



Kent Academic Repository

Chan, Joshua. C. C, Pettenuzzo, Davide, Poon, Aubrey and Zhu, Dan (2025)
Conditional forecasts in large Bayesian VARs with multiple equality and inequality constraints. *Journal of Economic Dynamics and Control*, 173 . ISSN 0165-1889.

Downloaded from

<https://kar.kent.ac.uk/108589/> The University of Kent's Academic Repository KAR

The version of record is available from

<https://doi.org/10.1016/j.jedc.2025.105061>

This document version

Publisher pdf

DOI for this version

Licence for this version

CC BY (Attribution)

Additional information

Versions of research works

Versions of Record

If this version is the version of record, it is the same as the published version available on the publisher's web site. Cite as the published version.

Author Accepted Manuscripts


If this document is identified as the Author Accepted Manuscript it is the version after peer review but before type setting, copy editing or publisher branding. Cite as Surname, Initial. (Year) 'Title of article'. To be published in **Title of Journal**, Volume and issue numbers [peer-reviewed accepted version]. Available at: DOI or URL (Accessed: date).

Enquiries

If you have questions about this document contact ResearchSupport@kent.ac.uk. Please include the URL of the record in KAR. If you believe that your, or a third party's rights have been compromised through this document please see our [Take Down policy](https://www.kent.ac.uk/guides/kar-the-kent-academic-repository#policies) (available from <https://www.kent.ac.uk/guides/kar-the-kent-academic-repository#policies>).



Conditional forecasts in large Bayesian VARs with multiple equality and inequality constraints [☆]

Joshua C.C. Chan ^a, Davide Pettenuzzo ^b, Aubrey Poon ^{c,d}, Dan Zhu ^{e,  *}

^a Purdue University, United States of America

^b Brandeis University, United States of America

^c University of Kent, United Kingdom

^d Orebro University, Sweden

^e Monash University, Australia

ARTICLE INFO

JEL classification:

C11
C32
C51

Keywords:

Precision-based method
Conditional forecast
Vector autoregression

ABSTRACT

Conditional forecasts, i.e. projections of a set of variables of interest on the future paths of some other variables, are used routinely by empirical macroeconomists in a number of applied settings. In spite of this, the existing algorithms used to generate conditional forecasts tend to be very computationally intensive, especially when working with large Vector Autoregressions or when multiple linear equality and inequality constraints are imposed at once. We introduce a novel precision-based sampler that is fast, scales well, and yields conditional forecasts from linear equality and inequality constraints. We show in a simulation study that the proposed method produces forecasts that are identical to those from the existing algorithms but in a fraction of the time. We then illustrate the performance of our method in a large Bayesian Vector Autoregression. Within this setting, we first highlight how we can simultaneously impose a mix of linear equality and inequality constraints on the future trajectories of several key US macroeconomic indicators over a forecast horizon spanning multiple years. Next, we test the benefits of using inequality constraints in an out-of-sample exercise spanning the period between 1995Q1 and 2022Q3 and find that imposing these constraints on the future path of Real GDP leads to significant improvement in point and density forecasts of the large BVAR model.

1. Introduction

Conditional forecasts are projections of a set of variables of interest on the future paths of some other variables. Since the seminal work of Waggoner and Zha (1999), conditional forecasts have become a popular tool for forecasters and policymakers within empirical macroeconomics. They are often paired with Vector Autoregressions (VAR) and are used routinely to project the future path of a set of macroeconomic variables after conditioning on a particular policy instrument or an important macroeconomic indicator, such as the Fed fund rate or real GDP.¹

[☆] We thank Dario Caldara, Dimitris Korobilis, Gary Koop, Giulia Mantoan, Michael McCracken, Hashem Pesaran, Ivan Petrella and Mike West for many constructive comments and useful suggestions.

* Corresponding author.

E-mail address: dan.zhu@monash.edu (D. Zhu).

¹ To produce conditional forecasts, one could alternatively treat the variable that is conditioned on as an exogenous deterministic process over the forecast period. When proceeding in this way, all existing methods, both frequentistic and Bayesian, can be applied. See, for example, the method of Sims and Zha (1998) for statistical

<https://doi.org/10.1016/j.jedc.2025.105061>

Received 2 July 2024; Received in revised form 14 November 2024; Accepted 30 January 2025

A first important consideration when generating conditional forecasts is whether one wants to compute classic (reduced-form) conditional forecasts or instead is interested in scenario analysis. In the former case, it is sufficient to rely on the empirical correlations between the variables in the system and not take a stance on the underlying causal mechanisms behind the results. This is equivalent to assuming that the conditions are generated by all the structural shocks of the model, and it is the most commonly adopted approach in the literature.² In practice, assuming that the conditions are generated by all structural shocks of the model may be undesirable and economically irrelevant, and there are, therefore, reasons to prefer a more structural approach, where the conditional forecasts are obtained from a sequence of specific shocks derived from a (point or set) identified structural VAR. See, for example, the discussion in Dieppe et al. (2016). An important contribution in this regard is the recent work of Antolín-Díaz et al. (2021), who introduced a unified framework for conditional forecasts and structural scenario analysis with VAR models.³ Their approach works by first specifying a set of conditions on observable variables and then pairing these conditions with the subset of structural shocks that are believed to be driving the forecasts.^{4,5}

A second and equally important consideration when producing conditional forecasts is whether one wants to fix the future paths of the conditioned variables at specific values (i.e., exact conditions) or instead prefers to allow the future values of the conditioned variables to lie within a certain range (i.e., equality or inequality conditions). The exact constraint case is the most commonly employed approach in the literature (see again Giannone et al. (2012, 2014) as well as Jarocinski and Smets (2008) and Lenza et al. (2010)), while to date there has been only limited work with equality and inequality constraints (in addition to Waggoner and Zha (1999), see also Andersson et al. (2010)). This is largely due to the computational challenges that the inequality constraints entail. For example, the algorithm of Waggoner and Zha (1999) for the inequality constraints is computationally very heavy, even when using the more efficient version of the algorithm proposed by Jarociński (2010). This is because the approach relies on a simple accept-reject algorithm that obtains candidates from the original, unconstrained distribution. Consequently, it requires a large number of candidate draws to obtain one that satisfies all the constraints.

The limited use of inequality constraint applications is unfortunate, as there are many good reasons why inequality constraints should appeal to forecasters and policymakers. For one, it is often the case that one does not know the actual future realization path of the constrained endogenous variables, and in these situations, it is much simpler to impose that the future values of the variables conditioned on will be within a range or an interval instead of an exact path. In addition, inequality constraints allow the forecaster to acknowledge the uncertainty surrounding the future realization path of the constrained endogenous variables. For example, Andersson et al. (2010) show that ignoring the uncertainty about the conditions can lead to density forecasts of the unrestricted variables that are too narrow.

In view of these considerations, we introduce a novel approach to conditional forecasting that generalizes and extends the existing methods available in the literature in a number of ways. First, much like Antolín-Díaz et al. (2021), our approach is closed-form and can be used for both conditional forecast and structural scenario analysis. Second, thanks to the way we derive the conditional forecasts' distribution, our approach is significantly more efficient and better suited to handling large dimensional VARs as well as situations in which we have a large number of conditioning variables and long forecast horizons. We accomplish this by building on the intuition of Bańbura et al. (2015), who shows that the conditional forecasts can be considered as time series with missing data. This, in turn, allows us to build on the efficient sampling approach developed in Chan et al. (2023)—designed for handling complex missing data patterns in linear Gaussian state space models—to our settings, where the missing data is subject to inequality constraints. Third, our approach can be used to produce conditional forecasts and structural scenario analysis under both equality and inequality conditions without compromising the computational efficiency and scalability of the algorithm. This is in contrast to the traditional methods available in the literature, such as Waggoner and Zha (1999) and Andersson et al. (2010).⁶ To accomplish this, we pair the use of the precision sampler of Chan and Jeliazkov (2009) to the exponential minimax tilting method of Botev (2017), which allows us to draw the inequality constraint conditional forecasts in a fast and efficient manner within the precision sampler.

We conduct a simulation study to compare the computation time and accuracy of our proposed precision-based conditional forecast sampler against the existing methods in the literature. Specifically, we compare our precision-based equality constraint conditional forecasts against the Waggoner and Zha (1999), Bańbura et al. (2015) and Antolín-Díaz et al. (2021) approaches. We show that our approach generates exactly the same conditional forecasts and credible sets as the three existing methods but is substantially less demanding computationally. Next, we investigate the accuracy and computational efficiency of our inequality constraint precision-based approach algorithm by comparing it against the Waggoner and Zha (1999), and Andersson et al. (2010) accept-reject methods

inferences on point forecasts. However, this approach has some limitations, as it is often advisable to include the variable conditioned on as endogenous in a VAR setting.

² Examples along this line of work include Andersson et al. (2010), Giannone et al. (2012, 2014), Altavilla et al. (2016), Aastveit et al. (2017), Giannone et al. (2019), and Tallman and Zaman (2020).

³ Antolín-Díaz et al. (2021) define a structural scenario as a combination of a path for one or more variables in the system and a restriction that only a subset of structural shocks can deviate from their unconditional distribution.

⁴ Relative to the existing methods in the literature (i.e., the seminal work of Waggoner and Zha (1999), the Kalman filter based approach of Clarida and Coyle (1984) and Bańbura et al. (2015), and the “forecast scenario” method of Baumeister and Kilian (2014)), the approach developed in Antolín-Díaz et al. (2021) is closed-form and valid for an arbitrary number of conditions.

⁵ Antolín-Díaz et al. (2021) also show that in the Gaussian case, their approach to conditional forecasting is equivalent to the entropic forecast tilting method of Robertson et al. (2005).

⁶ Andersson et al. (2010) show how to implement the inequality restrictions using a Gibbs sampler instead of the original accept-reject method introduced in Waggoner and Zha (1999). They rely on the Gibbs sampler proposed by Geweke (1996) and Genz (1992) to sample from a number of univariate truncated Gaussian distributions. However, their method is only suitable for a small-dimension VAR model with only one inequality constraint.

(both Bańbura et al. (2015) and Antolín-Díaz et al. (2021) are not designed to generate conditional forecasts in the presence of inequality constraints) and once again find that our inequality constraint algorithm is significantly faster than the existing alternatives. For example, using the Waggoner and Zha (1999) algorithm, it takes about 80 and 459 minutes to generate the in- and out-of-sample conditional forecasts from a small dimension VAR with only one inequality constraint. In contrast, our proposed approach only takes 1-2 minutes to generate the corresponding conditional forecasts from the same VAR model.

We next move to illustrate the benefits of our proposed precision-based conditional forecast sampler with two empirical applications, where we build on Crump et al. (2021) and estimate a large BVAR including 31 quarterly US macro and financial variables from 1976 to the end of 2022. We use the first application to illustrate the generality and flexibility of our method, simultaneously imposing a large number of equality and inequality constraints on the out-of-sample trajectories of CPI inflation, unemployment rate, and the 10-year Treasury rate. We stop the estimation sample of the BVAR in 2019Q4 and setup our constraints to mimic the baseline and severe scenarios prepared by the Federal Reserve Board for their 2020 stress test analysis, covering the 13 quarter periods ranging from 2020Q1 to 2023Q1. In our second application, we investigate whether having a scalable methodology to implement conditional forecasts with inequality constraints can be beneficial in an out-of-sample setting. We continue to work with the 31 variable BVAR of Crump et al. (2021), and carry out a real-time exercise where for each quarter between 1995Q1 and 2022Q3 we re-estimate the BVAR model using data up to that point in time and then use inequality constraints to generate conditional forecasts where we restrict the future path of Real GDP over the next four quarters to lie within the 25th and 75th percentiles of the SPF cross sectional dispersion. We focus our attention on six key variables (Real Consumption Growth, PCE Inflation, Non-farm employment growth, Industrial Production, Capacity Utilization and Fed Funds Rate) and compare the accuracy of the inequality constraint forecasts against both unconditional forecasts and multiple versions of conditional forecasts based on equality constraints. We find that, across the board, the inequality constraint conditional forecasts yield the most accurate point and density forecasts, with gains that at times are particularly large.

The remainder of the paper is organized as follows. Section 2 introduces the proposed precision-based conditional forecast sampler and illustrates how the approach can be used with equality and inequality constraints on observables as well as for scenario analysis and entropic tilting. Next, Section 3 presents the results of our simulation study, where we compare our proposed precision-based conditional forecast sampler against the existing methods available in the literature. Section 4 focuses on the empirical application, where we show how our approach can be used to generate conditional forecasts in a large Bayesian VAR when multiple inequality and equality constraints are at play simultaneously. Finally, Section 5 provides some concluding remarks.

2. Methodology

In this section, we introduce our approach to generating unconditional and conditional forecasts. Our starting point is a structural VAR (SVAR), a very flexible and general approach used routinely by empirical macroeconomists to produce forecasts and impulse response analysis. We follow the setups considered in both Waggoner and Zha (1999) and Antolín-Díaz et al. (2021), but at the same time we highlight that our framework is more general as we are able to accommodate within the same general methodology both equality and inequality constraints on observables, as well as constraints on structural shocks and scenario analysis. Most importantly, thanks to the use of precision-based sampling methods and the ability to exploit fast band matrix algorithms, we are able to achieve this generality without sacrificing computational efficiency, and this, in turn, makes our methods particularly suitable for handling large VARs, long forecast horizons, and multiple conditions.

We begin in subsection 2.1 with derivations of both unconditional and conditional predictive densities given the observed data and the model parameters. Next, we show in subsection 2.2.1 and subsection 2.2.2 how to apply this setup to producing equality- and inequality-constrained conditional forecasts. Finally, in subsection 2.3 we show how the approach can be extended to accommodate situations where the conditional forecasts depend only on a subset of structural shocks or to carry out structural scenario analysis.

2.1. Unconditional forecasts

To start, consider an $n \times 1$ vector of variables $\mathbf{y}_t = (y_{1,t}, \dots, y_{n,t})'$, and write the following SVAR with p lags:

$$\mathbf{A}_0 \mathbf{y}_t = \mathbf{a} + \mathbf{A}_1 \mathbf{y}_{t-1} + \dots + \mathbf{A}_p \mathbf{y}_{t-p} + \boldsymbol{\varepsilon}_t, \quad \boldsymbol{\varepsilon}_t \sim \mathcal{N}(\mathbf{0}_n, \mathbf{I}_n), \quad (1)$$

where \mathbf{a} is an $n \times 1$ vector of intercepts, $\mathbf{A}_1, \dots, \mathbf{A}_p$ are the $n \times n$ VAR coefficient matrices, \mathbf{A}_0 is a full-rank contemporaneous impact matrix, $\mathbf{0}_n$ is an $n \times 1$ vector of zeros and \mathbf{I}_n is the n -dimensional identity matrix.

