Submitted to *Unknown Journal*

Shadow-Rate VARs

ANDREA CARRIERO
School of Economics and Finance, Queen Mary University of London and Department of Economics, University of Bologna

TODD E. CLARK
Research Department, Federal Reserve Bank of Cleveland

MASSIMILIANO MARCELLINO
Department of Economics, Bocconi University

ELMAR MERTENS
Research Centre, Deutsche Bundesbank

VARs are popular for forecasting and structural analysis, but ill-suited to handle occasionally binding constraints, like the effective lower

---

Andrea Carriero: a.carriero@qmul.ac.uk
Todd E. Clark: todd.clark@researchfed.org
Massimiliano Marcellino: massimiliano.marcellino@unibocconi.it
Elmar Mertens: elmar.mertens@bundesbank.de

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 Boragan Aruoba, Sidd Chib, Refet Gürkaynak, Sophocles Mavroeidis, Emanuel Mönch, Serena Ng, Frank Schorfheide, Harald Uhlig and seminar participants at the Bank of England, Bocconi University, Cleveland Fed, Deutsche Bundesbank, Penn, 2023 IAAE meeting, 2023 MEG meeting, and 2023 NBER Summer Institute. 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. A supplementary online appendix with additional results and replication codes are available at https://github.com/elmarmertens/CCMMshadowrateVAR-code. Declarations of interest: none.

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 find benefits to including both actual and shadow rates serve as explanatory variables. Historical shadow-rate estimates indicate that the FOMC would have set the funds rate much lower than it could during recent ELB episodes. In historical forecast accuracy, when compared to a standard VAR, shadow-rate VARs generate superior predictions for interest rates and deliver some gains for macroeconomic variables. Our structural analysis of shocks to financial conditions show strong differences in the reaction of interest rates depending on whether the ELB binds or not, with implications for the response of economic activity to the shocks.

KEYWORDS. Macroeconomic forecasting, effective lower bound, term structure, censored observations.
JEL CLASSIFICATION. C34, E17, E43.

1. INTRODUCTION

Linear vector autoregressions (VARs), popular models for forecasting and structural analysis, 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. 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 et al. (2019), and Rogers et al. (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,

[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.

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 (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).

An upshot of the term structure literature is the availability of shadow-rate estimates, such as those regularly updated by Krippner (2013, 2015) and Wu and Xia (2016). Their availability has led some researchers to avoid dealing with censoring by taking a shortcut of plugging these shadow-rate estimates in as data for the nominal short-term interest rate during an ELB episode. 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. Indeed, as documented in an earlier working paper version of this manuscript (Carriero et al., 2023), such plug-in VARs also generate inferior out-of-sample forecasts compared to the approach proposed here. More broadly, 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 characterizing monetary policy in ELB episodes and conducting structural analysis, and we 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. Throughout, 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.[3]

Specifically, we consider two specifications. In the "fully-hybrid shadow-rate VAR," lags of shadow and actual rates enter all VAR equations. Thus, in this case, macroeconomic indicators and shadow rates can respond to lags of shadow and actual interest rates. In the spirit of the structural VAR developed in Mavroeidis (2021), the model allows intercepts and dynamics of all variables to be different when at the ELB than when away. In contrast, the "block-hybrid shadow-rate VAR" imposes certain block zero restrictions so that shadow and actual interest rates do not appear jointly in a given VAR equation, and ELB effects propagate more selectively. The distinction between the models is motivated by the following: For observations when the ELB is not binding, actual and shadow rates are perfectly collinear, and their roles can only be distinguished and identified with ELB data. The (yet) limited experience with ELB data thus poses a challenge for generating reliable out-of-sample forecasts when actual and shadow rates are jointly included as regressors in each VAR equation. In the block-hybrid shadow-rate VAR, in equations for macroeconomic variables, only lagged actual rates appear on the right-hand side, and in the equations for shadow rates, only shadow rates appear on the right-hand side. These restrictions are motivated by the idea that spending and investment decisions should be conditioned on actual interest rates. In contrast, having lagged shadow rates in the interest rate block of the VAR helps to capture lower-for-longer or make-up elements in the conduct of monetary policy when forming expectations about future interest rates. Through the effects of lagged actual rates on macroeconomic variables, the block-hybrid specification still features some, albeit more limited, ELB effects on the VAR's dynamics. Empirically, the block-hybrid model generates shadow rate estimates

[3]In our empirical application on US data, we consider the ELB to have a 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). The supplementary online appendix documents the robustness of our main findings to using a value of 12.5 basis points instead.

and structural inferences similar to those obtained with the full hybrid model, while providing more reliable forecasts.

Our approach extends the unobserved components model of Johannsen and Mertens (2021) to the general VAR setting. In their setting, the censoring of actual rates affected the model's measurement equation, but not its state dynamics. In contrast, by including actual rates as VAR regressors, the state dynamics of our models are also affected by the ELB. Moreover, 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 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 15-20 US macroeconomic and financial variables, with stochastic volatility, which has been shown to generate competitive forecasts when ignoring the ELB (e.g., Carriero et al. (2019)). Critically, we also demonstrate how to handle data in which the ELB binds for multiple interest rate instruments of different maturities.

We begin by using our shadow-rate VAR formulations to illustrate the consequences of the ELB constraint for monetary policy in the US since 2009. On the one hand, the FOMC's setting of the federal funds rate was constrained by the ELB; it could not set the rate as low as it would have based on historical relationships among macroeconomic and financial indicators and interest rates. On the other hand, the FOMC used forward guidance and large-scale asset purchases to try to lower interest rates at longer maturities to mitigate the ELB constraint's impact and provide monetary accommodation. Still, even taking account of these actions, the FOMC may have set a funds rate below the ELB had that been possible. Taking on board observed data and the historical comovement of interest rates with a range of macroeconomic indicators, our shadow-rate models can be used to describe what the path for the federal funds rate would have been during these episodes in the absence of the ELB.

We first consider a 15-variable specification with a range of macroeconomic indicators and the federal funds rate as the only interest rate. We obtain a shadow rate estimate that turns significantly negative (how negative varies across the fully-hybrid and block-hybrid models) following the Great Financial Crisis (GFC) and the outbreak of the COVID-19 pandemic. These estimates imply that, based on macroeconomic conditions and historical relationships, the FOMC would have set the funds rate much lower than it could. We then extend the VAR to include additional interest rates — the 6-month Treasury bill rate and yields for several other bond maturities.[4] The inclusion of term structure information (beyond a measure of the short-term policy rate) turns out to importantly influence estimated shadow rates, especially in the block-hybrid model. The inclusion of interest rates other than the federal funds rate generates less negative shadow rates, representing less outsized expectations of monetary stimulus (through the federal funds rate) in response to the GFC. 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. Taking account of these unconventional policy actions, our model estimates indicate that the FOMC's setting of the funds rate was significantly constrained by the ELB in these episodes, but by less than estimates that are not directly informed by other interest rates.

Having used our shadow-rate VARs to characterize policy in past ELB episodes, we turn to examining possible benefits in macroeconomic forecasting. In our discussion, we focus on the 20 variable specification because inclusion of these longer-term yields improves forecast accuracy. In out-of-sample simulations for the US since 2009, interest rate forecasts obtained from our block-hybrid shadow-rate VAR 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 block-hybrid

[4] Specifically, we include maturities ranging from the daily federal funds rate to 10-year Treasury yields and 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 et al. (2022) and Kulish et al. (2017).

shadow-rate VAR are broadly on par with those from a standard VAR that ignores the ELB, although with some advantages for a few variables. 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 block-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. Similar to Guerrón-Quintana et al. (2023), we compute non-linear impulses responses (IRFs) as in Goncalves et al. (2021, 2023). The IRF approach builds on Koop et al. (1996) and Hayashi and Koeda (2019), and compares a baseline forecast (where the shock does not occur) against an alternative forecast that is affected by the identified shock. Both forecasts reflect the non-linearities of our model, 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 Gilchrist and Zakrajšek (2012), incorporated as an additional variable in our hybrid shadow-rate VAR. We compare responses for a shock origin of December 2006, 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.

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. 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 were no greater than that achieved by a standard VAR ignoring the ELB. Our specifications can also be used to characterize monetary policy when constrained by the ELB and to conduct structural analysis, to capture, for example, the effects of shocks to financial conditions 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 discusses implications for the characterization of policy in ELB episodes. Section 6 provides the forecast evaluation, and Section 7 illustrates the use of our hybrid model for structural analysis. Section 8 concludes. A supplementary online appendix provides technical details on our estimation procedure and additional empirical results. (An even more extensive set of additional results can be found in an earlier working paper version of this manuscript (Carriero et al., 2023).)

