Submitted to *Quantitative Economics*

# Forecasting with Shadow-Rate VARs

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

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

MASSIMILIANO MARCELLINO
Department of Economics, Baffi, Bidsa, IGIER, Bocconi University and CEPR

ELMAR MERTENS
Research Centre, Deutsche Bundesbank

Vector autoregressions (VARs) are popular for forecasting, but ill-suited to handle occasionally binding constraints, like the effec-

Andrea Carriero: a.carriero@qmul.ac.uk
Todd E. Clark: todd.clark@researchfed.org
Massimiliano Marcellino: massimiliano.marcellino@unibocconi.it
Elmar Mertens: em@elmarmertens.com
We gratefully acknowledge helpful suggestions and comments received from Editor Jim Hamilton, three referees, 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, De Nederlandsche Bank, 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. An earlier version of this paper has also been circulated under the title "Shadow-rate VARs." Our shadow-rate estimates, a supplementary online appendix with additional results and replication codes are available at https://github.com/elmarmertens/CCMMshadowrateVAR-code. Declarations of interest: none.

tive lower bound on nominal interest rates. We examine reduced-form "shadow-rate VARs" that model interest rates as censored observations of a latent shadow-rate process and develop an efficient Bayesian estimation algorithm that accommodates large models. When compared to a standard VAR, our better-performing shadow-rate VARs generate superior predictions for interest rates and broadly similar predictions for macroeconomic variables. We obtain this result for shadow-rate VARs in which the federal funds rate is the only interest rate and in models including additional interest rates. Our shadow-rate VARs also deliver notable gains in forecast accuracy relative to a VAR that omits shorter-term interest rate data in order to avoid modeling the lower bound.

KEYWORDS. Macroeconomic forecasting, effective lower bound, term structure, censored observations.

JEL CLASSIFICATION. C34, E17, E43.

## 1. INTRODUCTION

Linear vector autoregressions (VARs) typically include 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. 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.

As one way around ELB constraints, 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. (2024), 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, 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, from which no-arbitrage restrictions are derived for yields of all maturities. The resulting restrictions are, however, nontrivial 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).

In this paper, we examine shadow-rate approaches for accommodating the ELB in commonly-used macroeconomic VARs, integrating the shadow-rate inference into the macroeconomic model. 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.[1] The shadow rates are assumed to be equal to observed rates when above the ELB. More specifically, we apply our shadow-rate approaches to a medium-scale Bayesian VAR (BVAR) for 15-20 US macroeconomic and financial variables, with stochastic volatility (SV), 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 rates of different maturities. As in other studies of shadow-rate VARs, we do not enforce any specific no-arbitrage restrictions like those featured in the term structure literature.[2]

[1] Studies including 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.

[2] In doing so, we follow previous literature that uses VARs to derive forecasts and expectational errors of financial and economic variables without imposing the restrictions of a specific structural 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

Our approaches include three VAR specifications. We develop Bayesian estimation methods for these models and extend existing shadow-rate approaches to the case of multiple rates. First, the "general shadow-rate VAR" corresponds to the reduced form specification that Mavroeidis (2021) derives from his structural VAR. This model, which includes a single interest rate, accommodates parameter shifts induced by hitting the ELB. Lags of both the shadow rate and actual interest rate enter all VAR equations; macroeconomic indicators and the shadow rate can respond to lags of both the shadow and actual interest rates. Second, the "non-structural shadow-rate VAR" imposes some restrictions on the extent of parameter change allowed at the ELB. Relative to the general model, this specification continues to accommodate some change at the ELB and to allow lags of shadow and actual rates to enter all VAR equations. In addition, this model allows the inclusion of multiple shadow and interest rates to which the ELB applies. Third, the "restricted version of the non-structural shadow-rate VAR" imposes certain block zero restrictions so that shadow and actual interest rates do not appear jointly in a given VAR equation: In equations for macroeconomic variables, only lagged actual rates appear on the right-hand side, whereas in equations for shadow rates, only shadow rates appear on the right-hand side. Through the effects of lagged actual rates on macroeconomic variables, the restricted specification still features some, more limited, ELB effects on the VAR's dynamics.

In our empirical analysis, we assess the benefits of our shadow-rate VARs in forecasting macroeconomic and financial variables in US data.[3] In out-of-sample simulations for the US since 2010, interest rate forecasts obtained from all of our shadow-rate VARs are clearly superior, in terms of both point and density accuracy, when compared to predictions from a standard VAR that ignores the ELB. In forecasting other economic indicators, our better-performing shadow-

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.

[3]Throughout, our analysis takes the level of the ELB as given at 25 basis points, consistent with other studies, such as Bauer and Rudebusch (2016), Johannsen and Mertens (2021), and Wu and Xia (2016). Our earlier working paper (Carriero et al., 2023) also shows that our main results are robust to instead setting the value of the $ELB$ to 12.5 basis points.

rate specifications match the accuracy of a standard VAR. In models with the federal funds rate as the only interest rate, the general shadow-rate VAR is comparable in macroeconomic forecast accuracy to the standard VAR, as is the non-structural shadow-rate VAR. With additional yields in the model, our restricted version of the non-structural shadow-rate VAR also matches the accuracy of the standard VAR in forecasting other economic indicators (while beating it in forecasting interest rates). Our shadow-rate VARs also deliver notable gains in forecast accuracy relative to a VAR that omits shorter-term interest rate data in order to avoid modeling the lower bound.

Examining full-sample estimates of the reduced-form shadow rates of our models, we begin with a 15-variable specification with the federal funds rate as the only interest rate, which yields a shadow rate estimate that quickly turned negative following the Great Financial Crisis (GFC) and the outbreak of the COVID-19 pandemic and eventually rose gradually back to and above the ELB. With multiple interest rates included in model, estimates from the restricted version of the non-structural shadow-rate VAR show a more sustained decline in the shadow rate.[4] The path of this reduced-form shadow rate estimate broadly resembles shadow rates from the affine term structure models of Krippner (2015) and Wu and Xia (2016). Collectively, our reduced-form shadow rate estimates imply that, based on macroeconomic conditions and historical relationships, the Federal Open Market Committee (FOMC) would have set the funds rate much lower than it could.

In the context of structural VAR models (SVARs), Aruoba et al. (2022), Ikeda et al. (2024), 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, reduced-form Bayesian VAR (with stochastic volatility), and we evaluate its application to forecasting. One of our key contributions is the development of a tractable Bayesian estimation algorithm that can be applied to re-

[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.

duced form models considerably larger than practical in existing structural VAR approaches. 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. Our Bayesian Gibbs sampler enables estimation of the reduced-form representation of the SVAR in Mavroeidis (2021), which applies to the case of a single interest rate subject to the ELB. Moreover, by adding some restrictions, we extend the model and its estimation to accommodate interest rates of multiple maturities at the ELB. Our approach builds on the work of Wei (1999) for a dynamic Tobit model as well as Chib (1992) and Chib and Greenberg (1998) for static Tobit and Probit models, which also featured data augmentation.

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 provides the forecast evaluation. Section 6 concludes. A supplementary online appendix provides technical details on our estimation procedure and additional empirical results.[5]

## 2. Related Literature

With some term structure researchers making updates of their shadow-rate estimates readily available (e.g., Krippner (2015) and Wu and Xia (2016)), some other researchers have taken a shortcut of plugging these shadow-rate estimates in as data for the nominal short-term interest rate during an ELB episode. In this vein, Francis et al. (2020) find that linear VARs estimated with shadow-rate estimates from Krippner (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. While convenient, this plug-in

[5] A more extensive set of additional results can be found in an earlier working paper version of this manuscript (Carriero et al., 2023).

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.[6]

Our approach extends the unobserved components model of Johannsen and Mertens (2021) to the general VAR setting.[7] In their setting, the censoring of actual rates affected the model's measurement equation, but not its state dynamics.[8] 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 (detailed below) to be able to estimate larger models. Sampling the shadow rate 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. Our algorithm development also includes a Bayesian complement to the maximum likelihood-based approach of Mavroeidis (2021) to tractably estimate a general shadow-rate VAR that accommodates the ELB-induced parameter change implied by his structural VAR.

Our approach also extends Carriero et al. (2021), which forecast bond yields with a BVAR-SV in just yields using a prior based on a no-arbitrage affine

---

[6]In addition, 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.

[7]Johannsen and Mertens (2021) provide an out-of-sample forecast evaluation for short- and long-term nominal interest rates in a model smaller than our VARs, and find their unobserved components shadow-rate model to be competitive with the no-arbitrage model of Wu and Xia (2016), but do not consider forecasts of other variables. Gonzalez-Astudillo and Laforte (2024) embed a shadow-rate model in an unobserved components model and report improved point forecasts for economic and financial variables from the shadow-rate approach.

