# Shadow-Rate VARs*

Andrea Carriero,[1] Todd E. Clark,[2]
Massimiliano Marcellino,[3] and Elmar Mertens[4]

[1]*Queen Mary University of London, and University of Bologna,*
[2]*Federal Reserve Bank of Cleveland,*
[3]*Bocconi University, IGIER, and CEPR,*
[4]*Deutsche Bundesbank.*

May 15, 2023

**Abstract**

VARs are a popular tool for forecasting and structural analysis, but ill-suited to handle occasionally binding constraints, like the effective lower bound on nominal interest rates. We extend the VAR framework by modeling interest rates as censored observations of a latent shadow-rate process, and propose an efficient sampler for Bayesian estimation of such "shadow-rate VARs." We analyze specifications where actual and shadow rates serve as explanatory variables and find benefits of including both. In comparison to a standard VAR, shadow-rate VARs generate superior predictions for short- and long-term interest rates, and deliver some gains for macroeconomic variables in US data. Our structural analysis estimates economic responses to shocks in financial conditions, showing strong differences in the reaction of interest rates depending on whether the ELB binds or not. After an adverse shock, our shadow-rate VAR sees a stronger decline of economic activity at the ELB rather than when not.

*Keywords: Macroeconomic forecasting, effective lower bound, term structure, censored observations*
*JEL classification codes: C34, C53, E17, E37, E43, E47*

*Earlier versions of this paper were also circulated under the title "Forecasting with Shadow-Rate VARs." We gratefully acknowledge helpful suggestions and comments received from the Editor Serena Ng, two anonymous referees, Sidd Chib, Refet Gürkaynak, Sophocles Mavroeidis, Emanuel Mönch, Frank Schorfheide, Harald Uhlig and seminar participants at the Bank of England, Bocconi University, Cleveland Fed, Deutsche Bundesbank, Penn, 2021 IAAE conference, 2021 European Economic Association meeting, and 2021 ESRC Virtual Workshop. The views expressed herein are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Cleveland, the Federal Reserve System, the Eurosystem, or the Deutsche Bundesbank. Marcellino thanks Ministero dell' Istruzione, Universita' e Ricerca - PRIN Bando 2017 – prot. 2017TA7TYC for financial support. Additional results are provided in an supplementary online appendix, and replication codes are available at `https://github.com/elmarmertens/CCMMshadowrateVAR-code`. Corresponding author: Todd E. Clark, todd.clark@researchfed.org.

# 1 Introduction

A popular class of models for forecasting and structural analysis is linear vector autoregressions (VARs) that typically include shorter- and longer-term interest rates due to their importance for macroeconomic forecasting and structural analysis. However, in a number of economies, shorter-term interest rates have been stuck for years at or near their effective lower bound (ELB), with longer-term rates drifting toward the constraint as well. In such an environment, linear econometric models that ignore the ELB constraint on nominal interest rates can be problematic along various dimensions.

For concreteness, we consider the case of the US, where the Federal Open Market Committee (FOMC) has set the target range for the federal funds rate no lower than 0-25 basis points. The Committee maintained this target range over the seven-year stretch from December 2008 through December 2015, after the Great Recession, and again maintained this target range from mid-2020 through mid-March 2022, in response to the recession triggered by the COVID-19 pandemic. Considering other economies, the ELB may even be a bit below zero, with several central banks pursuing so-called negative interest rate policies (NIRP), albeit still at levels close to zero.[1] Consistent with the existence of an ELB (albeit some at a non-positive level), policy rates observed under NIRP appear so far constrained to not fall much below zero; see, for example, estimates obtained for the Euro area by Wu and Xia (2020).

In such an environment, a fundamental challenge for econometric models is to appropriately capture the existence of an ELB on interest rates and the resulting asymmetries in predictive densities and impulse response functions not only for interest rates, but also likely for other economic variables. At a mechanical level, the existence of an ELB calls for treating nominal interest rates as variables whose observations are constrained not to fall below the lower bound.[2]

So far the literature has discussed a number of potential remedies to ELB complications. From a macroeconomic perspective, Swanson and Williams (2014) have argued that it may be sufficient to track longer-term nominal interest rates, as long as their dynamics have remained unaffected by a binding ELB on shorter-term rates, and this has been done by, for example, Crump, et al. (2021), Debortoli, Gali, and Gambetti (2019), and Rogers, Scotti, and Wright (2018). However, by 2020, even 10-year US Treasury yields had fallen below 1 percent, with 5-year yields hovering just above 25 basis points. In contrast, the finance literature has derived important implications of the ELB for the entire term structure of interest rates. Following the seminal work of Black (1995), the term structure literature views the ELB as a censoring constraint on nominal interest rates (as we do), from which no-arbitrage restrictions are derived for yields of all maturities

[1] For example, the Swiss National Bank targeted a level of $-75$ basis points for its policy rate until mid-June 2022, while the European Central Bank maintained a deposit rate of $-50$ basis points from September 2019 (which was the culmination of a series of steps starting in December 2011 to gradually lower the rate from 25 basis points) through July 2022. The repo (now policy) rate of the Swedish Riksbank was at or below 25 basis points from October 2014 through June 2022 (bottoming out at $-50$ basis points from February 2016 to January 2019, and remaining at 0 through late April 2022). One of the most extensive episodes of monetary policy near the ELB has occurred in Japan, where policy rates have been near zero since 2008, with the current policy rate at $-10$ basis points since 2016.

[2] For example, Bäurle, et al. (2020), Iwata and Wu (2006), and Nakajima (2011) have modeled the nominal interest rate as a censored (or bounded) variable in VAR systems featuring only lagged actual (but not shadow) rates on the right-hand side of equations for interest rates and other economic variables. Relatedly, Chan and Strachan (2014) model nominal rates as bounded processes. Under either of these approaches, nominal interest rates and other variables respond to lags of the actual interest rate but not the underlying (and unconstrained) shadow rate.

(which we do not). The resulting restrictions are, however, non-trivial and have mostly been implemented for models with state dynamics that are affine, homoskedastic, and time-invariant; see, for example, Bauer and Rudebusch (2016), Christensen and Rudebusch (2012, 2015, 2016), Krippner (2015), and Wu and Xia (2016).[3]

An upshot of the term structure literature is the availability of shadow-rate estimates, such as those regularly updated by Krippner (2013, 2015) or Wu and Xia (2016), which also led to a by now popular shortcut for applied work that avoids dealing with censoring. Indeed, an easy choice is to plug in these shadow-rate estimates as data points for the nominal short-term interest rate during an ELB episode. However, while convenient, this plug-in approach risks a generated regressor problem that could be substantial, as documented by, for example, Krippner (2020). Mavroeidis (2021) notes that a plug-in approach rules out consistent estimation and valid inference with a VAR, due to estimation error in the shadow rate that is often highly autocorrelated and not asymptotically negligible. Shadow-rate estimates are model-specific objects, fitted to best capture the dynamics of observed data through the lens of the model, and can be quite sensitive to model choices (Christensen and Rudebusch (2015); Krippner (2020)). An obvious remedy to these considerations is to integrate the shadow-rate inference into the model used for forecasting and structural analysis, as we do here.

In this paper, we develop shadow-rate approaches for accommodating the ELB in commonly-used macroeconomic VARs. We illustrate the use of shadow-rate VARs for structural analysis and assess their benefits in forecasting macroeconomic and financial variables in US data. To handle the ELB on interest rates, we model observed rates as censored observations of a latent shadow-rates process in an otherwise standard VAR setup. The shadow rates are assumed to be equal to observed rates when above the ELB. Motivated in part by other work in the literature described below, we consider two specifications of the relationship of macroeconomic variables to interest rates or the shadow rates. In a formulation we refer to as the "simple shadow-rate VAR," the model relates macroeconomic variables to lags of the shadow rates; in this case, the shadow rates enter all of the VAR's equations. The "hybrid shadow-rate VAR" instead relates macroeconomic variables to lags of actual interest rates; the shadow rates enter just the interest rate equations of the VAR. We prefer the hybrid specification for its richer dynamics that allow constraints on interest rates to directly feed into macroeconomic forecasts and generate some advantages in forecast accuracy. Our analysis takes the level of the ELB as given, as it appears reasonable to abstract from uncertainty about its level (or even drift therein) in the context of the US.[4]

Our approach extends the unobserved components model of Johannsen and Mertens (2021) to the general VAR setting, and we develop a computationally more efficient shadow-rate sampling algorithm to be able to estimate larger models. In particular, we use a Gibbs sampling procedure, which is embedded in a Markov chain Monte Carlo (MCMC) sampler, to generate posterior draws from the latent shadow-rate process. Drawing directly from the truncated posterior makes the procedure computationally efficient, and

[3] Kim and Singleton (2012) also consider a quadratic-Gaussian specification with a shadow rate and find that it fits data for Japan from 1995 to 2008 as well as a shadow-rate specification in the tradition of Black (1995).

[4] In our empirical application on US data, we consider the ELB to have a constant and known value of 25 basis points, consistent with other studies, such as Bauer and Rudebusch (2016), Johannsen and Mertens (2021), and Wu and Xia (2016). Considering the euro area, for example, Wu and Xia (2020) model and estimate a stochastic downward drift in the ELB level. We conduct robustness checks with values of 12.5 and 50 basis points and find our main findings to be robust to the choice.

using QR methods to construct positive definite second-moment matrices makes the procedure numerically reliable. We apply our shadow-rate approaches to a medium-scale Bayesian VAR (BVAR) for twenty US macroeconomic and financial variables, with stochastic volatility, which has been shown to generate competitive forecasts when ignoring the ELB (e.g., Carriero, Clark, and Marcellino (2019)).[5] Critically, we also demonstrate how to handle data in which the ELB binds for multiple interest rate instruments of different maturities. Our baseline specification includes a set of interest rates, consisting of not only the federal funds rate and 6-month Treasury bill rate but also yields for several other bond maturities.[6] The inclusion of term structure information (beyond a measure of the short-term policy rate) turns out to importantly influence estimated shadow rates. The inclusion of interest rates other than the federal funds rate generates less negative shadow rates, representing less outsized expectations of monetary stimulus in response to the Great Financial Crisis (GFC), which in turn leads to improved shadow rate-based forecasts for macroeconomic and financial variables. Bond yields and other financial indicators such as stock prices allow the shadow-rate VARs to capture the effects of asset purchases and forward guidance regarding the path of policy rates.

In out-of-sample simulations for the US since 2009, interest rate forecasts obtained from our shadow-rate VARs are clearly superior, in terms of both point and density accuracy, when compared to predictions from a standard VAR that ignores the ELB. For measures of economic activity and inflation, forecasts from the simple shadow-rate VAR are broadly on par with those from a standard VAR that ignores the ELB. Predictions from the hybrid shadow-rate VAR even perform a little better than the standard model for a few of the macroeconomic variables (and are otherwise on par). The close performance of the (simple) shadow-rate VAR and the standard VAR reflects that, in both models, forecasts of short-term interest rates that enter the forecasting equations of macroeconomic variables can and do go negative (albeit to varying degrees) during ELB periods. Moreover, as shown in our supplementary online appendix, when conditioned on interest rates of shorter- and longer-term maturities, our shadow-rate VARs considerably improve upon forecasts for macroeconomic variables obtained from a linear VAR that simply omits short-term interest rates to avoid ELB issues.

We also apply our hybrid shadow-rate VAR to structural analysis, examining the impacts of shocks to financial conditions. Generally defined, impulse responses measure a change in forecasts prompted by a specific shock. Thus, a good forecasting model is needed to properly measure impulse responses, and with interest rates affected by the ELB constraint, our shadow-rate VARs are promising candidates for this analysis. When the economy is near the ELB, the constraint on interest rates can affect the responses of macroeconomic variables to the identified shock, as interest rates will decline less than they would in the absence of the constraint, which in turn feeds into the macroeconomic projections of our hybrid shadow-rate VAR. We use the generalized impulse response approach of Koop, Pesaran, and Potter (1996), so that the estimated responses reflect any endogenous changes in the extent to which the ELB constraint binds due to the identified shock. The financial conditions shock is measured by the excess bond premium (EBP) of

[5]Other studies documenting the relevance of heteroskedasticity in VARs are Chan and Eisenstat (2018), Clark (2011), Clark and Ravazzolo (2015), and D'Agostino, Gambetti, and Giannone (2013).

[6]Specifically, we include maturities ranging from the daily federal funds rate to 10-year Treasury yields (as well as the yield on BAA-rated corporate bonds with maturities of at least 20 years). Other empirical studies that use medium- and longer-term yields include Crump, et al. (2021), who omit the federal funds rate to obviate the ELB problem, and Jones, Kulish, and Morley (2022) and Kulish, Morley, and Robinson (2017), who consider yields of up to 5 years in their analysis.

Gilchrist and Zakrajšek (2012), which is incorporated as an additional variable in our hybrid shadow-rate VAR. We compare responses for a shock origin of January 2007, when interest rates were far away from the ELB, to responses for a shock origin of December 2012, when short-term rates were at the ELB. According to our estimates, ELB constraints temper the decline of interest rates that normally occurs when the ELB does not bind, in turn yielding a sharper decline of economic activity and stock prices. A similar pattern is obtained when comparing estimated responses from a model entirely ignoring the ELB to those from our model accounting for the ELB binding when the shock hits. And the effects on yields and macro variables of a financial uncertainty shock, obtained by replacing the excess bond premium with a measure of financial uncertainty (based on the VXO index of options-implied volatility in stock returns), are qualitatively similar to those of the EBP shock, across periods and model specifications.

Overall, our shadow-rate specifications successfully address the ELB, which drastically improves interest rate forecasts (compared to a model that ignores the ELB), while preserving (and even marginally improving upon) the standard VAR's ability to forecast a range of other variables. In this respect, our proposed approaches could be seen as helpful tools for preserving the practical value of VARs for forecasting in the presence of the ELB. In practical settings, presented with forecasts from standard VARs in which interest rates fall below the ELB, consumers of forecasts could question the reliability or plausibility of the forecasts of the other variables of interest. To these consumers, forecasts of macroeconomic variables from shadow-rate VARs that obey the ELB could be seen as more coherent and therefore practically useful even if their historical accuracy was no greater than that achieved by a standard VAR ignoring the ELB. Our specifications can also be used for structural analysis, to capture, for example, the effects of shocks to financial conditions or uncertainty while ensuring the estimated responses obey lower bound constraints on interest rates.

The remainder of this paper is structured as follows. Section 2 relates our paper to other contributions regarding the modeling of the ELB and its consequences. Section 3 describes the modeling and estimation of our shadow-rate VARs. Section 4 details the data used in our empirical application. Section 5 presents shadow rate estimates and resulting interest rate projections. Section 6 provides the forecast evaluation, and Section 7 illustrates the use of our hybrid model for structural analysis. Section 8 concludes. The Appendix contains technical details on our estimation procedure.

## 2 Related Literature

In this section we relate our approach to other shadow-rate work. There is a term structure literature on shadow-rate models, with which we share the approach of modeling nominal interest rates as censored variables. But we do not enforce any specific no-arbitrage (or other structural) restrictions. As such, our approach is part of the literature that uses VARs to derive forecasts and expectational errors of financial and economic variables without imposing the restrictions of a specific structural model (such as an affine term structure or DSGE model). Should the data satisfy such restrictions, they will also be embodied in estimates derived from a more generic reduced-form model. The potential loss in the efficiency of forecasts that do not explicitly enforce such restrictions can be offset by a gain in robustness obtained from not imposing

restrictions that are false. In fact, as argued by Joslin, Le, and Singleton (2013), the possible gains for forecasting from imposing restrictions from the true term structure model may be small. Moreover, as in Johannsen and Mertens (2021), economic forecasters may be interested in using time series models that allow for features, such as time-varying parameters and stochastic volatilities, that may be harder to embed in a formal no-arbitrage model.[7] While the VAR framework shuns restrictions specific to a particular structural model, it is, of course, well suited to inference on the economic responses to structurally identified shocks. In this vein, we also extend our shadow-rate approach to structural VAR analysis as discussed below.

In the context of DSGE models, Sims and Wu (2021) show that asset purchase programs can be analogous to large reductions in the federal funds rate, and that asset purchases provided about as much stimulus during the Great Recession as a decline in the policy rate commensurate with the fall in shadow rate estimates such as in Wu and Xia (2016). In particular, Wu and Zhang (2019) derive conditions in which alternative policies can circumvent the ELB such that the economy retains a linear representation with a shadow rate capturing the effects of policy. Moreover, Wu and Zhang (2019) provide references to empirical work that has concluded that conventional and unconventional monetary policies work in a similar fashion. Other DSGE models formulate monetary policy in terms of censored prescriptions from a policy rule for shadow rates as in Aruoba, et al. (2021), Gust, et al. (2017), Jones, Kulish, and Morley (2022), and Kulish, Morley, and Robinson (2017).[8]

Mavroeidis (2021) discusses various specification choices for the underlying VAR model, similar to some that we also evaluate.[9] Our simple shadow-rate VAR assumes that only shadow rates, not actual rates, matter as explanatory variables in linear forecasting equations. As detailed in Section 3.4, the simple shadow-rate VAR corresponds in reduced form to the "censored" model of Mavroeidis (2021). In contrast, our hybrid shadow-rate VAR combines actual and shadow interest rates in its state dynamics, in a reduced form that combines elements of the "censored" and "kinked" VAR of Mavroeidis (2021).[10] Since actual and shadow rates are perfectly collinear when above the ELB, their joint inclusion in a VAR can create challenges in identification and estimation, in particular for the purpose of generating reliable out-of-sample forecasts based on the (yet) still limited experience with ELB data. Hence, our hybrid shadow-rate VAR proceeds selectively: We include shadow rates as VAR regressors only in forecasting equations for other (shadow) term structure variables, while including actual rates (and not shadow rates) as explanatory variables in the VAR equations of macroeconomic variables and other measures of financial conditions. While maintaining

[7] Johannsen and Mertens (2021) provide an out-of-sample forecast evaluation for short- and long-term nominal interest rates in a model smaller than our VARs, and find their unobserved components shadow-rate model to be competitive with the no-arbitrage model of Wu and Xia (2016), but do not consider forecasts of other variables. Gonzalez-Astudillo and Laforte (2020) embed a shadow-rate model in an unobserved components model and report improved point forecasts for economic and financial variables from the shadow-rate approach.

[8] The shadow-rate estimates of Aruoba, et al. (2022), Gust, et al. (2017), Jones, Kulish, and Morley (2022), and Kulish, Morley, and Robinson (2017) reflect data on real activity, inflation, and short- and (in case of Jones, Kulish, and Morley (2022) and Kulish, Morley, and Robinson (2017)) medium-term interest rates.

[9] Goncalves, et al. (2021) study local projections for nonlinearly transformed variables, including the case of censoring.

[10] All our shadow-rate specifications assign a role to lagged shadow (as opposed to actual) rates to serve as predictors of future outcomes, which is also in keeping with some of the policy-rule specifications in DSGE models such as Gust, et al. (2017). In contrast, Aruoba, et al. (2022) limit attention to settings where VAR forecasts depend on lagged actual rates, but not lagged shadow rates. In Aruoba, et al. (2022), the shadow rate arises only contemporaneously when the VAR vector is shocked. Similarly, Berg (2017), Chan and Strachan (2014), and Iwata and Wu (2006) consider only censoring of the VAR's left-hand side variables, without tracking the underlying, uncensored shadow rate as a potential predictor.

a well-identified VAR system, the hybrid allows for potentially relevant non-linearities through the joint effects of actual and shadow rates on future outcomes in all variables.[11]

In the context of structural VAR models (SVARs), Aruoba, et al. (2022) and Mavroeidis (2021) consider shadow-rate approaches to identify and estimate impulse responses to monetary policy shocks. We differ from these SVAR studies in focusing on the implementation of shadow-rate approaches in a medium-scale Bayesian VAR (with stochastic volatility), and we evaluate its application to forecasting and structural analysis. Our structural analysis conditions on the proximity of interest rates of different maturities to the ELB with the generalized impulse response approach of Koop, Pesaran, and Potter (1996), which can capture relevant state dependencies in the economic responses to structural shocks. Moreover, our structural analysis considers shocks to financial conditions, as measured by the excess bond premium of Gilchrist and Zakrajšek (2012), and the results are consistent with the model-based predictions of Gilchrist, et al. (2017). Importantly, in common with Aruoba, et al. (2022) and Mavroeidis (2021), we jointly estimate latent shadow rates and structural impulses within one coherent model. Krippner (2020) provides a critical review of plug-in approaches that estimate conventional impulse responses from VARs that condition on external shadow-rate estimates as data inputs.

# 3 Shadow-Rate VARs

This section contrasts our shadow-rate approaches with a standard VAR. Throughout, we take the value of the lower bound, denoted $ELB$, as a given and known constant. Our model considers a vector of interest rates to which the ELB applies. For brevity, we use the singular to refer to "the" nominal interest rate, $i_t$, and its associated shadow rate, $s_t$, while both $i_t$ and $s_t$ will generally be vectors of length $N_i = N_s$. A central element of our approach is to relate actual and shadow rates via a censoring equation known from Black (1995):

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

where the censoring of the shadow-rate vector $s_t$ is element-wise. As in the no-arbitrage literature on the term structure of interest rates (surveyed in Section 1), the censoring function (1) implies that the shadow rate is observed and equal to the actual interest rate when the latter is above the ELB.[12] When the ELB is binding, so that $i_t = ELB$, the shadow rate is a latent variable that can only take values below (or equal to) $ELB$, which will inform our inference about $s_t$. In the data, at a given point in time, the ELB may, of course, be binding for none, some, or all interest rate measures included in $i_t$.

In principle, our model pairs each interest rate measure with a separate shadow rate. However, the distinction between actual and shadow rates matters only for those measures for which the ELB has actually been binding in the data.[13] In the baseline specification of our empirical application, we include a total of

[11] At least indirectly, through effects on other economic and financial variables, both actual and shadow rates eventually influence forecasts of all variables in this hybrid VAR system.

[12] The property that the shadow rate is identical to the actual rate when above the ELB makes our approach based on Black (1995) distinct from others, like Lombardi and Zhu (2018), that defines the shadow rate more broadly as a common factor of interest rates and possibly other variables intended to capture the stance of monetary policy.

[13] For a given data set, the implementation of our model may thus allow for $N_s \leq N_i$ with $N_s$ reflecting the number of elements

$N_i = 6$ interest rates, covering shorter- and longer-term maturities, out of which the ELB has been binding for up to three measures in our data set.[14] In another specification we consider, the model includes only the federal funds rate (and a single shadow rate).[15] In principle, in the event of aggressive forward guidance or yield curve control, the ELB could bind for additional, longer maturities, which could also be handled by our model.