Given the whole history of observations $\mathbf{y}^T = (\mathbf{y}'_{1-p}, \dots, \mathbf{y}'_T)'$, the unconditional forecast of our n variables for the next h periods, $\mathbf{y}_{T+1,T+h} = (\mathbf{y}'_{T+1}, \dots, \mathbf{y}'_{T+h})'$, can be written as

$$\mathbf{H} \mathbf{y}_{T+1,T+h} = \mathbf{c} + \boldsymbol{\varepsilon}_{T+1,T+h}, \quad \boldsymbol{\varepsilon}_{T+1,T+h} \sim \mathcal{N}(\mathbf{0}_{nh}, \mathbf{I}_{nh}), \quad (2)$$

where

$$\mathbf{c} = \begin{bmatrix} \mathbf{a} + \sum_{j=1}^p \mathbf{A}_j \mathbf{y}_{T+1-j} \\ \mathbf{a} + \sum_{j=2}^p \mathbf{A}_j \mathbf{y}_{T+2-j} \\ \mathbf{a} + \sum_{j=3}^p \mathbf{A}_j \mathbf{y}_{T+3-j} \\ \vdots \\ \mathbf{a} + \mathbf{A}_p \mathbf{y}_T \\ \mathbf{a} \\ \vdots \\ \mathbf{a} \end{bmatrix}, \mathbf{H} = \begin{bmatrix} \mathbf{A}_0 & \mathbf{0}_{n \times n} & \cdots & \cdots & \cdots & \cdots & \cdots & \mathbf{0}_{n \times n} \\ -\mathbf{A}_1 & \mathbf{A}_0 & \mathbf{0}_{n \times n} & \cdots & \cdots & \cdots & \cdots & \mathbf{0}_{n \times n} \\ -\mathbf{A}_2 & -\mathbf{A}_1 & \mathbf{A}_0 & \mathbf{0}_{n \times n} & \cdots & \cdots & \cdots & \mathbf{0}_{n \times n} \\ \vdots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots & \vdots \\ -\mathbf{A}_{p-1} & \cdots & \cdots & -\mathbf{A}_1 & \mathbf{A}_0 & \mathbf{0}_{n \times n} & \cdots & \vdots \\ \mathbf{0}_{n \times n} & \cdots & \cdots & \cdots & \cdots & \ddots & \ddots & \vdots \\ \vdots & \cdots & \cdots & \cdots & \cdots & \ddots & \ddots & \vdots \\ \mathbf{0}_{n \times n} & \cdots & \mathbf{0}_{n \times n} & -\mathbf{A}_p & \cdots & -\mathbf{A}_2 & -\mathbf{A}_1 & \mathbf{A}_0 \end{bmatrix} \quad (3)$$

and $\mathbf{0}_{n \times n}$ denotes the $n \times n$ zero matrix. Since \mathbf{A}_0 is of full-rank and the determinant of \mathbf{H} is $|\mathbf{A}_0|^h \neq 0$, the inverse \mathbf{H}^{-1} exists. It follows from (2) that

$$\mathbf{y}_{T+1, T+h} \sim \mathcal{N}(\mathbf{H}^{-1} \mathbf{c}, (\mathbf{H}'\mathbf{H})^{-1}). \quad (4)$$

It's worth noting that since \mathbf{H} is an $nh \times nh$ band matrix with band width np , the precision-based sampling approach of Chan and Jeliazkov (2009) can be used to efficiently draw from the unconditional distribution in (4), and this becomes particularly convenient when either (or both) n and h are large.⁷ We will rely heavily on this result as we expand our methods below, as this will allow us to introduce a fast and scalable algorithm to generate conditional forecasts in large BVARs.

2.2. Conditional forecasts

With this in mind, we now move to describe our approach to conditional forecasting. As in Andersson et al. (2010) and Antolín-Díaz et al. (2021), we write the conditional forecasts as a set of linear restrictions on the path of future observables $\mathbf{y}_{T+1, T+h}$, i.e.

$$\mathbf{R} \mathbf{y}_{T+1, T+h} \sim \mathcal{N}(\mathbf{r}, \mathbf{\Omega}), \quad (5)$$

where \mathbf{R} is a $r \times nh$ constant matrix with full row rank (so that there are no redundant restrictions), while \mathbf{r} and $\mathbf{\Omega}$ are $r \times 1$ and $r \times r$ matrices representing the mean and covariance of the restrictions. As Antolín-Díaz et al. (2021) note, the setup in (5) is very general and can accommodate both the classic exact linear equality constraint case as defined in Waggoner and Zha (1999) (simply setting $\mathbf{\Omega} = \mathbf{0}_{r \times r}$) and the more general density forecast case as defined by Andersson et al. (2010). It also replicates *entropic tilting* for the Gaussian case, the approach introduced and popularized by Robertson et al. (2005), which looks for the density forecast distribution that meets the constraints in (5) while minimizing the relative entropy with the unconditional forecast distribution in (4) (see Antolín-Díaz et al. (2021) for a formal proof of this equivalence in the Gaussian case). Next, combining (2) and (5), we obtain

$$\mathbf{R} \mathbf{y}_{T+1, T+h} = \mathbf{R} \mathbf{H}^{-1} \mathbf{c} + \mathbf{R} \mathbf{H}^{-1} \boldsymbol{\varepsilon}_{T+1, T+h} \sim \mathcal{N}(\mathbf{r}, \mathbf{\Omega}). \quad (6)$$

In what follows, we first derive the set of restrictions on the future shocks implied by (5) and (6). In particular, following Antolín-Díaz et al. (2021), we let $\boldsymbol{\varepsilon}_{T+1, T+h} | \mathbf{R}, \mathbf{r}, \mathbf{\Omega}$ denote the restricted future shocks with the distribution

$$\boldsymbol{\varepsilon}_{T+1, T+h} | \mathbf{R}, \mathbf{r}, \mathbf{\Omega} \sim \mathcal{N}(\boldsymbol{\mu}_\varepsilon, \mathbf{I}_{nh} + \boldsymbol{\Psi}_\varepsilon), \quad (7)$$

where $\boldsymbol{\mu}_\varepsilon$ and $\boldsymbol{\Psi}_\varepsilon$ are, respectively, the deviations of the mean vector and covariance matrix of the restricted future shocks from their unconditional counterparts in (2). In turn, equations (6) and (7) combined imply the following restrictions on $\boldsymbol{\mu}_\varepsilon$ and $\boldsymbol{\Psi}_\varepsilon$:

$$\begin{aligned} \mathbf{R} \mathbf{H}^{-1} (\mathbf{c} + \boldsymbol{\mu}_\varepsilon) &= \mathbf{r} \\ \mathbf{R} \mathbf{H}^{-1} (\mathbf{I}_{nh} + \boldsymbol{\Psi}_\varepsilon) \mathbf{H}^{-1'} \mathbf{R}' &= \mathbf{\Omega}. \end{aligned} \quad (8)$$

In typical situations where $r < nh$, the system in (8) is underdetermined and has multiple solutions.⁸ Following Antolín-Díaz et al. (2021), we choose a solution that can be expressed in terms of the Moore-Penrose inverse of $\mathbf{R} \mathbf{H}^{-1}$, which we denote as $(\mathbf{R} \mathbf{H}^{-1})^+$:

$$\begin{aligned} \boldsymbol{\mu}_\varepsilon &= (\mathbf{R} \mathbf{H}^{-1})^+ (\mathbf{r} - \mathbf{R} \mathbf{H}^{-1} \mathbf{c}) \\ \boldsymbol{\Psi}_\varepsilon &= (\mathbf{R} \mathbf{H}^{-1})^+ (\mathbf{\Omega} - \mathbf{R} (\mathbf{H}'\mathbf{H})^{-1} \mathbf{R}') (\mathbf{R} \mathbf{H}^{-1})^{+'}. \end{aligned} \quad (9)$$

This solution minimizes the sum of the Frobenius norms of $\boldsymbol{\mu}_\varepsilon$ and $\boldsymbol{\Psi}_\varepsilon$. In other words, this solution represents the smallest deviations of the mean vector and covariance matrix of the conditional future shocks from the unconditional ones. Finally, we can map the constraints on the future shocks implied by (5) and (6) to the corresponding constraints on the forecasts. If we denote the conditional forecast distribution with

$$\mathbf{y}_{T+1, T+h} | \mathbf{R}, \mathbf{r}, \mathbf{\Omega} \sim \mathcal{N}(\boldsymbol{\mu}_y, \boldsymbol{\Sigma}_y), \quad (10)$$

⁷ More specifically, to obtain a draw from the Gaussian distribution in (4), we first obtain $\mathbf{u} = (u_1, \dots, u_{nh})'$, where $u_i \sim \mathcal{N}(0, 1), i = 1, \dots, nh$. Then, we solve the linear system $\mathbf{H} \mathbf{z} = \mathbf{c} + \mathbf{u}$ for \mathbf{z} by, e.g., Gaussian elimination. The last operation can be done very quickly as the matrix \mathbf{H} is banded. It can be easily verified that $\mathbf{z} \sim \mathcal{N}(\mathbf{H}^{-1} \mathbf{c}, (\mathbf{H}'\mathbf{H})^{-1})$.

⁸ If $r = nh$, the system is just determined and has a unique solution; if $r > nh$, the system is inconsistent and there are no solutions.

then, (2) and (9) imply that

$$\begin{aligned} \mu_y &= \mathbf{H}^{-1} \left[\mathbf{c} + (\mathbf{R}\mathbf{H}^{-1})^+ (\mathbf{r} - \mathbf{R}\mathbf{H}^{-1}\mathbf{c}) \right] \\ \Sigma_y &= \mathbf{H}^{-1} \left[\mathbf{I}_{nh} + (\mathbf{R}\mathbf{H}^{-1})^+ (\mathbf{\Omega} - \mathbf{R}(\mathbf{H}'\mathbf{H})^{-1}\mathbf{R}') (\mathbf{R}\mathbf{H}^{-1})^{+'} \right] \mathbf{H}'^{-1}. \end{aligned} \tag{11}$$

This result is extremely general and in fact encompasses a number of useful and popular applications of conditional forecasting. For example, if one wished to impose restrictions on the mean of the conditional forecasts μ_y while preserving the variance of the unconditional forecasts, this could be easily accomplished by setting $\mathbf{\Omega} = \mathbf{R}(\mathbf{H}'\mathbf{H})^{-1}\mathbf{R}'$. In that case, it is easy to show that $\Sigma_y = (\mathbf{H}'\mathbf{H})^{-1}$, i.e. the covariance matrix of the unconditional forecasts. Similarly, if one were to consider imposing restrictions on the second moment of the conditional forecasts while preserving the mean of the unconditional forecasts, this could be accomplished by setting $\mathbf{r} = \mathbf{R}\mathbf{H}^{-1}\mathbf{c}$. Also note, to conclude, that our result is equivalent to what is reported in Antolín-Díaz et al. (2021). The key difference is that they parameterize the unconditional and conditional forecast distributions in terms of the covariance matrix, while we work with the precision matrix (i.e., the inverse covariance matrix), and this, in turn, allows us to exploit fast band matrix algorithms such as those introduced by Chan and Jeliaskov (2009) to substantially speed up computations, especially as the dimension of the SVAR n or the dimension of the forecast horizon h (or both) growth. We will return to this point in the simulation section, where we will focus on the computational gains of the approach.

In many applications, the constraints in (5) may be considered too strong, or one may not have enough information to elicit both the mean and covariance of the restrictions, especially as the forecast horizon increases. Similarly, in other settings, one may question the choice of the normal distribution to express the constraint. As a concrete example of this, consider a situation where there is significant disagreement on the future path of some variable(s) we would like to condition our forecasts on, and this disagreement is best described via a bi-modal distribution. Examples where this may be the case are not hard to find, and this was in fact the case with the Survey of Professional Forecasters (SPF) for a number of macroeconomic variables right after the first wave of the Covid-19 pandemic in the second half of 2020.⁹ In these situations, it may be easier and more plausible for the forecaster to specify a range for the future path of the constrained endogenous variables. It is trivial to extend the setup above to accommodate the case of inequality conditioning, where one allows the future values of the conditioned variables to lie within a certain range. Let the inequality constraints be expressed as

$$\underline{\mathbf{c}} < \mathbf{S}\mathbf{y}_{T+1,T+h} < \bar{\mathbf{c}} \tag{12}$$

where \mathbf{S} is a $s \times nh$ pre-specified full-rank constant matrix, $\underline{\mathbf{c}}$ and $\bar{\mathbf{c}}$ are $s \times 1$ vectors of constants (with generic elements \underline{c}_i and \bar{c}_i in $\mathbb{R} \cup \{\pm\infty\}$) and the inequalities hold component-wise. In typical applications, \mathbf{S} would be a selection matrix, but our framework allows for inequality restrictions on any linear combinations of the variables. Combining (4) with (12) leads to a truncated multivariate normal distribution for $\mathbf{y}_{T+1,T+h}$,

$$\mathbf{y}_{T+1,T+h} | \underline{\mathbf{c}} < \mathbf{S}\mathbf{y}_{T+1,T+h} < \bar{\mathbf{c}} \sim \mathcal{N}(\mathbf{H}^{-1}\mathbf{c}, (\mathbf{H}'\mathbf{H})^{-1}) \mathbb{1}(\underline{\mathbf{c}} < \mathbf{S}\mathbf{y}_{T+1,T+h} < \bar{\mathbf{c}}), \tag{13}$$

where $\mathbb{1}(\cdot)$ is the indicator function. We note that this is a slightly more general formulation of inequality constraints than that in Waggoner and Zha (1999) and it reduces to the latter when setting $\mathbf{S} = \mathbf{I}_{nh}$. We also note that this approach is closely related to Andersson et al. (2010), but again since we are parameterizing the forecast distribution using the precision matrix $\mathbf{H}'\mathbf{H}$, our approach allows us to exploit the computational efficiency of the precision-based sampling approach of Chan and Jeliaskov (2009) and it will therefore lead to large computational gains.

It is also straightforward to combine the two types of constraints in (5) and (13). In particular, suppose we impose restrictions on μ_y while preserving the variance of the unconditional forecasts by setting $\mathbf{\Omega} = \mathbf{R}(\mathbf{H}'\mathbf{H})^{-1}\mathbf{R}'$. Then, using the same logic we applied above, we obtain

$$\mathbf{y}_{T+1,T+h} | \mathbf{R}, \mathbf{r}, \mathbf{\Omega}, \underline{\mathbf{c}} < \mathbf{S}\mathbf{y}_{T+1,T+h} < \bar{\mathbf{c}} \sim \mathcal{N}(\mu_y, (\mathbf{H}'\mathbf{H})^{-1}) \mathbb{1}(\underline{\mathbf{c}} < \mathbf{S}\mathbf{y}_{T+1,T+h} < \bar{\mathbf{c}}) \tag{14}$$

where the definition of μ_y is given in (11).¹⁰

Next, we showcase the generality of the approach with a broad set of examples. We begin with the classic exact or equality conditional forecast case introduced by Doan et al. (1984) and popularized by Waggoner and Zha (1999). Next, we move to the inequality constraint case, also discussed in Waggoner and Zha (1999). In both of these two cases, the conditional forecasts are in reduced form, i.e. rely on all shocks. While this is useful, it is often very relevant to conditions on a subset of structural shocks. We review that last.

