Conditional Forecasts in Large Bayesian VARs with Multiple Soft and Hard Constraints

Joshua C. C. Chan
Purdue University

Davide Pettenuzzo
Brandeis University

Aubrey Poon
Orebro University

Dan Zhu
Monash University

February 2023

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 soft and hard constraints are imposed at once. We introduce a novel precision-based sampler that is fast, scales well, and yields conditional forecasts from both hard and soft 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 where we simultaneously impose a mix of soft and hard constraints on the future trajectories of key US macroeconomic indicators over the 2020–2022 period.

Keywords: precision-based method, conditional forecast, vector autoregression

JEL classification: C11, C32, C51

1 Introduction

Conditional forecasts are projections of a set of variables of interest on the future paths of some other variables. This is in contrast to unconditional forecasts, where no knowledge of the future paths of any variables is assumed. 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]

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 stand 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. Examples along this line of work include Andersson, Palmqvist, and Waggoner (2010), Giannone, Lenza, Pill, and Reichlin (2012); Giannone, Lenza, Momferatou, and Onorante (2014), Altavilla, Giannone, and Lenzaa (2016), Aastveit, Carriero, Clark, and Marcellino (2017), Giannone, Lenza, and Reichlin (2019), and Tallman and Zaman (2020).

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, Legrand, and Van Roye (2016). An important contribution in this regard is the recent work of Antolín-Díaz, Petrella, and Rubio-Ramírez (2021), who introduced a unified framework for

---

[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 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.

conditional forecasts and structural scenario analysis with VAR models.[2] 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. 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, Giannone, and Lenza (2015), and the "forecast scenario" method of Baumeister and Kilian (2014)), the proposed approach is closed-form and valid for an arbitrary number of conditions.[3]

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., hard conditions) or instead prefers to allow the future values of the conditioned variables to lie within a certain range (i.e., soft conditions). The hard constraint case is the most commonly employed approach in the literature (see again Giannone, Lenza, Pill, and Reichlin (2012); Giannone, Lenza, Momferatou, and Onorante (2014) as well as Jarocinski and Smets (2008) and Lenza, Pill, and Reichlin (2010)), while to date there has been only limited work with soft constrained conditional forecasts (in addition to Waggoner and Zha (1999), see also Andersson, Palmqvist, and Waggoner (2010)). This is largely due to the computational challenges that the soft constraints entail. For example, the algorithm of Waggoner and Zha (1999) for the soft constraints is very computationally 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 soft constraint applications is unfortunate, as there are many good reasons why soft 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

---

[2]Antolín-Díaz, Petrella, and Rubio-Ramírez (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.

[3]Antolín-Díaz, Petrella, and Rubio-Ramírez (2021) also show that in the Gaussian case, their approach to conditional forecasting is equivalent to the entropic forecast tilting method of Robertson, Tallman, and Whiteman (2005).

impose that the future values of the variables conditioned on will be between a range or an interval instead of an exact path. In addition, soft constraints allow the forecaster to acknowledge the uncertainty surrounding the future realization path of the constrained endogenous variables. For example, Andersson, Palmqvist, and Waggoner (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, Petrella, and Rubio-Ramírez (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. As we show in our derivations, this result is made possible by how we parameterize the distribution of the conditional forecasts, which relies on the precision matrix (i.e., the inverse of the conditional forecasts' covariance matrix) and which in turns leads to exploiting fast band matrix algorithms such as those introduced by Chan and Jeliazkov (2009) to speed up computations substantially.

Third, our approach can be used to produce conditional forecasts and structural scenario analysis under both hard and soft 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, Palmqvist, and Waggoner (2010).[4] 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 soft constraint conditional forecasts in a fast and efficient manner within the precision sampler. More specifically, Botev (2017) provides an efficient method to generate random draws from high-dimensional Gaussian distributions under linear restrictions. Even though

---

[4]Andersson, Palmqvist, and Waggoner (2010) show how to implement the soft-condition 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 soft constraint.

the algorithm is based on an accept-reject sampler, it carefully constructs a suitable proposal distribution that satisfies all the constraints via exponential tilting. As such, it has good theoretical properties, and the acceptance probability of the candidate draws is typically high.

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 hard constraint conditional forecasts against the Waggoner and Zha (1999), Bańbura, Giannone, and Lenza (2015) and Antolín-Díaz, Petrella, and Rubio-Ramírez (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 soft constraint precision-based approach algorithm by comparing it against the Waggoner and Zha (1999), and Andersson, Palmqvist, and Waggoner (2010) accept-reject methods (both Bańbura, Giannone, and Lenza (2015) and Antolín-Díaz, Petrella, and Rubio-Ramírez (2021) are not designed to generate conditional forecasts in the presence of soft constraints) and once again find that our soft 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 soft 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 an empirical application where we follow Crump, Eusepi, Giannone, Qian, and Sbordone (2021) and estimate a large BVAR including 31 quarterly US macro and financial variables from 1976 to the end of 2019. We use this setting to investigate the effect of simultaneously imposing soft and hard constraints on the trajectories of CPI inflation, unemployment rate, and the 10-year Treasury rate over the following 13 quarters, setting our constraints to mimic the baseline and severe scenarios prepared by the Federal Reserve Board for their 2020 stress test analysis. We find that when we condition on the Fed's baseline projections for CPI inflation, unemployment and the 10-year Treasury rate, meant to approximate a moderate expansion over the next 13 quarters, we observe small temporary increases

in real GDP, industrial production, housing starts and the S&P 500 followed by a gradual convergence back to their original levels sometime around the end of 2022. On the contrary, the story that emerges from conditioning our forecasts on the adverse scenario, which the Fed describes as a severe global recession accompanied by a period of heightened stress in commercial real estate and corporate debt markets, is quite different. In this case, we observe a sharp decrease in real GDP, industrial production, and housing starts, 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 index and a sudden increase in the market volatility, as measured by the VIX. Interestingly, with the exception of the effect on the S&P 500 index, the initial responses we see in our conditional forecasts are in the same direction and, in most cases, of a magnitude similar to what the US economy experienced after the large and unexpected COVID-19 shocks that occurred in the first half of 2020. However, the sudden and unprecedented nature of the COVID-19 shock and the unfolding of events that followed the initial impact were such that the US economy's initial reaction was much more immediate and the recovery that followed much faster.