[8]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.

term structure model. Carriero et al. (2021) also accommodated the ELB with the shadow rate treatment of Johannsen and Mertens (2021), using lags of the shadow rate on the right-hand side of the equations in forming out-of-sample forecasts. Our paper differs in that we include macroeconomic variables in a larger VAR and in that we allow lags of both actual and shadow rates as predictors when estimating the model and forming forecasts. As noted above, our paper also makes use of a more efficient approach to sampling shadow rates. Importantly, as noted above, our approach extends Carriero et al. (2021) by accommodating parameter change included by the ELB, rather than treating reduced-form VAR dynamics as unchanged at the ELB.

Finally, a number of studies have developed or deployed DSGE models that formulate monetary policy in terms of censored prescriptions from a policy rule for shadow rates. Examples include Aruoba et al. (2021), Gust et al. (2017), Ikeda et al. (2024), Jones et al. (2022), Kulish et al. (2017), and Wu and Zhang (2019).

## 3. Shadow-Rate VARs

This section presents the shadow-rate models considered in our empirical analysis. Throughout, we take the value of the lower bound, denoted $ELB$, as a given and known constant. A central element of our approach is to relate actual and shadow rates via a censoring equation known from Black (1995):

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

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. When the ELB is binding, so that $i_t = ELB$, the shadow rate is a latent variable that can only take values below (or equal to) $ELB$, which will inform inference about $s_t$.

Let $y_t = (x_t', i_t')'$ denote the vector of $N_y$ observed variables, with $N_x$ macroeconomic and financial indicators, denoted $x_t$, and $N_i$ interest rates, $i_t$, to which the ELB applies. Throughout our exposition of the models, we omit intercepts for simplicity and use $p$ to denote the VAR lag order.

We begin with a general reduced form model patterned after Mavroeidis (2021). This reduced form arises from a structural VAR, in which the variable set is limited to contain only a single interest rate, the federal funds rate, with a corresponding single shadow rate. In the case of a single interest rate there are only two regimes (at and away from the ELB) for which coherency and completeness need to be verified to ensure the uniqueness of the reduced-form solution. Extending the general model to include multiple interest rates would lead to a richer set of ELB-related regimes and thus also a richer, more complex set of conditions for coherency and completeness — a case we leave as a subject for future research. Instead, to include multiple interest rates in a reduced-form shadow-rate VAR, we consider a slightly simplified version of the general model, which satisfies the conditions of coherency and completeness in a straightforward fashion. We will refer to this model as a non-structural shadow-rate VAR.

With multiple interest rates included, the non-structural shadow-rate VAR 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 $N_i = N_s$, where the censoring of the shadow-rate vector $s_t$ is element-wise. 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 our data, we include a total of $N_i = 6$ interest rates, and the ELB has been binding for three of those. Finally, we consider a restricted version of the non-structural shadow-rate VAR which imposes block zero restrictions so that no lagged shadow rates appear in equations for macroeconomic variables, whereas no lagged actual rates appear in equations for shadow rates.

### 3.1 *General shadow-rate VAR for the case of a single interest rate*

The structural VAR model of Mavroeidis (2021) applies to the case where $i_t$ is a scalar, which leads the model to track two states ($i_t$ at and away from the ELB). Each regime allows for different shock impacts, and these are mirrored by a set of regime-specific coefficients in the model's reduced form representation:

$$x_t = b_{xs} s_t^* + \sum_{j=1}^{p} \Pi_{xx,j}\, x_{t-j} + \sum_{j=1}^{p} \Pi_{xs,j}\, s_{t-j} + \sum_{j=1}^{p} \Pi_{xi,j}\, i_{t-j} + v_{x,t}\,, \quad (2)$$

$$s_t = \sum_{j=1}^{p} \Pi_{sx,j}\, x_{t-j} + \sum_{j=1}^{p} \Pi_{ss,j}\, s_{t-j} + \sum_{j=1}^{p} \Pi_{si,j}\, i_{t-j} + v_{s,t}\,, \quad (3)$$

$$\text{where} \quad s_t^* \equiv \mathbb{1}(s_t < ELB) \cdot (s_t - ELB) \quad (4)$$

and $b_{xs}$, $\Pi_{xx,j}$, $\Pi_{xs,j}$, $\Pi_{xi,j}$, $\Pi_{sx,j}$, $\Pi_{ss,j}$, $\Pi_{si,j}$ denote coefficient matrices of appropriate dimension.[9] With $N_i = 1$, this general model corresponds to the reduced form specification in Proposition 2 of Mavroeidis (2021), who establishes point identification of the coefficient matrices and the partially latent shadow rate $s_t$.[10]

We augment the model to feature heteroskedastic shocks, and model the regression residuals as conditionally normal, with time-varying volatility:

$$\begin{bmatrix} v_t^x \\ v_t^s \end{bmatrix} = Q^{-1} \begin{bmatrix} \varepsilon_t^x \\ \varepsilon_t^s \end{bmatrix}, \qquad \text{with} \quad \varepsilon_t \equiv \begin{bmatrix} \varepsilon_t^x \\ \varepsilon_t^s \end{bmatrix} \sim \mathcal{N}\left(\mathbf{0}, \Lambda_t\right), \quad (5)$$

where $Q$ is an upper triangular matrix with ones on the diagonal and $\Lambda_t$ is a diagonal matrix. Letting $\lambda_t$ denote the vector of diagonal elements of $\Lambda_t$, the log volatility process is $\log \lambda_t = \gamma_0 + \gamma_1 \log \lambda_{t-1} + \eta_t$, with $\gamma_0$ a vector of intercepts, $\gamma_1$ a diagonal matrix of AR(1) coefficients, and $\eta_t \sim \mathcal{N}(0, \Phi)$. This specification of multivariate stochastic volatility, used in many previous studies, implies a reduced-form innovation variance-covariance matrix of $\Sigma_t = Q^{-1}\Lambda_t(Q^{-1})'$.

The upper triangular assumption in this reduced form specification is a deliberate choice. It enables tractable Bayesian estimation of the general shadow-rate VAR, while providing a way to handle the simultaneity between $x_t$ and $s_t$ present

[9] Specifically, for all $j$, the dimensions of $b_{xs}$, $\Pi_{xx,j}$, $\Pi_{xs,j}$, $\Pi_{xi,j}$, $\Pi_{sx,j}$, $\Pi_{ss,j}$, $\Pi_{si,j}$ are $N_x \times N_s$, $N_x \times N_x$, $N_x \times N_i$, $N_x \times N_i$, $N_i \times N_x$, $N_i \times N_i$, and $N_i \times N_i$, respectively. Please recall that $N_i = N_s$.

[10] Note that our $b_{xs}$ and $s_t$ correspond to the reduced-form parameter $\tilde{\beta}$ and the non-structural shadow-rate $\bar{Y}_{2t}^*$ in Mavroeidis (2021), respectively. Identification of reduced-form parameters is discussed in Section 3.1.1, p.2866, of his paper.

in (2).[11] As detailed below, in estimation we break the VAR system of (2) and (3) into separate Gibbs steps, that can be separated due to the upper triangular specification of $Q$, and which build on ideas from Carriero et al. (2022). We should stress that we do not view the triangular specification of $Q$ as representing the impact responses of structural shocks, but merely as useful orthogonalization of the VAR's residuals for reduced-form estimation and out-of-sample forecasting.

### 3.2 *Non-structural shadow-rate VAR*

We also consider a non-structural shadow-rate VAR that extends to the case of multiple rates. The non-structural shadow-rate VAR restricts the general model of equations (2)-(3) by imposing $b_{xs} = 0$ so that the regime-specific shadow-rate regressor drops out of the $x_t$ equation:

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

The non-structural shadow-rate VAR consists of equations (3) and (6), while allowing also for $N_i \geq 1$.

The additional restrictions in the non-structural model mean that interest rates do not have to be be placed last in the VAR ordering with an upper triangular specification of $Q$, although for comparability, we maintain that specification and ordering in our empirical implementation. In addition, the non-structural model includes the same specification of its innovations and stochastic volatility indicated above for the general model. This non-structural model, like the general shadow-rate VAR, relates all variables (i.e., both $x_t$ and $s_t$) to lags of both the shadow rate and actual interest rates; 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 reflecting some effects of unconventional monetary policies (such as forward guidance or

[11] By contrast, Mavroeidis (2021) employs a maximum-likelihood evaluation for a smaller scale application with the likelihood computed using a particle filter.

asset purchases), whereas actual rates are paid (earned) by borrowers (lenders) and thereby enter in economic dynamics and predictions.

This non-structural specification permits us to extend the set of interest rates included beyond the federal funds rate, to include multiple bond yields, with multiple rates (not necessarily all) potentially constrained by the ELB. Studies of multivariate time series models for forecasting also commonly include multiple interest rates (e.g., Chan (2021), Giannone et al. (2015), Gonzalez-Astudillo and Laforte (2024), and Johannsen and Mertens (2021)). The accuracy of macroeconomic forecasts may 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. (2024) develop a large VAR intended to be useful for a range of forecasting questions faced by a central bank and include several bond yields and financial indicators.