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 et al. (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.[5] 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

[5] 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.

work that has concluded that conventional and unconventional monetary policies work in a similar fashion. Relatedly, Guerrón-Quintana et al. (2023) study non-linear dynamic factor models. When applied to a shadow-rate model for the term structure of interest rates, they find "little evidence of nonlinearities in the factor dynamics" (as opposed to non-linearities in the censored measurement equation), which is consistent with the VAR representation used in our paper. Several DSGE models formulate monetary policy in terms of censored prescriptions from a policy rule for shadow rates as in Ikeda et al. (2023), Aruoba et al. (2021), Gust et al. (2017), Jones et al. (2022), and Kulish et al. (2017).[6] In this spirit, our shadow-rate VARs combine actual and shadow interest rates in their state dynamics, in reduced forms that combine elements of the "censored" and "kinked" VAR of Mavroeidis (2021).[7]

In the context of structural VAR models (SVARs), Ikeda et al. (2023),Aruoba et al. (2022) and Mavroeidis (2021) study 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

[6]The shadow-rate estimates of Aruoba et al. (2022), Gust et al. (2017), reflect data on real activity, inflation, and a short-term interest; in addition to those variables, Ikeda et al. (2023), Jones et al. (2022), and Kulish et al. (2017) include also a medium- or long-term interest rate.

[7]Our shadow-rate specifications assign a role to lagged shadow 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.

on the proximity of interest rates of different maturities to the ELB with the nonlinear impulse response methods of Goncalves et al. (2021, 2023) to 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 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 term structure literature (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.[8] 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)

[8]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 Choi et al. (2022), Doh and Choi (2016), and Lombardi and Zhu (2018), who define the shadow rate more broadly as a common factor of interest rates and possibly other variables intended to capture monetary policy.

ELB, which will inform inference about s_t. In the data, at a given point in time, the ELB may 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.[9] In the baseline specification of our empirical application, we include a total of $N_i = 6$ 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.[10] In another specification we consider, the model includes only the federal funds rate (and a single shadow rate).

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 to simplify exposition, we have the following system of N_y equations for a VAR with p lags:

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

Anticipating our subsequent application, we assume time-invariant transition matrices, C_j, but allow for time-varying shock volatilities, Σ_t, as in Clark (2011),

[9]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 in the N_i-dimensional vector i_t for which the ELB has been binding in the sample.

[10]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.

among others.[11] 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. 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 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. However, as reported in the supplementary online appendix, we find the shadow-rate approach presented below to improve upon the historical forecast accuracy of a linear VAR that omits shorter-term yields (and is thus unencumbered by the ELB).

[11] In our empirical application, we follow Carriero et al. (2019) and assume that $v_t = A^{-1}\varepsilon_t$, where A is a lower unit-triangular matrix, $\varepsilon_t \sim \mathcal{N}(0, \Lambda_t)$, and Λ_t is a diagonal matrix. The vector of diagonal elements of Λ_t is denoted λ_t, with $\log \lambda_t = \log \lambda_{t-1} + \eta_t$, $\eta_t \sim \mathcal{N}(0, \Phi)$. Other forms of heteroskedasticity could also be specified.

3.2 *Fully-hybrid shadow-rate VAR*

The 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). In explaining the model and estimation, our notation will partition y_t in a vector of $N_x = N_y - N_s$ other variables (macroeconomic and financial indicators that are not interest rates), x_t, that have unbounded support, and the nominal interest rate i_t, with x_t ordered on top:

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

The fully-hybrid shadow-rate VAR relates all variables (i.e., both x_t and s_t) to lags of both the shadow rates and actual interest rates:

$$x_t = \sum_{j=1}^{p} C_{xx,j}\, x_{t-j} + \sum_{j=1}^{p} C_{xs,j}\, s_{t-j} + \sum_{j=1}^{p} C_{xi,j}\, i_{t-j} + v_{x,t}\,, \tag{4}$$

$$s_t = \sum_{j=1}^{p} C_{sx,j}\, x_{t-j} + \sum_{j=1}^{p} C_{ss,j}\, s_{t-j} + \sum_{j=1}^{p} C_{si,j}\, i_{t-j} + v_{s,t}. \tag{5}$$

In this fully-hybrid specification, both actual interest rates and shadow rates predict future macroeconomic variables and interest rates, with nominal interest rates modeled as censored processes. Shadow rates may be seen as capturing effects of unconventional monetary policies (such as forward guidance or asset purchases), whereas actual rates are paid (earned) by borrowers (lenders) and thereby enter in economic dynamics and predictions.

This general specification allows for some coefficient changes at the ELB. When all shadow rates are above the ELB, then $s_t = i_t$ and the VAR becomes

$$x_t = \sum_{j=1}^{p} C_{xx,j}\, x_{t-j} + \sum_{j=1}^{p} (C_{xs,j} + C_{xi,j})\, s_{t-j} + v_{x,t}\,,$$

$$s_t = \sum_{j=1}^{p} C_{sx,j}\, x_{t-j} + \sum_{j=1}^{p} (C_{ss,j} + C_{si,j})\, s_{t-j} + v_{s,t}.$$

When all shadow rates are below the ELB, then $i_t = ELB$, and the VAR takes the form

$$x_t = \sum_{j=1}^{p} C_{xx,j}\, x_{t-j} + \sum_{j=1}^{p} C_{xs,j}\, s_{t-j} + ELB \cdot \sum_{j=1}^{p} C_{xi,j} + v_{x,t}\,,$$

$$s_t = \sum_{j=1}^{p} C_{sx,j}\, x_{t-j} + \sum_{j=1}^{p} C_{ss,j}\, s_{t-j} + ELB \cdot \sum_{j=1}^{p} C_{si,j} + v_{s,t}.$$

As this indicates, when interest rates are constrained at the ELB, the fully-hybrid shadow-rate VAR includes parameter changes in the coefficients on s_{t-1} and an intercept shift relative to when interest rates are not constrained. When the lags of some shadow rates are above the ELB, and other lags are below, a linear combination of these changes occurs. Moreover, the intercept shift is not arbitrary, but reflects the difference in coefficient loadings on s_{t-j} (i.e., reflects $C_{xi,j}$ and $C_{si,j}$, $j = 1, \ldots, p$) when at or above the ELB. Of course, to separately identify C_{xs} and C_{ss} on the one hand, and C_{xi} and C_{si} on the other, we need ELB data where $i_t = ELB > s_t$.

3.3 *Block-hybrid shadow-rate VAR*

Because actual and shadow rates are perfectly collinear when above the ELB, their joint inclusion in the fully-hybrid shadow-rate VAR can create challenges in identification and estimation, in particular for the purpose of generating useful out-of-sample forecasts and impulse responses based on the (yet) still limited experience with ELB data. Hence, as a feasible alternative, we also consider a block-hybrid specification in which zero restrictions are imposed such that the 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, collected in x_t. In terms of the fully-hybrid VAR's equations (4) and (5), the restrictions are $C_{xs} = 0$ and $C_{si} = 0$, so that the block-hybrid model takes the following form:

$$x_t = \sum_{j=1}^{p} C_{xx,j}\, x_{t-j} + \sum_{j=1}^{p} C_{xi,j}\, i_{t-j} + v_{x,t}\,, \tag{6}$$

$$s_t = \sum_{j=1}^{p} C_{sx,j}\, x_{t-j} + \sum_{j=1}^{p} C_{ss,j}\, s_{t-j} + v_{s,t}. \tag{7}$$

In this specification, all coefficients are identifiable from sample periods not constrained by the ELB (when we have $s_t = i_t$). 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.

In the block-hybrid system of equations (6) and (7), the censored values of (lagged) actual interest rates are state variables that influence the evolution of macroeconomic 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 (but not shadow rate 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 relate our reduced-form shadow-rate VARs to the structural VAR 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

$$\underbrace{\begin{bmatrix} A_{11} & A_{12}^{*} & A_{12} \\ A_{21} & A_{22}^{*} & A_{22} \end{bmatrix}}_{=A} \begin{bmatrix} x_t \\ s_t \\ i_t \end{bmatrix} = \sum_{j=1}^{p} B_{x,j}\, x_{t-j} + \sum_{j=1}^{p} B_{s,j}\, s_{t-j} + \sum_{j=1}^{p} B_{i,j}\, i_{t-j} + \varepsilon_t. \tag{8}$$

The general structural form in (8) has time-invariant parameters. But, as discussed by Mavroeidis (2021), the corresponding reduced-form VAR has time-varying parameters due to the interplay of the censoring constraint on nominal interest rates, and the actual-rate terms on the left-hand side of (8) that constrain the impact responses. Compared with equation (8), our models omit the impact terms of the actual rate, without which (8) leads to the following reduced-form representation:

$$\begin{bmatrix} x_t \\ s_t \end{bmatrix} = \sum_{j=1}^{p} A^{-1} B_{x,j}\, x_{t-j} + \sum_{j=1}^{p} A^{-1} B_{s,j}\, s_{t-j} + \sum_{j=1}^{p} A^{-1} B_{i,j}\, i_{t-j} + A^{-1}\, \varepsilon_t\,, \tag{9}$$