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 hard and soft 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 soft and hard 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, Petrella, and Rubio-Ramírez (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 hard and soft 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 subsubsection 2.2.1 and subsubsection 2.2.2 how to apply this setup to producing hard- and soft-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

As our starting point, consider an $n \times 1$ vector of variables $\mathbf{y}_t = (y_{1,t}, \ldots, y_{n,t})'$, and write the following SVAR with p lags:

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

where $\mathbf{a}$ is an $n \times 1$ vector of intercepts, $\mathbf{A}_1, \ldots, \mathbf{A}_p$ 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 $\mathbf{y}^T = (\mathbf{y}'_{1-p}, \ldots, \mathbf{y}'_T)'$, the unconditional forecast of our n variables for the next h periods, $\mathbf{y}_{T+1,T+h} = (\mathbf{y}'_{T+1}, \ldots, \mathbf{y}'_{T+h})'$, can be

7

written as

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

where

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

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

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

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.[5] 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, Palmqvist, and Waggoner (2010) and Antolín-Díaz, Petrella, and Rubio-Ramírez (2021), we write the conditional forecasts as a set of linear restrictions

---

[5]More specifically, to obtain a draw from the Gaussian distribution in (4), we first obtain $\mathbf{u} = (u_1, \ldots, u_{nh})'$, where $u_i \sim \mathcal{N}(0,1), i = 1, \ldots, nh$. 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 $\mathbf{z} \sim \mathcal{N}\left(\mathbf{H}^{-1}\mathbf{c}, (\mathbf{H}'\mathbf{H})^{-1}\right)$.

on the path of future observables $\mathbf{y}_{T+1,T+h}$, i.e.

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

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, Petrella, and Rubio-Ramírez (2021) note, the setup in (5) is very general and can accommodate both the classic "hard" 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, Palmqvist, and Waggoner (2010). It also replicates *entropic tilting*, the approach introduced and popularized by Robertson, Tallman, and Whiteman (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, Petrella, and Rubio-Ramírez (2021) for a formal proof of this equivalence in the Gaussian case). Next, combining (2) and (5), we obtain

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

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

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

where $\boldsymbol{\mu}_{\varepsilon}$ and $\mathbf{\Psi}_{\varepsilon}$ 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 $\mathbf{\Psi}_{\varepsilon}$:

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

In typical situations where $r < nh$, the system in (8) is underdetermined and has multiple solutions.[6] Following Antolín-Díaz, Petrella, and Rubio-Ramírez (2021), we

---

[6]If $r = nh$, the system is just determined and has a unique solution; if $r > nh$, the system is

choose a solution that can be expressed in terms of the Moore-Penrose inverse of $\mathbf{R}\mathbf{H}^{-1}$, which we denote as $\left(\mathbf{R}\mathbf{H}^{-1}\right)^{+}$:

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

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

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

then, (2) and (9) imply that

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

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 $\boldsymbol{\mu}_{\mathbf{y}}$ while preserving the variance of the unconditional forecasts, this could be easily accomplished by setting $\boldsymbol{\Omega} = \mathbf{R}(\mathbf{H}'\mathbf{H})^{-1}\mathbf{R}'$. In that case, it is easy to show that $\boldsymbol{\Sigma}_{\mathbf{y}} = (\mathbf{H}'\mathbf{H})^{-1}$, 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, Petrella, and Rubio-Ramírez (2021). One important 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

---

inconsistent and there are no solutions.

band matrix algorithms such as those introduced by Chan and Jeliazkov (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. 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 soft conditioning, where one allows the future values of the conditioned variables to lie within a certain range. Let the soft 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 $\mathbf{y}_{T+1,T+h}$,

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

where $\mathbb{1}(\cdot)$ is the indicator function. We note that this is a slightly more general formulation of soft 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, Palmqvist, and Waggoner (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 Jeliazkov (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 $\boldsymbol{\mu_y}$ while preserving the variance of the unconditional forecasts by setting $\boldsymbol{\Omega} = \mathbf{R}(\mathbf{H}'\mathbf{H})^{-1}\mathbf{R}'$. Then, using the same logic we

applied above, we obtain

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

where the definition of $\boldsymbol{\mu}_\mathbf{y}$ is given in (11).[7]

We now turn to discuss a number of special cases to highlight the generality of the approach. We begin with the classic "hard" conditional forecast case introduced by Doan, Litterman, and Sims (1984) and popularized by Waggoner and Zha (1999). Next, we move to the soft 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.

2.2.1   Hard Constraints on Future Observables

The first special case is the classic conditional forecast setup introduced by Doan, Litterman, and Sims (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

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

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 $\boldsymbol{\Omega} = \mathbf{0}_{r_o \times r_o}$ and is sometimes referred to as conditional forecasting under hard constraints.

In this case, some of the computations described in the general framework become

---

[7]In the special case without soft constraints, i.e. when $\underline{c}_i = \infty$ and $\bar{c}_i = \infty, i = 1, \ldots, nh$, we revert back to the results presented in (5).

unnecessary and the algorithm can be streamlined. To that end, we partition the $nh \times 1$ vector $\mathbf{y}_{T+1,T+h}$ into $\mathbf{y}^o_{T+1,T+h}$ and $\mathbf{y}^u_{T+1,T+h}$, where the former is a $r_o \times 1$ vector including the hard-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 $\mathbf{y}^o_{T+1,T+h}$ is the $h \times 1$ vector of policy rates ($r_0 = h$) over the forecast horizons $t = T+1, \ldots, T+h$ and $\mathbf{y}^u_{T+1,T+h}$ is the $3h \times 1$ vector consisting of GDP, inflation and unemployment.

Let $\mathbf{R}^-_o$ denote the associated $(nh - r_o) \times nh$ selection matrix that selects $\mathbf{y}^u_{T+1,T+h}$. Then, we can write $\mathbf{y}_{T+1,T+h}$ as follows:

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

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 $\mathbf{y}^u_{T+1,T+h}$ given $\mathbf{y}^o_{T+1,T+h}$ and the model parameters $\mathbf{A}_0$ and $\mathbf{A} = (\mathbf{a}, \mathbf{A}_1, \ldots, \mathbf{A}_p)'$. To this end, substitute (16) into (2) to obtain

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

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

$$p(\mathbf{y}^u_{T+1,T+h} \mid \mathbf{y}^o_{T+1,T+h}, \mathbf{A}_0, \mathbf{A})$$
$$\propto \exp\left\{ -\frac{1}{2}(\mathbf{H}(\mathbf{M}_u \mathbf{y}^u_{T+1,T+h} + \mathbf{M}_o \mathbf{y}^o_{T+1,T+h}) - \mathbf{c})'(\mathbf{H}(\mathbf{M}_u \mathbf{y}^u_{T+1,T+h} + \mathbf{M}_o \mathbf{y}^o_{T+1,T+h}) - \mathbf{c}) \right\}$$
$$\propto \exp\left\{ -\frac{1}{2}\left(\mathbf{y}^{u'}_{T+1,T+h} \mathbf{M}'_u \mathbf{H}' \mathbf{H} \mathbf{M}_u \mathbf{y}^u_{T+1,T+h} - 2\mathbf{y}^{u'}_{T+1,T+h} \mathbf{M}'_u \mathbf{H}' (\mathbf{c} - \mathbf{H} \mathbf{M}_o \mathbf{y}^o_{T+1,T+h}) \right) \right\}$$
$$\propto \exp\left\{ -\frac{1}{2}(\mathbf{y}^u_{T+1,T+h} - \boldsymbol{\mu}_u)' \mathbf{K}_u (\mathbf{y}^u_{T+1,T+h} - \boldsymbol{\mu}_u) \right\},$$

13

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

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

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 $\mathbf{y}^u_{T+1,T+h}$ efficiently.

