Forecasting using a large number of predictors: Is Bayesian regression a valid alternative to principal components?

Christine de Mol, Domenico Giannone, Lucrezia Reichlin

► To cite this version:

Christine de Mol, Domenico Giannone, Lucrezia Reichlin. Forecasting using a large number of predictors: Is Bayesian regression a valid alternative to principal components?. *Econometrics*, 2008, 146 (2), pp.318. 10.1016/j.jeconom.2008.08.011 . hal-00529920

HAL Id: hal-00529920

<https://hal.science/hal-00529920>

Submitted on 27 Oct 2010

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Accepted Manuscript

Forecasting using a large number of predictors: Is Bayesian regression a valid alternative to principal components?

Christine De Mol, Domenico Giannone, Lucrezia Reichlin

PII: S0304-4076(08)00110-3

DOI: [10.1016/j.jeconom.2008.08.011](https://doi.org/10.1016/j.jeconom.2008.08.011)

Reference: ECONOM 3066

To appear in: *Journal of Econometrics*

Please cite this article as: De Mol, C., Giannone, D., Reichlin, L., Forecasting using a large number of predictors: Is Bayesian regression a valid alternative to principal components?. *Journal of Econometrics* (2008), doi:10.1016/j.jeconom.2008.08.011

This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

Forecasting using a large number of predictors: is Bayesian regression a valid alternative to principal components?

Christine De Mol, Université Libre de Bruxelles, ECARES
Domenico Giannone*, European Central Bank, ECARES and CEPR
Lucrezia Reichlin, European Central Bank, ECARES and CEPR

March 5, 2008

Abstract

This paper considers Bayesian regression with normal and double-exponential priors as forecasting methods based on large panels of time series. We show that, empirically, these forecasts are highly correlated with principal component forecasts and that they perform equally well for a wide range of prior choices. Moreover, we study conditions for consistency of the forecast based on Bayesian regression as the cross-section and the sample size become large. This analysis serves as a guide to establish a criterion for setting parameters in a large cross-section.

JEL Classification: C11,C13,C33,C53

Keywords: Bayesian VAR, ridge regression, Lasso regression, principal components, large cross-sections

*Corresponding author. European Central Bank, Directorate General Research, Kaiserstrasse 29, D-60311 Frankfurt am Main, Germany. Tel.: +49 69 1344 7849, Fax: +49 69 1344 8553.

E-mail address: Domenico.Giannone@ecb.int

1 Introduction

Many problems in economics require the exploitation of large panels of time series. Recent literature has shown the “value” of large information for signal extraction and forecasting, and new methods have been proposed to handle the large-dimensionality problem (Forni, Hallin, Lippi, and Reichlin, 2005; Giannone, Reichlin, and Sala, 2004; Stock and Watson, 2002a,b).

A related literature has explored the performance of Bayesian model averaging for forecasting (Koop and Potter, 2003; Stock and Watson, 2006, 2005a; Wright, 2003) but, surprisingly, few papers explore the performance of Bayesian regression in forecasting with high-dimensional data. Exceptions are Stock and Watson (2005a) who consider normal Bayes estimators for orthonormal regressors and Giacomini and White (2006) who provide an empirical example in which a Bayesian regression with a large number of predictors is compared with principal component regression (PCR).

Bayesian methods are part of the traditional econometrician toolbox and offer a natural solution to overcome the curse of dimensionality problem by shrinking the parameters via the imposition of priors. In particular, the Bayesian VAR has been advocated as a device for forecasting macroeconomic data (Doan, Litterman, and Sims, 1984; Litterman, 1986). It is then surprising that, in most applications, these methods have been applied to relatively small systems and that their empirical and theoretical properties for large panels have not been given more attention in the literature.

This paper is a first step towards filling this gap. We analyze Bayesian regression methods under Gaussian and double-exponential prior and study their forecasting performance on the standard “large” macroeconomic dataset that has been used to establish properties of principal-component-based forecast (Stock and Watson, 2002a,b). Moreover, we analyze the asymptotic properties of Gaus-

sian Bayesian regression for n , the size of the cross-section, and T , the sample size, going to infinity. The aim is to establish a connection between Bayesian regression and the classical literature on forecasting with large panels based on principal components.

Our two choices for the prior correspond to two interesting cases: variable aggregation and variable selection. Under Gaussian prior, the posterior mode solution is such that all variables in the panel are given non-zero coefficients. Regressors, as in PCR, are linear combinations of all variables in the panel, but while the Gaussian prior gives decreasing weight to the ordered eigenvalues of the covariance matrix of the data, principal components imply unit weight to the dominant ones and zero to the others. The double-exponential prior, on the other hand, favors *sparse* models since it puts more mass near zero and in the tails which induces a tendency of the coefficients maximizing the posterior density to be either large or zero. As a result, it favors the recovery of few large coefficients instead of many small ones and truly zero rather than small values. This case is interesting because it results in variable selection rather than in variable aggregation and, in principle, this should give results that are more interpretable from the economic point of view.

Under a Gaussian prior, it is easy to compute the maximizer of the posterior density. Under such prior with independent and identically distributed (i.i.d.) regression coefficients, the solution amounts to solving a penalized least-squares problem with a penalty proportional to the sum of the squares of the coefficients, i.e. to a so-called Ridge regression problem. Under a double-exponential prior, however, there is no analytical form for the maximizer of the posterior density, but we can exploit the fact that, under such prior with i.i.d. coefficients, the solution amounts to a Lasso regression problem, i.e. to penalized least-squares with a penalty proportional to the sum of the absolute values of the coefficients.

Several algorithms have been proposed for Lasso regression. In our empirical study, we have used two algorithms recently proposed which work without limitations of dimensionality: LARS (Least Angle Regression) developed by Efron, Hastie, Johnstone, and Tibshirani (2004) and the Iterative Landweber scheme with soft-thresholding at each iteration developed by De Mol and Defrise (2002) and Daubechies, Defrise, and De Mol (2004).

An interesting feature of Lasso regression is that it combines variable selection and parameter estimation. The estimator depends in a nonlinear way on the variable to be predicted and this may have advantages in some empirical situations. The availability of the algorithms mentioned above, which are computationally feasible, makes the double-exponential prior an attractive alternative to other priors used for variable selection and requiring computationally demanding algorithms, such as the one proposed by Fernandez, Ley, and Steel (2001) in the context of Bayesian Model Averaging and applied by Stock and Watson (2005a) to macroeconomic forecasting with large cross-sections.

Although Gaussian and double-exponential Bayesian regressions rely on different estimation strategies, an out-of-sample evaluation based on the Stock and Watson dataset, shows that, for a given range of the prior choice, the two methods produce forecasts which are highly correlated and are characterized by similar mean-square errors. Moreover, these forecasts are highly correlated with those produced by principal components, also with similar mean-square errors: they do well when PCR does well. Hence, although the Lasso prior leads to the selection of few variables, the forecasts obtained from these informative targeted predictors do not outperform PCR based on few principal components.

In order to understand these results, we study the asymptotic properties of the forecast based on Bayesian regression as the cross-section and the sample size become large. This double-asymptotic analysis has been applied by recent

literature to the case of PCR (Bai, 2003; Bai and Ng, 2002; Forni, Giannone, Lippi, and Reichlin, 2007; Forni, Hallin, Lippi, and Reichlin, 2004; Stock and Watson, 2002a,b) but never to Bayesian regression. This analysis is however important to understand performance of this method for large panels and also as a guide to set shrinkage parameters as the dimension of the panel changes. Here we will limit the analysis to the Bayesian regression based on Gaussian prior and show that, under very general conditions, consistency is achieved provided that the degree of shrinkage increases with the cross-sectional dimension. The conditions under which we show consistency require that most of the regressors are informative about the future of the variable to forecast. This condition is satisfied in the particular case in which the data follow an approximate factor structure, case for which the literature has shown consistency for PCR. The approximate factor structure imposes a high degree of collinearity in the data that persists as we add series to the panel. Intuitively, under those assumptions, if the prior is chosen appropriately in relation with n , Bayesian regression under normality will give larger weight to the principal components associated with the dominant eigenvalues and therefore will produce results which are similar to PCR.

Our empirical work shows, moreover, that Lasso forecasts, although based on regression on few variables are as accurate and as highly correlated with PCR forecasts as are those obtained under normality. This result may seem puzzling, but it can be explained by the fact that our panel is highly collinear. Under collinearity, few variables, if selected appropriately, should capture the essence of covariation. In this case, we expect them to be strongly correlated with principal components and, as the latter, to span the space of the pervasive common factors. Under collinearity, however, we expect the selection not to be stable and to be very sensitive to minor perturbation of the data. In this

sense, we do not expect variable selection to provide results which lead to clearer economic interpretation than principal components or ridge regression.

The paper is organized as follows. The second Section introduces the problem of forecasting using large cross sections. The third Section reports the results of the out-of-sample exercise for the three methods considered: principal components, Bayesian regression with normal and with double-exponential prior. The fourth Section reports asymptotic results for the Gaussian prior case. The fifth Section concludes and outlines problems for future research.

2 Three solutions to the “curse of dimensionality” problem

Consider the $(n \times 1)$ vector of covariance-stationary processes $Z_t = (z_{1t}, \dots, z_{nt})'$. We will assume that they all have mean zero and unitary variance.

We are interested in forecasting linear transformations of some elements of Z_t using all the variables as predictors. Precisely, we are interested in estimating the linear projection

$$y_{t+h|t} = \text{proj} \{y_{t+h} | \Omega_t\}$$

where $\Omega_t = \text{span} \{Z_{t-s}, s = 0, 1, 2, \dots\}$ is a potentially large information set at time t and $y_{t+h} = z_{i,t+h}^h = f_h(L)z_{i,t+h}$ is a filtered version of z_{it} , for a specific $i = 1, \dots, n$.