### 3.1 Standard linear VAR

Before turning to our VAR-based specification of a process for $s_t$, we first describe the standard VAR approach. A standard VAR is a linear model for the evolution of a vector of observed data, $y_t$. Omitting intercepts, we have the following system of $N_y$ equations for a VAR with $p$ lags:

$$y_t = \sum_{j=1}^{p} C_j \, y_{t-j} + v_t \,, \qquad \text{with} \quad v_t \sim N(0, \Sigma_t) \,, \qquad \text{and} \quad E_{t-1} v_t = 0 \,. \tag{2}$$

Anticipating our subsequent application, we assume time-invariant transition matrices, $C_j$, but allow for time-varying shock volatilities, $\Sigma_t$, as in Clark (2011), among others.[16] Critically, VAR errors are typically assumed to have a symmetric distribution with unbounded support. When $y_t$ includes the nominal interest rate, $i_t$, the resulting predictive densities will fail to incorporate the effects of the effective lower bound, with particularly detrimental effects when $i_t$ is close to $ELB$. As a special case of (2), consider a random walk process for the nominal interest rate, $i_t = i_{t-1} + v_t$. When $i_t = ELB$, the $k$-period-ahead point forecast still satisfies the ELB, since $E_t i_{t+k} = ELB$. But the associated density forecasts have 50 percent of their mass below $ELB$ as the linear model ignores the ELB constraint.

That being said, a standard linear VAR could yield plausible macroeconomic forecasts even in settings when monetary policy is constrained by the ELB. With short-term rates included, a conventional VAR may forecast at least adequately because, at any given forecast origin, projections of future short-term interest rates can turn negative.[17] To the extent that the historical behavior of monetary policy implies the central bank would have set the policy rate negative in an ELB episode but could not and took other steps to provide policy accommodation, the simple linear VAR's forecasts could be helped by being allowed to project negative rates over the forecast horizon. The accuracy of macroeconomic forecasts may also be helped by the inclusion of long-term bond yields and other financial indicators such as stock prices; these indicators reflect the effects of asset purchases and forward guidance from the central bank regarding the path of policy rates. Indeed, Crump, et al. (2021) develop a large VAR intended to be useful for a range

in the $N_i$-dimensional vector $i_t$ for which the ELB has been binding in the sample.

[14] In our data set, with an ELB value of 25 basis points, the constraint has been binding for the federal funds rate, the 6-month Treasury bill rate, and the 1-year yield. In addition, our baseline specification includes yields on 1-year, 5-year, and 10-year Treasuries as well as Moody's BAA yield on bonds with maturities of 20 years and above.

[15] Moreover, the supplementary online appendix provides further details about the effects of adding specific yields, ranging from 6-month to 20-year maturities, on estimated shadow rates and average forecast accuracy.

[16] In our empirical application, we follow Carriero, Clark, and Marcellino (2019) and assume that $v_t = A^{-1} \Lambda_t^{0.5} \varepsilon_t$, where $A$ is a lower unit-triangular matrix, $\Lambda_t$ is a diagonal matrix, and the vector of its diagonal elements is denoted $\lambda_t$, with $\log \lambda_t = \log \lambda_{t-1} + \eta_t$, $\eta_t \sim N(0, \Phi)$, and $\varepsilon_t \sim iid \, N(0, I)$. Other forms of heteroskedasticity could also be specified.

[17] As discussed below, Figures 4 and 5 provide forecasts to show that this is empirically the case for the US.

of forecasting questions faced by a central bank. Their model omits short-term interest rates out of ELB considerations, instead using the 2-year Treasury yield as an indicator of the stance of policy, with several other yields and financial indicators also included in the model.

### 3.2 Simple shadow-rate VAR

The simple shadow-rate approach posits a VAR for a hypothetical data vector, $z_t$, that is identical to $y_t$ except for replacing $i_t$ with $s_t$, and derives actual rates, $i_t$, from shadow rates $s_t$ via the censoring constraint (1). Without loss of generality, partition $y_t$ in a vector of $N_x = N_y - N_s$ other variables, $x_t$, that have unbounded support, and the nominal interest rate $i_t$, with $x_t$ ordered on top:

$$y_t = \begin{bmatrix} x_t \\ i_t \end{bmatrix} \quad \text{and let} \quad z_t = \begin{bmatrix} x_t \\ s_t \end{bmatrix} \quad \text{with} \quad i_t = \max\left(ELB, s_t\right). \tag{3}$$

In the simple shadow-rate VAR approach, we posit VAR dynamics for the partially latent vector $z_t$. Analogously to (2) we have:

$$z_t = \sum_{j=1}^{p} C_j \, z_{t-j} + v_t \,, \qquad \text{with} \quad v_t \sim N(0, \Sigma_t) \,, \qquad \text{and} \quad E_{t-1} v_t = 0 \,. \tag{4}$$

The simple shadow-rate VAR system is a non-linear state space model that consists of the kinked measurement equation (3) and the linear state evolution described by the VAR in (4).[18] The purely linear state evolution of the simple shadow-rate VAR also occurs in affine term structure models with shadow rates. Note that, although our specification assumes Gaussianity, this Gaussianity is in the *shadow-rate* system and innovations to the shadow rate, not innovations to interest rates constrained by the ELB. (The censoring equation implies that observed actual rates above the ELB are identical to shadow rates, which gives us observations for realizations from the linear Gaussian process for the shadow rates.)

### 3.3 Hybrid shadow-rate VAR

In our simple shadow-rate VAR, forecasts depend on lagged shadow rates, not actual interest rates; shadow rates are assumed to be the relevant interest rate measure for predicting all variables. However, while shadow rates may indeed capture effects of unconventional monetary policies (such as forward guidance or asset purchases), it is actual rates that are paid (earned) by borrowers (lenders). In this vein, we consider a hybrid model in which actual interest rates, not shadow rates, predict future macroeconomic variables, while nominal interest rates are still modeled as censored processes. Since actual and shadow rates are perfectly collinear when above the ELB, their joint inclusion in a VAR can create challenges in identification and estimation, in particular for the purpose of generating reliable out-of-sample forecasts based on the (yet) still limited experience with ELB data. Hence, our hybrid shadow-rate VAR includes shadow rates as VAR regressors only in forecasting equations for other (shadow) term structure variables, while using actual rates

[18] In addition, the simple shadow-rate VAR system includes any state equations needed to track parameter drift, such as the time-varying volatilities embedded in $\Sigma_t$ in the case of our application.

(and not shadow rates) as explanatory variables in the VAR equations of macroeconomic variables and other measures of financial conditions, collected in $x_t$. This hybrid model takes the form:

$$x_t = \sum_{j=1}^{p} C_{xx,j}\, x_{t-j} + \sum_{j=1}^{p} C_{xi,j}\, i_{t-j} + v_{x,t}\,, \quad \text{and} \quad s_t = \sum_{j=1}^{p} C_{sx,j}\, x_{t-j} + \sum_{j=1}^{p} C_{ss,j}\, s_{t-j} + v_{s,t}. \tag{5}$$

In this specification, all coefficients are identifiable from sample periods not constrained by the ELB. The model can be seen as capturing the dynamics of short-term rates that are implied by the historical behavior of monetary policy which would have prescribed pushing rates below the ELB (if possible) while modeling actual economic outcomes as a function of actual interest rates, not the shadow rates.[19]

Compared to this hybrid model, the simple shadow-rate VAR has the advantage of simplicity, in that its state dynamics are entirely linear (see equation (4)), and censoring affects only its measurement equation in (3). Such a structure would be consistent with work discussed in Sections 1 and 2 (and further below in Section 3.5) that sees unconventional monetary policies as having been capable of entirely offsetting the ELB on policy rates. In contrast, in the hybrid system of equation (5) the censored values of (lagged) actual interest rates are state variables that influence the evolution of all variables. Such a specification could be seen as advantageous since the decisions of households and firms are based on the actual (and not shadow) levels of interest rates, so that their levels should serve as predictors in the VAR system. Of course, the distinction is lessened when longer-term rates (for which the ELB has not been binding so far) are included in the vector $i_t$, as we consider in our empirical application. While longer-term rates may indeed be relevant for certain spending and investment categories, some lending rates (e.g., car loans) may be more tied to short-term interest rates than 5- or 10-year bond yields. In addition, deposit rates earned by some savers will also be more tied to short-term rates, making actual short-term rates relevant for macroeconomic forecasting even when long-term rates are included in the analysis.

### 3.4 Relationship to Mavroeidis (2021)

In this section, we sketch the relationship of our reduced-form shadow-rate and hybrid VARs to the specifications of Mavroeidis (2021). Adapted to our notation, the general "censored and kinked" model of Mavroeidis (2021) (CKSVAR, see his equation (20)) takes the form

$$\begin{bmatrix} A_{11} & A_{12}^{*} & A_{12} \\ A_{21} & A_{22}^{*} & A_{22} \end{bmatrix} \begin{bmatrix} x_t \\ s_t \\ i_t \end{bmatrix} = BX_t + B^{*}X_t^{*} + \varepsilon_t, \tag{6}$$

where $X_t$ contains the relevant lags of all variables in the model (including the policy rate), and $X_t^{*}$ contains the relevant lags of the shadow rate that may enter the right-hand side of the VAR. As discussed by Mavroei-

[19] We also investigated the use of a VAR featuring lagged actual and shadow rates in all equations, that is a VAR of the form $z_t = \sum_j (C_j\, z_{t-j} + F_j\, i_{t-j}) + v_t$. However, since in post-war data prior to 2009, the ELB has never been binding, we have $i_t = s_t$, leading to perfectly collinear regressors. Only ELB data can be used to identify separate slope coefficients on lagged actual and shadow rates in this case (and only with still considerable uncertainty, due to the short sample available and the large number of parameters to be estimated). To evaluate out-of-sample forecasts, such a specification is thus of (yet) limited use.

dis (2021), the general specification in (6) has time-invariant parameters, but leads to a reduced form with some time variation in parameters due to the effects of the censoring constraint on nominal interest rates.

Our shadow-rate VARs contain elements of the "censored" and "kinked" specifications of Mavroeidis (2021). In particular, both of our shadow-rate VARs imply the impact matrix restrictions $A_{12} = 0$ and $A_{22} = 0$ in (6). With those restrictions added to the Mavroeidis model, we obtain the following time-invariant reduced-form representation for (6):

$$\begin{bmatrix} x_t \\ s_t \end{bmatrix} = \begin{bmatrix} \bar{A}_{11} & \bar{A}^*_{12} \\ \bar{A}_{21} & \bar{A}^*_{22} \end{bmatrix} BX_t + \begin{bmatrix} \bar{A}_{11} & \bar{A}^*_{12} \\ \bar{A}_{21} & \bar{A}^*_{22} \end{bmatrix} B^* X^*_t + \begin{bmatrix} \bar{A}_{11} & \bar{A}^*_{12} \\ \bar{A}_{21} & \bar{A}^*_{22} \end{bmatrix} \varepsilon_t \,, \tag{7}$$

$$\text{with} \quad \bar{A} = \begin{bmatrix} \bar{A}_{11} & \bar{A}^*_{12} \\ \bar{A}_{21} & \bar{A}^*_{22} \end{bmatrix} \equiv \begin{bmatrix} A_{11} & A^*_{12} \\ A_{21} & A^*_{22} \end{bmatrix}^{-1} = A^{-1}. \tag{8}$$

The simple shadow-rate VAR corresponds in reduced form to the "censored" SVAR specification (CSVAR) in Mavroeidis (2021). The CSVAR assumes the absence of effects from (lagged or current) actual policy rates on outcomes, so that the last $p$ columns of the $B$ matrix, which are associated with coefficients on lags of the actual policy rate in the regressor vector $X_t$, are zero. Furthermore, the CSVAR has $A_{12} = 0$ and $A_{22} = 0$, as assumed above.[20]

Our hybrid model combines elements of the "censored" and "kinked" SVARs of Mavroeidis (2021) in a reduced-form setting. With the regressors blocked as in the Mavroeidis formulation (so lags of the policy rate are the last elements in $X_t$), our hybrid VAR can be expressed as

$$\begin{bmatrix} x_t \\ s_t \end{bmatrix} = \begin{bmatrix} C_{11} & C_{12} \\ C_{21} & 0 \end{bmatrix} X_t + \begin{bmatrix} 0 \\ C^*_2 \end{bmatrix} X^*_t + \begin{bmatrix} \bar{A}_{11} & \bar{A}_{12} \\ \bar{A}_{21} & \bar{A}^*_{22} \end{bmatrix} \varepsilon_t. \tag{9}$$

Multiplying through by $A = \bar{A}^{-1}$ and collecting terms yields the following structural representation:

$$\begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A^*_{22} \end{bmatrix} \begin{bmatrix} x_t \\ s_t \end{bmatrix} = \begin{bmatrix} A_{11}C_{11} + A_{12}C_{21} & A_{11}C_{12} \\ A_{21}C_{11} + A^*_{22}C_{21} & A_{21}C_{12} \end{bmatrix} X_t + \begin{bmatrix} A_{12}C^*_2 \\ A^*_{22}C^*_2 \end{bmatrix} X^*_t + \varepsilon_t. \tag{10}$$

It follows that our hybrid specification has a number of the elements included in the general CKSVAR model of Mavroeidis. Both feature the same regressors on the right-hand side. However, our specification restricts the general model of Mavroeidis (2021) by setting $A_{12} = 0$ and $A_{22} = 0$, which makes the model's parameters constant over time and omits dependence of $x_t$ and $s_t$ on $i_t$.

### 3.5 Shadow rate interpretation and constant-parameter assumptions

Considering a standard VAR, Bernanke and Blinder (1992) proposed interpreting the policy rate equation of the VAR as a feedback rule that describes monetary policy. In a similar spirit, the shadow-rate equation

[20] From our setting, the restrictions $A_{12} = 0$ and $A_{22} = 0$ lead to the coefficient values $\kappa = 1$ and $\tilde{\beta} = 0$ in the setup of Mavroeidis (2021). Mavroeidis directly indicates that, in the CSVAR, the term $\kappa$ defined in his equation (22) equals 1, and the term $\tilde{\beta}$ defined in his equation (28) equals 0. Mavroeidis also comments (p. 2864) that the CSVAR imposes restrictions such that $A_{12} = 0$ and "the elements of $B$ corresponding to (the) lagged" policy rate are equal to 0 (the latter implying that $A_{22} = 0$). Moreover, with $\tilde{\beta} = 0$, the VAR's coefficients do not change with the ELB (see his equation (24)), as in our shadow-rate VAR.

of the VAR model in (4) can be thought of as embedding a monetary policy reaction function that relates the shadow rate to the variables included in the VAR (4).[21] The actual policy rate follows the same reaction function, except that the actual rate is constrained to not fall below the ELB. As a result, the model's prescriptions for the federal funds rate — evident in our out-of-sample forecasts — obey the ELB on actual policy rates, even as the shadow rate will be below zero during the ELB episodes. In contrast, the reaction function implied by a standard VAR ignores the ELB and can be prone to prescribing policy rates that violate the ELB during severe downturns as seen in the last two decades.

During the Great Recession, the FOMC's Bluebook or Tealbook regularly included forecast and policy simulations treating the federal funds rate as unconstrained and allowed to fall below the ELB, either in simple prescriptions from benchmark Taylor-type rules or simulations of optimal policy conducted with the FRB/US model. Text in the March 2008 Bluebook explicitly suggested a negative funds rate prescription could measure the effect of balance sheet policies or forward guidance and reflect the overall stance of policy: "The additional monetary easing associated with the counterfactual unconstrained optimal policy can provide a useful benchmark in judging the stimulus provided by nontraditional policy actions . . . . . (S)ome of the benefits of such an unconstrained policy can likely be attained by a policy of large-scale purchases of long-term Treasury securities and agency MBS."[22]

In the spirit of this work, monetary policy may be partially or fully effective at the ELB, and shadow rates may or may not completely reflect the stance (however defined) of monetary policy at the ELB. In either case, our VAR estimates map observed economic and financial indicators into shadow rate measures that reflect the historical relationships between interest rates and other economic and financial indicators. As such, the shadow-rate estimates are designed to be state variables useful for modeling the evolution of interest rates and other economic variables (within otherwise linear VAR systems). When the ELB is binding and the Federal Reserve is using balance sheet policies or forward guidance on the future funds rate to provide policy accommodation, the shadow rate becomes latent and can fall below the actual interest rate to capture the effects of these policies. Other studies, including Billi (2020), Gust, et al. (2017), and Reifschneider and Williams (2000), have suggested that the inclusion of lagged shadow rates as VAR predictors could track make-up policies at the ELB. For all these reasons, we expect that the shadow rate of our VAR specifications broadly captures the overall stance of monetary policy when the ELB binds and monetary policy counters by using asset purchases and forward guidance on future rates to provide stimulus, but we cannot claim that it necessarily entirely captures these alternative policies.

A number of studies have developed structural models that have implications for the measurement of the stance of policy and the specifications of VARs. Sims and Wu (2021) develop a DSGE framework that shows that asset purchase programs can be analogous to large reductions in the federal funds rate and that asset purchases account for much of the fall in shadow-rate estimates such as Wu and Xia (2016) during the Great Recession. Motivated by various empirical studies that have concluded that conventional and

[21] Using a smaller model in an unobserved components form, Johannsen and Mertens (2021) identify monetary policy shocks from surprises to the shadow rate, using short-run restrictions.

[22] See p.30 of the March 2009 Bluebook. The Bluebook or Tealbook continued for several years to report unconstrained funds rate measures that were negative; as a later example, see the Tealbook Book B from January 2013. Both references are available at `https://www.federalreserve.gov/monetarypolicy/files/FOMC20090318bluebook20090313.pdf` and `https://www.federalreserve.gov/monetarypolicy/files/FOMC20130130tealbookb20130124.pdf`.

unconventional monetary policies work in a similar fashion, Wu and Zhang (2019) derive a DSGE model in which alternative policy tools circumvent the ELB. In their setup, the stance of monetary policy can be summarized with a shadow rate, and a linear model can represent the economy — without an ELB-induced structural break. Consistent with this reasoning, Francis, Jackson, and Owyang (2020) find that linear VARs estimated with shadow-rate estimates from Krippner (2013, 2015) and Wu and Xia (2016) as a measure of monetary policy (unlike models that instead use the federal funds rate) pass tests of parameter stability and yield stable impulse response estimates – as predicted by Wu and Zhang (2019).

On the other hand, as discussed in Aruoba, et al. (2022), various studies using DSGE models have established that occasionally binding ELB constraints result in non-linear decision rules for consumption and prices with respect to the underlying state variables. Related, in recent work with structural VARs, including Aruoba, et al. (2022) and Mavroeidis (2021), a binding ELB constraint leads to a regime shift in coefficients, implying reduced-form models with time-varying parameters. The results of Mavroeidis (2021) for a small-scale SVAR in inflation, unemployment, and the federal funds rate suggest that unconventional policy is "only partially effective" in mitigating the ELB. Relatedly, the models of Kulish, Morley, and Robinson (2017) and Mavroeidis (2021) also distinguish between reduced-form shadow rates (which drive the censoring of actual interest rates) and structural shadow rates, which reflect the stance of monetary policy.

To allow for potential time variation in parameters of the reduced-form VAR, our shadow rate models could be represented more generally as time-varying parameter VARs with stochastic volatility in the tradition of Cogley and Sargent (2005), Cogley, Primiceri, and Sargent (2010), and Primiceri (2005). However, these methods are commonly applied with small models and difficult to use with variable sets as large as ours. In addition, identification of time-varying slope coefficients may become an issue since the shadow-rate components of $z_t$ are latent when the ELB binds, in particular for the purpose of producing reliable out-of-sample forecasts while ELB data is yet scarce. In our forecasting results, to entertain ELB-driven parameter change while constraining its extent to make estimation and forecasting feasible, we consider a version of the hybrid VAR extended to allow the impact matrix $A$ to change at the ELB. However, in out-of-sample forecasting results, presented in the supplementary online appendix for brevity, the model accommodating parameter change in this way performs poorly compared to our baseline constant-parameter models with a shadow rate. Empirical evidence of time variation in the parameters controlling the shadow-rate VAR dynamics, also reported in the supplementary online appendix, is rather limited as well.

Overall, our baseline assumption of constant parameters can be seen as having justification in the data, perhaps due to the rather large size of our VARs. Moreover, the constant-parameter version of our shadow-rate VAR in equation (4) can be seen as consistent with previous work in the literature that has concluded that monetary policy was unconstrained by the ELB (for example, through the use of unconventional policies) so that economic dynamics remain unaffected by the ELB. In addition to the studies noted above, this work includes empirical analysis of the response of bond yields to economic news by Swanson and Williams (2014) and evidence on the stability of macroeconomic volatility and responses to shocks, along with consistency with a DSGE model specification, in Debortoli, Gali, and Gambetti (2019).

### 3.6 Estimation, forecasting, and impulse response computation

All of our models are estimated with an MCMC sampler that builds on the methods of Carriero, Clark, and Marcellino (2019) for large BVAR-SV models (as corrected in Carriero, et al. (2022a) and with details provided therein) and extended to handle the ELB. As in their work, we use a Minnesota prior for the VAR coefficients $C_j$ and follow their other choices for priors as far as applicable, too.[23]

Throughout, we use $p = 12$ lags in a monthly data set, which is described in further detail in Section 4. Here we briefly explain the specifics of the sampler that pertain to handling the shadow rate as a latent process whose posterior is truncated *from above* when the ELB binds. Further details of the sampling procedures are provided in Appendix A. For ease of illustration, we focus here on the simple shadow-rate VAR. As described in Appendix A.4, the extension to the hybrid shadow-rate VAR proceeds with minimal adjustments.

Provided that data on $s_t$ and thus $z_t$ were always observed, the simple shadow-rate VAR in (4) would be straightforward to estimate with existing Bayesian MCMC methods for VARs. However, when the data include observations for which the ELB is binding, not only does $s_t$ become a latent variable, but also it is subject to the constraint that $s_t \leq ELB$ when $i_t = ELB$.

