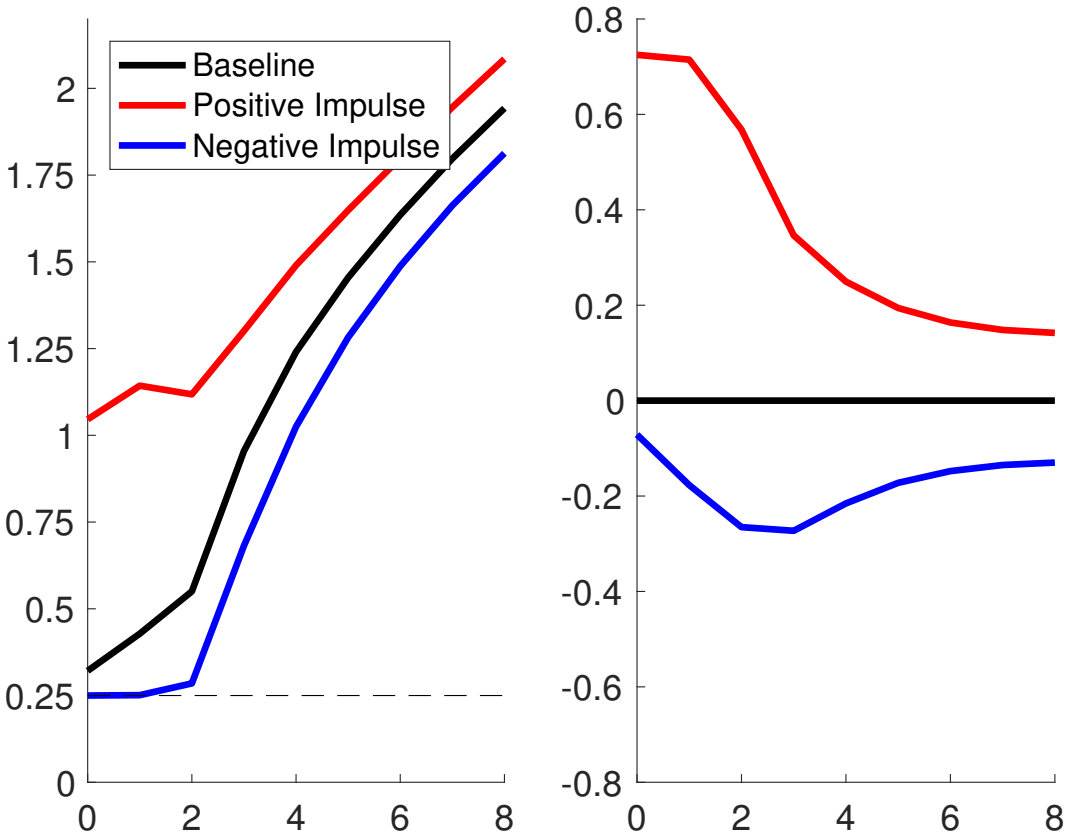

Note: The left panel depicts the baseline path projected for the actual rate at the end of 2015:Q3 (black) as well as two updated forecast trajectories that would result in 2015:Q4 from a surprise increase (red), or decrease (blue), respectively, in the shadow rate by one percentage point. The right panel reports the corresponding impulse responses, computed as differences between the updated forecast trajectory generated by the shadow-rate impulse and the baseline.

Figure 8: Shadow Rate Estimates (Based on Unemployment Rate Gap)

