G-7 Inflation forecasts.

Fabio Canova *

Universitat Pompeu Fabra, LBS and CEPR

This draft, March 2002

Abstract

This paper compares the forecasting performance of some leading models of inflation for the cross section of G-7 countries. We show that bivariate and trivariate models suggested by economic theory or statistical analysis are hardly better than univariate models. Phillips curve specifications fit well into this class. Significant improvements in both the MSE of the forecasts and turning point prediction are obtained with time varying coefficients models which exploit international interdependencies. The performance of the latter class of models is independent of the sample, while it is not the case for standard specifications.

Key words: Forecasting, Inflation, Panel VAR models, Markov Chain Monte Carlo Methods.

*Part of this work was conducted while the author was visiting the DG Research at the European Central Bank. I would like to thank David Hendry, Jim Stock, Lutz Kilian, the participants of the ECB workshop on "Forecasting Techniques", and of seminars at Tilburg, Oxford, Bocconi, Ente Einaudi, York and Bank of England for useful comments and suggestions.

1 Introduction

Since Central Banks have abandoned the control of monetary aggregates and implemented inflation targeting rules directly or indirectly, by means of aggressive Taylor rules, forecasting inflation rates has become crucial for both policy makers and private agents who try understand and react to Central Banks behavior. Several methods have been proposed to predict inflation rates but the overall performance has been, at best, mixed: the information provided by past inflation appears to suffice and very few other variables add marginal predictive content to univariate specifications (see e.g. Cecchetti, Chu and Steindel (2000)).

One of the most standard approaches employed in the literature builds on a traditional Phillips curve trade-off. Blinder (1997) attributes the popularity of this specification among practitioners to the stability of the relationship and its reliability as forecasting tool. Stock and Watson (1999) criticized this conventional wisdom, showing that in the US, the parameters of a standard Phillips curve have changed over time; that other measures of real economic activity forecasts inflation better than the unemployment rate, and that combining forecasts produced with different specifications improves the performance of a model which only uses the unemployment rate. Atkeson and Ohanian (2001) reinforce the argument by showing that Phillips curve forecast of US inflation over the last 15 years are no better than those obtained from a random walk model. Stock and Watson (2000) extend their study to consider the contribution of asset prices to forecast US inflation, and further demonstrate the instability of several forecasting relationships (see Goodhart and Hoffman (2000) for a similar exercise in the UK). Finally, Marcellino, Stock and Watson (2000) examine the problem of forecasting inflation rates in a number of industrialized countries.

In this paper we compare the forecasting performance of some leading models for inflation in G-7 countries. We focus the analysis on four different but interrelated questions. First, we would like to investigate how models which rely on a Phillips curve relationship fare relative to univariate models or to other bivariate models which use alternative variables suggested by economic theory. Second, we are interested in knowing whether there is information in other domestic variables (which may or may not have theoretical underpinning) useful for forecasting inflation. Third, we would like to know whether models which allow for drifts over time in the parameters improve the performance of specifications with fixed coefficients. Finally, we want to study whether the information coming from the cross section of the G-7 countries is useful in forecasting domestic inflation rates.

Our analysis concentrates forecasting CPI inflation but the results we present are qualitatively similar to those obtained using GDP deflator inflation. We are concerned with three forecasting horizons, which are of interest to policymakers: one quarter, one year and two years. We calculate statistics measuring the average magnitude of the forecast errors over time and the ability to predict turning points of inflation. Our sample covers quarterly data from 1980:1 to 2000:4 and the last five years are typically used to evaluate the performance of various specifications. All the comparisons are performed in real time: models are estimated recursively and forecasts computed using the information available at each t.

We consider several models in our horse race. First, as a benchmark, we estimate ARIMA models, identified with standard approaches, and simple AR models, where the lag length optimizes the trade-off between number of lags and the in-sample forecasting performance. To address the first question we consider nine alternative bivariate AR specifications: two with a standard Phillips curve relationship (standardized unemployment rate and employment growth) and one using a new-keynesian version of it (labor income share), two with a measure of real activity (output gap and GDP growth), two with monetary variables (nominal M2 growth and real M2 growth) and two with financial indicators (the slope of the nominal term structure and nominal stock returns). We show that, among bivariate models, those employing some version of a traditional Phillips curve are the preferable in terms of MSE in France, Canada and Italy, those which use output measures perform better in US, Japan and Germany and in the UK the best bivariate model includes inflation and stock returns. We find instances where bivariate models produce inflation forecasts which have a lower MSE than those produced by univariate models at some horizon. However there are also cases when the opposite is true. Furthermore, in those countries where a Phillips curve model was selected the MSE is broadly equivalent to the one of univariate models. As in Stock and Watson (2000) we find that conclusions concerning the relative ordering of models are somewhat tenuous: changing the forecasting sample at times alters the ranking of bivariate models and their MSE performance relative to univariate models. Finally, the performance of the collection of the best bivariate models in terms of turning point predictions is generally similar to that of univariate specifications, regardless of the horizon and the forecasting sample.

To examine whether there are other variables which help to forecast inflation, we run a battery of bivariate Granger causality tests and estimate trivariate VARs and BVAR models. We show that the MSE of the forecasts of bivariate and univariate models can be improved in some countries, but the gains are far from being systematic and they obtain mainly at longer horizons. Adding a Bayesian prior reduces the MSE of a VAR model in six of the seven countries but on average a BVAR is no better than univariate AR models. With the BVAR, turning point predictions improve at the one year horizon.

To study whether the presence of drifts in the coefficients improves the quality of the forecasts, we estimate univariate AR and BVAR models with time varying coefficients. These models easily adapt when slow or smooth changes occur and have the potential to improve over fixed time coefficient models. The added flexibility does not significantly alter the MSE of the forecasts of univariate AR models, except for the UK. However, turning point recognition improves significantly, primarily at the one quarter horizon. Adding time variation to the BVAR model, on the other hand, significantly reduces the MSE of the forecasts: the average gains are about 5% at long horizons, both relative to a fixed coefficient BVAR and relative to a univariate time varying coefficient model. The same improvement, although less systematic, is observable in turning point detection.

Finally, to examine whether international information helped to forecast domestic inflation, we consider three specifications: one which uses as predetermined in country-specific BVAR variables capturing the state of the world economy (commodity prices, US GDP growth and the median stock return in the G-7); a BVAR model which includes the seven inflation rates and the same three predetermined variables considered in the previous model; and an index version of panel BVAR model studied by Canova and Ciccarelli (1999) and (2001) where cross section information is a-priori shrunk into indices capturing common, country specific or variable specific effects. In all cases, time varying coefficients are allowed in the specification. Adding international variables to domestic models has some use in forecasting. In five of the seven countries the MSE obtained with the first specification is smaller than the MSE of a domestic model. The gains are clearly evident at long horizons where improvements of the order of 30% are not unusual. This improvement is also matched by more appropriate turning point recognition: this model produced the smallest number of missing signals of all the specifications at the one quarter horizon. The second specification improves over domestic models in some countries, primarily at longer horizons but, overall, its performance seems to be inferior to the previous one. Finally, the index model produces the best overall MSE of the forecasts at all horizons for Japan, Germany, France and the UK. MSE gains are more pronounced at medium-long horizons and robust across forecasting samples: in fact, forecasts produced by this model one and two years ahead track reasonably well actual inflation in the G-7 throughout the 1990's, with no visible change in the performance across the two halves. The record of turning point predictions is unimpressive but this is entirely due to the poor performance in forecasting US inflation rates.

In summary, the by now conventional wisdom that the simplest and the most naive models forecast inflation rates better it is not necessarily true. Standard Phillips curve models and other small scale specifications which use fixed coefficients and exploit only domestic interdependencies, are not much better than univariate AR models in several of the G-7 countries. However, multivariate models with time varying coefficients and some international linkages are superior to univariate models both with fixed and time varying coefficients. To a keen observer, this conclusions is perhaps not surprising: the process of inflation displayed instabilities over the period and there appears to be a common component in the international swings of inflation over time. Our results show that it is the combination of these two factors that helps in prediction: models with these features are useful for forecasting inflation in the medium term because they provide information about the time varying (local) trends present in almost all G-7 countries. While the performance of fixed coefficient models seems to be sample dependent, models like these provide a reliable and robust tool for tracking the behavior of time series with instabilities, clustering and varying local trends like inflation. For this reason they should be attractive to practitioners who are engaged in repeated out-of-sample predictions.

The rest of the paper is organized as follows: next section describes the specifications used and the data employed. Section 3 presents the statistics used. Section 4 examines the forecasting performance of the various specifications. Section 5 concludes.

2 The Specifications

Our sample covers seasonally adjusted quarterly data from 1980:1 to 2000:4 and we use the last five years to evaluate the performance of various forecasting specifications. In some exercises, we use the 1980's to estimate various models and the sample 1991:1-1995:4 to evaluate the forecasting performance. Details about the sources and the construction of the variables are in Appendix A. All the comparisons are performed using a real time methodology That is, models are estimated recursively and forecasts computed at one, four and eight quarters ahead using the information available at each t.

We collect two types of forecasting statistics, the MSE and the number of correct turning point signals. To avoid problems with the of units of measurement of variables in different models, we present Theil-U statistics (the ratio of the MSE of each model to the MSE produced by a no-change model) which immediately allows us to compare the performance relative to a random walk model (as in Atkeson and Ohanian (2001)). Results obtained using Mean Absolute Deviation (MAD) statistics are broadly similar and omitted for reason of space. Turning point predictions are calculated using a two quarter rule, as detailed in the next section. Complementing synthetic MSE measures with turning point detection is important because MSE reduction can be achieved by altering the timing and the direction of the forecasts. Therefore, an evaluation exercise based on both gives a better understanding of the properties of various specifications and an implicit robustness check on the conclusions.

2.1 Univariate Models

We consider two univariate specifications: Box-Jenkins and AR models. For the first type of models the specification is identified using plots of autocorrelation and the partial autocorrelation functions. Since some of the inflation series display residual seasonality, multiplicative unobservable component models, directly specifying the seasonal and nonseasonal components of the data, are selected. For AR models, we follow Ivanov and Kilian (2000) and use the Hannan and Quinn, the Akaike and the Schwarz criteria to select the lag length of the autoregression. Each of these methods trade-off the increase in the explanatory power obtained with additional lags with a penalty on the dimensionality of the model. When there were disparities across methods in selecting the lag length of the autoregression, we choose the one which was indicated by the majority of methods. When this was not possible, we picked the one which gave the smaller MSE one quarter ahead.

