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

Joshua Chan[1] Davide Pettenuzzo[2] Aubrey Poon[3] Dan Zhu[4]

June 2023

[1]Purdue University
[2]Brandeis University
[3]Orebro University
[4]Monash University

Motivation

- Conditional forecasts are projections of a set of variables of interest on the future paths of some other variables.
- They are used routinely by empirical macroeconomists in a number of applied settings.
- Since the seminal work of Waggoner and Zha (1999) ReStat paper, conditional forecasts are often paired with VARs.
- They are used 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 funds rate or real GDP.

Motivation (cont.)

- There are two types of conditional forecasts:
  1. Traditional (reduced-form) case
     - Conditions are generated by all the structural shocks of the model, e.g conditional on observables.
  2. Structural scenarios analysis
     - Conditions are generated from a sequence of specific shocks from a (point or set) identified VAR, e.g. conditional on shocks.
- A recent JME paper by Antolin-Diaz, Petrella, and Rubio-Ramirez (2021) introduce a unified framework for conditional forecasts and structural analysis within VAR models.

Motivation (cont.)

- In the traditional case, there are two types of conditional forecasts that one can make: **hard** and **soft** conditioned forecasts
  - Hard conditions: one wants to fix the future paths of the conditioned variables at specific values.
  - Soft conditions: one allows the future values of the conditioned variables to lie within a certain range.
- The hard conditions are the most commonly employed in the empirical 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)).

Motivation (cont.)

- The soft-constrained conditional forecasts are more scarce in the literature
  - This is due to the computational challenges associated with generating the conditional forecasts.
  - Waggoner and Zha (1999) rely on an acceptance-rejection algorithm that requires a large number of simulated draws to satisfy the constraints.
  - Andersson, Palmqvist and Waggoner (2010) is the only other known study in the literature.

Motivation (cont.)

- Often the case that one does not know the actual future realization path of the constrained endogenous variables.
- In these situations, it is much simpler to impose that the future values of the variables conditioned on will be between a range or an interval instead of an exact path.
- Soft constraints allow the forecaster to acknowledge the uncertainty surrounding the future realization path of the constrained endogenous variables.

Main contribution

- We introduce a novel precision-based approach to conditional forecasting that generalizes and extends the existing methods available in the literature in several ways.
- Similar to Antolin-Diaz, Petrella, and Rubio-Ramirez (2021), our precision-based approach is closed-form and can be used for both conditional forecast and structural scenario analysis.
- Our precision-based 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 (**hard** or **soft**) conditioning variables and long forecast horizons.
- Our proposed framework is similar to Chan, Poon and Zhu (2023) (*Forthcoming at Journal of Econometrics*) in which they apply the precision-based approach to high-dimensional state-space models with missing data.

Main contribution (cont.)

- We develop a fast and efficient method for simulating the conditional forecasts from a (or multiple) soft constraint(s).
  - To do this, we combine the precision sampler of Chan and Jeliazkov (2009) with the exponential minmax tilting method of Botev (2017).
  - The precision sampler exploits fast band matrix algorithms.
- Botev (2017) provides an efficient method to generate random draws from a high-dimensional Gaussian distribution under linear restrictions.
  - The algorithm is an accept-reject sampler, but it is constructed in a way that the proposal distribution satisfies all the constraints via exponential tilting.

Outline of the Presentation

- We first proposed a general framework for conditional forecasts.
- We derive both hard and soft constrained conditional forecast distributions.
- In a simulation study, we compare our novel precision-based conditional forecasting sampler to four existing approaches in the literature:
  - Waggoner and Zha (1999)
  - Banbura et al. (2015) - Filtering/Smoothing methods
  - Antolin-Diaz, Petrella, and Rubio-Ramirez (2021)
  - Andersson, Palmqvist and Waggoner (2010)
- Lastly, we apply our novel precision-based conditional forecast sampler to a Large BVAR with multiple hard and soft constraints in the empirical application.

