ie | BCRA

INVESTIGACIONES ECONÓMICAS

Economic Research
Working Papers | 2018 | N° 79

Forecasting Inflation in Argentina

Lorena Garegnani | Banco Central de la República Argentina
Mauricio Gómez Aguirre | Banco Central de la República Argentina

ie | BCRA
INVESTIGACIONES ECONÓMICAS

Forecasting Inflation in Argentina

Lorena Garegnani
Mauricio Gómez Aguirre
BCRA

May 2018

Resumen no técnico

Pregunta de la investigación

Uno de los objetivos del Banco Central de la República Argentina (BCRA) es inducir una baja sistemática y sostenible de la tasa de inflación, para lograr este objetivo, en 2016 el BCRA adoptó un esquema de metas de inflación. En este contexto proveer a las autoridades del BCRA proyecciones de las variables macroeconómicas relevantes, como es el caso de la inflación, es fundamental para la toma de decisiones.

Contribución

Con el fin de proyectar valores posibles de la inflación, se aplicaron distintas metodologías econométricas sobre diferentes modelos. De esta forma se utilizan la econometría frecuentista (tradicional) y la Bayesiana para realizar nuestras estimaciones sobre modelos de diferentes escalas o tamaños para luego obtener proyecciones de inflación. Estas proyecciones de la inflación son comparadas en términos de su capacidad predictiva con el fin de crear un criterio óptimo para seleccionar aquellos modelos que pronostican mejor.

Resultados

Los resultados obtenidos indican que el VAR Bayesiano presenta mejor capacidad predictiva comparado con el modelo autorregresivo considerado como base. Se han realizado test sobre distintos horizontes de tiempo para evaluar la capacidad predictiva de los distintos modelos estimados, encontrando que las estimaciones basadas en modelos VAR Bayesiano superan al modelo autorregresivo para todos los horizontes evaluados. Entre modelos VAR Bayesianos de diferentes tamaños no se encuentras diferencias significativas en su capacidad de pronóstico, no obstante, la inspección visual de los RMSE permite tener una idea de que los modelos de gran escala parecerían tener mejor performance predictiva en horizontes lejanos.

Forecasting Inflation in Argentina

Lorena Garegnani   Mauricio Gómez Aguirre

Abstract

During the year 2016, the Central Bank of Argentina has begun to announce inflation targets. In this context, providing the authorities of good estimates of relevant macroeconomic variables turns out to be crucial to make the pertinent corrections to reach the desired policy goals. This paper develops a group of models to forecast inflation for Argentina, which includes autoregressive models, and different scale Bayesian VARs (BVAR), and compares their relative accuracy. The results show that the BVAR model can improve the forecast ability of the univariate autoregressive benchmark's model of inflation. The Giacomini-White test indicates that a BVAR performs better than the benchmark in all forecast horizons. Statistical differences between the two BVAR model specifications (small and large-scale) are not found. However, looking at the RMSEs, one can see that the larger model seems to perform better for larger forecast horizons.

Promover la estabilidad monetaria significa para el BCRA inducir una baja sistemática y sostenible de la tasa de inflación, para lograr este objetivo, en 2016 el BCRA adoptó un esquema de metas de inflación. En este contexto proveer a las autoridades de proyecciones de las variables macroeconómicas relevantes como la inflación es fundamental para la toma de decisiones. En este trabajo se estima un grupo de modelos con el objetivo de generar pronósticos de inflación para Argentina. Se estiman modelos autorregresivos y modelos VAR Bayesianos con diferentes escalas y luego se compara su capacidad relativa de pronóstico. Los resultados indican que el VAR Bayesiano es quien presenta mejor capacidad predictiva comparado con el modelo autorregresivo considerado como base. El test de Giacomini-White muestra al modelo VAR Bayesiano superando al modelo autorregresivo en todos los horizontes. Entre modelos VAR Bayesianos de diferentes tamaños no se encuentran diferencias significativas en su capacidad de pronóstico, no obstante, la inspección visual de los RMSE permite tener una idea de que los modelos de gran escala parecerían tener mejor performance predictiva en horizontes lejanos.

---

1. <u>Introduction</u>

Several long-term nominal commitments such as labor contracts, mortgages and other debt are widespread features of modern economies. Forecasting how the general price level will evolve over the life of a commitment is an essential part of private sector decision-making.

The existence of long-term nominal obligations is also among the primary reasons economists generally believe that monetary policy is not neutral, at least over moderate horizons.

Central banks aim is to keep inflation stable, and perhaps also to keep output near an efficient level. With these objectives, the New Keynesian model makes explicit that optimal policy will depend on optimal forecasts (e.g., Svensson, 2005), and further, that policy will be most effective when it is well understood by the public.