The simple shadow-rate VAR system consisting of (3) and (4) belongs to a class of conditionally Gaussian unobserved components models, for which Johannsen and Mertens (2021) have derived a generic shadow-rate sampling approach that can be nested inside an otherwise standard MCMC sampler for the VAR estimation. The Johannsen-Mertens approach employs the conditionally linear, Gaussian structure of the model to derive a truncated normal posterior for the vector of unobserved shadow rates in the system, given draws of other model parameters, such as the VAR coefficients $C_j$, and the stochastic volatilities captured by $\Sigma_t$.[24] Crucially, this truncated normal posterior pertains to the entire trajectory of unobserved shadow rates (or the ensemble of trajectories in the case of multiple ELB periods), necessitating draws from a multivariate truncated normal. Johannsen and Mertens (2021) successfully employ rejection sampling to generate joint draws from this multivariate shadow-rate posterior. However, in more general applications, rejection sampling can become computationally tedious and highly inefficient.[25]

Specifically, consider the following setup for the simple shadow-rate VAR given by (3) and (4): Values for the VAR coefficients $\{C_j\}_{j=1}^{p}$ and error variances $\{\Sigma_t\}_{t=1}^{T}$ are given and the data for $\{x_t\}_{t=1}^{T}$ are known. We further assume that at $t = 1$, $p$ lags of the data for $x_t$ are known, and that the initial $p$ lags of the shadow-

[23] All VAR coefficients, $C_j$, have independent normal priors; all are centered around means of zero, except for the first-order own lags of certain variables as listed in Table 1. As usual, different degrees of shrinkage are applied to own- and cross-lag coefficients. Prior variances of the $j$th-order own lag are set to $\theta_1/j^{\theta_4}$. The cross-lag of the coefficient on variable $m$ in equation $n$ has prior variance equal to $\theta_1/j^{\theta_4} \cdot \theta_2 \cdot \hat{\sigma}_n^2/\hat{\sigma}_m^2$. The intercept of equation $n$ has prior variance $\theta_3 \cdot \hat{\sigma}_n^2$. In all of these settings, $\hat{\sigma}_n^2$ is the OLS estimate of the residual variance of variable $n$ in an AR(1) estimated over the entire sample. The shrinkage parameters are $\theta_1 = 0.2^2$, $\theta_2 = 0.5^2$, $\theta_3 = 100$, and $\theta_4 = 2$.

[24] For the remainder of this section, references to the shadow-rate posterior are understood as pointing to the posterior distribution of shadow rates conditional on other model parameters and other latent states, such as the sequence of the time-varying variance of covariance matrices for the residuals, $\{\Sigma_t\}_{t=1}^{T}$.

[25] For example, in an application like ours with monthly data for the US covering the years 2009 through 2015, the shadow-rate posterior is a multivariate truncated normal with 85 elements, necessitating a rejection whenever a single element out of these 85 should lie above the ELB. For illustrative purposes, consider the case where the shadow rate draws were $iid$ with an individual probability of being below $ELB$ of 0.95. The probability of all 85 draws lying below the ELB is then merely $0.95^{85} = 0.01$. Of course, positive serial correlation among adjacent shadow rates would generate an acceptance rate somewhat above 0.01.

rate vector, $s_0$, $s_{-1}$, ..., $s_{-p+1}$, are known.

The shadow-rate $s_t$ is unknown at least for some $t$. For ease of notation, we normalize time subscripts so that the first time the ELB is binding occurs at $t = 1$. In addition, denote the last ELB observation by $T^* \leq T$ (where $T$ is the length of the data sample), so that $s_t$ is unknown for $1 \leq t \leq T^*$. For simplicity we refer to the entire sequence $\{s_t\}_{t=1}^{T^*}$ as "unobserved," which corresponds to the case of a single ELB episode. However, the procedures described below also apply when multiple ELB episodes occur between $t = 1$ and $T^*$, so that only some, but not all, values of $s_t$ in this window are unobserved. In addition, we define the vector $\bar{y}_t$ that contains the observed data except for observations of the actual interest rate at the ELB; we have $\bar{y}_t = x_t$ when the ELB is binding, and $\bar{y}_t = [x_t' \, s_t']'$ otherwise. As noted above, the vector of all observed variables is $y_t$.

For ease of reference, we collect all unobserved shadow rates in a vector $\boldsymbol{S}$ and all observations of $\bar{y}_t$ in a vector $\bar{\boldsymbol{Y}}$, and observations of $y_t$ in a vector $\boldsymbol{Y}$:[26]

$$\boldsymbol{S} = \begin{bmatrix} s_{T^*} \\ s_{T^*-1} \\ \vdots \\ s_2 \\ s_1 \end{bmatrix}, \quad \text{and} \quad \bar{\boldsymbol{Y}} = \begin{bmatrix} \bar{y}_T \\ \bar{y}_{T-1} \\ \vdots \\ \bar{y}_0 \\ \bar{y}_{-p+1} \end{bmatrix}, \quad \text{and} \quad \boldsymbol{Y} = \begin{bmatrix} y_T \\ y_{T-1} \\ \vdots \\ y_0 \\ y_{-p+1} \end{bmatrix}. \tag{11}$$

The task of the shadow-rate sampler is then to sample $\boldsymbol{S} \mid \boldsymbol{Y}$, which includes the information that $\boldsymbol{S} \leq ELB$ (where the inequality is element-wise). Following Johannsen and Mertens (2021), the shadow-rate sampler builds on solving the "missing value" problem of characterizing $\boldsymbol{S} \mid \bar{\boldsymbol{Y}}$.[27] The missing-value problem does not condition on information that the ELB has been binding for certain observations, and thus does not impose $\boldsymbol{S} \leq ELB$. As shown in Johannsen and Mertens (2021), the linear structure of the model and its Gaussian error distribution results in a posterior of the missing-value problem that is a multivariate normal, and truncation of the missing-value posterior at the ELB leads to the solution of the shadow-rate sampler:[28]

$$\boldsymbol{S} \mid \bar{\boldsymbol{Y}} \sim N(\boldsymbol{\mu}, \boldsymbol{\Omega}) \tag{12}$$

$$\Rightarrow \quad \boldsymbol{S} \mid \boldsymbol{Y} \sim TN(\boldsymbol{\mu}, \boldsymbol{\Omega}, -\infty, ELB) \, . \tag{13}$$

The moments $\boldsymbol{\mu}$ and $\boldsymbol{\Omega}$ can be recursively computed using a standard Kalman smoother, and draws can be generated via a corresponding smoothing sampler. Our paper extends the Johannsen-Mertens approach to a generic VAR with details provided in Appendix A.

A further contribution of our paper is the implementation of the shadow-rate sampler via Gibbs sampling,

[26] The vector $\boldsymbol{S}$ is intended to capture only *unobserved* shadow rates. In the case of a single ELB episode lasting from $t = 1$ through $T^*$, $\boldsymbol{S}$ consists of the entire sequence $\{s_t\}_{t=1}^{T^*}$. In the case of multiple ELB episodes, observations where $s_t = i_t > ELB$ are excluded from entering $\boldsymbol{S}$.

[27] As described in Appendix A.4, the approach is easily extended to the case of a hybrid shadow-rate VAR, by including actual-rate values at the ELB as exogenous VAR regressors in the conditioning set of the missing value problem.

[28] The notation $\boldsymbol{S} \sim TN(\boldsymbol{\mu}, \boldsymbol{\Omega}, a, b)$ denotes a truncated multivariate normal distribution for the random vector $\boldsymbol{S}$, with typical elements $s_j$, where $a \leq s_j \leq b \; \forall \, j$, and where $\boldsymbol{\mu}$ and $\boldsymbol{\Omega}$ are the mean vector and variance-covariance matrix of the underlying normal distribution.

following Geweke (1991) and Park, Genton, and Ghosh (2007), and adapted to the variance-covariance structure of the VAR($p$) case, rather than the rejection sampling employed by Johannsen and Mertens (2021). Depending on parameter values, a (well-known) issue with rejection sampling from the truncated normal is a possibly low acceptance rate. In our case, the acceptance probability in sampling from (13) critically depends on VAR parameters and the observed data for macroeconomic and financial variables (other than short-term interest rates). As reported further below, when VAR parameters are drawn from the eventual posterior of our shadow-rate VAR, the acceptance probability that draws from the missing-value problem will lie below the ELB is fairly high. However, this need not be the case in general, and does not hold, for example, when our VAR is estimated while treating observations for short-term interest rates as missing (rather than censored) data when the ELB binds. Our adaptation of the Gibbs sampling approach of Geweke (1991) to the VAR($p$) case, with details described in Appendix A, provides a more efficient solution to the shadow-rate sampling problem. Moreover, our methods prove to be numerically robust due to the use of QR decomposition methods that assure positive definiteness in computation of second-moment matrices.

Our approach also builds on the work of Chib (1992) and Chib and Greenberg (1998) for Tobit and Probit models, respectively, which also relied on algorithms featuring data augmentation.[29]

Notably, non-linearities that arise from the censoring of actual rates do not impede shadow-rate inference based on Gaussian signal extraction methods. In the simple shadow-rate VAR, the dynamics of latent shadow-rate variables are described by a Gaussian VAR, which gives rise to a missing value problem with a (conditionally) normal posterior distribution. The moments of the Gaussian posterior for the missing value problem can be obtained from standard Kalman filtering and smoothing. Similarly, in the hybrid shadow-rate VAR, censored actual-rate values enter only as lagged explanatory variables (and only in non-shadow-rate equations), retaining the conditionally Gaussian structure of the missing value problem.

In out-of-sample forecasting, for every model considered, we generate draws from the predictive density of $y_{t+k}$ at forecast origin $t$ by recursive simulations. In each case, to generate draws from the $h$-step-ahead density, VAR residuals, $v_{t+k}$, are drawn for $k = 1, 2 \ldots, h$.[30] In the case of the standard VAR, conditional on current and lagged data for $y_t$, the simulation is standard and iterates over (2). In contrast, for the simple shadow-rate VAR, simulation of the predictive densities jumps off MCMC draws for $s_t$, $s_{t-1}$, $\ldots s_{t-p+1}$ that are used to initialize recursions over (4). Censoring of predicted interest rates is applied only at the level of the measurement equation (1), while uncensored draws of lagged shadow rates are fed into the VAR equation (4) to simulate subsequent predictions of $y_{t+k}$.

Our development of an efficient Gibbs sampler for Bayesian estimation of shadow-rate VARs — to obtain both shadow rate estimates and out-of-sample forecasts — represents another contribution of the paper. Mavroeidis (2021) instead uses particle filtering to obtain maximum likelihood estimates of a structural VAR in the face of ELB constraints. Aruoba, et al. (2022) develop a sequential Monte Carlo sampler for Bayesian estimation of a structural VAR with an occasionally-binding constraint that leads to shifts in coefficients.

[29] When the ELB binds, the censoring constraint on actual rates bounds the posterior distribution for the shadow rate. The censored shadow-rate problem is related, but different from settings where state variables have bounded support under the prior, as studied, for example, by Chan, Koop, and Potter (2013) and Koop and Potter (2011).

[30] As described in, for example, Carriero, Clark, and Marcellino (2019), draws from $v_{t+k} \sim N(0, \Sigma_{t+k})$ are conditioned on an MCMC draw of the underlying model parameters and SV states and involve forward simulation of the SV processes.

To estimate the responses of the economy to shocks, rather than computing standard impulse response functions conditional on whether the ELB binds, we rely on the generalized impulse response approach of Koop, Pesaran, and Potter (1996), in order to allow for the possibility that the shock effects lead to a change in whether the ELB binds. In these computations, we start with posterior distributions for VAR coefficients, shadow-rate values, and volatilities estimated with the full sample of data. At a given forecast origin $t$, we simulate a predictive density that conditions on a specific realization of the shock at $t+1$ and compare its predictive mean against the average of a baseline prediction, made at $t$ but otherwise unconditional. To illustrate the working of this procedure in our empirical application, we fix the size of the shock, denoted by $\sigma$, equal to the time-series average of the stochastic volatility estimates for 2005-2006, the two years preceding the Great Recession.[31] We then estimate the responses to a specific shock of size $\sigma$ that occurs at $t+1$.[32] Building on our earlier notation, let $v_t = A^{-1}\varepsilon_t$ denote the vector of VAR residuals, with $\varepsilon_t \sim N(0, \Lambda_t)$, $A$ a unit-lower-triangular matrix, and $\Lambda_t$ a diagonal matrix of stochastic variance values. With EBP ordered first in the VAR, $\varepsilon_{1,t}$ denotes the (scalar) shock to EBP at time $t$. Our density simulations take MCMC draws of parameters, as well as values for shadow rates and stochastic volatilities at time $t$ as given, and then simulate forward shocks to log volatility for period $t+1$ through $t+H$ (with $H$ = 48 months), as well as the standardized VAR residuals $\varepsilon_{t+h}$. For the baseline simulation, all shocks are drawn freely. For the alternative simulation, we set the shock to EBP at time $t+1$ equal to a value of $\sigma$ while residuals to all other variables are drawn conditional on that value. At each MCMC node $m$ we generate $J = 1000$ draws of the predictive densities under the baseline and the alternative, and denote their means for outcomes $h$ steps ahead by $E_t^{(m)}(y_{t+h})$ and $E_t^{(m)}(y_{t+h}|\varepsilon_{1,t} = \sigma)$. To compute the generalized impulse functions (GIRFs), we integrate over all MCMC draws, and obtain

$$\Psi_{t,h}(\sigma) = \sum_m \Psi_{t,h}^{(m)}(\sigma)\,, \tag{14}$$

$$\text{with} \quad \Psi_{t,h}^{(m)}(\sigma) \equiv E_t^{(m)}\left(y_{t+h} \,\middle|\, \varepsilon_{1,t} = \sigma\right) - E_t^{(m)}\left(y_{t+h}\right)\,. \tag{15}$$

Uncertainty bands for the GIRF reflect the distribution of $\Psi_{t,h}^{(m)}$ across MCMC nodes $m$ in (15).

# 4 Data

Our data set consists of monthly observations for either 15 or 20 macroeconomic and financial variables for the sample 1959:03 to 2022:08, taken from the September 2022 vintage of the FRED-MD database maintained by the Federal Reserve Bank of St. Louis. The baseline model's variables and their transformations to logs or log-differences are listed in Table 1. Reflecting the raw sample, transformations, and lag specification, the sample for model estimation always begins with 1960:04. Critically, the data set includes the federal funds rate, which was constrained by the ELB from late 2008 through late 2015 and from March 2020 through February 2022. In our main results using 20 variables, the data set includes the federal funds

[31] Estimated values of the EBP shock's stochastic volatility are given by its posterior median path, which averaged at a value of 0.11 over the years 2005-2006.

[32] As robustness checks, our supplementary online appendix also reports results for different shock sizes $k \times \sigma$, where $k$ = 1,5,10, as well as shocks that lead to a decrease in EBP (or VXO) rather than an increase as used here.

Table 1: List of variables

| Variable | FRED-MD code | Transformation | Minnesota prior |
|---|---|---|---|
| Real Income | RPI | $\Delta \log(x_t) \cdot 1200$ | 0 |
| Real Consumption | DPCERA3M086SBEA | $\Delta \log(x_t) \cdot 1200$ | 0 |
| IP | INDPRO | $\Delta \log(x_t) \cdot 1200$ | 0 |
| Capacity Utilization | CUMFNS | | 1 |
| Unemployment | UNRATE | | 1 |
| Nonfarm Payrolls | PAYEMS | $\Delta \log(x_t) \cdot 1200$ | 0 |
| Hours | CES0600000007 | | 0 |
| Hourly Earnings | CES0600000008 | $\Delta \log(x_t) \cdot 1200$ | 0 |
| PPI (Fin. Goods) | WPSFD49207 | $\Delta \log(x_t) \cdot 1200$ | 1 |
| PPI (Metals) | PPICMM | $\Delta \log(x_t) \cdot 1200$ | 1 |
| PCE Prices | PCEPI | $\Delta \log(x_t) \cdot 1200$ | 1 |
| Federal Funds Rate | FEDFUNDS | | 1 |
| Housing Starts | HOUST | $\log(x_t)$ | 1 |
| S&P 500 | SP500 | $\Delta \log(x_t) \cdot 1200$ | 0 |
| USD / GBP FX Rate | EXUSUKx | $\Delta \log(x_t) \cdot 1200$ | 0 |
| 6-Month Tbill | TB6MS | | 1 |
| 1-Year Yield | GS1 | | 1 |
| 5-Year Yield | GS5 | | 1 |
| 10-Year Yield | GS10 | | 1 |
| BAA Yield | BAA | | 1 |

Note: Data obtained from the 2022-09 vintage of FRED-MD. Monthly observations from 1959:03 to 2022:08. Entries in the column "Minnesota prior" report the prior mean on the first own-lag coefficient of the corresponding variable in each BVAR. Prior means on all other VAR coefficients are set to zero.

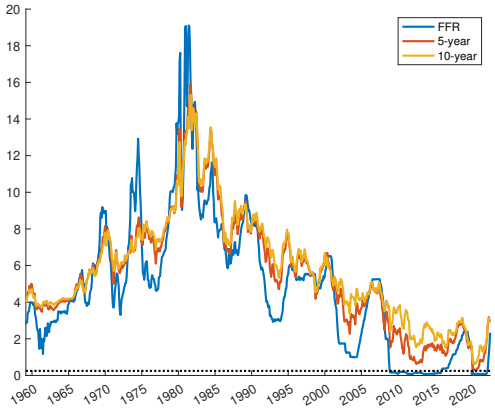
rate and five additional interest rates: two other rates constrained by the ELB, the 6-month Treasury bill rate and the yield on 1-year Treasuries, as well as three longer-maturity bond yields, including 5- and 10-year Treasuries and Moody's Seasoned BAA corporate bond yield. For comparability to some other work in the literature, we also provide shadow-rate VAR estimates for a 15-variable subset of the data listed in Table 1 in which the federal funds rate is the only interest rate measure; specifically, this subset omits the five yield measures listed at the bottom of the table. For the paper's application to structural analysis, we use the excess bond premium measure of Gilchrist and Zakrajšek (2012) from January 1973 to 2022:08, obtained from the website of the Federal Reserve Board of Governors.[33]

Data for short-term rates and longer-maturity Treasury yields (with some omissions for chart readability) are shown in Figure 1. During and following the Great Recession, while short-term interest rates were constrained by the ELB, longer-term bond yields remained solidly or well above the ELB. The 10-year (5-year) Treasury yield declined from 2.4 percent (1.5 percent) in December 2008 to a low of 1.5 percent (0.6 percent) in July 2012 and then moved higher. Following the COVID-19 outbreak and the FOMC's quick

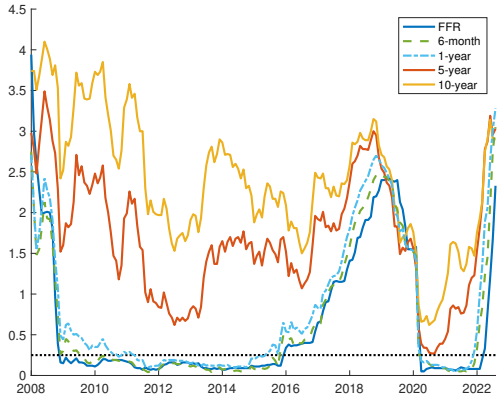
[33]The data are available at `https://www.federalreserve.gov/econresdata/notes/feds-notes/2016/files/ebp_csv.csv`.

Figure 1: Interest rate data

(a) Full data sample




(b) Since the GFC




Note: All interest rates quoted as annualized percentage rates. Data obtained from FRED-MD as listed in Table 1.

and substantial easing of monetary policy, bond yields were lower and much closer to the ELB than they were following the Great Recession. From April through December 2020, the 10-year (5-year) Treasury yield averaged 0.7 (0.3) percent. In 2021, the 10-year (5-year) Treasury yield averaged 1.4 (0.9) percent. In January 2022, the FOMC signaled its inclination to raise the federal funds rate off the ELB at its next meeting in March, after which bond yields moved higher.

In our application with US postwar data, the value of $ELB$ is set to 25 basis points, which was the upper end of the FOMC's target range for the federal funds rate between late 2008 and 2015, and was again from mid-March 2020 through mid-March 2022. We treat a given interest rate — at any maturity — as unconstrained unless it reaches the $ELB$. As a matter of consistency with this convention, we set readings for the federal funds rate, 6-month T-bill rate, and 1-year Treasury yield to 25 basis points when estimating shadow-rate VARs (not when including these rates in a standard VAR that ignores the lower bound constraint). Treasury yields with maturities of five years and longer and corporate bond yields stayed above 25 basis points in the data and can thus be treated as part of the vector $x_t$, defined in Section 3, for the purpose of model estimation.[34] The supplementary online appendix shows that our main results are robust to instead setting the value of the $ELB$ to 12.5 or 50 basis points. Admittedly, while some might argue that, even if longer-term bond yields did not actually hit the $ELB$, they were at least somewhat constrained when short-term rates hit the ELB, such a constraint falls outside our shadow-rate VAR framework. But a couple of considerations could be seen as supporting our approach: First, the hybrid model allows for changes in the predictive densities of all variables, when interest rates are above but close to the ELB (relative to the case when all rates were so high that the prospect of a binding ELB were negligible). The reason is that actual (and not shadow) rates matter in the forecasting equations for non-interest-rate variables (which in turn

[34]The lower bound constraint is an issue when simulating the predictive density for these yields, but it is not relevant for estimating the VAR.

also affect projections for future shadow and actual rates). These mechanics are present, and show some effect, in our forecasting results, but also our structural analysis based on generalized impulse responses. Second, a more generic, or even a more structural, approach could have been chosen to model interest-rate dynamics near (but above) the ELB, but both approaches have their drawbacks as well. A purely empirical approach, based on generic time variation in parameters, comes with its own issues regarding scalability and identification (as shadow rates are latent at the ELB). In contrast, a more structural approach could derive tighter restrictions on such behavior, but at the cost of having to impose more specific assumptions on economic structures.[35]

# 5 Shadow-Rate Estimates

We start our empirical analysis with an examination of the shadow-rate estimates, in order to understand the roles of the ELB and the various yields in the model and subsequently help interpreting the forecasting results presented in the next section.