Traditional time series methods approximate the projection using only a finite number, p , of lags of Z_t . In particular, they consider the following regression model:

$$y_{t+h} = Z_t' \beta_0 + \dots + Z_{t-p}' \beta_p + u_{t+h} = X_t' \beta + u_{t+h}$$

where $\beta = (\beta_0', \dots, \beta_p')'$ and $X_t = (Z_t', \dots, Z_{t-p}')'$. The implied forecast is given

by $y_{t+h|t} = X_t' \beta$ and the implied forecast error is $u_{t+h} = y_{t+h} - y_{t+h|t}$. The latter is assumed to be orthogonal to $z_{i,t-s}$ for $s = 0, 1, \dots, p$ and $i = 1, \dots, n$.

Given a sample of size T , we will denote by $X = (X_{p+1}, \dots, X_{T-h})'$ the $(T-h-p) \times n(p+1)$ matrix of observations for the predictors and by $y = (y_{p+1+h}, \dots, y_T)'$ the $(T-h-p) \times 1$ matrix of the observations on the dependent variable. The regression coefficients are typically estimated by Ordinary Least Squares (OLS), $\hat{\beta}^{LS} = (X'X)^{-1}X'y$, and the forecast is given by $\hat{y}_{T+h|T}^{LS} = X_T' \hat{\beta}^{LS}$. When the size of the information set, n , is large, such projection involves the estimation of a large number of parameters. This implies a loss of degrees of freedom and a poor forecast (“curse of dimensionality problem”). Moreover, if the number of regressors is larger than the sample size, $n(p+1) > T$, OLS is not feasible.

To solve this problem, the literature proposes to compute the forecast as a projection on the first few principal components (Forni, Hallin, Lippi, and Reichlin, 2005; Giannone, Reichlin, and Sala, 2004; Giannone, Reichlin, and Small, 2005; Stock and Watson, 2002a,b).

Consider the spectral decomposition of the sample covariance matrix of the regressors:

$$S_x V = V D \quad (1)$$

where $D = \text{diag}(d_1, \dots, d_{n(p+1)})$ is a diagonal matrix having on the diagonal the eigenvalues of $S_x = \frac{1}{T-h-p} X'X$ in decreasing order of magnitude and $V = (v_1, \dots, v_{n(p+1)})$ is the $n(p+1) \times n(p+1)$ matrix whose columns are the corresponding normalized eigenvectors¹. The normalized principal components

¹The eigenvalues and eigenvectors are typically computed on $\frac{1}{T-p} \sum_{t=p+1}^T X_t X_t'$ (see for example Stock and Watson, 2002a). We instead compute them on $\frac{1}{T-h-p} X'X = \frac{1}{T-h-p} \sum_{t=p+1}^{T-h} X_t X_t'$ for comparability with the other estimators considered in the paper.

(PC) are defined as:

$$\hat{f}_{it} = \frac{1}{\sqrt{d_i}} v_i' X_t \quad (2)$$

for $i = 1, \dots, N$ where N is the number of non-zero eigenvalues².

If most of the interactions among the variables in the information set is due to few common underlying factors, while there is limited cross-correlation among the variable-specific components of the series, the information content of the large number of predictors can indeed be summarized by few aggregates, while the part not explained by the common factors can be predicted by means of traditional univariate (or low-dimensional forecasting) methods and hence captured by projecting on the dependent variable itself (or on a small set of predictors). In such situations, few principal components provide a good approximation of the underlying factors. The principal component forecast is defined as:

$$y_{t+h|t}^{PC} = \text{proj} \left\{ y_{t+h} | \Omega_t^f \right\} \approx \text{proj} \left\{ y_{t+h} | \Omega_t \right\} \quad (3)$$

where $\Omega_t^f = \text{span} \left\{ \hat{f}_{1t}, \dots, \hat{f}_{rt} \right\}$, with $r \ll n(p+1)$, is a parsimonious representation of the information set. The parsimonious approximation of the information set makes the projection feasible, since it requires the estimation of a limited number of parameters.

The literature has studied rates of convergence of the principal component forecast to the efficient forecast under assumptions defining an approximate factor structure (see Section 4). Under those assumptions, once common factors are estimated via principal components, the projection is computed by OLS treating the estimated factors as if they were observables.

The Bayesian approach we follow consists instead in imposing limits on the

²Note that $N \leq \min\{n(p+1), T-h-p\}$.

length of β through priors and estimating the parameters as the posterior mode. The parameters are then used to compute the forecasts. Here we consider two alternatives: Gaussian and double-exponential prior.

Let us assume that $u_t \sim \text{i.i.d. } \mathcal{N}(0, \sigma_u^2)$, with known variance σ_u^2 ; then, under Gaussian prior $\beta \sim \mathcal{N}(\beta_0, \Phi_0)$, and assuming for simplicity that all parameters are shrunk to zero, i.e. $\beta_0 = 0$, we have:

$$\hat{\beta}^{bay} = (X'X + \sigma_u^2 \Phi_0^{-1})^{-1} X'y.$$

The corresponding forecast is then computed as:

$$\hat{y}_{T+h|T}^{bay} = X_T' \hat{\beta}^{bay}.$$

In the case in which the parameters are i.i.d.³, i.e. $\Phi_0 = \sigma_\beta^2 I$, the estimates are equivalent to those produced by penalized Ridge regression with parameter $\nu = \frac{\sigma_u^2}{\sigma_\beta^2}$. Precisely⁴:

$$\hat{\beta}^{bay} = \arg \min_{\beta} \{ \|y - X\beta\|^2 + \nu \|\beta\|^2 \}.$$

It is known that there exist close relationships between OLS, PCR, penalized and Bayesian regression – see e.g. the book by Hastie, Tibshirani, and Friedman (2001) for a more detailed discussion of the connections between the different methods. For example, if the prior belief on the regression coefficients is that they are i.i.d., the forecast can be represented as a weighted sum of the projections on the principal components:

³Homogenous variance and mean zero are very naive assumptions. In our case, they are justified by the fact that the variables in the panel we will consider for estimation are standardized and demeaned. This transformation is appropriate for comparison with principal components.

⁴In what follows we will denote by $\|\cdot\|$ the L^2 matrix norm, i.e. for every matrix A , $\|A\| = \sqrt{\lambda_{\max}(A'A)}$ where $\lambda_{\max}(A'A)$ is the maximum eigenvalue of $A'A$. For vectors $\|\cdot\|$ denotes the Euclidean norm.

$$X_T' \hat{\beta} = \sum_{i=1}^N w_i \hat{f}_{iT} \hat{\alpha}_i \quad (4)$$

where $\hat{\alpha}_i = \frac{1}{\sqrt{d_i}} v_i' X' y / (T - h - p)$ is the OLS regression coefficient of y on the i th principal component. For OLS we have $w_i = 1$ for all i . For the Bayesian estimates $w_i = \frac{d_i}{d_i + \frac{\nu}{T-h-p}}$, where $\nu = \frac{\sigma_u^2}{\sigma_\beta^2}$. For the PCR regression we have $w_i = 1$ for $i \leq r$, and zero otherwise.

OLS, PCR and Gaussian Bayesian regression give non-zero weight to all variables. An alternative is to select variables. For Bayesian regression, variable selection can be achieved by a double-exponential prior, which, in the case of a zero-mean i.i.d. prior, is equivalent to the method that is sometimes called Lasso regression (an acronym for “least absolute shrinkage and selection operator”)⁵. In this particular i.i.d. prior case the method can also be seen as a penalized regression with a penalty on the coefficients involving the L_1 norm instead of the L_2 norm. Precisely:

$$\hat{\beta}^{lasso} = \arg \min_{\beta} \left\{ \|y - X\beta\|^2 + \nu \sum_{i=1}^n |\beta_i| \right\} \quad (5)$$

where $\nu = 1/\tau$, τ being the scale parameter of the prior density⁶ (see e.g. Tibshirani, 1996; Fu, 1998).

Compared with the Gaussian density, the double-exponential puts more mass near zero and in the tails and this induces a tendency to produce estimates of the regression coefficients that are either large or zero. As a result, one favors the recovery of a few large coefficients instead of many fairly small ones. Moreover, as we shall see, the double-exponential prior favors truly zero values instead of small ones, i.e. it favors *sparse* regression coefficients (sparse mode).

To gain intuition about Lasso regression, let us consider, as an example, the

⁵It should be noted however that Lasso is actually the name of an algorithm proposed in Tibshirani (1996) for finding a minimizer of (5).

⁶We recall here that the variance of the prior density is proportional to $2\tau^2$.

case of orthogonal regressors, a case for which the posterior mode has a known analytical form. In particular, let us consider the case in which the regressors are the principal components of X . In this case, the Lasso solution can be cast in the form (4) with $w_i \hat{\alpha}_i$ replaced by $S_\nu(\hat{\alpha}_i)$, where S_ν is the *soft-thresholder* defined by

$$S_\nu(\alpha) = \begin{cases} \alpha + \nu/2 & \text{if } \alpha \leq -\nu/2 \\ 0 & \text{if } |\alpha| < \nu/2 \\ \alpha - \nu/2 & \text{if } \alpha \geq \nu/2. \end{cases} \quad (6)$$

Hence, this sparse solution is obtained by setting to zero all coefficients $\hat{\alpha}_i$ which in absolute value lie below the threshold $\nu/2$ and by shrinking the largest ones by an amount equal to the threshold. Let us remark that it would also be possible to leave the largest components untouched, as done in so-called *hard-thresholding*, but we do not consider this variant here since the lack of continuity of the hard-thresholding function makes the theoretical framework more complicated.

