


Working Paper

INTERNATIONAL MONETARY FUND



IMF Working Paper

Short-Term Forecasting: Projecting Italian GDP, One Quarter to Two Years Ahead

Matteo Iacoviello

IMF Working Paper

European I Department

**Short-Term Forecasting: Projecting
Italian GDP, One Quarter to Two Years Ahead**

Prepared by Matteo Iacoviello¹

Authorized for distribution by Thomas Krueger

August 2001

Abstract

The views expressed in this Working Paper are those of the author(s) and do not necessarily represent those of the IMF or IMF policy. Working Papers describe research in progress by the author(s) and are published to elicit comments and to further debate.

This paper presents a “bridge model” for short-run (one or two quarters ahead) forecasting of Italian GDP, relying on industrial production and survey indicators as key variables that can help in providing a real-time first GDP estimate. For a one- to two-year horizon, it formulates and estimates a Bayesian VAR (BVAR) model of the Italian economy. Both the “bridge” and the BVAR model can be of great help in supplementing traditional judgmental or structural econometric forecasts. Given their simplicity and their good forecasting power, the framework may be usefully extended to other variables as well as to other countries

JEL Classification Numbers: C11, C32, E32, E37

Keywords: Forecasting, Bayesian Vector Autoregressions, leading indicators

Author’s E-Mail Address: m.iacoviello@lse.ac.uk

¹ Author’s affiliation: London School of Economics. This paper was written while I was a summer intern at the Southern European I Division of the IMF, which I thank for hospitality and support. I am indebted to Messrs. Decressin, Krueger, and Monteforte for helpful comments and suggestions.

Contents	Page
I. Introduction	3
II. Forecasting Italian GDP Quarter by Quarter: The Bridge Model	4
A. Overview of the Model	4
B. Choice of the Model Variables	8
C. Rolling Regression Estimates	11
D. A Graphical Illustration of the Bridge Model	12
E. The Sensitivity of the Results to the Inclusions of a Trend	13
F. Two Quarters Ahead Forecasts	15
III. Vector Autoregressions	16
IV. Conclusions	21
References	22
Tables	
1. Estimation Results of the Bridge Model for the Italian Economy	10
2. Tests of the Accuracy of Forecast Models	15
3. One-Step Ahead Forecasts Given the Different Trend Specifications	15
4. Summary of the Prior for the Bayesian VAR Model	18
5. Comparison Between Coefficient Estimates, Bayesian Versus Standard VAR	19
Figures	
1. Autocorrelation Functions for Key Macro Variables	6
2. Cross-Correlations Between Italian GDP Growth and Key Leading Indicators	7
3. Cross-Correlations Between GDP Growth and Survey Variables	9
4. One-Step Ahead Forecasts, Model (II)	11
5. GDP Growth Forecasts for Different Values of Industrial Production Growth	12
6. GDP Growth Forecasts for Different Trend Specifications	14
7. GDP Growth Forecasts Two Quarters Ahead, BVAR Versus VAR Model	20
8. BVAR Projections Six Quarters Ahead: Implied Year-on-Year GDP Growth Rates	21

I. INTRODUCTION

Forecasting the economy is often a risky task. Although it represents one of the basic problems of statistical analysis, several difficulties arise in developing a model that successfully and consistently delivers good and reliable predictions. Data are limited, and often (especially the latest releases at each point in time) contaminated with measurement error. In addition, a variety of shocks hit the economy at every point in time. Lastly, the economic structure changes often, and with it the behavioural relationships that describe it.

This paper describes and presents applications for two widely used model based techniques² to obtain GDP forecasts over short horizons (one quarter to two years): the Indicator approach (Bridge Model) and the Econometric approach. Although the discussion will focus on the Italian GDP case, these approaches can be extended to other variables and other countries as well.