2.2 Bivariate Models

We considered nine bivariate AR models. Together with inflation we use unemployment, employment growth, the labor income share, output gap, output growth, money growth, real money growth, the slope of the term structure or stock returns. This choice of these variable is motivated by straightforward theoretical arguments. The first five specifications embody notions of traditional and New-Keynesian Phillips curve (see Mankiw (2000), Gali and Gertler (1999)). (Real) Money growth has been selected on the basis that inflation in the long run must be a monetary phenomena; the slope of the term structure is selected following work by Stock and Watson (1989) and Plosser and Rowenshort (1994) and the notion that monetary shocks are typically transmitted to the real economy via changes in the slope of the term structure at different horizons (see e.g. Evans and Marshall (1998)). Nominal stock returns are used since theoretically, they are considered a good hedge against inflation (see Fama (1970)). For all bivariate models the lag length was chosen using the same three criteria and the same selection rules used in the univariate cases.

2.3 VARs

The VAR models we consider include domestic variables only. The series are selected using in-sample bivariate Granger causality tests and the results of our previous bivariate forecasting exercises. We search for the best specification using GDP (growth and gap), stock returns, the slope of the term structure, wage and rent inflation, real and nominal exchange rate, unit labor costs inflation (when available), capacity utilization, unemployment rate and employment (rate and level) and the labor force. To avoid excessive data mining and large specifications, we restricted the models to have three variables and choose among all possible combinations of the 14 variables the two which have the best in-sample explanatory power for inflation in each of the G-7. The chosen variables differ across countries : for the US the trivariate model include inflation, output gap and the labor force; for Japan, inflation, the output gap and the real exchange rate; for Germany, inflation, employment growth and unit labor costs; for the UK, inflation, stock returns and unit labor costs; for France, inflation, rent and wage inflation; for Italy, inflation, the unemployment rate and rent inflation; for Canada the output gap and stock returns. The heterogeneity is more apparent than practical since many variables are closely correlated in the estimation sample. Once again, the lag length in each country is selected using the HQ, Akaike and Schwarz criteria.

2.4 BVARs

The Bayesian VAR (BVAR) models we consider here are standard and, at this stage, we keep the prior specification very simple. Since we are interested in a relative comparison, the variables for each country are the same used in the VAR. We write the BVAR model as:

$$y_t = B_t Y_{t-1} + C_t W_{t-1} + e_t \quad (1)$$
$$D_t = D_0 + u_t \quad (2)$$

where Y_{t-1} is a stacked version of the lags of the vector of depended variables y_t, W_{t-1} includes lags of the exogenous variables, $D_t = [B_t', C_t''] \; e_t \sim (0, \Sigma_e)$ and $u_t \sim (0, \Sigma_u)$. (2) describes the prior of the model. We follow the standard specification of the so-called Minnesota Prior and set $D_0 = 0$, except on the first lag of the dependent variable in each equation and use one parameter θ_1 to trade off the forecasting information in the mean of the first lag. The structure of Σ_u is standard: one parameter controls the general tightness of the specification (θ_2), one controls the decay of the prior variance as the lag length increases (θ_3), and one weights the relative contribution of lags of other variables in each equation (θ_4). The prior variance for the exogenous variables is diffuse ($\theta_5 = 1000$). Estimates of the four parameters can obtained informally, using a rough specification search over a grid of values, and formally, maximizing the predictive density of the model in an Empirical-Bayes fashion. Here, we use the second approach and we revert to informal searches whenever the maximization algorithm failed to converge. Estimates of the hyperparameter vector Θ are similar but not identical across countries. We report them in table 2.

2.5 Time varying coefficient models

We consider two types of time varying coefficients (TVC) models: a univariate AR and a VAR specification. Since their general structure is similar we describe them compactly in one single framework. The models have the following structure:

$$y_t = B_t Y_{t-1} + C_t W_{t-1} + e_t \quad (3)$$
$$D_t = P D_{t-1} + (I - P) D_0 + u_t \quad (4)$$

where P is a matrix of numbers, I the identity matrix, and all the other variables have been previous defined. y_t is a scalar in the univariate model and it is a 3×1 vector in the VAR. The major difference between (2) and (4) is that D_t is allowed to evolve over time in a geometric fashion and that the variance of u_t (denoted by Σ_{ut}) is allowed to vary over time. We parametrize these two additional features in a very simple fashion: we let $P = \theta_6 * I$ with $\theta_6 < 1$ and set $\Sigma_{ut} = \theta_7 * \Sigma_{u0}$ where Σ_{u0} has the same structure as Σ_u described in the previous subsection (for univariate models, the parameter controlling the importance of lag of other variables $\theta_4 \equiv 1$). Here θ_7 controls how much time variation there is in the evolution of the law of motion of the coefficients and θ_6 controls the extent of the mean reversion of the coefficient vector. Clearly, if $\theta_6 = 1$ and $\theta_7 = 0$, no time variation in the coefficients is necessary. In this specification we also allow a different structure for D_0, along the lines of

Canova (1992), (1993a) to take into account of predictability over the calendar year. That is, the prior mean of the first lag of the dependent variable is θ_1 and the prior mean of the fourth lag of the dependent variable is $1-\theta_1$. Estimates of the vector Θ for the two specifications are obtained by maximizing the predictive density and when the maximizing algorithm fail to convergence the analysis is supplemented with a rough grid search over the relevant hypercubes. As with BVARs, the final specification turns out to be similar although not identical across countries. Hyperparameter estimates for the two models are in table 2.

2.6 International Models

We consider three different specification in this class of models. The first one adds as predetermined to the previous BVAR model (with or without time varying coefficients, whichever is the best) variables which may control for demand and supply pressures in international markets. The variables we use are: a commodity price index, the US GDP growth and the median stock returns across the G-7. In all countries only one lag of these variables enters the specification. As with the other exogenous variables, the prior is assumed to be diffuse (the parameter controlling the prior variance is set to 1000).

The second specification is a BVAR with the G-7 inflation rates augmented with the same international factors used in the previous specification. The prior for this model depends on the same seven hyperparameters we have used for the basic BVAR with time varying coefficients and estimates are obtained in the same fashion (see again in table 2).

Since these two specifications take a rough short cut to the problem of specifying interdependencies which may leave out important aspects of the international transmission, the third specification we consider attempts to explicitly model these features. A multicountry panel VAR model can be written as :

$$y_{it} = B_{it,l}(L)Y_{t-l} + C_{it}(L)W_t + e_{it} \tag{5}$$

where $i = 1, ..., N;\ t = 1, ..., T;\ y_{it}$ is a $G \times 1$ vector for each country i, $Y_t = (y'_{1t}, y'_{2t}, \ldots y'_{Nt})'$, $B^j_{it,l}$ are $G \times G$ matrices, C^j_{it} is $G \times q$, W_t is a $q \times 1$ vector of exogenous variables, common to all units, and e_{it} is a $G \times 1$ vector of random disturbances. In this specification, we have p lags for the NG endogenous variables and l lags for the q exogenous variables. Here, interdependencies among units are allowed whenever $B^j_{it,r} \neq 0$ for $j \neq i$ and for any r. This generalization is not without costs: the number of parameters in the model is greatly increased (we have now $k = (NGp + ql)$ parameters each equation); furthermore the G variables entering the VAR must be the same for each unit.

Canova and Ciccarelli (1999) showed that cross country interdependencies are important in predicting output growth across countries. To solve the problem of the dimensionality of the parameter vector they assumed a two components hierarchical structure: the coefficients have a time invariant component, which randomly varies over the cross section, and a time varying one which drifts over time as in (4). In that specification the dimensionality of the time varying component may still be large. To circumvent this problem, Canova and Ciccarelli (2001) suggest an alternative index structure on the parameter vector, which drastically reduces the computational costs and provides posterior estimates which are, in general, more reliable. To adapt their structure to the present context, rewrite (5) as:

$$Y_t = M_t \delta_t + E_t \tag{6}$$

where $M_t = I_{NG} \otimes X_t'$; $X_t = \left(Y_{t-1}',\ Y_{t-2}', \ldots, Y_{t-p}',\ W_t', \ldots, W_{t-l}'\right)'$; $\delta_t = (\delta_{1t}', \ldots, \delta_{Nt}')'$ and $\delta_{it} = (\gamma_{it}^{1\prime}, \ldots, \gamma_{it}^{G\prime})'$. Here γ_{it}^{g} are $k \times 1$ vectors containing, stacked, the g rows of the matrices B_{it} and C_{it}, while Y_t and E_t are $NG \times 1$ vectors containing the endogenous variables and the random disturbances of the model.

Whenever δ_{it} vary with cross–sectional units in different time periods, it is impossible to obtain meaningful estimates using classical methods. Two short cuts are typically employed: either it is assumed that the coefficient vector does not depend on the unit, apart from a time invariant fixed effect, or that there are no interdependencies across units (see e.g. Holtz Eakin et al. (1988)). Neither of these assumptions is appealing in our context. Our approach views each coefficient as random with a given probability distribution (as we have done with TVC models). We assume that δ_t takes the form:

$$\delta_t = \Xi_1 \alpha + \Xi_2 \lambda_t + \Xi_3 \rho_t \tag{7}$$

where Ξ_1 is a matrix of ones and zeros of dimensions $N \cdot G \cdot k \times N$, Ξ_2 is a vector of ones of dimensions $N \cdot G \cdot k \times 1$, Ξ_3 is a $N \cdot G \cdot k \times G$ matrix of zeros and ones. We let $\beta_t = (\lambda_t, \rho_t)$. Furthermore, we assume

1. The $N \times 1$ vector α_i is normally distributed

$$\alpha \sim N\left(M_0 \mu, B_0\right) \tag{8}$$

   where M_0 is a $N \times m$, $\ m \leq N$ matrix and rank $(B_0) = N$. Also, the mean vector μ has a normal distribution

$$\mu \sim N\left(\bar{\mu}, \Sigma_\mu\right) \tag{9}$$

   The structure in (8)-(9) allows some degree of a-priori pooling of cross sectional information via an exchangeable prior on α_i. This may be useful when there are similarities in the characteristics of the vector of variables across units since coefficients of unit j may contain useful information for estimating the coefficients of unit i.

2. The $(1 + MG) \times 1$ vector β_t is normally distributed $\beta_t \sim N(M_1 \theta_t, B_1)$ where

$$\theta \sim N\left((I - C)\theta_0 + C\theta_{t-1}), B_{2t}\right) \tag{10}$$

and where M_1 is of dimension $(1+G) \times s$, $s \leq (G+1)$ and rank$(C) = s$. We also assume that $B_{2t} = \nu_1 B_{2t-1} + \nu_2 B_{20}$ and that B_1 and B_{20} are positive semidefinite, block diagonal, symmetric matrices of rank $(G+1)$ and s respectively.

In principle (10) allows some further dimensionality reduction if the vector ρ has components for which an exchangeable prior is appropriate. The structure underlying the variance structure of β_t is from Canova (1993b) and allows for nonlinearities in the moment structure (of both ARCH-M and Markov switching type) and nonnormalities in inflation process. Besides being useful to directly capture volatility clustering which are common across countries, time variations in the variance allow the model to quickly adapt when outliers or regime switches of short mean length are present.