Under "inflation targeting" the central banks generally released forecasts in quarterly "Inflation Reports" in a way to be more transparent in their actions. The costs and benefits of transparency are widely debated, but the need for a central bank to be concerned with inflation forecasting is broadly agreed. In short, inflation forecasting is of foremost importance to households, businesses, and policymakers.

During the year 2016, the Central Bank of Argentina has begun to announce inflation targets. In this context, providing the authorities of good estimates of relevant macroeconomic variables turns out to be crucial to make the pertinent corrections in order to reach the desired policy goals.

A standard tool in macroeconomics that is widely employed in forecasting is Vector Autoregressive (VAR) analysis. VARs are flexible time series models that can capture complex dynamic relationships among macroeconomic aggregates. However, their dense parameterization often leads to unstable inference and inaccurate out-of-sample forecasts, particularly for models with many variables, due to the estimation uncertainty of the parameters.

Litterman (1980) and Doan, Litterman, and Sims (1984) have proposed to combine the likelihood function (the data) with some informative prior distributions (the researcher's belief about the values of coefficients) to improve the forecasting performance of VAR models, introducing a Bayesian approach into VAR modeling.

In any Bayesian inference, a fundamental yet challenging step is prior specification, which influences posterior distributions of the unknown parameters and, consequently, the forecasts (Geweke, 2005). Fortunately, the literature has proposed some methodologies to set how informative the prior distributions should be.

Regarding prior selection, Litterman (1980) and Doan, Litterman, and Sims (1984) set the tightness of the prior by maximizing the out-of-sample forecasting performance of a small-scale model. Many authors follow this strategy, such as Robertson and Tallman (1999) and

Wright (2009), and Giannone et. al (2014), who minimize the Root Mean Square Error (RMSE) of the forecasts.

On the other hand, Banbura et al. (2008) propose to control the overfitting caused by the considerable number of variables in the model, by selecting the shrinkage of the coefficients in such a way as to give an adequate fitting "in-sample". Within this second selection strategy, we can find authors such as Giannone et al. (2012), Bloor and Mathenson (2009), Carriero et al. (2015) and Koop (2011).

Banbura, Giannone and Reichlin (2008) showed that by applying Bayesian VAR methodology, they were able to handle large unrestricted VARs models and therefore they demonstrated that VAR framework can be applied to empirical problems that require the analysis of more than a few set of time series. The authors showed that a Bayesian VAR is a viable alternative to factor models or panel VARs for analysis of large dynamic systems.

This paper develops a group of models to forecast inflation for Argentina, which includes autoregressive models, and different scale Bayesian VARs (BVAR), and compares their relative accuracy.

The paper is organized as follows: section (2) presents the methodological aspects related to the application of Bayesian analysis in a VAR framework, section (3) presents a brief description of the data, section (4) goes through the empirical results, and finally, section (5) concludes.

2. <u>Bayesian VAR Methodology</u>

A VAR model has the following structure

$$\boldsymbol{y}_t = \boldsymbol{c} + \boldsymbol{B_1}\boldsymbol{y}_{t-1} + \ldots + \boldsymbol{B_p}\boldsymbol{y}_{t-p} + \boldsymbol{\varepsilon}_t, \tag{1}$$

where $\mathbf{y}_t$ is a $n \times 1$ vector of endogenous variables, $\boldsymbol{\varepsilon}_t \sim N(\mathbf{0}, \boldsymbol{\Sigma})$ is a $n \times 1$ vector of exogenous shocks, $\mathbf{c}$ is a $n \times 1$ vector of constants, $\boldsymbol{B_1}$ to $\boldsymbol{B_p}$ are $n \times n$ matrices and $\boldsymbol{\Sigma}$ is $n \times n$ covariance matrix.

The BVAR coefficients are a weighted average of the prior mean (researcher's belief) and the maximum likelihood (ML) estimators (inferred from the data), with the inverse covariance of the prior and the ML estimators as weights.

Consider the following posterior distribution for the VAR coefficients

$$\boldsymbol{\beta}|\boldsymbol{\Omega} \sim N(\boldsymbol{\beta}_0, \boldsymbol{\Omega}^{-1}\xi) \tag{2}$$

where the vector $\boldsymbol{\beta}_0$ is the prior mean (whose elements will represent the coefficient in equation (1), the matrix $\boldsymbol{\Omega}$ is the known variance of the prior and ξ is a scalar parameter controlling the tightness of the prior information. Even though $\boldsymbol{\Omega}$ could have many shapes, gamma and Wishart

distributions are frequently used in the literature, since they ensure a normally distributed posterior.[2]