2.2.2 Soft Constraints on Future Observables

We now move to discussing the case of soft 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 or soft 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 soft 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, Palmqvist, and Waggoner (2010) instead use the approach in Geweke (1996) to draw from the truncated multivariate Gaussian distribution implied by the soft 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

14

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 soft constraints such that $\mathbf{S}$ is a $s_o \times nh$ selection matrix. Then, as in the hard-constraints-on-observables case, we can partition $\mathbf{y}_{T+1,T+h}$ into $\mathbf{y}^o_{T+1,T+h}$ and $\mathbf{y}^u_{T+1,T+h}$, where the former is a $s_o \times 1$ vector of the soft-constrained endogenous variables such that $\mathbf{y}^o_{T+1,T+h} = \mathbf{S}\mathbf{y}_{T+1,T+h}$ 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 $\mathbf{y}^o_{T+1,T+h}$ and the conditional distribution of $\mathbf{y}^u_{T+1,T+h}$ given $\mathbf{y}^o_{T+1,T+h}$. 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):

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

where $\mathbf{K}_u = \mathbf{M}'_u\mathbf{H}'\mathbf{H}\mathbf{M}_u$, $\boldsymbol{\mu}_u = \mathbf{K}_u^{-1}\mathbf{M}'_u\mathbf{H}'\mathbf{H}(\mathbf{H}^{-1}\mathbf{c} - \mathbf{M}_o\mathbf{y}^o_{T+1,T+h})$, $\mathbf{K}_o = (\mathbf{M}'_o(\mathbf{H}'\mathbf{H})^{-1}\mathbf{M}_o)^{-1}$, and $\boldsymbol{\mu}_o = \mathbf{M}'_o\mathbf{H}^{-1}\mathbf{c}$.[8] Hence, to obtain a draw for $\mathbf{y}_{T+1,T+h}$,

---

[8]Note that in general, a marginal distribution of a truncated Gaussian is not a truncated Gaussian. Here the marginal distribution of $\mathbf{y}^o_{T+1,T+h}$ is a truncated Gaussian because only $\mathbf{y}^o_{T+1,T+h}$ is restricted and $\mathbf{y}^u_{T+1,T+h}$ is not.

we can first sample $\mathbf{y}^o_{T+1,T+h}$ marginally from its truncated s_o-variate Gaussian distribution using the algorithm of Botev (2017). Given a draw for $\mathbf{y}^o_{T+1,T+h}$, we can then sample $\mathbf{y}^u_{T+1,T+h}$ 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 focused on the classic (or reduced-form) conditional forecast case, where one assumes that the hard or soft 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, Petrella, and Rubio-Ramírez (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:

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

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 subsubsection 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 hard constraints with $\mathbf{\Psi} = \mathbf{0}_{w \times w}$).

Finally, we consider the structural scenario case discussed in Antolín-Díaz, Petrella, and Rubio-Ramírez (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, Petrella, and Rubio-Ramírez (2021) call the driving shocks—can deviate from their unconditional distribution over the forecast horizon, while the remaining structural 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, Petrella, and Rubio-Ramírez (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

$$\underbrace{\begin{bmatrix} \mathbf{R}_o \\ \mathbf{W} \end{bmatrix}}_{\tilde{\mathbf{R}}} \mathbf{y}_{T+1,T+h} \sim \mathcal{N}\left(\underbrace{\begin{bmatrix} \mathbf{r}_o \\ \mathbf{W}\mathbf{c} \end{bmatrix}}_{\tilde{\mathbf{r}}}, \underbrace{\begin{bmatrix} \mathbf{\Omega}_o & \mathbf{0}_{r_0 \times w} \\ \mathbf{0}_{w \times r_0} & \mathbf{I}_w \end{bmatrix}}_{\tilde{\mathbf{\Omega}}} \right). \tag{20}$$

17

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

It is worth noting that while the cases reviewed in this section follow closely the discussion in Antolín-Díaz, Petrella, and Rubio-Ramírez (2021), by casting these restrictions within our general setup we are able to generalize and improve upon their approach. First, as noted earlier, both hard or soft constraints can be imposed on the observables, and this carries over to the structural scenario restrictions. For example, soft 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.

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, Palmqvist, and Waggoner (2010), Bańbura, Giannone, and Lenza (2015) and Antolín-Díaz, Petrella, and Rubio-Ramírez (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).

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

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

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 $\mathbf{\Sigma}$ from the inverse-Wishart distribution $\mathcal{IW}(n + 10, 0.07\mathbf{I}_n + 0.031_n\mathbf{1}_n')$.

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 $\mathbf{\Sigma}$, i.e. $\boldsymbol{\beta} \sim \mathcal{N}(\mathbf{0}_k, \mathbf{I}_k)$ and $\mathbf{\Sigma} \sim \mathcal{IW}(3 + n, \mathbf{I}_n)$.

3.1 Hard Constraints on Observables

In the first set of simulations, we work with hard 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 compare the proposed algorithms against the three conditional forecast algorithms in the literature: Waggoner and Zha (1999), Bańbura, Giannone, and Lenza (2015) and Antolín-Díaz, Petrella, and Rubio-Ramírez (2021). We consider VARs of different dimensions: a medium size VAR with a short forecast horizon ($n = 8$, $h = 5$) and a large VAR with a longer forecast horizon ($n = 15$, $h = 20$). In both cases, we set the number of constrained endogenous variables to $n_o = 3$ and retain 25,000 MCMC draws, after discarding the first 10,000 draws.[9]

Focusing on accuracy first, Figure 1 and Figure 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).