3. Conditional on M_t, the vector of random disturbances E_t has a normal distribution

$$E_t \sim N(0, \Sigma_e) \tag{11}$$

where $\Sigma_e = \Sigma \otimes H$, Σ is $N \times N$ and H is $G \times G$, both positive definite and symmetric matrices.

Given this structure, the Panel VAR model can be written as

$$Y_t = \mathbf{W}_t \lambda_t + \mathbf{A}_t \alpha + \mathbf{Z}_t \rho_t + \varepsilon_t \tag{12}$$

where clearly $\mathbf{W}_t = M_t \Xi_2$, $\mathbf{Z}_t = M_t \Xi_3$, and $\mathbf{A}_t = M_t \Xi_2$.

Note that in (12) the $NG \times 1$ vector Y_t depends on a $N \times 1$ vector of unit specific coefficients α, a common time coefficient λ_t and on an $G \times 1$ vector of variable specific coefficients (for a total of $N + G + 1$ coefficients as opposed to the original NGk coefficients). Additional indices relating for example to the lags, and/or combination of all the previous ones (e.g. country-variable, country-lag, etc.) can be similarly added to the specification. Note also, that because of the prior specification, the regressors of (12) are particular combinations of lags of the VAR variables. The coefficients λ_t, α, ρ_t can therefore be interpreted as "factor loadings" and measure the impact that different linear combinations of the lags of the endogenous and exogenous variables have on the current endogenous variables.

Hence, we have reparametrized the panel VAR model to have a hierarchical index structure composed of (12)-(11)-(8)-(9)-(10), which is very convenient to handle.

Several structures are nested in the general specification. For instance, when $B_0 = 0$, there is no cross sectional heterogeneity. If $C = I$, coefficients evolve over time as a random walk, while when $C = I$, and $B_1 = 0$ the model reduces to a standard dynamic panel model with no time–variation in the coefficient vector. Finally, when $B_0 = B_1 = 0$, $\nu_1 = \nu_2 = 0$, and $C = I$ neither heterogeneity nor time variation are present.

Posterior estimates for this model cannot be computed analytically, given the large number of nuisance parameters which need to be integrated out of the specification. We therefore resort to Monte Carlo Markov Chain methods to draw sequences from of the posterior distribution. The Gibbs sampling algorithm is useful in our case because it allows sampling from the joint distribution, once a vector of conditional posteriors for single parameters is available analytically. Given the assumptions made so far and a suitable specification for the distribution of the remaining parameters, conditional posterior distributions are easy to compute. Details on the algorithm are in appendix B.

Stock and Watson (1999), Forni et al. (1999) and others have examined macro panel models where either N or G or both are large. Their approach is to set up the problem so that it can be handled in the context of (dynamic) index models with classical methods. Three major features differentiate our approach from their: first, the indices are a-prior determined by the interest of the researcher through the specification of the hierarchical prior. Second, the coefficients on our indices are allowed to very over time. Third, our approach delivers exact estimates no matter what the dimension of N or G are while their inferential methods require asymptotic approximations.

3 The forecasting statistics

To compare the forecasting performance of the various specifications and the improvements obtained by the addition of each feature, we report two types of statistics: a linear type (Theil-U statistics at 1, 4, and 8 steps ahead, or ratio of MSE of the model to the best univariate model at these three horizons) and a nonlinear one (turning point prediction).

Turning point predictions can be computed from the predictive density of future observations. We follow Canova and Ciccarelli (1999) and define turning points as follows:

Definition 3.1 A downward turn for unit i at time $t+h+1$ occurs if S_{it+h}, the reference variable, satisfies for all h $S_{it+h-2},\ S_{it+h-1} < S_{it+h} > S_{it+h+1}$. An upward turn for unit i at time $t+h+1$ occurs if the reference variable satisfies $S_{it+h-2},\ S_{it+h-1} > S_{it+h} < S_{it+h+1}$.

Similarly, we define a non-downward turn and a non-upward turn:

Definition 3.2 A non-downward turn for unit i at time $t+h+1$ occurs if S_{it+h} satisfies for all h $S_{it+h-2},\ S_{it+h-1} < S_{it+h} \leq S_{it+h+1}$. A non-upward turn for unit i at time $t+h+1$ occurs if S_{it+h} satisfies $S_{it+h-2},\ S_{it+h-1} > S_{it+h} \geq S_{it+h+1}$.

Although there are other definitions in the literature (see e.g. Lahiri and Moore (1991)) this is the most used one and it suffices for our purposes. Let $\tilde{f}(Y_{i,t+h} \mid F_t) = \int_{Y_{p,t+h}} f(Y_{t+h} \mid F_t) dY_{p,t+h}$ be the marginal predictive density for the variables of unit i after integrating the remaining p variables and let $\mathcal{K}(S^1_{it+h} \mid F_t) = \int \ldots \int f(S^1_{it+h} \ldots S^G_{it+h} \mid F_t) dS^2_{it+h} \ \ldots dS^G_{it+h}$ be the marginal predictive density for inflation, which we order first in the list for unit i.

For our exercise we set $h = 1, 4, 8$. That is, we compute turning point predictions one, four and eight quarter ahead, using the information available at each t. To compute the probability of a turning point, say for $h = 1$, we have to calculate S^1_{it+1}. Given the marginal predictive density $\mathcal{K}$, the probability of a downturn in unit i is

$$\begin{aligned} P_{Dt} = Pr(S^1_{it+1} < S^1_{it} | S^1_{it-2}, S^1_{it-1} < S^1_{it}, F_t) &= \\ \int_{-\infty}^{S^1_{\text{it}}} \mathcal{K}\left(S^1_{it+1} \mid S^1_{it-2}, S^1_{it-1}, S^1_{it}, F_t\right) dS^1_{it} \end{aligned} \tag{13}$$

and the probability of an upturn is

$$\begin{aligned} P_{Ut} = Pr(S^1_{it+1} > S^1_{it} | S^1_{it-2}, S^1_{it-1} > S^1_{it}, F_t) &= \\ \int_{S^1_{\text{it}}}^{\infty} \mathcal{K}\left(S^1_{it+1} \mid S^1_{it-2}, S^1_{it-1}, S^1_{it}, F_t\right) dS^1_{it} \end{aligned} \tag{14}$$

Using a numerical sample from the predictive density satisfying $S^1_{it-2},\ S^1_{it-1} < S^1_{it}$, we can approximate these probabilities using the frequencies of realizations which are less then or greater then S_{it}. With a symmetric loss function, minimization of the expected loss leads us to predict the occurrence of turning point at $t + 1$ if $P_{Dt} > 0.5$ or $P_{Ut} > 0.5$.

Because some methods are univariate and other multivariate, the comparison country by country may not be adequate. Therefore, we report turning point prediction for the seven countries as a whole in each procedure. That is, for the 20 periods in the forecasting sample we compute forecasts and check whether at each date the forecasts satisfy the definition for each country. In table 4 we present the number of correct signals found.

4 The Results

Before describing the results in detail, we briefly discuss the time features of inflation rate across countries in the estimation and forecasting samples, as this may help to better understand the results we obtain. Figure 1 plots the inflation series and its correlation function while table 1 provides some statistics in the two samples.

Several features are worth commenting upon. First, in the estimation sample inflation tends to be persistent and up to five elements of the autocorrelation function are significant in the US, France, Italy and Canada. In the UK and Germany there is much less persistence and some residual seasonality is evident. Finally, in Japan the inflation rate has a clear and persistent two quarter pattern. Second, in the forecasting sample some of these tendencies are altered: for example, there is much less persistence in the first four countries; in Germany seasonality disappears and a two period pattern is becomes important; a two period pattern is added to seasonal variations in the UK and in Japan a very strong medium run effect is added to the two period pattern. Third, while in the estimation sample six of the seven inflation rates are skewed (the exception is Germany) and leptokurtosis is present in the US, Japan and the UK, in the forecasting sample inflation rates are approximately normally distributed with the exception of Japan. Fourth, the mean and both measures of dispersions are reduced in the forecasting sample. Fifth, there is a substantial drop in the typical value of the contemporaneous correlation of inflations across countries in the forecasting sample: while the majority of the correlations were significant in the first period, less than half of them are significant in the second. This is noticeable for US and Canadian inflation, which seem to have a peculiar idiosyncratic behavior. The drop is partially compensated by an increase in the first lagged correlation.

Overall, there appears to be instabilities in both the autocorrelation and the distribution of inflation rates for several countries in the two samples. While these differences may be due to the smaller size of the forecasting sample, they are suggestive of the problems that models may face in forecasting inflation over time.

AR models improve the MSE of the naive no-change model in six of the seven countries at all horizons, while for Italy results are mixed (see first column of Table 3). The gains are of the order of 10-25% and roughly uniform across horizons. Bootstrapping the forecasting residuals 1000 times, we find that the probability that the Theil-U at the one step horizon is greater or equal to one is negligible for all countries except for Italy. At longer horizons the results are less clear cut, and except for the US, Japan and Canada the value of one falls within the estimated one standard error band for the statistics at that horizon. ARIMA specifications produce uniformly smaller MSEs than those of the AR model in Japan (see second column in table 3). In the UK and Italy results depend on the horizon, and for the other four countries parsimonious ARIMA models do not capture well dynamics of the data. This feature is clearly confirmed in average and median MSE measures, which are uniformly larger for the ARIMA than for the AR model.

In terms of turning point predictions both models are surprisingly accurate (see the second and third rows of table 4). On average, between 65-75% of the turning points are recognized by both models with the ARIMA model doing significantly better at one step horizon and the AR model being more accurate at the four step horizon. At the eight step horizons, the two models have similar performance.

4.1 Should theory based models be used?

The answer is "it depends". There are countries where several theory-based bivariate models produce lower MSE which are smaller than univariate specifications, but there are also countries where none of the bivariate models improves over univariate approaches or, worse, the naive no-change model (see third column of table 3). For example, in the UK (Japan) at the one quarter horizon, six (four) of the nine bivariate models produce MSE which are smaller than those of univariate models while for the US none of the bivariate models beat a AR model. Results also depend on the horizon. The best model for the UK produces

MSE which is 10% lower at the one and four quarters horizons but it is equivalent to the ARIMA model at the eight step. On the other hand, in Japan, Italy, France and Canada the best bivariate models are superior to the best univariate specification at long horizons by 10-30%. On average or in the median the outcomes confirm to be far from uniform: AR models are superior to the best bivariate models at one step horizon, while at longer horizons the performance of bivariate models is little more encouraging.