⁹ See for example the projections of GDP and unemployment in the 2020Q2 SPF, released on May 15, 2020 and available here: <https://www.philadelphiafed.org/-/media/frbp/assets/surveys-and-data/survey-of-professional-forecasters/2020/spfq220.pdf?la=en>.

¹⁰ In the special case without inequality constraints, i.e. when $\underline{c}_i = \infty$ and $\bar{c}_i = \infty, i = 1, \dots, nh$, we revert back to the results presented in (5). More generally, to impose both types of constraints in (5) and (13), we first select $(\mu_\epsilon, \Sigma_\epsilon)$ that minimizes the deviation of future shocks $\epsilon_{T+1,T+h}$ from their unconditional distribution as in Antolín-Díaz et al. (2021); the explicit solution is given in (9). Then, we restrict the support of the joint distribution of future shocks implied by the inequality restrictions.

2.2.1. Equality constraints on future observables

The first special case is the classic conditional forecast setup introduced by Doan et al. (1984) and popularized by Waggoner and Zha (1999). It is used to produce forecasts of the variables of interest given the future path of a subset of other variables. As an example, a policymaker might be interested in the future path of GDP, inflation and unemployment that is conditioned on the scenario that future policy rate follows a fixed path across the forecast horizon.

This type of restriction can be represented as

$$\mathbf{R}_o \mathbf{y}_{T+1,T+h} = \mathbf{r}_o, \tag{15}$$

where \mathbf{R}_o is a $r_o \times nh$ pre-specified full-rank selection matrix—a matrix in which each row has exactly one element that is 1 and all other elements are 0—and \mathbf{r}_o is a $r_o \times 1$ vector of constants. The setup in (15) can be cast within our general framework in (5) by setting $\mathbf{R} = \mathbf{R}_o$, $\mathbf{r} = \mathbf{r}_o$ and $\mathbf{\Omega} = \mathbf{0}_{r_o \times r_o}$ and is sometimes referred to as conditional forecasting under exact or equality constraints.

In this case, some of the computations described in the general framework become unnecessary and the algorithm can be streamlined. In particular, the efficient sampling approach of Chan et al. (2023) designed for missing data with linear equality constraints can be directly applied. More specifically, we partition the $nh \times 1$ vector $\mathbf{y}_{T+1,T+h}$ into $\mathbf{y}_{T+1,T+h}^o$ and $\mathbf{y}_{T+1,T+h}^u$, where the former is a $r_o \times 1$ vector including the equality-constrained endogenous variables—the set of variables that are selected by \mathbf{R}_o —and the latter is a $(nh - r_o) \times 1$ vector of free or unconstrained variables. For example, if we are interested in the conditional forecasts of GDP, inflation and unemployment given that the policy rate follows a fixed path over the next h periods, then $\mathbf{y}_{T+1,T+h}^o$ is the $h \times 1$ vector of policy rates ($r_o = h$) over the forecast horizons $t = T + 1, \dots, T + h$ and $\mathbf{y}_{T+1,T+h}^u$ is the $3h \times 1$ vector consisting of GDP, inflation and unemployment. With this in mind, let \mathbf{R}_o^- denote the associated $(nh - r_o) \times nh$ selection matrix that selects $\mathbf{y}_{T+1,T+h}^u$. Then, we can write $\mathbf{y}_{T+1,T+h}$ as follows:

$$\mathbf{y}_{T+1,T+h} = \mathbf{M}_u \mathbf{y}_{T+1,T+h}^u + \mathbf{M}_o \mathbf{y}_{T+1,T+h}^o, \tag{16}$$

where $\mathbf{M}_u = (\mathbf{R}_o^-)'$ and $\mathbf{M}_o = \mathbf{R}_o'$. Note that both \mathbf{M}_u and \mathbf{M}_o have full column rank and are sparse with only, respectively, $nh - r_o$ and r_o non-zero elements.

With this setup in hand, we can now derive the joint conditional distribution of $\mathbf{y}_{T+1,T+h}^u$ given $\mathbf{y}_{T+1,T+h}^o$ and the model parameters \mathbf{A}_0 and $\mathbf{A} = (\mathbf{a}, \mathbf{A}_1, \dots, \mathbf{A}_p)'$. To this end, substitute (16) into (2) to obtain

$$\mathbf{H}(\mathbf{M}_u \mathbf{y}_{T+1,T+h}^u + \mathbf{M}_o \mathbf{y}_{T+1,T+h}^o) = \mathbf{c} + \boldsymbol{\varepsilon}_{T+1,T+h}, \quad \boldsymbol{\varepsilon}_{T+1,T+h} \sim \mathcal{N}(\mathbf{0}_n, \mathbf{I}_{nh}).$$

Next, the conditional density of $\mathbf{y}_{T+1,T+h}^u$ given $\mathbf{y}_{T+1,T+h}^o$ and the model parameters can be expressed as (recall \mathbf{H} and \mathbf{c} are functions of the model parameters \mathbf{A}_0 and \mathbf{A}):

$$\begin{aligned} p(\mathbf{y}_{T+1,T+h}^u | \mathbf{y}_{T+1,T+h}^o, \mathbf{A}_0, \mathbf{A}) & \propto \exp \left\{ -\frac{1}{2} (\mathbf{H}(\mathbf{M}_u \mathbf{y}_{T+1,T+h}^u + \mathbf{M}_o \mathbf{y}_{T+1,T+h}^o) - \mathbf{c})' (\mathbf{H}(\mathbf{M}_u \mathbf{y}_{T+1,T+h}^u + \mathbf{M}_o \mathbf{y}_{T+1,T+h}^o) - \mathbf{c}) \right\} \\ & \propto \exp \left\{ -\frac{1}{2} \left(\mathbf{y}_{T+1,T+h}^u{}' \mathbf{M}_u' \mathbf{H}' \mathbf{H} \mathbf{M}_u \mathbf{y}_{T+1,T+h}^u - 2 \mathbf{y}_{T+1,T+h}^u{}' \mathbf{M}_u' \mathbf{H}' (\mathbf{c} - \mathbf{H} \mathbf{M}_o \mathbf{y}_{T+1,T+h}^o) \right) \right\} \\ & \propto \exp \left\{ -\frac{1}{2} (\mathbf{y}_{T+1,T+h}^u - \boldsymbol{\mu}_u)' \mathbf{K}_u (\mathbf{y}_{T+1,T+h}^u - \boldsymbol{\mu}_u) \right\}, \end{aligned}$$

where $\mathbf{K}_u = \mathbf{M}_u' \mathbf{H}' \mathbf{H} \mathbf{M}_u$ and $\boldsymbol{\mu}_u = \mathbf{K}_u^{-1} \mathbf{M}_u' \mathbf{H}' \mathbf{H} (\mathbf{c} - \mathbf{H} \mathbf{M}_o \mathbf{y}_{T+1,T+h}^o)$. That is,

$$\mathbf{y}_{T+1,T+h}^u | \mathbf{y}_{T+1,T+h}^o, \mathbf{A}_0, \mathbf{A} \sim \mathcal{N}(\boldsymbol{\mu}_u, \mathbf{K}_u^{-1}).$$

Since \mathbf{H} and \mathbf{M}_u are band matrices, so is the precision matrix \mathbf{K}_u . Therefore, we can again use the precision sampler of Chan and Jeliazkov (2009) to draw $\mathbf{y}_{T+1,T+h}^u$ efficiently.

2.2.2. Inequality constraints on future observables

We now move to discussing the case of inequality constraints on observables. Continuing with the previous example, instead of conditioning on a fixed path of the future policy rate, we could be interested in restricting the future path of the policy rate to be, say, between 1% and 2% for the next 8 quarters and between 1.5% and 2.5% afterward. In terms of our general framework, this type of inequality constraints on the observables can be formulated via (13). In particular, this case can be specified simply by setting $\mathbf{S} = \mathbf{I}_{nh}$ and the appropriate elements in $\underline{\mathbf{c}}$ and $\bar{\mathbf{c}}$ to the desired values.

We note, however, that for large VARs with many inequality constraints, sampling from the truncated multivariate Gaussian distribution in (13) can be a computationally daunting task. For example, Waggoner and Zha (1999) implement a simple accept-reject algorithm where the proposal is the unrestricted multivariate Gaussian distribution. In their empirical application with a small VAR, they note that out of 185,000 simulated draws only 8,000 draws satisfy the inequality constraints. For high-dimensional problems with many endogenous variables and inequality constraints, it is clear that this approach is computationally infeasible. To circumvent this challenge, Andersson et al. (2010) instead use the approach in Geweke (1996) to draw from the truncated multivariate Gaussian distribution implied by the inequality constraints. Specifically, it relies on a Gibbs sampler that iteratively samples from each univariate conditional distribution, which in this case is a univariate truncated Gaussian distribution. Note that being a Gibbs sampler,

in high-dimensional settings where the components are highly correlated this approach would tend to generate highly auto-correlated MCMC draws.

Our proposed solution to this problem is to rely on the minimax tilting method of Botev (2017), a general algorithm that directly samples from a potentially high-dimensional Gaussian distribution under linear inequality restrictions. The key idea is to locate a proposal distribution by tilting the mean vector of the original unrestricted Gaussian distribution. This is done optimally by solving a minimax problem: it minimizes the worst-case behavior of the likelihood ratio, defined as the ratio of the original density divided by the proposal density. Due to the log-concavity of the Gaussian distribution, this optimization problem can be solved efficiently. The solution to this minimax optimization problem then provides a proposal distribution for the accept-reject algorithm with strong efficiency properties. As opposed to the Gibbs sampler of Geweke (1996), this approach directly samples from the target truncated multivariate Gaussian distribution.

While the minimax tilting method of Botev (2017) can be directly applied to draw from the truncated multivariate Gaussian distribution given in (13), we can further improve sampling efficiency by exploiting the special structure of our simulation problem. More specifically, suppose a subset of the future observables $\mathbf{y}_{T+1,T+h}$ is subjected to the inequality constraints such that \mathbf{S} is a $s_o \times nh$ selection matrix. Then, as in the equality-constraints-on-observables case, we can partition $\mathbf{y}_{T+1,T+h}$ into $\mathbf{y}_{T+1,T+h}^o$ and $\mathbf{y}_{T+1,T+h}^u$, where the former is a $s_o \times 1$ vector of the inequality-constrained endogenous variables such that $\mathbf{y}_{T+1,T+h}^o = \mathbf{S}\mathbf{y}_{T+1,T+h}$ and the latter is a $(nh - s_o) \times 1$ vector of unconstrained variables. Using the representation in (16) (with $\mathbf{M}_o = \mathbf{S}'$ in this case) and a similar derivation as before, we can decompose the joint distribution of $\mathbf{y}_{T+1,T+h}$ into the product of the marginal distribution of $\mathbf{y}_{T+1,T+h}^o$ and the conditional distribution of $\mathbf{y}_{T+1,T+h}^u$ given $\mathbf{y}_{T+1,T+h}^o$. In fact, the former is a truncated s_o -variate Gaussian distribution and the latter is a $(nh - s_o)$ -variate Gaussian distribution (with no restrictions):

$$\mathbf{y}_{T+1,T+h}^o | \mathbf{A}_0, \mathbf{A} \sim \mathcal{N}(\boldsymbol{\mu}_o, \mathbf{K}_o^{-1}) \mathbb{1}(\underline{\mathbf{c}} < \mathbf{y}_{T+1,T+h}^o < \bar{\mathbf{c}})$$

$$\mathbf{y}_{T+1,T+h}^u | \mathbf{y}_{T+1,T+h}^o, \mathbf{A}_0, \mathbf{A} \sim \mathcal{N}(\boldsymbol{\mu}_u, \mathbf{K}_u^{-1}),$$

where $\mathbf{K}_u = \mathbf{M}'_u \mathbf{H}' \mathbf{H} \mathbf{M}_u$, $\boldsymbol{\mu}_u = \mathbf{K}_u^{-1} \mathbf{M}'_u \mathbf{H}' \mathbf{H} (\mathbf{H}^{-1} \mathbf{c} - \mathbf{M}_o \mathbf{y}_{T+1,T+h}^o)$, $\mathbf{K}_o = (\mathbf{M}'_o (\mathbf{H}' \mathbf{H})^{-1} \mathbf{M}_o)^{-1}$, and $\boldsymbol{\mu}_o = \mathbf{M}'_o \mathbf{H}^{-1} \mathbf{c}$.¹¹ Hence, to obtain a draw for $\mathbf{y}_{T+1,T+h}$, we can first sample $\mathbf{y}_{T+1,T+h}^o$ marginally from its truncated s_o -variate Gaussian distribution using the algorithm of Botev (2017). Given a draw for $\mathbf{y}_{T+1,T+h}^o$, we can then sample $\mathbf{y}_{T+1,T+h}^u$ from its Gaussian conditional distribution using the precision sampler of Chan and Jeliazkov (2009), which can be done very quickly and the computational cost increases only linearly in the dimension. In typical applications where s_o is much smaller than nh , this approach based on the marginal-conditional decomposition is substantially more efficient, as it reduces the dimension of the more computationally intensive sampling step from the truncated Gaussian from nh to s_o .

2.3. Constraints on structural shocks and structural scenario analysis

While the discussion above has focused on the classic (or reduced-form) conditional forecast case, where one assumes that the equality or inequality conditions are generated by all the structural shocks of the model, it is worthwhile showing that our framework can be straightforwardly extended to accommodate the situation in which it is of interest to produce conditional forecasts by restricting the path of a subset of structural shocks over the forecast horizon. This specific case is discussed for example in Baumeister and Kilian (2014) and Antolín-Díaz et al. (2021). Using our notation, this type of restriction on structural shocks can be formulated as

$$\mathbf{W}\boldsymbol{\varepsilon}_{T+1,T+h} \sim \mathcal{N}(\mathbf{w}, \boldsymbol{\Psi}), \tag{17}$$

where \mathbf{W} is a $w \times nh$ full-rank selection matrix, \mathbf{w} is a $w \times 1$ vector of constants and $\boldsymbol{\Psi}$ is a $w \times w$ covariance matrix. To see how this case is also nested within our general framework, note that we can use (2) to write the structural shocks $\boldsymbol{\varepsilon}_{T+1,T+h}$ in terms of $\mathbf{y}_{T+1,T+h}$. Then, the restrictions on the structural shocks can be rewritten as restrictions on the observables:

$$\begin{aligned} \mathbf{W}\boldsymbol{\varepsilon}_{T+1,T+h} &= \mathbf{W}(\mathbf{H}\mathbf{y}_{T+1,T+h} - \mathbf{c}) \\ &= \mathbf{W}\mathbf{H}\mathbf{y}_{T+1,T+h} - \mathbf{W}\mathbf{c} \sim \mathcal{N}(\mathbf{w}, \boldsymbol{\Psi}) \end{aligned} \tag{18}$$

Comparing this expression with (5), we see that we can implement this type of restriction by setting $\mathbf{R} = \mathbf{W}\mathbf{H}$, $\mathbf{r} = \mathbf{W}\mathbf{c} + \mathbf{w}$ and $\boldsymbol{\Omega} = \boldsymbol{\Psi}$. And, as with the constraints-on-observables case discussed in subsection 2.2.1, simulating the conditional forecasts conditional on the structural shocks can often be streamlined by exploiting the special structure of the problem (e.g., in the case of equality constraints with $\boldsymbol{\Psi} = \mathbf{0}_{w \times w}$).¹²

Finally, we consider the structural scenario case discussed in Antolín-Díaz et al. (2021). A structural scenario combines restrictions on the path of future observations with the restriction that only a subset of the structural shocks—what Antolín-Díaz et al. (2021) call the driving shocks—can deviate from their unconditional distribution over the forecast horizon, while the remaining structural

¹¹ Note that in general, a marginal distribution of a truncated Gaussian is not a truncated Gaussian. Here the marginal distribution of $\mathbf{y}_{T+1,T+h}^o$ is a truncated Gaussian because only $\mathbf{y}_{T+1,T+h}^o$ is restricted and $\mathbf{y}_{T+1,T+h}^u$ is not.

