Macroeconomic Forecasting and Structural Change*

Antonello D'Agostino †
CBFSAI, ECB

Luca Gambetti‡
Universitat Autonoma de Barcelona

Domenico Giannone§
ULB-ECARES and CEPR

Abstract

The aim of this paper is to assess whether modeling structural change increases the accuracy of macroeconomic forecasts. We conduct a simulated real time out-of-sample exercise using a Time-Varying Coefficients VAR with Stochastic Volatility to predict the inflation rate, the unemployment rate and the interest rate in the United States. The model generates accurate predictions for the three variables. In particular the forecasts of inflation are much more accurate than those obtained with any other competing model, including fixed coefficients VARs, Time-Varying ARs and the naïve random walk model. The results hold true also after the mid 80s, a period in which forecasting inflation has been particularly hard.

JEL classification: C32, E37, E47.
Keywords: Forecasting, Forecast evaluation, Inflation, Stochastic Volatility, Time-Varying Coefficients Vector Autoregression, Predictive density.

---

*We would like to thank for very useful comments and suggestions Gianni Amisano, Michele Lenza, Kieran McQuinn and Lucrezia Reichlin, two anonomuos referees and the Editor, Herman van Dijk. The views expressed in this paper are those of the authors and do not necessarily reflect those of the Central Bank and Financial Services Authority of Ireland or the European Central Bank.

†Contact: European Central Bank - Kaiserstrasse 29, D-60311 Frankfurt am Main, Germany. E-mail: antonello.dagostino@ecb.int

‡The financial support from the Spanish Ministry of Science and Innovation through grant ECO2009-09847 and the Barcelona Graduate School Research Network is gratefully acknowledged. Contact: Office B3.174, Departament d'Economia i Historia Economica, Edifici B, Universitat Autonoma de Barcelona, Bellaterra 08193, Barcelona, Spain. Tel (+34) 935811289; e-mail: luca.gambetti@uab.cat

§Contact: Université libre de Bruxelles - ECARES, 50, Avenue Roosevelt CP 114 Brussels, Belgium; e-mail: dgiannon@ulb.ac.be

1 Introduction

The US economy has undergone many structural changes during the post-WWII period. Long run trends in many macroeconomic variables have changed. Average unemployment and inflation were particularly high during the 70s and low in the last decades, see e.g. Staiger, Stock and Watson (2001). Business cycle fluctuations have moderated substantially in the last twenty years, in particular the volatility of output growth has reduced sharply, the phenomenon typically referred to as the "Great Moderation" (Stock and Watson, 2004). Inflation has become less volatile and persistent (see Cogley and Sargent, 2001, and Cogley, Primiceri and Sargen, 2010).[1]

In addition to these series-specific changes many important shifts in the relationships between macroeconomic variables have been documented. For instance, some authors have argued that the Phillips curve is no longer a good characterization of the joint dynamics of inflation and unemployment. Such a claim is partly based on the result that the predictive content of unemployment for inflation has vanished since the mid 80s (Atkenson and Ohanian, 2001, Roberts, 2006 and Stock and Watson, 2008).[2] The same period has seen significant changes in the conduct of macroeconomic policy. For example, according to many observers, monetary policy has become much more transparent and aggressive against inflation since the early 80s (Clarida, Gali, and Gertler, 2000).

In this paper we address the following question: can the accuracy of macroeconomic forecasts be improved by explicitly modeling structural change? The answer to this question is far from trivial. On the one hand, clearly, if the structure of the economy has changed, a forecasting model that can account for such changes would be better suited and should deliver better forecasts. On the other hand, however, a richer model structure implying a higher number of parameters should increase the estimation errors and reduce the forecast accuracy.

The importance of modeling time variation for forecasting was originally stressed by Doan, Litterman, and Sims (1984), but surprisingly there are only a few papers aiming at exploring the issue systematically (see Stock and Watson, 1996, Canova, 2007, Clark and McCracken, 2007, Stock and Watson, 2007, Clark, 2009). Moreover, existing studies have two limitations. First, none of them consider both changes in the parameters and shock volatilities simultaneously, two features which have proven to be very relevant to correctly characterize structural changes in the US economy (see for instance Cogley and Sargent,

[1] Changes in persistence are still debated, for instance Pivetta and Reis (2007) finds that the changes are not significant.

[2] More generally, the ability to exploit macroeconomic linkages for predicting inflation and real activity seems to have declined remarkably since the mid-1980s, see D'Agostino, Giannone, and Surico (2006) and Rossi and Sekhposyan (2008).

2005). Second, none of them is a pure real time exercise.

In this paper we forecast in real-time three US macroeconomic variables: the unemployment rate, inflation and a short term interest rate, using a Time-Varying Coefficients VAR with Stochastic Volatility (TV-VAR henceforth) as specified by Primiceri (2005).[3] The model is very flexible. In particular it allows for a) changes in the predictable component (time-varying coefficients), which can be due to variations in the structural dynamic interrelations among macroeconomic variables; and b) changes in the unpredictable component (stochastic volatility), that is, variations in the size and correlation among forecast errors, which can be due to changes in the size of exogenous shocks or their impact on macroeconomic variables.[4]

In the forecasting exercise we aim at mimicking as close as possible the conditions faced by a forecaster in real-time. We use "real-time data" to compute predictions based only on the data that were available at the time the forecasts are made. We forecast up to 3 years ahead. Using mean square forecast errors and log predictive scores we compare the forecast accuracy of the TV-VAR based to that of other standard forecasting models: fixed coefficients VARs and ARs (estimated recursively or with rolling window), Time-Varying ARs and the naïve random walk model.

Our main findings show that the TV-VAR is the only model which systematically delivers accurate forecasts for the three variables. In particular, the forecasts of inflation generated by the TV-VAR are significantly more accurate than those obtained with any other model. For unemployment, the forecasting accuracy of the TV-VAR model is very similar to that of the fixed coefficient VAR, while forecasts for the interest rate are comparable to those obtained with the Time-Varying AR. The results are confirmed over the Great Moderation period, a period in which forecasting models are often found to have difficulties in outperforming simple naïve models in forecasting many macroeconomic variables especially inflation.

Results suggest that, on the one hand time varying models are "quicker" in recognizing structural changes in the permanent components of inflation and interest rate, and, on the other hand, that short term relationships among macroeconomic variables carry out important information, once structural changes are properly taken into account.

The remainder of the paper is organized as follows: section 2 describes the TV-VAR model; section 3 explains the forecasting exercise; section 4 describes the data and presents the results; section 5 concludes.

[3] An alternative strategy to cope with structural instability, although not pursued here, is represented by model averaging, see Bates and Granger (1968) and Strachan and van Dijk (2008).

[4] Allowing for the two sources of change is also important in the light of the ongoing debate about the relative importance of changes in the predictable and unpredictable components in the Great Moderation (Giannone, Lenza, and Reichlin, 2008).

2 The Time-Varying Vector Autoregressive Model

The model is the Time-Varying Coefficients Vector Autoregression with stochastic volatility of the residuals; it has become a quite popular tool in macroeconomics, over the last few years, to address questions related to the evolution of the structure of the economy and the volatility of the shocks (see Cogley and Sargent, 2005, Primiceri, 2005, Benati, 2008, Canova and Gambetti, 2009, Gali and Gambetti, 2009).

Let $y_t = (\pi_t, UR_t, IR_t)'$ where π_t is the inflation rate, UR_t the unemployment rate and IR_t a short term interest rate. We assume that y_t satisfies

$$y_t = A_{0,t} + A_{1,t}y_{t-1} + ... + A_{p,t}y_{t-p} + \varepsilon_t \tag{1}$$

where $A_{0,t}$ is a vector of time-varying intercepts, $A_{i,t}$ are matrices of time-varying coefficients, $i = 1, ..., p$ and ε_t is a Gaussian white noise with zero mean and time-varying covariance matrix Σ_t. Let $A_t = [A_{0,t}, A_{1,t} ..., A_{p,t}]$, and $\theta_t = vec(A_t')$, where $vec(\cdot)$ is the column stacking operator. Conditional on such an assumption, we postulate the following law of motion for θ_t:

$$\theta_t = \theta_{t-1} + \omega_t \tag{2}$$

where ω_t is a Gaussian white noise with zero mean and covariance Ω. We let $\Sigma_t = F_t D_t F_t'$, where F_t is lower triangular, with ones on the main diagonal, and D_t a diagonal matrix. Let σ_t be the vector of the diagonal elements of $D_t^{1/2}$ and $\phi_{i,t}$, $i = 1, ..., n-1$ the column vector formed by the non-zero and non-one elements of the $(i+1)$-th row of F_t^{-1}. We assume that

$$\log \sigma_t = \log \sigma_{t-1} + \xi_t \tag{3}$$

$$\phi_{i,t} = \phi_{i,t-1} + \psi_{i,t} \tag{4}$$

where ξ_t and $\psi_{i,t}$ are Gaussian white noises with zero mean and covariance matrix Ξ and Ψ_i, respectively. Let $\phi_t = [\phi_{1,t}', \ldots, \phi_{n-1,t}']$, $\psi_t = [\psi_{1,t}', \ldots, \psi_{n-1,t}']$, and Ψ be the covariance matrix of ψ_t. We assume that $\psi_{i,t}$ is independent of $\psi_{j,t}$, for $j \neq i$, and that ξ_t, ψ_t, ω_t, ε_t are mutually uncorrelated at all leads and lags.[5]

2.1 Forecasts

Let $y^T = [y_1' ... y_T']'$ be the vector of data up to time T and let $y^{T+1,T+h} = [y_{T+1}' ... y_{T+h}']'$ be the vector of variables we want to forecast. Forecasts are obtained using the posterior predictive density

$$p(y^{T+1,T+h}|y^T). \tag{5}$$

[5] Giordani and Villani (2009) and Amisano and Geweke (2007) relax the Gaussianity assumption and show that this can lead to gains in terms of forecasting performance.