Stock and Watson (2000) argued that when a bivariate model has an edge over other specifications in a sample, typically performs poorly in others, suggesting that forecasting relationship are tenuous. This lack of robustness is also partially confirmed here. In figure 2 we plot the MSE of each model relative to the MSE of the best univariate model for each of the three horizons. On the vertical axis we report the relative MSE obtained when data up to 1995 is used to estimate the models and data from 1996 used to forecast. On the horizontal axis we plot the relative MSE when data up to 1990 is use to estimate the models and data from 1991 to 1995 is used to forecast. Two main features can be noted. First, there are cases when the relative MSE ordering of the models change across the two subsamples (see e.g. the case of France and Italy at the one quarter horizon). Second, there are instances when a bivariate model is relatively good in a sample and relatively bad in the other (see e.g. the case of UK and US at the one quarter horizon). However, there are cases when relative MSEs line up approximately on the 45 degrees line (see e.g. Japan at one quarter horizon, Canada and Germany at eight quarter horizon). Hence, the instability of the forecasting relationship may depend on the country, the sample and the forecasting horizon.

The collection of the best bivariate models produce slight improvements in upturns recognitions at eight quarter horizon relative to AR models. However, the performance at shorter horizons is for all purposes equivalent. Since these results are obtained employing the best specification in each country, it should be clear that using the same specification for all countries would do no better.

Are Phillips curves useful for forecasting inflation? The evidence is mixed. Among the models we tried, we found that a specification which uses a traditional Phillips curve is preferred in France, Canada and Italy, while in the US, Germany and Japan pairing output growth or the output gap with inflation gives the best results. In the UK, traditional Phillips curve are systematically inferior at all horizons in terms of MSE to models which use financial market indicators to forecast inflation. From figure 2 we see that the best specification for all countries is the same in the two forecasting samples at the one step horizon except in the US (where a model which uses the slope of the term structure is preferable in the 1990-95 sample). However, except for Japan, none of the Phillips curve models produce MSEs which are smaller than those of univariate models and turning point recognition is no more accurate. Hence, we confirm the results of Atkeson and Ohanian (2001): Phillips curve models do not have an hedge over univariate models in forecasting inflation.

Furher and Moore (1995) indicated that the New-Keynesian Phillips curves have hard time to account for the in-sample dynamics of inflation. Gali and Gertler (1999) breathed a new life into the specification by pointing out that the theory has implication for the relationship between real marginal cost and inflation. Using ancillary assumptions to proxy for the unobservable marginal cost they concluded that inflation dynamics in the US are well approximated by a model with sticky prices (a-la Calvo) with a 5 quarters average period between price changes. Does real marginal cost help in forecasting inflation? The answer appears to be negative. Using the income labor share (as in Gali and Gertler) to proxy for these costs we find that in only 3 county such a bivariate model improves over univariate specifications, but it no case it is the preferred specification. Apparently, the predictive power of real marginal costs is negligible whenever lags of inflation are present.

Is money a leading indicator for inflation? In none of the G-7 countries money growth measures turn out to be the most useful variable to predict inflation when lags of inflation are included in the specification at short horizons. As a matter of fact, a specification which use money growth is the worst in six of the seven countries at one quarter horizon. However, at the longer horizons, specifications which use money growth or real balances growth are competitive with the best in five of the seven countries.

Three conclusions can be drawn from the above discussion. First, theoretically based bivariate models are not much better than atheoretical univariate models both in terms of MSE and turning point recognition. Superiority of one specification may be due to the sample used and to the horizon considered (see also Cecchetti, Chu and Steindel (2000)). Phillips curve relationships (both traditional and New-keynesian) are relative stable in the 90's but not very useful to forecast inflation across countries. Second, money has forecasting information for inflation only at long horizons. However, even at these horizons there are other real or financial indicators which are as good as money. Third, there is no unique model which is useful to forecast inflation in all G-7 countries and for both periods considered.

4.2 Is there information in other domestic variables?

Having examined the ability of variables indicated by theory, we next conducted a search over other variables, whose relationship with inflation has not been necessarily highlighted by any theory, to examine whether the forecasting performance of univariate models can be improved upon. The search was marginally successful: as shown in table 3, column 4, in Italy and Japan the performance of univariate models can be significantly improved at all horizons, while for Canada MSE improvements are visible at long horizons. These tendencies are not general: in fact, both median and mean measures are worse than those obtained with univariate models at all horizons. The BVAR, weakly restricting the statistical contribution of other variables to inflation forecasts, appears to be generically better than an unrestricted VAR (see column 5), and it is at medium-long horizons that the Bayesian prior helps most (the exception is the US). The record of turning point recognition of both the VAR and the BVAR model at one and eight steps horizons is practically identical to that of univariate AR models. However, at four step horizon, the BVAR model is significantly better than univariate models and four extra turning points are recognized (see row 5 in figure 4).

In sum, the addition of variables to univariate/bivariate specifications reduces the mean square error in a few instances but in others it introduces noise which increases the variability of the forecasts and clouds the ability of the model to track actual turning points of inflation. Statistically searching for good predictors of inflation is a worthwhile exercise when relationships are stable and the informational content of predictors robust. However, given the general characteristics of the inflation series we have highlighted, such a search may be doomed to failure. In this situation, atheoretical models which are flexible and adapt quickly to a changing environment and/or to structural breaks, may be more efficient forecasting tools than either theory based or statistically based fixed coefficient models.

4.3 Do time variations in coefficients help?

The instabilities in the distribution, autocorrelation and cross-correlation properties of inflation suggest that there are potential gains from using time varying coefficient models. The question of interest is whether these are gains quantitatively important. Table 3, column 6 indicates that for the AR model, MSE improvements are limited in size and not uniform across countries. A TVC-AR model produces better MSE than a fixed coefficient AR model for the UK and Italy at all three horizons and for France at the one quarter horizon, while for the US, Germany and Canada the addition of time varying coefficients does not help. This mixed performance is confirmed by cross sectional measures of relative MSE: a TVC is slightly better than a fixed coefficient models in terms of means but worse in terms of medians. The presence of time variations improves somewhat turning point predictions: at one step horizon four new turning points are recognized, and at the eight quarter horizon two new downturns are recorded.

The picture is different when a BVAR model is used (table 3, column 7): relative to a BVAR with fixed coefficients the MSE is significantly lower at the one and four step horizons in six of the seven countries, while at eight steps superior MSE performances is detectable in US, Germany, UK and France. The gain are quantitatively significant: for example, in the US, the average gain at the three horizons is larger than 10%. The exception here is Japan, but this is not surprising since the inflation rate in Japan does not display dramatic changes in the two subsamples. There are also improvements relative to TVC-AR models but they occur only in the US, Japan, and, at some horizons, in Germany, France and Canada. Gains are also present in mean and median MSE measures: a TVC-BVAR model gives the lowest average MSE at all three horizons among all the specifications we have considered so far. Relatively speaking, the improvement over fixed coefficient BVAR is of more than 10% and relative to a TVC-AR it varies from zero to 15%.

Time variations appear to be important in turning points detection (see rows 7 and 8 in table 4): the TVC-BVAR model recognizes essentially the same number of turning points as a fixed coefficient BVAR at one and four quarters horizons but picks up four more upturns at the eight quarter horizon. The performance relative to the TVC-AR model is less impressive at all three horizons.

The amount of time variation our procedure selects is small and only for the UK the best univariate model requires significant variations in the coefficients. These results should be contrasted with those of Stock and Watson (1996), who found that time variation does not help in improving forecasts of a multitude of US time series. There are two reasons for the differences. First, their exercise considers only the US but their sample is much longer. If there are long cycles in the data, what shows up as time variation in a sample of twenty years could be captured with mean reversion over longer periods of time. Second, our results indicate that a specification where coefficients evolve according to a random walk and the error has a small variance is appropriate in almost all G-7 countries. However, small deviations in value chosen for the variance produce dramatic changes in the forecasting performances of both models. Hence, a second reason for the differences may be due to larger and/or suboptimal amounts of time variations employed by these authors in their study.

We have also studied whether general time variation patterns where the variance of the coefficients is also allowed to evolve in an heteroskedastic fashion help in forecasting. Since there are countries where inflation displays leptokurtic behavior and volatility clustering appears to be pervasive, modelling heteroskedasticity as in Canova (1993b) has the potential to improve the forecasting performance in some of the countries. It turns out that this is not the case: the MSEs of models where parameters regulating heteroskedasticity in the coefficients are chosen to maximize the predictive density are practically identical to those reported in table 3. In fact, the estimated parameters imply that homoskedasticity is nearly optimal in six of the seven countries. Turning point detection is largely unchanged: the presence of heteroskedasticity only allows the recognition of one more upturn at the four step horizon at the cost of inducing a missing downturn signal at the same horizon.

To conclude, time variations in the coefficients do not significantly help in forecasting inflation when univariate AR models are employed. However, a tiny amount of time variation substantially improves the performance of VARs and largely ameliorate turning point detection. Once again, gains are more sizable at one and two years horizons.

4.4 Is information in the cross section?

Historically, inflation rates significantly comoved across G-7 countries. This was true of the high inflation decade of the 70's, of the subsequent deflationary period of the 80's and of the more stable period with declining trends experienced in the beginning of the 90's. Our estimates of the correlation function over the estimation period confirm this conventional wisdom: whatever helped, say, to forecast inflation in the UK should, in principle, have helped to forecast inflation also in France. This commonality seemed to have changed in the forecasting period and, except in a few cases, movements in inflation rates appear to be driven by idiosyncratic and national factors. Hence, while there seems to be room to improve the in-sample performance by adding international variables to domestic specifications, it is open to question whether cross country information will help in forecasting out of sample. Unless the model selected captures the time varying nature of the contemporaneous correlation of inflation rates across countries, it is unlikely that a specification with cross country interdependencies will improve the results we have so far obtained. Because the TVC-BVAR model is so far the best in the forecasting race in six of the seven countries, we will build cross country interdependencies over this basic structure. In Japan, a standard BVAR model is so far the best and that framework will be used as a starting point.

As mentioned the first two specifications we consider come short of fully modelling interdependencies across countries and variables and differ in the way international information is accounted for. The first specification, which treats international factors as predetermined in each of the seven domestic specification, represents a rough short cut in two ways: first, it assumes that cross sectional information can be summarized with a small number of observable variables. Second, it assumes that each G-7 country can be modelled as a small open economy, which takes international information as independent of the domestic one, an assumption which has somewhat less appealing content. In the second specification the seven inflation rates are directly interrelated with each other but only a small number of (hopefully common) factors captures the main features of domestic interdependencies.

Ideally, one would like to estimate a model were both national and international interdependencies are fully accounted for. However, there are computation costs; degrees of freedom limitations are important, given the short sample, and substantial noise may be added to the equations producing no clear forecasting value. Also, a fully fledged panel VAR requires that the same variables are used in each country. Hence, unless the dimensionality of the model is large, it maybe impossible to accommodate the idiosyncrasies of the national inflation rates we have found so far.