¹² In the case where we only impose inequality restrictions on the observables, one can characterize the distribution of the shocks implied by the bounds on the observables. More specifically, suppose $\mathbf{y}_{T+1,T+h}$ are bounded: $\underline{\mathbf{c}} < \mathbf{y}_{T+1,T+h} < \bar{\mathbf{c}}$. It follows from (2) that $\boldsymbol{\varepsilon}_{T+1,T+h} \sim \mathcal{N}(\mathbf{0}_{nh}, \mathbf{I}_{nh})$ and $\underline{\mathbf{c}} < \mathbf{H}^{-1} \mathbf{c} + \mathbf{H}^{-1} \boldsymbol{\varepsilon}_{T+1,T+h} < \bar{\mathbf{c}}$. Consequently, $\boldsymbol{\varepsilon}_{T+1,T+h}$ has a multivariate normal distribution $\mathcal{N}(\mathbf{0}_{nh}, \mathbf{I}_{nh})$ truncated to the set $\mathcal{R} = \{\boldsymbol{\varepsilon}_{T+1,T+h} \in \mathbb{R}^{nh} : \mathbf{H}\underline{\mathbf{c}} - \mathbf{c} < \boldsymbol{\varepsilon}_{T+1,T+h} < \mathbf{H}\bar{\mathbf{c}} - \mathbf{c}\}$.

shocks—the non-driving shocks—are restricted to retain their unconditional distribution. As such, this setup is more flexible and plausible than conditioning on a particular future path of structural shocks, as structural shocks are unobserved and it is typically difficult to elicit restrictions on their future path. It is also more appealing than restricting only the future path of observables since the user in this case can specify which structural shocks deviate from their unconditional distribution to drive the future observables.

Our setup can be easily extended to accommodate this particular case, as a structural scenario can be formulated by combining restrictions on observables and restrictions on structural shocks in a straightforward way. First, we can restrict the path of future observables using (15). We can then augment these restrictions by appending (17), where \mathbf{W} now denotes the $w \times nh$ full-rank selection matrix flagging the non-driving shocks. Over the forecasting horizon, these non-driving shocks, as Antolín-Díaz et al. (2021) discussed, should retain their unconditional distribution, i.e. $\mathbf{W}\boldsymbol{\varepsilon}_{T+1,T+h} \sim \mathcal{N}(\mathbf{0}_w, \mathbf{I}_w)$. This, in turn, implies

$$\mathbf{W}\mathbf{H}\mathbf{y}_{T+1,T+h} \sim \mathcal{N}(\mathbf{W}\mathbf{c}, \mathbf{I}_w). \tag{19}$$

Combining (15) with (19) yields

$$\underbrace{\begin{bmatrix} \mathbf{R}_o \\ \mathbf{W}\mathbf{H} \end{bmatrix}}_{\tilde{\mathbf{R}}} \mathbf{y}_{T+1,T+h} \sim \mathcal{N} \left(\underbrace{\begin{bmatrix} \mathbf{r}_o \\ \mathbf{W}\mathbf{c} \end{bmatrix}}_{\tilde{\mathbf{r}}}, \underbrace{\begin{bmatrix} \boldsymbol{\Omega}_o & \mathbf{0}_{r_o \times w} \\ \mathbf{0}_{w \times r_o} & \mathbf{I}_w \end{bmatrix}}_{\tilde{\boldsymbol{\Omega}}} \right). \tag{20}$$

Again, comparing this expression with (5), we see that this setting can be nested within our general framework by setting $\mathbf{R} = \tilde{\mathbf{R}}$, $\mathbf{r} = \tilde{\mathbf{r}}$ and $\boldsymbol{\Omega} = \tilde{\boldsymbol{\Omega}}$.

It is worth noting that while the cases reviewed in this section follow closely the discussion in Antolín-Díaz et al. (2021), by casting these restrictions within our general setup we are able to generalize and improve upon their approach. First, as noted earlier, both equality or inequality constraints can be imposed on the observables, and this carries over to the structural scenario restrictions. For example, inequality constraints on observables can be easily added by simply switching from (4) to (13). Second, the precision-based methods described in earlier sections can be used in both the structural shock restriction and the structural scenario cases to speed up the sampling of $\mathbf{y}_{T+1,T+h}$ and seamlessly apply these methods in large dimensional VARs or situations with many restrictions being imposed.

It is also worth pointing out that while the exposition in this section centers around the structural-form representation, it is straightforward to apply the methodology in reduced-form settings for simple conditional forecasting. This can be done, for example, by setting the impact matrix \mathbf{A}_0 as the inverse Cholesky factor of the reduced-form error covariance matrix, and computing the implied structural-form coefficient matrices $\mathbf{A}_1, \dots, \mathbf{A}_p$ from the reduced-form parameters.

3. A simulation study

To illustrate the performance of our precision-based conditional forecast samplers, we conduct a series of simulations in which we compare our proposed algorithms against existing conditional forecast algorithms in the literature. Our benchmarks in this exercise are the approaches of Waggoner and Zha (1999), Andersson et al. (2010), Bańbura et al. (2015) and Antolín-Díaz et al. (2021). We will focus on reporting both accuracy (i.e. how similar our conditional forecasts are to those produced by the various benchmarks) and speed (i.e. what is the computational cost of implementing the various methods). For all calculations in this section, we implement the algorithms using MATLAB and a standard desktop with an Intel Xeon W-2224 @3.60 GHz processor and 16 GB memory.

In most of our simulations, we consider a data-generating process (DGP) that follows an n -variable VAR with $p = 2$ lags:

$$\mathbf{y}_t = \mathbf{b} + \mathbf{B}_1 \mathbf{y}_{t-1} + \mathbf{B}_2 \mathbf{y}_{t-2} + \boldsymbol{\varepsilon}_t, \quad \boldsymbol{\varepsilon}_t \sim \mathcal{N}(\mathbf{0}_n, \boldsymbol{\Sigma}). \tag{21}$$

but as a robustness we will also investigate the case of $p = 4$. We set $T = 300$ and $\mathbf{b} = 0.01 \times \mathbf{1}_n$, where $\mathbf{1}_n$ is an $n \times 1$ column of ones. We then generate the diagonal elements of the first lag coefficient matrix from $\mathcal{U}(0, 0.5)$ and the off-diagonal elements from $\mathcal{U}(-0.2, 0.2)$. All the other elements of the higher VAR coefficient matrix are generated independently from $\mathcal{N}(0, 0.05^2/p^2)$. Finally, we generate the covariance matrix $\boldsymbol{\Sigma}$ from the inverse-Wishart distribution $\mathcal{IW}(n + 10, 0.07\mathbf{I}_n + 0.03\mathbf{1}_n\mathbf{1}_n')$.

We estimate the model above by eliciting a normal prior for the VAR coefficients and an inverse-Wishart prior for the covariance matrix. More specifically, let $\boldsymbol{\beta} = \text{vec} \left([\mathbf{b}_0, \mathbf{B}_1, \mathbf{B}_2]' \right)$ denote the $k \times 1$ vector of VAR coefficients stacked by rows. We then specify uninformative independent priors for $\boldsymbol{\beta}$ and $\boldsymbol{\Sigma}$, i.e. $\boldsymbol{\beta} \sim \mathcal{N}(\mathbf{0}_k, \mathbf{I}_k)$ and $\boldsymbol{\Sigma} \sim \mathcal{IW}(3 + n, \mathbf{I}_n)$.

3.1. Equality constraints on observables

In the first set of simulations, we work with equality constraints on observables. We specify a forecast horizon of length h and constrain the entire out-of-sample paths of n_o endogenous variables while leaving those of the remaining $n - n_o$ variables unconstrained. That is, we set the forecasts of the n_o variables to be the actual simulated data $\mathbf{y}_{T:T+h}^o$. We consider VARs of different dimensions: a medium size VAR with a short forecast horizon ($n = 8, h = 5$), a large VAR with a longer forecast horizon ($n = 15, h = 20$), and an extra-large VAR with an even longer forecast horizon ($n = 40, h = 30$). In all three cases, we let the number of constrained endogenous variables, n_o , range from one to five and retain 25,000 MCMC draws, after discarding the first 10,000 draws.

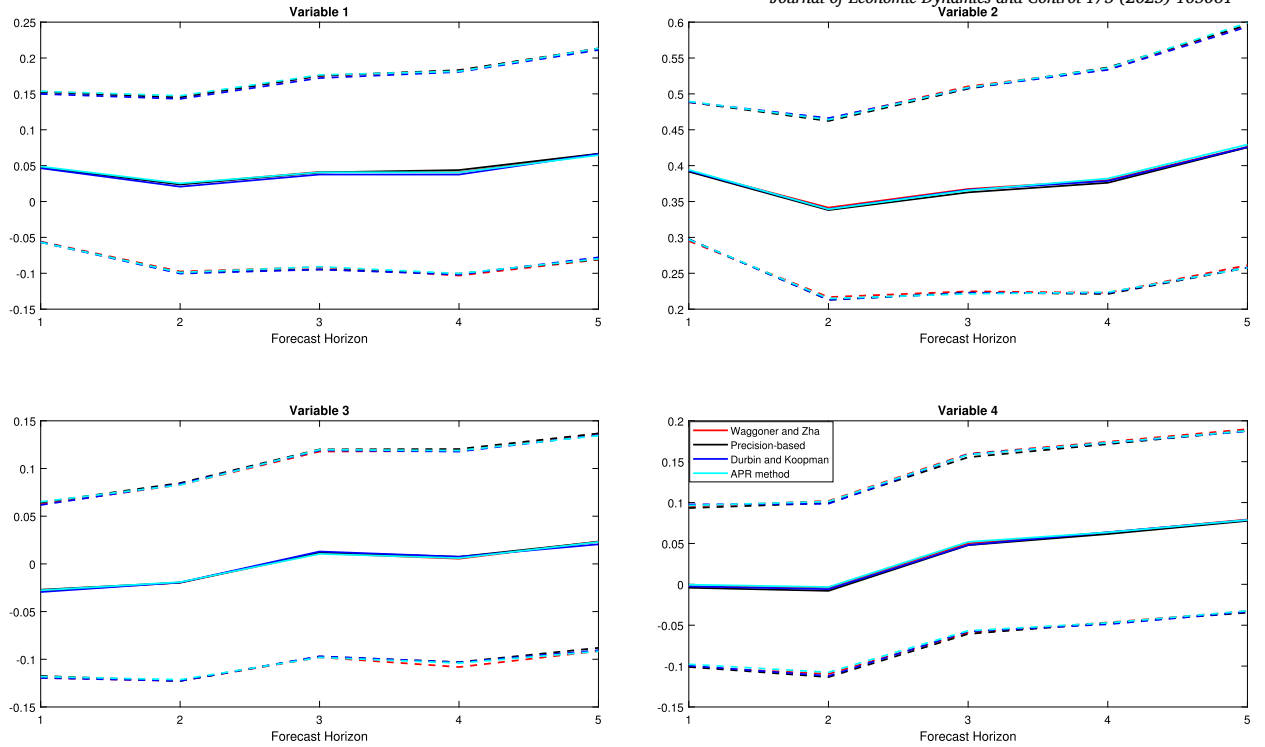


Fig. 1. Conditional forecasts from a medium VAR with three equality constraints. The thick black line is the posterior median estimate of the conditional forecast using our precision-based method. The thick red line is the posterior median estimates of the conditional forecast using the Waggoner and Zha (1999) approach. The thick dark blue line is the posterior median estimate of the conditional forecast from the approach of Bańbura et al. (2015) using the Durbin and Koopman (2002) smoother. The thick light blue line is the posterior median estimate of the conditional forecast using the Antolín-Díaz et al. (2021). The dashed lines are the corresponding 68% credible intervals for all four methods. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

We compare the proposed algorithms against the three conditional forecast algorithms in the literature: Waggoner and Zha (1999), Bańbura et al. (2015) and Antolín-Díaz et al. (2021). We note that although these Kalman filter based methods are implemented differently, they all aim to draw from the same Gaussian distribution. Consequently, the resulting samplers have the same mixing properties as ours, and it suffices to compare the computational speed. Focusing on accuracy first, Fig. 1 and Fig. 2 plot the conditional forecasts of the first four variables of both the medium and large VAR that we obtained from our method as well as the three benchmarks. We show both the median estimates (solid lines) as well as the associated 68 per cent credible intervals (dashed lines). As we can see from both figures, our proposed precision-based method produces virtually identical posterior conditional forecasts and credible sets as those of the other three benchmarks.

Having established that our approach produces conditional forecasts that are virtually identical to those obtained with the available methods in the literature, we turn to compare the computation time of the proposed precision-based conditional forecast sampler for the VAR models against the three other methods.¹³ We present these results in Table 1. The computation time reported is based on the computation cost of solely drawing the conditional forecasts across the four methods. We excluded the computation time of drawing the VAR parameters since these are the same across the four methods.¹⁴

Uniformly across all cases considered, our proposed precision-based method is computationally more efficient in producing conditional forecasts than the other three methods. We can also see that as the number of imposed equality constraints increases in the VAR, the proposed precision-based method becomes relatively more efficient in generating the conditional forecasts than the three other methods.

This implies that the proposed precision-based method is scalable in terms of the VAR dimensions and the number of imposed equality constraints. On the other hand, the conditional forecasting method proposed by Antolín-Díaz et al. (2021) does not ap-

¹³ For the Waggoner and Zha (1999) conditional forecasts, we use the approach of Jarociński (2010), who developed a more efficient algorithm to simulate the conditional forecast with one equality constraint compared to the original algorithm of Waggoner and Zha (1999).