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

Evidently, the fully-hybrid shadow-rate VAR in equations (4) and (5) corresponds to (9), and has time-invariant coefficients for the lags of x_t, s_t, and i_t. As illustrated before, actual rates are either constant (and equal to ELB) or identical to shadow rates, so that our VAR models feature slope coefficient changes (in their response to lagged actual and shadow rates) and intercept shifts at the ELB.[12] Our block-hybrid model also conforms to (9), but with certain block-zero

[12] In the context of his general "censored and kinked" (CKSVAR) model, Mavroeidis (2021) derives additional parameter changes at the ELB. These additional parameters disappear, however, when $A_{12} = 0$ and $A_{22} = 0$. This configuration leads to the coefficient values $\kappa = 1$ and $\tilde{\beta} = 0$ in his setup for which the time-varying parameters characterized by Proposition 2 of his paper disappear.

restrictions on lag responses of actual and shadow rates (as discussed above). Specifically, these zero restrictions are imposed on the upper block of coefficients in $C_{s,j} = A^{-1}B_{s,j}$ (for each j) that link lagged shadow rates to macro variables, as well as the lower block in $C_{i,j} = A^{-1}B_{i,j}$ (for each j), linking lagged actual rates to shadow rates. As such, our block-hybrid model combines elements of the "censored" and "kinked" SVARs of Mavroeidis (2021): For macroeconomic variables, the block-hybrid VAR borrows from his kinked specification, whereas for shadow rate equations it follows his purely censored model. These choices are deliberately intended to condition spending and investment decisions on actual interest rates. In contrast, having lagged shadow rates in the interest rate block of the VAR helps to capture lower-for-longer or make-up elements in the conduct of monetary policy when forming expectations about future interest rates.

All told, there is a close correspondence between our shadow rate VARs and the general CKSVAR model of Mavroeidis. As in his model, the censoring of lagged actual rates generates slope coefficient and intercept shifts at the ELB. Compared to his structural model, we omit the impact terms for actual rates on the left-hand side of this SVAR system in (8), which leads to a straightforward reduced form representation as in (9). As discussed by Mavroeidis (2021), the inclusion of all impact terms, with non-zero A_{12} and A_{22} as in (8), imposes over-identifying restrictions on all coefficients in A to ensure coherence and completeness of the SVAR (and thus uniqueness of its reduced form). When the terms involving i_t are dropped from the left-hand side of (8), the impact responses of i_t are solely determined by the impact responses of s_t and the censoring constraint. By forgoing these over-identifying restrictions, our model specification enables fairly straightforward estimation with MCMC and the application to larger VAR systems. Moreover, the added zero restrictions on lag coefficients embedded in the block-hybrid VAR enables the identification of all VAR coefficients from observations regardless whether they are at or away from the ELB. The latter feature enables our application to out-of-sample forecasting that is competitive relative to a standard VAR within the confines of the (yet) limited ELB episodes in US data.

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 of our models can be thought of as embedding a monetary policy reaction function that relates the shadow rate to the variables included in the VAR.[13] 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."[14]

---

[13]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.

[14]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.

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.[15] In any 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. In fact, as argued by Hamilton (2017), without the inclusion of lagged shadow rates in the (notional) policy rule, it is hard to explain persistent episodes at the ELB. For all these reasons, we expect that the shadow rate of our VAR specifications broadly captures overall 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 unconventional monetary policies work in a similar fashion, Wu and Zhang (2019) derive a DSGE model in which

[15] Hamilton (2018) provides a cogent review of the efficacy of large scale asset purchases.

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 et al. (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 et al. (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.

As noted above, our fully-hybrid shadow-rate VAR allows for discrete parameter changes at the ELB, but does not perform well in out-of-sample forecasting. As an additional check, in an earlier working paper version of this manuscript (Carriero et al., 2023), to entertain ELB-driven parameter change while constraining its extent to make estimation and forecasting feasible, we consider a version of the block-hybrid VAR extended to allow the impact matrix A to change at the ELB. However, in out-of-sample forecasting results, the resulting model performs poorly compared to our baseline constant-parameter models with a shadow rate. Empirical evidence of time variation in the parameters controlling the block-hybrid shadow-rate VAR dynamics, also reported in the supplementary online appendix, is rather limited as well.

The constant-parameter specification of our block-hybrid shadow-rate VAR 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 et al. (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 et al. (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.[16]

Throughout, we use $p = 12$ lags in a monthly data set, which is described in further detail in Section 4. While omitted above, our models also include intercepts. 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 the supplementary online appendix.

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 $i_t = ELB$. Our shadow-rate VAR systems belong to a class of

[16] 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 jth-order own lag are set to θ_1/j^{θ_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 $\theta_4 = 2$.

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 Σ_t.[17] 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.

Specifically, consider the following setup for the fully-hybrid shadow-rate VAR given by (3), (4), and (5): Values for the VAR coefficients $\{C_{\cdot,j}\}_{j=1}^{p}$ and error variances $\{\Sigma_t\}_{t=1}^{T}$ are given and the data for $\{x_t\}_{t=1}^{T}$ and $\{i_t\}_{t=1}^{T}$ 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 $t = 1$. We assume that at $t = 1$, p lags of the data for x_t, i_t, and thus also s_t are known (since for $t < 1$ the ELB has not been binding). 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.

[17] 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 (like $\{\Sigma_t\}_{t=1}^{T}$).

For ease of reference, we collect all *unobserved* shadow rates in a vector $\boldsymbol{S}$ and all observations of $y_t = [x_t' \ i_t']'$ in a vector $\boldsymbol{Y}$:

$$\boldsymbol{S} = \begin{bmatrix} s_{T^*} \\ s_{T^*-1} \\ \vdots \\ s_2 \\ s_1 \end{bmatrix}, \quad \text{and} \quad \boldsymbol{Y} = \begin{bmatrix} y_T \\ y_{T-1} \\ \vdots \\ y_0 \\ y_{-p+1} \end{bmatrix}. \tag{10}$$

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 case of multiple ELB episodes, observations with $s_t = i_t > ELB$ are excluded from $\boldsymbol{S}$.

The task of the shadow-rate sampler is then to sample $\boldsymbol{S} \,|\, \boldsymbol{Y}$, 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 a "missing value" problem for $\boldsymbol{S}$ that does not condition on information that the ELB has been binding for certain observations, and thus does not impose $\boldsymbol{S} \leq ELB$. We denote the missing value problem by $\boldsymbol{S} \,|\, \boldsymbol{Y}^*$. Linearity and Gaussianity of the missing-value problem result in a posterior that is multivariate normal, and truncation at the ELB leads to the following shadow-rate sampler:[18]

$$\boldsymbol{S} \,|\, \boldsymbol{Y}^* \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Omega}) \quad \Rightarrow \quad \boldsymbol{S} \,|\, \boldsymbol{Y} \sim \mathcal{TN}(\boldsymbol{\mu}, \boldsymbol{\Omega}, -\infty, ELB). \tag{11}$$

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 the supplementary online appendix.

A further contribution of our paper is the implementation of the shadow-rate sampler via Gibbs sampling, following Geweke (1991) and Park et al. (2007), and adapted to the variance-covariance structure of the VAR(p) case, rather than the

[18] The notation $\boldsymbol{S} \sim \mathcal{TN}(\boldsymbol{\mu}, \boldsymbol{\Omega}, a, b)$ 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.

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 (11) 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 the supplementary online appendix, 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.[19]

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

[19]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 et al. (2013) and Koop and Potter (2011).

(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 $k = 1, 2 \ldots, h$. 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 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 the VAR system in (4) and (5). At each forecast horizon, censoring of predicted interest rates is applied to generate actual rate values, which are fed into the VAR equations 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.

To estimate the responses of the economy to shocks, rather than computing standard impulse response functions conditional on whether the ELB binds or not, we let the shock affect whether (and when) the ELB binds over the course of the response horizon. To do so, we build on the generalized impulse response approach of Koop et al. (1996), and implement conditional nonlinear impulse responses as in Goncalves et al. (2021, 2023).[20] At a given forecast origin t, we simulate a predictive density that conditions on a specific realization of the shock at

[20] Goncalves et al. (2023) refer to these as conditional average responses.