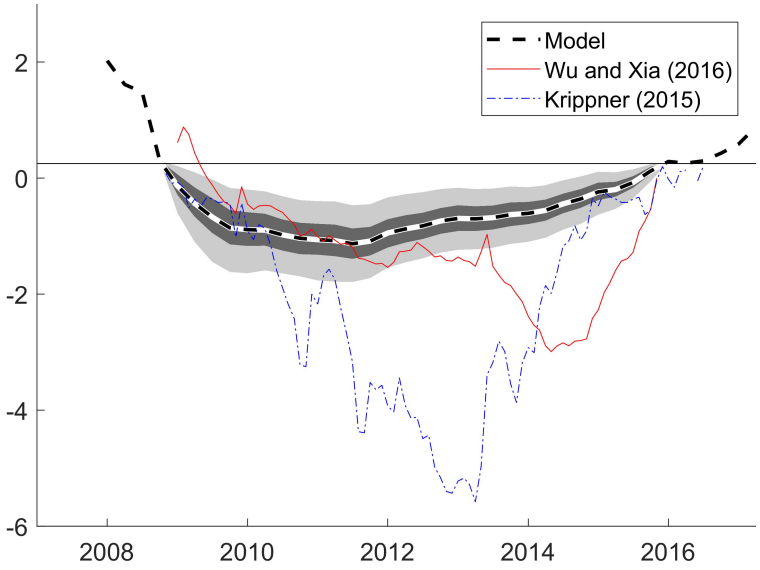

(a) Full-Sample Estimates and Other Estimates

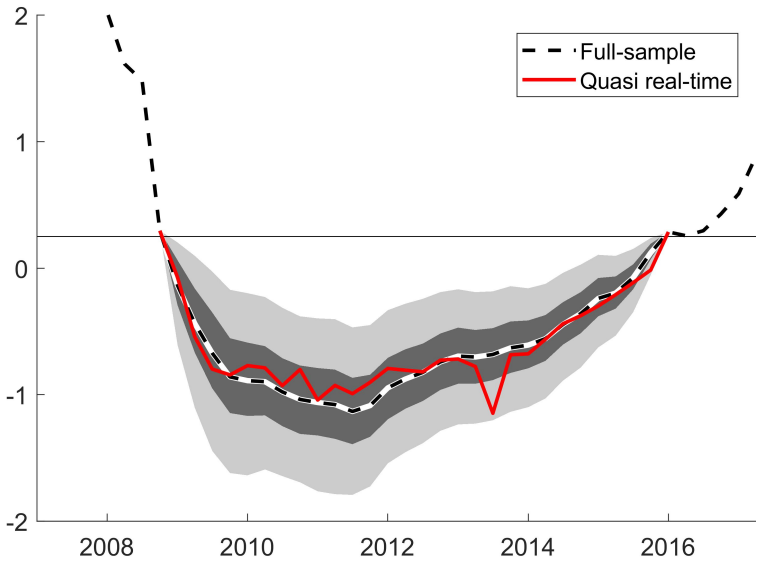

(b) Full-Sample and Quasi Real–Time Estimate

Note: Shaded areas indicate 50 and 90 percent uncertainty bands, dashed lines are posterior means. Results shown in Panel (a) reflect the endpoints of sequentially re-estimating the entire model over growing samples of quarterly observations starting in 1960:Q1, thus reflecting "filtered" estimates of the model's latent variables. Results shown in Panel (b) reflect "smoothed" estimates using all available observations from 1960:Q1 through 2017:Q2. Uncertainty bands reflect the joint uncertainty about model parameters and states.

Figure 9: The Real Rate in the Long Run (Based on Unemployment Rate Gap)

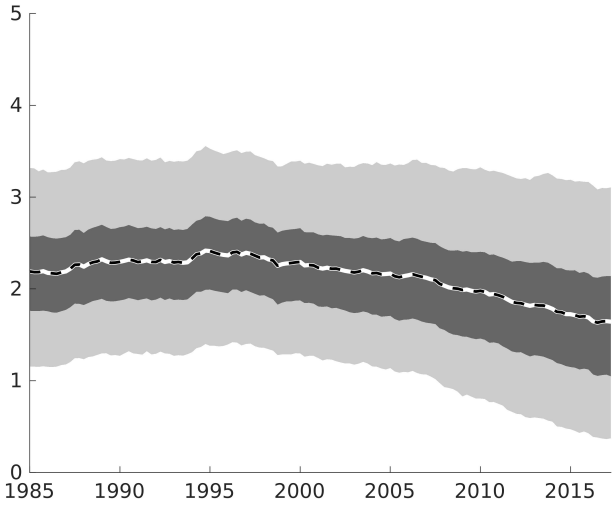

(a) Smoothed Estimates

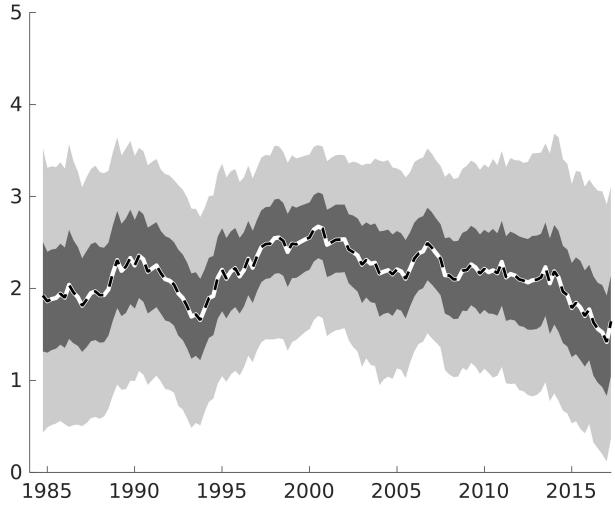

(b) Quasi Real–Time Estimate

Note: Shaded areas indicate 50 and 90 percent uncertainty bands, dashed lines are posterior means. Results shown in Panel (a) reflect the endpoints of sequentially re-estimating the entire model over growing samples of quarterly observations starting in 1960:Q1, thus reflecting "filtered" estimates of the model's latent variables. Results shown in Panel (b) reflect "smoothed" estimates using all available observations from 1960:Q1 through 2017:Q2. Uncertainty bands reflect the joint uncertainty about model parameters and states.