Relative to the general structural VAR of Mavroeidis (2021) (specifically, his "censored and kinked" ("CKSVAR") specification), the non-structural specification of equations (3) and (6) arises when certain blocks of the impact matrix of the underlying SVAR are zeroed out. In particular, our non-structural shadow-rate VAR zeros out the SVAR's contemporaneous responses of $x_t$ and $s_t$ to actual interest rates.[12]

While not capturing the full extent of changes in coefficients at the ELB allowed in the general model (as detailed in Mavroeidis (2021)), this non-structural 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

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

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

[12]Formally, in the notation of Mavroeidis (2021), this model imposes on his equation (20) that $A_{12} = 0$ and $A_{22} = 0$. From these restrictions, it follows that $\kappa = 1$ and $\tilde{\beta} = 0$ in his setup.

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

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

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

As this indicates, when interest rates are constrained at the ELB, the non-structural shadow-rate VAR includes parameter changes in the coefficients on $s_{t-j}$ 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 $\Pi_{xi,j}$ and $\Pi_{si,j}$, $j = 1, \ldots, p$) when at or above the ELB.

### 3.3 *Restricted version of the non-structural shadow-rate VAR*

For out-of-sample forecasting, in which parsimony is known to have important benefits, the inclusion of both actual interest rates and shadow rates as predictors in the general and non-structural shadow-rate VARs might not be fully conducive to forecast accuracy. Accordingly, we consider an additional specification that restricts which rates appear as predictors in the reduced-form VAR. Specifically, we consider a specification in which blocks of 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.[13] In terms of the non-structural shadow-rate VAR's equations (3) and(6), the restrictions are

[13]The block zero restrictions in this model mean that interest rates do not have to be be placed last in the VAR ordering with an upper triangular specification of $Q$, although for comparability, we maintain that specification and ordering in our model exposition and empirical implementation.

$\Pi_{xs,j} = 0$ and $\Pi_{si,j} = 0$, so that the restricted model takes the following form:

$$x_t = \sum_{j=1}^{p} \Pi_{xx,j}\, x_{t-j} + \sum_{j=1}^{p} \Pi_{xi,j}\, i_{t-j} + v_{x,t}\,, \tag{7}$$

$$s_t = \sum_{j=1}^{p} \Pi_{sx,j}\, x_{t-j} + \sum_{j=1}^{p} \Pi_{ss,j}\, s_{t-j} + v_{s,t}. \tag{8}$$

Of course, this model is nested within the non-structural shadow-rate VAR presented above, and includes the same specification of its innovations and stochastic volatility. In its predictive relationships, the restricted 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 restricted system of equations (7)-(8), the censored values of (lagged) actual interest rates are state variables that influence the evolution of macroeconomic variables. From a forecasting perspective, such a specification could be seen as advantageous since the decisions of households and firms are most directly connected to 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 some of our empirical applications. 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.

Relative to the structural VARs of Mavroeidis (2021) (specifically, his CKSVAR specification), our restricted version of the non-structural shadow-rate VAR combines elements of his "censored" and "kinked" SVARs: For macroeconomic variables, the restricted model borrows from his kinked specification, whereas for shadow rate equations it follows his purely censored model. In the interest of brevity, we omit details of the relevant restrictions (in his notation, on certain elements of the $C$ coefficient matrices that are products of $A$ and $B$ terms of the SVAR). This restricted model, like the non-structural shadow-rate VAR described in Section 3.2, satisfies the coherency and completeness conditions for a unique reduced form solution.

Of course, this restricted shadow-rate VAR treats economic dynamics as unchanged (apart from interest rate censoring) in the face of the ELB. While a departure from the emphasis of structural VAR work in studies such as Aruoba et al. (2022) and Mavroeidis (2021), the constant-parameter specification of the restricted shadow-rate VAR can be seen as consistent with some 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). In the context of DSGE models, 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.[14]

### 3.4 *Estimation and forecasting*

All of our models are estimated with an MCMC sampler that extends the methods of Carriero et al. (2019) and Carriero et al. (2022), henceforth "CCCM," for estimation of large BVAR-SV models to handle the ELB. The MCMC sampler is fairly standard, except for a step that draws values for shadow rates when actual rates are at the ELB.

[14] Wu and Zhang (2019) also provide references to empirical work that has concluded that conventional and unconventional monetary policies work in a similar fashion.

We collect all *unobserved* shadow rates in a vector $\boldsymbol{S}$ and all observations of $y_t = [x_t' \;\; i_t']'$ in a vector $\boldsymbol{Y}$. For ease of reference, suppose the ELB binds for all elements of $s_t$ at $t = t^*, t^*+1, \dots, T^*-1, T^*$, and we have:

$$\boldsymbol{S} = \begin{bmatrix} s_{T^*} \\ s_{T^*-1} \\ \vdots \\ s_{t^*+1} \\ s_{t^*} \end{bmatrix}, \quad \text{and} \quad \boldsymbol{Y} = \begin{bmatrix} y_T \\ y_{T-1} \\ \vdots \\ y_0 \\ y_{-p+1} \end{bmatrix}. \tag{9}$$

If the data contains more than one ELB episode, $\boldsymbol{S}$ contains non-consecutive observations since values above the ELB are excluded from $\boldsymbol{S}$. Similarly, in the case of multiple interest rates, so that $s_t$ is a vector, elements of $s_t$ with values above $ELB$ are excluded from $\boldsymbol{S}$.

Below we describe an MCMC algorithm for shadow-rate VAR models that features a shadow-rate sampling step that draws from $f(\boldsymbol{S} \mid \boldsymbol{Y}, \boldsymbol{\Pi}, \boldsymbol{\Sigma}; \boldsymbol{S} \leq ELB)$, where the inequality is element-wise, and $\boldsymbol{\Pi}$ and $\boldsymbol{\Sigma}$ refer to given values of the VAR parameters $\{\Pi_{\cdot,\cdot,j}\}_{j=1}^{p}$ (and in case of the general model also $b_{xs}$), and the volatility matrices $\{\Sigma_t\}_{t=1}^{T}$, respectively.[15] Importantly, shadow-rate sampling conditions on the restriction that $\boldsymbol{S} \leq ELB$, which reflects the known timing of when the ELB binds for the interest rates contained in the data vector $\boldsymbol{Y}$.

As in Johannsen and Mertens (2021), the shadow-rate sampler builds on solving a "missing value" problem, which ignores the restriction $\boldsymbol{S} \leq ELB$, and is thus characterized by the density $f(\boldsymbol{S} \mid \boldsymbol{Y}, \boldsymbol{\Pi}, \boldsymbol{\Sigma})$. In light of the conditionally linear and Gaussian structure of the model, the missing-value results in a multivariate normal posterior, and truncation at the ELB leads to the following posterior for the

[15]For sake of exposition, we let $\boldsymbol{\Sigma}$ not only include $Q$ and $\{\Lambda_t\}_{t=1}^{T}$ (which make up the volatility matrices $\Sigma_t = Q^{-1}\Lambda_t(Q^{-1})'$), but also the parameter values for the SV processes, $\gamma_0$, $\gamma_1$, and $\Phi$.

shadow rates:[16]

$$\boldsymbol{S} \,|\, (\boldsymbol{Y}, \boldsymbol{\Pi}, \boldsymbol{\Sigma}) \qquad \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Omega}) \tag{10}$$

$$\Rightarrow \quad \boldsymbol{S} \,|\, (\boldsymbol{Y}, \boldsymbol{\Pi}, \boldsymbol{\Sigma};\, \boldsymbol{S} \leq ELB) \sim \mathcal{TN}(\boldsymbol{\mu}, \boldsymbol{\Omega}, -\infty, ELB)\,. \tag{11}$$

Our MCMC sampler is summarized in Algorithm 1, where the $m$th MCMC draws of shadow rates, VAR parameters, and volatility matrices are denoted by $\boldsymbol{S}^{(m)}$, $\boldsymbol{\Pi}^{(m)}$, and $\boldsymbol{\Sigma}^{(m)}$, respectively, and $\boldsymbol{\Pi}$ collects the VAR's coefficients $\boldsymbol{\Pi}$:

MCMC ALGORITHM 1 (Shadow-rate VAR estimation). *Given initial values for $\boldsymbol{\Pi}^{(0)}$ and $\boldsymbol{\Sigma}^{(0)}$, iterate over the following blocks for $m = 1, 2, \ldots, M$:*

1. *Draw shadow rates from $f(\boldsymbol{S}^{(m)}|\boldsymbol{\Pi}^{(m-1)}, \boldsymbol{\Sigma}^{(m-1)}, \boldsymbol{Y};\, \boldsymbol{S} \leq ELB)$ as discussed above, and further detailed in the supplementary online appendix.*
2. *Draw VAR coefficients from $f(\boldsymbol{\Pi}^{(m)}|\boldsymbol{S}^{(m)}, \boldsymbol{\Sigma}^{(m-1)}, \boldsymbol{Y})$ following CCCM.*
3. *Draw stochastic volatility parameters from $f(\boldsymbol{\Sigma}^{(m)}|\boldsymbol{S}^{(m)}, \boldsymbol{\Pi}^{(m)}, \boldsymbol{Y})$ as in Cogley and Sargent (2005) and Kim et al. (1998).*