In the general case, i.e. with non-orthogonal regressors, the Lasso solution will enforce sparsity on the variables themselves rather than on the principal components, and this is an interesting feature of the method since it implies a regression on just a few observables rather than on a few linear combinations of the observables. Note that with such non-Gaussian priors the model is not invariant under orthogonal linear transformation of the data.

Notice also that, unlike in Ridge and PC regressions, where the regressors are weighted independently of the choice of the series to be forecasted, in the Lasso regression the selection and shrinkage depend on that choice.

Methods described by equation (4) will perform well provided that no truly relevant coefficients α_i are observed for $i > r$, because in principal component regression they will not be taken into account and in Ridge their influence will

be highly weakened. Bad performances are to be expected if, for example, we aim at forecasting a time series y_t , which by bad luck is just equal or close to a principal component $\hat{f}_i$ with $i > r$. Lasso solves this problem.

Unfortunately, in the general case, the mode of the posterior distribution has no analytical form and has to be computed using numerical methods such as the Lasso algorithm of Tibshirani (1996) or quadratic programming based on interior point methods as advocated in Chen, Donoho, and Saunders (2001). Two efficient alternatives to the Lasso algorithm, which work without limitations of dimensionality also for sample size T smaller than the number of regressors $n(p + 1)$, have been developed more recently by Efron, Hastie, Johnstone, and Tibshirani (2004) under the name LARS (Least Angle Regression)⁷ and by De Mol and Defrise (2002); Daubechies, Defrise, and De Mol (2004) who use instead an Iterative Landweber scheme with soft-thresholding applied at each iteration step⁸.

In the next Section we study the empirical performance of the three methods discussed in an out-of-sample forecast exercise based on a large panel of time series.

3 Empirics

The dataset employed for the out-of-sample forecasting analysis is the same as the one used in Stock and Watson (2005b). The panel includes real variables (sectoral industrial production, employment and hours worked), nominal variables (consumer and producer price indices, wages, money aggregates), asset prices (stock prices and exchange rates), the yield curve and surveys, for a total

⁷The LARS algorithm has also been used in econometric forecasting by Bai and Ng (2006b) who also use it for selecting variables to form principal components.

⁸The latter algorithm carries out most of the intuition of the orthogonal regression case and is described in Appendix B. For the LARS algorithm we refer to Efron, Hastie, Johnstone, and Tibshirani (2004).

of $n = 131$ variables⁹.

Series are transformed to obtain stationarity. In general, for real variables, such as employment, industrial production, sales, we take the monthly growth rate. We take first differences for series already expressed in rates: unemployment rate, capacity utilization, interest rate and some surveys. Prices and wages are transformed to first differences of annual inflation following Giannone, Reichlin, and Sala (2004); Giannone, Reichlin, and Small (2005).

Let us define IP as the monthly industrial production index and CPI as the consumer price index. The variables we forecast are

$$z_{IP,t+h}^h = (ip_{t+h} - ip_t) = z_{IP,t+h} + \dots + z_{IP,t+1}$$

and

$$z_{CPI,t+h}^h = (\pi_{t+h} - \pi_t) = z_{CPI,t+h} + \dots + z_{CPI,t+1}$$

where $ip_t = 100 \times \log IP_t$ is the (rescaled) log of IP and $\pi_t = 100 \times \log \frac{CPI_t}{CPI_{t-12}}$ is the annual CPI inflation (IP enters in the pre-transformed panel in first differences of the logarithm, while annual inflation enters in first differences).

The forecasts for the (log) IP and the level of inflation are recovered as:

$$\hat{ip}_{T+h|T} = \hat{z}_{IP,T+h|T}^h + ip_T$$

and

$$\hat{\pi}_{T+h|T} = \hat{z}_{CPI,T+h|T}^h + \pi_T .$$

The accuracy of predictions is evaluated using the mean-square forecast error

⁹A full description of our dataset is given in a separate appendix containing supplementary material about this paper and available on request or from the website <http://homepages.ulb.ac.be/~dgiannon/>.

(*MSFE*) metric, given by:

$$MSFE_{\pi}^h = \frac{1}{T_1 - T_0 - h + 1} \sum_{T=T_0}^{T_1-h} (\hat{\pi}_{T+h|T} - \pi_{T+h})^2$$

and

$$MSFE_{ip}^h = \frac{1}{T_1 - T_0 - h + 1} \sum_{T=T_0}^{T_1-h} (\hat{ip}_{T+h|T} - ip_{T+h})^2.$$

The sample has a monthly frequency and ranges from 1959:01 to 2003:12. The evaluation period is 1970:01 to 2002:12. $T_1=2003:12$ is the last available point in time, $T_0=1969:12$ and $h=12$. We consider rolling estimates with a window of 10 years, i.e. parameters are estimated at each time T using the most recent 10 years of data. For all methods we report results for $p=0$ (no lags of the regressor) which is the one typically considered in macroeconomic applications. Qualitative results are not affected by this choice¹⁰.

All the procedures are applied to standardized data. Mean and variance are re-attributed to the forecasts accordingly.

We report results for industrial production (IP) and the consumer price index (CPI).

Let us start from principal component regression. We report results for the choice of $r=1, 3, 5, 10, 25, 50, 75$ principal components. The case $r=0$ is the forecast implied from a random walk with drift on the log of IP and the annual CPI inflation.

We report MSFE relative to the random walk, and the variance of the forecasts relative to the variance of the series of interest¹¹. The MSFE is also reported for two sub-samples: the first half of the evaluation period 1970-1985,

¹⁰Empirical results supporting this claim are reported in a separate appendix containing supplementary material about this paper and available on request or from the website <http://homepages.ulb.ac.be/~dgiannon/>.

¹¹We limit the empirical evaluation to point forecasts as it is standard in the literature on principal components forecasts. The theoretical results derived in the next Section are also limited to point forecasts.

and the second half 1985-2002. These results help us understand the relative performance of the methods for the cases where the predictability of key macroeconomic time series has dramatically decreased (on this point, see D'Agostino, Giannone, and Surico, 2006). Results are reported in Table 1.

INSERT TABLE 1 OVER HERE

Let us start with the entire evaluation sample. Results show that principal components improve a lot over the random walk both for IP and CPI. The advantage is lost when taking too many PC, which implies loss of parsimony. Notice that, as the number of PC increases, the variance of the forecasts increases and can become even larger than the variance of the series itself. This is explained by the large sample uncertainty of the regression coefficients when there is a large number of regressors. Looking at the two sub-samples, we see that PCs perform very well in the first part of the sample, while in the most recent period they perform very poorly, worse than the random walk.

The empirical literature on principal component regression has also considered the inclusion of the past of the variable of interest to capture series specific dynamics. The inclusion of those additional regressors does not affect qualitative results and in particular does not significantly improve the accuracy of the forecasts¹².

Let us now do a similar exercise for the i.i.d. Gaussian prior (Ridge regression). Note, that, for $h = 1$, this case corresponds to a row of a VAR of order one. The Gaussian prior works well for the case $p = 0$ considered here¹³.

¹²Results are available on request. Similar results have also been reported in D'Agostino and Giannone (2007).

¹³Incidentally, for the case $p > 0$, let us observe that it might be useful to shrink more the coefficients of additional lagged regressors, as, for example, with the Minnesota prior (Doan, Litterman, and Sims, 1984; Litterman, 1986). An additional feature of the Litterman priors is to shrink less the coefficients associated with the variable to forecast. This can be helpful when series specific dynamics have significant forecasting power. The study of such more refined priors goes beyond the scope of the present empirical analysis which is meant as a first assessment of the general performance of the methods.

For the Bayesian Gaussian (Ridge) case, we run the regression using the first estimation sample 1959-1969 for a grid of priors. We then choose the priors for which the in-sample fit explains a given fraction $1 - \kappa$ of the variance of the variable to be forecast. We report results for different values of κ (the associated values of ν , which are kept fixed for the whole out-of-sample evaluation period, are also reported). Notice that $\kappa = 1$ corresponds to the random walk since, in this case, all coefficients are set to zero. The other extreme, κ close to 0, is associated with a quite uninformative prior and hence will be very close to the OLS. Results are reported in Table 2.

INSERT TABLE 2 OVER HERE

The Ridge forecast performs well for a range of κ between 30% and 70% that are associated with shrinkage parameters between half and ten times the cross-sectional dimension n . For the whole sample, the MSFE are close to those obtained with principal component regression. Moreover, the forecasts produced by Ridge regressions are smoother than the PC forecasts, which is a desirable property.

The last line of the table shows the correlation among Ridge forecasts and principal component forecasts¹⁴. Principal components and Ridge forecasts are highly correlated, particularly when the prior is such that the forecasting performances are good. The fact that correlation is maximal for parameters giving the best forecasts suggests that there is a common explanation for the good performance of the two methods.

As for the two sub-samples, results are also qualitatively similar to principal component forecasts. Ridge performs particularly well in the first sub-sample but loses all the advantage in the second. We can note, however, more stability

¹⁴For the principal component forecasts we use $r = 10$. We obtain similar results also for $r = 3, 5$, i.e. when PC forecasts perform well.

than in the principal components case. This is not surprising since Ridge uses all eigenvalues in decreasing importance instead of truncating after r as in the principal components case. Notice also that, for inflation, with ν in the intermediate range, even in the most recent sample there is a slight improvement over the random walk.