---

[9]For all calculations in this section, we used a standard desktop with an Intel Xeon W-2224 @3.60GHz processor and 16GB memory.
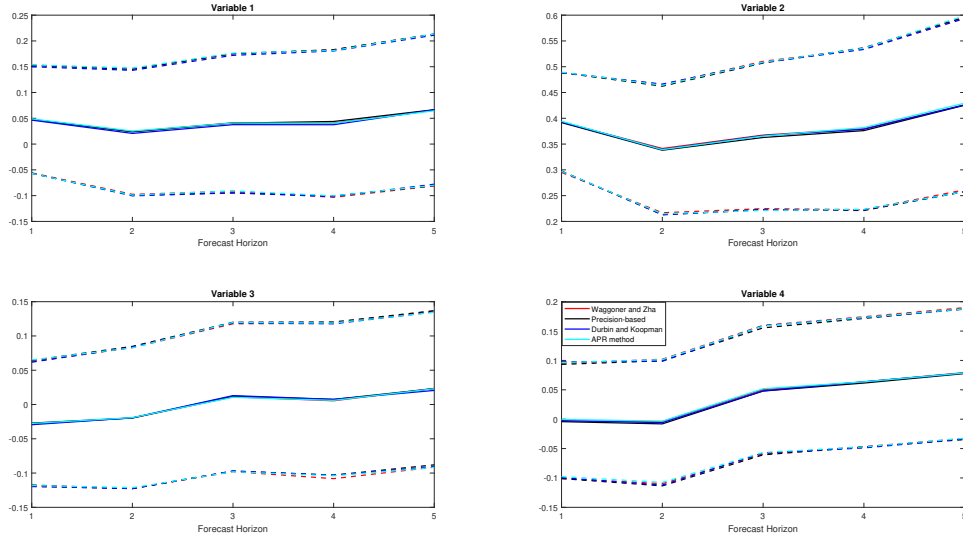
Figure 1: **Conditional forecasts from a medium VAR with three hard constraints.**
The thick black line is the posterior median estimate of the conditional forecast using
our precision-based method. The thick red line is the posterior median estimates
of the conditional forecast using the Waggoner and Zha (1999) approach. The thick
dark blue line is the posterior median estimate of the conditional forecast using the
Bańbura, Giannone, and Lenza (2015) Durbin and Koopman method. The thick light
blue line is the posterior median estimate of the conditional forecast using the Antolín-
Díaz, Petrella, and Rubio-Ramírez (2021). The dashed lines are the corresponding 68%
credible intervals for all four methods.

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
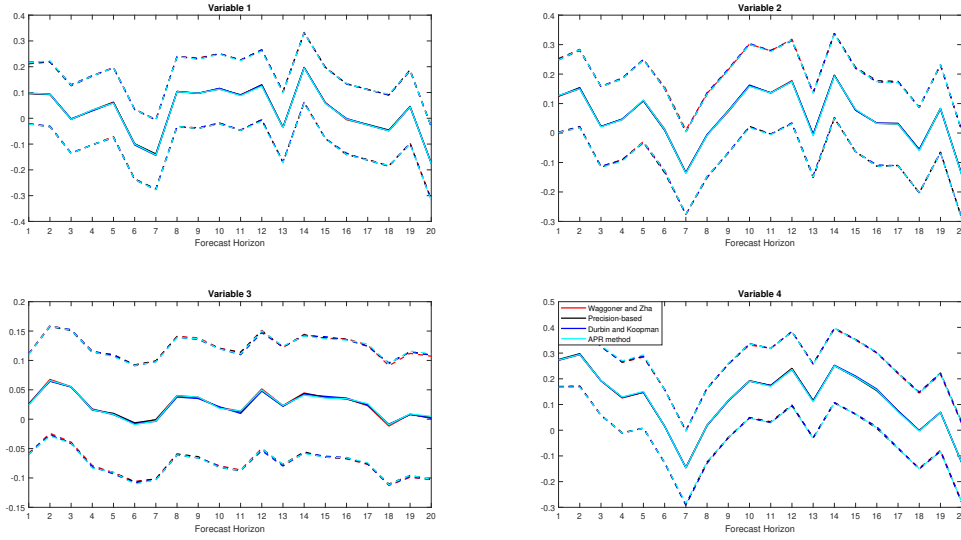other three benchmarks.

Figure 2: **Conditional forecasts from a large VAR with three hard 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 using the Bańbura, Giannone, and Lenza (2015) Durbin and Koopman method. The thick light blue line is the posterior median estimate of the conditional forecast using the Antolín-Díaz, Petrella, and Rubio-Ramírez (2021). The dashed lines are the corresponding 68% credible intervals for all four methods.

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.[10] We present these results in Table 1.[11] 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 hard constraints increases in the VAR, the proposed precision-based method

---

[10]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 from a hard constraint compared to the original algorithm of Waggoner and Zha (1999).

[11]The computation times reported in this table are 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.

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 hard constraints. On the other hand, the conditional forecasting method proposed by Antolín-Díaz, Petrella, and Rubio-Ramírez (2021) does not appear to scale well, becoming computationally very intensive when the dimension of the VAR or the number of hard constraints increase.

3.2 Soft Constraints on Observables

Next, we move to the soft-constraint case. We follow the same simulation structure as described above in the hard constraint case, except that we can only compare our conditional forecasts to those obtained using the Waggoner and Zha (1999) and Andersson, Palmqvist, and Waggoner (2010) algorithms since both the approaches of Bańbura, Giannone, and Lenza (2015) and Antolín-Díaz, Petrella, and Rubio-Ramírez (2021) are not designed to handle soft constraints. More specifically, we set $\bar{\mathbf{y}}^o_{T-h:T} - 0.1 < \mathbf{y}^o_{T:T+h} < \bar{\mathbf{y}}^o_{T-h:T} + 0.1$, where $\bar{\mathbf{y}}^o_{T-h:T} = \frac{1}{h}\sum_{t=T-h+1}^{T} \mathbf{y}^o_t$, is the average of the actual simulated data over the periods $t = T - h, \ldots, T$. In this simulation exercise, we consider a medium-sized VAR where $n = 8$ and a long forecast horizon $h = 20$ with only one soft constraint $n_o = 1$. 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 hard constraint exercise.

Figure 3 plots the conditional forecasts for the first four variables from the medium-sized VAR with one soft 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, Palmqvist, and Waggoner (2010) (denoted in the figure as APW) algorithms. However, as mentioned in the previous section, the Waggoner and Zha (1999) soft constraint algorithm is very computationally intensive. For example, when working with a medium VAR and one soft constraint, their accept-reject algorithm took approximately 2,542 minutes to compute the conditional forecasts. In contrast, our proposed precision-based method and the Andersson, Palmqvist, and Waggoner (2010) take about 62 and 70 seconds to obtain the corresponding conditional forecasts, respectively.