$t+1$ and compare its predictive mean against the average of a baseline prediction, made at t, that does not condition on specific shock realizations. In these computations, we start with posterior distributions for VAR coefficients, shadow-rate values, and volatilities estimated with the full sample of data. Importantly, our IRF calculations integrate over the entire posterior distribution of lagged shadow-rate values, instead of jumping off any specific estimate for current and past values of the shadow rate.[21]

In our application, we denote the size of the added shock by σ and set its size equal to the time-series average of the stochastic volatility estimates for 2005-2006, the two years preceding the Great Recession.[22] We then estimate the responses to a specific shock of size σ that occurs at $t+1$. Building on our earlier notation, let $v_t = A^{-1}\varepsilon_t$ denote the vector of VAR residuals, with $\varepsilon_t \sim \mathcal{N}(0, \Lambda_t)$, A a unit-lower-triangular matrix, and Λ_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 = 60 months), as well as the orthogonalized VAR residuals ε_{t+h}. For the baseline simulation, all shocks are drawn freely and denoted $\tilde{\varepsilon}_{t+h}$. For the alternative simulation, we add to these baseline shocks a shock to EBP at $t+1$. Our impulse-response definition thus corresponds to the setup of Goncalves et al. (2021, 2023). 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 $E_t^{(m)}(y_{t+h}|\varepsilon_{1,t+1} = \tilde{\varepsilon}_{1,t+1} + \sigma)$, respectively.[23] The impulse response function at a given MCMC node m follows as:

$$\Psi_{t,h}^{(m)}(\sigma) \equiv E_t^{(m)}\left(y_{t+h} \;\middle|\; \varepsilon_{1,t+1} = \tilde{\varepsilon}_{1,t+1} + \sigma\right) - E_t^{(m)}\left(y_{t+h}\right). \tag{12}$$

---

[21]By contrast, Ikeda et al. (2023) condition their IRFs on point estimates for lagged shadow rates.

[22]The posterior median path for the EBP shock's volatility averaged at a value of 0.11 in 2005-2006.

[23]To reduce Monte Carlo error, we generate these predictive densities with antithetic variates, meaning that for each path based on a given set of draws of orthogonalized VAR residuals, ε_t, and SV shocks, we simulate additional paths based on the negatives of these draws.

We report medians and uncertainty bands for the distribution of $\Psi_{t,h}^{(m)}$ across MCMC nodes.

4. DATA

Our data set consists of monthly observations for either 15 or 20 macroeconomic and financial variables for 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. Table 1 lists the baseline model's variables and their transformations to logs or log-differences. 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 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. In our initial assessment of policy in ELB episodes and in some subsequent comparisons, 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.[24]

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

[24] https://www.federalreserve.gov/econresdata/notes/feds-notes/2016/files/ebp_csv.csv.

TABLE 1. List of variables

Variable	FRED-MD code	transformation	Minnesota prior
PANEL A: Non-interest-rate variables (x_t)			
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
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
PANEL B: Nominal interest rates (i_t)			
Federal Funds Rate	FEDFUNDS		1
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 used in our BVARs (with prior means on all other VAR coefficients set to zero).

the COVID-19 outbreak and the FOMC's quick 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.

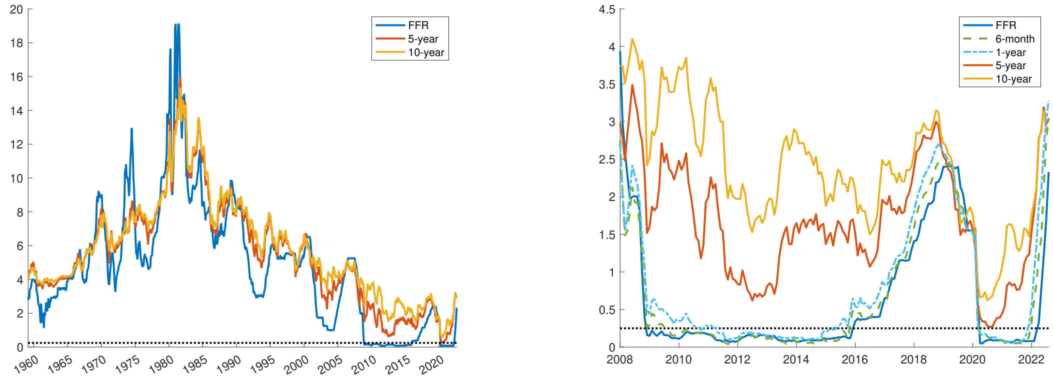

(a) Full data sample (b) Since the GFC

FIGURE 1. Interest rate data. All interest rates quoted as annualized percentage rates. Data obtained from FRED-MD as listed in Table 1.

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.[25] The supplementary online appendix shows that our main results are robust to instead setting the value of the ELB to 12.5 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

[25]For these yields, the lower bound constraint is an issue when simulating the predictive density, but not for estimating the VAR.

constraint falls outside our shadow-rate VAR framework. But a couple of considerations could be seen as supporting our approach: First, the shadow rate VARs allow 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 rates matter in the forecasting equations for non-interest-rate variables (which in turn 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.[26]

5. Shadow-Rate Estimates

We start our empirical analysis with an assessment of the shadow-rate estimates' implications for the description of monetary policy in ELB episodes. Figure 2 reports our shadow rate estimates (posterior medians and 90 percent credible sets) associated with the federal funds rate, along with comparisons to some other estimates. Panel (a) shows smoothed, full-sample estimates obtained from our fully-hybrid and block-hybrid shadow-rate VARs that omit bond yields and include only the federal funds rate as an interest rate measure in the data vector.[27] Panel (b) provides corresponding estimates for models including our full variable

[26]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 a pure finance model would still be silent on specific sensitivities of macroeconomic variables when interest rates are near the ELB, as captured by our shadow-rate VARs.

[27]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.

set with additional interest rates. Panel (c) compares our full-sample estimate from the fully-hybrid model with all 20 variables to the shadow-rate measures from Krippner (2013, 2015) and Wu and Xia (2016) based on affine term structure models. Finally, for the block-hybrid shadow-rate VAR including all 20 variables, Panel (d) compares the full-sample estimates that use data through August 2022 and quasi-real-time estimates that are the end-of-sample estimates produced by recursive estimation of the model starting in January 2009.

As indicated in Panel (a) of Figure 2, when the only interest rate included in the model is the federal funds rate, the fully-hybrid shadow-rate VAR generates a shadow rate estimate (black line with gray shading) that declines significantly starting in 2009, reaching a low point below -2 percent and remaining well below zero for a considerable period before gradually drifting up to the ELB. The shadow rate estimate from this model showed similar behavior during the first year of the COVID-19 pandemic, dropping quickly in the spring of 2020. The shadow rate estimate from the block-hybrid VAR (red lines) shows much sharper declines when the ELB binds, reaching a low of nearly -7 percent in the first ELB episode; the estimate is particularly negative in the 2009-2013 period that corresponds more closely to the recovery in real activity. The relatively sharper decline of the shadow rate from the block-hybrid model as compared to the fully-hybrid model reflects the fact that, in the latter specification, the shadow rate can inherit the persistence of the actual, constrained short-term interest rate on the right-hand side of the VAR. Used to characterize monetary policy, the estimates from both specifications imply that, based on macroeconomic conditions and historical relationships, the FOMC would have set the funds rate much lower than it could due to the ELB; the ELB sharply constrained policy. The contours of our block-hybrid shadow-rate estimates generated from macroeconomic and financial variables while excluding term structure data closely resemble unconstrained Taylor-rule prescriptions calculated for the Great Recession years by Eberly et al. (2020).

Adding other interest rates to the shadow-rate VARs leads to some changes in the characterization of monetary policy, but also makes estimates from the