Figure 10: Inflation Trend Estimates with re-ordered Gap-VAR

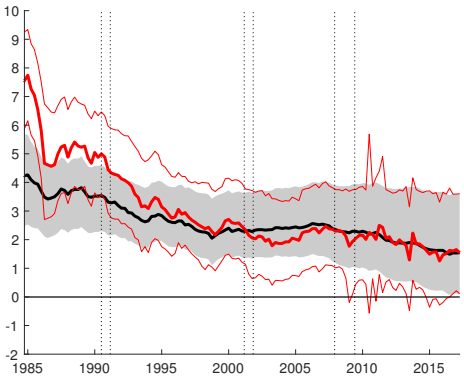

(a) π,c,s,y (baseline)

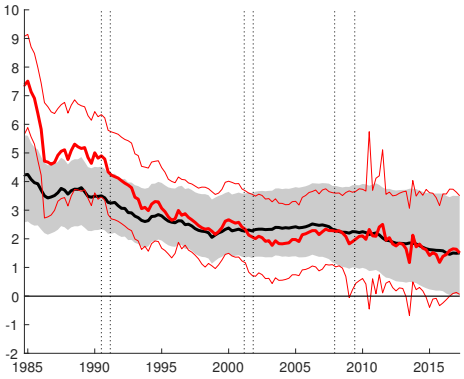

(b) π,s,c,y

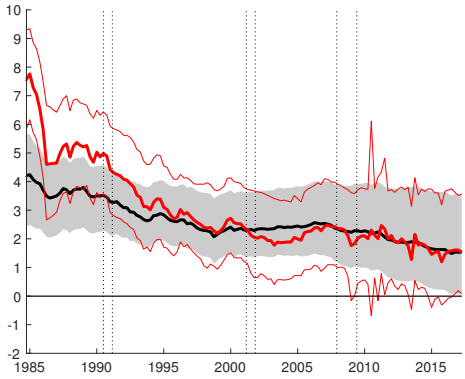

(c) c,π,s,y

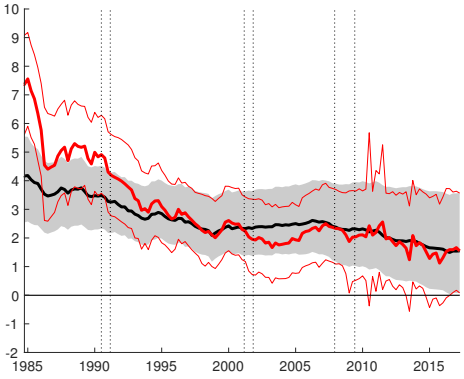

(d) s,c,π,y

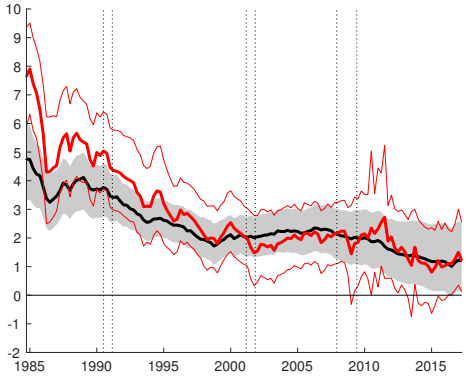

(e) y,s,c,π

Note:Estimates generated from alternative orderings of gap variables in the VAR described in equation (7). Panel 10a depicts baseline estimates; alternative orderings are as indicated above where "y" indicates the block of yield gaps y^2, y^5, y^{10}. Filtered estimates in red, smoothed estimates in black; both surrounded by 90% uncertainty bands (Filtered estimates reflect the endpoints of sequentially re-estimating the entire model over growing samples starting in 1960:Q1. Smoothed estimates use all available observations from 1960:Q1 through 2017:Q2.