Finally, we analyze the case of double-exponential priors. In this case, instead of fixing the values of the parameter ν , we select the prior that delivers a given number (k) of non zero coefficients at each estimation step in the out-of-sample evaluation period. We look at the cases of $k = 1, 3, 5, 10, 25, 50, 75$ non-zero coefficients¹⁵.

Results, reported in Table 3, show that good forecasts are obtained with a limited number of predictors, between 5 and 25. As for Ridge, maximal correlation with the principal component forecast is achieved for the selection of parameters that gives the best results.

INSERT TABLE 3 OVER HERE

Comparable MSFE for the three methods as well as high correlation of the forecasts suggest that all three methods are capturing similar features of the data. In particular, the correlation of the two Bayesian forecasts with the principal component forecast, for the priors that ensure good performance, implies that there must be a common explanation for the success of the three methods.

The similarity between forecasts based on PC and Ridge can be explained by collinearity among predictors. In fact, since the covariance of our data is characterized by few dominant eigenvalues, PC and Ridge, by keeping the largest ones and giving, respectively zero weight and small weight to the others, should

¹⁵An alternative, closer in spirit to the exercise with Gaussian prior, is to select the prior ν at the beginning of the evaluation and then keep it fixed over the evaluation sample. This alternative strategy provides qualitatively similar results. See the appendix with supplementary material.

perform similarly. This point will emerge more clearly in Section 4 on the basis of the asymptotic analysis.

The result for Lasso is less straightforward to interpret since it is a regression on few variables rather than on few aggregates of the variables. The high correlation of the Lasso forecast with the PC forecast suggests that our data are highly collinear. Under collinearity, few variables, if appropriately selected, should capture the essence of the covariation of the data and, as principal components, span approximately the space of the pervasive common factors. However, under these circumstances, we should also expect the selection to be unstable and very sensitive to minor perturbations of the data. With collinear data structure, variable selection methods are unlikely to provide results that are more interpretable than principal components or ridge regressions from the economic point of view.

We examined the variables selected for $k \approx 10$ at the beginning and at the end of the out-of-sample evaluation period¹⁶. Two main results emerge from this analysis. First, only some of the selected variables coincide with those typically included in small-medium size models: the commodity price indexes, the spreads, money aggregates and stock market variables. Some of the selected variables are sectoral (production, labor market and price indicators) or regional (housing). Second, the selection is different at different points in the sample, although selected variables generally belong to the same economic category.

We have two conjectures about these results. The fact that variables are not clearly interpretable and that the procedure selects different variables at different points of the sample is, as mentioned above, the consequence of collinearity. The latter result also suggests temporal instability. Notice, however, that temporal instability does not affect the relative performance of principal components

¹⁶These variables are reported in the last two columns of the table describing the database contained in the appendix containing supplementary material about this paper and available on request or from the website <http://homepages.ulb.ac.be/~dgiannon/>.

and Ridge with respect to Lasso. This suggests that principal components and Ridge, by aggregating all variables in the panel, stabilize results providing a sort of insurance against temporal instability. These conjectures will be explored in further work.

4 Theory

We have seen that Bayesian regression and PCR are methods that help us solve the curse of dimensionality problem which typically arises when trying to extract relevant information from a large number of predictors.

For PCR, the literature has analyzed the asymptotic properties for the size of the cross-section n and the sample size T going to infinity under assumptions that essentially impose that, as we increase the number of time series, the sources of common dynamics remain limited (Bai, 2003; Bai and Ng, 2002; Forni, Giannone, Lippi, and Reichlin, 2007; Forni, Hallin, Lippi, and Reichlin, 2005; Stock and Watson, 2002a,b). Double asymptotics for Bayesian regression, on the other hand, has never been studied and this is a relevant analysis for understanding its behaviour when using a large number of predictors. In what follows, we will consider double (n, T) asymptotics for the case of the Gaussian prior, under conditions that are more general than those considered for PCR in the literature mentioned above. As we will see, our assumptions impose that the optimal forecast and the observable predictors depend on a number of finite unobservable factors.

For the sake of simplicity, we will assume throughout this Section that no lags of the regressors are used in the forecasting regression. In the notation of Section 2, this means that we set $p = 0$. In this case, X_t coincides with the predictors at time t , $Z_t = (z_{1,t}, \dots, z_{n,t})'$. All results, however, apply straightforwardly to the case in which lagged predictors are also included, i.e. p different from zero.

Let us first assume that the forecast of y_t depends on a finite number of unobserved factors.

Assumption A

$y_{t+h} = \gamma F_t + v_{t+h}$, where v_{t+h} is orthogonal to X_t for all n and where the factors $F_t = (f_{1t}, \dots, f_{rt})'$ are a r -dimensional stationary process with covariance matrix $E F_t F_t' = I_r$.

Consider the forecast based on the projection on the unobserved factors F_t :

$$y_{t+h|t}^* = \gamma F_t.$$

Under Assumption A, the forecast $y_{t+h|t}^*$ is optimal in the sense that, due to the assumption of orthogonality between the residuals v_{t+h} and the observed predictors X_t , its forecast accuracy cannot be improved using the information available at time t . For fixed n , the optimal forecast is unfeasible, even with infinite sample size T , since the factors are unobserved. We assume that the observed predictors are related to the common factors as follows:

Assumption B

X_t has the following representation:

$$X_t = \Lambda F_t + \xi_t$$

where

- (i) the residuals ξ_t are a n -dimensional stationary process with covariance matrix $E \xi_t \xi_t' = \Psi$ of full rank for all n ;

- (ii) the matrix Λ loading the factors is a non-random matrix of dimension $n \times r$ and of full-rank r for each n ;
- (iii) the residuals ξ_t are orthogonal to the factors F_t .

In Assumption B the predictors are decomposed into two parts. One part (ΛF_t) is driven by the factors which are informative about the future of the target variable. The residuals (ξ_t) can be considered as the component of the predictors that is not informative. For convenience we assume that the two components are orthogonal. The assumption that the non-informative residuals are of full rank entails that there are no redundant predictors. This ensures that when we increase the number of predictors we do not duplicate information.

Under Assumptions A and B, we have $\Sigma_x = E(X_t X_t') = \Lambda \Lambda' + \Psi$ and $\Sigma_{xy} = E(X_t y_{t+h}) = \Lambda \gamma'$. Because of Assumption B (i), Σ_x is invertible for all n . Consequently, for given number n of predictors, the population OLS regression coefficient $\beta = \Sigma_x^{-1} \Sigma_{xy}$ is unique and the forecast is given by:

$$y_{t+h|t} = X_t' \beta = X_t' (\Lambda \Lambda' + \Psi)^{-1} \Lambda \gamma'.$$

Let us first derive conditions on the shrinkage parameter that will allow to obtain consistent forecasts from Bayesian regression under Gaussian priors. We will need the additional Assumption C that ensures that the elements of the sample covariances of X_t with itself and with y_t converge uniformly to their population counterpart; see Appendix A for details.

Let us consider the prediction based on the Gaussian prior, $u_t \sim \text{i.i.d. } \mathcal{N}(0, \sigma_u^2)$ and $\beta \sim \mathcal{N}(0, \Phi_0)$. We have the following result:

Proposition 1 Under Assumptions A, B and C, and if $\liminf_{n, T \rightarrow \infty} \frac{\lambda_{\min}(\Phi_0)}{\|\Phi_0\|} >$

0, we have for $n, T \rightarrow \infty$:

$$X_t' \hat{\beta}^{bay} - y_{t+h|t} = O_p \left(\frac{\sqrt{\delta_{1,n}}}{\delta_{2,n}^2} \frac{\sqrt{n}}{T \|\Phi_0\|} \right) + O_p \left(\left[1 + \frac{\sqrt{\delta_{1,n}}}{\delta_{2,n}} \sqrt{n} \right] n \sqrt{T} \|\Phi_0\| \right)$$

where $\delta_{1,n} = \lambda_{max}(\Lambda' \Lambda)$ and $\delta_{2,n} = \lambda_{min}(\Lambda' \Lambda)$.

Proof. See Appendix A.

Proposition 1 indicates that the behavior of the Bayesian forecast under Gaussian prior is governed by the quantities $\delta_{1,n}$ and $\delta_{2,n}$, which in turn are related to the information content of the observable predictors X_t with respect to the factors F_t . If the factors are pervasive throughout the predictors' cross-section with non decreasing weights, then $\delta_{1,n}$ and $\delta_{2,n}$ go to infinity with n . We assume that they increase linearly with n :

Assumption D

$$0 < \liminf_{n \rightarrow \infty} \frac{1}{n} \lambda_{min}(\Lambda' \Lambda) \leq \limsup_{n \rightarrow \infty} \frac{1}{n} \lambda_{max}(\Lambda' \Lambda) < \infty .$$

Under Assumption D all the predictors are informative for the factors F_t and hence they all help improve the forecast accuracy. In this case the sample forecast converges to its population counterpart. Precisely:

Corollary 1 Under the assumptions of Proposition 1, and if Assumption D holds, then:

$$X_t' \hat{\beta}^{bay} - y_{t+h|t} = O_p \left(\frac{1}{nT \|\Phi_0\|} \right) + O_p \left(n \sqrt{T} \|\Phi_0\| \right) \text{ as } n, T \rightarrow \infty.$$

To achieve consistency, a suitable choice for the prior is $\|\Phi_0\|^{-1} = cnT^{\frac{1}{2}+\delta}$, with $0 < \delta < 1/2$ and c an arbitrary constant. Under this condition on the prior we have

$$\Delta_{nT} \left(X_t' \hat{\beta}^{bay} - y_{t+h|t} \right) = O_p(1) \text{ as } n, T \rightarrow \infty,$$

where $\Delta_{nT} = \min \left\{ T^\delta, T^{\frac{1}{2}-\delta} \right\}$ and $0 < \delta < 1/2$. Let us stress here that no restriction on the relative path of divergence of T and n is needed in order to achieve consistency. In this sense the estimates are viable also when the size of the cross-section n is much larger than the sample size T .