Let Θ_t denote the vector containing all the drifting parameters of the model and Φ the vector of constant parameters. Let $\Theta^{T+1,T+h} = [\Theta'_{T+1}...\Theta'_{T+h}]'$ and $\Theta^T = [\Theta'_1...\Theta'_T]'$. The above density can be written as

$$p(y^{T+1,T+h}|y^T) = \int\int p(y^{T+1,T+h}, \Theta^{T+h}, \Phi|y^T) d\Theta^{T+h} d\Phi \tag{6}$$

where

$$p(y^{T+1,T+h}, \Theta^{T+h}, \Phi|y^T) = p(y^{T+1,T+h}|y^T, \Theta^{T+h}, \Phi)p(\Theta^{T+1,T+h}|y^T, \Theta^T, \Phi)p(\Theta^T, \Phi|y^T)$$

A draw from the last term of the above expression is obtained using the Primiceri (2005) MCMC algorithm described in Appendix. To draw from the second term, a draw of future shocks in the drifting parameters is made and a realization of the parameters is obtained iterating on (2)-(4). Finally using such a draw for the parameters together with a draw for future VAR residuals a draw for the vector of future variables is obtained,see Cogley, Morozov, and Sargent (2005) for details. The point estimate of the forecast is the median of the predictive density.

2.2 Priors Specification

The model is estimated using Bayesian methods. While the details of the posterior simulation are accurately described in the Appendix, in this section we briefly discuss the specification of our priors. Following Primiceri (2005), we make the following assumptions for the priors densities. First, the coefficients of the covariances of the log volatilities and the hyperparameters are assumed to be independent of each other. The priors for the initial states θ_0, ϕ_0 and $\log\sigma$ are assumed to be normally distributed. The priors for the hyperparameters, Ω, Ξ and Ψ are assumed to be distributed as independent inverse-Wishart. More precisely, we have the following priors:

- Time varying coefficients: $P(\theta_0) = N(\hat{\theta}, \hat{V}_\theta)$ and $P(\Omega) = IW(\Omega_0^{-1}, \rho_1)$;
- Diagonal elements: $P(\log\sigma_0) = N(\log\hat{\sigma}, I_n)$ and $P(\Psi_i) = IW(\Psi_{0i}^{-1}, \rho_{3i})$;
- Off-diagonal elements: $P(\phi_{i0}) = N(\hat{\phi}_i, \hat{V}_{\phi_i})$ and $P(\Xi) = IW(\Xi_0^{-1}, \rho_2)$;

where the scale matrices are parameterized as follows $\Omega_0^{-1} = \lambda_1\rho_1\hat{V}_\theta$, $\Psi_{0i} = \lambda_{3i}\rho_{3i}\hat{V}_{\phi_i}$ and $\Xi_0 = \lambda_2\rho_2 I_n$. The hyper-parameters are calibrated using a time invariant recursive VAR estimated using a sub-sample consisting of the first $T_0 = 32$ observations. For the initial states θ_0 and the contemporaneous relations ϕ_{i0}, we set the means, $\hat{\theta}$ and $\hat{\phi}_i$, and the variances, $\hat{V}_\theta$ and $\hat{V}_{\phi_i}$, to be the maximum likelihood point estimates and four times its variance. For the initial states of the log volatilities, $\log\sigma_0$, the mean of the distribution is chosen to be the logarithm of the point estimates of the standard errors of the residuals of the estimated time invariant VAR. The degrees of freedom for the covariance matrix of the drifting coefficient's innovations are set to be equal to T_0, the size of the initial-sample. The degrees of freedom for the priors on the covariance of the stochastic volatilities' innovations, are set to be equal to the minimum necessary to insure that the prior is proper. In particular, ρ_1 and ρ_2 are equal to the number of rows Ξ_0^{-1} and Ψ_{0i}^{-1} plus one respectively.

The parameters λ_i are very important since they control the degree of time variations in the unobserved states. The smaller such parameters are, the smoother and smaller are the changes in coefficients. The empirical literature has set the prior to be rather conservative in terms of the amount of time variations. In our benchmark specification explosive roots are discarded and the parameters are set as follows: $\lambda_1 = (0.01)^2$, $\lambda_2 = (0.1)^2$ and $\lambda_3 = (0.01)^2$. However, we also run some robustness checks to understand the sensitivity of the model to alternative specifications. In a first simulation, we set more conservative priors (less time variation), while in a second simulation we do not discard the explosive draws.

3 Real-Time Forecasting

As stressed above, three variables are included in our model: the unemployment rate, the interest rate and the inflation rate.

The interest rate IR_t is measured by the three-month Treasury bill rate, the unemployment rate UR_t is measured by the civilian rate of unemployment. Inflation is measured in terms of annualized quarterly growth rate of prices, $\pi_t = 400\log(\frac{P_t}{P_{t-1}})$, where P_t is the GDP deflator.

Following Cogley and Sargent (2001, 2005) and Cogley, Primiceri, and Sargent (2008), we transform the unemployment rate and the interest rate, available at monthly frequencies, into quarterly series by taking the value at the second month of the quarter for the unemployment rate and the value at the first month of the quarter for the interest rate.

We use quarterly real time data vintages from 1969:Q4 to 2007:Q4. Vintages can differ since new data on the most recent period are released, but also because old data get revised.[6] As a convention we date a vintage as the last quarter for which all data are available. For each vintage the sample starts in 1948:Q1.[7]

We forecast the interest rate and the unemployment rate, IR_{t+h} and UR_{t+h} respectively,

[6]Real-time data for unemployment rate and the deflator are available on the Federal Reserve Bank of Philadelphia website at: http://www.phil.frb.org/econ/forecast/reaindex.html. The interest rate series is not subject to revisions and is available on the FRED dataset of the Federal Reserve Bank of St. Louis (mnemonics TB3MS), at: http://research.stlouisfed.org/fred2/series/TB3MS.

[7]The vintages have a different time length, for example the sample span for the first vintage is 1948:Q1-1969:Q4, while the sample span for the last available vintage is 1948:Q1-2007:Q4.

h-quarter ahead. As far as prices are concerned, we forecast the h-period ahead annualized price changes of the GDP deflator, which corresponds to the average over the first h horizons of the forecast of the annualized quarterly inflation rate: $\pi^h_{t+h} = \frac{400}{h} log(\frac{P_{t+h}}{P_t}) = \frac{1}{h}(\pi_{t+1} + ... + \pi_{t+h})$.

3.1 The Forecasting Exercise

We perform an out-of-sample simulation exercise. The procedure consists of generating the forecasts using the predictive density (6) with the same information available to the econometrician at the time the forecasts are made. The simulation exercise begins in 1969:Q4 and, for such a vintage, parameters are estimated using the sample 1948:Q1 to 1969:Q4. The model is estimated with two lags. We compute the forecasts up to 12 quarters ahead outside the estimation window, from 1970:Q1 to 1972:Q4, and the results are stored.[8] Then, we move one quarter ahead and re-estimate the model using the data in vintage 1970:Q1. Forecasts from 1970:Q2 to 1973:Q1 are again computed and stored. This procedure is then repeated using all the available vintages. Predictions are compared with ex-post realized data vintages. Since data are continuously revised at each quarter, several vintages are available. Following Romer and Romer (2000), predictions are compared with the figures published after the next two subsequent quarters. These figures are conceptually similar to the series predicted in real time since they do not incorporate re-benchmarking and other definitional changes. Nonetheless the results are similar if we compare the forecasts with the data from the last vintage.

We evaluate the forecasting accuracy by considering both point forecasts (the median of the predictive density) and density forecasts. Specifically we focus on mean square forecast errors (MSFE) and log predictive scores. Accuracy is evaluated over two samples: the full sample, 1970:Q1-2007:Q4 and the post-84 sample, 1985:Q1-2007:Q4. The latter corresponds to the great moderation period.

3.2 Other Forecasting Models

We compare the forecasts obtained with the TV-VAR with those obtained using different standard forecasting models. We consider six competing models:

1) A Time-varying Coefficients Autoregressive process with Stochastic Volatility (TV-AR) estimated separately for each of the three series. This corresponds to the general TV-VAR model in which the matrices of the autoregressive parameters and the co-

[8] In the simulation exercise forecasts for horizon $h = 1$ correspond to nowcast, given that in real time data are available only up to the previous quarter.

variances of the residuals are diagonal. We keep the same specification and priors used for the TV-VAR.