Unconditional Forecast from a Structural VAR

- Let us first define $n \times 1$ vector of variables $\mathbf{y}_t = (y_{1,t}, \ldots, y_{n,t})'$ structural VAR with p lags
$$\mathbf{A}_0\mathbf{y}_t = \mathbf{a} + \mathbf{A}_1\mathbf{y}_{t-1} + \ldots + \mathbf{A}_p\mathbf{y}_{t-p} + \epsilon_t, \epsilon_t \sim N(\mathbf{0}_n, \mathbf{I}_n),$$
- 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 the observables for the next h periods, $\mathbf{y}_{T+1,T+h} = (\mathbf{y}'_{T+1}, \ldots, \mathbf{y}'_{T+h})'$, can be written as
$$\mathbf{H}\mathbf{y}_{T+1,T+h} = \mathbf{c} + \epsilon_{T+1,T+h}, \quad \epsilon_{T+1,T+h} \sim N(\mathbf{0}_n, \mathbf{I}_n),$$

Unconditional Forecast from a Structural VAR (cont.)

- Thus, we can define both $\mathbf{c}$ and $\mathbf{H}$ as

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

- where $\mathbf{0}_{n\times n}$ is an $n \times n$ zero matrix.

Unconditional Forecast from a Structural VAR (cont.)

- 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. Therefore, we can show that

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

- 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 above unconditional distribution, even when both n and h are large.

Conditional Forecasts

- We can write the conditional forecasts as a set of linear restrictions on the path of future observables $\mathbf{y}_{T+1,T+h}$

$$\mathbf{R}\mathbf{y}_{T+1,T+h} \sim N(\mathbf{r}, \Omega), \tag{2}$$

- where $\mathbf{R}$ is a $r \times nh$ constant matrix with full row rank (so that there are no redundant restrictions), $\mathbf{r}$ and Ω are $r \times 1$ and $r \times r$ matrices representing the mean and covariance of the restrictions.
- The above setup is very general and can accommodate both the hard, e.g setting $\Omega = \mathbf{0}_{r\times r}$, and soft constraints.

Conditional Forecasts (cont.)

- We can combine both (1) and (2)

$$\mathbf{R}\mathbf{y}_{T+1,T+h} = \mathbf{R}\mathbf{H}^{-1}\mathbf{c} + \mathbf{R}\mathbf{H}^{-1}\epsilon_{T+1,T+h}, \quad \epsilon_{T+1,T+h} \sim N(\mathbf{r}, \Omega), \tag{3}$$

- Next, we can derive the set of restrictions on the future shocks implied by both (2) and (3).

$$\epsilon_{T+1,T+h} | \mathbf{R}, \mathbf{r}, \Omega \sim N(\mu_\epsilon, \mathbf{I}_{nh} + \Psi_\epsilon), \tag{4}$$

- where μ_ϵ and Ψ_ϵ are, respectively, the deviations of the mean vector and covariance matrix of the restricted future shocks from their unconditional counterparts in (1).

Conditional Forecasts (cont.)

- Combining both (3) and (4) implies the following restriction on μ_ϵ and Ψ_ϵ

$$\begin{aligned} \mathbf{R}\mathbf{H}^{-1}(\mathbf{c}+\mu_\epsilon) &= \mathbf{r} \\ \mathbf{R}\mathbf{H}^{-1}(\mathbf{I}_{nh}+\Psi_\epsilon)\mathbf{H}^{-1'}\mathbf{R}' &= \Omega. \end{aligned} \tag{5}$$

- When $r < nh$, the above (5) is underdetermined and has multiple solution.