Corollary 1 tells us that, under the factor structure assumption, the Bayesian regression should use a prior that shrinks increasingly all regression coefficients to zero as the number of predictors rises. This is because, if the factors are pervasive, then all variables are informative for the common factors and we should give weight to all of them. Consequently, as the number of predictors increases, the magnitude of each regression coefficient has to decrease. The condition $\liminf_{n, T \rightarrow \infty} \frac{\lambda_{\min}(\Phi_0)}{\|\Phi_0\|} > 0$ requires that all the regression coefficients should be shrunk at the same asymptotic rate.

In the empirical exercise, the condition $\liminf_{n, T \rightarrow \infty} \frac{\lambda_{\min}(\Phi_0)}{\|\Phi_0\|} > 0$ is satisfied since we used the i.i.d. prior ($\Phi_0 = \sigma_\beta^2 I$). Moreover, from Corollary 1, consistency requires that the shrinkage parameter $\nu = \frac{\sigma_\eta^2}{\sigma_\beta^2}$ grows asymptotically at a rate equal to the number of predictors n . Although this is an asymptotic condition that is difficult to assess empirically on the basis of a finite cross-section and sample size, the empirical results appear to roughly confirm that, in order to get accurate forecasts, the parameter ν has to be in line with the above condition. Precisely, we found that the forecasts are accurate when the shrinkage parameter ranges between half and ten times the cross-sectional dimension n (see Table 2).

Let us now turn to the next question: under which conditions do we actually

obtain the optimal forecast? We will show that we can get it under further assumptions on the residuals ξ_t , i.e. the part of the predictors which is non-informative about the forecast. Loosely speaking, these assumptions will ensure that ξ_t is dominated in size by the common component driven by the factors. This will be the case when the errors are weakly correlated across predictors.

First notice that the population regression coefficient vector is given by

$$\beta = \Sigma_x^{-1} \Sigma_{xy} = (\Lambda \Lambda' + \Psi)^{-1} \Lambda \gamma' = \Psi^{-1} \Lambda (\Lambda' \Psi^{-1} \Lambda + I)^{-1} \gamma'.$$

Consider now

$$y_{t+h|t} = X_t' \beta = F_t' \Lambda' \Psi^{-1} \Lambda (\Lambda' \Psi^{-1} \Lambda + I)^{-1} \gamma' + \xi_t' \Psi^{-1} \Lambda (\Lambda' \Psi^{-1} \Lambda + I)^{-1} \gamma';$$

we have

$$\|\Lambda' \Psi^{-1} \Lambda (\Lambda' \Psi^{-1} \Lambda + I)^{-1} - I\| = \|(\Lambda' \Psi^{-1} \Lambda + I)^{-1}\| \leq \frac{1}{\lambda_{\min}(\Lambda' \Psi^{-1} \Lambda)} ;$$

moreover,

$$\begin{aligned} \mathbb{E} \left[(\xi_t' \beta)^2 \right] &= \beta' \Psi \beta = \gamma' (\Lambda' \Psi^{-1} \Lambda + I)^{-1} \Lambda' \Psi^{-1} \Psi \Psi^{-1} \Lambda (\Lambda' \Psi^{-1} \Lambda + I)^{-1} \gamma' \\ &\leq \|\gamma\|^2 \|(\Lambda' \Psi^{-1} \Lambda)^{-1}\| \|(\Lambda' \Psi^{-1} \Lambda) (\Lambda' \Psi^{-1} \Lambda + I)^{-1}\| \\ &\leq \|\gamma\|^2 \frac{1}{\lambda_{\min}(\Lambda' \Psi^{-1} \Lambda)} . \end{aligned}$$

Since $\|\gamma\| = O(1)$, by the Markov inequality we have: $\xi_t' \beta = O_p \left(\lambda_{\min}^{-1/2}(\Lambda' \Psi^{-1} \Lambda) \right)$.

Summarizing the results just derived above, we have

$$y_{t+h|t} - \gamma F_t = O_p \left(\lambda_{\min}^{-1/2}(\Lambda' \Psi^{-1} \Lambda) \right) \text{ as } n \rightarrow \infty$$

provided $\lambda_{\min}(\Lambda' \Psi^{-1} \Lambda) \geq 1$. We see that the key quantity involved in the

convergence to the optimal forecast is the minimum eigenvalue $\lambda_{\min}(\Lambda'\Psi^{-1}\Lambda)$. This is a measure of the importance of the informative component relative to the non-informative component of the predictors. If the informative component dominates, then the projection over the observables X_t converges to the optimal unfeasible forecast as the number of predictors goes to infinity.

Since the following bound holds

$$\lambda_{\min}(\Lambda'\Psi^{-1}\Lambda) \geq \lambda_{\min}(\Psi^{-1})\lambda_{\min}(\Lambda'\Lambda) = \lambda_{\max}^{-1}(\Psi)\lambda_{\min}(\Lambda'\Lambda)$$

we have the following result:

Proposition 2 Under Assumptions A, B and C, and if

$$\liminf_{n,T \rightarrow \infty} \frac{\lambda_{\min}(\Phi_0)}{\|\Phi_0\|} > 0 ,$$

we have for $n, T \rightarrow \infty$:

$$X_t' \hat{\beta}^{bay} - y_{t+h|t}^* = O_p \left(\sqrt{\frac{\delta_{3,n}}{\delta_{2,n}}} \right) + O_p \left(\frac{\sqrt{\delta_{1,n}}}{\delta_{2,n}^2} \frac{\sqrt{n}}{T \|\Phi_0\|} \right) \quad (7)$$

$$+ O_p \left(\left[1 + \frac{\sqrt{\delta_{1,n}}}{\delta_{2,n}} \sqrt{n} \right] n \sqrt{T} \|\Phi_0\| \right) \quad (8)$$

where $\delta_{1,n} = \lambda_{\max}(\Lambda'\Lambda)$, $\delta_{2,n} = \lambda_{\min}(\Lambda'\Lambda)$ and $\delta_{3,n} = \lambda_{\max}(\Psi)$ (provided $\delta_{2,n} \geq \delta_{3,n}$).

From Proposition 2 we can see that convergence to the optimal forecast, as the number of predictors n tends to infinity, is achieved when the maximum eigenvalue of Ψ is bounded or tends to infinity at a slower rate than the minimum eigenvalue of $\Lambda'\Lambda$, the latter rate being $O(n)$ when Assumption D holds true. Let us now introduce an assumption on the asymptotic growth rate of $\lambda_{\max}(\Psi)$:

Assumption E There exists $0 < \alpha \leq 1$ such that

$$\limsup_{n \rightarrow \infty} \frac{1}{n^{1-\alpha}} \lambda_{max}(\Psi) < \infty .$$

Assumptions D and E allow us to separate the two clusters of eigenvalues in the covariance matrix, one corresponding to the common components and the other to the residuals. This allows convergence to the optimal forecast.

Corollary 2 Under the assumptions of Proposition 2, and if Assumptions D and E hold, then:

$$X'_t \hat{\beta}^{bay} - y_{t+h|t}^* = O_p \left(n^{-\alpha/2} \right) + O_p \left(\frac{1}{nT \|\Phi_0\|} \right) + O_p \left(n\sqrt{T} \|\Phi_0\| \right) \text{ as } n, T \rightarrow \infty.$$

Under the assumptions of Corollary 2, when the prior is such that $\|\Phi_0\|^{-1} = cnT^{\frac{1}{2}+\delta}$, then the Bayesian forecast converges to the optimal forecast. Precisely,

$$\Delta_{nT} \left(X'_t \hat{\beta}^{bay} - y_{t+h|t}^* \right) = O_p(1) \text{ as } n, T \rightarrow \infty,$$

where $\Delta_{nT} = \min \left\{ n^{\alpha/2}, T^\delta, T^{\frac{1}{2}-\delta} \right\}$ and $0 < \delta < 1/2$.

Assumptions D and E require that the non-informative component ξ_t is driven by sources that are less pervasive than those driving the informative component ΛF_t . These conditions are satisfied under an approximate factor structure. The latter is defined by Assumptions A, B, C, D and the condition that $\lambda_{max}(\Psi)$ remains bounded as the number of predictors tends to infinity. If this is verified, the residuals ξ_t are idiosyncratic in the sense that they are weakly correlated across predictors (see Forni, Hallin, Lippi, and Reichlin, 2000,

2005; Forni, Giannone, Lippi, and Reichlin, 2007)¹⁷.

If the residuals ξ_t are idiosyncratic then Assumption E holds for $\alpha = 1$. Consequently,

$$y_{t+h|t} - y_{t+h|t}^* = O_p\left(\frac{1}{\sqrt{n}}\right) \text{ as } n \rightarrow \infty;$$

hence, from Corollary 2, when the prior is such that $\|\Phi_0\|^{-1} = cnT^{\frac{1}{2}+\delta}$, we have:

$$\Delta_{nT} \left(X_t' \hat{\beta}^{bay} - y_{t+h|t}^* \right) = O_p(1) \text{ as } n, T \rightarrow \infty,$$

where $\Delta_{nT} = \min \left\{ \sqrt{n}, T^\delta, T^{\frac{1}{2}-\delta} \right\}$ and $0 < \delta < 1/2$.