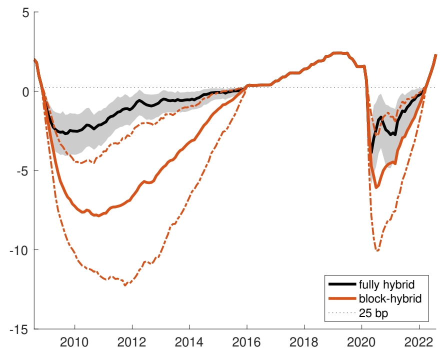

(a) Without yields: fully vs. block-hybrid

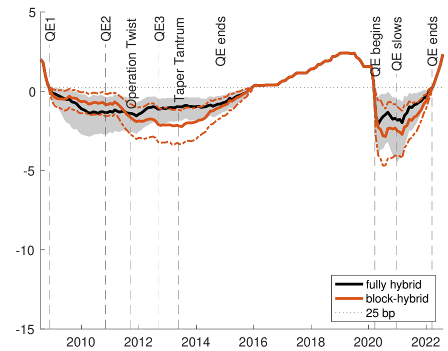

(b) With yields: fully vs. block-hybrid

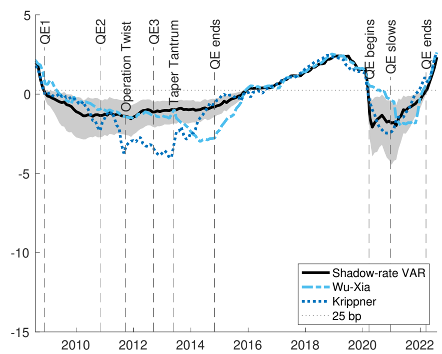

(c) Shadow-rate VAR & other estimates

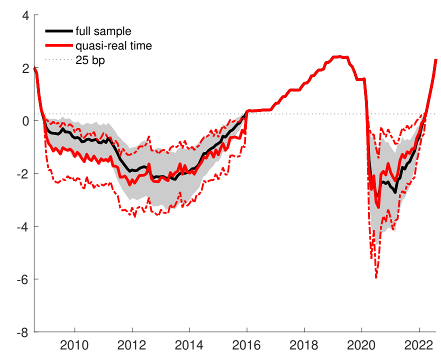

(d) Quasi-real time estimates

FIGURE 2. Shadow-rate estimates. Panel (a) compares the smoothed shadow-rate estimates from the fully-hybrid and block-hybrid specifications when data on interest rates other than the federal funds rate are omitted. Panel (b) compares the shadow-rate estimates from the fully-hybrid and block-hybrid specifications using all 20 variables listed in Table 1 in the estimation. Panel (c) compares the smoothed shadow-rate estimates from the fully-hybrid shadow-rate VAR (estimated with yields, also shown in Panel (b)) against updated estimates obtained from Krippner (2013, 2015) and Wu and Xia (2016). Panel (d) compares smoothed (also shown in Panel (b)) and quasi-real time shadow-rate estimates from our block-hybrid shadow-rate VAR using all 20 variables.

fully-hybrid (black line with gray shading) and block-hybrid (red lines) specifications more similar. As indicated in Panel (b), full-sample estimates from VARs including our baseline set of 20 variables including various interest rates show the shadow rate falling significantly starting in 2009, but neither as rapidly nor as sharply as in the estimates from the VAR without other bond yields.[28] For example, in the block-hybrid model informed by bond yields, the shadow rate estimate reached a nadir of -2.1 percent in mid-2013 (rather than a nadir of nearly -7 in the block-hybrid estimates without bond yields). 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 mid-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 mid-March 2022. From 2009 through 2011, the fully-hybrid and block-hybrid estimates of the shadow rate are very similar, whereas the block-hybrid estimate runs modestly below the fully-hybrid estimate from 2012 through 2014 before quickly converging when the Federal Reserve first raised the funds rate in December 2015. Overall, the estimates from our models with and without bond yields indicate that the shadow rate estimates are significantly informed by the interest rate equations.

Finally, the estimates for the block-hybrid model with all bond yields in Panel (d) indicate that, as might be expected, the quasi-real-time estimates (red lines) have more time variability than do the full sample estimates (black line with gray shading), but follow a quite similar contour. As might also be expected, the quasi-real-time estimates are less precise than the full sample estimates, with credible sets wider than those of the full sample estimate. However, in results omitted in the interest of brevity, comparisons to estimates from VARs including only the federal funds rate show that the inclusion of bond yields improves the precision of the shadow rate estimate in quasi-real time and reduces its variability over time.

[28]The supplementary online appendix provides estimates of shadow rates for the 6-month and 1-year Treasury maturities. The contours of these estimates follow those shown here for the shadow federal funds rate.

Our interpretation of the shadow rate estimates from the VARs with various bond yields as compared to the VARs with the federal funds rate as the only interest rate is that the inclusion of yields allows the shadow-rate VARs to capture the effects of asset purchases and forward guidance regarding the path of policy rates. Taking account of these unconventional policy actions, our estimates indicate that the FOMC's setting of the funds rate was significantly constrained by the ELB in these episodes, but by much less than is indicated by estimates that are not directly informed by other interest rates.

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 (c) of Figure 2, our fully-hybrid shadow-rate VAR estimate (black line with gray shading) 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 fully-hybrid 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.

One might wonder whether 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.[29] In the interest of brevity, the supplementary online appendix provides additional results to assess the effects of shadow-rate modeling and enforcement of the ELB in model estimation. A simple missing data approach 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

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

forecasting, conditioning estimates on information when the ELB was binding could (and should) embody non-trivial information about the relevant parameters of the VAR. Comparisons provided in the supplementary online appendix show that shadow-rate sampling, which takes into account observations of interest rates at the ELB, can have important effects on model estimates of parameters and SV, which in turn bear on the shadow rate estimates. In particular, without forcing the draws of missing interest rate observations to lie at or below the ELB (which is what one would do under a simple missing-data approach), the posterior credible set of the policy rate puts significant probability mass above the ELB for some periods in which, in fact, the federal funds rate was at the ELB. 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).

6. Forecast Evaluation

Having shown how our shadow-rate VAR can be used to characterize monetary policy in ELB episodes, we turn now to a second way in which our approach can be used: macroeconomic forecasting. 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 online 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. We should also note that, in view of the benefits of including other interest rates as described below, in these results we focus on 20-variable specifications that include not only the federal funds rate but also bond or bill rates.

6.1 *Average performance 2009–2017*

Comparing various model specifications discussed below, Table 2 provides 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 = 6, 12, and 24 months (unreported results for h = 3 months are similar). 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.[30]

Table 2 compares the accuracy of forecasts from our block-hybrid shadow-rate VAR to those from a standard VAR that simply takes the forecasts as given and does nothing to obey ELB constraints. As noted above, the hybrid specification relates the macroeconomic indicators to actual interest rates, not shadow rates. To facilitate comparisons, we report RMSE, MAE, and CRPS results for the 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.

[30] 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 $(1200/h)\sum_{s=1}^{s=h}\Delta\log(x_{t+s})$.

Overall, these results indicate that our block-hybrid 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 (and in some instances improving) 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 0.24 to 0.50, the MAE ratios range from 0.20 to 0.38, and the CRPS ratios range from 0.19 to 0.39, with statistical significance of all of the MAE and CRPS gains. For Treasury rates of maturities 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 0.41 to 0.68 (0.30 to 0.48); CRPS ratios are between 0.31 and 0.51, 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 0.50 to 0.56. At the 5- and 10-year maturities, the block-hybrid 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 0.83 to 0.94. 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 block-hybrid shadow-rate specifications are often near 1, particularly at the 6-month forecast horizon. At longer horizons, the block-hybrid specification improves on the longer-horizon accuracy of the standard VAR for some variables (e.g., income, consumption, IP, capacity utilization, and unemployment) and reduces longer-horizon accuracy for some other economic activity indicators (e.g., hours worked) as well as PCE inflation and S&P 500 returns.

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

TABLE 2. Forecast performance of standard vs. block-hybrid shadow-rate VAR

	RMSE			MAE			CRPS		
	6	12	24	6	12	24	6	12	24
Income	1.00	1.00	**0.65**	1.00	1.01	0.97	1.00	1.01	0.99
Consumption	1.01	0.99	**0.88*****	1.01	1.01	**0.90***	1.00	0.99	0.97**
IP	1.01**	1.02	0.96	1.01	1.01	0.98	1.01	1.01	1.00
Cap. Util.	0.99	1.02	0.96	0.99	1.00	**0.95**	1.00	1.01	0.97
Unemp.	1.00	0.99	**0.91*****	1.00	0.98	**0.91*****	1.00	0.99	**0.94*****
Nfm Pyrlls	0.99	1.04	0.99	1.00	1.01	**0.91****	1.00	1.01	0.97
Hours	1.01	1.02	1.03	1.01	1.02	1.03	1.01	1.01	1.01
H. Earnings	1.02*	1.01	1.03**	1.01	1.01**	1.01	1.01	1.01**	1.01**
PPI (Fin.)	0.99	0.97	1.00	0.99	0.99	1.01	0.99	0.99	1.00
PPI (Metals)	0.99	0.99	1.02	0.99	0.99	1.00	1.00	1.00	1.01
PCE Prices	0.98	0.97	**1.07***	0.97*	0.98	**1.08*****	0.99	0.99	**1.05*****
Hsng Strts	0.98	**0.91**	**0.90**	0.98	**0.91**	**0.90**	0.99	**0.94**	**0.93**
S&P 500	1.00	1.01**	**1.08***	0.99	1.01	1.02**	1.00	1.01*	1.01***
USD / GBP	1.00	1.00	1.02	1.01	1.02	**1.05**	1.00	1.00	1.02
FFR	**0.24**	**0.29**	**0.50**	**0.20****	**0.28***	**0.38****	**0.19****	**0.25****	**0.39*****
6m Tbill	**0.41**	**0.51**	**0.68**	**0.30****	**0.36****	**0.48*****	**0.31****	**0.40****	**0.51*****
1y Trsy	**0.62**	**0.66**	**0.78**	**0.44*****	**0.46****	**0.56*****	**0.50*****	**0.53*****	**0.56*****
5y Trsy	**0.90****	**0.82*****	**0.75****	**0.93**	**0.90**	**0.89**	**0.94**	**0.88****	**0.83*****
10y Trsy	**0.94**	**0.88***	**0.78*****	0.96	**0.88**	0.97	0.96	**0.93**	**0.95**
BAA Yld	0.97	1.00	0.99	0.96	1.01	**1.07**	0.99	1.02	**1.07**

Note: Comparison of "standard linear VAR" (baseline, in denominator) against "block-hybrid shadow-rate VAR" for horizons 36, 12, and 24 months. 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.