Figure 2 reports our shadow-rate estimates associated with the federal funds rate, along with comparisons to some other estimates. Panel (a) of the figure compares full-sample estimates from the simple shadow-rate VAR using data through August 2022 (black line with the credible set indicated by gray shading) to quasi-real-time estimates (solid red line with the credible set indicated by dotted lines).[36] The quasi-real-time estimates are the end-of-sample estimates produced by recursive estimation of the model starting in January 2009. Panel (b) shows full-sample and quasi-real-time estimates obtained from a simple shadow-rate VAR that omits bond yields and includes only the federal funds rate as an interest rate measure in its data vector. Panel (c) provides estimates obtained from the hybrid shadow-rate VAR discussed in Section 3.3. Finally, Panel (d) compares our full-sample estimate (posterior median shown as black line with the 90 percent credible set indicated by gray shading) to the shadow-rate measures from Krippner (2013, 2015) and Wu and Xia (2016) based on affine term structure models.

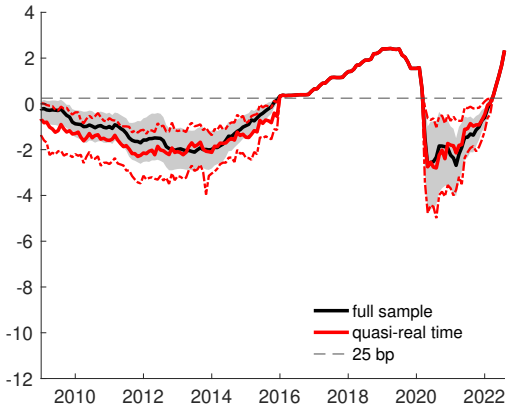
Our baseline full-sample estimates from the simple shadow-rate VAR show the shadow rate dropping sharply starting in 2009, reaching a nadir of $-2.1$ percent in mid-2013. The rate then gradually rose and crossed the ELB in early 2016, following the Federal Reserve's first increase in the federal funds rate in December 2015 (when the FOMC raised the target range from 0-25 basis points to 25-50 basis points). The rate dropped precipitously in the spring of 2020, with the posterior median reaching $-2.7$ percent in May 2020. The rate moved gradually higher starting in April 2021 and crossed the ELB in April 2022, following the FOMC's first increase in the federal funds rate in March 2022. As might be expected, the quasi-real-time estimates have more time variability than do the full sample estimates, but follow a quite similar contour. As might also be expected, the quasi-real-time estimates are less precise, with credible sets wider than those of the full sample estimate. The supplementary online appendix provides full-sample and quasi-real-time

[35] For example, a no-arbitrage shadow-rate model in finance, would commonly assume linear and time-invariant state dynamics, and a specific set of pricing factors. Moreover, for our purposes, such pure finance model would still be silent on specific sensitivities of macroeconomic variables when interest rates are near the ELB, as captured by our hybrid VAR.

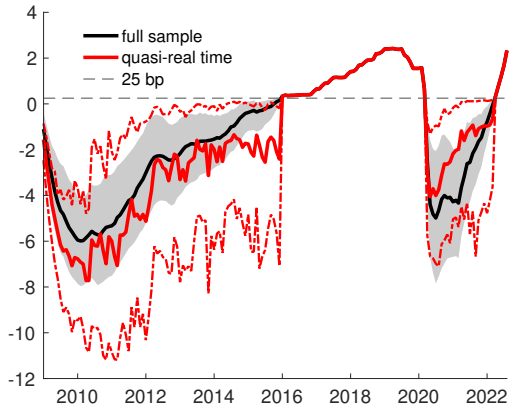
[36] To be clear, in the full sample case, the model is estimated with data for 1960:04 through 2022:08, but the figure omits the period of 1960-2008 during which the ELB did not bind.

Figure 2: Shadow-rate estimates

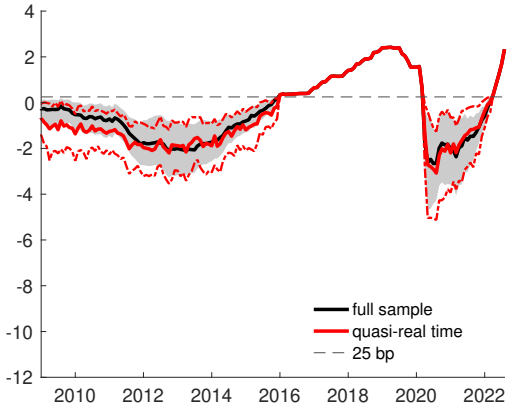
(a) Simple shadow-rate VAR (w/yields)



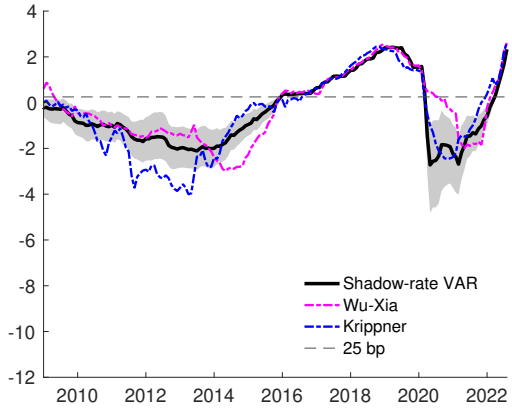

(b) Simple shadow-rate VAR (w/o yields)




(c) Hybrid shadow-rate VAR (w/yields)




(d) Simple shadow-rate VAR and other estimates




Note: Panel (a) compares smoothed and quasi-real time shadow-rate estimates from our simple shadow-rate VAR, while Panel (b) provides corresponding estimates from the simple shadow-rate VAR when data on interest rates other than the federal funds rate are omitted. Panel (c) depicts estimates from the hybrid shadow-rate VAR (using all twenty variables listed in Table 1 in the estimation). Panel (d) compares the smoothed shadow-rate estimates from the simple shadow-rate VAR (estimated with yields, also shown in Panel (a)) against updated estimates obtained from Krippner (2013, 2015) and Wu and Xia (2016). The quasi-real-time estimates are the end-of-sample estimates produced by recursive estimation of the model starting in January 2009. Each estimation conditions on available data since 1959:03, but the figure omits the period prior to 2008 during which the ELB did not bind. Posterior medians are shown as thick lines; grey shaded areas and thin lines depict 90 percent uncertainty bands.

estimates of shadow rates for the 6-month and 1-year Treasury maturities. The contours of these estimates follow those shown for the federal funds rate.

Panel (b) reports the shadow-rate estimate obtained from a VAR in which the only interest rate is the federal funds rate, which can be directly compared to the baseline estimate provided with the same chart scaling in Panel (a). Without bond yields, the simple shadow-rate VAR generates a much more negative

estimate of the shadow rate, particularly in the 2009-2013 period that corresponds more closely to the recovery in real activity and again during the first year of the COVID-19 pandemic. The contours of our shadow-rate estimates generated from macroeconomic and financial variables while excluding term structure data also resemble more closely estimates from Ikeda, et al. (2021), also obtained from macroeconomic data alone, or unconstrained Taylor-rule prescriptions calculated for the Great Recession years by Eberly, Stock, and Wright (2020). It is also notable that, in our model, the exclusion of bond yields makes the quasi-real-time estimate of the shadow rate more variable from quarter to quarter. With bond yields included, Panel (c) shows that estimates based on the hybrid shadow-rate VAR are rather similar to estimates from the simple shadow-rate VAR. As the main difference in the two VARs is in the specifications of the macro equations, this finding confirms that the shadow rate estimates are mainly informed by the interest rate equations.

Finally, although our shadow-rate VARs do not impose the restrictions of an affine term structure model, our shadow-rate estimates have some similarities to the Krippner and Wu-Xia measures based on affine term structure models. As indicated in Panel (d) of Figure 2, our simple shadow-rate VAR estimate and the Wu-Xia series move together from 2009 through 2013. Over the remainder of the ELB episode following the Great Recession, as our estimate gradually rose to the ELB over the course of 2014 and 2015, the Wu-Xia series fell and then rose sharply. Compared to the other estimates, the Krippner measure fell more sharply in 2012-2013 and then bounced back more quickly in 2014-2015. For the period from 2009 through 2015, our simple shadow-rate VAR estimate can be seen as akin to a measure one would obtain by averaging the Krippner and Wu-Xia measures. From 2020 through mid-2022, our shadow-rate VAR estimate is more similar to the Krippner measure than the Wu-Xia estimate in the period of falling rates but a little more similar to the Wu-Xia measure than the Krippner estimate in the period of rising (but still below the ELB) shadow rates.

Figure 3 provides some comparisons to assess the effects of shadow-rate modeling and enforcement of the ELB in model estimation. Focusing on the federal funds rate and its shadow rate, Panel (a) compares shadow-rate (black) and missing-data (red) draws for the shadow rate $s_t$ obtained from the posterior of our simple shadow-rate VAR. Shadow-rate draws are obtained from the truncated posterior for $s_t$ that satisfies the ELB, and described by the problem of drawing from $\boldsymbol{S}|\boldsymbol{Y}$ in (13). Missing-data draws are obtained from the underlying (and untruncated) posterior of the missing data problem, which ignores the ELB, and correspond to draws from $\boldsymbol{S}|\bar{\boldsymbol{Y}}$ in (12). As the chart comparison indicates, the posteriors obtained from these alternatives are very similar.[37] The supplementary online appendix shows similar results obtained from the hybrid shadow-rate VAR.

These results might suggest that an approach that treats observed policy rates at the ELB as missing values might be a close alternative to shadow-rate sampling that explicitly accounts for the ELB.[38] However, such a conclusion would neglect the effects of enforcing the ELB as part of the shadow-rate sampling on inference for other VAR parameters and state variables (like SV). As discussed by Waggoner and Zha (1999) in the context of conditional forecasting, conditioning estimates on information when the ELB was binding

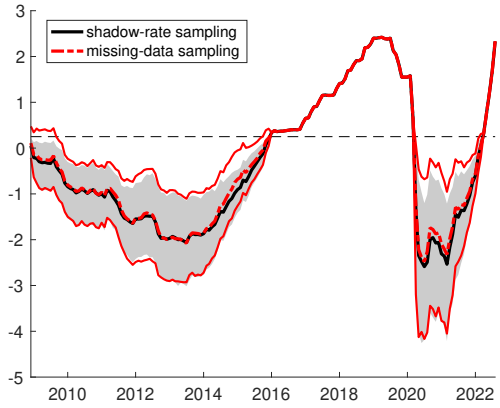
[37] Late in the Great Recession episode and in the early months of the COVID-19 episode, using a missing data approach without fully enforcing the ELB led to draws of interest rates above the ELB, whereas with the full shadow-rate treatment, the interest rate distributions remained at or below 25 basis points.

[38] Such a missing data approach has been used by, for example, Del Negro, et al. (2017).

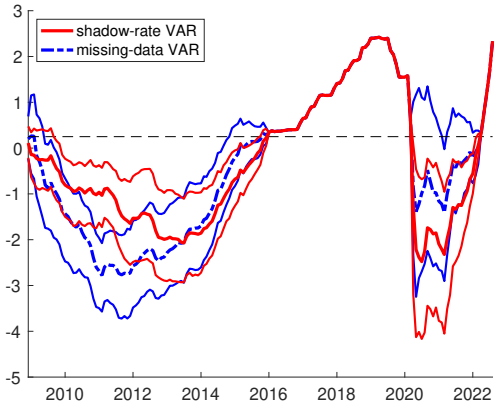
Figure 3: Effect of imposing ELB on shadow-rate estimates

**Simple shadow-rate VAR estimates from data including yields**

(a) Missing-data and shadow-rate draws

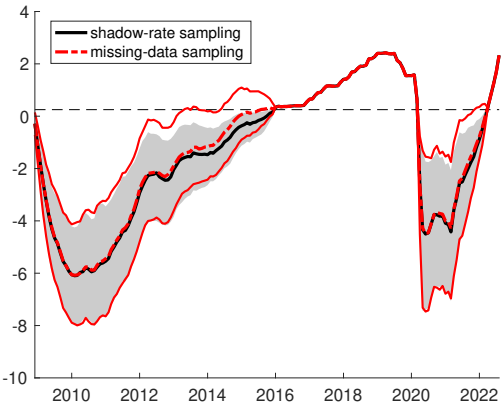



(b) Missing-data draws w/and w/o shadow-rate sampling

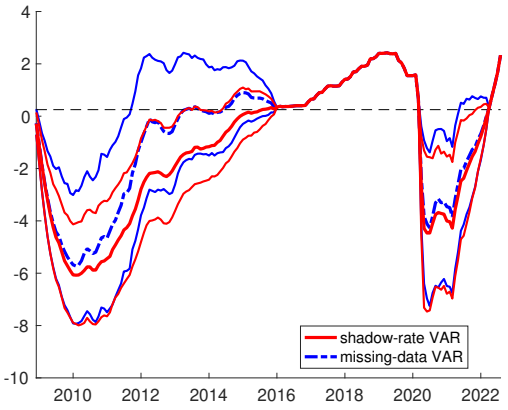



**Estimates from data excluding yields**

(c) Missing-data and shadow-rate draws




(d) Missing-data draws w/and w/o shadow-rate sampling




Note: Panel (a) compares shadow-rate (black) and missing-data (red) draws for $s_t$ obtained from the posterior of our simple shadow-rate VAR. Shadow-rate draws are obtained from the truncated posterior for $s_t$ that satisfies the ELB. Missing-data draws are obtained from the underlying (and untruncated) posterior of the missing data problem that ignores the ELB. Panel (b) displays missing-data posteriors obtained from two sets of simple shadow-rate VAR estimates: In the baseline (red), parameter and SV draws reflect shadow-rate sampling. In the alternative version (blue), parameters and SV are drawn while treating the policy rate at the ELB as missing data and without requiring that missing data draws lie below the ELB. Panels (c) and (d) provide corresponding estimates from a VAR that omits all interest rate data except for the federal funds rate. Each panel shows reports medians (thick lines) and 90 percent uncertainty bands (grey shaded areas or thin lines).

could (and should) embody non-trivial information about the relevant parameters of the VAR.

To illustrate these effects, Panel (b) compares missing-data posteriors from two sets of estimates obtained from simple shadow-rate VARs: In the baseline (red), parameter and SV draws reflect shadow-rate sampling

(as also shown in Panel (a)). In the alternative version (blue), parameters and SV are drawn while treating the policy rate at the ELB as missing data and without requiring that missing data draws lie below the ELB. The comparison highlights the non-negligible effects of shadow-rate sampling, which takes into account observations of interest rates at the ELB, on model estimates of parameters and SV. Without forcing the draws of missing interest rate observations to lie at or below the ELB, the upper bound of the posterior credible set rises above the ELB in 2014-2015 and again in 2020-2021, which contradicts observations of the federal funds rate that were at the ELB during those times. In contrast, the use of shadow-rate sampling, as opposed to a missing-data approach, leads to estimates of parameters and SV that increase the odds of obtaining missing-data draws for the shadow rate that lay below the ELB (for observations when the ELB binds).

A similar pattern is observed when the estimation excludes all data on yields (except for the federal funds rate), as shown in Panels (c) and (d) of the figure. With the federal funds rate the only interest rate in the model, when the draws of missing interest rate observations are not forced to lie at or below the ELB, the upper bound of the posterior credible set shows an even longer period above the ELB, from about 2012 through 2015 and again in 2020. With or without additional yields in the model, we find that shadow-rate sampling (as opposed to missing data sampling) increases the odds of drawing shadow-rate values below the ELB by adjusting the SV path for shocks to the shadow-rate equation, with only minor effects on the VARs' transition coefficients; details are provided in the supplementary online appendix.

# 6 Forecast Evaluation

We conduct an out-of-sample forecast evaluation in quasi-real time, where we simulate forecasts made from January 2009 through December 2017. For every forecast origin, each model is re-estimated based on growing samples of data that start in 1959:03. We stop forecasting in December 2017 to be sure the unusual volatility of the COVID-19 pandemic does not distort forecast comparisons; with a maximum forecast horizon of 24 months, the last outcome date in the evaluation sample is December 2019, so that our evaluation sample does not include any realizations from 2020 or later. However, as shown in the supplementary appendix, ending the evaluation sample in mid-2022 yields similar results. Of course, the evaluation window in our main results is relatively short and largely informed by a single ELB episode. Forecasts made prior to 2009 are not considered, due to the absence of observed interest rates at the ELB in postwar US data. All data are taken from the September 2022 vintage of FRED-MD; we abstract from issues related to real-time data collection. For sake of comparability, forecasts from linear and shadow-rate VARs are compared against realized interest-rate values that are censored at the ELB.

## 6.1 Average performance 2009–2017

Comparing various model specifications discussed below, Tables 2 and 3 provide results on point and density forecast accuracy, measured by root mean squared error (RMSE, computed around mean forecasts), mean absolute error (MAE, computed around median forecasts), and continuous ranked probability score (CRPS), respectively. We provide the MAE results in light of the concerns of Bauer and Rudebusch (2016) with the

use of mean forecasts for interest rates near the ELB constraint. The reported forecast horizons are $h$ = 3, 6, 12, and 24 months. For those variables that enter the model in monthly growth rates (e.g., real income and nonfarm payrolls), at horizons $h$ greater than 1 month, the $h$-step forecasts are transformed to average growth rates over $h$ periods.[39]

Table 2 compares the accuracy of forecasts from our simple shadow-rate VAR to those from a standard VAR that simply takes the forecasts as given and does nothing to obey ELB constraints. To facilitate comparisons, we report RMSE, MAE, and CRPS results for the simple shadow-rate model as relative to the standard VAR, so that entries of less (more) than 1 mean the shadow rate model's forecast is more (less) accurate than the baseline. To roughly gauge the significance of differences with respect to the baseline, we use $t$-tests as in Diebold and Mariano (1995) and West (1996), denoting significance in the tables with asterisks.

Overall, the results in Table 2 indicate that our simple shadow-rate VAR specification for accommodating the ELB performs better in forecasting than does a standard VAR. The shadow-rate specification significantly improves forecasts of not only the federal funds rate (FFR) but also most other interest rates, without harming the forecasts of indicators of economic activity, measures of inflation, and other financial indicators. RMSE ratios for federal funds rate forecasts for $h$ = 3, 6, 12, and 24 range from 0.22 to 0.55, the MAE ratios range from 0.13 to 0.40, and the CRPS ratios range from 0.15 to 0.44, with statistical significance of all of the MAE and CRPS gains and one-fourth of the RMSE gains. For Treasury rates from 6 months through 10 years, the shadow-rate specification also typically improves forecast accuracy, with gains greater at short maturities than at long maturities. For the 6-month T-bill yield, RMSE (MAE) ratios range from 0.37 to 0.77 (0.24 to 0.56); CRPS ratios are between 0.27 and 0.58, pointing to larger and more statistically significant density gains as compared to RMSE gains. Forecast improvements at the 1-year maturity are modestly smaller but can still be sizable; for example, in this case, CRPS ratios range from 0.54 to 0.64. At the 5- and 10-year maturities, the simple shadow-rate VAR consistently improves on the forecast accuracy of a standard VAR, but by smaller, less significant magnitudes. For instance, the CRPS ratios for the 5-year Treasury yield range from 0.86 to 0.97. Finally, for the indicators of economic activity, stock price returns, the BAA yield, and the exchange rate, the treatment of the ELB on interest rates does not seem to bear consistently and importantly on forecast accuracy. RMSE, MAE, and CRPS ratios for the standard VAR and simple shadow-rate specifications are often near 1, particularly at the 3- and 6-month forecast horizons. At longer horizons, forecasts from the simple shadow-rate VAR tend to be slightly more accurate than those from the standard VAR for housing starts and modestly less accurate for other indicators of economic activity.

Table 3 compares the accuracy of forecasts from the hybrid shadow-rate specification to the standard VAR. As explained above, the hybrid specification departs from the simple shadow-rate VAR by relating the macroeconomic indicators to actual interest rates rather than shadow rates. Over the full sample, the shadow-rate estimates obtained with the hybrid specification are quite similar to the estimates shown earlier for the baseline shadow-rate VAR. Of course, it could nonetheless be the case that predicting macroe-

[39] Specifically, for a variable that enters the model as $1200\Delta\log(x_t)$, at each forecast origin $t$ and multi-step horizon $h$, we average the forecasts of monthly growth rates as $(1200/h)\sum_{s=1}^{s=h}\Delta\log(x_{t+s})$.

Table 2: Forecast performance of standard vs. simple shadow-rate VAR

| | RMSE | | | | MAE | | | | CRPS | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Variable / Horizon | 3 | 6 | 12 | 24 | 3 | 6 | 12 | 24 | 3 | 6 | 12 | 24 |
| Real Income | 1.00 | 1.00 | 0.99 | **0.88** | 0.98** | 0.98 | 1.00 | 0.98 | 0.99* | 0.99* | 1.00 | 1.00 |
| Real Consumption | 1.00 | 1.00 | 1.02 | 0.98 | 1.00 | 1.00 | 1.03 | 0.99 | 1.00 | 1.00 | 1.01 | 1.00 |
| IP | 1.02 | 1.01 | **1.05***** | 1.00 | 1.03** | 1.02 | **1.05***** | 1.01 | 1.02 | 1.02 | 1.04*** | 1.01** |
| Capacity Utilization | 1.03* | **1.05** | **1.09**** | **1.10**** | **1.06***** | **1.08*** | **1.15***** | **1.11*** | 1.03** | **1.05**** | **1.09***** | **1.09***** |
| Unemployment | 1.00 | 1.00 | 1.01 | 1.00 | 1.00 | 0.99 | 0.96 | 0.96 | 1.00 | 1.00 | 1.01 | 1.00 |
| Nonfarm Payrolls | 0.98 | 0.99 | 1.04 | 1.04 | 1.00 | 1.02 | **1.08** | 1.01 | 1.00 | 1.02 | **1.05** | 1.02 |
| Hours | 1.00 | 1.01 | 1.00 | 1.03 | 1.00 | 1.00 | 1.01 | 1.04 | 1.01 | 1.00 | 1.01 | 1.02 |
| Hourly Earnings | 1.00 | 1.01 | 1.00 | 1.01 | 1.00 | 1.01 | 1.00 | 1.01** | 1.00 | 1.00 | 0.99 | 1.00 |
| PPI (Fin. Goods) | 1.00 | 0.99 | 1.00 | 0.99 | 0.99 | 0.98** | 1.00 | 1.01 | 0.99 | 0.99* | 1.00 | 1.00 |
| PPI (Metals) | 1.00 | 1.00 | 1.00 | 1.02 | 0.99 | 1.00 | 1.00 | 1.01** | 1.00 | 1.00 | 1.00 | 1.01* |
| PCE Prices | 1.00 | 1.00 | 1.01 | 1.01 | 1.01 | 1.00 | 1.01 | 1.02* | 1.00 | 1.00 | 1.00 | 1.01* |
| Federal Funds Rate | **0.22**** | **0.24** | **0.31** | **0.55** | **0.13***** | **0.22**** | **0.30*** | **0.40**** | **0.15***** | **0.20**** | **0.27*** | **0.44**** |
| Housing Starts | 1.01 | 0.99 | **0.95** | **0.90** | 1.00 | 0.98 | 0.98 | **0.92** | 1.01 | 0.99 | 0.96 | **0.92** |
| S&P 500 | 0.99 | 1.00 | 1.01 | 1.04 | 0.98 | 0.99 | 0.99 | 1.00 | 0.99 | 1.00 | 1.00 | 1.01* |
| USD / GBP FX Rate | 0.99* | 0.99* | 0.99* | 0.97*** | 0.98* | 0.99 | 0.98* | 0.97** | 0.99** | 0.99* | 0.99* | 0.99*** |
| 6-Month Tbill | **0.37*** | **0.43** | **0.55** | **0.77** | **0.24***** | **0.32**** | **0.40*** | **0.56**** | **0.27***** | **0.33**** | **0.45**** | **0.58***** |
| 1-Year Yield | **0.66** | **0.66** | **0.72** | **0.87** | **0.52***** | **0.48***** | **0.54**** | **0.67**** | **0.54***** | **0.54***** | **0.59**** | **0.64***** |
| 5-Year Yield | **0.93*** | **0.90***** | **0.85***** | **0.84** | **0.94*** | **0.91**** | **0.89***** | **0.90***** | 0.97 | **0.94*** | **0.90***** | **0.86***** |
| 10-Year Yield | 0.96 | **0.93**** | **0.88***** | **0.82**** | **0.95** | **0.95** | **0.89**** | **0.91*** | 0.98 | 0.96 | **0.93*** | **0.94** |
| BAA Yield | 0.98 | 0.99 | 1.02 | **0.93** | 1.00 | 0.99 | 1.03 | 0.98 | 1.00 | 1.01 | 1.02 | 1.04 |

Note: Comparison of "standard linear VAR" (baseline, in denominator) against "simple shadow-rate VAR." Values below 1 indicate improvement over baseline. Evaluation window with forecast origins from 2009:01 through 2017:12 (and outcome data as far as available). Significance assessed by Diebold-Mariano-West test using Newey-West standard errors with $h+1$ lags. Performance differences of 5 percent and more (relative to baseline) are indicated by bold face numbers.