The conditional posterior of $\boldsymbol{\beta}$ can be obtained by multiplying the prior by the likelihood function. The posterior takes the form

$$\boldsymbol{\beta}|\boldsymbol{\Omega},\boldsymbol{y} \sim N\left(\widehat{\boldsymbol{\beta}}(\xi),\widehat{\boldsymbol{V}}(\xi)\right), \tag{3}$$

where

$$\widehat{\boldsymbol{\beta}}(\xi) \equiv vec\left(\widehat{\boldsymbol{B}}(\xi)\right), \tag{4}$$

and

$$\widehat{\boldsymbol{B}}(\xi) \equiv (\boldsymbol{x}'\boldsymbol{x}\,\boldsymbol{\Sigma}^{-1} + (\boldsymbol{\Omega}\xi)^{-1})^{-1}(\boldsymbol{x}'\boldsymbol{y}\boldsymbol{\Sigma}^{-1} + (\boldsymbol{\Omega}\xi)^{-1}\boldsymbol{\beta}_0), \tag{5}$$

$$\widehat{\boldsymbol{V}}(\xi) \equiv (\boldsymbol{x}'\boldsymbol{x}\,\boldsymbol{\Sigma}^{-1} + (\boldsymbol{\Omega}\xi)^{-1})^{-1}. \tag{6}$$

Vectors **y** and **x** represent observed data while $\boldsymbol{\beta}_0$ is a matrix where each column corresponds to the prior mean of each equation.

It is important to note that if we choose a large value for ξ, the prior will have little weight into the posterior. This translates to a large volatility of the prior and not enough information coming from the prior. On the other hand, if the ξ is set to a small value (i.e. close to zero), the prior becomes more informative and the posterior mean moves towards the prior mean. To see this point, we can express (5) as follows:

$$\widehat{\boldsymbol{B}}(\xi) \equiv \widehat{\boldsymbol{\Omega}}[\boldsymbol{\Omega}_0^{-1}\boldsymbol{\beta}_0 + (\boldsymbol{\Sigma}^{-1} \otimes \boldsymbol{x}')\boldsymbol{y}] \tag{7}$$

and

$$\widehat{\boldsymbol{\Omega}} = [\boldsymbol{\Omega}_0^{-1} + \boldsymbol{\Sigma}^{-1} \otimes \boldsymbol{x}'\boldsymbol{x}]^{-1} \tag{8}$$

If the second element between brackets in equation (7) is multiplied by $(\boldsymbol{x}'\boldsymbol{x})^{-1}(\boldsymbol{x}'\boldsymbol{x})$, we obtain the following equation

$$\widehat{\boldsymbol{B}}(\xi) \equiv \widehat{\boldsymbol{\Omega}}[\boldsymbol{\Omega}_0^{-1}\boldsymbol{\beta}_0] + \widehat{\boldsymbol{\Omega}}[\boldsymbol{\Sigma}^{-1} \otimes \boldsymbol{x}'\boldsymbol{x}\,(\boldsymbol{x}'\boldsymbol{x})^{-1}\boldsymbol{x}'\boldsymbol{y}] \tag{9}$$

$$\widehat{\boldsymbol{B}}(\xi) \equiv \widehat{\boldsymbol{\Omega}}[\boldsymbol{\Omega}_0^{-1}\boldsymbol{\beta}_0] + \widehat{\boldsymbol{\Omega}}[\boldsymbol{\Sigma}^{-1} \otimes \boldsymbol{x}'\boldsymbol{x}\,\boldsymbol{\beta}_{ols}] \tag{10}$$

As can be seen, the posterior is a weighted average between the prior and the Ordinary Least Square (OLS) estimators[3], where the weights are the reciprocal of the prior covariance matrix and the reciprocal of the OLS covariance matrix respectively. As a result, if the information contained in the data is good enough to describe the process behind it, the posterior will move towards the OLS estimators. However, it is important to underscore that, even if the available series are adequate to describe the data generating process, the researcher could still formulate hypothesis about the distribution of the parameters based on his own beliefs. That would imply

---

[2] If the posterior distributions are in the same family as the prior probability distribution, the prior and posterior are then called conjugate distributions.

[3] The OLS estimators of a VAR coincide exactly with the ML estimators conditional on the initial values.

ignoring the information contained in the data, and usually that kind of decisions are based on strong beliefs.

The issue mentioned in the last paragraph demonstrates the need to be cautious about choosing the prior mean and the hyperpriors. In the following subsections, these aspects are discussed in more detail.

2.1 <u>Level or Growth Rate</u>

It is unclear a priori whether transforming variables into their growth rates can enhance the forecast performance of a BVAR model. On one hand, the level specification can better accommodate the existence of long-run (cointegrating) relationships across the variables, which would be omitted in a VAR in differences. On the other hand, Clements and Hendry (1996) have shown that in a classical framework, differencing can improve forecasting performance in the presence of instability.