The third specification allows, in principle, for a rich set of interdependencies across variables in all country, but a-priori collapses complicated feedbacks into a index structure with time varying coefficients. The index structure solves some of the problems: it averages out noise; it limits the dimensionality of the parameter vector; and it allows the estimation of a large scale model. However, it is important to remember, that the index structure forces substantial similarities across countries and may be suboptimal when significant heterogeneities are present.

The first specification (BVAR-I in the tables) produces uniformly superior forecasting results in several countries. For example, in Canada MSE gains are on average of the order of 5-8%. In Japan and Germany, the introduction of foreign factors produces different results at different horizons. The only country where the presence of international factors makes the MSE of the forecasts uniformly worse is the UK and this is somewhat surprising given that it fits the prototype of small open economy that this model is trying to capture. Because of the improvements obtained in a number of countries, both average MSE measures decrease at all horizons: the BVAR-I model improves on average the MSE of a domestic

TVC-BVAR at one step horizon by 3% and at long horizon by even more. The fact that international information helps in medium run forecasts of inflation is not surprising: while short term inflation movements tend to be dominated by national factors, it is international conditions which play the largest role in determining inflation dynamics in the medium run. Adding international factors is important also in turning point detection: the BVAR-I model produces the best pattern of turning point recognition of all models at one quarter horizon.

The performance of the second specification (7-BVAR in the tables) is somewhat less appealing. The MSE of the forecasts is similar across countries and only for the UK and Canada improvements over the BVAR-I specification at all horizons are visible. Since these are two small open economies with similar trade links, it is comforting to find that a model where inflation interdependencies are fully modelled is appropriate. For the other countries one quarter ahead horizon are not very encouraging. At long horizons modelling inflation interdependencies directly is beneficial for UK and Germany - the countries with the worse eight step ahead performance in the BVAR-I model- and the US and Canada. Improvements are also visible in average terms at the two years horizon but the median MSE is somewhat worse than the one obtained with the BVAR-I model. The record of turning point predictions is mixed: there are slight improvements over a model without international interdependencies but the model is somewhat worse than a BVAR-I at the one step horizon.

The performance of index specification is encouraging. For Japan, Germany, UK and France substantial MSE improvements are evident at some or all horizons and MSE reductions of 10% are common. The major improvements obtain for Japan and Germany, which are the countries where previous international specifications failed. For Italy and Canada the performance is mixed while for the US the MSE greatly increases. Apparently, averaging the information contained in lagged G-7 variables has little forecasting use in the US, where the evolution of inflation has substantial idiosyncratic components in the forecasting sample.

The cross section provides important information for forecasting medium run trends of inflation: in fact, in all countries the MSE at the eight quarter horizon is lower than the MSE at the one quarter horizon. The polarization in the MSE performance is evident in the large discrepancy of median and average measures: the median MSE is lower than with all other specifications, while the mean is larger. These improvements are not reflected in more accurate turning point predictions. This result however is entirely due to the poor forecasting performance of the model for US inflation (see figure 3 where the actual values of inflation and one year ahead forecasts obtained with the index model are plotted). Inflation forecast are off both in terms of size and timing of peaks and troughs for the US but they tract actual inflation reasonably well in the other countries.

Figure 3 further indicates that the performance of the index model is robust to changes in the forecasting sample and its ability to track inflation in the G-7 is similar in the first and second half of the 90's. This is not necessarily the case, e.g., of the BVAR-I model: had we used the first part of the 90's as forecasting sample, the MSE of the model would have been worse in some cases by up to 5%.

To summarize, adding international interdependencies to domestic models helps to forecast national inflation rates, primarily one or two years ahead. However, there is not one superior way to take into account these interdependencies. In the US and Italy, controlling for international demand and supply factors appears to be sufficient; in Canada modelling the interdependencies of inflation rates is more appropriate; in the other four countries averaging national information provides a more clear forecasting signal. Finally, while the forecasting power of several is sample dependent, the performance of panel index specification is consistent throughout the 90's.

5 Conclusions

This paper addressed four interrelated questions related with the problem of forecasting inflation rates in the G-7. First, we studied how models which rely on the existence of a Phillips curve relationship fare relative to univariate models or to bivariate models which use other variables suggested by economic theory. Second, we investigated whether there is information in other domestic variables (which may or may not have theoretical underpinning) useful for forecasting inflation. Third, we examined whether models which allow the parameters to drift over time improve the forecasting performance of specifications with fixed coefficients. Finally, we evaluated whether there are forecasting gains when information coming from the cross section of the G-7 countries is used to forecast domestic inflation.

We show that among the bivariate specifications, models employing some version of traditional Phillips curve are preferred in terms of MSE in three countries, models which use output measures perform better in three countries and in the UK the best bivariate model includes inflation and stock returns. We show that there are instances when bivariate models produce inflation forecasts with a smaller MSE than the one of univariate models. However, there are also cases when the opposite is true. Furthermore, in those countries where a Phillips curve model was selected the MSE of the forecasts is similar to the one of univariate models, confirming the conclusions of Atkeson and Ohanian (2001). As in Stock and Watson (2000) we find that conclusions concerning the relative ordering of MSEs are somewhat tenuous: changing the forecasting sample at times alters the ranking of bivariate models and their MSE performance relative to univariate models. In terms of turning point predictions, the performance of the collection of the best bivariate models is no better than the one of univariate specifications, regardless of the horizon and the sample we considered.

The MSE of the forecasts of univariate models can be somewhat improved selecting variables on the basis of their in-sample predictive power. However, gains are not generalized across countries and obtain only at longer horizons. Adding a Bayesian prior ameliorates the performance of a VAR for six of the seven countries but no systematic MSE improvements over univariate models are noticeable. Turning point predictions are similar to those obtained with univariate models except at the one year horizon where BVAR recognizes a significantly larger number of upturns and downturns.

The addition of time variation does not change significantly the MSE of the forecasts of univariate AR models except in a few countries. However, turning point recognition improves, primarily at the one quarter horizon. Adding time variation to the BVAR model, on the other hand, significantly reduces the MSE of the forecasts in several countries and gains are visible both over univariate TVC specifications and over fixed coefficient BVAR models. The same improvement, although less systematic, is observable in turning point detection: gains are more significant at one and two years horizons and for upturns.

Adding international variables to domestic models has some use in forecasting. With the simplest specification MSE reductions are obtained in five countries and gains of the order of 30% are not unusual. Turning point recognition also improves and this model produced the smallest number of missing signals of all the specifications at the one quarter horizon. The model which directly captures inflation interdependencies produces MSE improvement in some countries, primarily at longer horizons but, overall, the performance of this specification seems to be inferior to the previous one. Finally, the index model produces the best overall MSE of the forecasts at all horizons for Germany, Japan and France and the UK. MSE gains produced are more pronounced at medium-long horizons. The gains are not matched by more successful turning point predictions because of the poor performance of the model for US inflation. In general, the forecasts produced by this model track well actual inflation throughout the 1990's using the information available one or even two years ahead, with no visible change in the performance across the two halves.

We conclude that the conventional wisdom that the simplest and the most naïve models are best to forecast inflation rates is not necessarily true. Phillips curves or other small scale domestic models with fixed coefficients are not much better than univariate AR models. However, multivariate models with time varying coefficients and some international linkages are superior to univariate models both in terms of MSE of the forecasts and turning point recognition at several horizons. Models with these features are useful for medium term forecasts of inflation because they exploit information about the time varying (local) common trend present in almost all G-7 countries. Because of their flexibility and the nature of the inflation process in the G-7, time varying models with international interdependencies appear to be robust forecasting tools throughout the 90's.

One may be curious to know how these models would fare in next few years and whether they will be able to capture the changes in the pattern of inflation rates that starts creeping up in the last few quarters. Given their characteristics we expect them to be able to adapt easily to changing environments and to be able to employ more cross-country information as international interdependencies increase.

There are few open avenues of research that this paper leaves untouched. For example none of the models considered allow different variables to have different forecasting power for inflation at different point in time and general regime switching specifications are useful features one may want to include in a successful model.

Finally, from the academics and policymakers point of view, it may be interesting to conduct not only unconditional but also conditional forecasts. The results obtained here, together with the machinery developed in Wagonner and Zha (1999), make the task feasible even in complicated hierarchical Bayesian panel VAR models.

Appendix A: The Data

The data used comes from the OECD data base. Inflation rates are calculated annualizing quarterly rate of growth of the consumer price index. We have also considered, as alternative measure, inflation rates computed using GDP deflator and we have verified that both the qualitative and the quantitative conclusions we have reached are robust. Furthermore we have also computed annual growth rates by taking the rate of change in the consumer price level over four quarters. Although some of the quantitative conclusions we have reached are altered by this alternative measure, the results are also robust to this change. The rate of growth of money and of output are similarly computed. We use in all cases a broad definition of money: M2 is employed in all countries but France where M3 is used. Note that money supply data for European countries are available only up to 1998:4. Therefore the forecasting statistics produced by bivariate models where money growth is used cover only the years 1996:1 to 1998:4. The output gap is computed using an exponential smoothing algorithm. We choose this specification, as opposed to a more standard Hodrick and Prescott (HP) filter, because the time series for the gap retains the timing of information of the original series, a feature not produced by the HP filter because of its two-sided nature. The slope of the term structure is calculated by taking the difference the annualized long term interest rate (of maturities varying from 5 to 10 years, depending on the country) and the annualized short term one (typically a three month rate). The unemployment rate measures are standardized unemployment figures as provided by the OECD. The employment growth series measures the number of workers seeking jobs which have found a job in the quarter under consideration, annualized as the other variables. The rent price index measures the costs of housing and its growth rate is computed in a similar fashion. The labor share is computed using wages divided by labor productivity. All the other variables we have tried (unit labor costs, capacity utilization, employment, wage rate, nominal exchange rates, real effective exchange rates in levels and growth rates) are the standard ones provided by the OECD data bank. Stock prices are obtained from the BIS data bank. We use in the exercise SP500 index (US), the Nikkei 225 index (Japan), the Toronto 300 composite index (Canada), the DAX composite index (Germany), the CAC 40 index (France), the FT 100 (UK) and the MIB index (Italy).

Appendix B: Gibbs sampling densities

The prior has the following hierarchical structure:

$$\begin{aligned} \alpha &= M_o\mu + v, & v &\sim N(0, B_o) \\ \beta_t &= M_1\theta_t + u_t, & u_t &\sim N(0, B_1) \\ \theta_t &= (I - C)\,\theta_o + C\theta_{t-1} + \eta_t & \eta_t &\sim N(0, B_{2t}) \end{aligned}$$

where B_o is $(N \times N)$, $B_1 = diag\,(b_1 \;\; b_2)$ is block diagonal (b_1 of dimensions 1×1, and b_2 of dimensions $G \times G$), M_o is a matrix of dimension $N \times m$, with $m \leq N$, M_1 is a matrix of dimensions $(1 + G) \times s$, where $s \leq 1+G$, the vector $\theta_t = (\theta_t^1, \;\; \theta_t^2)'$ is of dimensions $s \times 1$, and C is $s \times s$. The matrix B_{2t} follows the process $B_{2t} = \delta_t * B_{20}$ with $\delta_t = \nu_1^t + \nu_2\,(1 - \nu_1^t)\,/\,(1 - \nu_1)$.

