International Journal of Forecasting 31 (2015) 739–756

Contents lists available at ScienceDirect

International Journal of Forecasting

journal homepage: www.elsevier.com/locate/ijforecast

Conditional forecasts and scenario analysis with vector autoregressions for large cross-sections

Marta Bańbura [a], Domenico Giannone [b,c,d,e,*], Michele Lenza [a,c]

[a] European Central Bank, Germany
[b] LUISS, Italy
[c] ULB-ECARES, Belgium
[d] CEPR, London, United Kingdom
[e] Federal Reserve Bank of New York, USA

ARTICLE INFO

Keywords:
Vector autoregression
Bayesian shrinkage
Dynamic factor model
Conditional forecast
Large cross-sections

ABSTRACT

This paper describes an algorithm for computing the distribution of conditional forecasts, i.e., projections of a set of variables of interest on future paths of some other variables, in dynamic systems. The algorithm is based on Kalman filtering methods and is computationally viable for large models that can be cast in a linear state space representation. We build large vector autoregressions (VARs) and a large dynamic factor model (DFM) for a quarterly data set of 26 euro area macroeconomic and financial indicators. The two approaches deliver similar forecasts and scenario assessments. In addition, conditional forecasts shed light on the stability of the dynamic relationships in the euro area during the recent episodes of financial turmoil, and indicate that only a small number of sources drive the bulk of the fluctuations in the euro area economy.

1. Introduction

Vector autoregressions (VARs) are very flexible and general models, and provide a reliable empirical benchmark for alternative econometric representations such as dynamic stochastic general equilibrium (DSGE) models, which are more grounded in theory but, at the same time, impose more structure on the data (see, for example, Christiano, Eichenbaum, & Evans, 1999, Chapter 2).

Recently, the literature has shown that VARs are also viable tools for large sets of data (see Bańbura, Giannone, & Reichlin, 2010). In this paper, we construct a large VAR for the euro area and apply it to unconditional forecasting, as well as to conditional forecasts and scenarios. These, along with structural analysis (assessing the effects of structural shocks), have been the main applications of VARs to date. While large VARs have been used for unconditional forecasting and structural analysis,[1] limited attention has been devoted thus far to conditional forecasting. This is because popular algorithms for deriving conditional forecasts have been challenging computationally for large data sets. We overcome this problem by computing the conditional forecasts recursively using Kalman filtering techniques.

Conditional forecasts, and scenarios in particular, are projections of a set of variables (that are of interest) on the future paths of some other variables. This is in contrast

* Corresponding author at: Federal Reserve Bank of New York, USA.
E-mail addresses: marta.banbura@ecb.europa.eu (M. Bańbura), domenico.giannnone@ny.frb.org (D. Giannone), michele.lenza@ecb.europa.eu (M. Lenza).

[1] See for example Giannone, Lenza, Momferatou, and Onorante (2014); Koop (2013); Giannone, Lenza, and Reichlin (2012); Paciello (2011); and Giannone, Lenza, and Primiceri (in press).

http://dx.doi.org/10.1016/j.ijforecast.2014.08.013
0169-2070/

to unconditional forecasts, where no knowledge of the future paths of any variables is assumed. The prior knowledge, albeit imperfect, of the future evolution of some economic variables may contain information for the outlooks of other variables. For example, future fiscal packages would affect the future evolution of economic activity, and thus, might provide important off-model information. Moreover, it may be of interest to assess the impacts of specific future events on a set of variables, i.e., to conduct scenario analyses. Notable examples of the latter are the stress tests recently conducted in the US and the euro area in order to assess the vulnerability of their banking systems. For recent examples of conditional forecasts, see Bloor and Matheson (2011), Giannone et al. (2014), Giannone, Lenza, Pill, and Reichlin (2012), Giannone, Lenza, and Reichlin (2010), Jarociński and Smets (2008), Lenza, Pill, and Reichlin (2010) and Stock and Watson (2012a). Recently, Clark and McCracken (2014) proposed and evaluated a range of tests of predictive ability for conditional forecasts from estimated models.