Figure 11: Real Rate Trend Estimates with re-ordered Gap-VAR

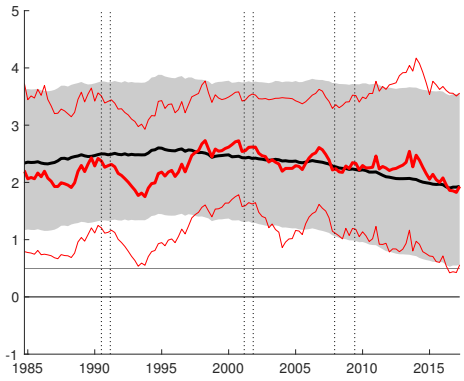

(a) π,c,s,y (baseline)

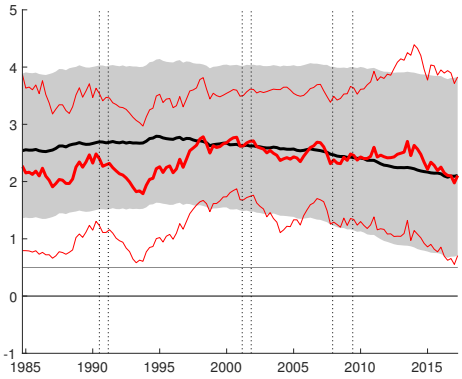

(b) π,s,c,y

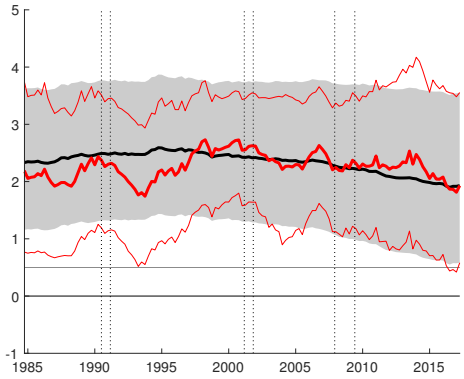

(c) c,π,s,y

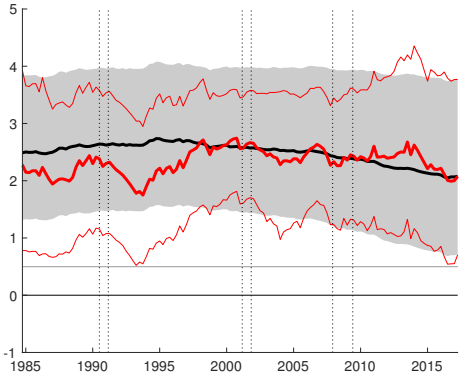

(d) s,c,π,y

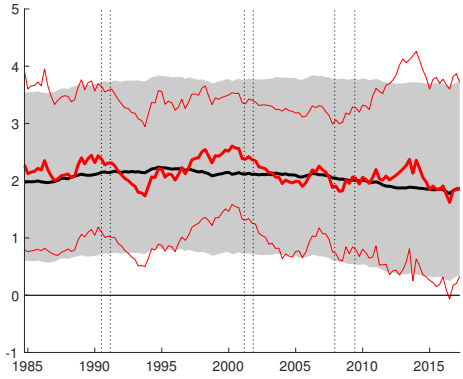

(e) y,s,c,π

Note:Estimates generated from alternative orderings of gap variables in the VAR described in equation (7). Panel 11a depicts baseline estimates; alternative orderings are as indicated above where "y" indicates the block of yield gaps y^2, y^5, y^{10}. Filtered estimates in red, smoothed estimates in black; both surrounded by 90% uncertainty bands (Filtered estimates reflect the endpoints of sequentially re-estimating the entire model over growing samples starting in 1960:Q1. Smoothed estimates use all available observations from 1960:Q1 through 2017:Q2.

Figure 12: Shadow Rate Estimates with re-ordered Gap-VAR

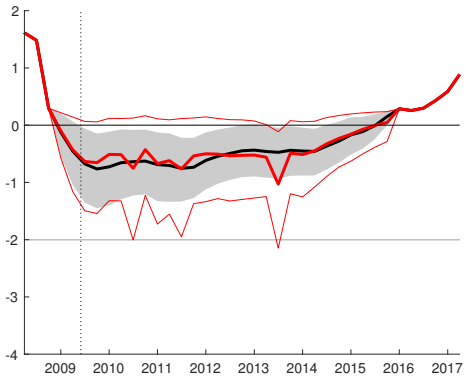

(a) π,c,s,y (baseline)