There has been little effort in the BVAR literature to compare specifications in levels versus differences. Carriero et al. (2015) work with this specific topic and found that models in growth rates generally yield more accurate forecasts than those obtained from the models in levels. However, we can find both approaches in the literature. Following the Litterman (1986) tradition, some authors considered BVARs with variables in levels (e.g., Banbura et. al (2008), Giannone et al. (2014), and Giannone et al. (2012)). Other authors used BVARs with variables in differences or growth rates (e.g., Clark and McCracken (2007) and Del Negro et al. (2004)).

As mentioned above, there is no apparent reason to opt for series in levels or in differences to work with; nevertheless, choosing a representation ex-ante, gives us information about the characteristics of the prior distribution (values of the mean prior). For example, working with variables in differences, implies that the persistence of those variables should be low, and that one should impose a number close to zero as a prior mean of the first lag, denoting low persistence in the series.

Since it is a good practice to start with some idea about the value that the prior could take and encouraged by the evidence found by Carriero et al. (2015), we have opted to work with variables in differences.

In the next subsection, we will treat the variance of the prior as another aspect of prior distribution.

2.2 <u>Choice of Hyperparameters and lag length strategy</u>

To select the hyperparameters and the lag length we will follow the strategy suggested by Banbura et al. (2008); Carriero, et al. (2015) and Giannone et al. (2012). Suppose, that a model is described by a likelihood function $p(\mathbf{y}|\boldsymbol{\theta})$ and a prior distribution $p_{\boldsymbol{\gamma}}(\boldsymbol{\theta})$, where $\boldsymbol{\theta}$ is a vector parameter of the model and $\boldsymbol{\gamma}$ is a vector of hyperparameters affecting the distribution of all the

priors of the model. It is natural to choose these hyperparameters by interpreting the model as a hierarchical one, i.e. replacing $p_\gamma(\theta)$ with $p(\theta|\gamma)$ and evaluating their posterior (Berger, 1985; Koop, 2003). In this way, the posterior can be obtained by applying Bayes' law

$$p(\gamma|y) \approx p(y|\gamma)p(\gamma), \tag{11}$$

where $p(\gamma)$ is the density of the hyperparameters and $p(y|\gamma)$ is the marginal likelihood. In turn, the marginal likelihood is the density that comes from the data when the hyperparameters change, in other words, the marginal likelihood can be obtained after integrating out the uncertainty about the parameters in the model,

$$p(\gamma|y) = \int p(y|\theta, \gamma)\, p(\theta|\gamma) d\theta. \tag{12}$$

For every conjugate prior, the density $p(\gamma|y)$ can be computed in closed form. To obtain the Bayesian hierarchical structure, it is necessary to obtain the distribution of $p(\theta)$ by integrating out the hyperparameters

$$p(\theta) = \int p(\theta, \gamma)\, p(\gamma) d\gamma. \tag{13}$$

More precisely, we can find different values of the prior distribution from different hyperparameter values and, in this way, we can represent the posterior as:

$$p(\theta, \gamma|y) = p(y|\theta, \gamma)p(\theta, \gamma)p(\gamma). \tag{14}$$

The marginal likelihood should be sufficient to discriminate among models; in this sense, we can choose models with different hyperparameters and different likelihood specification (more precisely, lags length structure). To make this point operational, we estimate different models, following Giannone et al. (2012) who introduce a procedure allowing to optimize the values of the hyperparameters that maximize the value of the marginal likelihood of the model. This implies that the hyperparameter values are not set a priori but are estimated.

Then the marginal likelihood can be estimated for every combination of hyperparameter values within specified ranges and for different lag length structures, and the optimal combination is retained as the one that maximizes that value.

2.3 Comparison Strategy

In this subsection we present some details about our strategy for model comparison. We will mention the steps that we will follow to do it and then give more details about the predictive ability tests used for comparison:

  a) Estimate a univariate AR model.
  b) Compute the relative RMSE to the AR from a).
  c) Compute the relative RMSE to the BVAR.[4]

---

[4] The mean of the predictive density is considered.

d) Run the test of Giacomini and White (2006) to compare both models.

Our benchmark is a univariate model. This means that we have at hand different statistical measures that cover both the frequentist and the Bayesian approaches. While frequentist literature tends to compare the forecasts with actual values, Bayesian literature compares the realized values with the whole posterior predictive density.