¹⁴ For proper Bayesian inference in conditional forecasting, the usual Gibbs sampler for the VAR needs to be adjusted so that the conditions are taken into account when drawing the model parameters. While this point has been emphasized by Waggoner and Zha (1999), the bulk of the literature instead implements a two-step procedure: first estimate the model parameters using a standard Gibbs sampler; then, compute the conditional forecasts given the parameter estimates. We follow the latter approach in this paper and side-step the estimation issue, since our main contribution lies in improving the computation of conditional forecasts.

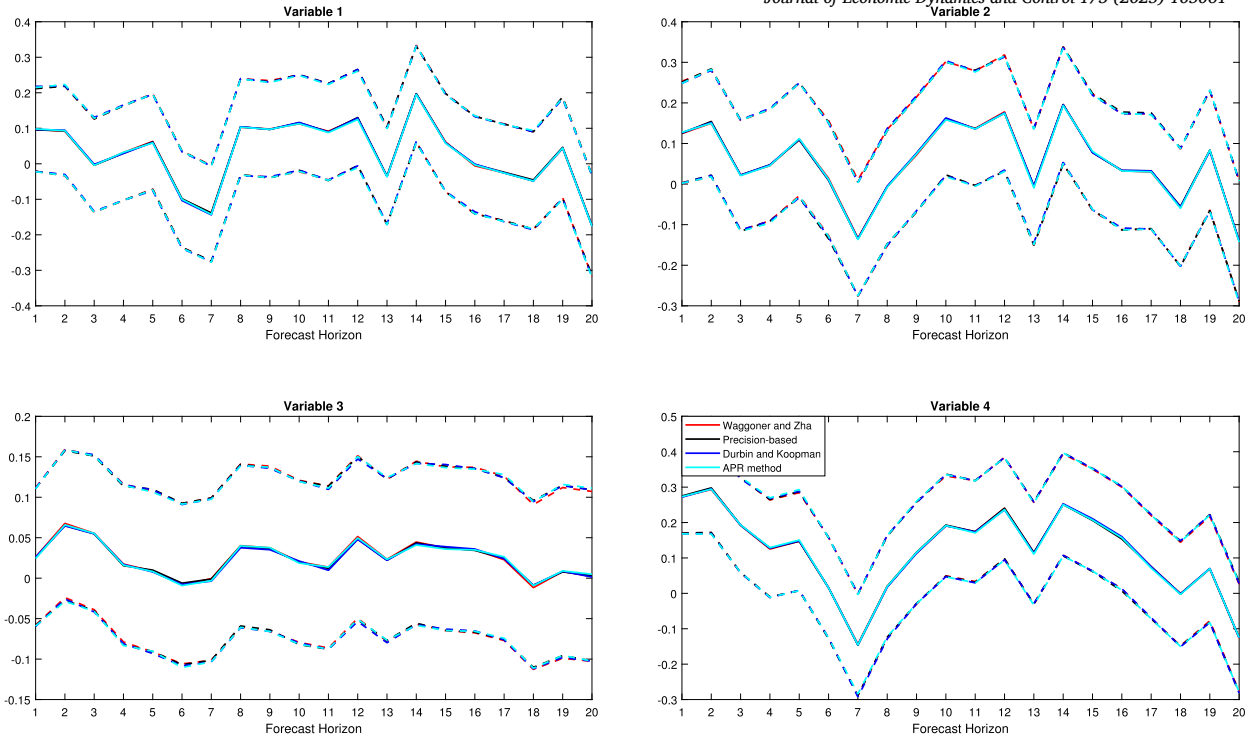


Fig. 2. Conditional forecasts from a large VAR with three equality constraints. The thick black line is the posterior median estimate of the conditional forecast using our precision-based method. The thick red line is the posterior median estimates of the conditional forecast using the Waggoner and Zha (1999) approach. The thick dark blue line is the posterior median estimate of the conditional forecast from the approach of Bańbura et al. (2015) using the Durbin and Koopman (2002) smoother. The thick light blue line is the posterior median estimate of the conditional forecast using the Antolín-Díaz et al. (2021). The dashed lines are the corresponding 68% credible intervals for all four methods.

pear to scale well, becoming computationally very intensive when the dimension of the VAR or the number of equality constraints increase.¹⁵

3.2. Inequality constraints on observables

Next, we move to the inequality-constraint case. We follow the same simulation structure as described above in the equality constraint case, except that we can only compare our conditional forecasts to those obtained using the Waggoner and Zha (1999) and Andersson et al. (2010) algorithms since both the approaches of Bańbura et al. (2015) and Antolín-Díaz et al. (2021) are not designed to handle inequality constraints. In terms of sampling efficiency, the proposed method has the same mixing properties as the approach in Waggoner and Zha (1999), as both aim to directly draw from the same multivariate truncated distribution. It is in principle more efficient than the algorithm in Andersson et al. (2010), since they implement a Gibbs sampling step to draw from the multivariate truncated distribution that induces additional autocorrelation in the MCMC draws. However, we find that the inefficiency factors of the posterior draws of the conditional forecasts for each of the three methods are very similar, and we therefore focus on comparing the computational speed. We set $\bar{\mathbf{y}}_{T-h:T}^o - 0.1 < \mathbf{y}_{T:T+h}^o < \bar{\mathbf{y}}_{T-h:T}^o + 0.1$, where $\bar{\mathbf{y}}_{T-h:T}^o = \frac{1}{h} \sum_{t=T-h+1}^T \mathbf{y}_t^o$ is the average of the actual simulated data over the periods $t = T - h, \dots, T$. In this simulation exercise, we focus on the medium- and large-sized VARs, where $n = 8$ and $n = 15$, and a long forecast horizon $h = 20$. As in the equality constraint case, we left the number of constrained variables, n_o , range from one to five. As before, we estimate the medium VAR using 25,000 MCMC draws with a burn-in period of 10,000 draws and implement the same priors as in the equality constraint exercise.

Fig. 3 plots the conditional forecasts for the first four variables from the medium-sized VAR with one inequality constraint. As before, our proposed precision-based method produces virtually identical posterior conditional forecast estimates and credible sets as both the Waggoner and Zha (1999) and Andersson et al. (2010) (denoted in the figure as APW) algorithms. However, as mentioned in the previous section, the Waggoner and Zha (1999) inequality constraint algorithm is very computationally intensive. For example, when working with a medium VAR and one inequality constraint, their accept-reject algorithm took approximately 2,542 minutes to

¹⁵ As discussed in Section 2.2, the proposed method parameterizes the conditional forecast distribution in terms of its precision matrix as opposed to the dense covariance matrix as in Antolín-Díaz et al. (2021). Furthermore, under the proposed method, simulations of the conditional forecasts given the equality constraints involve only primitive matrices such as \mathbf{M}_y and \mathbf{H} , which are banded. By contrast, Antolín-Díaz et al. (2021) simulate the conditional forecasts using the MATLAB built-in function `mvnrnd`, which requires additional computations such as computing the Cholesky decomposition or the singular value decomposition of the covariance matrix.

Table 1

Computation time for the equality constraints case. This table reports the computation time (in seconds) required to generate conditional forecasts from the precision-based sampler, the Waggoner and Zha (1999) (WZ), Bańbura et al. (2015) implemented using the smoother of Durbin and Koopman (2002) (DK) and Antolí-Díaz et al. (2021) (APR) methods. The computation times are based on 25,000 MCMC draws with a 10,000 burn-in period.

| Dimension | | Precision-based | WZ | DK | APR | |
|-------------------------|----------|-----------------|-----|------|------|-----|
| Model with $p = 2$ lags | | | | | | |
| Medium | $h = 5$ | $n_o = 1$ | 3 | 6 | 9 | 10 |
| | | $n_o = 3$ | 3 | 6 | 9 | 11 |
| | | $n_o = 5$ | 2 | 7 | 9 | 12 |
| Large | $h = 20$ | $n_o = 1$ | 24 | 71 | 50 | 141 |
| | | $n_o = 3$ | 23 | 73 | 50 | 147 |
| | | $n_o = 5$ | 22 | 88 | 51 | 174 |
| Extra Large | $h = 30$ | $n_o = 1$ | 338 | 1554 | 410 | - |
| | | $n_o = 3$ | 332 | 1865 | 405 | - |
| | | $n_o = 5$ | 291 | 1873 | 419 | - |
| Model with $p = 4$ lags | | | | | | |
| Medium | $h = 5$ | $n_o = 1$ | 4 | 7 | 10 | 15 |
| | | $n_o = 3$ | 3 | 8 | 11 | 15 |
| | | $n_o = 5$ | 3 | 9 | 10 | 17 |
| Large | $h = 20$ | $n_o = 1$ | 44 | 104 | 68 | 238 |
| | | $n_o = 3$ | 39 | 118 | 70 | 254 |
| | | $n_o = 5$ | 34 | 122 | 66 | 466 |
| Extra Large | $h = 30$ | $n_o = 1$ | 937 | 3066 | 1375 | - |
| | | $n_o = 3$ | 834 | 3340 | 1341 | - |
| | | $n_o = 5$ | 799 | 3340 | 1260 | - |

Table 2

Computation time for inequality constraints case. This table report the computation time (in seconds) required to simulate 1,000 draws for the precision-based sampler and the Andersson et al. (2010) (APW) method in the case of a medium VAR and a $h = 20$ forecast horizon.

| | No. of inequality constraints | | |
|-------------------------|-------------------------------|-----------|-----------|
| | $n_o = 1$ | $n_o = 3$ | $n_o = 5$ |
| Medium VAR with $p = 2$ | | | |
| Precision-based | 2 | 5 | 8 |
| APW | 3 | 9 | 21 |
| Large VAR with $p = 2$ | | | |
| Precision-based | 4 | 7 | 9 |
| APW | 4 | 9 | 22 |
| Medium VAR with $p = 4$ | | | |
| Precision-based | 3 | 5 | 9 |
| APW | 4 | 9 | 21 |
| Large VAR with $p = 4$ | | | |
| Precision-based | 7 | 9 | 11 |
| APW | 6 | 12 | 23 |

compute the conditional forecasts. In contrast, our proposed precision-based method and the Andersson et al. (2010) take about 62 and 70 seconds to obtain the corresponding conditional forecasts, respectively.

Note that while the computation costs of our precision-based and the Andersson et al. (2010) methods are relatively similar when focusing on a single inequality constraint, Table 2 shows that as the number of inequality constraints increase our precision-based method becomes more computationally efficient than the Andersson et al. (2010). This is true for both a VAR with $p = 2$ (top half of the table) and $p = 4$ lags (bottom half of the table).

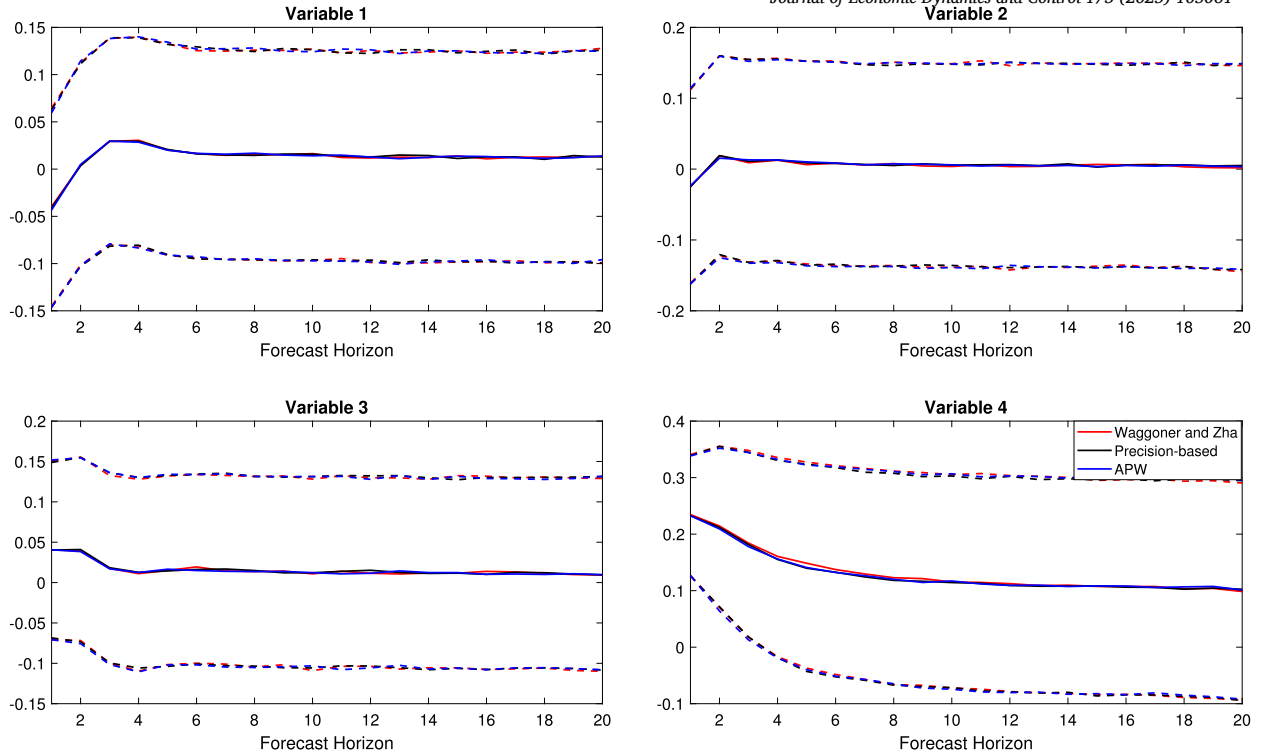


Fig. 3. Conditional forecast from a medium VAR with one inequality constraint. The thick black line is the posterior median estimate of the conditional forecast using the proposed precision-based method. The thick red line is the posterior median estimates of the conditional forecast using the Waggoner and Zha (1999) method. The thick blue line is the posterior median estimates of the conditional forecasting using the Andersson et al. (2010) (APW) method. The dash lines are the corresponding 68 per cent credible intervals for all three methods.