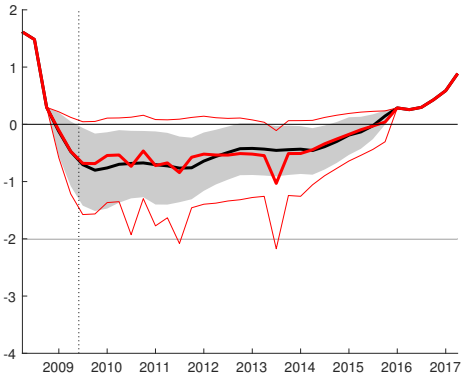

(b) π,s,c,y

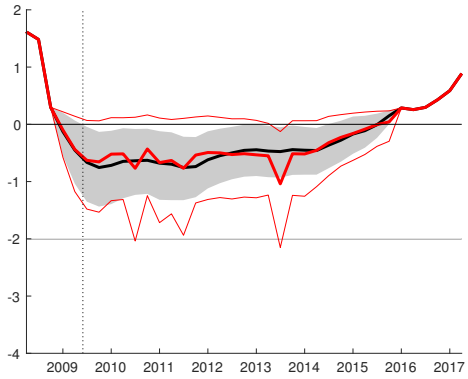

(c) c,π,s,y

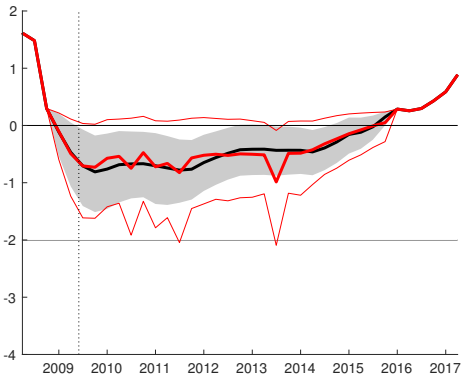

(d) s,c,π,y

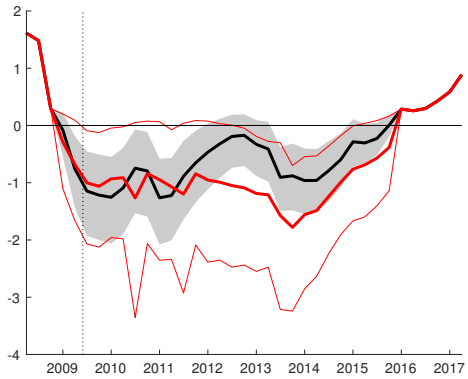

(e) y,s,c,π

Note:Estimates generated from alternative orderings of gap variables in the VAR described in equation (7). Panel 12a depicts baseline estimates; alternative orderings are as indicated above where "y" indicates the block of yield gaps y^2, y^5, y^{10}. Filtered estimates in red, smoothed estimates in black; both surrounded by 90% uncertainty bands (Filtered estimates reflect the endpoints of sequentially re-estimating the entire model over growing samples starting in 1960:Q1. Smoothed estimates use all available observations from 1960:Q1 through 2017:Q2.

Figure 13: Inflation Trend Estimates with re-ordered Gap-VAR (w/Unemployment Rate Gap)

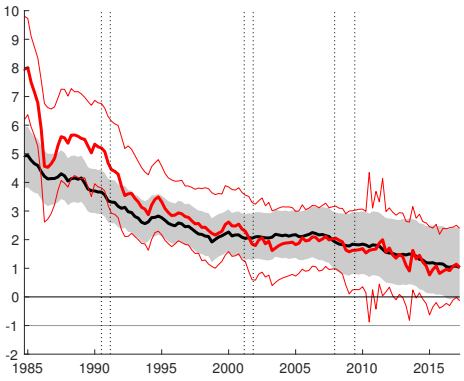

(a) π,c,s,y (baseline)