2-3) A stochastic volatility VAR (SV-VAR) and AR (SV-AR) with fixed coefficients. We keep the same prior specification for the stochastic volatilities and assume a flat prior for the VAR coefficients.

4-5) Fixed coefficients VARs estimated using both a rolling estimation scheme (VAR-ROL), i.e. using only the ten most recent years of data available at each time the forecasts are made and a recursive estimation scheme (VAR-REC), i.e. using all the data available when the forecasts are made.

6-7) Fixed coefficients ARs estimated separately for each variables using a recursive estimation scheme (AR-REC) and a rolling estimation scheme (AR-ROL).

8) Random walk or naïve models (RW). According to these models, the unemployment rate and the interest rate forecasts, for all future horizons, are equal to the value observed in the current quarter. In the case of inflation, we follow Atkeson and Ohanian (2001) which proposes a model of "no change" for annual inflation. More precisely, the naïve model predicts quarterly inflation to be constant and equal to the average inflation over the most recent past four quarters. It is worth noticing that all the random walk models are restricted ARs. In particular, for the interest rate and the unemployment rate we have an AR(1) with autoregressive coefficient equal to 1. For inflation we have an AR(4) with all coefficients equal to 1/4.

All the models are estimated using Bayesian techniques and two lags. However for the time invariant AR and VAR we specify a Jeffreys prior which implies that the forecasts coincide with those obtained using maximum likelihood (ML). For the naïve model the only unknown parameter is the residual variance, for which we also use the Jefferys prior.

4 Results

This section discusses the main findings of the forecasting exercise.

4.1 How much time variation?

As a preliminary step of our analysis, in order to understand whether time variation is an important characteristic of the dataset, we estimate the TV-VAR model over the whole sample.

Figure 1 shows the evolution of the coefficients over time. The solid lines are the posterior means of the coefficients and the dashed lines are the 68% confidence bands. Many coefficients display constant patterns, while four parameters are characterized by remarkable fluctuations over time. Figure 2 shows the evolution of the standard deviations of the residuals. Again the solid lines are the posterior means of the coefficients and the dashed lines the 68% confidence bands. All the volatilities exhibit substantial time variation over the sample. The figure also shows that, concomitant with the great moderation period, around the mid 80s, there is a sharp drop in the volatilities.

All in all, the results show that time variation is an important feature of the data.

4.2 Forecast Accuracy: Point Forecast

Table 1 summarizes the results of the real time forecast evaluation based on the MSFE over the whole sample for the three variables (inflation rate π_t, unemployment rate UR_t and the interest rate IR_t), and for forecast horizons of one quarter, one year, two years and three years. To facilitate the comparison between various models, the results are reported in terms of relative MSFE, that is the ratio between the MSFE of a particular model and the MSFE of the naïve model, used as benchmark. When the relative MSFE is less than one, the forecasts of a given model are, on average, more accurate that those produced with the benchmark model. For example, a value of 0.8 indicates that the model under consideration improves upon the benchmark by 20%. For the naïve models we report the MSFE.

Overall the TV-VAR produces very accurate forecasts for all the variables and, on average, performs better than any other model considered. In particular it outperforms the naïve benchmark for all the variables at all horizons, except for unemployment at an horizon of one quarter, with gains ranging from 5 to 28 percent.

The best relative performances of the TV-VAR model is obtained for inflation. For this variable, the TV-VAR model produces the best forecast with an average (over the horizons) improvement upon the benchmark of about 30%. A relative good performance is also observed for the TV-AR with improvements of about 10% at horizons of 1 and 2 years. The other time invariant specifications, univariate and multivariate, fail to improve upon the benchmark in terms of forecasting accuracy.

As far as the interest rate is concerned, univariate and multivariate time-varying models both perform better than the constant parameters models. For unemployment, multivariate models tend to deliver better results than the random walk model and other univariate models especially at long horizons.

In conclusion, our main findings show that the TV-VAR is the only model which systematically delivers accurate forecasts for the three variables. In particular, the forecasts of inflation generated by the TV-VAR are significantly more accurate than those obtained with any other model.

To understand whether predictability comes from stochastic volatilities, as Clark (2011) suggests, or rather from time-varying coefficients we have investigated the forecasting performance of a stochastic volatility VAR and AR with fixed autoregressive parameters, SV-VAR and SV-AR respectively. In general both models improve in terms of forecasting accuracy upon constant parameters models. Specifically the forecasts of the unemployment rate obtained with the SV-VAR are found to be particularly accurate. Nonetheless the TV-VAR still produces better forecasts for inflation. All in all findings suggest an important role for both time-varying coefficient and volatilities.

Table 2 shows the results for the "Great Moderation" period. Most of the previous findings are confirmed. First, the TV-VAR model generates the most accurate forecasts for all the variables. Second, the TV-VAR performs particularly well in forecasting inflation. The average (over the horizons) improvement is about 30%. At a 3-year horizon the improvement in forecasting accuracy is almost two times that obtained in the full sample, being now about 52%. Third, the forecasts of the interest rate and the unemployment rate obtained with time varying models are more accurate than those obtained in the full sample.

The results obtained for the shorter sample are particularly interesting given that several authors have shown that over the post-84 period standard fixed-coefficients models fail to improve upon the simple naïve random walk. Here we show that, once structural changes, modeled by means of time-varying parameters and stochastic volatility, are accounted for, inflation becomes much more predictable.

Finally the higher predictability of the interest rate might reflect the increased importance of the systematic predictable component of monetary policy in the last two decades.

We also check of the robustness of the results with two alternative model specifications. First we tried with a different parametrizations of the λ's, namely $\lambda_1 = 0.00001$, $\lambda_2 = 0.001$, $\lambda_3 = 0.00001$. Second, we repeated the forecasting exercise keeping the explosive draws obtained in the Gibbs sampling algorithm. In the first case results are comparable, in terms of accuracy, with those obtained with the previous specification, although the magnitudes of the gains are slightly smaller. In the second exercise the accuracy of the forecasts deteriorates for all the variables and in particular for the unemployment rate and the interest rate.[9]

4.3 Forecast Accuracy: Predictive density

In this section we asses the accuracy of the models by looking at the entire predictive density. In a situation in which monetary policy and the economy are subject to ongoing changes, the

[9]Results are available in the ECB working paper version of the article.

ability of correctly predicting not only the central tendency but also the uncertainty of the forecasts has become an important aspect for assessing the quality of different models (See Clark, 2009; Amisano and Giacomini, 2007; Jore, Mitchell, and Vahey, 2010). By allowing the variance of the forecast errors to change over time, the TV-VAR model is in principle well suited to characterizing not only point forecasts but also forecasting uncertainty (See Cogley, Morozov, and Sargent, 2005).

Figure 3 plots the 2-year ahead forecast of unemployment, inflation and the interest rate. Shaded areas represent the 68% and 90% confidence bands and the solid lines are the true series. The model appears to correctly capture not only short-run movements of the series but also lower frequency fluctuations such as the rise and the fall of trend inflation.

To gauge the quality of the density intervals shown in Figure 3 we first perform PITs tests. The PITs are the value of the predictive cumulative distribution evaluated at the true realized values of the time series, see for instance Mitchell and Wallis (2010). PITS are widely used to assess the calibration of density forecasts (most recent works include Jore et al., 2009, Mitchell and Wallis, 2010, Geweke and Amisano, 2010, Huurman et al., 2010, and Clark, 2011). If the density forecasts are well calibrated, then the PITs should be uniformly distributed in the interval zero-one. Testing uniformity of the PITs is equivalent to test that the inverse normal transformation of the PITs is standard normal. This is done by testing whether the third and fourth moments of the inverse normal transformation of the PITs are equal to zero and three respectively.

In Tables 3 we report the third and fourth moments along with standard deviations in brackets.[10] Results suggest that the density forecasts produced by the TV-VAR model are well calibrated for inflation, since we do not find evidence of excess skewness or excess kurtosis. Evidence is more mixed for unemployment and the interest rate for which we find evidence of excess kurtosis at some horizons. In general, the calibration of the density forecasts produced by using TV-VAR model compares well with the ones obtained using the other models.

The accuracy of the density forecast is assessed by using the log-predictive scores of the model under consideration. The log-predictive score is simply the logarithm of the value of the predictive density evaluated at the true realized value of the time series[11]. A high value is obtained when high probability is assigned to the actual outcome. At the limit, the maximum value of the log score (zero) is obtained if a probability of 100% is assigned to the actual outcome.

[10] Following Bai and Ng (2005) we report the heteroskedasticity and autocorrelation consistent (HAC) standard deviation estimator. The estimates of the standard deviation should, however, be taken with caution since parameters estimation uncertainty is ignored.

[11] The predictive density is estimated by smoothing the empirical distribution of the forecasts obtained with the Gibbs sampler using a normal kernel function.

Using forecast densities we compute the log predictive score for every model, horizon and time period in the evaluation sample. Then we take the differences between the log score of the TV-VAR and the log score of any other model and regress them on a constant term. The constant represents the average difference between the log predictive scores. A constant greater than zero indicate that the density forecasts produced by the TV-VAR are more accurate.