Table 3: Forecast performance of standard vs. hybrid shadow-rate VAR

| | RMSE | | | | MAE | | | | CRPS | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Variable / Horizon | 3 | 6 | 12 | 24 | 3 | 6 | 12 | 24 | 3 | 6 | 12 | 24 |
| Real Income | 1.00 | 1.00 | 1.00 | 1.00 | 0.99 | 0.99 | 1.01 | 0.96 | 1.00 | 1.00 | 1.01 | 0.99 |
| Real Consumption | 1.00 | 1.01 | 1.00 | **0.90**$^{**}$ | 0.99 | 1.02 | 1.03 | **0.90**$^{*}$ | 1.00 | 1.01$^{*}$ | 1.00 | 0.97$^{*}$ |
| IP | 1.01 | 0.99 | 1.01 | 0.98 | 1.02$^{**}$ | 0.99 | 1.00 | 0.97 | 1.01 | 1.00 | 1.00 | 1.00 |
| Capacity Utilization | 1.00 | 0.98 | 0.99 | **0.93**$^{**}$ | 1.00 | 0.98 | 0.98 | **0.93** | 1.00 | 0.99 | 0.99 | 0.96$^{*}$ |
| Unemployment | 1.00 | 1.00 | 0.99 | **0.91**$^{***}$ | 1.00 | 1.00 | 0.99 | **0.91**$^{**}$ | 1.00 | 1.00 | 0.99 | **0.94**$^{***}$ |
| Nonfarm Payrolls | 0.98$^{***}$ | 0.99 | 1.04 | 0.98 | 0.98$^{***}$ | 0.99 | 1.00 | **0.92** | 0.99$^{***}$ | 0.99 | 1.00 | 0.97 |
| Hours | 1.01 | 1.02$^{*}$ | 1.02 | 1.04 | 1.00 | 1.01 | 1.02 | 1.03 | 1.01$^{**}$ | 1.01 | 1.02 | 1.01 |
| Hourly Earnings | 1.01 | 1.01 | 1.01 | **1.05**$^{***}$ | 1.00 | 1.01 | 1.01 | 1.03$^{***}$ | 1.01 | 1.01 | 1.00 | 1.02$^{***}$ |
| PPI (Fin. Goods) | 1.00 | 0.98 | 0.97 | 0.99 | 1.00 | 0.99 | 0.98 | 1.00 | 1.00 | 0.99 | 0.99 | 1.00 |
| PPI (Metals) | 0.99$^{*}$ | 1.00 | 0.99$^{*}$ | 1.01 | 0.99 | 0.99 | 0.99 | 1.00 | 1.00 | 0.99 | 0.99 | 1.00 |
| PCE Prices | 1.00 | 0.98 | 0.97 | **1.06**$^{**}$ | 1.00 | 0.98$^{*}$ | 0.99 | **1.08**$^{***}$ | 1.00 | 0.99 | 0.99 | **1.05**$^{***}$ |
| Federal Funds Rate | **0.22**$^{**}$ | **0.24** | **0.29** | **0.51** | **0.13**$^{***}$ | **0.20**$^{**}$ | **0.28**$^{*}$ | **0.39**$^{**}$ | **0.15**$^{***}$ | **0.19**$^{**}$ | **0.25**$^{**}$ | **0.39**$^{***}$ |
| Housing Starts | 1.00 | 0.98 | **0.90** | **0.91** | 1.00 | 0.97 | **0.91** | **0.90** | 1.00 | 0.99 | **0.94** | **0.94** |
| S&P 500 | 1.00 | 1.00 | 1.01 | **1.05** | 0.99 | 0.99 | 1.01 | 1.01 | 1.00 | 1.01 | 1.01 | 1.01$^{**}$ |
| USD / GBP FX Rate | 1.01 | 1.01 | 1.00 | 1.01 | 1.01 | 1.01 | 1.01 | 1.04 | 1.01 | 1.00 | 1.00 | 1.01 |
| 6-Month Tbill | **0.37**$^{*}$ | **0.41** | **0.52** | **0.69** | **0.24**$^{***}$ | **0.30**$^{**}$ | **0.36**$^{**}$ | **0.48**$^{***}$ | **0.26**$^{***}$ | **0.31**$^{**}$ | **0.41**$^{**}$ | **0.50**$^{***}$ |
| 1-Year Yield | **0.64** | **0.63** | **0.67** | **0.78** | **0.52**$^{***}$ | **0.44**$^{***}$ | **0.45**$^{***}$ | **0.56**$^{***}$ | **0.53**$^{***}$ | **0.50**$^{***}$ | **0.53**$^{***}$ | **0.56**$^{***}$ |
| 5-Year Yield | **0.95** | **0.90**$^{***}$ | **0.81**$^{***}$ | **0.75**$^{*}$ | 0.96 | **0.94** | **0.89** | **0.86**$^{*}$ | 0.98 | **0.94** | **0.88**$^{**}$ | **0.81**$^{***}$ |
| 10-Year Yield | 0.97 | **0.94** | **0.87**$^{**}$ | **0.77**$^{***}$ | 0.97 | 0.97 | **0.88**$^{*}$ | **0.91** | 0.98 | 0.96 | **0.92** | **0.93** |
| BAA Yield | 0.97 | 0.98 | 1.01 | 0.97 | 0.99 | 0.97 | 1.03 | **1.07** | 0.99 | 1.00 | 1.02 | **1.06** |

Note: Comparison of "standard linear VAR" (baseline, in denominator) against "hybrid shadow-rate VAR." Values below 1 indicate improvement over baseline. Evaluation window with forecast origins from 2009:01 through 2017:12 (and outcome data as far as available). Significance assessed by Diebold-Mariano-West test using Newey-West standard errors with $h+1$ lags. Performance differences of 5 percent and more (relative to baseline) are indicated by bold face numbers.

conomic outcomes with actual interest rates instead of shadow rates could make a difference in forecast accuracy. Overall, compared to the benchmark standard VAR, the hybrid shadow-rate VAR is comparable to the simple shadow-rate VAR in forecast accuracy, but with some advantages. The RMSE, MAE, and CRPS ratios in Table 3 show that the hybrid specification significantly improves forecasts of interest rates (relative to the standard VAR), with results very similar to those for the simple shadow-rate VAR. For example, for the 6-month T-bill yield, CRPS ratios are between 0.26 and 0.50, compared to ratios between 0.27 and 0.58 with the simple shadow-rate VAR. For macroeconomic variables, the hybrid specification improves on the longer-horizon accuracy of the standard VAR for some variables (e.g., IP, capacity utilization, and unemployment) and reduces longer-horizon accuracy for some other economic activity indicators (e.g., hours worked). If the hybrid and shadow-rate VAR forecasts are compared directly (results in the supplementary online appendix for brevity), the hybrid forecasts are modestly better for most economic activity indicators (e.g., IP, capacity utilization, and unemployment) at longer horizons and modestly worse only for a few indicators (e.g., hourly earnings and PCE prices), so that overall we prefer the hybrid over the simple shadow-rate VAR as a forecasting model in the presence of the ELB.

In addition, we compare point and density forecasts from our shadow-rate VARs against those obtained from a plug-in approach, where external shadow-rate estimates, specifically updated estimates from Krippner (2013, 2015) and Wu and Xia (2016), are used as "data," in place of the actual short-term interest rate, in an otherwise standard VAR. Following the spirit of the shadow-rate literature, forecasts from the plug-in VAR are simulated without censoring the resulting (shadow) interest rate projections. Similar to the case of our simple shadow-rate VAR, forecasts for nominal interest rates are censored only after the dynamic simulations for all variables (and all horizons) have been done. For brevity, detailed results obtained from such plug-in VARs are tabulated only in the supplementary online appendix. Using shadow-rate "data" either from Krippner or Wu and Xia, we find notable benefits for point and density forecasting from our shadow-rate and hybrid VARs. Forecasts obtained from the plug-in VAR are based on the latest vintages of full-sample estimates for the shadow rates from Krippner and Wu-Xia, so that the shadow-rate estimates abstract from one-sided filtering challenges (and other revisions). Thus, the results reported in the supplementary online appendix likely understate the gains that could be achieved in (quasi-)real-time forecasting from our simple shadow-rate VAR when compared to a plug-in approach.

Our supplementary online appendix reports various robustness checks, with fairly similar results to those reported here. In particular, when extending the evaluation period to end in August 2022, it turns out that while the economic effects of the pandemic left a heavy mark on readings of macroeconomic and financial variables in 2020 and 2021 — see, for example, our companion work in Carriero, et al. (2022b) — they did not materially affect the relative comparisons reported here. Moreover, we consider the effects of assuming either a lower (12.5 basis points) or higher (50 basis points) ELB value. Overall, contours of the estimated shadow rates, and relative forecast performance of the simple shadow-rate VAR are broadly similar to our baseline results, indicating robustness to the alternative choice in ELB value. The hybrid shadow-rate VAR continues to forecast well relative to the linear case, when the ELB value is lowered to by 12.5 basis points. Not surprisingly, however, the choice of a higher value for the censoring constraint on actual rates affects forecast simulations from the hybrid VAR more noticeably, since forecasts of macroeconomic variables

are now conditioned on actual rates that are assumed to not fall below 50 basis points (instead of 25 basis points). For good parts of our evaluation window, this choice conditions macroeconomic forecasts on policy rate projections that are about 25 basis points higher than they otherwise would have been, with negative impacts for longer-horizon forecasts.

### 6.2 Expected policy rates and forecasts of macroeconomic variables

The key differences between the VAR specifications considered here (standard/simple shadow-rate/hybrid shadow-rate) is the treatment of nominal interest rates near the ELB and the resulting forecasts for other, mainly macroeconomic, variables.[40] In each VAR, interest rate forecasts are not only an object of interest by themselves, but also serve as conditioning variables in the dynamic equations of other variables.[41] A discussion of forecasts for the actual interest rate in their own right is relegated to Section 6.3. Instead, here we highlight the differences in the predicted paths of the short-term policy rate used by each VAR to construct forecasts for macroeconomic variables; for brevity we also refer to these as "forecast-relevant" policy rates. Specifically, in the simple shadow-rate VAR, the forecast-relevant policy path used to generate macro forecasts is based on uncensored predictions of future shadow rates. In contrast, the hybrid shadow-rate model uses actual interest rates, which cannot fall below the ELB, as right-hand side variables for its VAR equations of $x_t$ in (5).[42] Critically, the standard linear VAR does not distinguish between actual and shadow rates. Ignoring the ELB in forming predictions of future interest rates, the linear VAR uses its uncensored interest rate forecasts in forming predictions of other variables.

To illustrate salient differences in the forecast-relevant policy rates that underlay predictions of macroeconomic (and other non-interest rate) variables, Figure 4 compares 12-month ahead averages of the forecast-relevant policy rates expected by each model out of sample. The left panel of the figure shows policy predictions generated from our full 20-variable system including yields, and the right panel shows corresponding predictions obtained when yields other than the federal funds rate are omitted from the VARs.

Using either data set, the three VARs differ notably in their forecasts of policy rates that enter as inputs in forecasts of macroeconomic conditions. By construction, the relevant policy rates of the hybrid shadow-rate VAR are constrained not to fall below the ELB, and turn out to hover closely above it for most of the post-GFC period. In contrast, the simple shadow-rate VAR jumps off the (largely) negative shadow-rate posteriors discussed above in Section 5 and predicts forecast-relevant policy rates to also remain fairly negative over the 12-months-ahead (time-averaged) horizon. The standard linear VAR borrows some negative rate outlook from the shadow-rate VAR early in the GFC, but with its forecasts jumping off actual-rate data remaining at the ELB, the standard model sees forecast-relevant policy rates hovering near the ELB for most of the

[40] Formally, using notation introduced in Section 3, by "other" variables we mean data contained in the partition $x_t$ of the VAR vector in (3). For brevity, we refer to these henceforth as "macroeconomic" variables, while noting that our application also includes some financial variables, like stock returns, but not term structure data, within $x_t$.

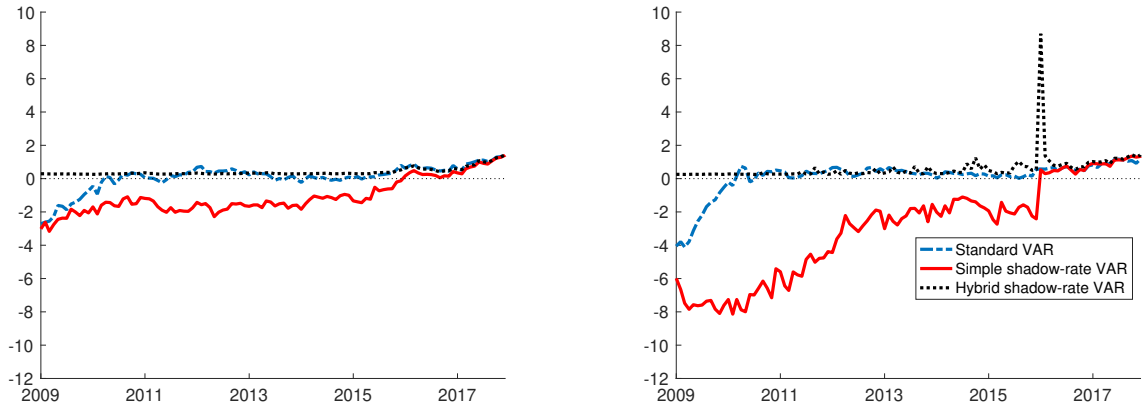
[41] As throughout this paper, we are concerned with dynamic simulations of so-called unconditional forecasts, $E(y_{t+h}|y^t)$, where $y^t$ denotes the history of available data at the forecast origin $t$, as opposed to conditional forecasts that take a fixed path for, say, future policy rates as given, such as $E(y_{t+h}|y^t, i^{t+h})$, as discussed by, among others, Antolín-Díaz, Petrella, and Rubio-Ramírez (2021), Jarocinski (2010), and Waggoner and Zha (1999). Instead, our present discussion is concerned with the (averages of) policy path expectations that are simulated by each model alongside its (unconditional) macroeconomic forecasts.

[42] The actual-rate values used in dynamic forecast simulations of the hybrid shadow-rate VAR reflect draws from the model's censored density for shadow-rate predictions; for an illustration see Section 6.3.

Figure 4: Expected averages of forecast-relevant policy rates (12-months ahead)

(a) Generated from data with yields

(b) Generated from data excluding yields




Note: 12-month ahead expected averages of federal funds rates (actual or shadow), generated out-of-sample at forecast origins 2009:01 – 2017:12. For each VAR model, the interest rate measures shown correspond to the forecast-relevant policy rate variables used in the model equations for macroeconomic variables, $x_t$: For the simple shadow-rate VAR, shadow-rate expectations are shown. For the hybrid shadow-rate VAR, predictions of actual rates are reported, and for the standard VAR, which ignores the ELB, (uncensored) federal funds rate predictions are shown.

post-2010 recovery.[43] Taken at face value in this reduced-form setting, the more negative rate forecasts of the simple shadow-rate VAR imply the most monetary policy stimulus, the positive forecasts of the hybrid shadow-rate VAR imply the least, and the standard linear VAR's interest rate projections fall in between.

While the inclusion of term structure data does not change the qualitative patterns in the forecast-relevant policy rate forecasts generated by the different models, the right panel of Figure 4 displays important quantitative differences. When the federal funds rate is the only interest rate included, both the linear and the simple shadow-rate VARs expect short rates to run almost twice as low during the ELB period as in the baseline models with additional interest rates included (this pattern is consistent with the historical estimates of shadow rates presented earlier). In the smaller specifications with the FFR as the only interest rate, this pattern in interest rate forecasts leads the simple shadow-rate VAR to over-predict the strength of the recovery. Consequently, as reported in the supplementary online appendix, a simple shadow-rate VAR without term structure data performs worse in average forecasting than a standard linear VAR.

In contrast, the hybrid shadow-rate VAR specification relates predictions of macro variables to lags of predicted actual interest rates (not shadow rates), which obey the ELB by construction. As shown in Figure 4, for most of the post-GFC recovery, the year-ahead averages of the policy rates relevant for forecasting macro variables were expected to stay very close to the ELB, and mostly quite a bit higher than in the simple shadow-rate VAR. As a result, even when term structure data is omitted in estimation, macro forecasts from the hybrid model are on par with a linear VAR (and thus more accurate than those from the simple

[43] Corresponding figures for the 6-month and 1-year yields, reported in the supplementary online appendix, are very similar.

shadow-rate VAR); details are provided in the supplementary online appendix. One possible interpretation of this pattern is that a shadow-rate VAR without term structure information may overstate the amount of monetary accommodation (either in quantity or efficacy) that was actually provided, at least during the ELB episode prompted by the GFC, which in turn harms the accuracy of forecasts made around ELB periods.

Our preferred models do include information about the term structure of interest rates, and add rates or yields from maturities of 6 months through 10 years. Although we omit direct comparisons in the interest of brevity, macro forecasts from these larger models are more accurate than those from the corresponding specifications in which the federal funds rate is the only interest rate. When term structure data is included, forecasts from (simple and hybrid) shadow-rate VARs improve in an absolute sense as well as relative to the linear VAR. In addition, shadow-rate estimates and FFR forecasts informed by term structure data become less negative around the 2009-2015 ELB episode and during the first year or so of the COVID-19 pandemic. One interpretation is that, by adding yields to the VAR vector, the shadow-rate specifications are less prone to over-estimating stimulus from unconventional monetary policy (at least during the initial ELB episode prompted by the GFC).

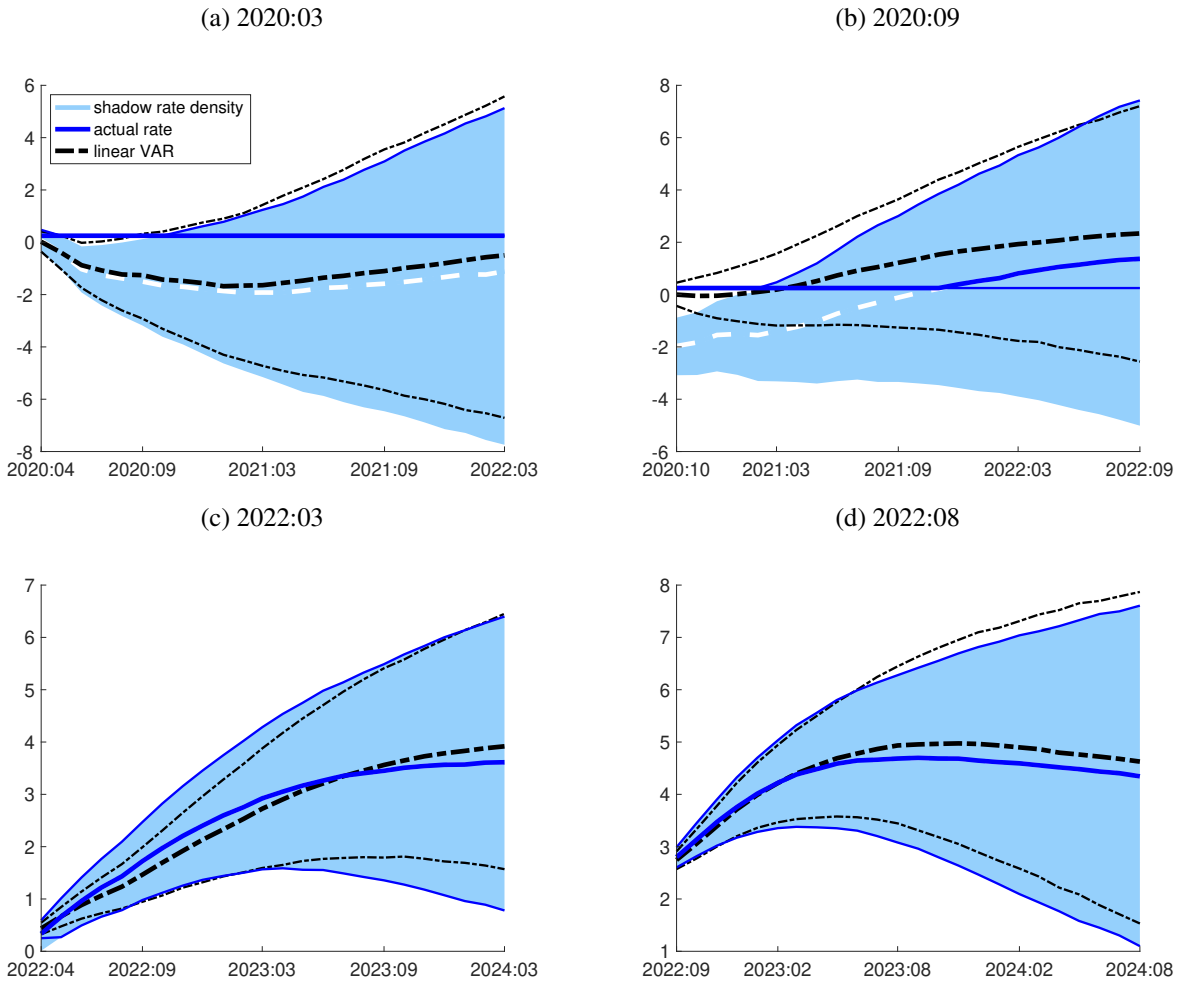
Overall, these comparisons suggest two broad takeaways about capturing the effects of monetary policy in reduced-form VARs for forecasting. First, for forecasting macro variables, it is helpful to include a range of interest rates and not just the federal funds rate. In particular, as shown in the supplementary online appendix, our shadow-rate VARs (using all yields) considerably improve upon forecasts for macro variables obtained from a linear VAR that simply omits short-term interest rates to avoid ELB issues. Second, a shadow-rate specification provides an internally consistent forecasting model (whereas the standard does not). Moreover, the shadow-rate VARs ensure interest rates obey the ELB, which improves the accuracy of their interest rate forecasts. In addition, shadow-rate VARs that distinguish between effects from actual and shadow rates as predictors for economic outcomes, improve the accuracy of some macro forecasts, while remaining otherwise comparable to a linear VAR.