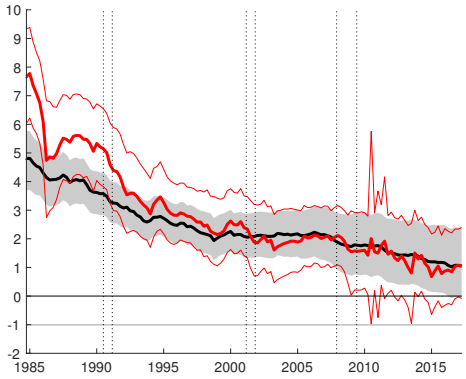

(b) π,s,c,y

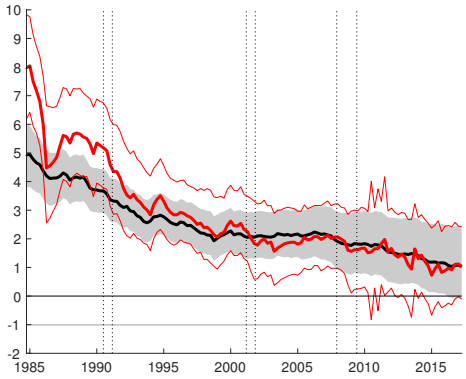

(c) c,π,s,y

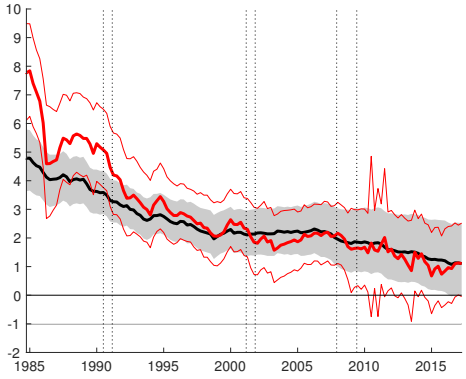

(d) s,c,π,y

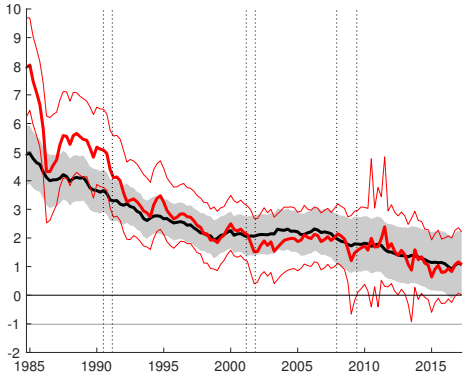

(e) y,s,c,π

Note:Estimates generated from alternative orderings of gap variables in the VAR described in equation (7). Panel 13a depicts baseline estimates; alternative orderings are as indicated above where "y" indicates the block of yield gaps y^2, y^5, y^{10}. Filtered estimates in red, smoothed estimates in black; both surrounded by 90% uncertainty bands (Filtered estimates reflect the endpoints of sequentially re-estimating the entire model over growing samples starting in 1960:Q1. Smoothed estimates use all available observations from 1960:Q1 through 2017:Q2.

Figure 14: Real Rate Trend Estimates with re-ordered Gap-VAR (w/Unemployment Rate Gap)

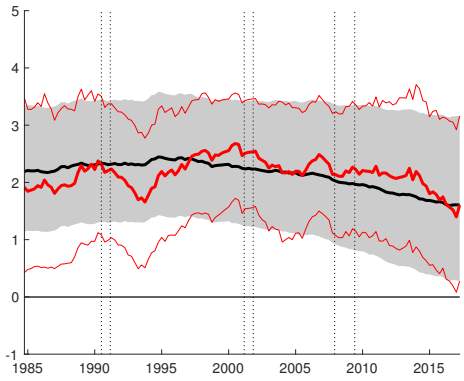

(a) π,c,s,y (baseline)

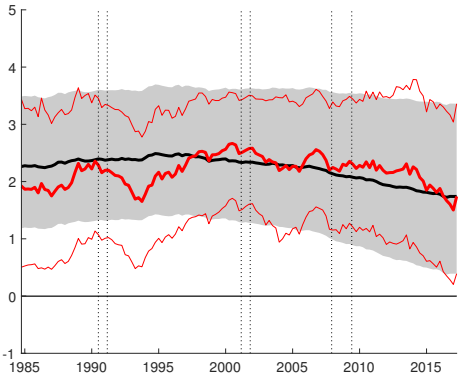

(b) π,s,c,y

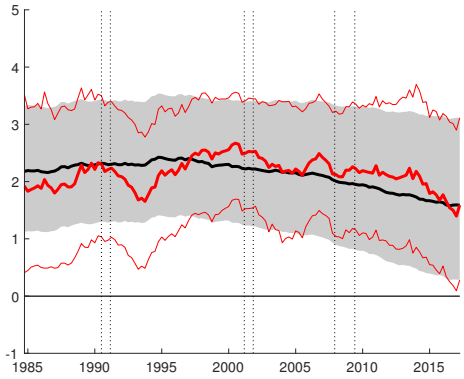

(c) c,π,s,y

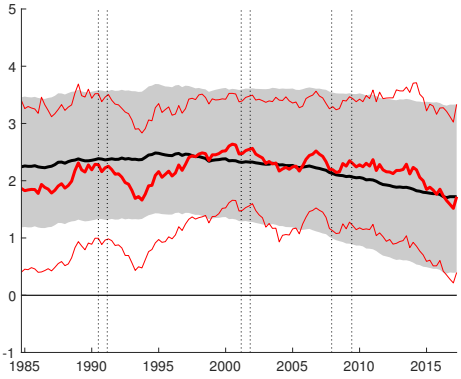

(d) s,c,π,y

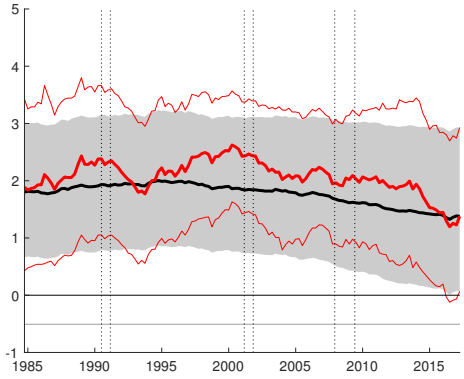

(e) y,s,c,π

Note:Estimates generated from alternative orderings of gap variables in the VAR described in equation (7). Panel 14a depicts baseline estimates; alternative orderings are as indicated above where "y" indicates the block of yield gaps y^2, y^5, y^{10}. Filtered estimates in red, smoothed estimates in black; both surrounded by 90% uncertainty bands (Filtered estimates reflect the endpoints of sequentially re-estimating the entire model over growing samples starting in 1960:Q1. Smoothed estimates use all available observations from 1960:Q1 through 2017:Q2.