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, the supplementary online appendix document the benefits in forecast performance from including longer-term yields in the VAR. In addition, we report improved out-of-sample forecasts generated from our block-hybrid VAR compared to those from a linear VAR that omits shorter-term yields (and is thus unencumbered by the ELB). We also consider the effects of assuming a lower ELB value of 12.5 basis points (which corresponds to the middle of the FOMC's target range for the federal funds rate during the recent ELB episodes). Overall, contours of the estimated shadow rates, and relative forecast performance of the 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.

6.2 *Expected policy rates and forecasts of macroeconomic variables*

The key differences between the VAR specifications considered here (standard vs. hybrid shadow-rate) is the treatment of nominal interest rates near the ELB and the resulting forecasts for other, mainly macroeconomic, variables.[31] 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. 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 the block-hybrid shadow-rate VAR to construct forecasts for macroeconomic variables; for brevity we also refer to these as "forecast-relevant" policy rates. The block-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 (6).[32] 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 when predicting other variables.

[31] 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 within x_t.

[32] 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.

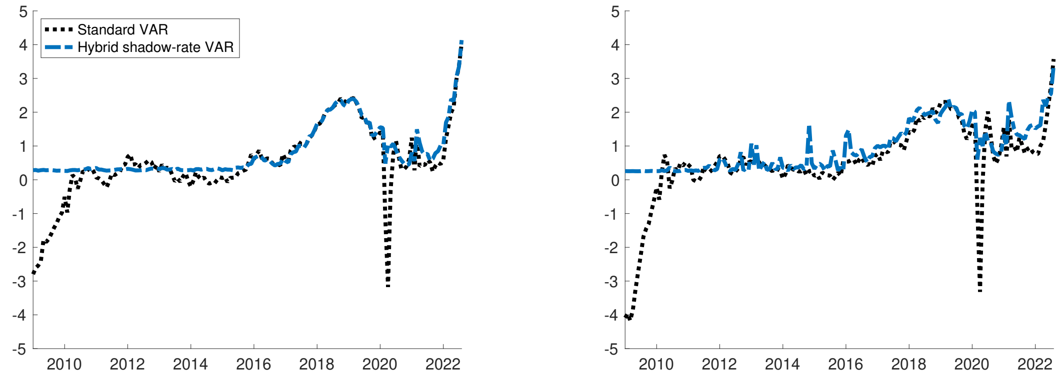

(a) Generated from data with yields (b) Generated from data excluding yields

FIGURE 3. Expected averages of forecast-relevant policy rates (12-months ahead). 12-month ahead expected averages of federal funds rates (actual or shadow), generated out-of-sample at forecast origins 2009:01 – 2022:08. 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 block-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.

To illustrate salient differences in the forecast-relevant policy rates that underlay predictions of macroeconomic (and other non-interest rate) variables, Figure 3 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 two VARs can differ — more so early in an ELB episode than late in the episode — 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 standard VAR predicts sharply negative interest rates in 2009 and some of 2010 and again briefly in the pandemic. 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 post-2010 recovery. Taken at face value in this reduced-form setting, the more negative rate forecasts of standard VAR imply more monetary policy stimulus as compared to the positive rate forecasts of the hybrid shadow-rate VAR. However, this distinction applies for relatively short portions of the ELB episodes, with predicted rates fairly comparable for substantial portions of the ELB episodes.

While the inclusion of term structure data does not change the broad qualitative patterns in the forecast-relevant policy rate forecasts generated by the different models, the right panel of Figure 3 displays some quantitative differences compared to the left panel. When the federal funds rate is the only interest rate included, the standard VAR predicts even lower interest rates in 2009 (a nadir of roughly -4 percent without yields as compared to about -3 percent with yields); 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 standard VAR to over-predict the strength of the recovery in its early stages. In addition, comparing the forecasts from the block-hybrid models without and with bond yields included, omitting yields leads the forecast of the policy rate to show more variability — rising above the ELB briefly before returning to it — in the second halves of the ELB episodes.

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 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 VAR 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. These comparisons focus here on the block-hybrid shadow-rate VAR as compared to a standard VAR. Figure 4 shows the evolution of federal funds rate forecasts over selected origins between March 2020 and the end of our sample in August 2022.[33] 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

[33] 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.

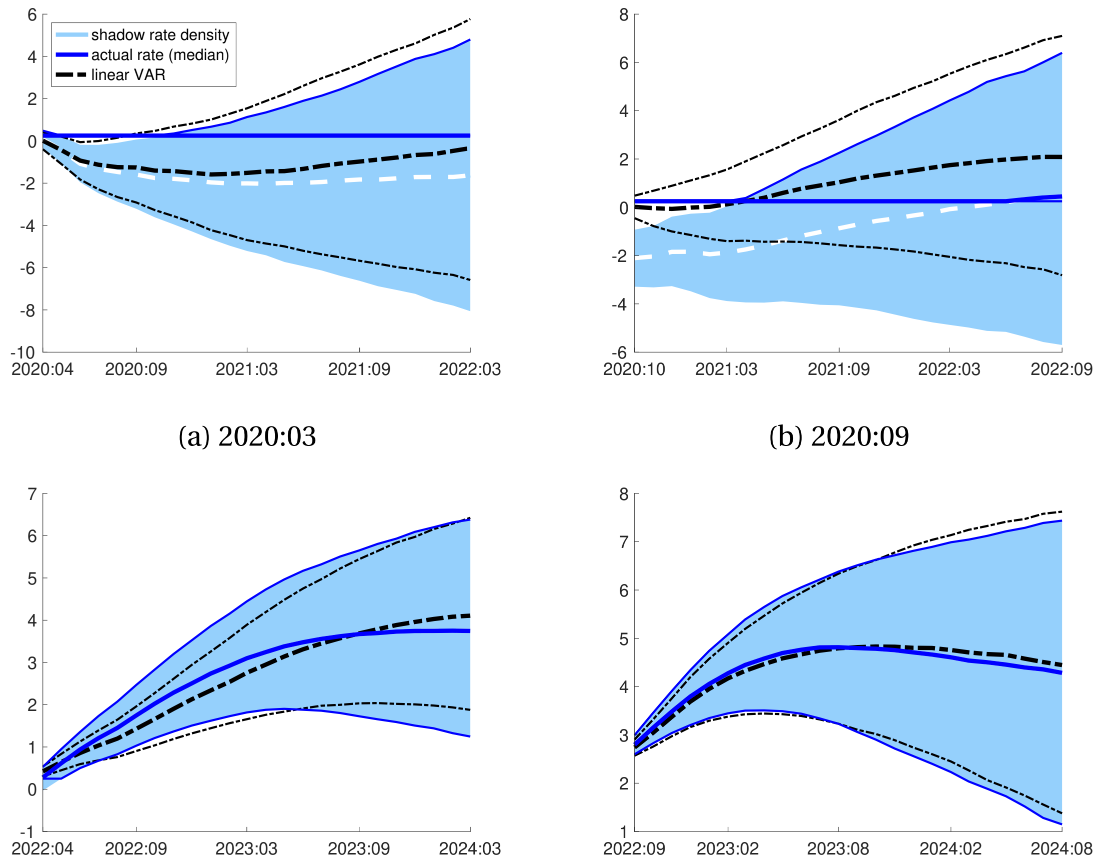

(a) 2020:03 (b) 2020:09

(c) 2022:03 (d) 2022:08

FIGURE 4. Predictive densities for actual and shadow values of the federal funds rate. 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 block-hybrid 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.

throughout the forecast 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 block-hybrid shadow-rate VAR was very similar 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 a little below -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.[34] As a result, until the second 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 shadow-rate VAR) were even more similar, with the funds rate rising from about 3 percent to nearly 5 percent before gradually declining.

Forecasts of the federal funds rate from the shadow-rate VARs reflect the predicted evolution of the shadow rate. As noted before, our shadow rates reflect the unconstrained policy rate prescriptions of the feedback rule for monetary policy that is implied by the shadow rate VAR. 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

[34]The shadow-rate VARs' predictive densities integrate 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.

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.