- (1) The indicator approach exploits early cyclical indicators, most of which are available on a monthly basis, in order to provide real time estimates of quarterly changes in GDP with an advance of a few months prior to the first official release from the National Statistical Office. The approach, first developed by Klein and Sojo (1989) for the U.S. economy and later applied to Italy by Parigi and Schlitzer (1995), consists of estimating a functional relationship between the variable that one wishes to estimate (say, quarterly GDP) and others that contain useful reference information for its short-term movements, such as industrial production and business cycle survey indicators. This approach is called “bridge” model (since it normally links monthly variables, such as industrial production, with quarterly ones, such as GDP). Its rationale is to adhere as closely as possible to prevailing practice used in the actual construction of the Quarterly National Accounts (QNA), which uses a (sizeable) number of indicators to provide first estimates in GDP movements. In other words,

² Of course, another common type is the so-called *judgement-based forecast*. This type of forecast is predominantly the result of a particular forecaster’s skill at reading the economic tea leaves, interpreting anecdotal evidence, and his or her experience at spotting empirical regularities in the economy. There are a number of obvious shortcomings in this, as highlighted by Robertson and Tallman (1999): (i) their accuracy can be evaluated only after a track record is established; (ii) given the element of subjectivity in such forecasts, changes in the forecasting staff will affect the accuracy of these forecasts; (iii) they are impossible to replicate or validate by independent forecasters; (iv) they normally do not come with a probabilistic assessment of a range of alternative outcomes; (v) they are deemed unable to predict recessions or strong booms.

Of course, the distinction between model-based and judgement-based forecasts cannot be pushed too far. “Successful model specifications also depend heavily on the skill and ingenuity of particular individuals. No model can be left on automatic pilot for long,” (Robertson and Tallman, 1999, page 21).

the bridge equations are not theoretical, but just practical expressions of what is presumably done, and can be justified in terms of simple economic or accounting (i.e., not behavioural or structural) considerations. Section II presents such a bridge model for the Italian economy.

- (2) Rather than focusing on a single equation, the “econometric-based” approach traditionally assumes some structure in the economy, either in the form of behavioural relationships describing the linkages across some key macroeconomic variables, or in a more atheoretical way, such as in the case of vector autoregressions. This paper will mainly focus on vector autoregressions to describe how GDP forecasts can be made over a one- to two-year horizon, partly because of their relative simplicity, partly because in the last decade VAR models have become a widely used tool for forecasting macroeconomic time series. In particular, considerable attention has been devoted to Bayesian VARs. Section III presents a specific example with applications to Italian data.

The above taxonomy does not explicitly consider *ARIMA models* and *leading indicators models* as other widely used forecasting tools. ARIMA models can be seen as a less general case of a vector autoregression model; therefore, they are not treated separately. Leading indicators are variables seen as informative about future movements in the variable of interest: models based on such indicators have been seen partly as a reaction to perceived failures by macro-econometric systems (Emerson and Hendry, 1996), and have been used in a variety of ways, also as a part of a wider VAR. Yet the fact that leading indicators are frequently altered casts some doubt on the post-sample performance of models based on these indicators.

II. FORECASTING ITALIAN GDP QUARTER BY QUARTER: THE BRIDGE MODEL

A. Overview of the Model

The bridge model is mainly based on two sets of variables: monthly indicators (I_{it}), most of which are available with very short delay,³ and quarterly series, mostly coming from the

³ Exchange rate and interest rate data are available in real time. The industrial production index is released from the Italian Statistical Office (ISTAT) on a monthly basis with a delay of about 45 days (i.e., industrial production data for month x are released mid-month $x+2$). This compares with the 80 days delay for the first detailed QNA estimates published by ISTAT (that break GDP down into its demand components); from November 2000, ISTAT also publishes a preliminary GDP first estimate 45 days after the end of the quarter. This first estimate is obtained using “statistical techniques of integration” (ISTAT, *Stima Preliminare del PIL*, third-quarter 2000) and normally subject to bigger revisions than the following QNA estimates. Another potentially important coincident indicator for Italian GDP, quarterly German GDP, is released with a delay of 60 days. There are on average 20 days in a quarter in which one might use current quarter information on German GDP in order to estimate the Italian GDP figure.