22

Dimension			Estimation-Conditional-on-forecast			
			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 1: **Computation time for the hard 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, Giannone, and Lenza (2015) (DK) and Antolín-Díaz, Petrella, and Rubio-Ramírez (2021) (APR) methods. The computation times are based on 25,000 MCMC draws with a 10,000 burn-in period.
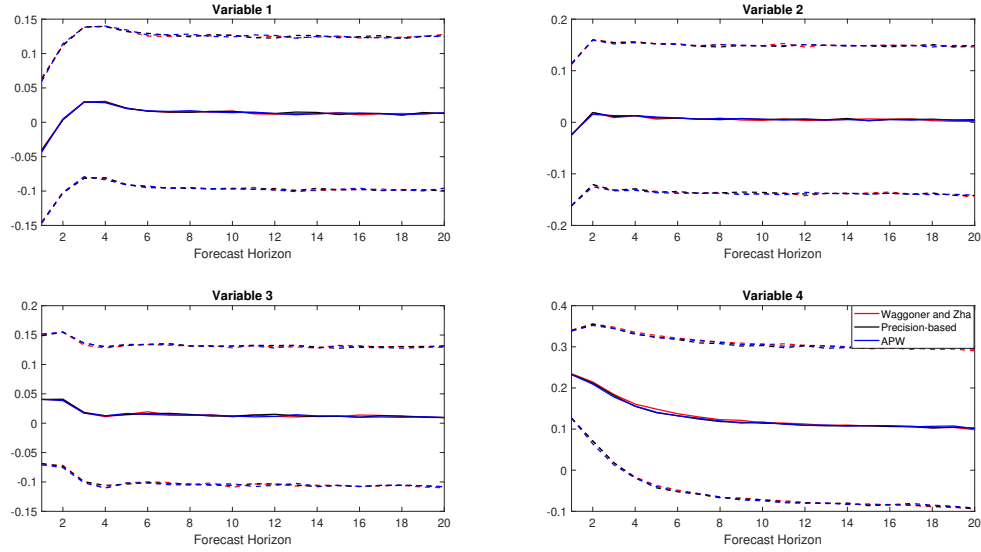
Figure 3: **Conditional forecast from a medium VAR with one soft constraint**. The thick black line is the posterior median estimate of the conditional forecast using the proposed precision-based method. The thick red line is the posterior median estimates of the conditional forecast using the Waggoner and Zha (1999) method. The thick blue line is the posterior median estimates of the conditional forecasting using the Andersson, Palmqvist, and Waggoner (2010) (APW) method. The dash lines are the corresponding 68 per cent credible intervals for all three methods.

Note that while the computation costs of our precision-based and the Andersson, Palmqvist, and Waggoner (2010) methods are relatively similar when focusing on a single soft constraint, Table 2 shows that as the number of soft constraints increase our precision-based method becomes more computationally efficient than the Andersson, Palmqvist, and Waggoner (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).

	No. of soft 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

Table 2: **Computation time for soft constraints case**. This table report the computation time (in seconds) required to simulate 1,000 draws for the precision-based sampler and the Andersson, Palmqvist, and Waggoner (2010) (APW) method in the case of a medium VAR and a $h = 20$ forecast horizon.

4   Empirical Application

To illustrate the performance of the 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. The complete list of the variables used along with their transformations is available in Appendix A and our sample goes from 1971Q1 to 2022Q3. Our analysis is inspired by the recent work of Crump, Eusepi, Giannone, Qian, and Sbordone (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.[12] We exploit the computational efficiency and generality of the methods presented in the previous sections to expand on the analysis of Crump, Eusepi, Giannone, Qian, and Sbordone (2021) and investigate the macroeconomic impact of a combination of multiple soft and hard constraints at once. 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 hard and soft constraints.[13]

Our estimation strategy 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,

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

where $\mathbf{A}_0$ is a lower triangular matrix with ones on its main diagonal and $\boldsymbol{\Sigma} = diag\left(\sigma_1^2, ..., \sigma_n^2\right)$ 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 $\boldsymbol{\alpha}_i$ denote the $i-1$ free elements from the i-th row of $\mathbf{A}_0$, and define $\boldsymbol{\beta}_i = \left(a_i, \mathbf{a}_{1,i}, ..., \mathbf{a}_{p,i}\right)'$. With this notation in hand, we can rewrite the i-th equation of

---

[12]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.

[13]Crump, Eusepi, Giannone, Qian, and Sbordone (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.

(22) as follows:

$$y_{it} = \mathbf{w}_{it}\boldsymbol{\alpha}_i + \mathbf{z}_{it}\boldsymbol{\beta}_i + \varepsilon_{it} \tag{23}$$

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

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

and

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

where s_i^2 denotes the sample variance of the residuals from an AR(p) model estimated on variable i, $i = 1, ..., n$. 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 $\mathbf{V}_i^{\alpha} = diag\left(1/s_1^2, ..., 1/s_{i-1}^2\right)$. Next, we set all elements of $\mathbf{m}_i^{\beta}$ to zero, except for the coefficient associated with the first own lag, which we set to one. As for $\mathbf{V}_i^{\beta}$, we specify it to be a diagonal matrix with the k-th diagonal element, $(\mathbf{V}_i^{\beta})_k$, inspired by the Minnesota prior,

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

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, Eusepi, Giannone, Qian, and Sbordone (2021) used to explore a number of their scenarios), after setting the lag length of the VAR to $p = 4$.[14] Before turning to the conditional forecasts, we quickly review the estimation

---

[14]As in Crump, Eusepi, Giannone, Qian, and Sbordone (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.

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, Eusepi, Giannone, Qian, and Sbordone (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.

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

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

Next, we move to studying the model's forecasts, both conditional and unconditional. We consider two exercises. In the first one, we replicate one of the analysis in Crump, Eusepi, Giannone, Qian, and Sbordone (2021), where we estimate the BVAR in (22) using data from 1976Q4 to 2019Q3 and increase the one-period ahead forecasts for real GDP by one percent, while leaving all other variables unconstrained. We then compute, for each of the 31 variables in the VAR, the difference between the conditional and unconditional forecasts over the next three years. This, as Crump, Eusepi, Giannone, Qian, and Sbordone (2021) note, is equivalent to computing an unorthogonalized impulse response function (IRF). Figure 4 displays the responses of 16 selected macroeconomic and financial variables to a one percent increase in real GDP. Along with the posterior mean of the IRFs, shown in the figure with solid lines, we report the 68% credibile intervals. As in Crump, Eusepi, Giannone, Qian, and

Sbordone (2021), a one percent increase in the level of real GDP produces higher real GDP, higher PCE inflation, higher 1-year Treasury yield, and higher exports and imports. On the other end, the unemployment rate falls at first, to then revert back to the level of the unconditional forecasts by the beginning of 2022.

In the second exercise, we leverage the generality of our approach and condition the VAR forecasts on a combination of hard and soft constraints. In particular, after updating the BVAR estimates to include data from 2019Q4, we set a number of hard constraints on the future paths of unemployment rate and the 10-year Treasury rate by relying on 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 and businesses even in a severe recession.[15] We consider two scenarios, namely (i) baseline and (ii) severely adverse and as in the previous exercise 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 hard constraints described above with 13 more soft 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 match these with a measure of disagreement that we extract from the Survey of Professional Forecasters (SPF).