The testing methodology of Giacomini and White (2006) consists on evaluating relative forecast accuracy with a Diebold-Mariano (1995) like test, but with one central difference: the size of the in-sample estimation window is kept fixed, instead of expanding. Using the sample observations available at time "t", forecasts of $\mathbf{y}_{t+\tau}$ are produced for different "t" for given "τ" periods into the future, with rolling windows of estimation with the two models that are being compared. The sequences of forecasts are then evaluated according to some loss function and then the difference of forecast losses is computed. This way, a time series of differences of forecast losses $\Delta L_{t+\tau}(\hat{\theta})$ that depends on the estimated parameters is constructed. The test then consists on a Wald test on the coefficients of the regression of that series against a constant, the unconditional version of the test in equation (15), or against other explanatory variables, the conditional version in equation (16),

$$\Delta L_{t+\tau}(\hat{\theta}) = \mu + \varepsilon_t, \tag{15}$$

$$\Delta L_{t+\tau}(\hat{\theta}) = \beta' X_t + \varepsilon_t. \tag{16}$$

Standard errors may be calculated using the Newey-West covariances estimator, controlling for heteroskedasticity and autocorrelation. In this paper, the unconditional version is used.

The Giacomini-White test[5] has many advantages: it captures the effect of estimation uncertainty on relative forecast performance, it allows for comparison between either nested or non-nested models, and, finally, it is quite easy to compute.

2.4 **Model Specification**

We follow Banbura et al. (2008) and analyze two VAR models that incorporate variables of special interest, including indicators of real economic activity, consumer prices and monetary variables. We consider the following two alternative models:

Small-Scale Model. This is a small monetary VAR including three key variables:

a) Prices: we used the consumer price index constructed by the National Institute of Statistics and Census of Argentina (INDEC). After December 2006 until July 2012, the previous series is linked with the evolution of the consumer price index provided Provincial Institute of Statistics and Census of the Province of San Luis and, after July 2012, series is again linked with the evolution of the consumer price index of the City of Buenos Aires.[6]

---

[5] See chapter 17 of the book by Hashimzade and Thornton (2013) for a detailed discussion about this test.

[6] From December 2006 to October 2015, the index by the INDEC presented severe discrepancies with provincial and private price index, and hence was discard for that period.

b) Economic Activity: we used a monthly economic activity indicator known as EMAE (*Estimador Mensual de Actividad Económica*) published by the INDEC. The EMAE is based on the value added for each activity at a base price plus net taxes (without subsidies), and it uses weights provided by Argentina's National Accounts (2004). It tries to replicate quarterly GDP at a monthly frequency.

c) Interest Rate: we used data from the Central Bank of Argentina (BCRA) on 30 to 59-day fixed term deposit rates.

Large-Scale Model. In addition to the variables included in the small-scale model, this version also includes the rest of variables in the data set. These are detailed in the next section.

In September 2016, Argentina transitioned to an inflation targeting regime. This could generate a structural break in the mean or/and variance. To account for this possible change in the mean of the process, we incorporate a dummy variable in both specifications (Marcelino and Mizon 2000).[7]

As we compare models of different sizes, we need a strategy on how to choose the shrinkage hyperparameter as models become larger. As the dimension increases, we want more shrinkage, as suggested by the analysis in De Mol et al. (2008) to control for overfitting. We set the tightness of the prior for the model to have better in-sample fit; in this way, we are shrinking more in a larger dimension model.

3. Data

Our data set is composed of a group of 16 monthly macroeconomic variables of Argentina available on monthly frequency. Sources of the series, the transformations did on them and their stationarity characteristics are described in the appendix.

4. Results

4.1 Estimation of the BVAR Model

4.1.1 The Optimal Hyperparameters

We work with a Normal-Wishart BVAR specification. In this type of specification there are two hyperparameters and two parameters. We estimate the overall tightness λ_1, lag decay λ_3

---

[7] In the Appendix, we show the posterior estimation of the whole sample to see the effect of this. We controlled the change in the mean due the transition to an inflation targeting regime and indeed obtained a significant coefficient in both models.

and the lag length as we have described in section (2.2), and then we impose the value of the prior mean (the autoregressive coefficient) equal to zero as discussed earlier.

The hyperparameter of the overall tightness λ_1 is the standard deviation of the prior of all the coefficients in the system different from the constant. In other words, it determines how all the coefficients are concentrated around their prior means.

The term λ_3 is a decay factor, and $1/(L^{\lambda_3})$ controls the tightness on lag "L" relative to the first lag. Since the coefficients of higher order lags are more likely to be close to zero than those of lower order lags, the prior for the standard deviations of the coefficients decrease as the lag length increases. The values usually used in the literature are 1 or 2, so we settle for $\lambda_3 = 2$.

The prior variance of the parameters of $\widehat{\boldsymbol{\beta}}(\xi)$ is set according to:

$$\sigma_{ij}^2 = \left(\frac{1}{\sigma_j^2}\right)\left(\frac{\lambda_1}{L^{\lambda_3}}\right)^2 , \tag{17}$$

where σ_j^2 denotes the OLS residual variance of the autoregressive coefficient for variable "j", λ_1 is an overall tightness parameter, "L" is the current lag, and λ_3 is a scaling coefficient controlling the speed at which coefficients for lags greater than 1 converge to 0.