The $t-$statistics associated to the constant can be interpreted as a Diebold-Mariano test of equal predictive accuracy. Results should be taken with caution, since the test ignores estimation uncertainty. [12]

The results of the evaluation of density forecasts over the full sample are reported in Table 3. We report the estimated constant and the heteroskedasticity and autocorrelation consistent (HAC) estimator of its standard deviation in brackets. The evaluation for the post-1984 period is reported in Table 4. Conclusions based on the magnitude of the relative MSFE in the previous section are largely confirmed when looking at the predictive density. In general the TV-VAR tends to significantly dominate all the other models, for all the horizons and all the variables.

Summing up, the good performance of the TV-VAR in forecasting inflation is confirmed when considering density forecasts. Improvements with respect to all the other models are large and significant especially for inflation. For interest rate the TV-VAR together with the TV-AR are the models that perform the best, whereas as far as unemployment is concerned the TV-VAR is the best model together with VAR-ROL.[13]

5 Conclusions

The US economy has changed substantially during the post-WWII period. This paper tries to assess whether explicitly modeling these changes can improve the forecasting accuracy of key macroeconomic time series.

We produce real time out-of sample forecasts for inflation, the unemployment rate and a short term interest rate using a time-varying coefficients VAR with stochastic volatility and we compare its forecasting performance to that of other standard models. Our findings show that the TV-VAR is the only model which systematically delivers accurate forecasts for the three variables. In particular, the forecasts of inflation generated by the TV-VAR are much more accurate than those obtained with any other model. These results hold for the Great Moderation period, after the mid 80's. The last result is particularly relevant given

[12]The $t-$statistics corresponds to the Amisano and Giacomini (2007) test when the models are estimated using a rolling scheme. This is not the case in our exercise since we also use a recursive estimation procedure.

[13]Results are confirmed also when testing for equal predictive accuracy of point forecasts. The results are available upon request.

that previous studies have found that over such a period forecasting models have considerable difficulty in outperforming simple naïve models in predicting many macroeconomic variables, in particular inflation.

We draw two main conclusions. First, taking into account structural economic change is important for forecasting. Second, the TV-VAR model is a powerful tool for real-time forecasting since it incorporates in a flexible but parsimonious manner the prominent features of a time-varying economy.

References

[1] Amisano G, and Geweke J. (2007). Hierarchical Markov normal mixture models with applications to financial asset returns. Working Paper Series 831, European Central Bank.

[2] Amisano G, and Geweke J. (2010). Comparing and evaluating Bayesian predictive distributions of asset returns. International Journal of Forecasting 26(2):216-230.

[3] Amisano G, and Giacomini R. (2007). Comparing Density Forecasts via Weighted Likelihood Ratio Tests. Journal of Business & Economic Statistics, 25: 177-190.

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

[5] Bai, J, and Ng S. (2005). Tests for Skewness, Kurtosis, and Normality for Time Series Data. Journal of Business & Economic Statistics, 23:49-60.

[6] Bates J, and Granger C. W. J.(1968). The Combination of Forecasts. Operational Research Quarterly, 20:451-468.

[7] Benati L. (2008). The Great Moderation in the United Kingdom. Journal of Money, Credit and Banking, 40(1): 121-147.

[8] Canova F. (2007). G-7 Inflation Forecasts: Random Walk, Phillips Curve Or What Else?. Macroeconomic Dynamics, 11(01):1-30.

[9] Canova F, and Gambetti L. (2009). Structural Changes in the US Economy: Is There a Role for Monetary Policy?. Journal of Economic Dynamics and Control, 33(2):477-490.

[10] Carter C. K, and Kohn R. (1994). On Gibbs Sampling for State Space Models. Biometrika, 81:541-553.

[11] Clarida R, Gali J, and Gertler M. (2000). Monetary Policy Rules And Macroeconomic Stability: Evidence And Some Theory. Quarterly Journal of Economics, 115(1):147-180.

[12] Clark T. E. (2009). Real-time density forecasts from VARs with stochastic volatility. Research Working Paper RWP 09-08, Federal Reserve Bank of Kansas City.

[13] Clark T. E, and McCracken M. W. (2007). Forecasting with small macroeconomic VARs in the presence of instabilities. Journal of Business Economics and Satistics, forthcoming.

[14] Cogley T, Morozov S, and Sargent T. J. (2005). Bayesian fan charts for U.K. infl ation: Forecasting and sources of uncertainty in an evolving monetary system. Journal of Economic Dynamics and Control, 29(11):1893-1925.

[15] Cogley T, Primiceri G. E, and Sargent T. J. (2008). Inflation-Gap Persistence in the U.S. American Economic Journal: Macroeconomics 2(1):43-69.

[16] Cogley T, and Sargent T. J. (2001). Evolving Post WWII U.S. Inflation Dynamics. in Macroeconomics Annual, 331-373. MIT Press.

[17] Cogley T, and Sargent T. J. (2005). Drifts and Volatilities: Monetary Policies and Outcomes in the Post WWII US. Review of Economic Dynamics, 8: 262-302.

[18] D'Agostino A, Giannone D, and Surico P. (2006). (Un)Predictability and Macroeconomic Stability. Working Paper Series 605, European Central Bank.

[19] Doan T, Litterman R, and Sims C. A. (1984). Forecasting and Conditional Projection Using Realistic Prior Distributions. Econometric Reviews, 3:1-100.

[20] Gali J, and Gambetti L. (2009). On the sources of the Great Moderation. American Economic Journal: Macroeconomics, 1(1):26-57.

[21] Giannone D, Lenza M, and Reichlin L. (2008). Explaining The Great Moderation: It Is Not The Shocks. Journal of the European Economic Association, 6(2-3):621-633.

[22] Giordani P, and Villani M. (2009). Forecasting Macroeconomic Time Series With Locally Adaptive Signal Extraction. Working Paper Series 234, Sveriges Riksbank (Central Bank of Sweden).

[23] Huurman C, Ravazzolo F, and Zhou, C. (2010). The power of weather. Computational Statistics and Data Analysis, forthcoming.

[24] Jore A. S, Mitchell J, and Vahey S. P. (2010). Combining forecast densities from VARs with uncertain instabilities. Journal of Applied Econometrics, 25(4):621-634.

[25] Mitchell J, and Wallis K. (2010). Evaluating Density Forecasts: Forecast Combinations, Model Mixtures, Calibration and Sharpness. Journal of Applied Econometrics, forthcoming.

[26] Pivetta F, and Reis R. (2007). The Persistence of In ation in the United States. Journal of Economic Dynamics and Control, 31(4): 1326-1358.

[27] Primiceri G. E. (2005). Time Varying Structural Vector Autoregressions and Monetary Policy. Review of Economic Studies, 72:821-852. 16

[28] Roberts J. M. (2006). Monetary Policy and In ation Dynamics. International Journal of Central Banking, 2(3):193-230.

[29] Romer D. H, and Romer C. (2000). Federal Reserve Information and the Behavior of Interest Rates. American Economic Review, 90:429-457.

[30] Rossi B, and Sekhposyan T. (2008). Has models' forecasting performance for US output growth and in ation changed over time, and when?. Manuscript, Duke University.

[31] Staiger D, Stock J. H, and Watson M. W. (2001). Prices, Wages, and the U.S. NAIRU in the 1990s. in In The Roaring Nineties: Can Full Employment Be Sustained?, ed. by A. B. Krueger, and R. M. Solow. Russell Sage Foundation.

[32] Stock J. H, and Watson M. W. (1996). Evidence on Structural Instability in Macroeconomic Time Series Relations. Journal of Business & Economic Statistics, 14(1), 1130.

[33] Stock J. H, and Watson M. W. (2004). Has the Business Cycle Changed and Why?. in NBER Macroeconomics Annual 2003, ed. by M. Gertler, and K. Rogoff. MIT Press.

[34] Stock J. H, and Watson M. W. (2007). Why Has U.S. In ation Become Harder to Forecast? Journal of Money, Credit and Banking, 39(1):

[35] Stock J. H, and Watson M. W. (2008). Phillips Curve In ation Forecasts. NBER Working Paper 14322, NBER.