Block 1 of the MCMC sampler in Algorithm 1 requires drawing from a truncated multivariate normal distribution, for which no direct method exists. We combine a sequential Gibbs sampler, for steps $m$ that are early in the MCMC routine, with rejection sampling for steps $m$ late in the chain. Following Johannsen and Mertens (2021), rejection sampling is straightforward to do based on the missing value problem in (10), retaining only draws that adhere to $\boldsymbol{S} \leq ELB$. However, the inequality restriction pertains to the entire trajectory of unobserved shadow rates, $\boldsymbol{S}$, which can lead to low acceptance rates. In our experience, the acceptance probability in rejection sampling $\boldsymbol{S} \leq ELB$ from (10) can be quite low , but mostly during the early stages of the MCMC sampler that is described above. Once the MCMC sampler has passed its burnin state, and draws for $\boldsymbol{\Pi}$ and $\boldsymbol{\Sigma}$ are

[16]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.

converging towards their eventual posteriors, rejection sampling turns out to be computationally efficient, in particular when implemented via a precision-based approach as in Chan et al. (2023) or Mertens (2023). For the burnin period of the MCMC sampler (and when rejection sampling gets stuck post burnin), we employ instead a sequential Gibbs sampling approach for sampling from the truncated shadow-rate distribution in equation (11). This Gibbs sampler efficiently adapts the methods of Geweke (1991) and Park et al. (2007) to the variance-covariance structure of the VAR($p$) case. The computations for the Gibbs sampler solve a sequential Kalman filtering problem. As discussed by, among others, Durbin and Koopman (2012), Kalman filtering calculations are susceptible to non-positive-definite results for second-moment matrices that arise from round off errors. We circumvent these issues by using the array methods of Kailath et al. (2000).

In Block 2 of the MCMC algorithm, we build on CCCM and break down estimation of the VAR system into a sequence of separate Gibbs steps, one for each equation's coefficients. The Gibbs steps are enabled by the triangular factorization of the VAR residuals' variance covariance matrices, $\Sigma_t$, as illustrated in equation (5) above. For estimation of the general model in equations (2) and (3), it is relevant that an upper-triangular $Q$ in equation (5) (and a VAR vector that orders shadow rates last) leads the CCCM approach to take the residuals of the shadow-rate equation as given when estimating the coefficients of equations for $x_t$. This feature enables the CCCM approach to estimate the responses of $x_t$ to contemporaneous shadow rates, $b_{xs}$, as part of a standard draw for the regression coefficients in equation (2). For model variants other than the general model, $Q$ may also be taken to be upper triangular. As in CCCM, we use a Minnesota prior for the VAR coefficients and follow their other choices for priors as far as applicable.[17]

---

[17] All VAR coefficients in $\mathbf{\Pi}$ 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. (The implications of unit roots for trends differ in our non-linear specifications as compared to linear models for which the Minnesota prior was designed (see, e.g., Duffy et al. (2023))). As usual, different degrees of shrinkage are applied to own- and cross-lag coefficients. Prior variances of the $j$th-order own lag are set to $\theta_1/j^{\theta_4}$. The cross-lag of the coefficient on variable $m$ in equation $n$ has prior variance equal to

For the purpose of out-of-sample forecasting, we generate from every model 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 $k = 1, 2 \ldots, h$. For the shadow-rate VARs, 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 (2) and (3) (or their more restricted counterparts). 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 $y_{t+k}$.

## 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 (McCracken and Ng, 2016). Reflecting the raw sample, transformations to growth rates for some variable, and lag specification, the sample for model estimation always begins with 1960:04. We first describe results from models containing 15 variables, with the funds rate as the only interest rate, which was constrained by the ELB from late 2008 through late 2015 and from March 2020 through February 2022. We also report some results for the non-structural and restricted shadow-rate VARs containing 20 variables, including 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. Table 1 lists the variable set and transformations.

---

$\theta_1/j^{\theta_4} \cdot \theta_2 \cdot \hat{\sigma}_n^2/\hat{\sigma}_m^2$. In shadow-rate equations that feature lags of actual and shadow rates as regressors, the shadow-rate lags are considered as own lags, while actual-rate lags are treated as cross-lag coefficients. Similarly, the prior for $b_{xs}$ in the general shadow-rate VAR is identical to the prior for a first-order cross lag. 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 $\theta_4 = 2$.

TABLE 1. List of variables

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

In our application with monthly data, we use $p = 12$ lags. 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 in previous ELB episodes. We treat a given interest rate — at any maturity — as unconstrained unless it reaches the value of $ELB$. 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.[18] Our working paper (Carriero et al., 2023) 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 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. 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.[19]

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

[19]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.

## 5. Forecast Evaluation

This section presents our empirical evaluation of the shadow-rate VARs for macroeconomic forecasting, beginning with an evaluation of forecast accuracy and proceeding with a closer look at historical estimates of the shadow rate estimates and interest rate forecasts from our models.

### 5.1 *Average performance 2010–2017*

We conduct an out-of-sample forecast evaluation in quasi-real time, where we simulate forecasts made from January 2010 through December 2017. For every forecast origin, each model is re-estimated based on growing samples of data that start in 1959:03. We begin the forecast evaluation in 2010 in order to concentrate on a sample in which the ELB had already been binding for about a year and continued to do so for most of the period.[20] 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. All data are taken from the September 2022 vintage of FRED-MD; we abstract from issues related to real-time data collection.

To evaluate our shadow rate models, we compare their forecast accuracy to that of a benchmark linear VAR (also with stochastic volatility) in $y_t$ that has no shadow rate treatment of the ELB. 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

[20]With the funds rate hitting the ELB in December 2008, our general and (unrestricted) non-structural specifications are challenged to forecast well in 2009. The restricted non-structural shadow-rate VAR shows better forecast accuracy in 2009.

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. However, for sake of comparability in assessments of interest rate projections, forecasts from linear and shadow-rate VARs are compared against realized interest-rate values that are censored at the ELB.

Comparing various model specifications discussed below, Tables 2 and 3 provide results on point and density forecast accuracy, measured by mean absolute error (MAE, computed around median forecasts) and continuous ranked probability score (CRPS), respectively. For point forecasts, we provide the MAE results rather than root mean squared error estimates (computed around mean forecasts) 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), the $h$-step forecasts are transformed to average growth rates over $h$ periods.

We begin with forecasts from models in which the federal funds rate is the only interest rate (which allows us to consider the general shadow-rate VAR). To facilitate comparisons, we report MAE and CRPS results for the general and non-structural shadow-rate models as relative to the standard VAR that ignores the ELB; 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 (at 10 percent or better) with asterisks.

Overall, the results in Table 2 indicate that, in models with the funds rate as the only included interest rate, the general shadow-rate VAR specification for accommodating the ELB performs better in forecasting — both point and density — than does a standard VAR. For federal funds rate forecasts, the MAE (CRPS) ratios for $h$ = 6, 12, and 24 range from 0.51 to 0.63 (0.46 to 0.60), nearly all significant. For most of the indicators of economic activity, stock price returns, and the exchange

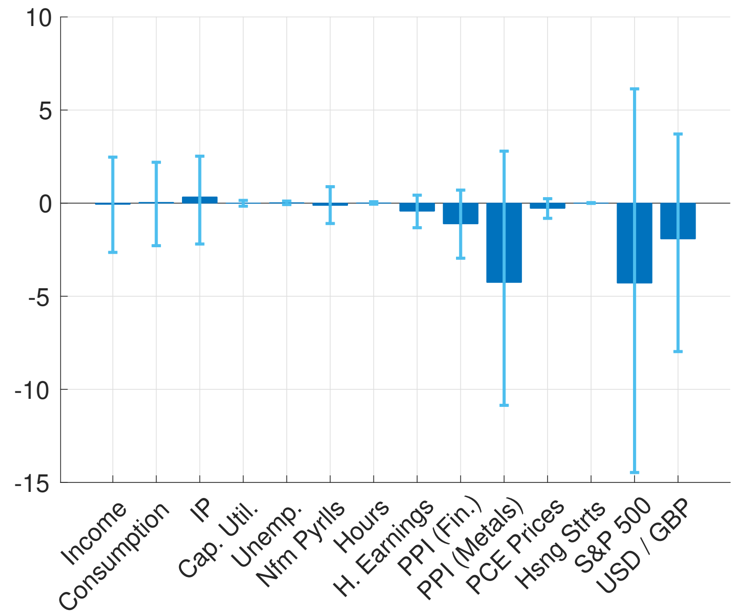


FIGURE 1. Estimates of ELB-specific coefficient vector $b_{xs}$. Posterior medians and posterior 90% uncertainty of estimates obtained using the full data sample through 2022:08.

rate, the treatment of the ELB on interest rates does not seem to bear consistently and importantly on forecast accuracy. In most cases, the MAE and CRPS ratios for the general model relative to the linear VAR are close to 1. However, there are exceptions in both directions: The general model modestly and consistently improves forecasts of unemployment and housing starts. On the other hand, the linear VAR is more accurate for PCE inflation and hourly earnings.