For exogenous variables, we define the variances as:

$$\sigma^2 = (\lambda_1 \lambda_4)^2 \tag{18}$$

The results for the hyperparameters and prior means of the small and the big scale model are shown in the table 2; as you can see, all the hyperparameters are equal for both type of models excepted for the hyperparameter λ_1.

The characteristics of our hyperparameters after the optimization procedure is as follow:

Table 2: List of hyperparameter values

Hyperparameters Values	Large-Scale Model	Small-Scale Model
Autoregressive Coefficient:	0	0
Overall Tightness (λ_1):	0.05	0.23
Lag Decay (λ_3):	2	2
Exogenous Variable Tightness	1	1
Lag Length	1	1

The hyperparameter λ_1 is equal to 0.05 for the large-scale model while the hyperparameter λ_1 for the small-scale is 0.23. From a practical point of view, this means that the "true" value of the coefficients estimated (posterior) are probably to be farther from the prior mean in the small-scale model than in the large-scale one.

Another aspect to consider about λ_1, is the fact that this hyperparameter impacts on the distribution of the parameters of lagged endogenous and exogenous variables of each equation in the system. In this sense, with more shrinkage for example, it is less probable that the posterior coefficients of the lagged endogenous and exogenous variables depart from the prior.

As can see in table 2, the posterior coefficients of the variables in the large-scale model are less probable to depart from the prior than the small-scale ones. Models with lots of variables will tend to have a better in-sample fit even when λ_1 is set to a "loose" value.

The posteriors obtained for the small- and the large-scale model of the inflation equation in each type of model are shown in the Appendix.

4.1.2 Forecasting Exercise

Our forecasting exercise is conducted in the following way. We estimate the hyperparameters considering the whole sample, through the maximization of the marginal likelihood; and then, we compute the forecasts.

As we mention before, the data set goes from January 2004 to July 2017. We compute one-, three- and six-step-ahead forecasts with rolling windows. The size of the estimation sample is the same for each forecast horizon. Out-of-sample forecast accuracy is measured in terms of RMSE of the forecasts. Therefore, we obtained three RMSEs for each model.

Relative forecast accuracy is analyzed in table 3, by computed the different combinations of RMSE ratios. On average, the BVAR presents better accuracy than the benchmark independently of the forecast horizon. For immediate horizons, the small-scale model slightly outperforms the larger one, but the large-scale model outperforms the small one for further forecast horizons.

Table 3: Relative forecast accuracy

One Step Ahead			Three Step Ahead			Six Step Ahead		
Ratio Small Model-Benchmark	Ratio Large Model-Benchmark	Ratio Large Model-Small Model	Ratio Small Model-Benchmark	Ratio Large Model-Benchmark	Ratio Large Model-Small Model	Ratio Small Model-Benchmark	Ratio Large Model-Benchmark	Ratio Large Model-Small Model
0.77	0.90	1.69	0.78	0.77	1.02	0.87	0.82	0.94

In the next subsection, we analyze these results with a Giacomini-White test.

4.2 Forecast Evaluation

To evaluate the predictive performance of the different models, we used the tests described earlier. Each column of the table 4 contains the probability value of Giacomini-White test statistic for the different models.

Table 4: Giacomini-White test

Forecast Horizon	Large BVAR vs. Benchmark	Small BVAR vs. Benchmark	Difference Between BVAR Models
One-Step-Ahead	0.03	0.01	0.29
Three-Steps-Ahead	0.00	0.00	0.49
Six-Steps-Ahead	0.09	0.05	0.41

The result of the Giacomini-White test show that at a 5% of significance level, the large BVAR model outperforms the benchmark for one step and three step ahead forecast horizon; while the small BVAR outperforms the benchmark at a 5% significance level for all forecast horizons. The last column of the table shows the Giacomini-White test applied to the differences in predictive ability between the small- and large-scale BVAR models, but in this case the differences are not significant for all forecast horizons.

5. <u>Conclusions</u>

This paper assesses the performance of Bayesian VAR to forecast inflation in Argentina. We considered a Normal Wishart BVAR specification for a small- and a large-scale model of differentiated variables setting the prior mean according to standard recommendations in previous studies. The overall tightness hyperprior and the lag length of the different models were set by optimization of the marginal likelihood. We found that large-scale models have narrower priors, giving more weight to the priors mean than small-scale models.