As in principal component regression, if the informative factors are pervasive and the non-informative residuals are idiosyncratic, then we converge to the optimal forecast. The intuition of this result is simple. The factor structure implies that there are few r dominant eigenvalues that diverge faster than the remaining ones as the cross-section dimension increases. The parameters' prior chosen as above ensures that the effect of the factors associated with the dominant eigenvalues is not distorted asymptotically whereas the effect of the smallest ones goes to zero asymptotically.

Notice that the rates of consistency of the Bayesian forecasts are slower than the ones derived for principal components by Bai (2003) under the assumption that the non-informative component ξ_t is idiosyncratic. The reason is that, as we have seen, the assumptions required to achieve consistency of the forecast based on Bayesian regression are more general than those implied by an approximate factor structure. In particular, the convergence of the Bayesian to the optimal forecast is achieved in the case in which Assumption E holds for $0 < \alpha \leq 1$. This can be viewed as a sort of “weak factor structure” since $\lambda_{min}(\Psi)$ can be un-

¹⁷This is a generalization to the dynamic case of the assumptions defining an approximate factor structure given by Chamberlain and Rothschild (1983). Bai (2003), Bai and Ng (2002) and Stock and Watson (2002a) give similar conditions.

bounded as $n \rightarrow \infty$. More interestingly, convergence to the population forecast based on the observed predictors $y_{t+h|t}$ holds under arbitrary correlation structure among the non-informative component ξ_t . Such generality gives flexibility to the method but at the price of a slower rate of convergence.

This result suggests that the properties of alternative methods suitable for forecasting with a large number of predictors can be studied under more general conditions than those used in the recent literature on principal components¹⁸.

Let us remark in concluding this Section that we have only studied theoretical properties of point forecasts. Under the assumption that the data follow an approximate factor structure, prediction intervals for principal components regressions are derived in Bai and Ng (2006a). For the Bayesian regression, predictive intervals can be computed from the posterior distribution, although theoretical properties for large cross-sections are not known.

5 Conclusions and open questions

This paper has analyzed the properties of Bayesian regression in large panels of time series and compared them to PCR.

We have considered the Gaussian and the double-exponential prior and showed that they offer a valid alternative to principal components. For the macroeconomic panel considered, the forecast they provide is very correlated to that of PCR and implies similar mean-square forecast errors.

This exercise should be understood as rather stylized. For the Bayesian case there is room for improvement, in particular by using developments in BVAR (Doan, Litterman, and Sims, 1984; Litterman, 1986) and related literature. We explore this conjecture in a related paper (Banbura, Giannone, and Reichlin,

¹⁸Within the framework of principal component regression, this line of research has been pursued by Onatski (2006) who studies the properties of principal components when factors are weak.

2007).

In the asymptotic analysis, we have considered the Gaussian prior case. For that case, we have shown (n, T) rates of convergence to the efficient forecast under an approximate factor structure. This analysis guides us in the setting of the prior, also interpreted as a Ridge penalization parameter. The empirical analysis reports results for the optimal parameter and for a larger range of parameter choice.

The setting of the parameters for the double-exponential case, on the other hand, has been exclusively empirical. It is designed to deliver a given number of non-zero coefficients at each estimation step in the out-of-sample evaluation period. The algorithm provides good results by selecting few variables in the regression. Selected variables, however, are not clearly interpretable, typically not the ones that a macroeconomist would include in a VAR. Moreover, the selected variables change over time. These results suggest that our data, which correspond to the typical macroeconomic dataset used for macroeconomic policy analysis, is characterized by collinearity. Under collinearity we should expect both that few appropriately selected variables capture the bulk of the covariation and that the selection is sensitive to minor perturbations of the data. In this circumstances we should not expect to obtain results that are more interpretable, from the economic point of view, than principal components or ridge regression, but we should expect comparable forecasting performance. To explore in more depth these conjectures, we should extend the double-asymptotic analysis that we have provided for the Gaussian case to the double-exponential Bayesian regression. We intend to do this in further work.

Appendix A: Proof of Proposition 1

Throughout this Section we will maintain the simplifying assumption $p = 0$ introduced in Section 4. All results still hold for the case where p lags are included, by simply replacing n by $n(p + 1)$.

Denote:

- by y_t the generic variable to be forecast as $y_t = z_{it}^h$;
- the covariance matrix of the regressors as $\Sigma_x = E(X_t X_t')$. The sample equivalent will be denoted by $S_x = X'X/T$ and the estimation error by $E_x = \Sigma_x - S_x$. These matrices are of dimension $n \times n$;
- the covariance matrix of the regressors and the variable to be predicted by $\Sigma_{xy} = E(X_t y_{t+h}')$. The sample equivalent will be denoted by $S_{xy} = X'y/T$ and the estimation error by $E_{xy} = \Sigma_{xy} - S_{xy}$. These matrices are of dimension $n \times 1$.

We assume stationarity. Moreover, we need the following assumption:

Assumption C

There exists a finite constant K , such that for all $T \in \mathbb{N}$ and $i, j \in \mathbb{N}$

$$T E[(e_{x,ij})^2] < K \quad \text{and} \quad T E[(e_{xy,i})^2] < K$$

as $T \rightarrow \infty$, where $e_{x,ij}$ denotes the i, j th entry of E_x and $e_{xy,i}$ denotes the i th entry of E_{xy} . Sufficient conditions can be found in Forni, Giannone, Lippi, and Reichlin (2007).

We consider here only the case of an i.i.d. Gaussian prior on the coefficients and we denote by $\tilde{\nu} = \frac{\sigma_u^2}{T\|\Phi_0\|}$ the rescaled penalization in the Ridge regression.

Remark 1 Notice that this does not imply that we lose in generality. Indeed, in the case of non-i.i.d. prior, we can always redefine the regression in terms of $\tilde{X}_t = \frac{\Phi_0^{1/2} X_t}{\sqrt{\|\Phi_0\|}}$. Then the corresponding rescaled regression coefficients, $\tilde{\beta} = \sqrt{\|\Phi_0\|} \Phi_0^{-1/2} \beta$, will be i.i.d. with prior variance $\|\Phi_0\|$. Moreover, under Assumption B, the transformed regressors $\tilde{X}_t$ have the representation

$$\tilde{X}_t = \tilde{\Lambda} F_t + \tilde{\xi}_t$$

where $\tilde{\Lambda} = \frac{\Phi_0^{1/2} \Lambda}{\sqrt{\|\Phi_0\|}}$ and $\tilde{\xi}_t = \frac{\Phi_0^{1/2} \xi_t}{\sqrt{\|\Phi_0\|}}$. The assumption $\liminf_{n, T \rightarrow \infty} \frac{\lambda_{\min}(\Phi_0)}{\|\Phi_0\|} > 0$ ensures that the transformed regressors still satisfy Assumptions D and E when the original regressors do.

Defining $\Sigma_x(\tilde{\nu}) = \Sigma_x + \tilde{\nu} I_n$ and the sample equivalent $S_x(\tilde{\nu}) = S_x + \tilde{\nu} I_n$, we are interested in the properties of $\beta(\tilde{\nu})$ and $\hat{\beta}(\tilde{\nu})$ which are solutions of the following linear system of equations:

$$\begin{aligned} \Sigma_x(\tilde{\nu}) \beta(\tilde{\nu}) &= \Sigma_{xy} \\ S_x(\tilde{\nu}) \hat{\beta}(\tilde{\nu}) &= S_{xy} . \end{aligned} \tag{9}$$

Notice that $\beta(0) = \beta$ is the population OLS regression coefficient and $\hat{\beta}(0) = \hat{\beta}$ its sample counterpart. For $\tilde{\nu} > 0$, we have the Ridge regression coefficients.

Lemma 1 Under the assumptions of Proposition 1 we have

$$\|\beta(\tilde{\nu})\| \leq \|\beta\| = O\left(\frac{\sqrt{\delta_{1,n}}}{\delta_{2,n}}\right) \tag{10}$$

and

$$\|\beta - \beta(\tilde{\nu})\| = O\left(\tilde{\nu} \frac{\sqrt{\delta_{1,n}}}{\delta_{2,n}^2}\right) \text{ as } n \rightarrow \infty . \tag{11}$$

Proof. We have

$$\|\beta\| = \|(\Lambda\Lambda' + \Psi)^{-1}\Lambda\gamma'\|.$$

First notice that for any vector v of the form $\Lambda\gamma'$, i.e. orthogonal to the null-space of Λ , we have the inequality $\lambda_{\min}(\Lambda'\Lambda)\|v\| \leq \|(\Lambda\Lambda' + \Psi)v\|$ since Ψ is positive definite. Then taking $w = (\Lambda\Lambda' + \Psi)v$, the inequality becomes

$$\lambda_{\min}(\Lambda'\Lambda)\|(\Lambda\Lambda' + \Psi)^{-1}w\| \leq \|w\| \quad (12)$$

and holds for any w in the range of Λ , i.e. of the form $\Lambda\gamma'$. Now, replacing w in (12) by $\Lambda\gamma'$, we get, for any γ'

$$\|(\Lambda\Lambda' + \Psi)^{-1}\Lambda\gamma'\| \leq \frac{\|\Lambda\gamma'\|}{\lambda_{\min}(\Lambda'\Lambda)} \leq \frac{\sqrt{\delta_{1,n}}}{\delta_{2,n}} \|\gamma'\|$$

where $\delta_{1,n} = \lambda_{\max}(\Lambda'\Lambda)$ and $\delta_{2,n} = \lambda_{\min}(\Lambda'\Lambda)$. Since $\|\gamma'\| = O(1)$, we get (10). The inequality $\|\beta(\tilde{\nu})\| \leq \|\beta\|$ is straightforward. Now,

$$\begin{aligned} \beta - \beta(\tilde{\nu}) &= [(\Lambda\Lambda' + \Psi)^{-1} - (\Lambda\Lambda' + \Psi + \tilde{\nu}I_n)^{-1}] \Lambda\gamma' \\ &= (\Lambda\Lambda' + \Psi)^{-1} \tilde{\nu}I_n (\Lambda\Lambda' + \Psi + \tilde{\nu}I_n)^{-1} \Lambda\gamma' \end{aligned}$$

thanks to the matrix identity

$$A^{-1} - B^{-1} = A^{-1}(B - A)B^{-1}. \quad (13)$$

Hence

$$\|\beta - \beta(\tilde{\nu})\| \leq \tilde{\nu} \|(\Lambda\Lambda' + \Psi)^{-2}\Lambda\gamma'\|.$$

Replacing w in (12) by β and using the bound (10), we easily obtain the bound (11). Q.E.D.