4. Empirical application

To illustrate the performance of our proposed precision-based conditional forecast method in an empirical setting, we next setup a large BVAR with 31 quarterly variables to model the dynamic evolution over time of the US macroeconomic and financial sectors. Our analysis is inspired by the recent work of Crump et al. (2021), who estimate a quarterly BVAR with variables extracted from those included in the Federal Reserve Board of Governor’s Tealbook A Greensheets and selected to provide a broad picture the US macroeconomic and financial conditions. They then consider a number of conditional forecasting and tilting exercises to produce counterfactual scenarios and evaluate their impact on the US economy both retrospectively and prospectively.¹⁶ We exploit the computational efficiency and generality of the methods presented in the previous sections to expand on the analysis of Crump et al. (2021) and investigate the macroeconomic impact of a combination of multiple inequality and equality constraints at once. The complete list of the variables used along with their transformations is available in Appendix A. To the best of our knowledge, this is the first study within the literature that considers conditional forecasting in a large VAR setting with multiple equality and inequality constraints.¹⁷

Our estimation strategy for the VAR parameters follows closely Chan (2022), who proposes a novel asymmetric natural conjugate Minnesota-type prior for large VARs. Unlike the more traditional natural conjugate prior, which does not allow for asymmetric shrinkage of own lags and lags of other variables, this new prior is more flexible and can be used, for example, to enforce stronger shrinkage on the other variables’ lags. At the same time, this prior maintains many useful analytical results of the traditional conjugate prior, such as an extremely fast and scalable posterior simulator and a closed-form expression of the marginal likelihood. As in subsection 2.1, we start with a VAR(p) in structural form,

$$\mathbf{A}_0 \mathbf{y}_t = \mathbf{a} + \mathbf{A}_1 \mathbf{y}_{t-1} + \dots + \mathbf{A}_p \mathbf{y}_{t-p} + \boldsymbol{\varepsilon}_t, \quad \boldsymbol{\varepsilon}_t \sim \mathcal{N}(\mathbf{0}_n, \boldsymbol{\Sigma}), \tag{22}$$

¹⁶ The Federal Reserve Tealbook A, officially subtitled “Economic and Financial Conditions: Current Situation and Outlook”, is produced by the staff of the Board of Governors and provides in-depth analysis and forecasts of the U.S. and international economy. See <https://www.federalreserve.gov/monetarypolicy> for additional details on the variables being tracked.

¹⁷ Crump et al. (2021) employ a filtering-based method to generate their conditional forecasts, simulating the conditional forecasts one period at a time. In contrast, our proposed precision-based method allows us to simulate all the conditional forecasts in one step and as we showed in Section 3 leads to large computational savings.

Table 3
Optimal shrinkage. This table reports the values of the hyperparameters under the symmetric natural conjugate prior and the asymmetric conjugate prior.

| | Symmetric Prior | Asymmetric Prior |
|------------|-----------------|------------------|
| κ_1 | 0.0045 | 0.083 |
| κ_2 | 0.0045 | 0.0024 |
| Log-ML | -6121.4 | -6078.4 |

where \mathbf{A}_0 is a lower triangular matrix with ones on its main diagonal and $\mathbf{\Sigma} = \text{diag}(\sigma_1^2, \dots, \sigma_n^2)$ is diagonal. This, in turn, allows us to estimate the n -dimensional VAR above in a recursive manner, one equation at a time. For this purpose, let a_i denote the i -th element of \mathbf{a} and let $\mathbf{a}_{j,i}$ represent the i -th row of \mathbf{A}_j . In addition, let α_i denote the $i - 1$ free elements from the i -th row of \mathbf{A}_0 , and define $\beta_i = (a_i, \mathbf{a}_{1,i}, \dots, \mathbf{a}_{p,i})'$. With this notation in hand, we can rewrite the i -th equation of (22) as follows:

$$y_{it} = \mathbf{w}_{it}\alpha_i + \mathbf{z}_{it}\beta_i + \varepsilon_{it} \tag{23}$$

where $\mathbf{w}_{it} = (-y_{1t}, \dots, -y_{i-1,t})$, $\mathbf{z}_{it} = (1, \mathbf{y}'_{i-1}, \dots, \mathbf{y}'_{i-p})$ and $\varepsilon_{it} \sim \mathcal{N}(0, \sigma_i^2)$. We pair the model in (23) with the following normal-inverse-gamma prior for $(\alpha_i, \beta_i, \sigma_i^2)$,

$$\begin{aligned} \alpha_i | \sigma_i^2 &\sim \mathcal{N}(\mathbf{m}_i^\alpha, \sigma_i^2 \mathbf{V}_i^\alpha) \\ \beta_i | \sigma_i^2 &\sim \mathcal{N}(\mathbf{m}_i^\beta, \sigma_i^2 \mathbf{V}_i^\beta) \end{aligned} \tag{24}$$

and

$$\sigma_i^2 \sim IG\left(\frac{v_0 + i - n}{2}, \frac{s_i}{2}\right), \tag{25}$$

where s_i^2 denotes the sample variance of the residuals from an AR(p) model estimated on variable i , $i = 1, \dots, n$. We specify the prior hyperparameters by mirroring the choices in Chan (2022). In particular, we set $v_0 = n + 2$, $\mathbf{m}_i^\alpha = \mathbf{0}_{i-1}$ and $\mathbf{V}_i^\alpha = \text{diag}(1/s_1^2, \dots, 1/s_{i-1}^2)$. Next, we set all elements of \mathbf{m}_i^β to zero, except for the coefficient associated with the first own lag, which we set to one. As for \mathbf{V}_i^β , we specify it to be a diagonal matrix with the k -th diagonal element, $(\mathbf{V}_i^\beta)_k$, inspired by the Minnesota prior,

$$(\mathbf{V}_i^\beta)_k = \begin{cases} \frac{\kappa_1}{l^2 s_i^2} & \text{for the coefficient on the } l\text{th lag of variable } i \\ \frac{\kappa_2}{l^2 s_j^2} & \text{for the coefficient on the } l\text{th lag of variable } j, \quad j \neq i \\ 100 & \text{for the intercept} \end{cases} \tag{26}$$

and where the hyperparameters κ_1 and κ_2 control the overall shrinkage intensity for the coefficients on the own lags and other variables' lags, respectively. As in Chan (2022), we estimate the optimal degrees of shrinkage from the data.

As a first step, we estimate the model in (22) on data ranging from 1976Q3 to 2019Q3 (this is the same sample Crump et al. (2021) used to explore a number of their scenarios), after setting the lag length of the VAR to $p = 4$.¹⁸ Before turning to the conditional forecasts, we quickly review the estimation results and in particular we focus on whether the asymmetric conjugate priors we employ allow to fit the data better than the traditional natural conjugate priors (the latter is the prior adopted by Crump et al. (2021)). Table 3 reports the results of this comparison, showing both the posterior means of the κ_1 and κ_2 hyperparameters (in the case of the natural conjugate prior, with symmetric shrinkage on the own and other variables' lags, $\kappa_1 = \kappa_2$), as well as the marginal likelihoods. Under the symmetric prior, the optimal hyperparameter value is 0.0045, while when we allow the two hyperparameters to differ the results are quite different, with the data clearly pointing to a much stronger shrinkage on the coefficients of the other variables' lags. In addition, allowing for asymmetric shrinkage increases the marginal log-likelihood by about 43 points. Taken all together, these results overwhelmingly point to the data favoring the asymmetric prior specification.

4.1. Conditional forecasts

Next, we shift our attention to the model's conditional forecasts. To showcase the generality and scalability of our approach, we focus on producing conditional forecasts in the presence of multiple equality and inequality constraints. We re-estimate the BVAR above on data up to 2019Q4 and begin by setting a number of equality constraints on the future paths of unemployment rate and the 10-year Treasury rate, mirroring the projections reported by the Federal Reserve Board in their 2020 stress tests. These tests are known as the Dodd-Frank Act stress test (DFAST) and the Comprehensive Capital Analysis and Review (CCAR) and are used on an annual basis to ensure that large bank holding companies operating in the United States will be able to lend to households

¹⁸ As in Crump et al. (2021), rates enter the VAR in levels, while most other variables enter in log-levels. See Appendix A for the full list of all variable transformations.

Table 4

Summary of equality and inequality constraints. This table reports the full set of inequality and equality constraints over the 2020Q1–2023Q1 time period that we use to generate the BVAR conditional forecasts. We consider two different scenarios from the Federal Reserve Board stress tests, namely (i) Baseline and (ii) Severely Adverse.

| Date | CPI Inflation (Inequality Constraints) | | | Equality Constraints | |
|--------------------------|--|------------------|-------------|----------------------|------|
| | Lower Bound | Fed's projection | Upper Bound | UNRATE | GS10 |
| Baseline Scenario | | | | | |
| 2020Q1 | 1.69 | 2.20 | 2.71 | 3.60 | 1.80 |
| 2020Q2 | 1.55 | 2.10 | 2.65 | 3.60 | 1.90 |
| 2020Q3 | 1.58 | 2.00 | 2.42 | 3.60 | 1.90 |
| 2020Q4 | 1.47 | 1.90 | 2.33 | 3.70 | 2.00 |
| 2021Q1 | 1.57 | 2.10 | 2.63 | 3.70 | 2.00 |
| 2021Q2 | 1.57 | 2.10 | 2.63 | 3.70 | 2.10 |
| 2021Q3 | 1.57 | 2.10 | 2.63 | 3.80 | 2.10 |
| 2021Q4 | 1.57 | 2.10 | 2.63 | 3.80 | 2.20 |
| 2022Q1 | 1.77 | 2.30 | 2.83 | 3.90 | 2.20 |
| 2022Q2 | 1.67 | 2.20 | 2.73 | 3.90 | 2.40 |
| 2022Q3 | 1.67 | 2.20 | 2.73 | 3.90 | 2.50 |
| 2022Q4 | 1.67 | 2.20 | 2.73 | 3.90 | 2.60 |
| 2023Q1 | 1.67 | 2.20 | 2.73 | 3.90 | 2.70 |
| Adverse Scenario | | | | | |
| 2020Q1 | 1.19 | 1.70 | 2.21 | 4.50 | 0.70 |
| 2020Q2 | 0.55 | 1.10 | 1.65 | 6.10 | 0.90 |
| 2020Q3 | 0.58 | 1.00 | 1.42 | 7.40 | 1.00 |
| 2020Q4 | 0.67 | 1.10 | 1.53 | 8.40 | 1.10 |
| 2021Q1 | 0.77 | 1.30 | 1.83 | 9.20 | 1.20 |
| 2021Q2 | 0.87 | 1.40 | 1.93 | 9.70 | 1.30 |
| 2021Q3 | 0.97 | 1.50 | 2.03 | 10.00 | 1.40 |
| 2021Q4 | 1.17 | 1.70 | 2.23 | 9.90 | 1.50 |
| 2022Q1 | 1.27 | 1.80 | 2.33 | 9.70 | 1.60 |
| 2022Q2 | 1.27 | 1.80 | 2.33 | 9.50 | 1.80 |
| 2022Q3 | 1.27 | 1.80 | 2.33 | 9.20 | 1.90 |
| 2022Q4 | 1.27 | 1.80 | 2.33 | 8.80 | 2.10 |
| 2023Q1 | 1.17 | 1.70 | 2.23 | 8.50 | 2.20 |

and businesses even in a severe recession.¹⁹ We consider two scenarios, namely (i) Baseline and (ii) Severely Adverse and restrict our attention to the 13 quarter period following the end of the estimation sample, i.e. 2020Q1 through 2023Q1. As noted in the Fed publications, the 2020 baseline scenario is a moderate economic expansion over the 13-quarter stress-test period and mimic closely the January 2020 consensus projections from Blue Chip Economic Indicators. The 2020 severely adverse scenario is instead characterized by a severe global recession accompanied by a period of heightened stress in commercial real estate and corporate debt markets, and is designed to assess the strength of banking organizations and their resilience to unfavorable economic conditions. We complement the 26 equality constraints described above with 13 more inequality constraints on the path of CPI inflation over the same out-of-sample period. To design these constraints, we once again start from the baseline and severely adverse scenarios provided by the Federal Reserve Board stress tests and complement these with a measure of disagreement that we extract from the Survey of Professional Forecasters (SPF). The Federal Reserve of Philadelphia reports the cross-sectional dispersion, defined as the difference between the 75th percentile and the 25th percentile of the SPF projections, for almost all variables included in their surveys for each quarterly vintage, going back to the beginning of 1968.²⁰ More specifically, for each quarter past the end of our estimation sample, we use the difference between the P75 and P25 of the SPF forecasts for CPI inflation to create an offset to the stress test projections, such that our inequality constraints for CPI inflation are centered on the FED stress test numbers and have a range that reflects the P25-P75 dispersion we see in the SPF forecasts. Two clarifications are in order. First, since the SPF does not provide projections under different scenarios, we use the same measure of disagreement in the baseline and severely adverse scenarios. Second, since the SPF cross-sectional dispersion is only available for the first five quarters following the end of the sample, we fix the inequality constraint intervals from the sixth quarter (i.e., from 2021Q2) onward to match the forecast dispersion we observed in 2021Q1.²¹ Table 4 summarizes the inequality and equality constraints for the baseline and adversely severe scenarios.

¹⁹ The Federal Reserve Board annual stress tests from 2013 onward can be found at <https://www.federalreserve.gov/publications/dodd-frank-act-stress-test-publications.htm>.

²⁰ The cross-sectional dispersion from the survey of professional forecasters is available at <https://www.philadelphiafed.org/surveys-and-data/real-time-data-research/dispersion-forecasts>.

²¹ As a sensitivity analysis, we also experimented with the idea of gradually widening the inequality constraint intervals around the corresponding FED stress test numbers from 2021Q2 onward as we progress through the forecast horizon, as a way of modeling the higher degree of uncertainty that could potentially arise on future inflation. The results from this alternative scenario are very similar to the ones we have reported.

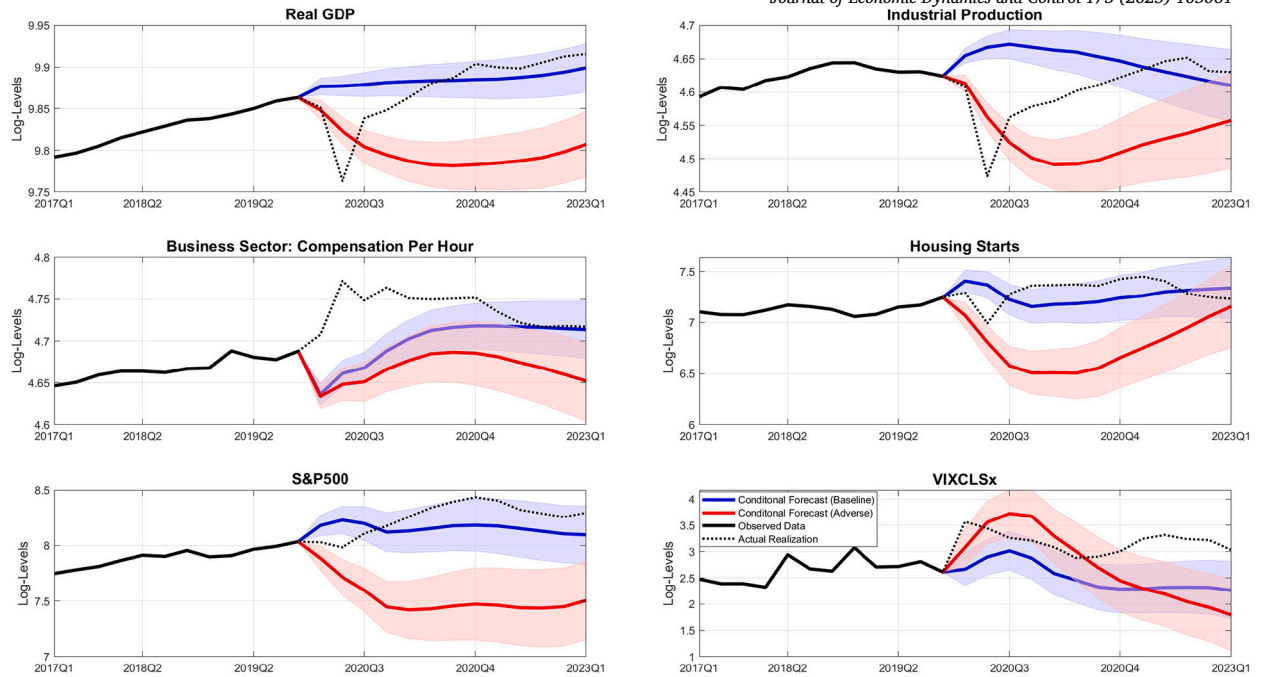


Fig. 4. Conditional forecasts when CPI inflation, unemployment, and the 10-year Treasury rate in 2020Q1–2023Q1 match the Fed’s stress test Baseline and Severely Adverse projections. The thick blue line represents the conditional forecasts under the Baseline scenario, while the thick red line represents the conditional forecasts under the Severely Adverse scenario. Shaded bands correspond to the associated 68% coverage intervals, while the dotted black lines denote the in-sample and out-of-sample values.