To complete the prior information we specify a prior for $(\mu, \Omega, B_o, B_1, B_{20})$. The following is assumed:

$$p\left(\mu, \Omega^{-1}, B_o^{-1}, B_1^{-1}, B_{20}^{-1}\right) = p(\mu)\, p\left(\Omega^{-1}\right) p\left(B_o^{-1}\right) p\left(B_1^{-1}\right) p\left(B_{20}^{-1}\right)$$

with

$$\begin{aligned} p(\mu) &= N(\bar{\mu}, \Sigma_\mu) \\ p\left(\Omega^{-1}\right) &= W(\omega_1, Q_1) \\ p\left(B_o^{-1}\right) &= W(\omega_2, Q_2) \end{aligned}$$

$B_1 = diag\,(b_1, b_2)$, where

$$\begin{aligned} p(b_1) &= IG(v_0, \delta_0) \\ p\left(b_2^{-1}\right) &= W(\omega_3, Q_3) \end{aligned}$$

and $B_{20} = diag\,(\phi_1, \Phi_2)$ where ϕ_1 is of dimensions 1×1 and $\Phi_2 = \phi_2 * I_s$. We assume a diffuse prior on both ϕ_1 and ϕ_2. Here W indicates the Wishart distribution and IG the inverted gamma distribution. We assume that $\chi = [\nu_1, \nu_2, \mu, vec(\Sigma_\mu), \omega_1, \omega_2, \omega_3, vec(Q_1, Q_2, Q_3)]$ are known.

Posterior densities of the parameters of interest are obtained by combining the likelihood of the data, which is proportional to

$$\propto |\Omega|^{-T/2} \exp\left\{-\frac{1}{2}\sum_{t=1}^{T}(y_t - \mathbf{R}_t\beta_t - \mathbf{A}_t\alpha)'\,\Omega^{-1}\,(y_t - \mathbf{R}_t\beta_t - \mathbf{A}_t\alpha)\right\}.$$

with the prior distributions above. Letting $Y_T = (y_1, ..., y_T)$ denote the sample data and $\psi = (\{\beta_t\}_t\,, vec(\Omega), \{\theta_t\}_t\,, \alpha, \mu, vec(B_o),\, b_1, vec(b_2), \phi_1,\, \phi_2)$ denote the parameters whose joint distribution needs to be found, and $\psi_{-\zeta}$ the vector ψ excluding the parameter ζ, it is easy to derive the following conditional posteriors:

$$\beta_t \mid Y_T, \psi_{-\beta_t} \sim N\left(\hat{\beta}_t, V_{1t}\right), \quad t \leq T;$$

$$\Omega^{-1} \mid Y_T, \psi_{-\Omega} \sim W\left(\omega_1 + T, Q_{1T}\right);$$
$$\alpha \mid Y_T, \psi_{-\alpha} \sim N\left(\hat{\alpha}, V_2\right);$$
$$\mu \mid Y_T, \psi_{-\mu} \sim N\left(\hat{\mu}, V_3\right);$$
$$B_o^{-1} \mid Y_T, \psi_{-B_o} \sim W\left(\omega_2 + 1, Q_{2T}\right);$$
$$b_1 \mid Y_T, \psi_{-b_1} \sim IG\left(\frac{(v_o + T)}{2}, \frac{\delta_o + \sum_t \left(\lambda_t - \theta_t^1\right)' \left(\lambda_t - \theta_t^1\right)}{2}\right)$$
$$b_2^{-1} \mid Y_T, \psi_{-b_2} \sim W\left(\omega_3 + T, Q_{3T}\right);$$
$$\phi_1 \mid Y_T, \psi_{-\phi_1} \sim IG\left(\frac{T}{2}, \frac{\sum_t \left(\theta_t^1 - \theta_{t-1}^1\right)' \left(\theta_t^1 - \theta_{t-1}^1\right)}{2}\right);$$

$$\phi_2 \mid Y_T, \psi_{-\phi_2} \sim IG\left(\frac{T\left(s-1\right)}{2}, \frac{\sum_t \left(\theta_t^2 - \theta_{t-1}^2\right)' \left(\theta_t^2 - \theta_{t-1}^2\right)}{2}\right)$$

where $\hat{\beta}_t = V_{1t}\left(B_1^{-1} M_1 \theta_t + \mathbf{R}_t' \Omega^{-1}\left(y_t - A_t \alpha\right)\right)$; $V_{1t} = \left(B_1^{-1} + \mathbf{R}_t' \Omega^{-1} \mathbf{R}_t\right)^{-1}$;
$Q_{1T} = \left[Q_1^{-1} + \sum_t \left(y_t - \mathbf{R}_t \beta_t - \mathbf{A}_t \alpha\right)\left(y_t - \mathbf{R}_t \beta_t - \mathbf{A}_t \alpha\right)'\right]^{-1}$; $\hat{\alpha} = V_2\left(B_o^{-1} M_o \mu + \sum_t A_t' \Omega^{-1}\left(y_t - \mathbf{R}_t \beta_t\right)\right)$;
$V_2 = \left(B_o^{-1} + \sum_t A_t' \Omega^{-1} A_t\right)^{-1}$; $\hat{\mu} = V_3\left(M_o' B_o^{-1} \alpha\right)$; $V_3 = \left(M_o' B_o^{-1} M_o\right)^{-1}$;
$Q_{2T} = \left[Q_2^{-1} + \left(\alpha - M_o \mu\right)\left(\alpha - M_o \mu\right)'\right]^{-1}$; $Q_{3T} = \left[Q_3^{-1} + \sum_t \left(\rho_t - M_{1\_22} \theta_t^2\right)\left(\rho_t - M_{1\_22} \theta_t^2\right)'\right]^{-1}$
where M_{1_22} is the diagonal block of M_1 controlling the tightness of the vector θ_t^2.

The Gibbs sampling iterates across these conditional distributions until convergence is achieved.

Assuming $C = I$, the joint conditional posterior distribution of $\theta_0, \theta_1, ..., \theta_T \mid Y_T, \psi_{-\theta_t}$, is obtained in two steps as in Chib and Greenberg (1995). First, obtain $\{\theta_t\}_t$ for each t recursively with the Kalman filter, i.e. :

$$\begin{aligned} \hat{\theta}_{t|t} &= \hat{\theta}_{t|t-1} + K_t\left(\beta_t - M_1 \hat{\theta}_{t|t-1}\right) \\ R_{t|t} &= \left(I - K_t M_1\right) R_{t|t-1} \\ K_t &= R_{t|t-1} M_1 F_{t|t-1}^{-1} \\ F_{t|t-1} &= M_1 R_{t|t-1} M_1' + B_1 \end{aligned} \tag{15}$$

where $\hat{\theta}_{t|t-1} = \hat{\theta}_{t-1|t-1}$ and $R_{t|t-1} = R_{t-1|t-1} + \delta_t B_{20}$. Second, the joint conditional posterior distribution $\theta_0, \theta_1, ..., \theta_T \mid Y_T, \psi_{-\theta_t}$ is sampled in reverse time order from

$$\begin{aligned} \theta_T &\sim N\left(\hat{\theta}_{T|T}, R_{T|T}\right) \\ \theta_{T-1} &\sim N\left(\hat{\theta}_{T-1}, R_{T-1}\right) \\ &\vdots \\ \theta_0 &\sim N\left(\hat{\theta}_0, R_0\right) \end{aligned} \tag{16}$$

where $\hat{\theta}_t = \hat{\theta}_{t|t} + \mathbf{M}_t\left(\theta_{t+1} - \hat{\theta}_{t|t}\right)$, $R_t = R_{t|t} - \mathbf{M}_t R_{t+1|t}\mathbf{M}_t'$, and $\mathbf{M}_t = R_{t|t}R_{t+1|t}^{-1}$.

To make the updating scheme described in (15)-(16) operational, initial values for B_{20}, R_0, and the vector $\hat{\theta}_0$ at time $t = 1$ (the first period of the sample) must be assigned. Here we set $B_{20} = R_0$ arbitrarily diagonal and $\phi_1 = \phi_2 = 0.5$. $\hat{\theta}_0$ is initialized by running a VAR for each country and taking the constant.

In the same way, Q_1 is taken as the variance covariance matrix of a pooled VAR, and Ω is initialized equal to the same value. Q_2 is initialized by running a VAR for each variable across countries and averaging the variance covariance matrix, while Q_3 is initialized by running a VAR for each country and averaging the variance covariance matrix. Finally, we set $\Sigma_\mu^{-1} = 0$; $\nu_1 = 0$, $\nu_2 = 0.0001$; $v_0 = 6.$; $\delta_0 = 1$.

References

[1] Atkeson, A. and Ohanian, L. (2001), "Are Phillips Curves Useful for Forecasting Inflation?", Federal Reserve Bank of Minneapolis, Quarterly Review, 25 (1), 2-11.

[2] Blinder, A. (1997), " Is there a Core of Practical Macroeconomics we Should all Believe", American Economic Review, 87, 240-243.

[3] Canova, F. (1992), "An Alternative Approach to Modelling and Forecasting Seasonal Time Series", Journal of Business and Economic Statistics, 10, 97-108.

[4] Canova, F. (1993a) "Forecasting a Multitude of Time Series with Common Seasonal Pattern",Journal of Econometrics, 55, 173-202

[5] Canova, F. (1993b) "Modelling and Forecasting Exchange Rates using a Bayesian Time varying coefficient model", Journal of Economic Dynamics and Control, 17, 233-262.

[6] Canova, F. and Ciccarelli, M. (1999) "Forecasting and Turning Point Prediction in a Bayesian Panel VAR Model", UPF working paper 443, available at www.econ.upf.es.

[7] Canova, F. and Ciccarelli, M. (2001) "Bayesian leading indicators" in progress.

[8] Cecchetti, S., Chu, R. and Steindel, C. (2001), "The Unreliability of Inflation Indicators", Current Issues in Economics and Finance, Federal Reserve Bank of New York, 6(4), 1-6.

[9] Evans, C. and Marshall, D. (1998), "The Economic Determinants of the Term structure of Nominal Interest rates",Carnegie Rochester Conference Series in Public Policy, 49, 53-101.

[10] Fama, E. (1970) Efficient Capital markets: a review of theory and empirical work,Journal of Finance, 25, 383-417.