[36] Strachan R. W, and van Dijk H. K. (2008). Bayesian Averaging over Many Dynamic Model Structures with Evidence on the Great Ratios and Liquidity Trap Risk. Tinbergen Institute Discussion Papers 08-096/4, Tinbergen Institute. 17

Appendix

Estimation is done using Bayesian methods. To draw from the joint posterior distribution of model parameters we use a Gibbs sampling algorithm along the lines described in Primiceri (2005). The basic idea of the algorithm is to draw sets of coefficients from known conditional posterior distributions. The algorithm is initialized at some values and, under some regularity conditions, the draws converge to a draw from the joint posterior after a burn in period. Let z be $(q \times 1)$ vector, we denote z^T the sequence $[z_1', ..., z_T']'$. Each repetition is composed of the following steps:

1. $p(\sigma^T | x^T, \theta^T, \phi^T, \Omega, \Xi, \Psi, s^T)$
2. $p(s^T | x^T, \theta^T, \sigma^T, \phi^T, \Omega, \Xi, \Psi)$[14]
3. $p(\phi^T | x^T, \theta^T, \sigma^T, \Omega, \Xi, \Psi, s^T)$
4. $p(\theta^T | x^T, \sigma^T, \phi^T, \Omega, \Xi, \Psi, s^T)$
5. $p(\Omega | x^T, \theta^T, \sigma^T, \phi^T, \Xi, \Psi, s^T)$
6. $p(\Xi | x^T, \theta^T, \sigma^T, \phi^T, \Omega, \Psi, s^T)$
7. $p(\Psi | x^T, \theta^T, \sigma^T, \phi^T, \Omega, \Xi, s^T)$

Gibbs sampling algorithm

• Step 1: sample from $p(\sigma^T | y^T, \theta^T, \phi^T, \Omega, \Xi, \Psi, s^T)$

To draw σ^T we use the algorithm of Kim, Shephard and Chibb (KSC) (1998). Consider the system of equations $y_t^* \equiv F_t^{-1}(y_t - X_t'\theta_t) = D_t^{1/2} u_t$, where $u_t \sim N(0, I)$, $X_t = (I_n \otimes x_t')$, and $x_t = [1_n, y_{t-1} ... y_{t-p}]$. Conditional on y^T, θ^T, and ϕ^T, y_t^* is observable. Squaring and taking the logarithm, we obtain

$$y_t^{**} = 2r_t + \upsilon_t \tag{7}$$

$$r_t = r_{t-1} + \xi_t \tag{8}$$

where $y_{i,t}^{**} = \log((y_{i,t}^*)^2 + 0.001)$ - the constant (0.001) is added to make estimation more robust - $\upsilon_{i,t} = \log(u_{i,t}^2)$ and $r_t = \log \sigma_{i,t}$. Since, the innovation in (7) is distributed as $\log \chi^2(1)$, we use, following KSC, a mixture of 7 normal densities with component probabilities q_j, means $m_j - 1.2704$, and variances v_j^2 (j=1,...,7) to transform the system in a Gaussian one, where $\{q_j, m_j, v_j^2\}$ are chosen to match the moments of the $\log \chi^2(1)$ distribution. The values are:

[14] See below the definition of s^T.

Table A1: Parameters Specification

j	q_j	m_j	v_j^2
1.0000	0.0073	-10.1300	5.7960
2.0000	0.1056	-3.9728	2.6137
3.0000	0.0000	-8.5669	5.1795
4.0000	0.0440	2.7779	0.1674
5.0000	0.3400	0.6194	0.6401
6.0000	0.2457	1.7952	0.3402
7.0000	0.2575	-1.0882	1.2626

Let $s^T = [s_1, ..., s_T]'$ be a matrix of indicators selecting the member of the mixture to be used for each element of υ_t at each point in time. Conditional on s^T, $(\upsilon_{i,t}|s_{i,t} = j) \sim N(m_j - 1.2704, v_j^2)$. Therefore we can use the algorithm of Carter and Kohn (1994) to draw r_t (t=1,...,T) from $N(r_{t|t+1}, R_{t|t+1})$, where $r_{t|t+1} = E(r_t|r_{t+1}, y^t, \theta^T, \phi^T, \Omega, \Xi, \Psi, s^T,)$ and $R_{t|t+1} = Var(r_t|r_{t+1}, y^t, \theta^T, \phi^T, \Omega, \Xi, \Psi, s^T)$.

• Step 2: sample from $p(s^T|y^T, \theta^T, \sigma^T, \phi^T, \Omega, \Xi, \Psi)$

Conditional on $y_{i,t}^{**}$ and r^T, we independently sample each $s_{i,t}$ from the discrete density defined by $Pr(s_{i,t} = j|y_{i,t}^{**}, r_{i,t}) \propto f_N(y_{i,t}^{**}|2r_{i,t} + m_j - 1.2704, v_j^2)$, where $f_N(y|\mu, \sigma^2)$ denotes a normal density with mean μ and variance σ^2.

• Step 3: sample from $p(\phi^T|y^T, \theta^T, \sigma^T, \Omega, \Xi, \Psi, s^T)$

Consider again the system of equations $F_t^{-1}(y_t - X_t'\theta_t) = F_t^{-1}\hat{y}_t = D_t^{1/2}u_t$. Conditional on θ^T, $\hat{y}_t$ is observable. Since F_t^{-1} is lower triangular with ones in the main diagonal, each equation in the above system can be written as

$$\hat{y}_{1,t} = \sigma_{1,t}u_{1,t} \tag{9}$$

$$\hat{y}_{i,t} = -\hat{y}_{[1,i-1],t}\phi_{i,t} + \sigma_{i,t}u_{i,t} \quad i = 2, ..., n \tag{10}$$

where $\sigma_{i,t}$ and $u_{i,t}$ are the ith elements of σ_t and u_t respectively, $\hat{y}_{[1,i-1],t} = [\hat{y}_{1,t}, ..., \hat{y}_{i-1,t}]$. Under the block diagonality of Ψ, the algorithm of Carter and Kohn (1994) can be applied equation by equation, obtaining draws for $\phi_{i,t}$ from a $N(\phi_{i,t|t+1}, \Phi_{i,t|t+1})$, where $\phi_{i,t|t+1} = E(\phi_{i,t}|\phi_{i,t+1}, y^t, \theta^T, \sigma^T, \Omega, \Xi, \Psi)$ and $\Phi_{i,t|t+1} = Var(\phi_{i,t}|\phi_{i,t+1}, y^t, \theta^T, \sigma^T, \Omega, \Xi, \Psi)$.

• Step 4: sample from $p(\theta^T|y^T, \sigma^T, \phi^T, \Omega, \Xi, \Psi, s^T)$

Conditional on all other parameters and the observables we have

$$y_t = X_t'\theta_t + \varepsilon_t \tag{11}$$

$$\theta_t = \theta_{t-1} + \omega_t \tag{12}$$

Draws for θ_t can be obtained from a $N(\theta_{t|t+1}, P_{t|t+1})$, where $\theta_{t|t+1} = E(\theta_t|\theta_{t+1}, y^T, \sigma^T, \phi^T, \Omega, \Xi, \Psi)$ and $P_{t|t+1} = Var(\theta_t|\theta_{t+1}, y^T, \sigma^T, \phi^T, \Omega, \Xi, \Psi)$ are obtained with the algorithm of Carter and Kohn (1994).

- Step 5: sample from $p(\Omega|y^T, \theta^T, \sigma^T, \phi^T, \Xi, \Psi, s^T)$

Conditional on the other coefficients and the data, Ω has an Inverse-Wishart posterior density with scale matrix $\Omega_1^{-1} = (\Omega_0 + \sum_{t=1}^{T} \Delta\theta_t(\Delta\theta_t)')^{-1}$ and degrees of freedom $df_{\Omega_1} = df_{\Omega_0} + T$, where Ω_0^{-1} is the prior scale matrix, df_{Ω_0} are the prior degrees of freedom and T is length of the sample use for estimation. To draw a realization for Ω make df_{Ω_1} independent draws z_i (i=1,...,df_{Ω_1}) from $N(0, \Omega_1^{-1})$ and compute $\Omega = (\sum_{i=1}^{df_{\Omega_1}} z_i z_i')^{-1}$ (see Gelman et. al., 1995).

- Step 6: sample from $p(\Xi_{i,i}|y^T, \theta^T, \sigma^T, \phi^T, \Omega, \Psi, s^T)$

Conditional the other coefficients and the data, Ξ has an Inverse-Wishart posterior density with scale matrix $\Xi_1^{-1} = (\Xi_0 + \sum_{t=1}^{T} \Delta \log \sigma_t(\Delta \log \sigma_t)')^{-1}$ and degrees of freedom $df_{\Xi_1} = df_{\Xi_0} + T$ where Ξ_0^{-1} is the prior scale matrix and df_{Ξ_0} the prior degrees of freedom. Draws are obtained as in step 5.

- Step 7: sample from $p(\Psi|y^T, \theta^T, \sigma^T, \phi^T, \Omega, \Xi, s^T)$.

Conditional on the other coefficients and the data, Ψ_i has an Inverse-Wishart posterior density with scale matrix $\Psi_{i,1}^{-1} = (\Psi_{i,0} + \sum_{t=1}^{T} \Delta\phi_{i,t}(\Delta\phi_{i,t})')^{-1}$ and degrees of freedom $df_{\Psi_{i,1}} = df_{\Psi_{i,0}} + T$ where $\Psi_{i,0}^{-1}$ is the prior scale matrix and $df_{\Psi_{i,0}}$ the prior degrees of freedom. Draws are obtained as in step 5 for all i.

In the first estimation (the first out-of-sample forecast iteration), we make 12000 repetitions discarding the first 10000 and collecting one out of five draws. For the other estimates, we initialize the coefficients with the medians obtained in the previous estimation, and we make 2500 repetitions discarding the first 500 and collecting one out of five draws.

Tables

Table 1: Forecasting Accuracy over the sample 1970-2007: mean square forecast errors.