- We follow Antolin-Diaz, Petrella, and Rubio-Ramirez (2021) and choose a solution that can be expressed in terms of the Moore-Penrose inverse of $\mathbf{RH}^{-1}$, which we denote as $(\mathbf{RH}^{-1})^{+}$

$$\begin{aligned} \mu_\epsilon &= (\mathbf{RH}^{-1})^{+}(\mathbf{r} - \mathbf{RH}^{-1}\mathbf{c}) \\ \Psi_\epsilon &= (\mathbf{RH}^{-1})^{+}(\Omega - \mathbf{R}(\mathbf{H}'\mathbf{H})^{-1}\mathbf{R}')(\mathbf{RH}^{-1})^{+\prime}. \end{aligned} \tag{6}$$

- This means the solution represents the smallest deviations of the mean vector and covariance matrix of the conditional future shocks from the unconditional ones.

Conditional Forecasts (cont.)

- Using the previous result, we can map the constraints on the future shocks implied by (2) and (3) to the corresponding constraints on the forecasts.
- Let us denote the conditional forecast distribution as

$$\mathbf{y}_{T+1,T+h}|\mathbf{R},\mathbf{r},\Omega \sim N(\mu_{\mathbf{y}},\Sigma_{\mathbf{y}}),$$

- Then, (1) and (6) imply that

$$\mu_{\mathbf{y}} = \mathbf{H}^{-1}\left[\mathbf{c} + (\mathbf{R}\mathbf{H}^{-1})^{+}(\mathbf{r} - \mathbf{R}\mathbf{H}^{-1}\mathbf{c})\right],$$

$$\Sigma_{\mathbf{y}} = \mathbf{H}^{-1}\left[\mathbf{I}_{nh} + (\mathbf{R}\mathbf{H}^{-1})^{+}(\Omega - \mathbf{R}(\mathbf{H}'\mathbf{H})^{-1}\mathbf{R}')(\mathbf{R}\mathbf{H}^{-1})^{+'}\right]\mathbf{H}^{-1'}.$$

- This result is extremely general and, in fact, encompasses a number of useful and popular applications of conditional forecasting.

Conditional Forecasts (cont.)

- We can extend our general conditional forecasting framework to accommodate soft conditioning, where one allows the future values of the conditioned variables to lie within a certain range.

$$\underline{\mathbf{c}} < \mathbf{S}\mathbf{y}_{T+1,T+h} < \bar{\mathbf{c}}, \tag{7}$$

- where $\mathbf{S}$ is a $s \times nh$ pre-specified full-rank constant matrix, $\underline{\mathbf{c}}$ and $\bar{\mathbf{c}}$ are $s \times 1$ constant vectors (with elements in $\mathbb{R} \cup \{\pm\infty\}$) and the inequalities hold component-wise.
- $\mathbf{S}$ can be a selection matrix and also allows for inequality restrictions on any linear combinations of the variables.
- Thus, from equations (1) and (7) implies $\mathbf{y}_{T+1,T+h}$ has a truncated multivariate normal distribution

$$\mathbf{y}_{T+1,T+h} | \underline{\mathbf{c}} < \mathbf{S}\mathbf{y}_{T+1,T+h} < \bar{\mathbf{c}} \sim N(\mathbf{H}^{-1}\mathbf{c}, (\mathbf{H}'\mathbf{H})^{-1}) 1(\underline{\mathbf{c}} < \mathbf{S}\mathbf{y}_{T+1,T+h} < \bar{\mathbf{c}}), \tag{8}$$

- where $1(\cdot)$ is the indicator function.

Special Case: Hard Constraint on Future Observables

- Hard constraint example: A policymaker might be interested in the the future path of GDP conditioned on the scenario that the future policy rate follows a fixed path across the forecast horizon.
- This type of restriction can be written as

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

- 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.
- This setting can be nested within our general framework by setting $\mathbf{R} = \mathbf{R}_o$, $\mathbf{r} = \mathbf{r}_o$ and $\Omega = \mathbf{0}_{r_o \times r_o}$.

Special Case: Hard Constraint on Future Observables (cont.)