Fig. 4 present the results of this exercise for six key variables, namely (i) Real GDP, (ii) Industrial Production, (iii) Hourly compensation in the business sector, (iv) Housing starts, (v) S&P 500 index, and (vi) CBOE Volatility index. Within each panel, we show the posterior means of both the Baseline and Severely Adverse conditional forecasts, along with their 68% coverage intervals. Not surprisingly, the two sets of conditional forecasts produce markedly different paths for the six variables in question, in line with what one would have expected. Conditioning on the Fed’s adverse scenario projections for CPI inflation, unemployment and the 10-year Treasury rate leads to a large decrease in real GDP and industrial production, bottoming up around the second half of 2021 before a gradual recovery begins. Similarly, we see a significant drop in housing starts and the S&P 500 and a spike in the market volatility, as picked up the by VIX measure. Within each panel we also overlay, with black solid and dotted lines, the actual realization of each series, both in-sample and out-of-sample. This is mainly to provide a reference point for how the economy actually evolved over the particularly tumultuous period that followed the initial COVID-19 shock and its unfolding. Finally, to further help quantifying the differences between the two sets of conditional forecasts, Fig. 5 translates the results shown in Fig. 4 into percentage points (relative to the baseline scenario paths), displaying both posterior means and the 68% coverage intervals for these differences.

4.2. Out-of-sample analysis

In the previous section, we have shown how our proposed approach can be used to efficiently combine multiple equality and inequality constraints within a large BVAR. In this section, we narrow our focus to the inequality constraints, as implemented with the approach described in equation (12). While we have discussed earlier how in some settings the inequality constraints can be more practical and advantageous to use (relative to exact constraints or equality constraints), we now turn our attention to investigate whether inequality constraints can be leveraged to produce more accurate conditional forecasts. We continue to work with the 31 variable VAR described in the previous section and carry out a pseudo real-time forecasting exercise where for each quarter between 1995Q1 and 2022Q3 we re-estimate the BVAR parameters in (22) using all data up to that point in time and use them to produce both unconditional and conditional forecasts for up to four quarters ahead. As for the conditional forecasts, we implement them by constraining the future path of (the log of) Real GDP to align with the SPF projections that would have been available at each forecast date within our evaluation sample. More specifically, we implement the inequality constraints using the same approach adopted in the previous section, i.e. setting \underline{c} and \bar{c} in (12) at the 25th and 75th percentiles of Real GDP’s SPF cross sectional dispersion. We then consider three versions of equality constraint forecasts, as implemented in (5). In all three cases, we set \mathbf{r} at the median of the individual SPF Real GDP’s forecasts.²² As for the covariance matrix $\mathbf{\Omega}$, we consider three scenarios of the form $\mathbf{\Omega} = \omega \times \mathbf{I}_h$. Our first choice uses $\omega = 0.66$, a value we chose so that, on average, the inter-quantile range of the normal distribution used in (5) aligns with the P25-P75

²² Results produced using the mean instead of the median of Real GDP’s individual forecasts are almost identical.

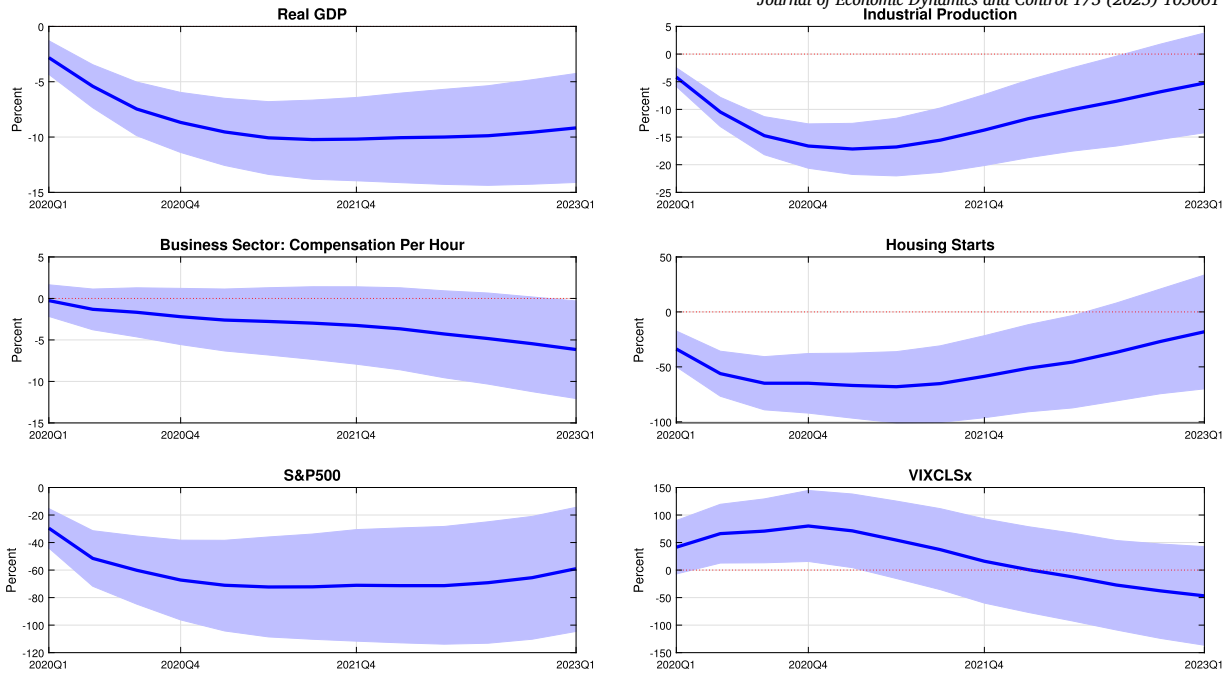


Fig. 5. Percentage difference between the two sets of conditional forecasts (Severely Adverse minus Baseline). The shaded bands represent the 68% coverage intervals.

range of the SPF Real GDP’s cross sectional dispersion. We label this our Middle case. Our two other alternatives are built off the Middle case so that we can properly assess the impact that a different Ω plays on the accuracy of the conditional forecasts. In particular, our first alternative implements narrower distributions, using $\omega = 0.2$, while the second one enforces wider distributions, using $\omega = 2$.

Table 5 reports the results of this exercise for forecast horizons ranging from one to four quarters ahead, and focusing on six key variables, namely Real Consumption Growth (PCECC96), PCE Inflation (PCECTPI), Non-farm employment growth (PAYEMS), Industrial Production (INDPRO), Capacity Utilization (CUMFNS) and Fed Funds Rate (FEDFUNDS).²³ All results are relative to the unconditional forecast benchmark. In particular, columns two to five focus on point forecast accuracy, showing the relative Root Mean Squared Forecast Error (RMSFE) of both the inequality (INEQ) and equality (EQ) constraint conditional forecasts. Columns six to nine report the accuracy of the probabilistic forecasts, assessed through the relative Continuous Ranked Probability Score (CRPS). In all instances, numbers smaller than one identify situations when the conditional forecasts improve upon the unconditional ones. In addition, stars next to the table entries identify differences that are statistically significant, calculated using the Diebold and Mariano (1995) t-statistic and standard normal critical values, with one, two, or three stars correspond to significance levels of 10%, 5%, and 1%, respectively.²⁴ Inspecting the table, the story that emerges is one where the inequality constraint forecasts dominate, and this is true both in terms of point and density forecast accuracy. Starting with the relative RMSFEs, it is remarkable how the inequality constraint forecasts are the race winner in all 24 instances considered (six variables, four forecast horizons). In addition, except for Real Consumption Growth at the longer horizons, these differences are statistically significant and often quite large. Particularly noteworthy is the fact that the improvements relative to the Middle equality forecast case, which we remind here was setup in a way to resemble the support of the inequality constraint forecasts, are in most cases quite large. This is particularly true for Real Consumption Growth, Employment, Industrial Production, Capacity Utilization, and the Federal fund rate. The story we see when moving on to the CRPS metric is a very similar one. In all but one instance, the inequality constraints lead to the most accurate density forecasts, and in all but a handful of cases the improvements over the unconditional forecast benchmark are statistically significant. We also continue to see large gains relative to the Middle scenario of the equality forecasts.

5. Conclusions

We introduced a novel precision-based approach that can be used for conditional forecasting, scenario analysis and entropic tilting and can handle both equality and inequality constraints. Thanks to the way we have derived the conditional forecasts’ distribution, our approach is computationally very efficient and particularly well suited to handle large dimensional VARs as well as situations in which we have a large number of conditioning variables and long forecast horizons. We have showed in a simulation study that

²³ As a robustness check, to rule out that the results are driven by the most recent period, which includes the onset of the COVID pandemic and the ensuing recovery, Table B.1 in Appendix B reports the results when we stop right before the beginning of the COVID pandemic.

²⁴ The underlying p-values are based on t-statistics computed with a serial correlation-robust variance, using a rectangular kernel, $h - 1$ lags, and the small-sample adjustment of Harvey et al. (1997).

Table 5

RMSFE and CRPS statistics for selected variables, relative to the BVAR unconditional forecast benchmark. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively, for the Diebold-Mariano test. The evaluation sample is 1995Q1 through 2022Q3. Boldfaced entries denote, within each row and for each statistics, the best performing model.

| Forc. <i>h</i> | RMSFE | | | | CRPS | | | |
|--|----------------|-------------|-------------|-----------|----------------|-------------|-------------|-----------|
| | INEQ | EQ (Middle) | EQ (Narrow) | EQ (Wide) | INEQ | EQ (Middle) | EQ (Narrow) | EQ (Wide) |
| Real Consumption Growth (PCECC96) | | | | | | | | |
| <i>h</i> = 1 | 0.55*** | 0.82*** | 0.90*** | 0.75*** | 0.32*** | 0.75*** | 0.88*** | 0.59*** |
| <i>h</i> = 2 | 0.76** | 0.88 | 0.92* | 0.89 | 0.58*** | 0.71*** | 0.85*** | 0.64*** |
| <i>h</i> = 3 | 0.94 | 1.05 | 1.04 | 1.05 | 0.91 | 0.90 | 0.93 | 0.91 |
| <i>h</i> = 4 | 0.97 | 1.13 | 1.12 | 1.11 | 1.03 | 1.04 | 1.02 | 1.06 |
| PCE Inflation (PCECTPI) | | | | | | | | |
| <i>h</i> = 1 | 0.94*** | 0.99*** | 0.99*** | 0.97*** | 0.93*** | 0.99*** | 0.99*** | 0.97*** |
| <i>h</i> = 2 | 0.92*** | 0.97*** | 0.99** | 0.96*** | 0.90*** | 0.97*** | 0.99** | 0.95*** |
| <i>h</i> = 3 | 0.90*** | 0.96** | 0.98** | 0.94** | 0.87*** | 0.95*** | 0.98** | 0.92*** |
| <i>h</i> = 4 | 0.92*** | 0.96** | 0.99 | 0.95** | 0.89*** | 0.94*** | 0.98** | 0.92*** |
| Non-farm employment growth (PAYEMS) | | | | | | | | |
| <i>h</i> = 1 | 0.75*** | 0.90** | 0.94** | 0.87** | 0.65*** | 0.88*** | 0.94*** | 0.82*** |
| <i>h</i> = 2 | 0.73*** | 0.93 | 0.94 | 0.95 | 0.51*** | 0.84** | 0.92*** | 0.75*** |
| <i>h</i> = 3 | 0.77** | 1.04 | 1.02 | 1.06 | 0.55*** | 0.82* | 0.92 | 0.72** |
| <i>h</i> = 4 | 0.83* | 1.18 | 1.15 | 1.17 | 0.65*** | 0.85 | 0.94 | 0.78 |
| Industrial Production (INDPRO) | | | | | | | | |
| <i>h</i> = 1 | 0.63*** | 0.90*** | 0.95*** | 0.84*** | 0.53*** | 0.88*** | 0.94*** | 0.79*** |
| <i>h</i> = 2 | 0.62*** | 0.85** | 0.92*** | 0.80** | 0.49*** | 0.79*** | 0.90*** | 0.67*** |
| <i>h</i> = 3 | 0.78** | 0.93 | 0.96 | 0.93 | 0.69*** | 0.81** | 0.90* | 0.76** |
| <i>h</i> = 4 | 0.89 | 1.07 | 1.06 | 1.05 | 0.90 | 0.96 | 0.97 | 0.96 |
| Capacity Utilization (CUMFNS) | | | | | | | | |
| <i>h</i> = 1 | 0.64*** | 0.89*** | 0.95*** | 0.83*** | 0.53*** | 0.87*** | 0.93*** | 0.78*** |
| <i>h</i> = 2 | 0.59*** | 0.85** | 0.92** | 0.80* | 0.45*** | 0.77*** | 0.89*** | 0.64*** |
| <i>h</i> = 3 | 0.70*** | 0.92 | 0.96 | 0.91 | 0.59*** | 0.77** | 0.88** | 0.69*** |
| <i>h</i> = 4 | 0.77*** | 1.04 | 1.06 | 0.99 | 0.73*** | 0.88 | 0.94 | 0.84* |
| Fed Funds Rate (FEDFUNDS) | | | | | | | | |
| <i>h</i> = 1 | 0.82*** | 0.96*** | 0.99*** | 0.93*** | 0.77*** | 0.96*** | 0.98*** | 0.91*** |
| <i>h</i> = 2 | 0.66*** | 0.89*** | 0.95*** | 0.82*** | 0.58*** | 0.87*** | 0.94*** | 0.78*** |
| <i>h</i> = 3 | 0.75*** | 0.88*** | 0.94*** | 0.83*** | 0.68*** | 0.84*** | 0.92*** | 0.76*** |
| <i>h</i> = 4 | 0.77*** | 0.88** | 0.94** | 0.85** | 0.69*** | 0.81*** | 0.90*** | 0.75*** |

the proposed approach generates exactly the same conditional forecasts and credible sets as those from Waggoner and Zha (1999), Bańbura et al. (2015), and Antolín-Díaz et al. (2021), but is substantially less demanding computationally. Finally, we conducted two empirical exercises to highlight the flexibility and usefulness of the method. In the first application, we estimated a Bayesian VAR featuring 31 quarterly macroeconomic and financial series, and we used our approach to investigate the effect of simultaneously imposing a number of inequality and equality constraints on the trajectories of CPI inflation, the unemployment rate, and the 10-year Treasury rate over the 2020–2022 period. In our second application, we performed a forecasting exercise where, using the same 31 variable VAR, we separately generated pseudo real-time conditional forecasts with either equality or inequality constraints built off the SPF Real GDP projections and found that conditional forecasts constructed using the inequality constraints were on average more accurate than their equality constraint counterparts. In the future, we believe it would be useful to generalize the current approach to also work with Bayesian VARs with time-varying parameters and stochastic volatility, and it would also be interesting to extend the precision-based sampler to produce conditional forecasts for binary variables, along the lines of McCracken et al. (2021).