---

[15]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.
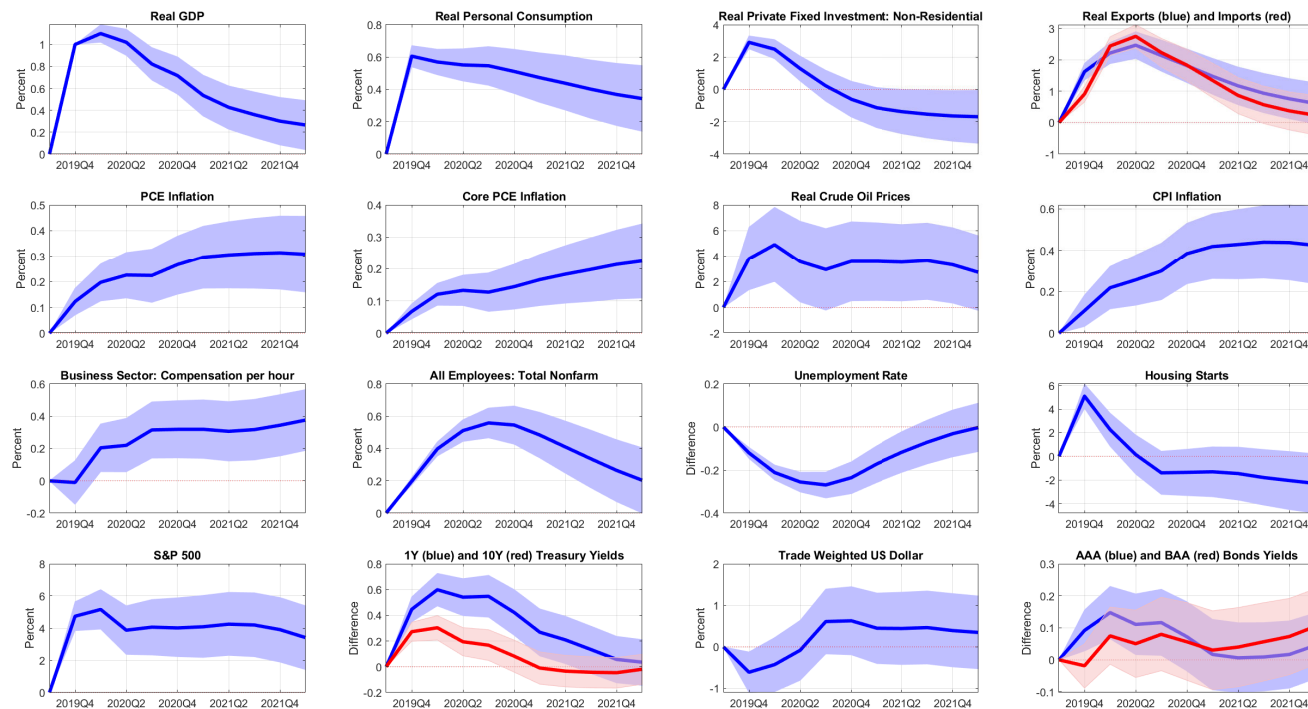
Figure 4: **Responses to a one percent increase in the level of real GDP one quarter ahead**. Within each panel, the solid line(s) shows the posterior mean while the shaded area depicts the 68% coverage intervals.

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.[16] 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 soft constraints for CPI inflation are centered on the FED stress test numbers and have a P25-P75 range that reflects the disagreement 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, after the sixth quarter (or from 2021Q2 onwards), we gradually widen the soft constraint interval around the corresponding FED stress test number as we progress through the forecast horizon to reflect the high degree of uncertainty that we have about future inflation. Table 4 summarizes the soft and hard constraints for the baseline and adversely severe scenarios.

Figure 5 to Figure 8 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) the S&P 500 index, and (vi) the CBOE Volatility index. Starting with the baseline scenario, Figure 5 displays the unconditional forecasts (blue bands and solid line) and conditional forecasts (red bands and solid line), where the bands correspond to 68% coverage intervals and the solid lines correspond to the posterior means. 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. The latter provides a reference point for how the economy actually evolved over the particularly tumultuous period that followed the initial COVID-19 shock and its unfolding.

---

[16]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.

	CPI Inflation (Soft Constraint)			Hard Constraints	
Date	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	1.40	2.10	3.00	3.70	2.10
2021Q3	1.40	2.10	3.00	3.80	2.10
2021Q4	1.25	2.10	4.00	3.80	2.20
2022Q1	1.25	2.30	4.00	3.90	2.20
2022Q2	1.10	2.20	5.00	3.90	2.40
2022Q3	1.10	2.20	5.00	3.90	2.50
2022Q4	1.00	2.20	6.00	3.90	2.60
2023Q1	1.00	2.20	6.00	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	0.90	1.40	2.00	9.70	1.30
2021Q3	0.90	1.50	2.00	10.00	1.40
2021Q4	0.95	1.70	3.00	9.90	1.50
2022Q1	0.95	1.80	3.00	9.70	1.60
2022Q2	0.97	1.80	4.00	9.50	1.80
2022Q3	0.97	1.80	4.00	9.20	1.90
2022Q4	1.00	1.80	5.00	8.80	2.10
2023Q1	1.00	1.70	5.00	8.50	2.20

Table 4: **Summary of hard and soft constraints**. This table reports the full set of soft and hard constraints over the 2021Q1–2023Q1 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) adverse.

Next, Figure 6 plots the differences between the unconditional and conditional forecasts, once again showing the posterior means and the 68% coverage intervals. As the figure shows, relative to the unconditional forecasts, conditioning on the Fed's baseline projections for CPI inflation, unemployment and the 10-year Treasury rate leads to temporary increases in most of the variables, with the most notable case being the industrial production which increases as much as 5% over the first half of 2021, before gradually returning back to its original level by the end of 2022.

The story that emerges from the Fed's adverse scenario, shown in Figure 7 and Figure 8 is instead quite different. 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.
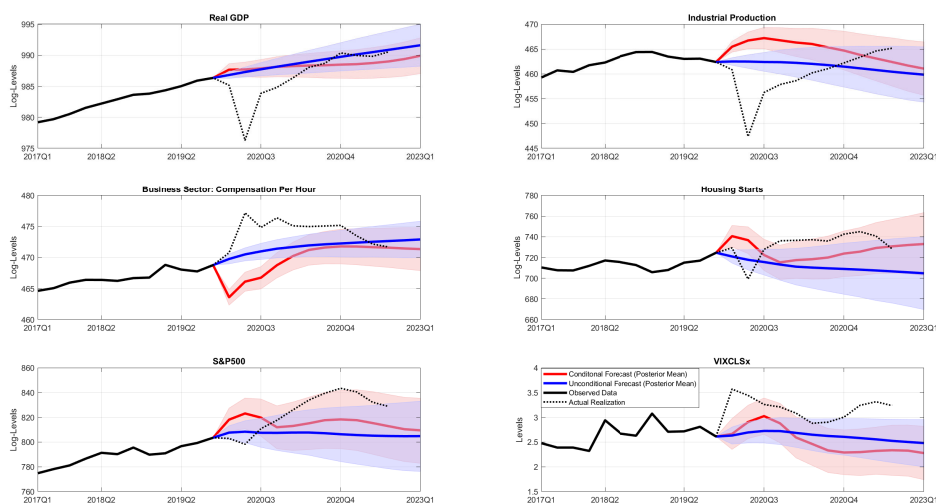