- Next, we can first partition the $nh \times 1$ vector $\mathbf{y}_{T+1,T+h}$ into
  - $\mathbf{y}^o_{T+1,T+h}$ is a $r_o \times 1$ vector of the hard-constrained endogenous variables—the set of variables that are selected by $\mathbf{R}_o$.
  - $\mathbf{y}^u_{T+1,T+h}$ is a $(nh - r_o) \times 1$ vector of free or unconstrained variables.
- Let $\mathbf{R}^-_o$ denote the associated $(nn - r_o) \times nh$ selection matrix that selects $\mathbf{y}^u_{T+1,T+h}$. Then, we can write $\mathbf{y}_{T+1,T+h}$ as follows:

$$\mathbf{y}_{T+1,T+h} = \mathbf{M}_u \mathbf{y}^u_{T+1,T+h} + \mathbf{M}_o \mathbf{y}^o_{T+1,T+h}, \tag{9}$$

- 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.

Special Case: Hard Constraint on Future Observables (cont.)

- Next, we can 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 $\mathbf{A} = (\mathbf{a}, \mathbf{A}_1, \ldots, \mathbf{A}_p)'$, by substituting (9) into (1)

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

- Hence, 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

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

Special Case: Hard Constraint on Future Observables (cont.)

- Thus,

$$\mathbf{K}_u = \mathbf{M}_u'\mathbf{H}'\mathbf{H}\mathbf{M}_u \ , \ \mu_u = \mathbf{K}_u^{-1}\mathbf{M}_u'\mathbf{H}'\mathbf{H}(\mathbf{H}^{-1}\mathbf{c} - \mathbf{M}_o\mathbf{y}^o_{T+1,T+h}).$$

- That is,

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

- 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.

Special Case: Soft Constraint on Future Observables

- Returning to the previous example, instead of conditioning on a fixed path of the future policy rate, we can restrict the future path of the policy rate to be, say, between 1% and 2% in the next 8 quarters and between 1.5% and 2.5% afterward.
- Recall from (8)

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

- How do we draw from this truncated multivariate normal distribution?

Special Case: Soft Constraint on Future Observables (cont.)

- To draw from this truncated multivariate normal distribution, we employed the minmax tilting method of Botev (2017), which is a generic algorithm to efficiently sample from a potentially high-dimensional Gaussian distribution under linear inequality restrictions.
- We combined this minmax tilting method with the precision sampler of Chan and Jeliazkov (2009).

Special Case: Soft Constraint on Future Observables (cont.)

- How do we simulate a draw of the conditional forecasts $\mathbf{y}_{T+1,T+h}$ subject to a soft constraint:
  1. We first sample $\mathbf{y}^o_{T+1,T+h}$ marginally from its truncated s_o-variate Gaussian distribution using the algorithm of Botev (2017).
  2. 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 step of sampling from the truncated Gaussian from nh to s_o.

Simulation study

- We conduct two simulation studies.
- Hard constraint
  - Precision-based method
  - Waggoner and Zha (1999) (WZ)
  - Banbura et al. (2015) - Filtering/Smoothing methods (DK)
  - Antolin-Diaz, Petrella, and Rubio-Ramirez (2021) (APR)
- Soft constraint
  - Precision-based method
  - Waggoner and Zha (1999)
  - Andersson, Palmqvist and Waggoner (2010) (APW)

Simulation study (cont.)

- We consider a data-generating process (DGP) that follows an n-variables VAR structure with $p = 2$ lags

$$\mathbf{y}_t = \mathbf{b} + \mathbf{B}_1 \mathbf{y}_{t-1} + \mathbf{B}_2 \mathbf{y}_{t-2} + \epsilon_t, \epsilon_t \sim N(0, \Sigma), \tag{10}$$

- We set $T = 300$, $\mathbf{b} = 0.01 \times \mathbf{1}_n$, where $\mathbf{1}_n$ is an $n \times 1$ column of ones.
- We generate the diagonal of elements of $\mathbf{B}_1$ from $U(0, 0.5)$ and off diagonal elements from $U(-0.2, 0.2)$.
- All the other elements of the higher VAR coefficients are generated independently from $N(0, 0.05^2/p^2)$.
- Finally, we generate the covariance matrix from $IW(n + 10, 0.07\mathbf{I}_n + 0.03\mathbf{1}_n\mathbf{1}_n^{'})$.
- We estimate the above model using standard uninformative normal priors for the VAR coefficients and inverse-Wishart prior for the covariance matrix.