Our non-structural shadow-rate VAR also performs better in forecasting than does a standard VAR. While imposing some restrictions relative to the general shadow-rate VAR specification, the non-structural VAR effectively matches its forecast accuracy. MAE and CRPS ratios relative to the linear VAR baseline are broadly similar for the non-structural and general shadow-rate VARs. In forecast accuracy, the restrictions imposed with the non-structural model have little cost or benefit. Relative to the baseline, this model continues to improve the accuracy of forecasts for the federal funds rate and some economic indicators (e.g., PCE inflation) while generally matching the baseline accuracy for most other variables.

TABLE 2. Forecast performance of shadow-rate VARs with single interest rate

| | MAE | | | | | | CRPS | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | General | | | Non-structural | | | General | | | Non-structural | | |
| | 6 | 12 | 24 | 6 | 12 | 24 | 6 | 12 | 24 | 6 | 12 | 24 |
| FX $/£ | 0.99 | 0.99* | 0.99 | 0.99 | 0.98* | 1.00 | 0.99 | 0.99* | 1.00 | 0.99 | 0.99* | 1.00 |
| S&P500 | 1.00 | 1.00 | 1.00 | 1.00 | 0.99 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 |
| Hsng St | 0.97 | **0.89** | **0.77*** | 0.99 | 0.99 | 0.99 | 0.98 | **0.94** | **0.89*** | 0.99 | 0.99 | 1.02 |
| PCE Inf | **1.07*** | **1.11*** | **1.13*** | 0.97 | 0.96 | 0.96 | **1.06*** | **1.09*** | **1.11*** | 0.97 | **0.94** | **0.94** |
| PPI Met | 0.99 | 1.00 | 1.01 | 1.00 | 1.00 | 1.00 | 0.99 | 1.00 | 1.00 | 1.00 | 0.99 | 1.00 |
| PPI Fin | 1.01 | 1.02 | 1.02 | 1.00 | 0.98 | 0.97 | 1.01 | 1.01 | 1.02 | 1.00 | 0.97 | 0.97 |
| H. Earn | 1.03* | **1.06*** | **1.08*** | 0.99 | 0.99 | 0.99 | 1.03* | 1.04* | **1.08*** | 0.99 | 1.00 | 0.99 |
| Hours | 1.01 | 1.04 | **1.05** | 1.02 | 1.03 | **1.06** | 1.01 | 1.02 | 1.04 | 1.02 | 1.03 | **1.05** |
| Nfm Pyr | **1.05*** | 0.99 | 0.98 | 1.03 | 1.00 | 1.00 | 1.03* | 1.00 | 1.01 | 1.03 | 1.02 | 1.01 |
| Unemp | **0.95** | 0.96 | **0.93** | 0.98 | 0.98 | 0.97 | 0.97 | 0.96 | **0.93** | 1.00 | 0.99 | 0.98 |
| CapUtil | 1.00 | 0.96 | **0.95** | 1.03 | 0.98 | 0.96 | 1.00 | 0.97 | 0.96 | 1.01 | 0.98 | 0.96 |
| IP | 1.04* | 1.04 | 1.04 | 1.01 | 0.99 | 0.99 | 1.04* | 1.03* | 1.03 | 1.02* | 1.00 | 1.00 |
| Cons | 1.01 | 1.00 | 1.00 | 0.99 | 0.97* | 0.97* | 1.00 | 1.00 | 1.00 | 1.00 | 0.99 | 0.99 |
| Income | 1.01* | 1.01 | 1.01 | 1.01* | 1.01 | 1.00 | 1.01* | 1.01 | 1.00 | 1.01 | 1.01 | 1.00 |
| FFR | **0.51*** | **0.63** | **0.57*** | **0.44*** | **0.60** | **0.55*** | **0.46*** | **0.57*** | **0.60*** | **0.39*** | **0.51*** | **0.56*** |

*Note*: Comparison of General shadow-rate VAR and Non-structural shadow-rate VAR against standard linear VAR (baseline, in denominator) for horizons 6, 12, and 24. Values below 1 indicate improvement over baseline. Evaluation window with forecast origins from 2010:01 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, and stars indicating $p$ values of 10% and below. Relative differences of 5 percent and more (compared to baseline) are indicated by bold face numbers.

The fact that the non-structural shadow-rate VAR forecasts about as well as the general model suggests that the contemporaneous shadow-rate terms included in the general model's equations for macroeconomic variables may not have much predictive content. As a simple check, Figure 1 reports estimates of the coefficients of the vector $b_{xs}$ from the general model estimated with data through 2022:08. For most variables, the posterior medians of the coefficients are small and imprecisely estimated. For a few variables — PPI inflation, S&P500 returns, and growth in the exchange rate — the posterior medians are more sizable but still imprecisely estimated. At least in our model and sample, the ELB-specific

TABLE 3. Forecast performance of shadow-rate VARs with multiple interest rates

| | MAE | | | | | | CRPS | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | Non-structural | | | Restricted | | | Non-structural | | | Restricted | | |
| | 6 | 12 | 24 | 6 | 12 | 24 | 6 | 12 | 24 | 6 | 12 | 24 |
| FX $/£ | 0.99 | 1.00 | **1.05** | 1.00 | 1.01 | 1.01 | 1.00 | 1.01 | **1.17**$^{*}$ | 1.00 | 1.01 | 1.01 |
| S&P500 | 0.99 | 1.01 | **1.08** | 1.01 | 1.01 | 1.02 | 1.00 | 1.03 | **1.26**$^{*}$ | 1.01$^{*}$ | 1.01 | 1.02$^{*}$ |
| Hsng St | 1.01 | 1.04 | 1.00 | 1.00 | 1.00 | 1.03 | 1.01 | 1.02 | **1.12** | 1.01 | 1.00 | 1.00 |
| PCE Inf | 1.02 | 1.01 | **1.08** | 1.00 | 1.03 | **1.08**$^{*}$ | 1.01 | 1.01 | **1.20** | 1.00 | 1.01 | **1.06**$^{*}$ |
| PPI Met | 1.00 | 1.03 | **1.08** | 1.00 | 1.01 | 1.01 | 1.00 | 1.04$^{*}$ | **1.24**$^{*}$ | 1.00 | 1.00 | 1.01$^{*}$ |
| PPI Fin | 1.00 | **1.05** | **1.14** | 1.00 | 1.01 | 1.00 | 0.99 | 1.03 | **1.26** | 1.00 | 1.00 | 0.99 |
| H. Earn | 1.01 | 1.02 | **1.13** | 1.00 | 1.02$^{*}$ | 1.03$^{*}$ | 1.02 | **1.05**$^{*}$ | **1.28**$^{*}$ | 1.01 | 1.01$^{*}$ | 1.02$^{*}$ |
| Hours | **1.05** | 1.04 | **1.12** | 0.99 | 1.00 | 1.01 | **1.05** | **1.08** | **1.29**$^{*}$ | 0.99 | 1.00 | 1.01 |
| Nfm Pyr | **1.05** | **1.09** | **1.12** | 1.00 | **0.95**$^{*}$ | **0.85**$^{*}$ | **1.05** | **1.12**$^{*}$ | **1.35**$^{*}$ | 0.99 | 0.97$^{*}$ | **0.91**$^{*}$ |
| Unemp | 0.98 | 1.03 | 1.00 | 0.99 | 0.97 | **0.86**$^{*}$ | 1.01 | **1.07** | **1.13** | 0.99 | 0.98 | **0.90**$^{*}$ |
| CapUtil | 1.03 | 1.04 | 1.01 | 1.01 | 0.98 | **0.91**$^{*}$ | 1.04$^{*}$ | **1.09** | **1.19**$^{*}$ | 1.00 | 0.99 | **0.95** |
| IP | 0.96 | 1.04 | **1.33** | 1.01 | 0.99 | 1.02 | 0.98 | **1.06** | **1.49** | 1.01 | 0.99 | 1.03 |
| Cons | **1.05** | 1.04 | **1.17** | 1.00 | 0.97$^{*}$ | **0.94** | 1.03$^{*}$ | **1.06**$^{*}$ | **1.27**$^{*}$ | 1.00 | 0.99$^{*}$ | 0.97$^{*}$ |
| Income | **1.07** | **1.13** | **1.31** | 1.01$^{*}$ | 1.00 | 0.97 | **1.05** | **1.12** | **1.36** | 1.01$^{*}$ | 1.00 | 0.98 |
| BAA | **1.07**$^{*}$ | **1.13**$^{*}$ | **1.08** | 0.99 | 1.03 | 1.00 | **1.07**$^{*}$ | **1.14**$^{*}$ | **1.25**$^{*}$ | 1.02 | 1.03 | **1.06** |
| 10y Tsy | 1.02 | 1.01 | **1.09** | **0.94** | **0.85**$^{*}$ | **0.83**$^{*}$ | 1.03 | 1.03 | **1.14**$^{*}$ | **0.93** | **0.89** | **0.92** |
| 5y Tsy | 0.96 | 0.96 | 0.99 | **0.89**$^{*}$ | **0.86**$^{*}$ | **0.83**$^{*}$ | 1.00 | 0.99 | 1.02 | **0.91** | **0.86**$^{*}$ | **0.85**$^{*}$ |
| 1y Tsy | **0.74**$^{*}$ | **0.81** | **0.77**$^{*}$ | **0.61**$^{*}$ | **0.69**$^{*}$ | **0.79**$^{*}$ | **0.77**$^{*}$ | **0.82** | **0.77**$^{*}$ | **0.67**$^{*}$ | **0.73**$^{*}$ | **0.73**$^{*}$ |
| 6m Tsy | **0.58**$^{*}$ | **0.65**$^{*}$ | **0.70**$^{*}$ | **0.53**$^{*}$ | **0.58**$^{*}$ | **0.71**$^{*}$ | **0.58**$^{*}$ | **0.67**$^{*}$ | **0.70**$^{*}$ | **0.51**$^{*}$ | **0.62**$^{*}$ | **0.68**$^{*}$ |
| FFR | **0.47**$^{*}$ | **0.60** | **0.61**$^{*}$ | **0.40**$^{*}$ | **0.55**$^{*}$ | **0.60**$^{*}$ | **0.40**$^{*}$ | **0.49**$^{*}$ | **0.60**$^{*}$ | **0.35**$^{*}$ | **0.43**$^{*}$ | **0.55**$^{*}$ |