National Accounts (N_{it}). The high serial correlation in the monthly series provides a basis for extrapolation whenever a monthly series does not cover the entire quarter. As observed by Klein and Park (1993) for the U.S. case, “most, if not all, major economic variables can be projected from fitted ARIMA equations quite well over six-months horizons.” Once the predicted values of I_{it} are obtained, one can construct a variable with a quarterly frequency, for instance by taking three-months quarterly averages.

Quarterly indicators can then be related to the (seasonally adjusted) GDP (Y) figure that one wishes to estimate according to a simple functional relationship (the bridge equation), taking the general form (where Δ indicates the first difference operator):

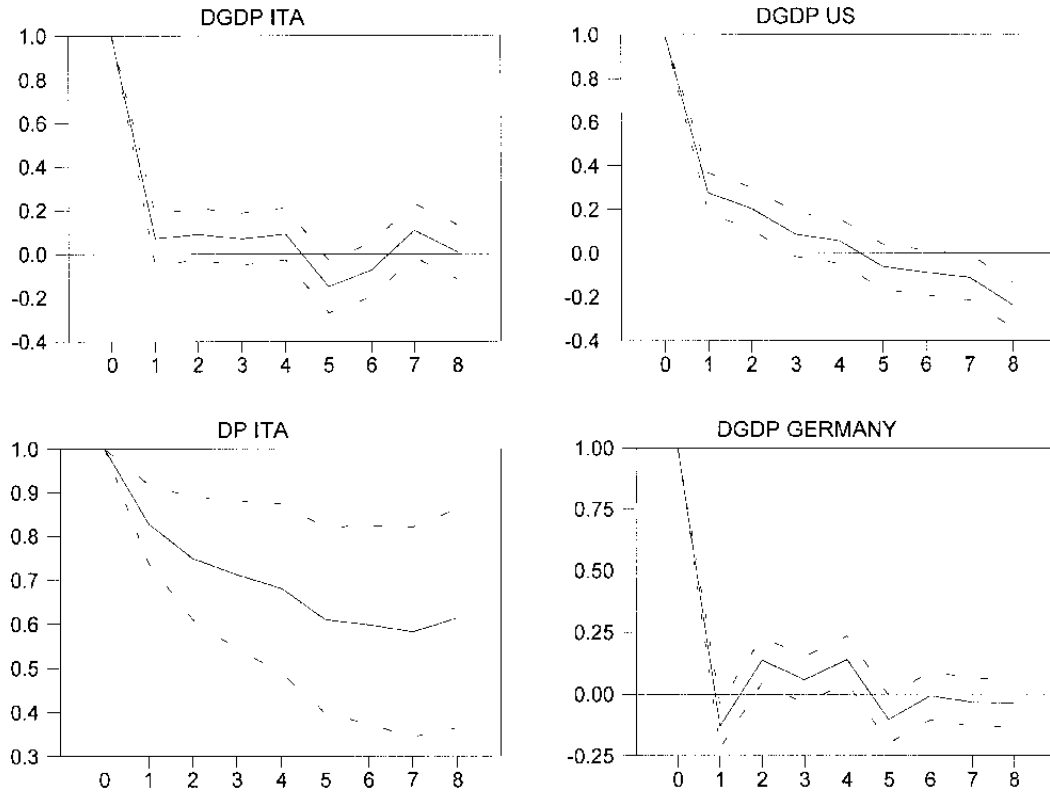
$$\Delta Y_t = f(I_t, I_{t-1}, Y_{t-1}, X_t) + e_t \quad (1)$$

where X denotes a vector of predetermined variables (such as trend terms). In contrast with the typical ARIMA model, equation (1) puts additional emphasis on coincident indicators I_t that are released prior to the GDP figure rather than on GDP's past history. That these indicators are especially important for the Italian case can be justified as time series for Italian quarterly GDP growth appear to follow an erratic pattern over the last 20 years. This is striking not only by comparing Italian GDP with its U.S. counterpart, but also by comparing GDP to Italian CPI inflation. As shown in Figure 1, the autocorrelation function (ACF) for Italy's GDP (quarter-on-quarter) growth is positively but insignificantly autocorrelated over short horizons. This stands in sharp contrast with the ACF for U.S. GDP (Cogley and Nason, 1995): at lags of one and two quarters, sample autocorrelations are positive and statistically significant, thus offering a good basis for some fruitful extrapolation.⁴

⁴ It is interesting to note that a persistency pattern holds instead for *yearly* Italian GDP growth, which shows a serial correlation of 41 percent year on year.