Horizon (quarters)	Variable	RW (MSFE)	AR-REC (RMSFE)	AR-ROL (RMSFE)	SV-AR (RMSFE)	TV-AR (RMSFE)	VAR-REC (RMSFE)	VAR-ROL (RMSFE)	SV-VAR (RMSFE)	TV-VAR (RMSFE)
	π	2.15	1.13	1.08	1.05	1.03	1.15	1.01	1.14	0.85
1	UR	0.15	1.00	1.08	0.98	1.00	0.99	1.18	0.96	1.04
	IR	0.87	1.12	1.23	1.01	1.04	0.99	1.09	0.94	0.99
	Avg.		1.08	1.13	1.01	1.02	1.04	1.09	1.01	0.96
	π	2.24	1.17	1.03	0.82	0.88	1.37	1.22	1.01	0.62
4	UR	1.07	1.03	1.24	0.95	1.01	0.67	0.91	0.67	0.77
	IR	3.46	1.05	1.20	0.93	0.95	0.96	1.39	0.93	0.93
	Avg.		1.08	1.16	0.90	0.95	1.00	1.17	0.87	0.77
	π	3.06	1.19	1.13	0.89	0.93	1.6	1.38	1.11	0.65
8	UR	2.39	0.95	1.14	0.88	0.95	0.45	0.63	0.42	0.61
	IR	7.54	1.05	1.18	1.11	0.92	0.99	1.44	0.85	0.89
	Avg.		1.06	1.15	0.88	0.93	1.01	1.15	0.80	0.72
	π	3.31	1.28	1.24	0.95	1.00	1.93	1.60	1.26	0.69
12	UR	3.22	0.85	1.12	0.79	0.86	0.47	0.85	0.40	0.51
	IR	10.28	1.08	1.15	0.82	0.91	1.03	1.32	0.79	0.86
	Avg.		1.07	1.17	0.85	0.92	1.14	1.26	0.82	0.72

The table reports the results relative to the forecasting accuracy using point forecasts. The variable we forecast are inflation (π_t), the unemployment rate (UR_t) and the interest rate (IR_t). The forecasting models are: RW - random walk; AR-REC - AR estimated recursively; AR-ROL - AR estimated with a rolling window; TV-VAR - time-varying VAR; VAR-REC - VAR estimated recursively; VAR-ROL - VAR estimated with a rolling window. For the random walk model we report the mean square forecast error (MSFE). For the other models we report the relative mean square forecast error (RMSFE), i.e. the ratio of the MSFE of a particular model to the MSFE of the naïve model. For each horizon it is also reported the average of the RMSFE across variables (Avg.).

Table 2: Forecasting Accuracy over the sample 1985-2007: Mean square forecast errors.

Horizon (quarters)	Variable	RW (MSFE)	AR-REC (RMSFE)	AR-ROL (RMSFE)	SV-AR (RMSFE)	TV-AR (RMSFE)	VAR-REC (RMSFE)	VAR-ROL (RMSFE)	SV-VAR (RMSFE)	TV-VAR (RMSFE)
	π	0.93	2.61	1.19	1.23	1.21	1.29	1.35	1.28	0.98
1	UR	0.05	2.80	1.16	1.05	1.07	1.09	1.17	0.99	1.03
	IR	0.27	3.64	1.08	0.85	0.83	0.87	1.02	0.77	0.83
	Avg.		3.02	1.14	1.05	1.04	1.08	1.18	1.01	0.94
	π	0.45	5.76	1.54	1.19	1.16	2.22	2.64	1.42	0.94
4	UR	0.37	3.00	1.15	0.82	0.82	0.97	1.23	0.77	0.89
	IR	2.09	1.74	1.17	0.78	0.81	0.78	1.20	0.74	0.81
	Avg.		3.50	1.29	0.93	0.93	1.32	1.69	0.97	0.87
	π	0.57	6.39	2.09	1.10	1.08	3.03	3.11	1.55	0.72
8	UR	1.33	1.72	0.86	0.61	0.56	0.42	0.72	0.38	0.57
	IR	5.16	1.53	1.05	0.68	0.74	0.67	1.20	0.63	0.75
	Avg.		3.21	1.33	0.80	0.79	1.37	1.68	0.85	0.69
	π	0.92	4.61	2.10	0.91	0.86	3.47	2.51	1.34	0.46
12	UR	2.25	1.22	0.72	0.48	0.43	0.35	0.73	0.27	0.48
	IR	7.69	1.44	0.89	0.55	0.63	0.70	1.13	0.51	0.61
	Avg.		2.42	1.24	0.65	0.64	1.51	1.46	0.71	0.51

The table reports the results relative to the forecasting accuracy using point forecasts. The variable we forecast are inflation (π_t), the unemployment rate (UR_t) and the interest rate (IR_t). The forecasting models are: RW - random walk; AR-REC - AR estimated recursively; AR-ROL - AR estimated with a rolling window; TV-VAR - time-varying VAR; VAR-REC - VAR estimated recursively; VAR-ROL - VAR estimated with a rolling window. For the random walk model we report the mean square forecast error (MSFE). For the other models we report the relative mean square forecast error (RMSFE), i.e. the ratio of the MSFE of a particular model to the MSFE of the naïve model. For each horizon it is also reported the average of the RMSFE across variables (Avg.).

Table 3: PITs Tests.

Horizon (quarters)	Variable	Moment	RW	AR-REC	AR-ROL	SV-AR	TV-AR	VAR-REC	VAR-ROL	SV-VAR	TV-VAR
	π	3rd	-0.05	0.44	0.65	0.37	0.50	0.59	-0.10	0.76	0.24
			(0.15)	(0.21)	(0.59)	(0.36)	(0.40)	(0.49)	(0.18)	(0.45)	(0.27)
		4th	0.94	1.50	7.73	4.00	4.25	6.15	1.25	4.89	2.68
			(0.34)	(0.53)	(1.67)	(0.88)	(1.03)	(1.32)	(0.39)	(1.25)	(0.61)
1	UR	3rd	0.55	0.86	0.70	0.79	1.01	0.76	0.32	0.42	0.16
			(0.31)	(0.38)	(0.58)	(0.38)	(0.42)	(0.52)	(0.26)	(0.28)	(0.29)
		4th	2.25	2.96	7.05	3.55	4.11	6.27	2.16	2.52	2.64
			(0.88)	(1.10)	(1.67)	(1.04)	(1.20)	(1.48)	(0.62)	(0.68)	(0.72)
	IR	3rd	0.18	0.08	-0.34	-0.24	-0.18	-0.07	0.18	0.31	0.15
			(0.61)	(0.61)	(0.61)	(0.40)	(0.38)	(0.51)	(0.59)	(0.40)	(0.45)
		4th	7.00	6.99	6.94	3.74	3.64	5.77	6.68	4.05	4.65
			(1.81)	(1.83)	(1.81)	(1.14)	(1.00)	(1.43)	(1.73)	(1.07)	(1.25)
	π	3rd	0.18	1.12	0.52	0.76	1.01	0.61	-0.17	1.09	0.53
			(0.51)	(0.69)	(0.93)	(0.72)	(0.72)	(0.89)	(0.55)	(0.86)	(0.50)
		4th	2.51	3.00	7.45	3.84	3.86	7.06	2.86	5.22	2.68
			(1.25)	(1.94)	(2.54)	(2.06)	(2.00)	(2.41)	(1.32)	(2.44)	(1.29)
4	UR	3rd	0.78	1.22	1.67	1.37	1.52	1.13	0.42	0.43	0.34
			(0.64)	(0.60)	(0.86)	(0.67)	(0.72)	(0.67)	(0.37)	(0.26)	(0.25)
		4th	4.00	3.04	5.83	3.93	4.18	4.61	1.46	1.52	1.30
			(1.65)	(1.58)	(2.47)	(1.87)	(2.00)	(1.82)	(1.00)	(0.54)	(0.53)
	IR	3rd	0.54	0.42	-0.17	0.30	0.61	0.15	0.35	1.31	1.04
			(1.00)	(0.89)	(0.80)	(0.53)	(0.42)	(0.82)	(0.71)	(0.95)	(0.76)
		4th	8.63	7.34	6.47	3.66	2.93	5.98	5.52	8.03	5.97
			(2.73)	(2.37)	(2.02)	(1.22)	(0.82)	(2.19)	(1.78)	(2.52)	(1.92)
	π	3rd	0.09	1.84	1.15	1.04	1.59	1.02	-0.21	1.26	0.97
			(0.75)	(1.12)	(1.41)	(0.92)	(1.02)	(1.33)	(1.16)	(1.11)	(0.77)
		4th	3.70	4.73	9.30	4.29	4.91	8.73	5.72	5.43	3.09
			(2.19)	(3.10)	(3.93)	(2.70)	(2.77)	(3.51)	(3.15)	(3.08)	(2.15)
8	UR	3rd	0.73	1.80	2.16	1.77	2.10	1.70	0.35	0.44	0.39
			(0.69)	(0.99)	(1.20)	(0.99)	(1.14)	(1.07)	(0.34)	(0.27)	(0.24)
		4th	3.80	4.51	7.00	4.95	5.65	5.10	1.11	1.04	0.89
			(1.68)	(2.52)	(3.23)	(2.62)	(3.03)	(2.85)	(0.69)	(0.46)	(0.37)
	IR	3rd	0.99	1.02	-0.01	0.89	1.03	0.09	0.57	1.21	1.50
			(1.49)	(1.30)	(1.38)	(1.05)	(0.80)	(1.58)	(1.14)	(1.59)	(1.16)
		4th	10.80	8.61	9.94	6.29	4.52	10.25	7.28	10.60	7.36
			(4.35)	(3.55)	(3.69)	(2.56)	(1.84)	(4.43)	(3.05)	(4.31)	(2.92)
	π	3rd	0.21	2.66	1.81	1.43	2.26	1.41	-0.38	1.79	1.10
			(0.52)	(1.61)	(1.91)	(1.06)	(1.41)	(1.67)	(1.57)	(1.44)	(0.75)
		4th	2.26	6.86	12.72	4.52	6.54	10.35	7.61	6.57	2.82
			(1.44)	(4.38)	(5.42)	(3.00)	(3.79)	(4.43)	(4.27)	(3.97)	(1.95)
12	UR	3rd	0.51	2.16	1.59	1.73	2.11	1.69	0.46	0.37	0.46
			(0.57)	(1.31)	(1.06)	(1.11)	(1.30)	(1.15)	(0.53)	(0.34)	(0.30)
		4th	2.72	5.54	5.47	4.98	5.87	5.23	1.78	1.13	0.96
			(1.29)	(3.32)	(2.44)	(2.72)	(3.26)	(3.04)	(1.17)	(0.64)	(0.49)
	IR	3rd	1.20	1.38	0.53	1.06	0.99	-0.07	1.06	1.28	1.63
			(1.51)	(1.56)	(1.74)	(1.36)	(0.84)	(1.67)	(1.28)	(1.99)	(1.42)
		4th	8.34	8.07	10.11	6.75	3.98	9.25	6.58	11.49	7.35
			(4.63)	(4.51)	(4.81)	(3.52)	(1.80)	(4.53)	(3.56)	(5.50)	(3.73)