In contrast to what happens for $\tilde{\nu} = 0$, notice that the Ridge parameter

$\tilde{\nu}$ introduces a bias which tends to zero for large cross-sectional dimensions provided that the asymptotic behavior of $\tilde{\nu}$ as $n \rightarrow \infty$ is appropriately tuned.

Let us now consider the sample estimates and investigate the relationship between $\beta(\tilde{\nu})$ and $\hat{\beta}(\tilde{\nu})$. We first need the following lemma:

Lemma 2

- (i) $\|E_x\| = O_p\left(\frac{n}{\sqrt{T}}\right)$
- (ii) $\|E_{xy}\| = O_p\left(\frac{\sqrt{n}}{\sqrt{T}}\right)$

Proof. We have:

$$\|E_x\|^2 \leq \text{trace}[E'_x E_x] = \sum_{i=1}^n \sum_{j=1}^n e_{x,ij}^2$$

Taking expectations, we obtain:

$$\mathbb{E} \left[\sum_{i=1}^n \sum_{j=1}^n e_{x,ij}^2 \right] = \sum_{i=1}^n \sum_{j=1}^n \mathbb{E}[e_{x,ij}^2] \leq \frac{n^2 K}{T} = O\left(\frac{n^2}{T}\right)$$

We also have $\|E_{xy}\|^2 = \sum_{i=1}^n e_{xy,i}^2$. Taking expectations, we get:

$$\mathbb{E} \left[\sum_{i=1}^n e_{xy,i}^2 \right] = \sum_{i=1}^n \mathbb{E}[e_{xy,i}^2] \leq \frac{nK}{T} = O\left(\frac{n}{T}\right)$$

The results follow from the Markov inequality. Q.E.D.

Lemma 3 Under the Assumptions A, B and C, we have

$$\|\hat{\beta}(\tilde{\nu}) - \beta(\tilde{\nu})\| = O\left(\frac{\sqrt{n}}{\tilde{\nu}\sqrt{T}} \left[1 + \sqrt{n} \frac{\sqrt{\delta_{1,n}}}{\delta_{2,n}}\right]\right) \text{ as } n, T \rightarrow \infty.$$

Proof. From (9) we have

$$\hat{\beta}(\tilde{\nu}) - \beta(\tilde{\nu}) = S_x(\tilde{\nu})^{-1} S_{xy} - \Sigma_x(\tilde{\nu})^{-1} \Sigma_{xy}$$

and hence also

$$\hat{\beta}(\tilde{\nu}) - \beta(\tilde{\nu}) = S_x(\tilde{\nu})^{-1} [S_{xy} - \Sigma_{xy}] + S_x(\tilde{\nu})^{-1} \Sigma_{xy} - \Sigma_x(\tilde{\nu})^{-1} \Sigma_{xy}$$

Using again the identity (13), we get

$$\hat{\beta}(\tilde{\nu}) - \beta(\tilde{\nu}) = S_x(\tilde{\nu})^{-1} [S_{xy} - \Sigma_{xy}] + S_x(\tilde{\nu})^{-1} [\Sigma_x(\tilde{\nu}) - S_x(\tilde{\nu})] \Sigma_x(\tilde{\nu})^{-1} \Sigma_{xy}$$

whence

$$\|\hat{\beta}(\tilde{\nu}) - \beta(\tilde{\nu})\| \leq \|S_x(\tilde{\nu})^{-1}\| (\|S_{xy} - \Sigma_{xy}\| + \|\Sigma_x(\tilde{\nu}) - S_x(\tilde{\nu})\| \|\beta(\tilde{\nu})\|)$$

Using Lemma 2, the bound (10) and the fact that $\|S_x(\tilde{\nu})^{-1}\| \leq 1/\tilde{\nu}$, we get the desired result. Q.E.D.

We can now combine the results of Lemma 1 and Lemma 3, by means of the triangular inequality, and use the fact that $\|X_t\| = O_p(\sqrt{n})$ to establish the following Lemma, which is a simple corollary of the previous ones.

Lemma 4 Under Assumptions A, B and C, we have, as $n, T \rightarrow \infty$,

$$X'_t \hat{\beta}(\tilde{\nu}) = X'_t \beta + O_p \left(\tilde{\nu} \sqrt{n} \frac{\sqrt{\delta_{1,n}}}{\delta_{2,n}^2} \right) + O_p \left(\frac{n}{\tilde{\nu} \sqrt{T}} \left[1 + \sqrt{n} \frac{\sqrt{\delta_{1,n}}}{\delta_{2,n}} \right] \right).$$

Proposition 1 is now established using Lemma 4 and the definition of $\tilde{\nu} = \sigma_u^2/(T\|\Phi_0\|)$.

By a proper choice of the regularization parameter $\tilde{\nu}$ as a function of n and T , we will have to ensure that both terms tend to zero as $n, T \rightarrow \infty$. This is done in Corollary 1, under a supplementary assumption on the asymptotic behaviour of $\delta_{1,n}$ and $\delta_{2,n}$ (Assumption D).

Appendix B

An alternative to matrix inversion for computing regression estimates is provided by iterative methods as, for example, the so-called *Landweber iteration* scheme, which can be modified to cope with the penalties used in Ridge and Lasso regression.

To ensure convergence of this algorithm the norm of the sample matrices X must be smaller than 1. Since our regressors are standardized, this condition is fulfilled when using the rescaled regressors $\tilde{X} = X/\sqrt{n(p+1)(T-h-p)}$, and hence estimating the corresponding regression coefficients as

$$\tilde{\beta} = \sqrt{n(p+1)(T-h-p)} \beta.$$

Starting from the normal equation of the ordinary least squares, we can rewrite it as $\tilde{\beta} = \tilde{\beta} + \tilde{X}'y - \tilde{X}'\tilde{X}\tilde{\beta}$ and, starting from arbitrary $\tilde{\beta}^{(0)}$, try to solve it through the successive approximations scheme

$$\tilde{\beta}^{(j+1)} = \tilde{\beta}^{(j)} + \tilde{X}'y - \tilde{X}'\tilde{X}\tilde{\beta}^{(j)}; \quad j = 0, 1, \dots \quad (14)$$

which is the standard Landweber iteration. A nice feature of this scheme is that it can be easily extended to cope with additional constraints or penalties, and in particular with those used in Ridge or Lasso regression. As concerns the Lasso

functional (5), Daubechies, Defrise, and De Mol (2004) have recently proposed the following *thresholded Landweber iteration*

$$\tilde{\beta}^{(j+1)} = \mathbf{S}_\nu(\tilde{\beta}^{(j)} + \tilde{X}'y - \tilde{X}'\tilde{X}\tilde{\beta}^{(j)}); \quad j = 0, 1, \dots \quad (15)$$

where the thresholding operator is acting componentwise on a vector by performing the soft-thresholding operation defined by (6) and is thus defined by

$$\mathbf{S}_\nu(\tilde{\beta}) = [S_\nu(\tilde{\beta}_i)]_{i=1, \dots, n} . \quad (16)$$

This operation enforces the sparsity of the regression coefficients in the sense that all coefficients below the threshold $\nu/2$ are set to zero. The scheme (15) has been proved in Daubechies, Defrise, and De Mol (2004) to converge to a minimizer of the Lasso functional (5). Let us remark that this functional fails to be strictly convex when the null-space of $\tilde{X}$ is not reduced to zero, in which case the minimizer of (5) is not necessarily unique.

Acknowledgments

The paper has been prepared for the conference to honor the 25th anniversary of Beveridge and Nelson's JME paper, in Atlanta March 31st-April 1st, 2006. We would like to thank an anonymous referee, Marta Banbura, Michel Defrise, James Hamilton, James Nason, Christian Schumacher, Farshid Vahid, Peter Vlaar, Mark Watson for useful comments, and also seminar participants at the Atlanta Federal Reserve, the 8th Bundesbank spring conference, the 5th IAP workshop, Louvain-la-Neuve, the conference on Macroeconometrics and Model Uncertainty at the Reserve Bank of New Zealand, the 2006 Australasian meeting of the Econometric Society, the 26th International Symposium on Forecasting. The opinions in this paper are those of the authors and do not necessarily re-

flect the views of the European Central Bank. Support by the grants “Action de Recherche Concertée” Nb 02/07-281 and IAP-network in Statistics P5/24 is gratefully acknowledged. Replication files as well as an appendix containing supplementary material are available at: <http://homepages.ulb.ac.be/~dgiannon/> or <http://homepages.ulb.ac.be/~reichli/>.