### 6.3 Interest rate forecasts made since the outbreak of COVID-19

The period following the outbreak of the COVID-19 pandemic in the US and the aggressive easing of monetary policy by the FOMC provides an opportunity for a case study of predicted interest rate dynamics from our shadow-rate VARs as compared to a standard VAR that ignores the ELB. In these comparisons, we focus here on the simple shadow-rate VAR as compared to a standard VAR; the supplementary online appendix provides similar results for a comparison between the hybrid shadow-rate VAR and the standard VAR, as well as estimates generated from data that excludes all interest rate measures except for the federal funds rate. Figure 5 shows the evolution of federal funds rate forecasts over selected origins between March 2020 and the end of our sample in August 2022.[44] With data available through March 2020, as the outbreak of COVID-19 hit the US economy, the point forecast from a standard VAR put the funds rate well below the ELB for the entire forecast horizon, with substantial probability mass on very negative rates. As of September 2020, the point forecast was close to the ELB for several months, but throughout the forecast

[44]Forecasts generated at a given forecast origin, say March 2020, reflect model estimates based on data up to and including the month of the forecast origin.

Figure 5: Predictive densities for actual and shadow values of the federal funds rate




Note: Predictive density for the actual and shadow values of the federal funds rate, simulated out of sample at different jump-off dates. Dashed-dotted (black) lines depict the predictive density for the actual rate as generated from the standard VAR. The shaded (light blue) area with dashed (white) lines represent the simple shadow-rate VAR's predictive density of the shadow rate, while solid lines (dark blue) reflect the corresponding censored density for the actual interest rate. Posterior medians and 68 percent bands.

horizon substantial mass in the predictive distribution remained in negative territory.

At the March 2020 forecast origin, past data for the US economy was still well above the ELB, and the predictive density for the shadow rate generated from the simple shadow-rate VAR was essentially identical to the (uncensored) actual rate distribution obtained from the standard VAR, as shown in Panel (a) of the figure. However, things changed as the economy stayed at the ELB in subsequent months. As previously described in Section 5, the rapid deterioration in economic conditions led to a decline in the (median) shadow rate which stabilized around $-2$ percent by the second half of 2020. As shown in Panel (b), as of September 2020, negative levels of the shadow rate at the forecast origin pulled down the predictive densities for the

shadow rate.[45] As a result, until the first half of 2021 (detailed results omitted in the interest of brevity), the shadow rate was expected to cross above the ELB quite a bit later than implied by uncensored federal funds rate predictions generated from the standard VAR.

Jumping forward to March 2022, when the FOMC raised the funds rate off the ELB, the predictive densities of the funds rate from both the standard and shadow-rate VARs were solidly above the ELB throughout the forecast horizon. The median projections were similar, with the funds rate rising steadily to about 3.5 percent over about 18 months, but with the projection from the shadow-rate VAR slightly higher than the standard VAR's forecast over most of the period. At the end of our sample, in August 2022, the median forecasts from the two approaches (standard and simple shadow-rate VAR) were also similar, with the standard VAR's forecast slightly higher; in both forecasts, the funds rate rises from about 3 percent to nearly 5 percent before gradually declining.

Forecasts of the federal funds rate from the simple shadow-rate VARs reflect the predicted evolution of the shadow rate. As noted before, shadow rates reflect the unconstrained policy rate prescriptions of the feedback rule for monetary policy that is implied by the VAR in (4). From an economic perspective, shadow-rate VARs can capture lower-for-longer or make-up elements of the Federal Reserve's monetary policy strategy through the dependence of predicted interest rates on lagged notional rates as suggested by, for example, the models of Billi (2020), Gust, et al. (2017), and Reifschneider and Williams (2000). Moreover, the shadow-rate estimates are informed by observed data on longer-term yields and economic conditions, which enables the estimates to pick up on the effects of unconventional policies, such as forward guidance and asset purchases, through these channels.

As a reference point, the Federal Reserve Bank of Cleveland regularly publishes a set of policy path prescriptions obtained from monetary policy rules in the tradition of Taylor (1993). Prescriptions are derived from seven different rules and for three alternative sets of forecasts for economic conditions; similar to our shadow-rate concept, all prescriptions ignore the ELB.[46] Some of these rules are so-called "inertial" rules that generate prescriptions with strong lagged dependence on past policy rates. At the ELB, inertia with respect to lagged interest rate values as obtained from unconstrained prescriptions (or so-called notional rates) can capture lower-for-longer policies as discussed by, for example, Billi (2020), and resemble the form of the feedback rule embedded in our shadow-rate VARs.

Our shadow-rate estimates for the COVID-19 period were negative from April 2020 through February 2022 (the estimates were around -2 percent from May 2020 through May 2021), which broadly aligns with unconstrained prescriptions of common policy rules. For example, with data available as of December 1, 2020, the median rule prescription calculated by the Federal Reserve Bank of Cleveland put the federal funds rate at about $-50$ basis points in 2020:Q4 and $-70$ basis points for the first three quarters of 2021. These simple rule predictions were jumping off the actual funds rate level in late 2020 (not the negative level of a notional rate). As a result, the simple prescriptions were a little less negative than those derived from

[45] The shadow-rate VARs' predictive densities integrated over the entire posterior distribution of shadow-rate values, as depicted, for example, in Figure 2, instead of jumping off any specific estimate for current and past values of the shadow rate.

[46] The rule results and documentation are available at `https://www.clevelandfed.org/en/our-research/indicators-and-data/simple-monetary-policy-rules.aspx`. Prescriptions from a similar set of policy rules are also computed by staff at the Board of Governors and presented to the FOMC ahead of each of its meetings as part of Tealbook Book B; see also Board of Governors of the Federal Reserve System (2023).

our simple shadow-rate VAR at the end of 2020:Q3 as shown in Panel (b) of the figure, but both predicted a similar timing of the first rate hike above the ELB by spring 2022. Reflecting stronger economic conditions towards the end of our sample, the median rule prescription calculated by the Federal Reserve Bank of Cleveland by June 1, 2021, put the federal funds rate for 2021:Q2 a bit above the ELB, at 72 basis points, while (as of June 2021) our simple shadow-rate VAR did not predict rates to rise above the ELB until early 2022, which is more in line with prescriptions from the lower end of the range of rules prescription reported by the Cleveland Fed in June 2021. However, the shallow federal funds rate path predicted by the simple shadow-rate VAR was much closer to survey expectations obtained from Blue Chip Financial Forecasts and the Survey of Professional Forecasters (SPF).[47]

# 7 Structural Analysis

This section illustrates the use of our models in structural analysis, estimating the response of the economy to a shock to financial conditions with our hybrid shadow-rate VAR. We focus here on the hybrid shadow-rate VAR, since in the simple shadow-rate VAR, the responses of macroeconomic variables to shocks are invariant to whether the ELB binds or not. The model takes the form described above, adding the excess bond premium of Gilchrist and Zakrajšek (2012) to our baseline variable set. For simplicity, we use a recursive ordering for identification with the EBP ordered first. As detailed in Section 3.6, we rely on the generalized impulse response function approach of Koop, Pesaran, and Potter (1996) to estimate the responses, with full-sample estimates of the model. The size of the shock reflects an average estimate of the standard deviation in EBP shocks during 2005-2006, the two years preceding the Great Recession.

To illustrate the influence of binding ELB constraints, Figure 6 reports estimated GIRFs for January 2007 (blue lines) and December 2012 (red lines).[48] In the former case, the economy was still expanding, and interest rates were well above the ELB (e.g., the federal funds rate was 5.25 percent), whereas in the latter, the ELB was binding for short-term rates (with readings for the federal funds, 6-month, and 1-year Treasury rates at 16, 12, and 16 basis points, respectively). For chart readability, we report estimates of a key subset of the model's variables. These GIRFs do not assume the economy to either remain at or away from the ELB over the response horizon, but endogenously account for changes of interest rates towards or away from the ELB as the simulations move forward in time.

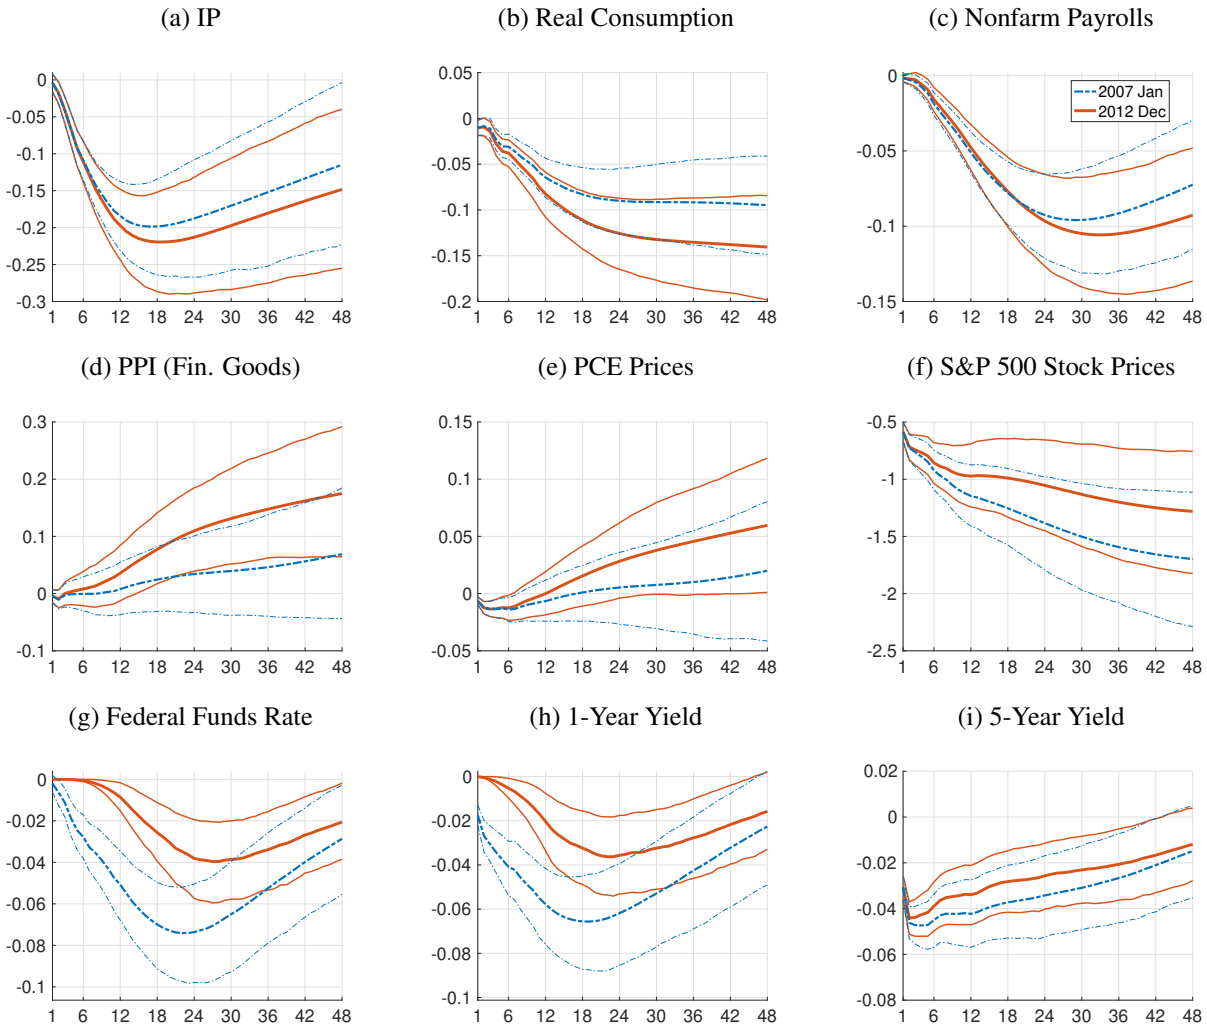
Our estimates of the shock's effects in 2007 resemble the estimates of Gilchrist and Zakrajšek (2012) based on quarterly data. The EBP shock depresses economic activity (in our figure, industrial production, consumption, and employment) and stock prices. In our estimates, the shock boosts measures of the aggregate price level, whereas it lowers the price level in Gilchrist and Zakrajšek (2012); but in both sets of estimates, the impacts of the shock on price levels are sufficiently imprecise as to not differ significantly from 0. The shock leads to a reduction of the federal funds rate, as well as declines in other yields.

When the same-sized shock hits financial conditions after the end of 2012, short-term rates are at the ELB, while policy can still affect longer-term yields (such as the 5- and 10-year yields in our model). With

[47] For example, the 2021:Q2 SPF did not see any significant rise in the 3-month Treasury rate before 2023.

[48] In keeping with the timing assumptions described in Section 3.6, the GIRF origins represent the jump-off point $t$ of baseline forecasts against which an alternative forecast is compared that conditions on a specific shock occurring at $t + 1$.

Figure 6: Generalized Impulse Responses to an EBP Shock in 2007 and 2012




Note: Estimates obtained from full-sample posterior hybrid shadow-rate VAR at two different GIRF origins with 48-month horizons. Size of EBP shock corresponds to average SV level of 0.11 for the years 2005-2006. Posterior means and 68 percent bands. Interest rates in levels, and cumulated log levels (times 100) for all other variables.

short-term rates already at the ELB just prior to impact, the shock leads to a much more muted decline in their forecasts than in the case of the 2007 shock. For example, the ELB constraint prevents federal funds rate predictions from falling for more than six months after the shock occurred. To be clear, these responses (as in standard impulse responses) are declines in average predictions for future funds rates relative to a baseline forecast (that is often rising), and not expected declines relative to a current level of the funds rate that is ostensibly at its effective lower bound. With the funds rate starting at the ELB in December 2012, the model expects the shadow rate to gradually rise from negative to positive territory, eventually returning to its historical mean, in turn implying a baseline path of the federal funds rate that after some months lifts off the ELB and subsequently gradually increases.With the shock, the funds rate lifts off later

and follows a lower trajectory. Accordingly, reflecting these baseline and shocked paths of the federal funds rate, Figure 6's GIRF estimate shows the federal funds rate eventually declining, albeit much less than it would in the absence of the ELB constraint. As might be expected given the term structure, the differences between the 2007 and 2012 responses of the 5-year bond yield are smaller than the differences in responses of the federal funds rate. With interest rates moving less in the 2012 GIRFs as compared to 2007, some measures of economic activity (e.g., industrial production) show sharper declines in response to the 2012 shock, albeit with overlap in the point-wise uncertainty bands from each simulation. Measures of aggregate prices faced by firms and households show stronger increases following the 2012 shock as compared to the 2007 shock. This pattern in price level responses could be seen as consistent with the model-based predictions in Gilchrist, et al. (2017) in which an adverse demand shock causes inflation to rise rather than fall in a model in which firms facing financial frictions raise their prices in response to adverse shocks to demand or financial conditions, driven by a desire to preserve internal liquidity and avoid external financing.

In the interest of brevity, we relegated to the supplementary online appendix another comparison, of GIRFs obtained from a linear VAR that ignores the ELB against the responses from our hybrid model as shown here. These estimates differ from each other in ways that are broadly similar to those shown in Figure 6's hybrid-model GIRF with origins in 2007 and 2012. For example, at times when the ELB binds, our model expects a more severe decline in economic activity than a linear VAR, which ignores the ELB. This comparison confirms the impact on structural inference from our approach to modeling the ELB through the shadow-rate VAR as opposed to a purely linear model.

Together, the GIRF estimates presented in this section illustrate the use of our (hybrid) shadow-rate VAR for structural analysis. Our approach to modeling ELB constraints on interest rates shows impacts of incorporating the ELB in the VAR through the shadow rate as opposed to a purely linear VAR. The results presented here are also confirmed by additional GIRFs, obtained for origins at other ELB dates, shocks of other sizes and signs, as well as replacing shocks to EBP by those to VXO, and a smaller-sized VAR model, as reported in the supplementary online appendix. Among others, we studied GIRFs obtained from a small-size VAR model that only contains EBP, unemployment, inflation and the federal funds rate. The smaller model generates responses that are qualitatively similar to those for our medium-size model, in the sense that at the ELB the interest rate decreases less so that unemployment increases more and inflation decreases less. Yet, the differences emerge more slowly than in the larger model, which provides a more granular representation of the economy that better captures the shock transmission mechanism. These results underscore the importance of working with such larger shadow-rate models, and use of our shadow-rate sampling methods assures their computational feasibility.

# 8 Conclusion

Motivated by the prevalence of lower bound constraints on nominal interest rates, this paper develops a tractable approach to including a shadow-rate specification in medium-scale VARs commonly used in macroeconomic forecasting. Our models treat interest rates as censored observations of a latent shadow-rate process in a VAR setup. As in a classic Tobit model, the shadow rate is assumed to run below the ELB when

the actual interest rate is at the ELB, and equal to the observed interest rate when the ELB is not binding. Our shadow-rate approaches extend the specific unobserved components model of Johannsen and Mertens (2021) to the general VAR setting and build on the data augmentation methods developed by Chib (1992) and Chib and Greenberg (1998) for Tobit and Probit models. By using a computationally more efficient shadow-rate sampling algorithm, together with the recursive methods of Carriero, Clark, and Marcellino (2019) and Carriero, et al. (2022a) for efficient estimation of Bayesian VARs with stochastic volatility, our approaches are easily applied to a medium-scale VAR system, and they successfully handle data in which the ELB is binding for multiple interest rates of different maturities.

We consider two specifications of the shadow-rate VAR. The "simple shadow-rate VAR" assumes that only shadow rates, not actual rates, matter as explanatory variables in linear forecasting equations. In contrast, the "hybrid shadow-rate VAR" combines actual and shadow interest rates in its state dynamics. Since actual and shadow rates are perfectly collinear when above the ELB, their joint inclusion in a VAR can create challenges in identification and estimation, in particular for the purpose of generating reliable out-of-sample forecasts based on the (yet) still limited experience with ELB data. Hence, our hybrid shadow-rate VAR includes shadow rates as VAR regressors only in forecasting equations for other (shadow) term structure variables, while using actual rates (and not shadow rates) as explanatory variables in the VAR equations of macroeconomic variables and other measures of financial conditions. As in the simple shadow-rate VAR, the hybrid model uses shadow rates as dependent variables to describe the evolution of (shadow) term structure variables. Through direct and indirect effects, both actual and shadow rates influence multi-step forecasts of all variables in the hybrid VAR system.

We use our shadow-rate models to form forecasts and responses to structural shocks from medium-scale BVARs with stochastic volatility. In our results, forecasts for interest rates obtained from a simple shadow-rate VAR for the US since 2009 are clearly superior, in terms of both point and density forecasts, to predictions from a standard VAR that ignores the ELB. These interest rates include not only the federal funds rate but also longer-term bond yields. For some macroeconomic variables, the shadow-rate VAR generates gains in predictive performance, in particular when actual (instead of shadow) rates serve as explanatory variables. However, the models are quite similar in accuracy for forecasts for many macroeconomic variables, as the standard VAR can generate a negative-rate outlook (for actual rates) that influences its economic forecasts in ways that are similar to a shadow-rate VAR. For forecasting financial and macro variables, it is, of course, helpful to include in the models a range of yields from the term structure of interest rates and not just the federal funds rate. Critically, when conditioned on interest rates of shorter- and longer-term maturities, our shadow-rate VARs considerably improve upon forecasts for macroeconomic variables obtained from a linear VAR that simply omits short-term interest rates to avoid ELB issues. Our forecasts are also superior to those from a VAR that replaces the federal funds rate with an external shadow-rate estimate.

Finally, to illustrate the use of our shadow-rate VARs for structural analysis, we estimate the response of the economy to a shock to financial conditions and find that accounting for the ELB matters for estimated responses, as the latter differ both when at or away from the ELB and with respect to the same responses from a linear VAR.

# Appendix

## A Shadow-Rate Sampling

This appendix provides further details on the application of a shadow-rate sampler in a VAR context that can be embedded in an otherwise standard MCMC estimation. Throughout, we take as given values of all parameters (including SV) of the shadow-rate VAR in (4). These parameter values can be obtained from standard MCMC steps (and based on previously sampled "data" for $\{z_t\}_{t=1}^T$) as described in, for example, Carriero, Clark, and Marcellino (2019)). Here we focus on the MCMC step concerned with sampling from the shadow-rate problem stated in (13) for given values of the VAR parameters $\{C_j\}_{j=1}^p$ and $\{\Sigma_t\}_{t=1}^T$. For ease of notation, references to $\{C_j\}_{j=1}^p$ and $\{\Sigma_t\}_{t=1}^T$ will be suppressed from the conditioning sets described below. Moreover, the value of $ELB$ is taken as a known constant.