Overall, the results show that the BVAR model can improve the forecast ability of the univariate autoregressive benchmark's model of inflation. The Giacomini-White test indicates

that a BVAR performs better than the benchmark in all forecast horizons. Statistical differences between the two BVAR model specifications (small and large-scale) are not found. However, looking at the RMSEs, one can see that the larger model seems to perform better for larger forecast horizons.

Appendix

Data Characteristics

Table 1A: List of endogenous variables

	Source	Description	Transf.	Characteristics		
1	INDEC	EMAE	log	SA		Unit-Root
2	INDEC	CPI Inflation	-	-		Trend
3	INDEC	Core CPI Inflation (ex. seasonal and regulated)	-	-		Trend
4	INDEC	Industrial Employment	log	SA		Unit-Root
5	INDEC	Construction Employment	log	SA		Unit-Root
6	INDEC	Retail Trade Employment	log	SA		Stationary
7	BCRA	M2 Monetary Aggregate	log	SA		Unit-Root
8	BCRA	Multilateral Nominal Exchange Rate	log	-		Unit-Root
9	BCRA	30 to 59-day Deposit Rate	-	-		Unit-Root
10	INDEC	Imports of Intermediate Goods	log	SA		Unit-Root
11	INDEC	Total Exports	log	SA		Unit-Root
12	UTDT	Consumer Confidence Index	-	-		Unit-Root
13	INDEC	Monthly Supermarket Sales	log	SA		Unit-Root
14	AFCP	Cement Sales	log	SA		Unit-Root
15	MINEM	Asphalt Sales	log	-		Stationary
16	MERVAL	Stock Market Index	log	-		Unit-Root

Results characteristics

Table 2A: Small BVAR characteristics

Endogenous variables:	Inflation, Interest Rate, Real Activity
Exogenous variables:	Constant, Dummy 2016-11
Estimation sample:	July 2004 to July 2017
Sample size (omitting initial conditions):	156
Number of lags included in regression:	1
Prior:	Normal-Wishart
Autoregressive coefficient:	0
Overall tightness:	0.23
Lag decay:	2
Exogenous variable tightness:	1

Table 3A: Small BVAR inflation equation coefficient values

	Median	SD	LB	UB
INF(-1)	0.468	0.066	0.338	0.598
I(-1)	0.901	0.640	-0.356	2.157
Y(-1)	2.631	3.500	-4.237	9.499
Constant	0.280	0.071	0.140	0.420
d112016	-0.197	0.144	-0.479	0.086

Sum of squared residuals: 91.05
R-squared: 0.291
Adj. R-squared: 0.272

Table 4A: Large BVAR characteristics

Endogenous variables	Inflation, Interest Rate, Real Activity, Multilateral Exchange Rate, Industrial Employment, Cement Sales, Asphalts Sales, Imports of Intermediate Goods, Total Exports, M2, Core Inflation, Construction Employment, Consumer Confidence Index, Supermarket Sales, Stock Market Index
Exogenous variables:	Constant, Dummy 2016-11
Estimation sample:	July 2004 to July 2017
Sample size:	156
Number of lags:	1
Prior:	Normal-Wishart
Autoregressive coefficient:	0
Overall tightness:	0.05
Lag decay:	2
Exogenous variable tightness:	1

Table 5A: Large BVAR inflation equation coefficient values

	Median	SD	LB	UB
INF(-1)	0.145	0.045	0.057	0.234
I(-1)	0.436	0.407	-0.362	1.235
Y(-1)	1.177	2.131	-3.005	5.359
E(-1)	7.261	3.431	0.528	13.994
EMPI(-1)	16.644	11.611	-6.143	39.431
CEM(-1)	-0.680	0.556	-1.771	0.410
ASPH(-1)	0.083	0.411	-0.723	0.888
IMP(-1)	0.125	0.477	-0.810	1.061
EXP(-1)	0.091	0.491	-0.873	1.055
M2(-1)	4.093	2.410	-0.637	8.823
INFC(-1)	0.183	0.047	0.091	0.275

EMPC(-1)	-1.452	2.933	-7.207	4.303
ICC(-1)	-0.011	0.013	-0.036	0.013
SUP(-1)	2.243	1.322	-0.351	4.837
STK(-1)	0.133	1.110	-2.045	2.310
Constant	0.056	0.039	-0.021	0.132
d112016	-0.014	0.042	-0.096	0.067

Sum of squared residuals: 89.33
R-squared: 0.304
Adj. R-squared: 0.224

References

Andrew Blake and Haroon Mumtaz. (2012). "Center of Central Banking Studies Technical Handbook: Applied Bayesian econometrics for central bankers." Technical Handbook N°4.

Banbura M., Giannone D. and Reichlin L. (2008). "Large Bayesian VARs." Working Paper Series N° 996. European Central Bank.