Simulation study - Hard constraint

- We consider
  - A medium VAR with a short forecast horizon ($n = 8$, $h = 5$, and $r_o = 3$) with three constrained variables.
  - A large VAR with a long forecast horizon ($n = 15$, $h = 20$, and $r_o = 3$) with three constrained variables.
- We estimate the VAR models using 25000 MCMC draws with a burn-in period of 10000 draws.

Simulation study - Hard constraint (cont.)

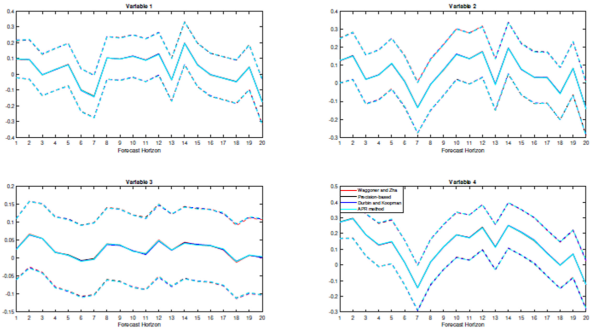

Figure 2: **Conditional forecasts from a large VAR with three hard constraints.** The thick black line is the posterior median estimate of the conditional forecast using our precision-based method. The thick red line is the posterior median estimates of the conditional forecast using the Waggoner and 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.

Simulation study - Hard constraint (cont.)

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.

Simulation study - Soft constraint

- We follow the same simulation structure as described above in the hard constraint case, except that for the soft constraint, we restrict the actual constrained (conditioned) variable $\mathbf{y}^{o}_{t,T-h:T}$ to be between some intervals.
- More specifically, we set $\bar{\mathbf{y}}^{o}_{t,T-h:T} - 0.1 \preceq \mathbf{y}^{o}_{t,T-h:T} \preceq \bar{\mathbf{y}}^{o}_{t,T-h:T} + 0.1$ and $\bar{\mathbf{y}}^{o}_{t,T-h:T} = \frac{1}{h}\sum_{t=T-h}^{T} \mathbf{y}^{o}_{t,}$ is the average of the actual simulated data across this forecast period.
- In this simulation exercise, we consider an eight variables VAR where $n = 8$ and a long forecast horizon $h = 20$ with only one soft constraint $n_o = 1$.
- We also estimate the eight variables VAR using 25000 MCMC draws with a burn-in period of 10000 draws and implement the same priors as in the hard constraint exercise.

Simulation study - Soft constraint (cont.)

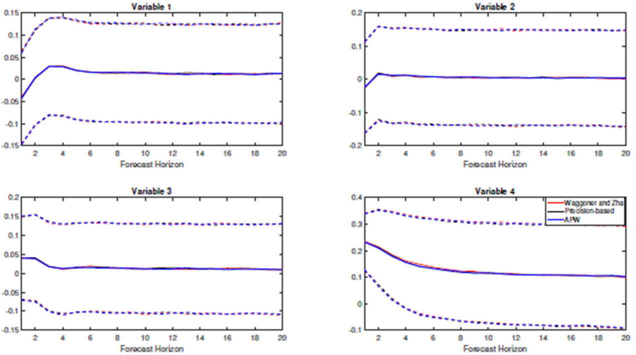

Figure 3: **Conditional forecast from a medium VAR with one soft constraint.** The thick black line is the posterior median estimate of the conditional forecast using the proposed precision-based method. The thick red line is the posterior median estimates of the conditional forecast using the Waggoner 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.

Simulation study - Soft constraint (cont.)