The scenario analysis described above and studied in this paper can be considered as "reduced form", in the sense that the forecasts are conditional on observables, and the identification of structural shocks is not required.[2] Note that, if necessary, the structural shocks that are "compatible" with the scenario can be retrieved from the reduced form innovations using some identifying assumptions. An alternative approach consists of constructing scenarios by manipulating specific structural shocks so that the resulting paths of the observed variables are consistent with the conditioning information (see also Adolfson, Laséen, Lindé, & Villani, 2005; Christoffel, Coenen, & Warne, 2007; Leeper & Zha, 2003; Luciani, in press). Along similar lines, Baumeister and Kilian (2013) construct scenarios for the real price of oil from a VAR by conditioning on a sequence of appropriately derived structural shocks rather than on a pre-specified path for observables.

For VAR models, the conditional forecasts are typically computed using the algorithm developed by Waggoner and Zha (1999). Roughly speaking, the methodology involves drawing (the entire) paths of reduced form shocks which are compatible with the conditioning path on the observables. The computational burden of this approach means that it can easily become impractical or unfeasible for high dimensional data and long forecast horizons, even if the computationally more efficient version of Jarociński (2010) is employed. However, many problems in macroeconomics and finance can only be addressed by looking at the joint dynamic behaviors of a large number of time series. For example, business cycle research, as in the NBER tradition, typically involves the analysis of many macroeconomic variables. Professional forecasters and policymakers look at a variety of different indicators in order to predict key variables of interest and to make their decisions. Investors analyze the joint behaviors of many asset returns in order to choose the optimal portfolios. More generally, contemporary science relies more and more on the availability and exploitation of large data sets.

In this paper, building on an old insight of Clarida and Coyle (1984), we propose an algorithm based on Kalman filtering techniques for computing the conditional forecasts. Since the Kalman filter works recursively, i.e., period by period, this algorithm reduces the computational burden significantly for longer forecast horizons, and is particularly well suited for empirical approaches where large data sets are being handled. Using a simulation smoother (for examples of simulation smoothers, see Carter & Kohn, 1994; de Jong & Shephard, 1995; Durbin & Koopman, 2002) allows for the computation of the full distribution of conditional forecasts. The algorithm applies to any model which can be cast in a linear state space representation. For the VAR framework, we compare the computational efficiencies of different simulation smoothers and find that, for large systems, the simulation smoother of Durbin and Koopman (2002) can offer substantial computational gains over the more popular algorithm of Carter and Kohn (1994).

The interest in issues which are best addressed by considering large information sets raises a trade-off between the excessive simplicity of the models, leading to misspecification due to omitted variables, and their excessive complexity, with many free parameters leading to a large estimation uncertainty. Recent developments in macroeconometrics have suggested two approaches for dealing with the complexity of large sets of data without losing their salient features: Bayesian VARs (BVARs) and dynamic factor models (DFMs).

The aforementioned flexibility of VARs comes at the cost of a large number of free parameters requiring estimation. Specifically, for a generic VAR(p) model for a vector of n variables $y_t = (y_{1,t}, \dots, y_{n,t})'$:

$$\begin{aligned} y_t &= c + A_1 y_{t-1} + \cdots + A_p y_{t-p} + \varepsilon_t, \\ \varepsilon_t &\sim WN(0, \Sigma), \end{aligned} \tag{1}$$

where $WN(0, \Sigma)$ refers to a white noise process with mean 0 and covariance matrix Σ, we count: (i) pn^2 parameters in autoregressive matrices, $A_1, \dots, A_p$, that are of dimension $n \times n$ each; (ii) $n(n+1)/2$ free parameters in the $n \times n$ covariance matrix of residuals Σ; and (iii) n parameters in the constant term c. The number of parameters proliferates as the number of variables in the model increases, making estimation unreliable or unfeasible. For example, when the number of variables in a VAR with four lags increases from six, as in the original VAR model proposed by Sims (1980), to 20, 50 or 100, the total number of parameters to be estimated goes from 171 to numbers in the order of 2, 10 and 50 thousands, respectively. Such large numbers of parameters cannot be estimated well by ordinary least squares, for example, since the typical macroeconomic sample involves limited numbers of data points (in the best case, 50–60 years of data, i.e., 200–250 data points with quarterly data). This parameter proliferation problem that prevents econometricians from conducting reliable inference with large dimensional systems is also known as the "curse of dimensionality".

One solution to the curse of dimensionality in the VAR framework is to adopt Bayesian shrinkage. The idea of this

[2] For a discussion of the invariance of the conditional forecast distribution to alternative identification assumptions for structural shocks, see Waggoner and Zha (1999).