The table reports the estimates of the third and fourth moments of the distribution of the inverse normal transformation of the PITs. HAC corrected standard deviations are reported in brackets.

Table 4: Forecasting Accuracy over the sample 1970-2007: log scores.

Horizon (quarters)	Variable	RW	AR-REC	AR-ROL	SV-AR	TV-AR	VAR-REC	VAR-ROL	SV-VAR
	π	0.20	0.19	0.13	0.07	0.08	0.20	0.11	0.14
		(0.04)	(0.04)	(0.06)	(0.03)	(0.04)	(0.04)	(0.05)	(0.05)
1	UR	0.14	0.08	0.20	0.03	0.02	0.10	0.05	-0.05
		(0.05)	(0.04)	(0.14)	(0.08)	(0.06)	(0.03)	(0.06)	(0.02)
	IR	1.38	1.50	0.41	-0.05	-0.06	0.05	0.93	-0.06
		(0.87)	(1.15)	(0.21)	(0.11)	(0.06)	(0.51)	(0.11)	(0.05)
	π	0.23	0.39	0.29	0.22	0.21	0.43	0.76	0.52
		(0.07)	(0.10)	(0.13)	(0.11)	(0.10)	(0.10)	(0.41)	(0.24)
4	UR	0.16	0.20	0.21	0.20	0.23	0.13	0.02	-0.08
		(0.08)	(0.09)	(0.13)	(0.18)	(0.18)	(0.05)	(0.11)	(0.03)
	IR	0.30	0.20	0.26	-0.04	-0.02	0.47	0.03	0.07
		(0.16)	(0.15)	(0.16)	(0.05)	(0.04)	(0.08)	(0.26)	(0.08)
	π	0.21	0.58	0.78	0.41	0.62	0.47	2.66	1.32
		(0.13)	(0.23)	(0.39)	(0.24)	(0.41)	(0.18)	(2.04)	(0.94)
8	UR	0.21	0.24	0.29	0.22	0.25	0.10	-0.14	-0.16
		(0.08)	(0.14)	(0.15)	(0.15)	(0.17)	(0.08)	(0.14)	(0.05)
	IR	0.19	0.01	0.25	-0.09	-0.03	0.69	-0.01	0.10
		(0.11)	(0.08)	(0.14)	(0.06)	(0.06)	(0.10)	(0.36)	(0.13)
	π	0.25	0.66	0.71	0.33	0.46	0.62	2.85	1.00
		(0.12)	(0.30)	(0.32)	(0.15)	(0.23)	(0.25)	(2.06)	(0.63)
12	UR	0.18	0.19	0.26	0.14	0.19	0.22	-0.18	-0.21
		(0.08)	(0.17)	(0.15)	(0.13)	(0.17)	(0.08)	(0.18)	(0.08)
	IR	0.22	0.07	0.17	-0.04	-0.11	0.48	-0.05	0.10
		(0.16)	(0.11)	(0.15)	(0.09)	(0.16)	(0.08)	(0.41)	(0.21)

The table reports the results relative to the forecasting accuracy using predictive densities. The variable we forecast are inflation (π_t), the unemployment rate (UR_t) and the interest rate (IR_t). The forecasting models are: RW - random walk; AR-REC - AR estimated recursively; AR-ROL - AR estimated with a rolling window; TV-VAR - time-varying VAR; VAR-REC - VAR estimated recursively; VAR-ROL - VAR estimated with a rolling window. For each model we report the sample average of the difference between the log score of the TV-VAR and the log score of that model. Standard deviation are reported in brackets.

Table 5: Forecasting Accuracy over the sample 1985-2007: log scores.

Horizon (quarters)	Variable	RW	AR-REC	AR-ROL	SV-AR	TV-AR	VAR-REC	VAR-ROL	SV-VAR
	π	0.27	0.23	0.05	0.06	0.05	0.25	0.11	0.09
		(0.05)	(0.05)	(0.06)	(0.03)	(0.03)	(0.05)	(0.05)	(0.03)
1	UR	0.22	0.13	0.02	-0.03	-0.03	0.00	0.12	-0.05
		(0.04)	(0.03)	(0.05)	(0.03)	(0.03)	(0.03)	(0.08)	(0.02)
	IR	0.30	0.29	0.16	0.10	0.02	0.16	0.24	-0.03
		(0.07)	(0.06)	(0.05)	(0.12)	(0.04)	(0.06)	(0.06)	(0.03)
	π	0.25	0.39	0.23	0.13	0.11	0.43	0.38	0.18
		(0.07)	(0.10)	(0.17)	(0.05)	(0.05)	(0.10)	(0.13)	(0.05)
4	UR	0.07	0.16	0.08	-0.02	0.00	0.18	0.09	-0.10
		(0.05)	(0.07)	(0.13)	(0.05)	(0.06)	(0.07)	(0.18)	(0.04)
	IR	0.05	-0.01	0.16	-0.06	-0.02	0.25	-0.06	-0.05
		(0.10)	(0.10)	(0.12)	(0.07)	(0.05)	(0.10)	(0.15)	(0.08)
	π	0.35	0.42	0.68	0.21	0.19	0.52	0.65	0.31
		(0.12)	(0.14)	(0.45)	(0.08)	(0.08)	(0.16)	(0.27)	(0.10)
8	UR	0.16	0.10	0.15	0.04	0.03	0.21	-0.10	-0.16
		(0.09)	(0.10)	(0.13)	(0.08)	(0.08)	(0.10)	(0.22)	(0.07)
	IR	0.01	-0.10	0.22	-0.06	0.01	0.60	-0.18	-0.08
		(0.11)	(0.10)	(0.15)	(0.07)	(0.07)	(0.13)	(0.34)	(0.14)
	π	0.43	0.38	0.79	0.25	0.21	0.63	0.81	0.36
		(0.17)	(0.18)	(0.45)	(0.11)	(0.11)	(0.25)	(0.36)	(0.15)
12	UR	0.22	0.00	0.15	-0.01	-0.04	0.26	-0.19	-0.24
		(0.11)	(0.10)	(0.13)	(0.08)	(0.09)	(0.11)	(0.30)	(0.11)
	IR	0.09	-0.07	0.20	-0.06	0.05	0.89	-0.12	-0.13
		(0.10)	(0.07)	(0.16)	(0.09)	(0.07)	(0.10)	(0.63)	(0.15)

The table reports the results relative to the forecasting accuracy using predictive densities. The variable we forecast are inflation (π_t), the unemployment rate (UR_t) and the interest rate (IR_t). The forecasting models are: RW - random walk; AR-REC - AR estimated recursively; AR-ROL - AR estimated with a rolling window; TV-VAR - time-varying VAR; VAR-REC - VAR estimated recursively; VAR-ROL - VAR estimated with a rolling window. For each model we report the sample average of the difference between the log score of the TV-VAR and the log score of that model. Standard deviation are reported in brackets.