Figure 15: Shadow Rate Estimates with re-ordered Gap-VAR (w/Unemployment Rate Gap)

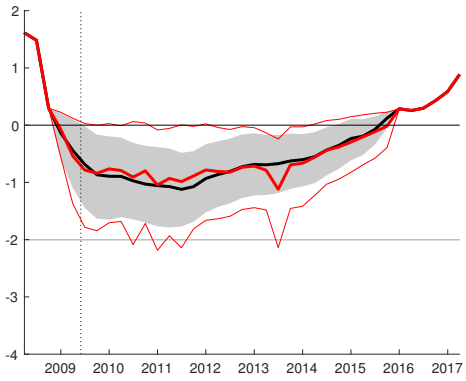

(a) π,c,s,y (baseline)

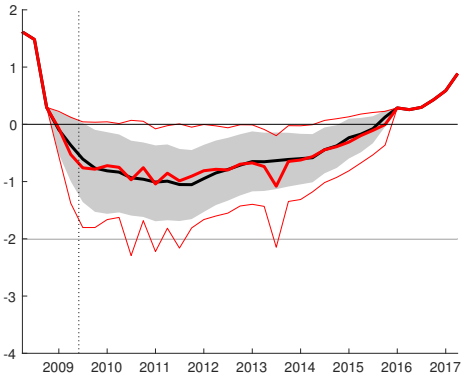

(b) π,s,c,y

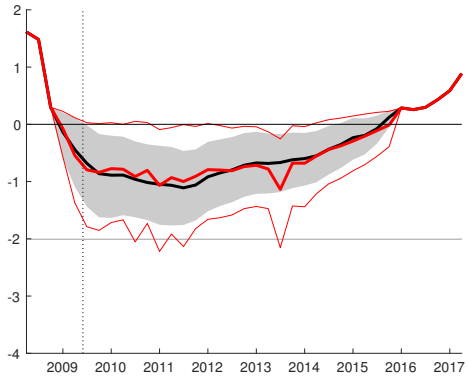

(c) c,π,s,y

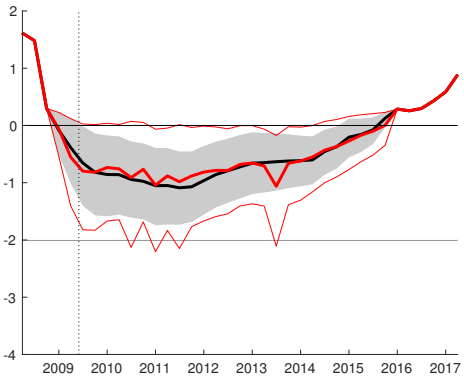

(d) s,c,π,y

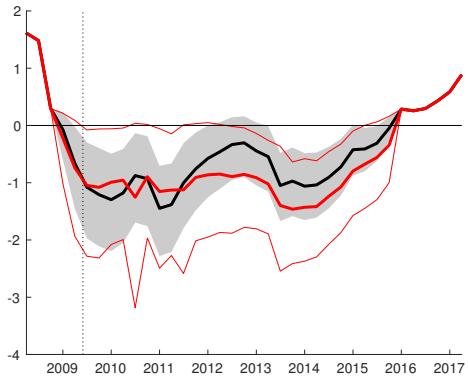

(e) y,s,c,π

Note:Estimates generated from alternative orderings of gap variables in the VAR described in equation (7). Panel 15a depicts baseline estimates; alternative orderings are as indicated above where "y" indicates the block of yield gaps y^2, y^5, y^{10}. Filtered estimates in red, smoothed estimates in black; both surrounded by 90% uncertainty bands (Filtered estimates reflect the endpoints of sequentially re-estimating the entire model over growing samples starting in 1960:Q1. Smoothed estimates use all available observations from 1960:Q1 through 2017:Q2.