Appendix A. Data and transformations

See Table A.1.

Table A.1

This table lists all the variables in our empirical application. For each series, we report the FRED-QD mnemonic and the transformation we adopt. We extract all variables from the FRED-QD dataset, available at <https://research.stlouisfed.org/econ/mccracken/fred-databases/>. The list of variables and transformations follow closely the choices from Crump et al. (2021).

| Variable | Mnemonic | Transformation |
|--|----------|----------------|
| Real Gross Domestic Product | GDPC1 | 100ln(x_t) |
| Real Personal Consumption Expenditures | PCECC96 | 100ln(x_t) |
| Real Private Fixed Investment: Residential | PRFIx | 100ln(x_t) |

Table A.1 (continued)

| Variable | Mnemonic | Transformation |
|---|-----------------|----------------|
| Real Private Fixed Investment: Non-Residential | PNFIx | 100ln(x_t) |
| Real Exports of Goods and Services | EXPGSC1 | 100ln(x_t) |
| Real Imports of Goods and Services | IMPGSC1 | 100ln(x_t) |
| Real Government Consumption Expenditures and Gross Investment | GCEC1 | 100ln(x_t) |
| Real Government Consumption Expenditures and Gross Investment: Federal (Chain-Type Quantity Index) | B823RA3Q086SBEA | 100ln(x_t) |
| Gross Domestic Product: Chain-Type Price Index | GDPCTPI | 100ln(x_t) |
| Producer Price Index by Commodity: All Commodities | PPIACO | 100ln(x_t) |
| Personal Consumption Expenditures Excluding Food and Energy (Chain-Type Price Index) | PCEPILFE | 100ln(x_t) |
| Consumer Price Index | CPIAUCSL | 100ln(x_t) |
| Consumer Price Index for All Urban Consumers: All Items Less Food and Energy in U.S. City Average | CPILFESL | 100ln(x_t) |
| Business Sector: Real Hourly Compensation for All Employed Persons | RCPHBS | 100ln(x_t) |
| All Employees, Total Nonfarm | PAYEMS | 100ln(x_t) |
| Unemployment Rate | UNRATE | Level |
| Industrial Production: Total Index | INDPRO | 100ln(x_t) |
| Capacity Utilization: Manufacturing | CUMFNS | 100ln(x_t) |
| New Privately-Owned Housing Units Started: Total Units | HOUST | 100ln(x_t) |
| Real Disposable Personal Income | DPIC96 | 100ln(x_t) |
| University of Michigan: Consumer Sentiment | UMCSENTx | Level |
| Market Yield on U.S. Treasury Securities at 1-Year Constant Maturity | GS1 | Level |
| Market Yield on U.S. Treasury Securities at 10-Year Constant Maturity | GS10 | Level |
| Moody's Seasoned Aaa Corporate Bond Yield | AAA | Level |
| Moody's Seasoned Baa Corporate Bond Yield | BAA | Level |
| Trade Weighted U.S. Dollar Index: Advanced Foreign Currencies | TWEXAFEGSMTHx | 100ln(x_t) |
| S&P 500 | S&P 500 | 100ln(x_t) |
| CBOE Volatility Index: VIX | VIXCLSx | ln(x_t) |
| Personal Consumption Expenditures: Chain-type Price Index | PCECTPI | 100ln(x_t) |
| Real Crude Oil Prices: West Texas Intermediate (WTI) | OILPRICEx | 100ln(x_t) |
| Federal Funds Effective Rate | FEDFUNDS | Level |

Appendix B. Additional results

Table B.1

RMSFE and CRPS statistics for selected variables, relative to the BVAR unconditional forecast benchmark. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively, for the Diebold-Mariano test. The evaluation sample is 1995Q1 through 2019Q4.

| Forc. h | RMSFE | | | | CRPS | | | |
|----------------|----------------|----------------|-------------|-----------|----------------|-------------|----------------|-------------|
| | INEQ | EQ (Middle) | EQ (Narrow) | EQ (Wide) | INEQ | EQ (Middle) | EQ (Narrow) | EQ (Wide) |
| PCECC96 | | | | | | | | |
| $h = 1$ | 0.33*** | 0.76*** | 0.89*** | 0.61*** | 0.25*** | 0.71*** | 0.87*** | 0.52*** |
| $h = 2$ | 0.54*** | 0.66*** | 0.82*** | 0.56*** | 0.48*** | 0.60*** | 0.79*** | 0.49*** |
| $h = 3$ | 0.83** | 0.79*** | 0.86*** | 0.78** | 0.83 | 0.76*** | 0.83*** | 0.77** |
| $h = 4$ | 0.93 | 0.92 | 0.94** | 0.92 | 0.99 | 0.92 | 0.92* | 0.95 |
| PCECTPI | | | | | | | | |
| $h = 1$ | 0.95*** | 0.99*** | 1.00 | 0.98*** | 0.93*** | 0.99*** | 0.99*** | 0.97*** |
| $h = 2$ | 0.92*** | 0.97*** | 0.99*** | 0.95*** | 0.90*** | 0.97*** | 0.99*** | 0.94*** |
| $h = 3$ | 0.90*** | 0.95*** | 0.98*** | 0.93*** | 0.87*** | 0.94*** | 0.97*** | 0.91*** |
| $h = 4$ | 0.91*** | 0.95*** | 0.98*** | 0.93*** | 0.89*** | 0.94*** | 0.97*** | 0.91*** |
| PAYEMS | | | | | | | | |
| $h = 1$ | 0.72*** | 0.92*** | 0.97*** | 0.87*** | 0.66*** | 0.91*** | 0.96*** | 0.84*** |
| $h = 2$ | 0.54*** | 0.81*** | 0.91*** | 0.70*** | 0.45*** | 0.79*** | 0.90*** | 0.65*** |
| $h = 3$ | 0.54*** | 0.71*** | 0.85*** | 0.60*** | 0.46*** | 0.69*** | 0.84*** | 0.55*** |
| $h = 4$ | 0.63*** | 0.71*** | 0.83*** | 0.63*** | 0.54*** | 0.66*** | 0.82*** | 0.56*** |
| INDPRO | | | | | | | | |
| $h = 1$ | 0.60*** | 0.90*** | 0.96*** | 0.82*** | 0.52*** | 0.88*** | 0.95*** | 0.79*** |
| $h = 2$ | 0.52*** | 0.77*** | 0.89*** | 0.65*** | 0.43*** | 0.74*** | 0.88*** | 0.60*** |
| $h = 3$ | 0.67*** | 0.74*** | 0.86*** | 0.68*** | 0.62*** | 0.71*** | 0.84*** | 0.63*** |
| $h = 4$ | 0.83* | 0.87** | 0.91*** | 0.85* | 0.84 | 0.85* | 0.89** | 0.84 |

Table B.1 (continued)

| Forc. h | RMSFE | | | | CRPS | | | |
|-----------------|----------------|-------------|-------------|-----------|----------------|-------------|-------------|-----------|
| | INEQ | EQ (Middle) | EQ (Narrow) | EQ (Wide) | INEQ | EQ (Middle) | EQ (Narrow) | EQ (Wide) |
| CUMFNS | | | | | | | | |
| $h = 1$ | 0.63*** | 0.90*** | 0.96*** | 0.82*** | 0.53*** | 0.87*** | 0.94*** | 0.77*** |
| $h = 2$ | 0.51*** | 0.77*** | 0.89*** | 0.65*** | 0.41*** | 0.73*** | 0.87*** | 0.58*** |
| $h = 3$ | 0.62*** | 0.73*** | 0.85*** | 0.65*** | 0.55*** | 0.69*** | 0.83*** | 0.59*** |
| $h = 4$ | 0.73*** | 0.84*** | 0.91*** | 0.79*** | 0.69*** | 0.79*** | 0.87*** | 0.75*** |
| FEDFUNDS | | | | | | | | |
| $h = 1$ | 0.82*** | 0.97*** | 0.99*** | 0.93*** | 0.78*** | 0.96*** | 0.99*** | 0.92*** |
| $h = 2$ | 0.65*** | 0.89*** | 0.95*** | 0.81*** | 0.58*** | 0.87*** | 0.94*** | 0.77*** |
| $h = 3$ | 0.75*** | 0.86*** | 0.93*** | 0.80*** | 0.68*** | 0.83*** | 0.92*** | 0.75*** |
| $h = 4$ | 0.77*** | 0.85*** | 0.92*** | 0.80*** | 0.70*** | 0.80*** | 0.89*** | 0.73*** |

References

- Aastveit, K.A., Carriero, A., Clark, T.E., Marcellino, M., 2017. Have standard VARs remained stable since the crisis? *J. Appl. Econom.* 32 (5), 931–951.
- Altavilla, C., Giannone, D., Lenza, M., 2016. The financial and macroeconomic effects of the OMT announcements. *Int. J. Cent. Bank.* 12 (3), 29–57.
- Andersson, M.K., Palmqvist, S., Waggoner, D.F., 2010. Density-conditional forecasts in dynamic multivariate models. Discussion paper, Sveriges Riksbank Working Paper Series.
- Antolín-Díaz, J., Petrella, I., Rubio-Ramírez, J.F., 2021. Structural scenario analysis with SVARs. *J. Monet. Econ.* 117, 798–815.
- Bañura, M., Giannone, D., Lenza, M., 2015. Conditional forecasts and scenario analysis with vector autoregressions for large cross-sections. *Int. J. Forecast.* 31 (3), 739–756.
- Baumeister, C., Kilian, L., 2014. Real-time analysis of oil price risks using forecast scenarios. *IMF Econ. Rev.* 62 (1), 119–145.
- Botev, Z.I., 2017. The normal law under linear restrictions: simulation and estimation via minimax tilting. *J. R. Stat. Soc., Ser. B, Stat. Methodol.* 79 (1), 125–148.
- Chan, J.C.C., 2022. Asymmetric conjugate priors for large Bayesian VARs. *Quant. Econ.* 13 (3), 1145–1169.
- Chan, J.C.C., Jeliazkov, I., 2009. Efficient simulation and integrated likelihood estimation in state space models. *Int. J. Math. Model. Numer. Optim.* 1 (1–2), 101–120.
- Chan, J.C.C., Poon, A., Zhu, D., 2023. High-dimensional conditionally Gaussian state space models with missing data. *J. Econom.* 236 (1), 105468.
- Clarida, R.H., Coyle, D., 1984. Conditional Projection by Means of Kalman Filtering. Discussion paper, NBER Technical Working Papers 0036. National Bureau of Economic Research, Inc.
- Crump, R.K., Eusepi, S., Giannone, D., Qian, E., Sbordone, A.M., 2021. A Large Bayesian VAR of the United States Economy. FRB of New York Staff Report No. 976.
- Diebold, F.X., Mariano, R.S., 1995. Comparing predictive accuracy. *J. Bus. Econ. Stat.* 13 (3), 253–263.
- Dieppe, A., Legrand, R., Van Roye, B., 2016. The BEAR toolbox. ECB working paper.
- Doan, T., Litterman, R., Sims, C., 1984. Forecasting and conditional projection using realistic prior distributions. *Econom. Rev.* 3 (1), 1–100.
- Durbin, J., Koopman, S., 2002. A simple and efficient simulation smoother for state space time series analysis. *Biometrika* 89 (3), 603–615.
- Genz, A., 1992. Numerical computation of multivariate normal probabilities. *J. Comput. Graph. Stat.* 1 (2), 141–149.
- Geweke, J.F., 1996. Bayesian inference for linear models subject to linear inequality constraints. In: *Modelling and Prediction Honoring Seymour Geisser*. Springer, pp. 248–263.
- Giannone, D., Lenza, M., Momferatou, D., Onorante, L., 2014. Short-term inflation projections: a Bayesian vector autoregressive approach. *Int. J. Forecast.* 30 (3), 635–644.
- Giannone, D., Lenza, M., Pill, H., Reichlin, L., 2012. The ECB and the interbank market. *Econ. J.* 122 (564), F467–F486.
- Giannone, D., Lenza, M., Reichlin, L., 2019. Money, credit, monetary policy, and the business cycle in the euro area: what has changed since the crisis? *Int. J. Cent. Bank.* 15 (5), 137–173.
- Harvey, D., Leybourne, S., Newbold, P., 1997. Testing the equality of prediction mean squared errors. *Int. J. Forecast.* 13 (2), 281–291.
- Jarociński, M., 2010. Conditional forecasts and uncertainty about forecast revisions in vector autoregressions. *Econ. Lett.* 108 (3), 257–259.
- Jarociński, M., Smets, F., 2008. House prices and the stance of monetary policy. ECB Working Paper.
- Lenza, M., Pill, H., Reichlin, L., 2010. Monetary policy in exceptional times. *Econ. Policy* 25 (62), 295–339.
- McCracken, M.W., McGillicuddy, J.T., Owyang, M.T., 2021. Binary conditional forecasts. *J. Bus. Econ. Stat.* 40 (3), 1246–1258.
- Robertson, J.C., Tallman, E.W., Whiteman, C.H., 2005. Forecasting using relative entropy. *J. Money Credit Bank.* 37 (3), 383–401.
- Sims, C.A., Zha, T., 1998. Bayesian methods for dynamic multivariate models. *Int. Econ. Rev.* 39 (4), 949–968.
- Tallman, E.W., Zaman, S., 2020. Combining survey long-run forecasts and nowcasts with BVAR forecasts using relative entropy. *Int. J. Forecast.* 36 (2), 373–398.
- Waggoner, D.F., Zha, T., 1999. Conditional forecasts in dynamic multivariate models. *Rev. Econ. Stat.* 81 (4), 639–651.