*Note*: Comparison of Non-structural shadow-rate VAR and Restricted non-structural shadow-rate VAR against standard linear VAR (baseline, in denominator) for horizons 6, 12, and 24. Values below 1 indicate improvement over baseline. Evaluation window with forecast origins from 2010:01 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, and stars indicating $p$ values of 10% and below. Relative differences of 5 percent and more (compared to baseline) are indicated by bold face numbers.

coefficient estimates and their precision appear to be too limited to make this term helpful to out-of-sample forecast accuracy.

Turning to forecasts from models including some additional interest rates, Table 3 reports MAE and CRPS results for the non-structural shadow-rate model and its restricted version relative to a standard VAR that ignores the ELB. Starting with the non-structural specification, it significantly improves forecasts of not only the federal funds rate (FFR) but also some other, shorter-maturity interest rates. More specifically, in interest rate results, MAE (CRPS) ratios for federal funds rate forecasts for $h = 6$, 12, and 24 range from 0.47 to 0.61 (0.40 to 0.60). The model also offers sizable gains in forecasts of the 6-month T-bill rate (with continued statistical significance in this case) and smaller gains at longer maturities (statistically significant at only the 1-year maturity). In the case of indicators of economic activity, measures of inflation, and other financial indicators, the model's forecast performance is more mixed, often broadly comparable to the accuracy of the linear VAR, except for a noticeable deterioration in accuracy at the $h = 24$ horizon, sharper for CRPS than MAE. For macroeconomic indicators, MAE and CRPS ratios for the non-structural shadow-rate VAR as compared to the linear VAR are sometimes close to 1 (e.g., PCE and PPI inflation) and other times modestly above 1 (e.g., nonfarm payrolls). However, reflecting the effects of uncertainty about its ELB-specific coefficients compounding at longer horizons, density forecasts from the non-structural shadow-rate specification fall short of the linear VAR's accuracy at $h = 24$, with CRPS ratios consistently at about 1.3.

Our restricted version of the non-structural shadow-rate VAR improves on the overall accuracy of the non-structural specification and, in turn, the linear VAR. For the federal funds rate, MAE (CRPS) ratios for this shadow-rate VAR range from 0.40 to 0.60 (0.35 to 0.55) for $h = 6$, 12, and 24, similar to the ratios for the non-structural model. However, unlike the non-structural model, the restricted specification also significantly improves the accuracy of forecasts of most other interest rates (e.g., 5-year Treasury yields at $h = 12$ and 24). For other variables, the restricted specification broadly matches the MAE and CRPS accuracy of the linear VAR, with ratios sometimes slightly to modestly lower than those seen for the non-structural shadow-rate VAR and other times a little higher. For example, for nonfarm payrolls, MAE ratios for $h = 6$ through 24 range from 0.85 to 1.05 for the restricted model, compared to 1.05 to 1.12 for the non-structural specification. Comparing the two shadow rate models covered in the table, the superiority of the restricted model is sharpest at the $h = 24$ horizon. These ratios for non-interest rate variables indicate that the restricted model is overall very similar to the linear VAR, at all horizons, whereas ratios for the non-structural speci-

TABLE 4. Shadow-rate VARs compared against a linear VAR w/o short-term interest rates

| | MAE | | | | | | CRPS | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | w/o yields | | | w/yields | | | w/o yields | | | w/yields | | |
| | 6 | 12 | 24 | 6 | 12 | 24 | 6 | 12 | 24 | 6 | 12 | 24 |
| FX $/£ | 1.00 | 1.00 | 1.01 | 1.02 | 1.02 | 1.04 | 1.00 | 1.01 | 1.02 | 1.01 | 1.02 | 1.02 |
| S&P500 | 1.00 | 0.97$^{*}$ | 0.99$^{*}$ | 0.99 | 0.99 | 1.03 | 1.00 | 0.99 | 1.00 | 1.00 | 1.00 | 1.02$^{*}$ |
| Hsng St | **0.92**$^{*}$ | **0.85**$^{*}$ | **0.82** | 1.04 | **1.15** | **1.55**$^{*}$ | **0.93** | **0.88**$^{*}$ | **0.89** | 1.03 | **1.14** | **1.44**$^{*}$ |
| PCE Inf | **0.95** | **0.91** | **0.88** | **0.94** | **0.95** | 0.97 | **0.95** | **0.91** | **0.88** | **0.94** | **0.93** | 0.97 |
| PPI Met | 1.00 | 1.00 | 1.00 | 1.00 | 1.01 | 1.00 | 1.00 | 1.00 | 1.00 | 1.00 | 1.01 | 1.01 |
| PPI Fin | 1.01 | 0.99 | 0.97 | 0.98 | 0.99 | 0.97 | 1.01 | 0.99 | 0.96 | 0.97 | 0.97 | 0.97 |
| H. Earn | 1.03 | 1.03 | 1.01 | 1.00 | 1.00 | 0.98 | 1.01 | 1.01 | 1.00 | 1.01 | 1.01 | 0.99 |
| Hours | 1.01 | **0.93** | **0.86** | 1.00 | 0.96 | 0.97 | 1.04 | 0.98 | **0.88** | 1.01 | 0.98 | 0.97 |
| Nfm Pyr | 1.01 | **0.92** | **0.90** | **0.93** | **0.82**$^{*}$ | **0.83** | 1.00 | **0.95** | **0.91** | **0.94**$^{*}$ | **0.88**$^{*}$ | **0.87** |
| Unemp | **1.05** | 0.98 | **0.92** | 1.04 | 0.99 | 0.98 | 1.02 | 1.01 | **0.92** | 1.01 | 0.99 | 0.96 |
| CapUtil | 1.04 | **0.93** | **0.84**$^{*}$ | **0.79**$^{*}$ | **0.74**$^{*}$ | **0.72**$^{*}$ | 1.04 | **0.95** | **0.84**$^{*}$ | **0.86**$^{*}$ | **0.80**$^{*}$ | **0.77**$^{*}$ |
| IP | 0.98 | **0.95**$^{*}$ | **0.92** | 0.96 | **0.91**$^{*}$ | **0.91** | 0.99 | **0.95**$^{*}$ | **0.93** | 0.97 | **0.93**$^{*}$ | **0.93** |
| Cons | 1.01 | 1.00 | 0.97 | **0.94** | **0.91** | 0.97 | 1.00 | 0.99 | 0.98 | 0.99 | 0.96 | 0.98 |
| Income | 1.01 | 0.99 | 0.99$^{*}$ | 1.01 | 1.00 | 1.01 | 1.01 | 1.00 | 0.99$^{*}$ | 1.00 | 0.99 | 1.00 |
| BAA | −.− | −.− | −.− | 0.96 | **0.94** | **0.92** | −.− | −.− | −.− | 0.97 | **0.95** | **0.95** |
| 10y Tsy | −.− | −.− | −.− | 0.97 | 0.98 | 0.99 | −.− | −.− | −.− | 0.99 | 1.00 | 1.04 |
| 5y Tsy | −.− | −.− | −.− | 1.04 | **1.06** | 1.02 | −.− | −.− | −.− | 1.02 | 1.01 | **0.95** |

*Note*: Comparison of "Non-structural shadow-rate VAR (w/o yields)" and "Restricted non-structural shadow-rate VAR (w/yields)" against "Linear VAR without short-term yields' (baseline, in denominator) for horizons 6, 12, and 24. Values below 1 indicate improvement over baseline. Evaluation window with forecast origins from 2010:01 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, and stars indicating $p$ values of 10% and below. Relative differences of 5 percent and more (compared to baseline) are indicated by bold face numbers. In some cases, due to strong performance of the baseline model, relative MAE may involve divisions by zero. These cases are reported as blank entries.

fication show more advantage to the linear model, especially at longer horizons. From these patterns, it appears that, in forecasting with reduced-form shadow-rate VARs, there is some benefit to distinguishing between effects from actual and shadow rates as predictors for economic outcomes (i.e., to imposing zero restrictions on $\Pi_{xs,j}$ and $\Pi_{si,j}$, $j = 1, \ldots, p$).