Berger, J. (1985). "Statistical Decision Theory and Bayesian Analysis." Springer Series in Statistics.

Bloor, C. & Matheson, T. (2009). "Real-time conditional forecasts with Bayesian VARs: An application to New Zealand." Reserve Bank of New Zealand Discussion Paper Series DP2009/02, Reserve Bank of New Zealand.

Carriero A., Clark, T. and Marcellino M. (2015). "Bayesian VARs: Specification Choices and Forecast Accuracy." Journal of Applied Econometrics; Volume 30, Issue 1, Pages 46–73.

Clark, T. and McCracken, M. (2007). "Forecasting with small macroeconomic VARs in the presence of instabilities." Finance and Economics Discussion Series 2007-41, Board of Governors of the Federal Reserve System (U.S.).

Clements, Michael P. and Hendry David F. (1996). "Multi-Step Estimation for Forecasting." Oxford Bulleting of Economics and Statistics. Volume 58, Issue 4, Pages 657–684.

De Mol, C., D. Giannone, and L. Reichlin (2008): "Forecasting using a large number of predictors - Is Bayesian regression a valid alternative to principal components? "Journal of Econometrics, forthcoming.

Del Negro, Marco, Schorfheide, F., Smets, F. & Wouters, R. 2004. "On the fit and forecasting performance of New Keynesian models." FRB Atlanta Working Paper 2004-37, Federal Reserve Bank of Atlanta.

Diebold, F.X. and R.S. Mariano. (1995). "Comparing Predictive Accuracy." Journal of Business and Economic Statistics, 13: 253-63.

Doan, T., R. Litterman, and C. A. Sims (1984). "Forecasting and Conditional Projection Using Realistic Prior Distributions." Econometric Reviews, 3, 1–100.

Koop, Gary, 2011. "Forecasting with Medium and Large Bayesian VARs." SIRE Discussion Papers 2011-38, Scottish Institute for Research in Economics (SIRE).

Geweke J. (2005). "Contemporary Bayesian Econometrics and Statistics". Hoboken: Wiley.

Gelman, A. (2006). "Prior distributions for variance parameters in hierarchical models (Comment on Article by Browne and Draper)". Bayesian Analysis, 1, Number 3, pp. 515–534.

Giacomini, R. and H. White (2006). "Test of Conditional Predictive Ability." Econometrica 74, pp. 1545-1578.

Giannone, D., Lenza, M., and Primiceri, G. E. (2012). "Prior Selection for Vector Autoregressions." Working Papers ECARES ECARES 2012-002, ULB – Universite Libre de Bruxelles.

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.

Gneiting, T. and Raftery, A. (2007). "Strict Proper Scoring Rules, Prediction, and Estimation." Journal of the American Statistical Association, 102(477).

Gneiting, T. (2011). "Making and evaluating point forecasts". Journal of the American Statistical Association 106, 746–762.

Hashimzade, N. and Thornton M. A. (2013). "Handbooks of Research Methods and Applications in Empirical Macroeconomics." Ed. Edward Elgar

Koop, G. (2003). "Bayesian Econometrics." Ed. Wiley.

Litterman, R. B. (1979). "Techniques Of Forecasting Using Vector Autoregressions." Ph. D. thesis, University of Minnesota.

Litterman, R. B. (1980). "A Bayesian Procedure for Forecasting with Vector Autoregression." Working paper, Massachusetts Institute of Technology, Department of Economics.

Litterman, R. B. (1986). "Forecasting with bayesian vector autoregressions: Five years of experience." Journal of Business and Economic Statistics 4 (1), 25–38.

Mao, W. (2010). "Bayesian multivariate predictions." PhD Thesis, University of Iowa.

Marcellino, M. and Mizon, G. (2000). "Modelling shifts in the wage–price and unemployment–inflation relationships in Italy, Poland and the UK". Economic Modelling Vol. 17, Pages 387-413.

Robertson, J. C., and E. W. Tallman (1999). "Vector autoregressions: forecasting and reality." Economic Review, (Q1), 4–18.

Svensson, Lars E.O., 2005. "Monetary policy with judgment: forecast targeting." International Journal of Central Banking 1, 1–54.

Waggoner, D. F. and Zha, T. (1999). "Conditional Forecasts In Dynamic Multivariate Models." The Review of Economics and Statistics, 81(4):639–651.

Warne, A., Coenen, G., and Christoffel, K. (2013). "Marginalised Predictive Likelihood Comparisons with Application to DSGE, DSGE-VAR, and VAR Models." Technical report, Center for Financial Studies.

Wright, J. H. (2009). "Forecasting US inflation by Bayesian model averaging." Journal of Forecasting, 28(2), 131–144.

ie | BCRA
INVESTIGACIONES ECONÓMICAS