- To compute the conditional forecasts:
  - Our precision-based method took about 62 seconds.
  - Andersson, Palmqvist and Waggoner (2010) took about 70 seconds.
  - Waggoner and Zha (1999) accept-reject algorithm took approximately 2,542 minutes.

Simulation study - Soft constraint (cont.)

	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.

Empirical Application

- We apply our proposed method on a large BVAR model with $p = 4$ lags for 31 quarterly variables.
- The 31 quarterly variables are similar to the dataset chosen in Crump et al. (2021) NY Fed Working paper.
- We use a new asymmetric natural conjugate Minnesota-type prior for large BVARs proposed by Chan (2021) QE.
- We 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.

Empirical Application (cont.)

- Our sample period spans from 1976 to the end of 2019.
- We 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 2020Q1-2023Q1.
- We implement our constraints to mimic the baseline and adverse scenarios prepared by the Federal Reserve Board for their 2020 stress test analysis.

Summary of hard and soft constraints – Baseline scenario

	CPI Inflation (Soft Constraint)			Hard Constraints	
Date	Lower Bound	Fed's projection	Upper Bound	UNRATE	GS10
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

Summary of hard and soft constraints – Adverse scenario

	CPI Inflation (Soft Constraint)			Hard Constraints	
Date	Lower Bound	Fed's projection	Upper Bound	UNRATE	GS10
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

Empirical Application (cont.)

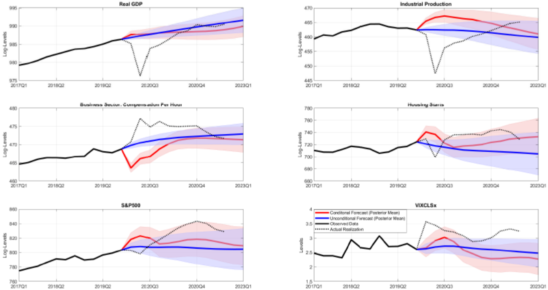

Figure 5: Conditional and unconditional forecasts when CPI inflation, unemployment, and the 10-year Treasury rate in 2021Q1–2023Q1 match the Fed's stress test baseline projections. The shaded bands correspond to the 68% coverage intervals while the solid and dotted black lines denote the in-sample and out-of-sample values.

Empirical Application (cont.)

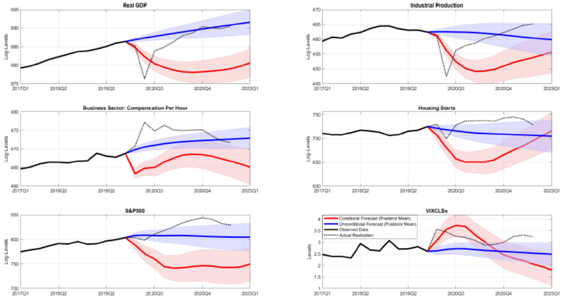

Figure 7: Conditional and unconditional forecasts when CPI inflation, unemployment, and the 10-year Treasury rate in 2021Q1–2023Q1 match the Fed's adverse projections. The shaded bands correspond to the 68% coverage intervals while the solid and dotted black lines denote the in-sample and out-of-sample values.

Conclusion

- 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.
- 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 a long forecast horizons.
- We have shown in a simulation study that the proposed approach generates exactly the same conditional forecasts and credible sets as those from Waggoner and Zha (1999), Banbura et al. (2015), Andersson, Palmqvist and Waggoner (2010), and Antolin-Diaz, Petrella, and Rubio-Ramirez (2021), but it is substantially less demanding computationally.

Conclusion (cont.)

- We conducted an empirical exercise where we estimated a Bayesian VAR featuring 31 quarterly macroeconomic and financial series.
- 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.
- Next steps:
  - Extend our framework to non-linear models, e.g. BVAR with Stochastic Volatility
  - Create a conditional forecasting toolbox for central bank policymakers and researchers.

Thank you

- Thank you for listening.
- Does anyone have any questions?