[11] Forni, M., Hennin,M. , Lippi, M, Reichlin, L (2000), "The Generalized Factor Model: Identification and Estimation", ULB manuscript.

[12] Furher, J. and Moore, G. (1995) "Inflation Persistence", Quarterly Journal of Economics, 110, 127-159

[13] Goodhart, C. and Hoffman, B.(2000) "Asset Prices and the Conduct of Monetary Policy", LSE manuscript.

[14] Gali, J. and Gertler, M. (1999) "Inflation Dynamics: A Structural Econometric Analysis", Journal of Monetary Economics, 44, 195-222.

[15] Holtz–Eakin D., W. Newey and H. Rosen (1988), Estimating vector autoregressions with panel data, Econometrica, 56(6) , 1371–1395.

[16] Ivanov, and Kilian, L (2000), "A Practitioner's Guide to Lag Order Selection for Vector Autoregressions", University of Michigan, manuscript.

[17] Lahari, K. and Moore, J. (1991) Leading Economic Indicators: New Approaches and Forecasting Records, Chicago, IL: University of Chicago Press.

[18] Mankiw, G. (2000), "The inesorable trade-off between inflation and unemployment", Economic Journal, 111, C45-60.

[19] Marcellino, M., Stock, J. and Watson. M. (2000) "Macroeconomic Forecasting in the Euro Area: Country Specific versus Area-wide Information ", forthcoming, European Economic Review.

[20] Plosser, C. and Rouwenshort (1994) , "International Term Structure and Real Economic Growth", Journal of Monetary Economics, 33, 133-156.

[21] Stock, J. and Watson. M (1989) "New Indexes of Coincident and Leading Economic Indicators" in Blanchard O. and Fisher, S. (eds) NBER Macroeconomic Annual, 352-394.

[22] Stock, J. and Watson. M (1996), "Evidence on the Structural Instability in Macroeconomic time Series Relationship", Journal of Business and Economic Statistics, 14, 11-30.

[23] Stock, J. and Watson. M (1999), "Forecasting Inflation", Journal of Monetary Economics, 44, 293-335.

[24] Stock, J. and Watson. M (2000) "Forecasting Output and Inflation: The Role of Asset prices", manuscript

[25] Wagonner, D. and Zha, T. (1999) "Conditional Forecasts in Dynamic Multivariate Models ", Review of Economics and Statistics, 81(4), 639-651 .

Table 1: Statistics of Inflation

Moment	US	Japan	Germany	UK	France	Italy	Canada
Sample 1980:1-1995:4							
mean	4.00	2.02	2.97	5.77	5.13	8.31	4.75
s.d.	2.97	2.98	2.55	5.02	4.07	5.15	3.55
lower 25%	2.90	-0.12	1.40	2.00	2.40	4.67	2.76
upper 25%	4.94	4.19	4.38	7.95	7.25	11.60	5.75
Skewness	1.69*	0.94*	0.56	1.42*	1.12*	1.31*	0.77*
Kurtosis	3.92*	1.36*	0.45	2.54*	0.08	1.14*	0.11
Sample 1996:1-2000:4							
mean	2.49	0.34	1.44	2.73	1.21	2.26	1.89
s.d.	0.91	2.69	1.53	2.30	1.17	1.05	1.21
lower 25%	1.81	-1.08	0.41	1.29	0.41	1.67	1.00
upper 25%	-3.25	1.12	2.37	3.99	2.05	2.75	2.71
Skewness	-0.20	1.38*	0.28	0.55	-0.37	0.29	0.27
Kurtosis	-0.84	3.53*	-0.48	0.34	-0.67	0.03	-0.60

Contemporaneous Correlations

	Japan	Germany	UK	France	Italy	Canada
Sample 1980:1-1995:4						
US	0.53	0.41	0.54	0.59	0.58	0.70
Germany		0.22	0.05	0.33	0.31	0.48
Japan			0.43	0.42	0.35	0.31
UK				0.44	0.31	0.57
France					0.76	0.76
Italy						0.72
Sample 1996:1-2000:4						
US	-0.02	0.17	0.00	0.32	0.33	0.45
Germany		0.31	-0.44	-0.48	-0.55	-0.39
Japan			0.47	0.14	0.30	-0.00
UK				-0.52	-0.32	-0.26
France					0.19	0.23
Italy						0.15

Notes: In the sample 1980-1995, correlations are significantly different from zero if greater than 0.25 in absolute value. In the sample 1996-2000 they are significant if greater than 0.43 in absolute value.

Table 2: Hyperparameter Estimates

Parameter	US	Japan	Germany	UK	France	Italy	Canada
BVAR							
θ_2	0.05	0.25	0.20	0.55	0.20	0.15	0.1
θ_3	-1.5	2.0	2.0	1.0	4.0	2.0	2.0
θ_4	1.0	0.5	1.0	0.5	0.01	1.0	0.05
Univariate-TVC							
θ_1	1.0	1.0	1.0	1.0	1.0	0.0	1.0
θ_2	0.02	0.122	0.062	0.062	0.062	0.022	0.062
θ_3	0.6	2.0	2.0	2.0	1.0	1.0	1.0
θ_5	10000	100	100	10000	10000	10000	10000
θ_6	1.0	1.0	1.0	1.0	1.0	1.0	1.0
θ_7	1.0e-6	1.0e-7	1.0e-7	3.6e-5	4.5e-5	1.46e-4	1.0e-7
BVAR-TVC							
θ_1	1.0	1.0	0.0	0.0	0.0	0.0	1.0
θ_2	0.1	0.15	0.05	0.2	0.1	0.1	0.05
θ_3	-1.5	3.0	2.0	0.0	-0.5	0.4	4.0
θ_4	0.02	0.05	0.5	0.02	0.07	0.30	0.05
θ_5	10	10	10	10	10	10	10
θ_6	0.97	1.0	1.0	1.0	1.0	0.98	0.98
θ_7	1.0e-7	0.0	1.0e-7	0.0	1.0e-6	1.0e-8	1.0e-6
BVAR-I							
θ_5	30	10	10	0.1	1.0	35	0.01

Notes: θ_1 controls the mean of the first (and fourth) lag; θ_2 the general tightness, θ_3 the harmonic decay of the variance of the lags, θ_4 the weight on variables other the dependent one, θ_5 the prior variance on the exogenous variables, θ_6 the decay toward the mean in the law of motion of the coefficients, θ_7 the time variations in the law of motion of the coefficients. The parameters for the BVAR-I are the same as those of a BVAR-TVC except for θ_5.

Table 3: Theil-U Statistics, Sample1996:1-2000:4

Country	Step	ARIMA	AR	Bivariate	VAR	BVAR	TVC-AR	TVC-BVAR	BVAR-I	7-BVAR	Index
US	1	1.21	0.82	1.12	1.20	0.86	0.92	0.77	0.64	0.82	1.46
	4	1.04	0.90	0.94	1.40	1.02	0.90	0.84	0.68	0.57	1.12
	8	0.94	0.77	0.84	1.36	1.02	0.80	0.73	0.58	0.50	0.96
JP	1	0.68	0.73	0.70	0.67	0.59	0.73	0.67	0.66	0.79	0.53
	4	0.72	0.80	0.61	0.58	0.56	0.80	0.68	0.67	0.83	0.53
	8	0.53	0.82	0.61	0.60	0.51	0.82	0.68	0.63	0.74	0.55
GE	1	0.83	0.71	0.76	1.05	0.85	0.79	0.78	0.82	0.84	0.76
	4	1.05	0.94	0.79	1.32	0.86	1.00	0.87	1.11	0.96	0.62
	8	1.03	0.95	0.89	1.23	1.01	1.05	0.87	1.54	0.99	0.64
UK	1	0.90	0.92	0.81	1.00	0.96	0.83	0.98	0.98	0.86	0.53
	4	1.01	0.90	0.87	1.05	0.97	0.92	1.00	1.05	0.88	0.49
	8	0.90	1.08	0.88	1.20	1.13	0.84	1.06	1.45	0.94	0.55
FR	1	0.92	0.82	0.95	0.86	0.93	0.76	0.82	0.79	0.83	0.79
	4	1.11	0.88	0.85	1.37	1.00	0.94	0.86	0.70	0.85	0.60
	8	1.25	0.86	0.78	1.35	1.16	0.93	0.77	0.59	0.76	0.60
IT	1	0.93	0.97	0.98	0.93	0.91	0.80	0.83	0.77	0.89	0.89
	4	1.04	1.17	1.10	1.47	1.09	1.07	0.99	0.56	1.19	0.71
	8	1.63	1.32	1.04	2.07	1.36	1.12	1.46	0.56	1.42	0.70
CA	1	1.15	0.83	0.87	1.78	0.95	0.85	0.93	0.87	0.72	1.01
	4	1.74	0.81	0.91	1.73	0.87	0.88	0.83	0.81	0.64	0.77
	8	2.09	0.81	0.75	1.84	0.77	0.89	0.83	0.73	0.59	0.77
Median	1	0.92	0.82	0.87	1.00	0.93	0.83	0.82	0.79	0.83	0.79
	4	1.04	0.90	0.87	1.37	0.97	0.92	0.86	0.70	0.85	0.62
	8	1.03	0.86	0.84	1.35	1.02	0.89	0.83	0.73	0.76	0.64
Mean	1	0.93	0.84	0.88	1.09	0.86	0.81	0.82	0.79	0.82	0.86
	4	1.09	0.90	0.87	1.27	0.91	1.09	0.87	0.79	0.84	0.69
	8	1.19	0.94	0.82	1.34	0.99	0.92	0.91	0.89	0.82	0.66

Notes: In the bivariate column we report results obtained with inflation and GDP growth for US and Germany; inflation and GDP gap for Japan; inflation and unemployment for France and Canada, inflation and employment growth for Italy; inflation and stock returns for the UK. BVAR-I is the BVAR with international variables; 7-BVAR is a model with inflation interdependencies; Index is a Panel VAR model with an index structure. The median and the mean refer to two different average measures of Theil-U across all countries.

Table 4: Turning point predictions, 1996:1-2000:4

	1-step		**4-steps**		**8-steps**	
Model	DT&NDT	UP &NUP	DT&NDT	UP &NUP	DT&NDT	UP &NUP
TRUE	47	42	37	38	30	25
Box-Jenkins	38	34	27	27	23	18
AR	37	33	31	27	20	20
Bivariate	36	32	29	26	22	19
VAR	36	32	27	24	21	17
BVAR	36	32	31	30	21	19
TVC-AR	37	36	28	30	20	22
TVC-BVAR	35	33	32	28	24	19
BVAR-I	40	36	33	31	25	21
7-BVAR	36	35	34	31	25	22
Index	32	32	21	29	21	20

Notes: DT&NDT indicates downturns and no-downturns, UT&NUT indicates upturns and no-upturns as defined in the paper. Under bivariate we report the number turning point recognized by the collection of the best bivariate models for each country. TRUE refers to the actual number of turning points.