7. Structural Analysis

This section illustrates a third use of our model, for structural analysis. In this application, we estimate the response of the economy to a shock to financial conditions with our hybrid shadow-rate VAR. We focus here on the block-hybrid shadow-rate VAR, in which case the responses of macroeconomic variables to shocks depend on whether the ELB binds or not. Although this model does not allow for regime change in ELB episodes, the supplementary online appendix shows that the fully-hybrid shadow-rate VAR — which does allow for regime change — yields qualitatively similar impulse response estimates.[35] 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 estimate non-linear impulse responses based on the approaches of Koop et al. (1996) and Goncalves et al. (2021, 2023). 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 5 reports estimated impulse response functions for December 2006 (blue lines) and December 2012 (red lines).[36] 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

[35] In particular, posterior median estimates of responses are very similar across models. However, likely reflecting some collinearity of actual and shadow rates, credible sets are much wider for the fully-hybrid estimates compared to those from the block-hybrid model.

[36] In keeping with the timing assumptions described in Section 3.6, the IRF origins represent the jump-off point t of baseline forecasts against with the alternative forecast conditioned on a specific shock occurring at $t+1$.

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 IRFs 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 2006 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 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 2006 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 5's IRFs show 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,

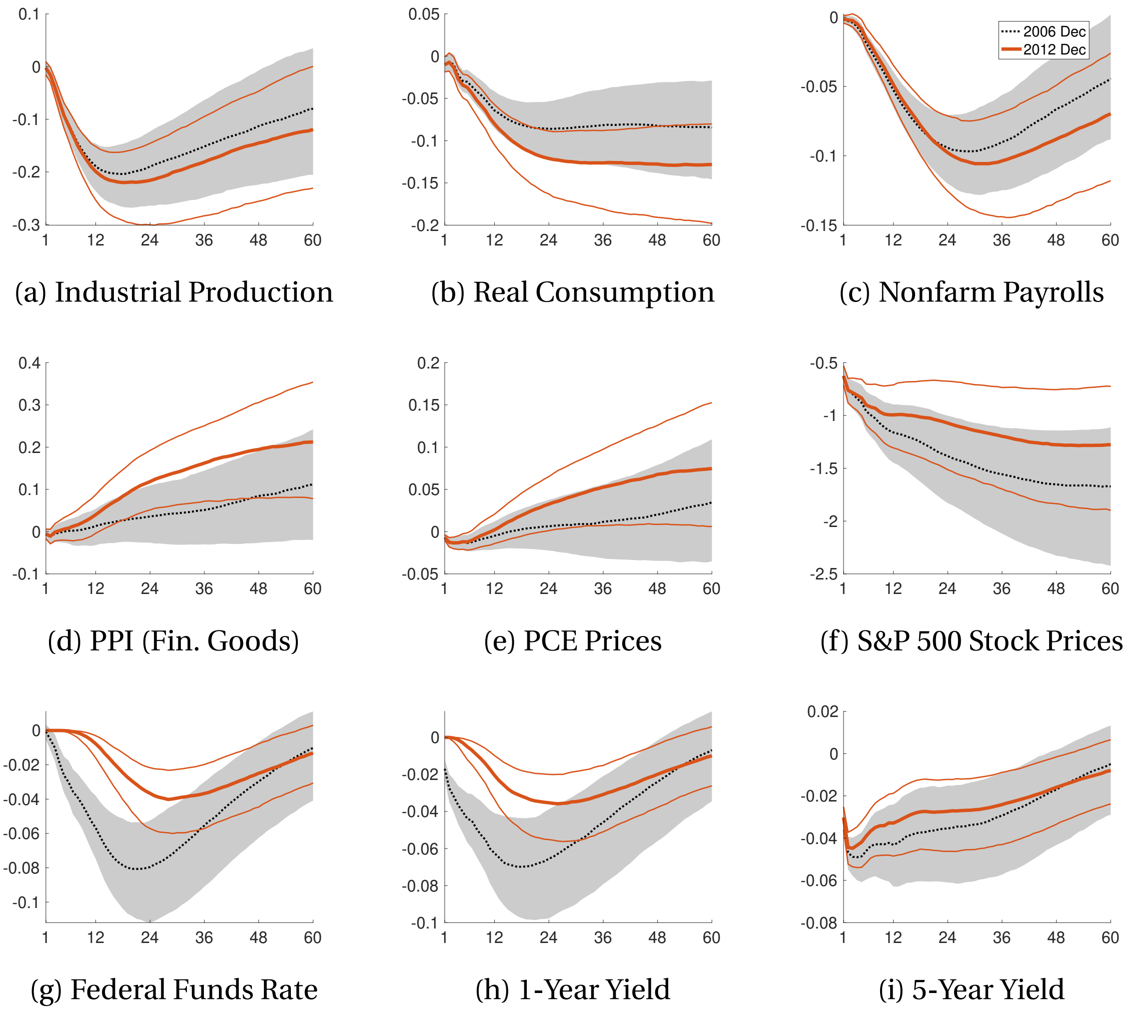

(a) Industrial Production (b) Real Consumption (c) Nonfarm Payrolls

(d) PPI (Fin. Goods) (e) PCE Prices (f) S&P 500 Stock Prices

(g) Federal Funds Rate (h) 1-Year Yield (i) 5-Year Yield

FIGURE 5. Impulse Responses to an EBP Shock in 2006 and 2012. Estimates obtained from full-sample posterior of the block-hybrid shadow-rate VAR at two different IRF origins with 60-month horizons. Size of EBP shock corresponds to average SV level of 0.11 for the years 2005-2006. Posterior medians and 68 percent bands. Interest rates in levels, and cumulated log levels (times 100) for all other variables.

the differences between the 2006 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 IRFs as compared to 2006, 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 2006 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 leads to a rise in inflation (rather than fall) due to financially constraint firms being forced to raise their prices in order to preserve internal liquidity and avoid external financing.

In the interest of brevity, we relegated to the supplementary online appendix another comparison, of IRFs obtained from a linear VAR that ignores the ELB against the responses from our hybrid model as shown here. The differences in the estimates are broadly similar to those shown in Figure 5's hybrid-model IRFs with origins in 2006 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. The supplementary online appendix also presents IRF estimates obtained from our fully-hybrid shadow-rate VAR, which are qualitatively similar to those shown here for the block-hybrid case. However, reflecting the relative scarcity of ELB data in our sample, the uncertainty bands around the fully-hybrid IRFs *at the ELB* are markedly wider than IRFs with shock origins away from the ELB as well as the IRFs in the block-hybrid case (where all parameters are identified from data at and away from the ELB).

Together, the IRF estimates presented in this section illustrate the use of our (block-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.[37]

[37] The results presented here are also confirmed by additional IRFs, obtained with a smaller-sized VAR model that only contains EBP, unemployment, inflation, and the federal funds rate, as reported in an earlier working paper version of this manuscript (Carriero et al., 2023). 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

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 et al. (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 that combine actual and shadow interest rates in their state dynamics and use shadow rates as dependent variables to describe the evolution of (shadow) term structure variables. The fully-hybrid shadow-rate VAR model relates all variables (both macroeconomic indicators and shadow rates) to lags of both the shadow rates and actual interest rates. This model allows intercepts and dynamics to be different when at the ELB than when away. Because 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

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.

data. Hence, the block-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. Through direct and indirect effects, both actual and shadow rates influence multi-step forecasts of all variables in the hybrid VAR.

We apply our shadow-rate models in three applications, using medium-scale BVARs with stochastic volatility: characterizing monetary policy in ELB episodes, forecasting a range of macroeconomic indicators over a long historical sample including ELB episodes, and conducting a structural analysis of the impacts of financial shocks.

In our results, historical estimates of the shadow rate from a block-hybrid model in which the federal funds rate is the only interest rate indicate that, based on macroeconomic conditions and historical relationships, the FOMC would have set the funds rate far, far lower than it could during the GFC and COVID-19 pandemic — on the order of -7 percent. However, once longer-term bond yields are included in the model, capturing some of the FOMC's unconventional policy actions, our estimates indicate that the FOMC's setting of the funds rate was significantly constrained by the ELB in these episodes, but by much less than estimates that are not directly informed by other interest rates.

In our forecast application, forecasts for interest rates obtained from a block-hybrid 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. However, especially at shorter horizons, 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 (1) when at or away from the ELB and (2) with respect to the same responses from a linear VAR.