As noted above, to mitigate ELB constraints some studies have omitted short-term interest rates from VARs. Table 4 compares the accuracy of forecasts from shadow-rate VARs that include short-term interest rates to forecasts from a linear BVAR that excludes short-term interest rates (specifically, the rates that hit the ELB in our sample: the federal funds rate and Treasury yields at 6-months and 1-year maturities) but includes longer-term yields. We begin with the non-structural shadow rate model in which the federal funds rate is the only interest rate (results are similar for the general model). At the $h$=6 forecast horizon, the MAE and CRPS ratios indicate that, for most variables, point and density forecasts from this model match or slightly the accuracy of corresponding forecasts from the linear VAR omitting short-term interest rates. The advantage of the shadow-rate model increases as the horizon lengthens, achieving gains as large as 18 percent. With additional interest rates included in the restricted version of the shadow-rate VAR, the MAE and CRPS ratios of Table 4 indicate that, in a few cases (e.g., housing starts) the linear VAR yields better forecasts than the shadow-rate specification. However, for most variables, the shadow-rate VAR matches or exceeds the forecast accuracy of the linear VAR that excludes short-term rates. The gains are more consistent across horizons than in the case of the shadow-rate VAR without longer-term interest rates, but of comparable magnitudes. Collectively, in our data, for out-of-sample forecasting the shadow-rate "solution" to the ELB is better than the alternative of simply excluding short-term interest rates.

Our supplementary online appendix and an earlier working paper version of this paper (Carriero et al., 2023) report various robustness checks, with similar results to those reported here. In particular, when extending the evaluation period to end in August 2022, it turns out that while the economic effects of the pandemic left a heavy mark on readings of macroeconomic and financial variables in 2020 and 2021 — see, for example, our companion work in Carriero et al. (2024) — they did not materially affect the relative comparisons reported here.

### 5.2 *Shadow-Rate Estimates*

The forecasting performance of the shadow-rate VARs of course reflects the underlying reduced-form estimates of the shadow rate. This section presents full-sample estimates of the shadow rate from a few of the specifications included in the out-of-sample forecast evaluation. Of course, it is well known from the term structure literature that shadow rate estimates are often sensitive to model specification (see, e.g., Christensen and Rudebusch (2015)). Our reported reduced-form estimates of shadow rates also show some sensitivity, to the model and whether the model includes multiple interest rates or just the federal funds rate.

Figure 2 reports our reduced-form shadow rate estimates (posterior medians and 90 percent credible sets) associated with the federal funds rate, along with comparisons to some other estimates.[21] Panel (a) shows smoothed, full-sample estimates obtained from models with the federal funds rate as the only interest rate: (1) the general shadow-rate VAR and (2) the non-structural shadow-rate VAR, which were shown above to have comparable forecast accuracy. Panel (b) provides corresponding estimates for the restricted version of the non-structural shadow-rate VAR that includes multiple interest rates, along with (for comparison) the shadow-rate measures from Krippner (2015) and Wu and Xia (2016) based on affine term structure models. For historical perspective, the chart samples begin in January 2006. For 2006 through November 2008, with the federal funds rate above the ELB, the shadow rate from our models is simply the observed funds rate.

As indicated in Panel (a), when the only interest rate included in the model is the federal funds rate, the general shadow-rate VAR generates a shadow rate estimate (dashed red line) that fell to just about -70 basis points in 2009. Thereafter, the shadow rate very slowly edged up over time, before finally rising above the ELB by 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

[21]To be clear, in the full sample case, the model is estimated with data for 1960:04 through 2022:08, but the figure omits most of the period of 1960-2008 during which the ELB did not bind. While not reported in the interest of brevity, quasi-real-time estimates have more time variability than do the full sample estimates, but follow a quite similar contour. As might also be expected, the quasi-real-time estimates are less precise than the full sample estimates, with credible sets wider than those of the full sample estimate. However, the inclusion of bond yields improves the precision of the shadow rate estimate in quasi-real time and reduces its variability over time.

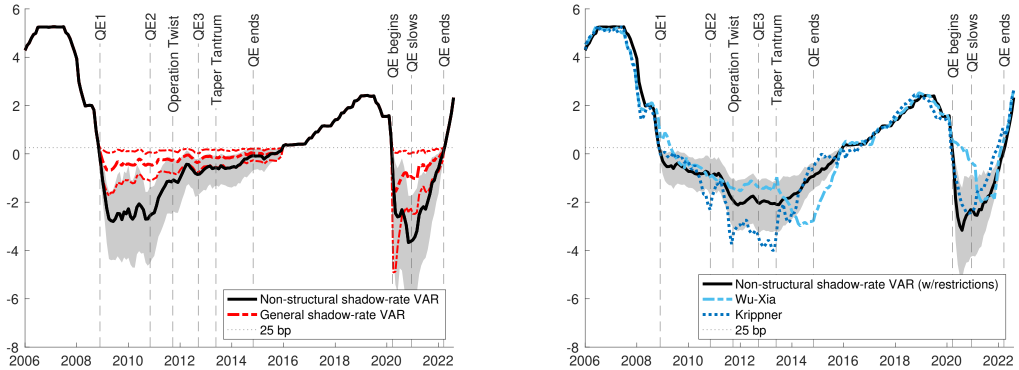


(a) Without yields

(b) With yields

FIGURE 2. Shadow-rate estimates (posterior medians and 90% bands). Panel (a) reports estimates generated with models including the federal funds rate as the only interest rate: (1) the general shadow-rate VAR and (2) the non-structural shadow-rate VAR. Panel (b) reports corresponding estimates from the restricted version of the non-structural shadow-rate VAR estimated using all variables, including longer-term yields. Panel (b) also includes current estimates of the shadow rates of Krippner (2015) and Wu and Xia (2016).

0-25 basis points to 25-50 basis points). During the first year of the COVID-19 pandemic, this shadow rate from the general model declined very quickly, briefly hitting a nadir of about -1.75 percent, before rising back up to about -1 percent and gradually climbing further to the ELB.

By comparison, as shown in the same panel, the estimate of the shadow rate from the non-structural shadow-rate VAR (solid black line) posted a notably sharper decline over the course of 2009, falling below -2 percent by 2010 and remaining well below zero for a considerable period before drifting back up again by 2011. The estimate is particularly negative in the 2009-2011 period, and corresponds to the initially slow recovery in real activity after the GFC. The shadow rate estimate from this model showed similar behavior (as compared to its behavior in the earlier ELB episode) during the first year of the COVID-19 pandemic, dropping quickly in the spring of 2020. With the model including a range of macroeconomic and financial variables but excluding term structure data, the contours of shadow rate estimates resemble unconstrained Taylor-rule prescriptions calculated for the Great Recession years by Eberly et al. (2020).

As indicated in Panel (b), the estimates from the restricted version of the non-structural shadow-rate VAR including multiple interest rates display contours with some similarities to the shadow rate estimates from the non-structural model with the funds rate as the only interest rate, but also with noticeable differences. With the additional yields (and added restrictions) in the model, the estimated shadow rate declined more slowly than in the non-structural model without bond yields. [22] This estimate of the reduced-form shadow rate bottomed out only during the years 2011 and 2013 (and just below −2 percent), compared to the earlier trough during the years 2009 and 2010 in the model estimates without bond yields. The rate then gradually rose and crossed the ELB in early 2016. The rate dropped precipitously in the spring of 2020, with the posterior median reaching −3.4 percent in August 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. 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, although our shadow-rate VARs do not impose the restrictions of an affine term structure model, our shadow-rate estimates from the restricted version of the non-structural model including multiple interest rates have some similarities to the Krippner and Wu-Xia measures based on affine term structure models. As indicated in Panel (b) of Figure 2, our restricted version of the non-structural 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.

[22] Our earlier working paper (Carriero et al., 2023) 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.