References

- BAI, J. (2003): “Inferential Theory for Factor Models of Large Dimensions,” *Econometrica*, 71(1), 135–171.
- BAI, J., AND S. NG (2002): “Determining the Number of Factors in Approximate Factor Models,” *Econometrica*, 70(1), 191–221.
- (2006a): “Confidence Intervals for Diffusion Index Forecasts and Inference for Factor-Augmented Regressions,” *Econometrica*, 74(4), 1133–1150.
- (2006b): “Forecasting Economic Time Series Using Targeted Predictors,” *Journal of Econometrics*, this issue.
- BANBURA, M., D. GIANNONE, AND L. REICHLIN (2007): “Bayesian VARs with Large Panels,” Discussion Paper 6326, Center for Economic Policy Research.
- CHAMBERLAIN, G., AND M. ROTHCHILD (1983): “Arbitrage, factor structure and mean-variance analysis in large asset markets,” *Econometrica*, 51, 1305–1324.
- CHEN, S. S., D. DONOHO, AND M. SAUNDERS (2001): “Atomic Decomposition by Basis Pursuit,” *SIAM Review*, 43, 129–159.

- D'AGOSTINO, A., AND D. GIANNONE (2007): "Comparing Alternative Predictors Based on Large-Panel Factor Models," Discussion Paper 6564, Center for Economic Policy Research.
- D'AGOSTINO, A., D. GIANNONE, AND P. SURICO (2006): "(Un)Predictability and Macroeconomic Stability," Working Paper Series 605, European Central Bank.
- DAUBECHIES, I., M. DEFRISE, AND C. DE MOL (2004): "An iterative thresholding algorithm for linear inverse problems with a sparsity constraint," *Communications on Pure and Applied Mathematics*, 57, 1416–1457.
- DE MOL, C., AND M. DEFRISE (2002): "A note on wavelet-based inversion methods," in *Inverse Problems, Image Analysis and Medical Imaging*, ed. by M. Z. Nashed, and O. Scherzer, pp. 85–96. American Mathematical Society.
- DOAN, T., R. LITTERMAN, AND C. A. SIMS (1984): "Forecasting and Conditional Projection Using Realistic Prior Distributions," *Econometric Reviews*, 3, 1–100.
- EFRON, B., T. HASTIE, I. JOHNSTONE, AND R. TIBSHIRANI (2004): "Least angle regression," *Annals of Statistics*, 32, 407–499.
- FERNANDEZ, C., E. LEY, AND M. F. J. STEEL (2001): "Benchmark priors for Bayesian model averaging," *Journal of Econometrics*, 100(2), 381–427.
- FORNI, M., D. GIANNONE, M. LIPPI, AND L. REICHLIN (2007): "Opening the Black Box: Structural Factor Models with large cross-sections," Working Paper Series 712, European Central Bank.
- FORNI, M., M. HALLIN, M. LIPPI, AND L. REICHLIN (2000): "The Generalized Dynamic Factor Model: identification and estimation," *Review of Economics and Statistics*, 82, 540–554.

- (2004): “The generalized dynamic factor model consistency and rates,” *Journal of Econometrics*, 119(2), 231–255.
- (2005): “The Generalized Dynamic Factor Model: one-sided estimation and forecasting,” *Journal of the American Statistical Association*, 100, 830–840.
- FU, W. J. (1998): “Penalized Regressions: The Bridge Versus the Lasso,” *Journal of Computational and Graphical Statistics*, 7, 497–416.
- GIACOMINI, R., AND H. WHITE (2006): “Tests of Conditional Predictive Ability,” *Econometrica*, 74(6), 1545–1578.
- GIANNONE, D., L. REICHLIN, AND L. SALA (2004): “Monetary Policy in Real Time,” in *NBER Macroeconomics Annual*, ed. by M. Gertler, and K. Rogoff, pp. 161–200. MIT Press.
- GIANNONE, D., L. REICHLIN, AND D. SMALL (2005): “Nowcasting GDP and inflation: the real-time informational content of macroeconomic data releases,” Finance and Economics Discussion Series 2005-42, Board of Governors of the Federal Reserve System (U.S.), *Journal of Monetary Economics*, forthcoming.
- HASTIE, T., R. TIBSHIRANI, AND J. FRIEDMAN (2001): *The Elements of Statistical Learning*. Springer-Verlag.
- KOOP, G., AND S. POTTER (2003): “Forecasting in large macroeconomic panels using Bayesian Model Averaging,” Staff Reports 163, Federal Reserve Bank of New York.
- LITTERMAN, R. (1986): “Forecasting With Bayesian Vector Autoregressions – Five Years of Experience,” *Journal of Business and Economic Statistics*, 4, 25–38.

- ONATSKI, A. (2006): “Asymptotic distribution of the principal components estimator of large factor models when factors are relatively weak,” Manuscript, Columbia University.
- STOCK, J. H., AND M. W. WATSON (2002a): “Forecasting Using Principal Components from a Large Number of Predictors,” *Journal of the American Statistical Association*, 97, 147–162.
- (2002b): “Macroeconomic Forecasting Using Diffusion Indexes,” *Journal of Business and Economics Statistics*, 20, 147–162.
- (2005a): “An Empirical Comparison Of Methods For Forecasting Using Many Predictors,” Manuscript, Princeton University.
- STOCK, J. H., AND M. W. WATSON (2005b): “Implications of Dynamic Factor Models for VAR Analysis,” NBER Working Papers 11467, National Bureau of Economic Research, Inc.
- (2006): *Forecasting with Many Predictors* vol. 1 of *Handbook of Economic Forecasting*, chap. 10, pp. 515–554. Elsevier.
- TIBSHIRANI, R. (1996): “Regression shrinkage and selection via the lasso,” *Journal of the Royal Statistical Society: Series B*, 58, 267–288.
- WRIGHT, J. H. (2003): “Forecasting U.S. inflation by Bayesian Model Averaging,” International Finance Discussion Papers 780, Board of Governors of the Federal Reserve System (U.S.).

Table 1: Principal component forecasts

Industrial Production							
	Number of Principal Components						
	1	3	6	10	25	50	75
MFSE 1971-2002	0.91	0.62	0.56	0.54	0.65	0.93	1.56
MFSE 1971-1984	0.89	0.45	0.35	0.34	0.46	0.70	1.18
MFSE 1985-2002	0.98	1.13	1.16	1.13	1.21	1.60	2.68
Variance*	0.23	0.70	0.79	0.97	1.28	1.43	1.78

Consumer Price Index							
	Number of Principal Components						
	1	3	6	10	25	50	75
MFSE 1971-2002	0.57	0.55	0.57	0.69	0.83	1.17	1.69
MFSE 1971-1984	0.48	0.40	0.39	0.48	0.56	0.89	1.23
MFSE 1985-2002	1.03	1.28	1.43	1.71	2.11	2.47	3.83
Variance*	0.36	0.55	0.61	0.63	0.69	0.89	1.69

MSFE are relative to a the Naive, Random Walk, forecast. *The variance of the forecast relative to the variance of the series.

Table 2: Bayesian forecasts with Gaussian prior

Industrial Production									
	In-sample Residual variance								
	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
ν	6	25	64	141	292	582	1141	2339	6025
MFSE 1971-2002	0.96	0.70	0.60	0.56	0.56	0.58	0.64	0.72	0.83
MFSE 1971-1984	0.74	0.50	0.41	0.38	0.40	0.44	0.52	0.63	0.78
MFSE 1985-2002	1.59	1.31	1.16	1.08	1.03	1.00	0.98	0.98	0.98
Variance*	0.71	0.63	0.57	0.49	0.39	0.29	0.19	0.12	0.07
Correlation with PC forecasts (r=10)	0.62	0.81	0.89	0.92	0.93	0.91	0.85	0.74	0.48

Consumer Price Index									
	In-sample Residual variance								
	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
ν	16	60	143	288	528	949	1751	3532	9210
MFSE 1971-2002	0.88	0.72	0.66	0.63	0.62	0.63	0.66	0.73	0.84
MFSE 1971-1984	0.72	0.58	0.52	0.51	0.51	0.54	0.59	0.68	0.82
MFSE 1985-2002	1.60	1.41	1.29	1.19	1.11	1.04	0.98	0.95	0.95
Variance*	0.41	0.35	0.32	0.28	0.24	0.19	0.13	0.08	0.05
Correlation with PC forecasts (r=10)	0.68	0.86	0.92	0.94	0.92	0.89	0.83	0.69	0.33

MSFE are relative to a the Naive, Random Walk, forecast. *The variance of the forecast relative to the variance of the series.

Table 3: Lasso forecasts

Industrial Production							
	Number of non-zero coefficients						
	1	3	5	10	25	50	75
MFSE 1971-2002	0.86	0.69	0.64	0.60	0.64	0.77	1.10
MFSE 1971-1984	0.80	0.56	0.50	0.44	0.47	0.58	0.91
MFSE 1985-2002	1.05	1.05	1.05	1.07	1.14	1.32	1.67
Variance*	0.07	0.16	0.24	0.40	0.53	0.65	0.79
Correlation with PC forecasts (r=10)	0.05	0.64	0.81	0.85	0.84	0.68	0.44

Consumer Price Index							
	Number of non-zero coefficients						
	1	3	5	10	25	50	75
MFSE 1971-2002	0.90	0.76	0.62	0.59	0.68	0.86	1.06
MFSE 1971-1984	0.88	0.70	0.54	0.48	0.52	0.70	0.93
MFSE 1985-2002	1.00	1.04	1.02	1.14	1.44	1.65	1.68
Variance*	0.05	0.09	0.18	0.26	0.33	0.39	0.50
Correlation with PC forecasts (r=10)	0.05	0.64	0.81	0.85	0.84	0.68	0.44

MSFE are relative to a the Naive, Random Walk, forecast. *The variance of the forecast relative to the variance of the series.