Figure 5: Conditional and unconditional forecasts when CPI inflation, unemployment, and the 10-year Treasury rate in 2021Q1–2023Q1 match the Fed's stress test baseline projections. The shaded bands correspond to the 68% coverage intervals while the solid and dotted black lines denote the in-sample and out-of-sample values.
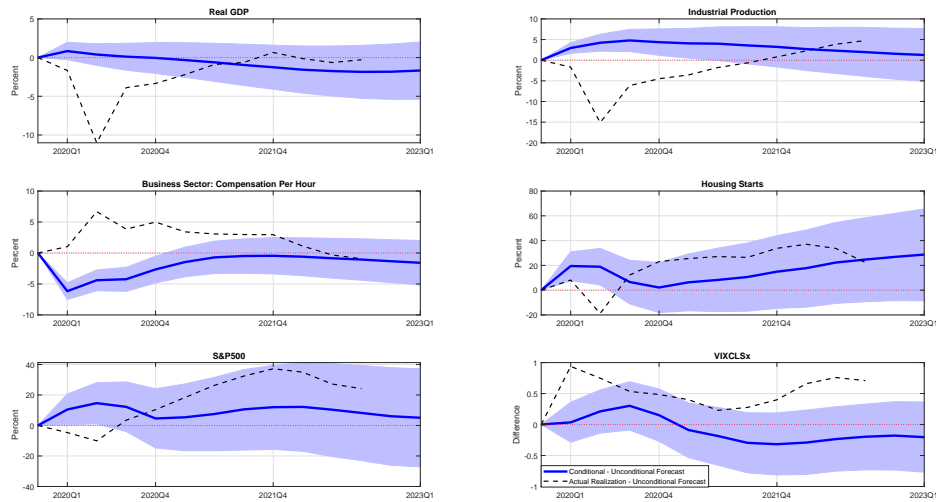
Figure 6: Difference between the conditional and unconditional forecasts when CPI inflation, unemployment, and the 10-year Treasury rate in 2021Q1–2023Q1 match the Fed's baseline projections. Bands denote to the 68% coverage intervals. The black dashed line shows the difference between the actual realizations and the posterior mean of the unconditional forecasts.

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. Interestingly, with the exception of the effect on the S&P 500 index, the initial responses we see in our conditional forecasts are in the same direction and in most cases of a magnitude similar to what the US economy experienced after the large and unexpected COVID-19 shocks that occurred during the first half of 2020. However, the sudden and unprecedented nature of the COVID-19 shock, as well as the unfolding of events that followed the initial impact, were such that the US economy's initial reaction was much more immediate and the recovery that followed much faster.
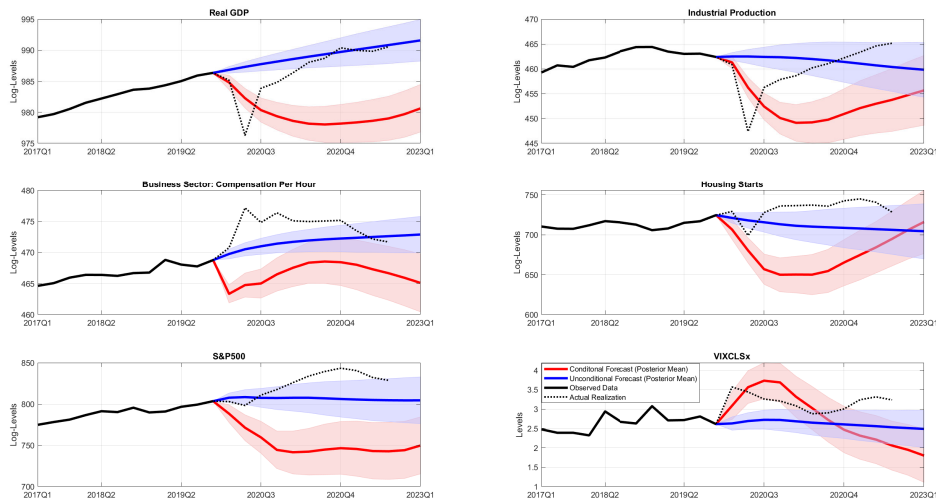
Figure 7: Conditional and unconditional forecasts when CPI inflation, unemployment, and the 10-year Treasury rate in 2021Q1–2023Q1 match the Fed's adverse projections. The shaded bands correspond to the 68% coverage intervals while the solid and dotted black lines denote the in-sample and out-of-sample values.
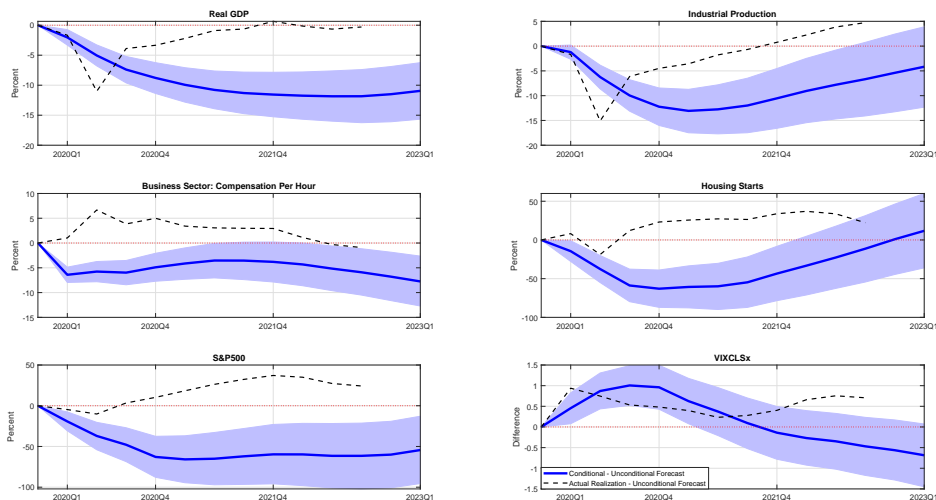

Figure 8: Difference between the conditional and unconditional forecasts when CPI inflation, unemployment, and the 10-year Treasury rate in 2021Q1–2023Q1 match the Fed's stress test adverse projections. Bands denote to the 68% coverage intervals. The black dashed line denotes the difference between the actual realizations and the posterior mean of the unconditional 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 hard and soft 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 the proposed approach generates exactly the same conditional forecasts and credible sets as those from Waggoner and Zha (1999), Bańbura, Giannone, and Lenza (2015), and Antolín-Díaz, Petrella, and Rubio-Ramírez (2021), but is substantially less demanding computationally. Finally, we conducted an empirical exercise where 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 soft and hard constraints on the trajectories of CPI inflation, the unemployment rate, and the 10-year Treasury rate over the 2020–2022 period. 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, McGillicuddy, and Owyang (2021).

References

AASTVEIT, K. A., A. CARRIERO, T. E. CLARK, AND M. MARCELLINO (2017): "Have Standard VARS Remained Stable Since the Crisis?," *Journal of Applied Econometrics*, 32(5), 931–951.

ALTAVILLA, C., D. GIANNONE, AND M. LENZAA (2016): "The financial and macroeconomic effects of the OMT announcements," *International Journal of Central Banking*, 12(3), 29 – 57.