### A.1 Gibbs sampling from the truncated multivariate normal

While the joint density from a multivariate normal (MVN) distribution can be factorized into a product of univariate normal densities, the same property does not generally extend to the truncated MVN density, as discussed by, for example, Geweke (1991). In our context, this means that draws from the missing-value problem, $\boldsymbol{S} \,|\, \bar{\boldsymbol{Y}} \sim N(\boldsymbol{\mu}, \boldsymbol{\Omega})$ in (12), could be recursively obtained by a sequence of univariate normal draws, but not so for the corresponding shadow-rate problem $\boldsymbol{S} \,|\, \boldsymbol{Y} \sim TN(\boldsymbol{\mu}, \boldsymbol{\Omega}, -\infty, ELB)$. However, consider a single element of $\boldsymbol{S}$, denoted $s_t$, and let $s_{1:t-1}$ and $s_{t+1:T}$ denote the vectors of all elements of $\boldsymbol{S}$ that precede and follow $s_t$, respectively.[49] Conditional on $s_{1:t-1}$ and $s_{t+1:T}$ (as well as $\boldsymbol{Y}$), $s_t$ has a univariate truncated normal distribution, with moment parameters $\mu_{1:t-1,t+1:T}$ and $\omega_{1:t-1,t+1:T}$ identical to those obtained from the corresponding missing-value problem:[50]

$$s_t \,\big|\, s_{1:t-1}, s_{t+1:T}, \boldsymbol{Y} \sim TN(\mu_{1:t-1,t+1:T}, \omega_{1:t-1,t+1:T}, -\infty, ELB) \,, \tag{A.1}$$

$$s_t \,\big|\, s_{1:t-1}, s_{t+1:T}, \bar{\boldsymbol{Y}} \sim N(\mu_{1:t-1,t+1:T}, \omega_{1:t-1,t+1:T}) \,. \tag{A.2}$$

As discussed by Geweke (1991), Gelfand, Smith, and Lee (1992), and references therein, the fact that conditional distributions of the truncated MVN are also truncated normals means that the problem of obtaining draws from the truncated MVN can be addressed with a Gibbs sampler, which we also pursue. For now, we abstract from estimating the VAR, and consider solely the shadow-rate sampling problem for given VAR parameters. Adopting language from Geweke (1991), a Gibbs sampler generates a draw from (13) by performing $N$ passes over (A.1), at each pass iterating over $t = 1, 2, \ldots T$.[51] Specifically, let $\boldsymbol{S}^{(n)}$ denote a

[49] In the case of $N_s > 1$, so that $s_t$ is not scalar, the argument made here applies to a single scalar element of $s_t$ conditional on values for the remainder of the shadow-rate vector, as well as $s_{1:t-1}$, $s_{t+1:T}$, and $\boldsymbol{Y}$.

[50] For recent treatments, see, for example, Horrace (2005) and Chopin (2011). The argument also extends to the case in which the truncation bounds vary from one element of the vector to another, which also allows us to handle where the sequence $s_{1:T}$ covers observations where the ELB does not bind, so that $s_t = i_t > ELB$. To handle those cases, the shadow-rate sampling problem can be stated more generally as $\boldsymbol{S} \,|\, \boldsymbol{Y} \sim TN(\boldsymbol{\mu}, \boldsymbol{\Omega}, -\infty, \boldsymbol{I})$, where $\boldsymbol{I}$ is the vector of all actual rates, and the inequality $\boldsymbol{S} \leq \boldsymbol{I}$ applies element-wise. Note that observations of $i_t$ for which the ELB does not bind are included in $\bar{\boldsymbol{Y}}$ so that the distribution of the missing-value problem in (12) collapses on a point mass, $s_t = i_t$, for those observations.

[51] When the sequence $s_{1:T}$ includes observations $t$ where the ELB does not bind, the posterior in (A.2) collapses on a point

draw of $\boldsymbol{S}$ from the $n$th pass of the Gibbs sampler, with typical element $s_t^{(n)}$. At each pass $n = 1, 2, \ldots, N$, the Gibbs sampler iterates from $t = 1, 2 \ldots, T$ and draws

$$s_t^{(n)} \mid s_{1:t-1}^{(n)}, s_{t+1:T}^{(n-1)}, \boldsymbol{Y} \,. \tag{A.3}$$

We keep $\boldsymbol{S}^{(N)} = \{s_t^{(N)}\}_{t=1}^T$ as the draw from (13). Appendix B describes in further detail how this Gibbs sampler for the shadow-rate problem is embedded in our MCMC sampler for the shadow-rate VAR.

Our implementation of the Gibbs sampling steps for the shadow-rate problem in (A.1) exploits the particular structure of the VAR($p$) setting to derive the conditional moments $\mu_{1:t-1,t+1:T}$, $\omega_{1:t-1,t+1:T}$ without ever having to compute the entire mean vector, $\boldsymbol{\mu}$, and variance-covariance matrix, $\boldsymbol{\Omega}$, of the full shadow-rate problem in (13), which can be substantial in size because $\boldsymbol{S}$ is a vector of size $T \cdot N_s$.[52] For example, in our application to monthly US data, and considering only the ELB episode witnessed from 2008:12 through 2015:12, we have $T = 85$ and $N_s = 3$ (as the ELB has been binding for the federal funds rate, 6-month T-bill rate, and 1-year Treasury yield in our data set over this period).

Drawing from (A.1) requires a draw from the truncated normal distribution. With $N_s = 3$ in our application, (A.1) represents a draw from the multivariate truncated normal, which can be broken down further into a sequence of Gibbs sampling steps consisting of draws of univariate truncated normals.[53] We implement this draw by application of uniform-inverse-transform sampling as follows. Dropping time subscripts, consider the problem of drawing the scalar $s \sim TN(\mu, \omega^2, -\infty, ELB)$, which is equivalent to drawing $v \sim TN(0, 1, -\infty, \bar{v})$ with $\bar{v} = (ELB - \mu)/\sigma$ and $s = \mu + \sigma\, v$. The support of $v$ is $v \leq \bar{v}$, over which its probability density function (*pdf*) and cumulative distribution function (*cdf*), $f_v(v)$ and $F_v(v)$, are:

$$f_v(v) = \frac{\phi(v)}{\Phi(\bar{v})}\,, \qquad\qquad F_v(v) = \frac{\Phi(v)}{\Phi(\bar{v})}, \tag{A.4}$$

where $\phi(\cdot)$ and $\Phi(\cdot)$ are the standard normal *pdf* and *cdf*, respectively. We construct $v = F_v^{-1}(u) = \Phi^{-1}\left(u \cdot \Phi(\bar{v})\right)$, using a draw $u \sim U(0, 1)$ of a uniform random variable with support $[0, 1]$. Alternatively, rejection sampling could be used, or a combination of both approaches; see, for example, Geweke (1991), Chopin (2011), and Botev (2017).[54]

mass, $s_t = i_t > ELB$, and those cases can be skipped. (Formally, the support of the truncated normal in (A.1) could be viewed as being bounded from above by $i_t$ for all observations $t$.)

[52] By exploiting sparsities and the recursive structure of the VAR's state space representation, our approach echoes recent advances in the field of sampling from the truncated MVN distribution made by Cong, Chen, and Zhou (2017), albeit specialized to the VAR($p$) that we intend to investigate further. Other advances in Gibbs sampling from the truncated MVN distribution are discussed by Botev (2017), Park, Genton, and Ghosh (2007), Chopin (2011), Damien and Walker (2001), and Robert (1995).

[53] In case of $N_s > 1$, the missing value problem gives rise to an $N_s$ dimensional multivariate normal with $N_s \times 1$ mean vector, $\mu_{1:t-1,t+1:T}$, and $N_s \times N_s$ dimensional variance-covariance-matrix $\omega_{1:t-1,t+1:T}$, based on which we can derive the conditional moments of a given scalar element of $s_t$ for given values of its remaining elements, as described, among others, by Geweke (1991).

[54] In our application, potential gains from applying a combination approach have been limited. Direct application of the MATLAB routine `trandn.m` from Botev (2017) underperformed relative to uniform-inverse-transform sampling. A likely cause appears to be our use of large pre-generated random arrays as opposed to generating pseudo-random values one-at-a-time.

## A.2 Shadow-rate VAR in companion form

To derive $\mu_{1:t-1,t+1:T}$, $\omega_{1:t-1,t+1:T}$, we now focus on the moments of the (unconstrained) missing-value problem stated in (A.2), understanding that draws for $s_t$ are to be generated from the constrained shadow-rate problem in (A.1). Written in companion form, the VAR($p$) in (4) is characterized by Markov dynamics of a state vector that tracks $z_t$ and $p-1$ of its lags. Consequently, it is sufficient to consider no more than $p$ lags and leads of $z_t$ in the derivation of $\mu_{1:t-1,t+1:T}$, $\omega_{1:t-1,t+1:T}$. We employ the following companion-form notation for the VAR (omitting intercepts), adapted to the partitioning of $z_t$ into $x_t$ and $s_t$:

$$Z_t = \boldsymbol{C} Z_{t-1} + \boldsymbol{\Psi}_t w_t \,, \qquad\qquad w_t \sim N(0, I) \,, \tag{A.5}$$

and let $D_x$ and $D_s$ be selection matrices so that $x_t = D_x Z_t$ and $s_t = D_s Z_t$.

To construct this companion form, consider for concreteness the case of a second-order system ($p = 2$):

$$X_t = \begin{bmatrix} x_t \\ x_{t-1} \end{bmatrix}, \qquad S_t = \begin{bmatrix} s_t \\ s_{t-1} \end{bmatrix}, \qquad Z_t = \begin{bmatrix} X_t \\ S_t \end{bmatrix} = \begin{bmatrix} x_t \\ x_{t-1} \\ s_t \\ s_{t-1} \end{bmatrix}, \tag{A.6}$$

$$\text{with} \quad \boldsymbol{C} = \begin{bmatrix} C^1_{xx} & C^2_{xx} & C^1_{xs} & C^2_{xs} \\ I & 0 & 0 & 0 \\ C^1_{sx} & C^2_{sx} & C^1_{ss} & C^2_{ss} \\ 0 & 0 & I & 0 \end{bmatrix}, \qquad \boldsymbol{\Psi}_t = \begin{bmatrix} \Psi_{x,t} \\ 0 \\ \Psi_{s,t} \\ 0 \end{bmatrix}, \tag{A.7}$$

where $\Psi_{x,t}$ and $\Psi_{s,t}$ are conformable partitions of a factorization $\Sigma_t^{0.5}$ of the variance-covariance matrix of VAR residuals in (4), so that $\Sigma_t = \Sigma_t^{0.5} \left(\Sigma_t^{0.5}\right)'$ and $\Sigma_t^{0.5} = \begin{bmatrix} \Psi'_{x,t} & \Psi'_{s,t} \end{bmatrix}'$.

## A.3 Moments of the missing-value problem

Given the VAR($p$) structure of the model, and using the companion-form notation introduced above, the posterior density of the Gaussian missing-value problem in (A.2) simplifies as follows:

$$\begin{aligned} f\left(s_t \,\middle|\, s_{1:t-1}, s_{t+1:T}, \bar{\boldsymbol{Y}}\right) &= f\left(s_t \,\middle|\, s_{t-p:t-1}, s_{t+1:t+p}, x_{t-p:t+p}\right) \\ &= f\left(s_t \,\middle|\, Z_{t-1}, Z_{t+p}, x_t\right) \\ &= f\left(s_t \,\middle|\, Z_{t-1}, Z_{t+p} - Z_{t+p|t-1}, v_t^x\right), \end{aligned} \tag{A.8}$$

where $Z_{t+p|t-1} = E(Z_{t+p}|Z_{t-1}) = \boldsymbol{C}^{p+1} Z_{t-1}$ and $v_t^x = x_t - E(x_t|Z_{t-1}) = D_x v_t$.

**Observations $t$ with $t+p \leq T$:** As stated above, we assume that we have observations for at least $p$ initial lags of $s_t$ at $t = 1$, so that we can condition on $s_{t-p:t-1}$. For $t + p \leq T$, we define a signal vector, $\boldsymbol{\mathcal{Z}}_{t+p}$:

$$\boldsymbol{\mathcal{Z}}_{t+p} = \begin{bmatrix} Z_{t+p} - Z_{t+p|t-1} \\ v_t^x \end{bmatrix} = \begin{bmatrix} \sum_{j=0}^{p} \boldsymbol{C}^{p-j}\, \boldsymbol{\Psi}_{t+j}\, w_{t+j} \\ D_x\, \boldsymbol{\Psi}_t\, w_t \end{bmatrix}.$$

The moments $\mu_{1:t-1,t+1:T}$ and $\omega_{1:t-1,t+1:T}$ follow from Gaussian signal extraction:

$$\begin{aligned} \mu_{1:t-1,t+1:T} &= E(s_t|Z_{t-1}, Z_{t+1}, x_t) = E(s_t|Z_{t-1}) + \boldsymbol{J}_{t-1}\boldsymbol{\mathcal{Z}}_{t+p}, \\ \text{with} \quad \boldsymbol{J}_{t-1} &= \operatorname{Cov}\left(s_t, \boldsymbol{\mathcal{Z}}_{t+p} \,\middle|\, Z_{t-1}\right) \left(\operatorname{Var}\left(\boldsymbol{\mathcal{Z}}_{t+p} \,\middle|\, Z_{t-1}\right)\right)^{-1}, \qquad \text{(A.9)} \\ \operatorname{Var}\left(\boldsymbol{\mathcal{Z}}_{t+p} \,\middle|\, Z_{t-1}\right) &= \begin{bmatrix} \sum_{j=0}^{p} \boldsymbol{C}^{p-j}\, \boldsymbol{\Psi}_{t+j} \boldsymbol{\Psi}_{t+j}' \left(\boldsymbol{C}^{p-j}\right)' & \boldsymbol{C}^{p}\, \boldsymbol{\Psi}_t \boldsymbol{\Psi}_t' D_x' \\ D_x\, \boldsymbol{\Psi}_t \boldsymbol{\Psi}_t' \left(\boldsymbol{C}^{p}\right)' & D_x\, \boldsymbol{\Psi}_t \boldsymbol{\Psi}_t' D_x' \end{bmatrix}, \\ \operatorname{Cov}\left(s_t, \boldsymbol{\mathcal{Z}}_{t+p} \,\middle|\, Z_{t-1}\right) &= \begin{bmatrix} D_s\, \boldsymbol{\Psi}_t \boldsymbol{\Psi}_t' \left(\boldsymbol{C}^{p}\right)' \\ D_s\, \boldsymbol{\Psi}_t \boldsymbol{\Psi}_t' D_x' \end{bmatrix}, \end{aligned}$$

$$\begin{aligned} \text{and} \quad \omega_{1:t-1,t+1:T} &= \operatorname{Var}\left(s_t|Z_{t-1}, Z_{t+1}, x_t\right) = D_s \boldsymbol{\Psi}_t \boldsymbol{\Psi}_t' D_s' \\ &\quad - \operatorname{Cov}\left(s_t, \boldsymbol{\mathcal{Z}}_{t+p} \,\middle|\, Z_{t-1}\right) \left(\operatorname{Var}\left(\boldsymbol{\mathcal{Z}}_{t+p} \,\middle|\, Z_{t-1}\right)\right)^{-1} \operatorname{Cov}\left(s_t, \boldsymbol{\mathcal{Z}}_{t+p} \,\middle|\, Z_{t-1}\right). \qquad \text{(A.10)} \end{aligned}$$

To compute the Kalman-smoothing gain $\boldsymbol{J}_{t-1}$ and residual variance $\operatorname{Var}\left(s_t|Z_{t-1}, Z_{t+1}, x_t\right)$ efficiently, and robustly to numerical round-off errors (which could otherwise imply non-positive-definite values for variance-covariance matrices), we employ a QR factorization that builds on some of the fast-array algorithms presented by Kailath, Sayed, and Hassibi (2000). Specifically, consider a factorization of the joint variance-covariance matrix of the signal, $\boldsymbol{\mathcal{Z}}_{t+p}$, and unknown state, $s_t$:[55]

$$\operatorname{Var}_{t-1}\left(\begin{bmatrix} \boldsymbol{\mathcal{Z}}_{t+p} \\ s_t \end{bmatrix}\right) = \boldsymbol{L}_{t-1}\, \boldsymbol{L}_{t-1}' = \boldsymbol{M}_{t-1}\, \boldsymbol{M}_{t-1}',$$

where $\boldsymbol{L}_{t-1}$ and $\boldsymbol{M}_{t-1}$ are the following square matrices of length $N_z \cdot p + N_s$ along each dimension:

$$\begin{aligned} \boldsymbol{L}_{t-1} &= \begin{bmatrix} \operatorname{Var}_{t-1}\left(\boldsymbol{\mathcal{Z}}_{t+p}\right)^{0.5} & \boldsymbol{0} \\ \boldsymbol{J}_{t-1} \left(\operatorname{Var}_{t-1}\left(\boldsymbol{\mathcal{Z}}_{t+p}\right)^{0.5}\right)' & \operatorname{Var}_{t-1}\left(s_t|\boldsymbol{\mathcal{Z}}_{t+p}\right)^{0.5} \end{bmatrix} \\ \boldsymbol{M}_{t-1} &= \begin{bmatrix} \boldsymbol{C}^{p}\boldsymbol{\Psi}_t & \boldsymbol{C}^{p-1}\boldsymbol{\Psi}_{t+1} & \dots & \boldsymbol{C}^{1}\boldsymbol{\Psi}_{t+p-1} & \boldsymbol{\Psi}_{t+p} \\ D_x\boldsymbol{\Psi}_t & 0 & \dots & \dots & 0 \\ D_s\boldsymbol{\Psi}_t & 0 & \dots & \dots & 0 \end{bmatrix}. \end{aligned}$$

The matrix $\boldsymbol{M}_{t-1}$ is straightforward to construct. Crucially, $\boldsymbol{L}_{t-1}$ is a lower triangular matrix, and its transpose, $\boldsymbol{L}_{t-1}'$, can directly be obtained from a QR factorization of $\boldsymbol{M}_{t-1}'$, $\boldsymbol{M}_{t-1}' = \boldsymbol{Q}_{t-1}\boldsymbol{L}_{t-1}'$ (where $\boldsymbol{Q}_{t-1}$ is an orthogonal matrix that can be ignored for our purpose). The Kalman-smoothing gain, $\boldsymbol{J}_{t-1}$, and

[55] For brevity, we let $\operatorname{Var}_{t-1}(\cdot)$ denote $\operatorname{Var}(\cdot|Z_{t-1})$. In addition, for any positive-(semi)definite square matrix $\boldsymbol{P}$, $\boldsymbol{P}^{0.5}$ denotes a (not necessarily positive-(semi)definite) lower-triangular factorization, such that $\boldsymbol{P} = \boldsymbol{P}^{0.5}\left(\boldsymbol{P}^{0.5}\right)'$.

the square root of the residual variance of $s_t$ are contained in appropriate partitions of $\boldsymbol{L}_{t-1}$. Based on (A.8), it follows that $\text{Var}_{t-1}\left(s_t|\boldsymbol{\mathcal{Z}}_{t+p}\right) = \omega_{1:t-1,t+1:T}$, which is the second moment needed for the shadow-rate sampling problem in (A.1) of drawing $s_t$ conditional on the remainder of the shadow-rate path.

**Observations $t$ with $t+p > T$:** When the ELB is binding near the end of the data sample, in particular for $t$ with $t + p > T$, the signal vector must be limited to include only leads of $s_t$ and $x_t$ up to $T$. Specifically, letting $t = T - k$ (with $k < p$), the adapted signal vector is shortened to a length of $N_z \cdot k + N_x$ as follows:

$$\boldsymbol{\mathcal{Z}}_{T,k} = \begin{bmatrix} Z_T - Z_{T|T-k-1} \\ v_t^x \end{bmatrix} = \begin{bmatrix} \sum_{j=0}^{k} \boldsymbol{H}_k \, \boldsymbol{C}^{k-j} \, \boldsymbol{\Psi}_{t+j} \, w_{t+j} \\ D_x \, \boldsymbol{\Psi}_t \, w_t \end{bmatrix}, \tag{A.11}$$

where $\boldsymbol{H}_k = \begin{bmatrix} I_{N_z \cdot k} & 0_{N_z \cdot (p-k)} \end{bmatrix}$ is a selection matrix selecting the first $N_z \cdot k$ elements out of a vector of length $N_z \cdot p$. In the case of $t = T$ (and thus $k = 0$), the signal vector collapses to $\boldsymbol{\mathcal{Z}}_{T,0} = v_t^x$. The expressions for $\boldsymbol{J}_{t-1}$ and $\text{Var}\left(s_t|Z_{t-1}, Z_{t+1}, x_t\right)$ in (A.9) and (A.10) are adjusted accordingly.

## A.4 Extension to hybrid shadow-rate VAR

The shadow-rate sampling framework described above is straightforward to extend to the case of a hybrid VAR that features lagged shadow and actual rates as right-hand side variables. To augment the companion-form representation of the shadow-rate VAR in (A.5), let $I_t = \max\left(S_t, ELB\right)$ denote the companion-form vector of actual rates (with element-wise application of the $\max$ operator) to obtain the following companion form for a hybrid shadow-rate VAR:

$$Z_t = \boldsymbol{C} Z_{t-1} + \boldsymbol{F} I_{t-1} + \boldsymbol{\Psi}_t w_t \,, \qquad\qquad w_t \sim N(0, I). \tag{A.12}$$

As a slight generalization of the hybrid shadow-rate VAR described in Section 3, this representation allows actual and shadow rates to enter each VAR equation concurrently, which nests the hybrid shadow-rate VAR described in (5) as a special case with zero restrictions on the coefficients on lagged $i_t$ and $s_t$ in certain equations. In addition, let $\boldsymbol{I} = \boldsymbol{1} \cdot ELB$ denote the actual-rate counterpart of the vector of shadow-rate values, $\boldsymbol{S}$, as defined in (11).