Inflation and Correlation function

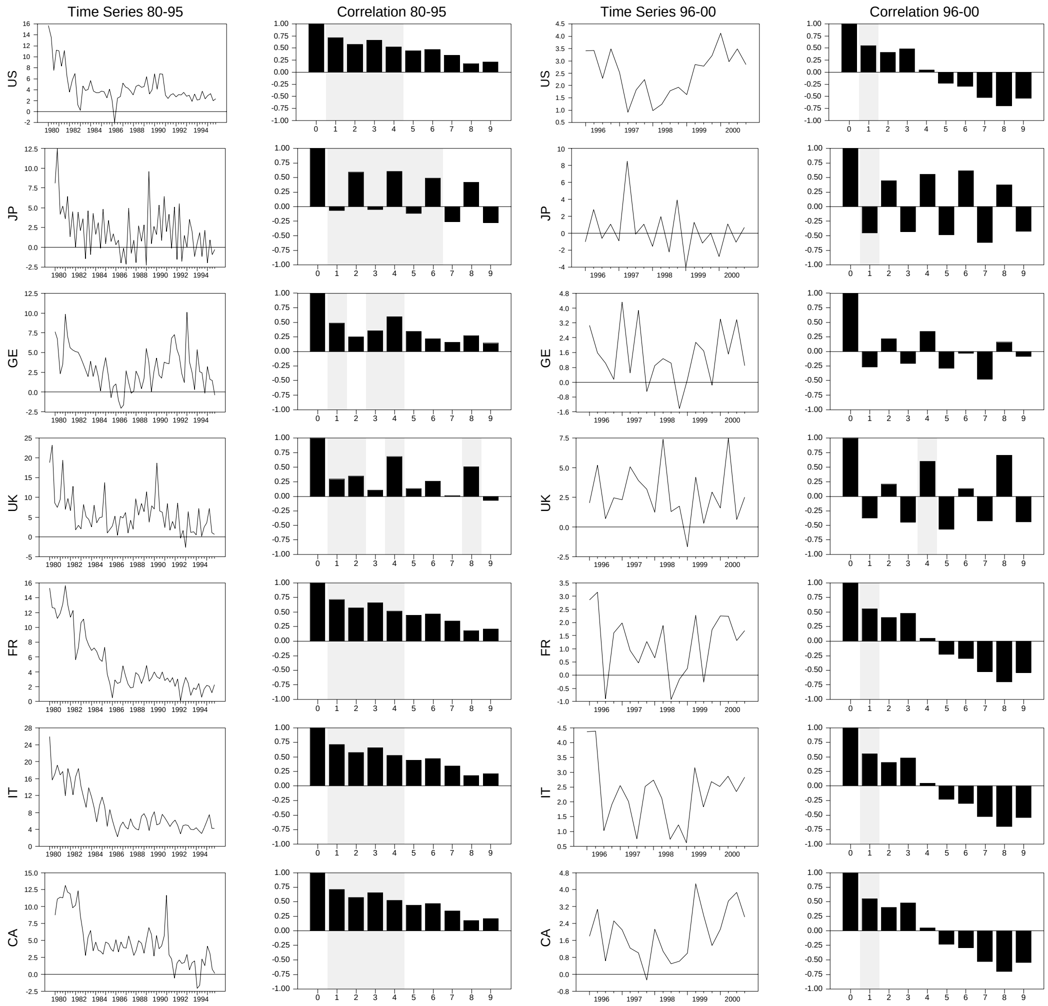

Figure 1:

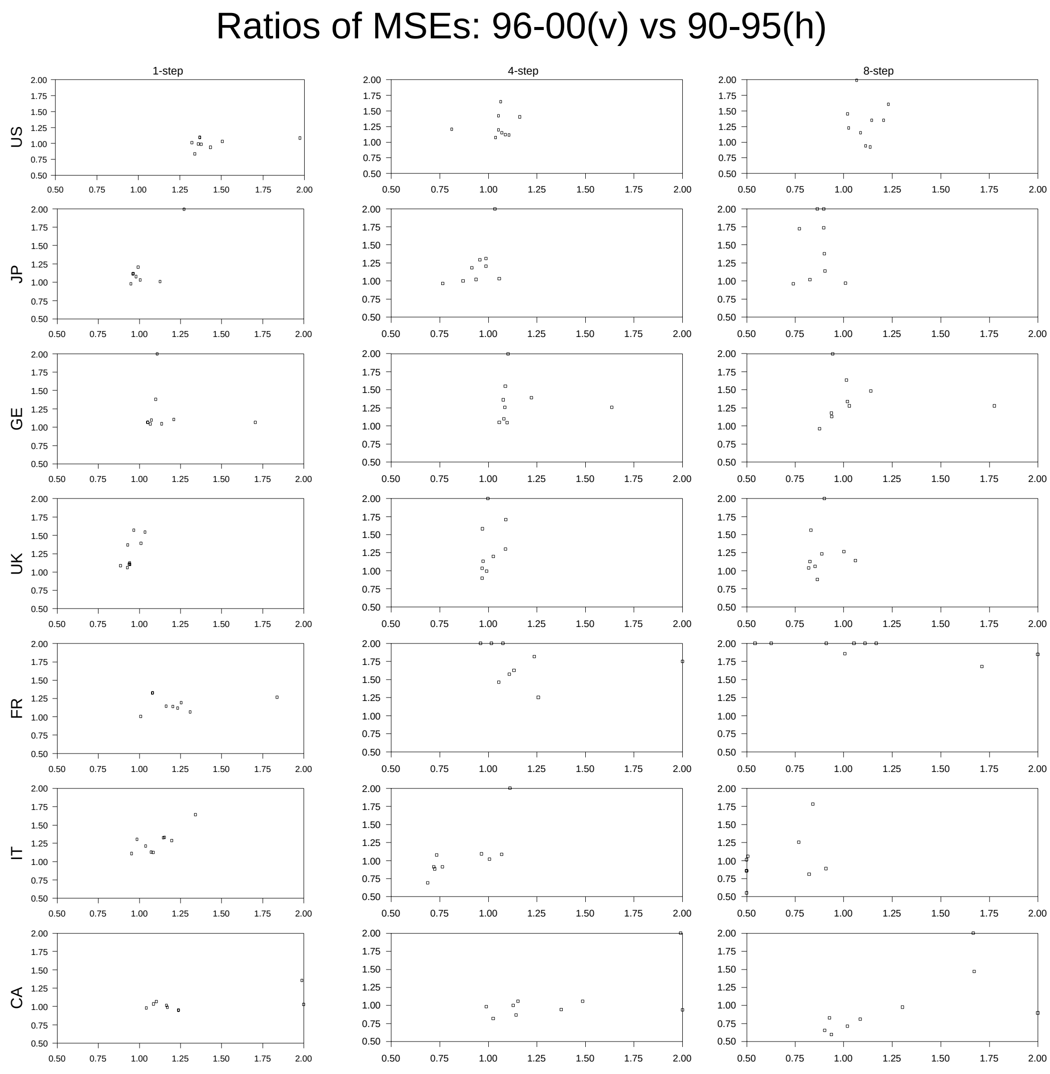

Figure 2: Bivariate Models

INFLATION FORECASTS

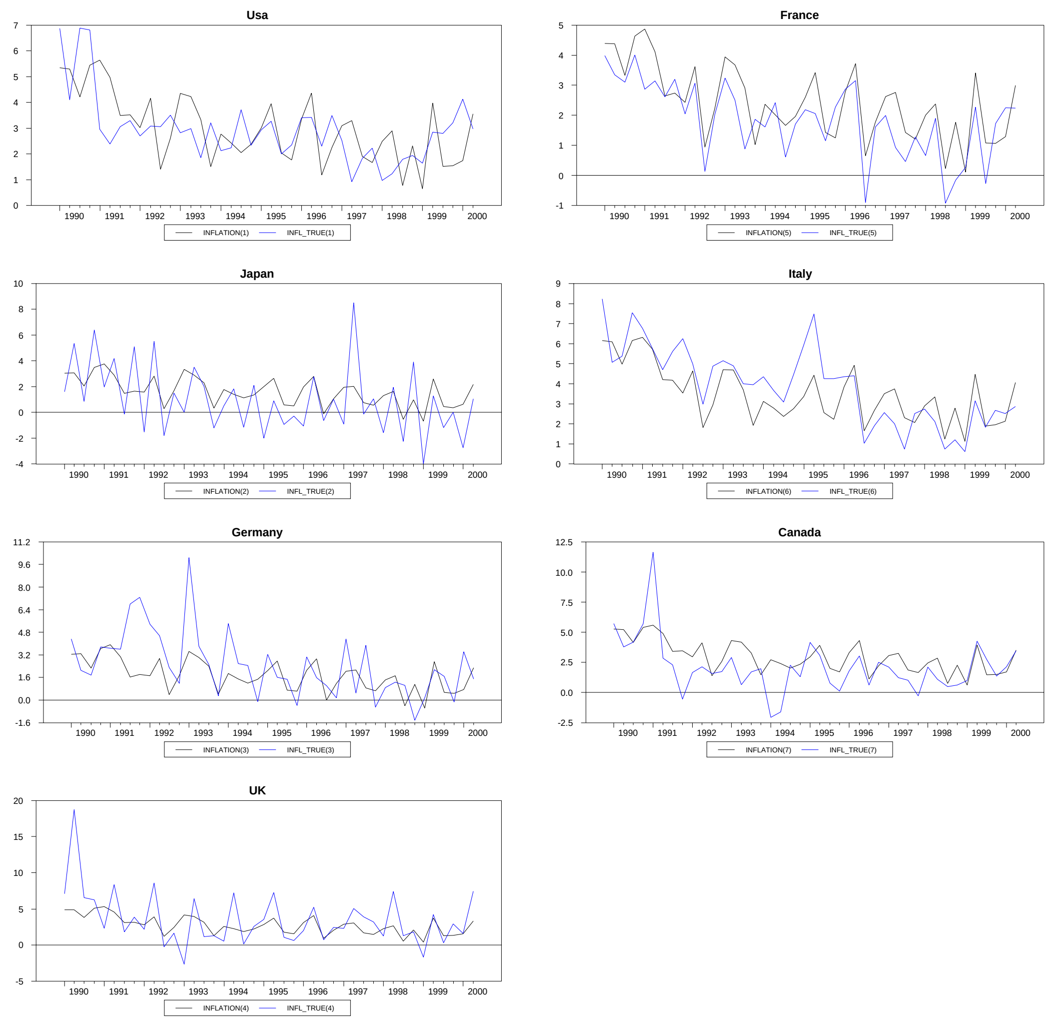

Figure 3: Index Model, 1 Year ahead Forecasts