Table 1: Forecast Evaluation vs. SPF (Post 2008)

	Forecast horizon h				
	1	2	3	4	5
Panel A: Short-term interest rate i_{t+h}					
	Model (output gap)				
MAD	0.01	0.02	0.06	0.14	0.29
RMSE	0.00	0.00	0.01	0.04	0.12
	SPF rel. to model (output gap)				
rel. MAD	0.34	1.17	1.55	1.32	1.12
rel. RMSE	0.36	1.03	1.26	1.30	1.25
	Model (unemployment rate gap)				
MAD	0.02	0.03	0.07	0.11	0.22
RMSE	0.00	0.01	0.02	0.05	0.10
	SPF rel. to model (unemployment rate gap)				
rel. MAD	0.27	0.94	1.29	1.67	1.49
rel. RMSE	0.27	0.72	1.00	1.24	1.42
Panel B: 10-year interest rate y^{10}_{t+h}					
	Model (output gap)				
MAD	0.26	0.40	0.50	0.57	0.65
RMSE	0.11	0.26	0.36	0.47	0.59
	SPF rel. to model (output gap)				
rel. MAD	0.68***	1.12	1.23*	1.39**	1.55***
rel. RMSE	0.67***	1.07	1.25**	1.35***	1.49***
	Model (unemployment rate gap)				
MAD	0.26	0.41	0.51	0.59	0.67
RMSE	0.11	0.27	0.38	0.49	0.62
	SPF rel. to model (unemployment rate gap)				
rel. MAD	0.69***	1.11	1.21*	1.35***	1.50***
rel. RMSE	0.67***	1.06	1.23**	1.33***	1.46***

Note: *RMSE* are root-mean-squared errors computed from using the medians of our model's and the mean forecast from the SPF as forecasts; *MAD* are mean absolute deviations obtained from using the same forecasts. Relative RMSE and MAD are expressed as ratios relative to the corresponding statistics from the baseline model (values below unity denoting better performance than our model). Predictive densities are re-estimated over growing samples that all start in 1990:Q1 for our model. For the forecast evaluation, the first forecast jumps off in 2009:Q1 and the last in 2017:Q1. Stars indicate significant differences, relative to baseline, in squared losses, absolute losses and density scores, respectively, as assessed by the test of Diebold and Mariano (1995) and West (1996); ***, ** and * denote significance at the 1%, 5% respectively 10% level.

Table 2: Forecast Evaluation vs. WX-SRTSM (Post 2008)

	Forecast horizon h					
	1	2	3	4	5	8
Panel A: Short-term interest rate i_{t+h}						
	Model (output gap)					
MAD	0.09	0.10	0.11	0.12	0.11	0.47
RMSE	0.18	0.21	0.24	0.25	0.22	0.56
	WX-SRTSM rel. to model (output gap)					
rel. MAD	0.78	0.96	1.05	1.30	1.64	0.76
rel. RMSE	0.68	0.94	1.01	1.10	1.44	0.94
	Model (unemployment rate gap)					
MAD	0.08	0.09	0.10	0.10	0.13	0.49
RMSE	0.17	0.18	0.20	0.19	0.21	0.63
	WX-SRTSM rel. to model (unemployment rate gap)					
rel. MAD	0.85	1.11	1.24	1.51	1.45	0.72^{**}
rel. RMSE	0.72	1.07	1.22	1.48	1.48^{*}	0.84^{*}
Panel B: 10-year interest rate y^{10}_{t+h}						
	Model (output gap)					
MAD	0.39	0.54	0.68	0.68	0.71	0.88
RMSE	0.51	0.70	0.80	0.84	0.90	1.14
	WX-SRTSM rel. to model (output gap)					
rel. MAD	1.00	1.04	1.03	1.16^{**}	1.18^{**}	1.22^{***}
rel. RMSE	0.99	1.03	1.06	1.10^{*}	1.12^{***}	1.18^{**}
	Model (unemployment rate gap)					
MAD	0.40	0.52	0.69	0.71	0.80	0.96
RMSE	0.52	0.69	0.79	0.86	0.99	1.30
	WX-SRTSM rel. to model (unemployment rate gap)					
rel. MAD	0.96	1.06	1.02	1.10^{**}	1.04	1.12^{*}
rel. RMSE	0.97	1.05	1.07^{**}	1.07	1.03	1.04

Note: *RMSE* are root-mean-squared errors computed from using the medians of our model's and the WX-SRTSM's predictive densities as forecasts; *MAD* are mean absolute deviations obtained from using the same forecasts. Relative RMSE and MAD are expressed as ratios relative to the corresponding statistics from the baseline model (values below unity denoting better performance than our model). Predictive densities are re-estimated over growing samples that all start in 1990:Q1 for our model and 1990:M1 for the WX-SRTSM. For the forecast evaluation, the first forecast jumps off in 2009:Q1 and the last in 2017:Q1. Stars indicate significant differences, relative to baseline, in squared losses, absolute losses and density scores, respectively, as assessed by the test of Diebold and Mariano (1995) and West (1996); ***, ** and * denote significance at the 1%, 5% respectively 10% level.