For shadow-rate sampling in the hybrid shadow-rate VAR, the values of actual rates can be treated as known, exogenous regressors. To restate the missing value problem, $\boldsymbol{I}$ is added to the conditioning set (as VAR regressors but without the censoring constraint) so that the missing value problem becomes to draw from $\boldsymbol{S}|\left(\bar{\boldsymbol{Y}}, \boldsymbol{I}\right) \sim N(\boldsymbol{\mu}, \boldsymbol{\Omega})$. As before, the requirement that $\boldsymbol{S} < ELB$ is imposed as part of drawing from $\boldsymbol{S}|\boldsymbol{Y}$. (Recall that $\boldsymbol{Y}$ also includes $\boldsymbol{I}$.) The moments of the missing-value problem can be restated in a similar form as before, by adjusting the conditioning set for the inclusion of $\boldsymbol{I}$ as follows:

$$f\left(s_t \,\middle|\, s_{1:t-1}, s_{t+1:T}, \bar{\boldsymbol{Y}}, \boldsymbol{I}\right) = f\left(s_t \,\middle|\, Z_{t-1}, Z_{t+p} - Z_{t+p|t-1}, v_t^x\right), \tag{A.13}$$

where we now have $Z_{t+p|t-1} = E(Z_{t+p}|Z_{t-1}, \boldsymbol{I}) = \boldsymbol{C}^{p+1} Z_{t-1} + \sum_{j=0}^{p} \boldsymbol{C}^j \boldsymbol{F} I_{t+p-j-1}$ and $v_t^x = x_t - E(x_t|Z_{t-1}, \boldsymbol{I}) = D_x v_t$.

# B MCMC Estimation of Shadow-Rate VAR

Appendix A describes a Gibbs sampler that generates draws for the shadow rate, according to (13), taking specific values for VAR parameters, $\{C_j\}_{j=1}^p$ and $\{\Sigma_t\}_{t=1}^T$, as given.[56] Here we describe how the shadow-rate Gibbs sampler is embedded into the MCMC estimation of the shadow-rate VAR system.

Denoting the $m$th MCMC draws of the VAR parameters $\{C_j\}_{j=1}^p$, $\{\Sigma_t\}_{t=1}^T$, and the shadow rates $\{s_t\}_{t=1}^T$, by $\boldsymbol{C}^m$, $\boldsymbol{\Sigma}^m$ and $\boldsymbol{S}^m$, respectively, and denoting (as before) the observed data by $\boldsymbol{Y}$, the MCMC sampler iterates over the following three blocks, for $m = 1, 2, \ldots, M$.

$$\boldsymbol{S}^{(m)} | \boldsymbol{C}^{(m-1)}, \boldsymbol{\Sigma}^{(m-1)}, \boldsymbol{Y} \tag{A.14}$$

$$\boldsymbol{A}^{(m)} | \boldsymbol{S}^{(m)}, \boldsymbol{\Sigma}^{(m-1)}, \boldsymbol{Y} \tag{A.15}$$

$$\boldsymbol{\Sigma}^{(m)} | \boldsymbol{S}^{(m)}, \boldsymbol{C}^{(m)}, \boldsymbol{Y} \tag{A.16}$$

Henceforth, we will refer to iterations over (A.14)– (A.16), as "the MCMC sampler." We use $M = 2'000$ draws, of which an initial $1'000$ burn-in draws are discarded. Convergence of the MCMC chain is assessed by inefficiency factors and potential scale reduction factors of Gelman and Rubin (1992) (applied to first and last third of the MCMC chain), with values for the latter typically falling below 1.2 and often close to 1.0.

The first block of the MCMC sampler, given by (A.14), consists of a sequence of Gibbs sampling steps, iterating over (A.1) for $t = 1, 2, \ldots, T$, with details described earlier in this appendix. The second and third block of the sampler, in (A.15) and (A.16), correspond to the blocks of a conventional sampler for heteroskedastic VAR models, given "data" for $\boldsymbol{Y}$ and $\boldsymbol{S}$. Our specification of the heteroskedasticity that is captured by $\{\Sigma_t\}_{t=1}^T$ follows Carriero, Clark, and Marcellino (2019) and Carriero, et al. (2022a); details of generating draws from $\{C_j\}_{j=1}^p$ and $\{\Sigma_t\}_{t=1}^T$ follow their procedures.[57]

The Gibbs sampler for the truncated MVN in the first block of the sampler in (A.14) is a single-move sampler that draws one observation of $s_t$ at a time (conditional on previously sampled values for all others). Consequently, a single pass from the Gibbs sampler for the truncated normal does not generate a direct draw from $\boldsymbol{S}^{(m)} | \boldsymbol{C}^{(m-1)}, \boldsymbol{\Sigma}^{(m-1)}, \boldsymbol{Y}$.[58] Nevertheless, repeated iterations over (A.14), (A.15), and (A.16), with each pass over the shadow-rate block in (A.14) captured by a single pass of the Gibbs sampler in (A.3), will eventually generate draws from the joint posterior density of $\boldsymbol{C}$, $\boldsymbol{\Sigma}$, and $\boldsymbol{S}$.[59]

[56] Given shadow-rate data, our implementation of the stochastic-volatility VAR (4) follows closely the setup of Carriero, Clark, and Marcellino (2019), where $\Sigma_t$ can be broken further down into a set of slope parameters of a Cholesky factorization and a latent vector process representing stochastic volatilities. Nevertheless, for brevity, we refer here to $\{\Sigma_t\}_{t=1}^T$ as a set of VAR "parameters."

[57] Specifically, we let $v_t = A^{-1}\Lambda_t^{0.5}\varepsilon_t$, where $A$ is a lower unit-triangular matrix, $\Lambda_t$ is a diagonal matrix, and the vector of its diagonal elements is denoted $\lambda_t$, with $\log \lambda_t = \log \lambda_{t-1} + \eta_t$, $\eta_t \sim N(0, \Phi)$, and $\varepsilon_t \sim N(0, I)$. Other forms of heteroskedasticity could also be specified. For the purpose of our discussion, we subsume the slope parameters $A$ and variance parameters $\Omega$ in the block of parameters denoted by $\{\Sigma_t\}_{t=1}^T$. Our MCMC sampler also reflects the ordering of steps in SV estimation recommended by Del Negro and Primiceri (2015).

[58] In contrast, in the case of the linear missing-value problem in (12), multi-move sampling is feasible. The missing-value problem has a linear Gaussian state space representation, and a Kalman-smoothing sampler can directly draw from the (untruncated) multivariate normal distribution of the problem, for example, by employing the methods described in Durbin and Koopman (2002).

[59] Formally, the $m$th draw from the shadow-rate block in (A.14) is then obtained by a single iteration over $s_t^{(m)} \,\big|\, s_{1:t-1}^{(m)}, s_{t+1:T}^{(m-1)}, \boldsymbol{C}^{(m-1)}, \boldsymbol{\Sigma}^{(m-1)}, \boldsymbol{Y}$ for $T = 1, 2, \ldots, T$ and holding $m$ fixed.

In order to achieve higher computational efficiency, we conduct multiple passes of the Gibbs sampler at every iteration, $m$, of the MCMC sampler.[60] In doing so, we exploit the fact that the Kalman-smoothing gains, $\boldsymbol{J}_t$, and residual variances $\text{Var}\left(s_t|Z_{t-1}, Z_{t+1}, x_t\right)$, described in (A.9) and (A.10) above, depend only on prior draws of the VAR parameters $\boldsymbol{C}^{(m-1)}$ and $\boldsymbol{\Sigma}^{(m-1)}$, but not the sampled path for the shadow rate, $\boldsymbol{S}^{(m)}$. As noted already by Geweke (1991), the second-moment matrices required for multiple passes of the Gibbs sampler for the truncated MVN need to be computed only once (given $\boldsymbol{C}^{(m-1)}$ and $\boldsymbol{\Sigma}^{(m-1)}$), which makes it computationally relatively cheap to conduct multiple Gibbs passes. Denoting the number of Gibbs passes by $N$, we retain only the output sampled in the $N$th pass, treating the initial $N-1$ as burn-in passes. Similar to Waggoner and Zha (1999), the motivation behind our approach is to hand over a draw $\boldsymbol{S}^{(m)}$ to the remaining MCMC steps that avoids the higher serial dependence between MCMC steps resulting from the previously described single-pass approach, while also being computationally relatively cheap to produce compared to other elements of the MCMC setup.

Formally, for every MCMC draw $m$, we implement the shadow-rate sampling block in (A.14) with $N$ Gibbs passes as follows: For each $n = 1, 2, \ldots, N$, (and holding $m$ fixed) iterate over

$$s_t^{(m,n)} \mid s_{1:t-1}^{(m,n)}, s_{t+1:T}^{(m,n-1)}, \boldsymbol{Y} \qquad \text{for } t = 1, 2, \ldots, T\,. \tag{A.17}$$

For each $m$, we initialize the first Gibbs pass with $s_t^{(m,0)} = s_t^{(m-1,N)}, \forall\, t$, and we retain $\boldsymbol{S}^{(m)} = \{s_t^{(m,N)}\}_{t=1}^T$ as the $m$th draw of the sequence of shadow rates. In our application, we employ $N = 101$ Gibbs passes over (A.17), and thus 100 burn-in passes for every step, $m$, of the MCMC sampler. To further enhance the efficiency of the MCMC sampler in estimation of the simple shadow-rate VAR, we attempt to draw the sequence $\boldsymbol{S}$ via rejection sampling from a draw of the missing value problem with 10 attempts.[61] We do so at every iteration $m$ after burnin. However, depending on data sample and model specification, acceptance rates for entire trajectories of $\boldsymbol{S}$ to lie below $ELB$ tend to be at most moderate, necessitating frequent use of the Gibbs sampling techniques described above.

# References


Antolín-Díaz, Juan, Ivan Petrella, and Juan F. Rubio-Ramírez (2021), "Structural scenario analysis with SVARs," *Journal of Monetary Economics*, 117, 798–815.

Aruoba, S. Boragan, Pablo Cuba-Borda, Kenji Higa-Flores, Frank Schorfheide, and Sergio Villalvazo (2021), "Piecewise-linear approximations and filtering for DSGE models with occasionally-binding constraints," *Review of Economic Dynamics*, 41, 96–120.

Aruoba, S. Boragan, Marko Mlikota, Frank Schorfheide, and Sergio Villalvazo (2022), "SVARs with occasionally-binding constraints," *Journal of Econometrics*, 231, 477–499.


[60] Our approach of embedding the procedures for drawing from the truncated MVN into the MCMC sampler with multiple passes at each MCMC step builds on ideas developed by Waggoner and Zha (1999) for forecast simulations with soft restrictions.

[61] Multiple draws from the missing value problem are efficiently obtained from the disturbance smoothing sampler of Durbin and Koopman (2002), which allows to generate the draws in parallel within the same Kalman filtering/smoothing iterations.

Bauer, Michael D., and Glenn D. Rudebusch (2016), "Monetary policy expectations at the zero lower bound," *Journal of Money, Credit and Banking*, 48, 1439–1465.

Bäurle, Gregor, Daniel Kaufmann, Sylvia Kaufmann, and Rodney Strachan (2020), "Constrained interest rates and changing dynamics at the zero lower bound," *Studies in Nonlinear Dynamics & Econometrics*, 24, 2017–0098.

Berg, Tim Oliver (2017), "Forecast accuracy of a BVAR under alternative specifications of the zero lower bound," *Studies in Nonlinear Dynamics & Econometrics*, 21, 1–29.

Bernanke, Ben S., and Alan S. Blinder (1992), "The federal funds rate and the channels of monetary transmission," *The American Economic Review*, 82, 901–21.

Billi, Roberto M. (2020), "Output gaps and robust monetary policy rules," *International Journal of Central Banking*, 16, 125–152.

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

Board of Governors of the Federal Reserve System (2023), "Monetary Policy Report," March.

Botev, Z. I. (2017), "The normal law under linear restrictions: simulation and estimation via minimax tilting," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 79, 125–148.

Carriero, Andrea, Joshua C. C. Chan, Todd E. Clark, and Massimiliano Marcellino (2022a), "Corrigendum to: Large Bayesian vector autoregressions with stochastic volatility and non-conjugate priors," *Journal of Econometrics*, 227, 506–512.

Carriero, Andrea, Todd E. Clark, and Massimiliano Marcellino (2019), "Large Bayesian vector autoregressions with stochastic volatility and non-conjugate priors," *Journal of Econometrics*, 212, 137–154.

Carriero, Andrea, Todd E. Clark, Massimiliano Marcellino, and Elmar Mertens (2022b), "Addressing COVID-19 outliers in BVARs with stochastic volatility," *Review of Economics and Statistics*, forthcoming.

Chan, Joshua C. C., and Eric Eisenstat (2018), "Bayesian model comparison for time-varying parameter VARs with stochastic volatility," *Journal of Applied Econometrics*, 33, 509–532.

Chan, Joshua C. C., Gary Koop, and Simon M. Potter (2013), "A new model of trend inflation," *Journal of Business & Economic Statistics*, 31, 94–106.

Chan, Joshua C. C., and Rodney Strachan (2014), "The zero lower bound: Implications for modelling the interest rate," Working Paper series 42-14, Rimini Centre for Economic Analysis.

Chib, Siddhartha (1992), "Bayes inference in the Tobit censored regression model," *Journal of Econometrics*, 51, 79–99.

Chib, Siddhartha, and Edward Greenberg (1998), "Analysis of multivariate probit models," *Biometrika*, 85, 347–361.

Chopin, Nicolas (2011), "Fast simulation of truncated Gaussian distributions," *Statistics and Computing*, 21, 275–288.

Christensen, Jens H. E., and Glenn D. Rudebusch (2012), "The Response of Interest Rates to US and UK Quantitative Easing," *Economic Journal*, 122, 385–414.

——— (2016), "Modeling yields at the zero lower bound: Are shadow rates the solution?" in *Dynamic Factor Models* eds. by Eric Hillebrand, and Siem Jan Koopman, 35 of Advances in Econometrics: Emerald Publishing Ltd, 75–125.

Christensen, Jens H.E., and Glenn D. Rudebusch (2015), "Estimating shadow-rate term structure models with near-zero yields," *Journal of Financial Econometrics*, 13, 226–259.

Clark, Todd E. (2011), "Real-time density forecasts from Bayesian vector autoregressions with stochastic volatility," *Journal of Business and Economic Statistics*, 29, 327–341.

Clark, Todd E., and Francesco Ravazzolo (2015), "Macroeconomic forecasting performance under alternative specifications of time-varying volatility," *Journal of Applied Econometrics*, 30, 551–575.

Cogley, Timothy, Giorgio E. Primiceri, and Thomas J. Sargent (2010), "Inflation-gap persistence in the US," *American Economic Journal: Macroeconomics*, 2, 43–69.

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

Cong, Yulai, Bo Chen, and Mingyuan Zhou (2017), "Fast simulation of hyperplane-truncated multivariate normal distributions," *Bayesian Analysis*, 12, 1017–1037.

Crump, Richard K., Stefano Eusepi, Domenico Giannone, Eric Qian, and Argia Sbordone (2021), "A large Bayesian VAR of the United States economy," Staff Report 976, Federal Reserve Bank of New York.

D'Agostino, Antonello, Luca Gambetti, and Domenico Giannone (2013), "Macroeconomic forecasting and structural change," *Journal of Applied Econometrics*, 28, 82–101.

Damien, Paul, and Stephen G. Walker (2001), "Sampling truncated normal, beta, and gamma densities," *Journal of Computational and Graphical Statistics*, 10, 206–215.

Debortoli, Davide, Jordi Gali, and Luca Gambetti (2019), "On the empirical (ir)relevance of the zero lower bound constraint," in *NBER Macroeconomics Annual 2019, Volume 34*: National Bureau of Economic Research, Inc.

Del Negro, Marco, Domenico Giannone, Marc P. Giannoni, and Andrea Tambalotti (2017), "Safety, liquidity, and the natural rate of interest," *Brookings Papers on Economic Activity*, 48, 235–316.

Del Negro, Marco, and Giorgio E. Primiceri (2015), "Time varying structural vector autoregressions and monetary policy: A corrigendum," *Review of Economic Studies*, 82, 1342–1345.

Diebold, Francis X., and Roberto S. Mariano (1995), "Comparing predictive accuracy," *Journal of Business and Economic Statistics*, 13, 253–63.

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

Eberly, Janice C., James H. Stock, and Jonathan H. Wright (2020), "The Federal Reserve's current framework for monetary policy: A review and assessment," *International Journal of Central Banking*, 16, 5–71.

Francis, Neville R., Laura E. Jackson, and Michael T. Owyang (2020), "How has empirical monetary policy analysis changed after the financial crisis?" *Economic Modelling*, 84, 309–321.

Gelfand, Alan E., Adrian F. M. Smith, and Tai-Ming Lee (1992), "Bayesian analysis of constrained parameter and truncated data problems using Gibbs sampling," *Journal of the American Statistical Association*, 87, 523–532.

Gelman, Andrew, and Donald B. Rubin (1992), "Inference from iterative simulation using multiple sequences," *Statistical Science*, 7, 457–472.

Geweke, John (1991), "Efficient simulation from the multivariate normal and student-t distributions subject to linear constraints and the evaluation of constraint probabilities," in *Computing Science and Statistics: Proceedings of the Twenty-Third Symposium on the Interface* ed. by E. M. Keramidas, 571–578, Fairfax: Interface Foundation of North America, Inc.

Gilchrist, Simon, Raphael Schoenle, Jae Sim, and Egon Zakrajšek (2017), "Inflation dynamics during the financial crisis," *American Economic Review*, 107, 785–823.

Gilchrist, Simon, and Egon Zakrajšek (2012), "Credit spreads and business cycle fluctuations," *American Economic Review*, 102, 1692–1720.

Goncalves, Silvia, Ana Maria Herrera, Lutz Kilian, and Elena Pesavento (2021), "Impulse response analysis for structural dynamic models with nonlinear regressors," *Journal of Econometrics*, 225, 107–130.

Gonzalez-Astudillo, Manuel, and Jean-Philippe Laforte (2020), "Estimates of r* consistent with a supply-side structure and a monetary policy rule for the U.S. economy," Finance and Economics Discussion Series 2020-085, Board of Governors of the Federal Reserve System.

Gust, Christopher, Edward Herbst, David López-Salido, and Matthew E. Smith (2017), "The empirical implications of the interest-rate lower bound," *American Economic Review*, 107, 1971–2006.

Horrace, William C. (2005), "Some results on the multivariate truncated normal distribution," *Journal of Multivariate Analysis*, 94, 209–221.

Ikeda, Daisuke, Shangshang Li, Sophocles Mavroeidis, and Francesco Zanetti (2021), "Testing the effectiveness of unconventional monetary policy in Japan and the United States," Papers 2012.15158, arXiv.org.

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

Jarocinski, Marek (2010), "Conditional forecasts and uncertainty about forecast revisions in vector autoregressions," *Economics Letters*, 108, 257–259.

Johannsen, Benjamin K., and Elmar Mertens (2021), "A time series model of interest rates with the effective lower bound," *Journal of Money, Credit and Banking*, 53, 1005–1046.

Jones, Callum, Mariano Kulish, and James Morley (2022), "A structural measure of the shadow federal funds rate," Centre for Applied Macroeconomic Analysis 61/2022, Australian National University.

Joslin, Scott, Anh Le, and Kenneth J. Singleton (2013), "Why Gaussian macro-finance term structure models are (nearly) unconstrained factor-VARs," *Journal of Financial Economics*, 109, 604–622.

Kailath, Thomas, Ali H. Sayed, and Babak Hassibi (2000), *Linear Estimation*, Prentice Hall Information and System Sciences Series: Pearson Publishing.

Kim, Don, and Kenneth J. Singleton (2012), "Term structure models and the zero bound: An empirical investigation of Japanese yields," *Journal of Econometrics*, 170, 32–49.

Koop, Gary, M. Hashem Pesaran, and Simon M. Potter (1996), "Impulse response analysis in nonlinear multivariate models," *Journal of Econometrics*, 74, 119–147.

Koop, Gary, and Simon M. Potter (2011), "Time varying VARs with inequality restrictions," *Journal of Economic Dynamics and Control*, 35, 1126–1138.

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

——— (2015), *Zero Lower Bound Term Structure Modeling: A Practitioner's Guide*: Palgrave Macmillan.

——— (2020), "A note of caution on shadow rate estimates," *Journal of Money, Credit and Banking*, 52, 951–962.

Kulish, Mariano, James Morley, and Tim Robinson (2017), "Estimating DSGE models with zero interest rate policy," *Journal of Monetary Economics*, 88, 35–49.

Lombardi, Marco J., and Feng Zhu (2018), "A shadow policy rate to calibrate U.S. monetary policy at the zero lower bound," *International Journal of Central Banking*, 14, 305–346.

Mavroeidis, Sophocles (2021), "Identification at the zero lower bound," *Econometrica*, 89, 2855–2885.

Nakajima, Jouchi (2011), "Monetary policy transmission under zero interest rates: An extended time-varying parameter vector autoregression approach," *The B.E. Journal of Macroeconomics*, 11.

Park, Jung Wook, Marc G. Genton, and Sujit K. Ghosh (2007), "Censored time series analysis with autoregressive moving average models," *The Canadian Journal of Statistics / La Revue Canadienne de Statistique*, 35, 151–168.

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

Reifschneider, David, and John C. Williams (2000), "Three lessons for monetary policy in a low-inflation era," *Journal of Money, Credit and Banking*, 32, 936–966.

Robert, Christian P. (1995), "Simulation of truncated normal variables," *Statistics and Computing*, 5, 121–125.

Rogers, John H., Chiara Scotti, and Jonathan H. Wright (2018), "Unconventional monetary policy and international risk premia," *Journal of Money, Credit and Banking*, 50, 1827–1850.

Sims, Eric, and Jing Cynthia Wu (2021), "Evaluating central banks' tool kit: Past, present, and future," *Journal of Monetary Economics*, 118, 135–160.

Swanson, Eric T., and John C. Williams (2014), "Measuring the effect of the zero lower bound on medium- and longer-term interest rates," *American Economic Review*, 104, 3154–3185.

Taylor, John B. (1993), "Discretion versus policy rules in practice," *Carnegie-Rochester Conference Series on Public Policy*, 39, 195–214.

Waggoner, Daniel F., and Tao Zha (1999), "Conditional forecasts in dynamic multivariate models," *The Review of Economics and Statistics*, 81, 639–651.

West, Kenneth D. (1996), "Asymptotic inference about predictive ability," *Econometrica*, 64, 1067–1084.

Wu, Jing Cynthia, and Fan Dora Xia (2016), "Measuring the macroeconomic impact of monetary policy at the zero lower bound," *Journal of Money, Credit and Banking*, 48, 253–291.

——— (2020), "Negative interest rate policy and the yield curve," *Journal of Applied Econometrics*, 35, 653–672.

Wu, Jing Cynthia, and Ji Zhang (2019), "A shadow rate New Keynesian model," *Journal of Economic Dynamics and Control*, 107, p. 103728.