Figures

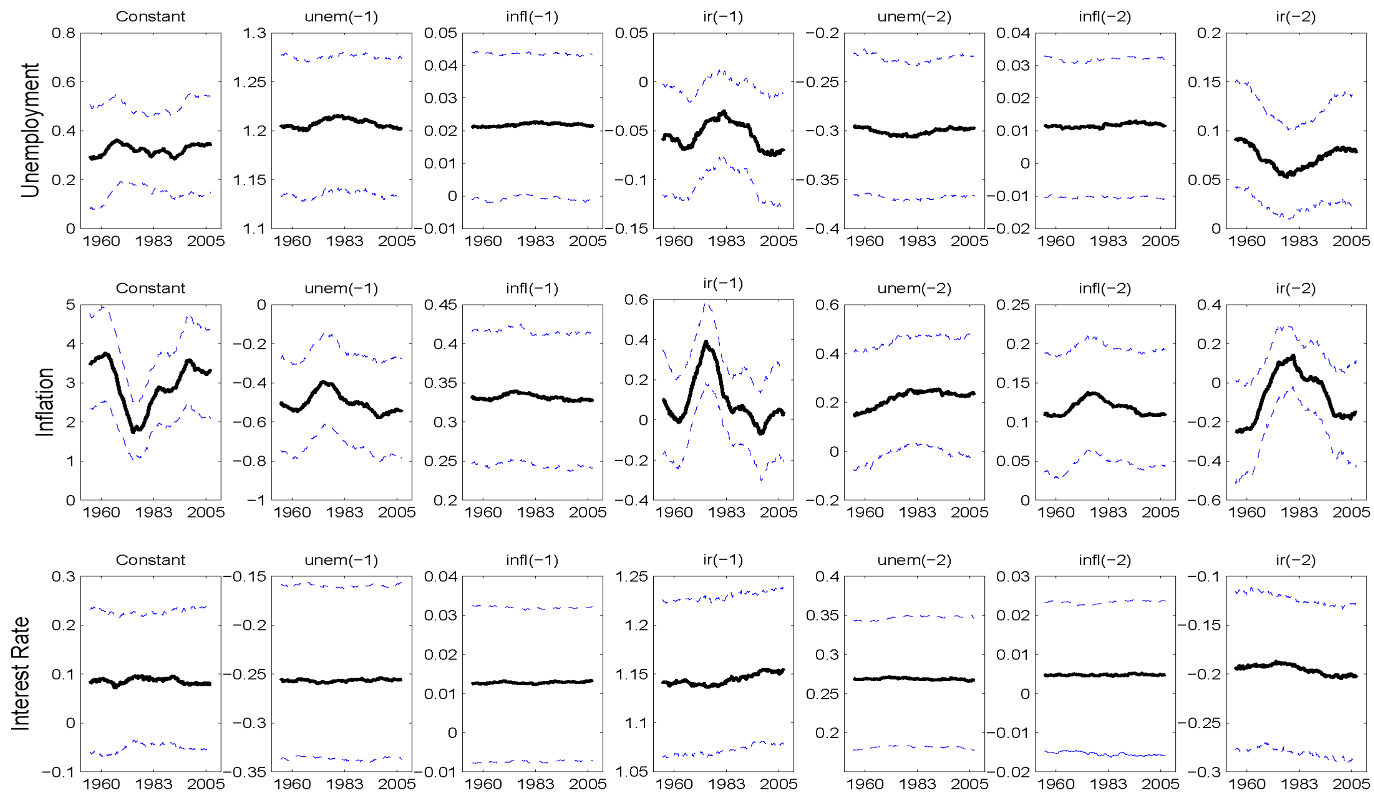

Figure 1: Time-varying parameters estimated with the TV-VAR model. Each row refers to a particular equation. In each column we report the coefficient associated to the indicated regressor. Solid lines are the posterior means, dashed lines are the 68% confidence bands.

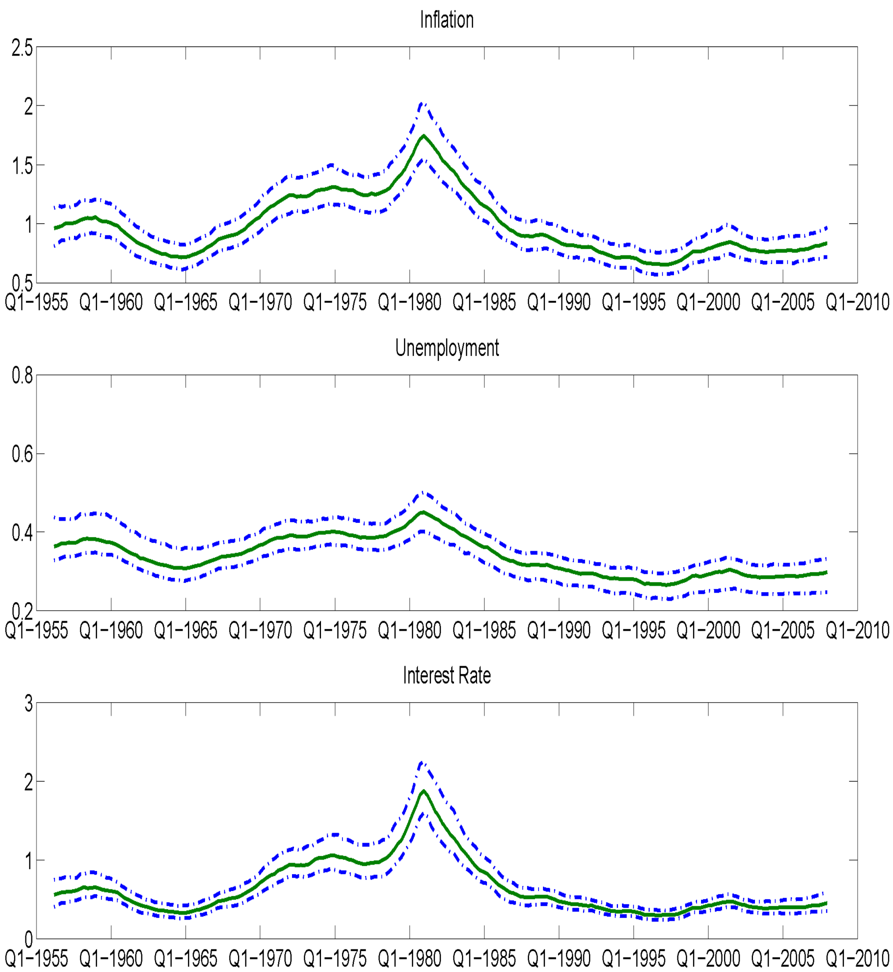

Figure 2: Stochastic Volatilities estimated with the TV-VAR model. Solid lines are the posterior means of the variances of the reduced form residuals, dashed lines are the 68% confidence bands.

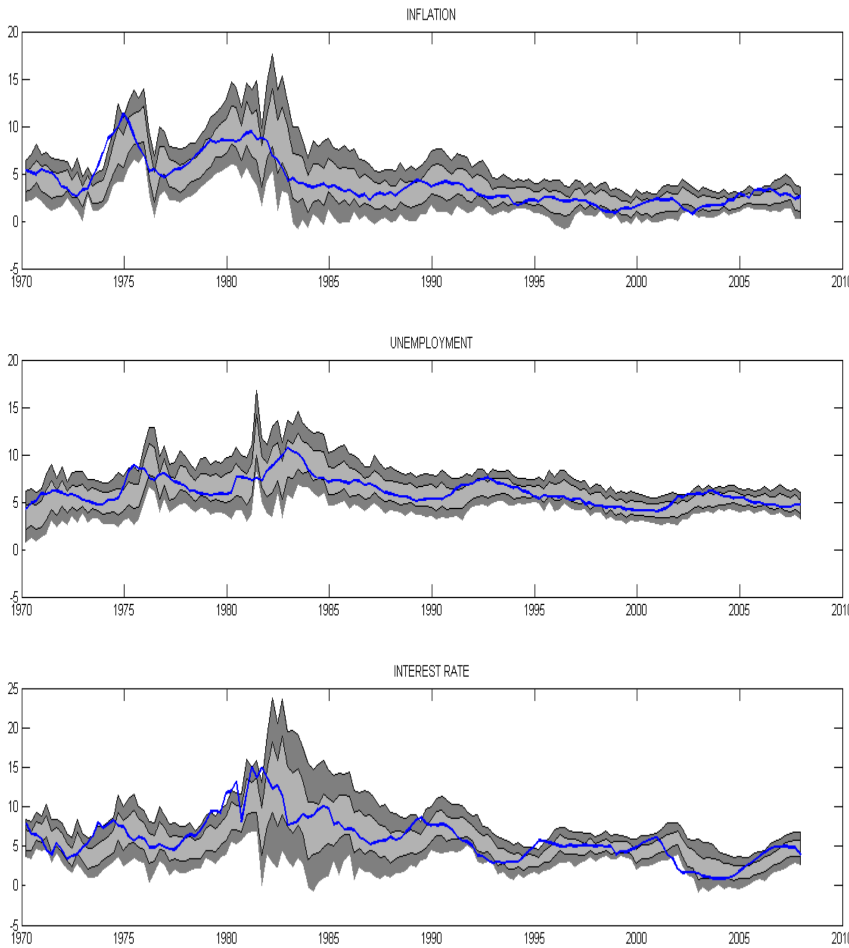

Figure 3: Two-years ahead forecast of inflation, the unemployment rate and the interest rate. Shaded areas represent the 68% and 90% confidence bands and the solid line is the true series.