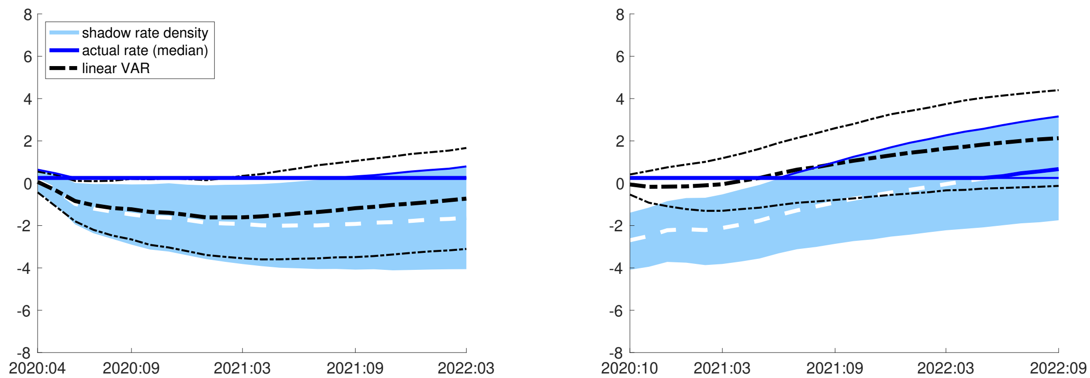


(a) 2020:03 (b) 2020:09

FIGURE 3. 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 restricted version of the non-structural 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.

Compared to the other estimates, the Krippner measure fell more sharply in 2012-2013 and then bounced back more quickly in 2014-2015.

### 5.3 *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 on the restricted version of the non-structural shadow-rate VAR that includes multiple interest rates as compared to a standard VAR. The key differences between the VAR specifications considered here (standard vs. shadow-rate) is the treatment of nominal interest rates near the ELB, which (while not shown) has implications for the resulting forecasts for other, mainly macroeconomic, variables.

Figure 3 shows the evolution of federal funds rate forecasts as of March 2020 and September 2020 origins.[23] With data available through March 2020, as the outbreak of COVID-19 hit the US economy, the point forecast from a standard VAR put the funds rate well below the ELB for the entire forecast horizon, with substantial probability mass on very negative rates. As of September 2020, the point forecast was close to the ELB for several months, but throughout the forecast 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 restricted version of the non-structural shadow-rate VAR was quite 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 noted earlier in this section, the rapid deterioration in economic conditions led to a decline in the (median) shadow rate, which stabilized a little below $-2.5$ 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.[24] 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.

In results not pictured in the interest of brevity, jumping forward in time a couple of years yields fewer differences in the models' predictive densities. In March 2022, when the FOMC raised the funds rate off the ELB, predictive densities of the funds rate from both the standard and shadow-rate VARs were solidly above the ELB throughout the forecast horizon, with comparable median projections.

---

[23] 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.

[24] 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.

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.

## 6. 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.

Overall, our shadow-rate specifications successfully address the ELB, which drastically improves interest rate forecasts (compared to a model that ignores the ELB), while broadly matching 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.

## 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. [8]

Aruoba, S. Boragan, Marko Mlikota, Frank Schorfheide, and Sergio Villalvazo (2022), "SVARs with occasionally-binding constraints." *Journal of Econometrics*, 231 (2), 477–499. [5, 6, 15]

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, 23]

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. [3]

Black, Fischer (1995), "Interest rates as options." *The Journal of Finance*, 50 (5), 1371–1376. [3, 8]

Carriero, Andrea, Joshua C. C. Chan, Todd E. Clark, and Massimiliano Marcellino (2022), "Corrigendum to: Large Bayesian vector autoregressions with stochastic volatility and non-conjugate priors." *Journal of Econometrics*, 227 (2), 506–512. [11, 15]

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. [3, 15]

Carriero, Andrea, Todd E. Clark, and Massimiliano Marcellino (2021), "No-arbitrage priors, drifting volatilities, and the term structure of interest rates." *Journal of Applied Econometrics*, 36 (5), 495–516. [7, 8]

Carriero, Andrea, Todd E. Clark, Massimiliano Marcellino, and Elmar Mertens (2023), "Shadow-rate VARs." Discussion Papers 14/2023, Deutsche Bundesbank. [4, 6, 7, 21, 29, 32]

Carriero, Andrea, Todd E. Clark, Massimiliano Marcellino, and Elmar Mertens (2024), "Addressing COVID-19 outliers in BVARs with stochastic volatility." *The Review of Economics and Statistics*, 106 (5), 1403–1417. [29]

Chan, Joshua C. C. (2021), "Minnesota-type adaptive hierarchical priors for large Bayesian VARs." *International Journal of Forecasting*, 37 (3), 1212–1226. [12]

Chan, Joshua C. C., Aubrey Poon, and Dan Zhu (2023), "High-dimensional conditionally Gaussian state space models with missing data." *Journal of Econometrics*, 236 (1), Article 105468. [18]

Chib, Siddhartha (1992), "Bayes inference in the Tobit censored regression model." *Journal of Econometrics*, 51 (1-2), 79–99. [6]

Chib, Siddhartha and Edward Greenberg (1998), "Analysis of multivariate probit models." *Biometrika*, 85, 347–361. [6]

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, 30]

Cogley, Timothy and Thomas J. Sargent (2005), "Drifts and volatilities: Monetary policies and outcomes in the post WWII US." *Review of Economic Dynamics*, 8 (2), 262–302. [17]

Crump, Richard K., Stefano Eusepi, Domenico Giannone, Eric Qian, and Argia Sbordone (2024), "A large Bayesian VAR of the United States economy." *International Journal of Central Banking*, forthcoming. [2, 12]

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, 15]

Diebold, Francis X. and Roberto S. Mariano (1995), "Comparing predictive accuracy." *Journal of Business and Economic Statistics*, 13 (3), 253–263. [23]

Duffy, James A., Sophocles Mavroeidis, and Sam Wycherley (2023), "Cointegration with occasionally binding constraints." Papers 2211.09604, arXiv.org. [18]

Durbin, J. and S. J. Koopman (2012), *Time Series Analysis by State Space Methods*, 2nd edition. Oxford Statistical Science Series, Oxford University Press. [18]

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. [6]

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. [18]

Giannone, Domenico, Michele Lenza, and Giorgio E. Primiceri (2015), "Prior selection for vector autoregressions." *The Review of Economics and Statistics*, 97 (2), 436–451. [12]

Gonzalez-Astudillo, Manuel and Jean-Philippe Laforte (2024), "Estimates of the natural rate of interest consistent with a supply-side structure and a monetary policy rule for the U.S. economy." *International Journal of Central Banking*, forthcoming. [7, 12]

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]

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. [8]

Ikeda, Daisuke, Shangshang Li, Sophocles Mavroeidis, and Francesco Zanetti (2024), "Testing the effectiveness of unconventional monetary policy in Japan and the United States." *American Economic Journal: Macroeconomics*, 16 (2), 250–286. [5, 8]

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. [3]

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, 7, 8, 12, 16, 17]

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. [8]

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. [4]

Kailath, Thomas, Ali H. Sayed, and Babak Hassibi (2000), *Linear Estimation*. Prentice Hall Information and System Sciences Series, Pearson Publishing. [18]

Kim, Sangjoon, Neil Shephard, and Siddhartha Chib (1998), "Stochastic volatility: Likelihood inference and comparison with ARCH models." *The Review of Economic Studies*, 65 (3), 361–393. [17]

Krippner, Leo (2015), *Zero Lower Bound Term Structure Modeling: A Practitioner's Guide*. Palgrave Macmillan. [3, 5, 6, 30]

Krippner, Leo (2020), "A note of caution on shadow rate estimates." *Journal of Money, Credit and Banking*, 52 (4), 951–962. [7]

Kulish, Mariano, James Morley, and Tim Robinson (2017), "Estimating DSGE models with zero interest rate policy." *Journal of Monetary Economics*, 88 (C), 35–49. [8]

Mavroeidis, Sophocles (2021), "Identification at the zero lower bound." *Econometrica*, 89 (6), 2855–2885. [4, 5, 6, 7, 9, 10, 11, 12, 14, 15]

McCracken, Michael W. and Serena Ng (2016), "FRED-MD: A monthly database for macroeconomic research." *Journal of Business & Economic Statistics*, 34 (4), 574–589. [19]

Mertens, Elmar (2023), "Precision-based sampling for state space models that have no measurement error." *Journal of Economic Dynamics and Control*, 154, 104720. [18]

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). [3]

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. [18]

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]

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, 15]

Wei, Steven X. (1999), "A Bayesian approach to dynamic Tobit models." *Econometric Reviews*, 18 (4), 417–439. [6]

West, Kenneth D. (1996), "Asymptotic inference about predictive ability." *Econometrica*, 64 (5), 1067–1084. [23]

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, 5, 6, 7, 30]

Wu, Jing Cynthia and Ji Zhang (2019), "A shadow rate New Keynesian model." *Journal of Economic Dynamics and Control*, 107, 103728. [8, 15]