ANDERSSON, M. K., S. PALMQVIST, AND D. F. WAGGONER (2010): "Density-

conditional forecasts in dynamic multivariate models," Discussion paper, Sveriges Riksbank Working Paper Series.

ANTOLÍN-DÍAZ, J., I. PETRELLA, AND J. F. RUBIO-RAMÍREZ (2021): "Structural scenario analysis with SVARs," *Journal of Monetary Economics*, 117, 798–815.

BAŃBURA, M., D. GIANNONE, AND M. LENZA (2015): "Conditional forecasts and scenario analysis with vector autoregressions for large cross-sections," *International Journal of forecasting*, 31(3), 739–756.

BAUMEISTER, C., AND L. KILIAN (2014): "Real-time analysis of oil price risks using forecast scenarios," *IMF Economic Review*, 62(1), 119–145.

BOTEV, Z. I. (2017): "The normal law under linear restrictions: simulation and estimation via minimax tilting," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 79(1), 125–148.

CHAN, J. C. C. (2022): "Asymmetric conjugate priors for large Bayesian VARs," *Quantitative Economics*, 13(3), 1145–1169.

CHAN, J. C. C., AND I. JELIAZKOV (2009): "Efficient simulation and integrated likelihood estimation in state space models," *International Journal of Mathematical Modelling and Numerical Optimisation*, 1(1-2), 101–120.

CLARIDA, R. H., AND D. COYLE (1984): "Conditional Projection by Means of Kalman Filtering," Discussion paper, NBER Technical Working Papers 0036. National Bureau of Economic Research, Inc.

CRUMP, R. K., S. EUSEPI, D. GIANNONE, E. QIAN, AND A. M. SBORDONE (2021): "A Large Bayesian VAR of the United States Economy," *FRB of New York Staff Report No. 976*.

DIEPPE, A., R. LEGRAND, AND B. VAN ROYE (2016): "The BEAR toolbox," *ECB working paper*.

DOAN, T., R. LITTERMAN, AND C. SIMS (1984): "Forecasting and conditional projection using realistic prior distributions," *Econometric reviews*, 3(1), 1–100.

GENZ, A. (1992): "Numerical computation of multivariate normal probabilities," *Journal of computational and graphical statistics*, 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*, pp. 248–263. Springer.

GIANNONE, D., M. LENZA, D. MOMFERATOU, AND L. ONORANTE (2014): "Short-term inflation projections: A Bayesian vector autoregressive approach," *International journal of forecasting*, 30(3), 635–644.

GIANNONE, D., M. LENZA, H. PILL, AND L. REICHLIN (2012): "The ECB and the interbank market," *The Economic Journal*, 122(564), F467–F486.

GIANNONE, D., M. LENZA, AND L. REICHLIN (2019): "Money, credit, monetary policy, and the business cycle in the euro area: What has changed since the crisis?," *International Journal of Central Banking*, 15(5), 137 – 173.

JAROCIŃSKI, M. (2010): "Conditional forecasts and uncertainty about forecast revisions in vector autoregressions," *Economics Letters*, 108(3), 257–259.

JAROCINSKI, M., AND F. SMETS (2008): "House prices and the stance of monetary policy," *ECB Working Paper*.

LENZA, M., H. PILL, AND L. REICHLIN (2010): "Monetary policy in exceptional times," *Economic Policy*, 25(62), 295–339.

MCCRACKEN, M. W., J. T. MCGILLICUDDY, AND M. T. OWYANG (2021): "Binary Conditional Forecasts," *Journal of Business & Economic Statistics*, 0(0), 1–13.

ROBERTSON, J. C., E. W. TALLMAN, AND C. H. WHITEMAN (2005): "Forecasting using relative entropy," *Journal of Money, Credit, and Banking*, 37(3), 383–401.

SIMS, C. A., AND T. ZHA (1998): "Bayesian Methods for Dynamic Multivariate Models," *International Economic Review*, 39(4), 949–968.

TALLMAN, E. W., AND S. ZAMAN (2020): "Combining survey long-run forecasts and nowcasts with BVAR forecasts using relative entropy," *International Journal of Forecasting*, 36(2), 373–398.

WAGGONER, D. F., AND T. ZHA (1999): "Conditional forecasts in dynamic multivariate models," *Review of Economics and Statistics*, 81(4), 639–651.

Appendices

A Data and Transformations

Table A.1: This table lists all the variables we included in our empirical application. For each series, we report the FRED-QD mnemonic and the transformation we adopted. 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, Eusepi, Giannone, Qian, and Sbordone (2021).

Variable	Mnemonic	Transformation
Real Gross Domestic Product	GDPC1	$100\ln(x_t)$
Real Personal Consumption Expenditures	PCECC96	$100\ln(x_t)$
Real Private Fixed Investment: Residential	PRFIx	$100\ln(x_t)$
Real Private Fixed Investment: Non-Residential	PNFIx	$100\ln(x_t)$
Real Exports of Goods and Services	EXPGSC1	$100\ln(x_t)$
Real Imports of Goods and Services	IMPGSC1	$100\ln(x_t)$
Real Government Consumption Expenditures and Gross Investment	GCEC1	$100\ln(x_t)$
Real Government Consumption Expenditures and Gross Investment: Federal (Chain-Type Quantity Index)	B823RA3Q086SBEA	$100\ln(x_t)$
Gross Domestic Product: Chain-Type Price Index	GDPCTPI	$100\ln(x_t)$
Producer Price Index by Commodity: All Commodities	PPIACO	$100\ln(x_t)$
Personal Consumption Expenditures Excluding Food and Energy (Chain-Type Price Index)	PCEPILFE	$100\ln(x_t)$
Consumer Price Index	CPIAUCSL	$100\ln(x_t)$
Consumer Price Index for All Urban Consumers: All Items Less Food and Energy in U.S. City Average	CPILFESL	$100\ln(x_t)$
Business Sector: Real Hourly Compensation for All Employed Persons	RCPHBS	$100\ln(x_t)$
All Employees, Total Nonfarm	PAYEMS	$100\ln(x_t)$
Unemployment Rate	UNRATE	Level
Industrial Production: Total Index	INDPRO	$100\ln(x_t)$
Capacity Utilization: Manufacturing	CUMFNS	$100\ln(x_t)$
New Privately-Owned Housing Units Started: Total Units	HOUST	$100\ln(x_t)$
Real Disposable Personal Income	DPIC96	$100\ln(x_t)$
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	$100\ln(x_t)$
S&P 500	S&P 500	$100\ln(x_t)$
CBOE Volatility Index: VIX	VIXCLSx	Level
Personal Consumption Expenditures: Chain-type Price Index	PCECTPI	$100\ln(x_t)$
Real Crude Oil Prices: West Texas Intermediate (WTI)	OILPRICEx	$100\ln(x_t)$
Federal Funds Effective Rate	FEDFUNDS	Level