References

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. [10]

Aruoba, S. Boragan, Marko Mlikota, Frank Schorfheide, and Sergio Villalvazo (2022), "SVARs with occasionally-binding constraints." *Journal of Econometrics*, 231 (2), 477–499. [10, 11, 21, 26]

Bauer, Michael D. and Glenn D. Rudebusch (2016), "Monetary policy expectations at the zero lower bound." *Journal of Money, Credit and Banking*, 48 (7), 1439–1465. [3, 4, 37]

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 (2), 2017–0098. [2]

Berg, Tim Oliver (2017), "Forecast accuracy of a BVAR under alternative specifications of the zero lower bound." *Studies in Nonlinear Dynamics & Econometrics*, 21 (2), 1–29. [10]

Bernanke, Ben S. and Alan S. Blinder (1992), "The federal funds rate and the channels of monetary transmission." *The American Economic Review*, 82 (4), 901–21. [19]

Billi, Roberto M. (2020), "Output gaps and robust monetary policy rules." *International Journal of Central Banking*, 16 (2), 125–152. [20, 45]

Black, Fischer (1995), "Interest rates as options." *The Journal of Finance*, 50 (5), 1371–1376. [3, 11]

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 (2), 506–512. [22, 50]

Carriero, Andrea, Todd E. Clark, and Massimiliano Marcellino (2019), "Large Bayesian vector autoregressions with stochastic volatility and non-conjugate priors." *Journal of Econometrics*, 212 (1), 137–154. [5, 13, 22, 50]

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. [39]

Carriero, Andrea, Todd E. Clark, Massimiliano Marcellino, and Elmar Mertens (2023), "Shadow-rate VARs." Discussion Papers 14/2023, Deutsche Bundesbank. [3, 8, 21, 49]

Chan, Joshua C. C., Gary Koop, and Simon M. Potter (2013), "A new model of trend inflation." *Journal of Business & Economic Statistics*, 31 (1), 94–106. [25]

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. [2, 10]

Chib, Siddhartha (1992), "Bayes inference in the Tobit censored regression model." *Journal of Econometrics*, 51 (1-2), 79–99. [25, 50]

Chib, Siddhartha and Edward Greenberg (1998), "Analysis of multivariate probit models." *Biometrika*, 85, 347–361. [25, 50]

Choi, Jason, Taeyoung Doh, Andrew Foerster, and Zinnia Martinez (2022), "Monetary policy stance is tighter than federal funds rate." Economic Letter 2022-30, Federal Reserve Bank of San Francisco. [11]

Christensen, Jens H. E. and Glenn D. Rudebusch (2012), "The response of interest rates to us and uk quantitative easing." *Economic Journal*, 122 (564), 385–414. [3]

Christensen, Jens H. E. and Glenn D. Rudebusch (2016), "Modeling yields at the zero lower bound: Are shadow rates the solution?" In *Dynamic Factor Models* (Eric Hillebrand and Siem Jan Koopman, eds.), volume 35 of *Advances in Econometrics*, 75–125, Emerald Publishing Ltd. [3]

Christensen, Jens H.E. and Glenn D. Rudebusch (2015), "Estimating shadow-rate term structure models with near-zero yields." *Journal of Financial Econometrics*, 13 (2), 226–259. [3]

Clark, Todd E. (2011), "Real-time density forecasts from Bayesian vector autoregressions with stochastic volatility." *Journal of Business and Economic Statistics*, 29 (3), 327–341. [12]

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. [2, 6, 13]

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* (Martin S. Eichenbaum, Erik Hurst, and Jonathan A. Parker, eds.), 141–170, National Bureau of Economic Research, Inc. [2, 22]

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 (1 (Spring)), 235–316. [35]

Diebold, Francis X. and Roberto S. Mariano (1995), "Comparing predictive accuracy." *Journal of Business and Economic Statistics*, 13 (3), 253–63. [37]

Doh, Taeyoung and Jason Choi (2016), "Measuring the stance of monetary policy on and off the zero lower bound." Economic Review 101(3), Federal Reserve Bank of Kansas City. [11]

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 (1), 5–71. [32]

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. [21]

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* (E. M. Keramidas, ed.), 571–578, Fairfax: Interface Foundation of North America, Inc. [24, 25]

Gilchrist, Simon, Raphael Schoenle, Jae Sim, and Egon Zakrajšek (2017), "Inflation dynamics during the financial crisis." *American Economic Review*, 107 (3), 785–823. [11, 49]

Gilchrist, Simon and Egon Zakrajšek (2012), "Credit spreads and business cycle fluctuations." *American Economic Review*, 102 (4), 1692–1720. [7, 11, 28, 46, 47]

Goncalves, Silvia, Ana María Herrera, Lutz Kilian, and Elena Pesavento (2021), "Impulse response analysis for structural dynamic models with nonlinear regressors." *Journal of Econometrics*, 225, 107–130. [7, 11, 26, 27, 46]

Goncalves, Silvia, Ana María Herrera, Lutz Kilian, and Elena Pesavento (2023), "State-dependent local projections." Working Papers 2302, Federal Reserve Bank of Dallas. [7, 11, 26, 27, 46]

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. [9]

Guerrón-Quintana, Pablo, Alexey Khazanov, and Molin Zhong (2023), "Financial and Macroeconomic Data Through the Lens of a Nonlinear Dynamic Factor Model." Finance and Economics Discussion Series 2023-027, Board of Governors of the Federal Reserve System (U.S.). [7, 10]

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 (7), 1971–2006. [10, 20, 45]

Hamilton, James D. (2017), "Commentary: Lower-Bound Beliefs and Long-Term Interest Rates." *International Journal of Central Banking*, 13 (3), 203–212. [20]

Hamilton, James D. (2018), "The efficacy of large-scale asset purchases when the short-term interest rate is at its effective lower bound." *Brookings Papers on Economic Activity*, 49 (2 (Fall)), 543–554. [20]

Hayashi, Fumio and Junko Koeda (2019), "Exiting from quantitative easing." *Quantitative Economics*, 10 (3), 1069–1107. [7]

Ikeda, Daisuke, Shangshang Li, Sophocles Mavroeidis, and Francesco Zanetti (2023), "Testing the effectiveness of unconventional monetary policy in Japan and the United States." *American Economic Journal: Macroeconomic*, forthcoming. [10, 27]

Iwata, Shigeru and Shu Wu (2006), "Estimating monetary policy effects when interest rates are close to zero." *Journal of Monetary Economics*, 53 (7), 1395–1408. [2, 10]

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. [4, 5, 9, 19, 23, 24, 25, 50]

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. [6, 10]

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 (3), 604–622. [9]

Koop, Gary, M. Hashem Pesaran, and Simon M. Potter (1996), "Impulse response analysis in nonlinear multivariate models." *Journal of Econometrics*, 74 (1), 119–147. [7, 26, 46]

Koop, Gary and Simon M. Potter (2011), "Time varying VARs with inequality restrictions." *Journal of Economic Dynamics and Control*, 35 (7), 1126–1138. [25]

Krippner, Leo (2013), "Measuring the stance of monetary policy in zero lower bound environments." *Economics Letters*, 118 (1), 135–138. [3, 21, 32]

Krippner, Leo (2015), *Zero Lower Bound Term Structure Modeling: A Practitioner's Guide*. Palgrave Macmillan. [3, 21, 32]

Krippner, Leo (2020), "A note of caution on shadow rate estimates." *Journal of Money, Credit and Banking*, 52 (4), 951–962. [3, 11]

Kulish, Mariano, James Morley, and Tim Robinson (2017), "Estimating DSGE models with zero interest rate policy." *Journal of Monetary Economics*, 88 (C), 35–49. [6, 10, 21]

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 (5), 305–346. [11]

Mavroeidis, Sophocles (2021), "Identification at the zero lower bound." *Econometrica*, 89 (6), 2855–2885. [3, 4, 10, 11, 17, 18, 21, 26]

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 (1). [2]

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 (1), 151–168. [24]

Reifschneider, David and John C. Williams (2000), "Three lessons for monetary policy in a low-inflation era." *Journal of Money, Credit and Banking*, 32 (4), 936–966. [20, 45]

Rogers, John H., Chiara Scotti, and Jonathan H. Wright (2018), "Unconventional monetary policy and international risk premia." *Journal of Money, Credit and Banking*, 50 (8), 1827–1850. [2]

Sims, Eric and Jing Cynthia Wu (2021), "Evaluating central banks' tool kit: Past, present, and future." *Journal of Monetary Economics*, 118, 135–160. [9, 20]

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 (10), 3154–3185. [2, 22]

Waggoner, Daniel F. and Tao Zha (1999), "Conditional forecasts in dynamic multivariate models." *The Review of Economics and Statistics*, 81 (4), 639–651. [35]

West, Kenneth D. (1996), "Asymptotic inference about predictive ability." *Econometrica*, 64 (5), 1067–1084. [37]

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 (2-3), 253–291. [3, 4, 9, 20, 21, 32]

Wu, Jing Cynthia and Fan Dora Xia (2020), "Negative interest rate policy and the yield curve." *Journal of Applied Econometrics*, 35 (6), 653–672. [2]

Wu, Jing Cynthia and Ji Zhang (2019), "A shadow rate New Keynesian model." *Journal of Economic Dynamics and Control*, 107, 103728. [9, 20, 21]