Figure 1: Autocorrelation Functions for Key Macro Variables

AUTOCORRELATIONS FOR ITALY GDP AND INFLATION AND U.S. AND GERMANY GDP AND INFLATION
(Quarter-on-quarter percentage changes)




Note: The horizontal axis measures quarters; the confidence bands are 95 percent. DGDP is quarter-on-quarter GDP growth, DP is quarter-on-quarter CPI inflation.

It is therefore apparent that forecasting Italian GDP over the short run relying only on its own past history might prove quite a hard task, unless one relies on timely coincident indicators that are significantly correlated with it over high frequencies. Figure 2 provides an insight into these indicators by showing sample cross-correlations between quarterly Italian GDP growth and its German counterpart (ΔY_{GER}), the industrial production index (ΔIP , quarterly changes) and a (first differenced) Coincident Survey Indicator (ΔCSI) obtained as simple

average of survey indicators provided by the EU Business Survey⁵ (sum of production expectations months ahead index, *PEMA*, selling price expectations months ahead, *SPEMA*, building construction index, *BCI*, assessment of stock of finished products, expressed with inverse sign, *ASFP*).⁶ All the three variables display a strong significant contemporaneous correlation with Italian GDP (in the 50 percent range).

While the information contained in German GDP cannot be systematically exploited for forecasting purposes since its release occurs only 20 days before the Italian figure, industrial production and survey indicators are available for the entire quarter more than one month in advance of GDP, for two-thirds of the indicators with two months advance, and so on. It is on these indicators that we rely in the choice of our model variables.

Figure 2: Cross-Correlations Between Italian GDP Growth and Key Leading Indicators



Note: Quarterly cross-correlations between Italian quarter-on-quarter (qoq) GDP growth and qoq Industrial Production growth (ΔIP), qoq German GDP growth ($\Delta YGER$), first difference of the Coincident Survey Indicator (ΔCSI).

⁵ The EU monthly business survey comprises qualitative questions aimed at obtaining information on the current situation and on the short-term (three–four months) trend of the main firm variables (such as order-books, production, finished products stocks, selling prices) as well as on expectations on the general economic trend.

⁶ The above-mentioned indicators outperform the consumer confidence indicator in terms of its cyclical properties and its forecasting power.

B. Choice of the Model Variables

An estimate/forecast of GDP can in principle be obtained in two ways:

- (1) Predicting and summing the estimates of all the components of internal demand minus imports, as done for instance in Parigi and Schlitzer (1995);
- (2) Exploiting the information coming from the supply side, relying on the prompt availability of the *IP* index and the survey data. Parigi and Schlitzer (1995) show that a “supply” model can do a good job in providing estimates for GDP one- and two- quarters ahead.⁷ Their specification is cast in the form of an error-correction model with the inclusion of a (linear and quadratic) trend in order to proxy the trend in the output of the services sector. Let Y and IP denote the log of seasonally adjusted real GDP and the Industrial Production Index⁸ (the latter is the average across months in the quarter). Let also ΔY and ΔIP denote quarter on quarter changes of Y and IP . The specification by Parigi and Schlitzer (1995) is:

$$\Delta Y_t = b_0 + b_1 Y_{t-1} + b_2 IP_{t-1} + b_3 \Delta IP_t + b_5 TREND + b_6 TREND^2 \quad (PS)$$

I take the (PS) specification as a natural benchmark in choosing the model variables, and I consider the survey indicators as a natural way of refining it. The final choice includes the variables discussed below.

- (1) The role of the **industrial production index** is not in dispute as a short-term indicator and predictor of economic activity. In a horse race with several macroeconomic variables, industrial production is by far the most highly (contemporaneous) correlated variable with GDP growth. Since it is conceivable to assume that industrial production and GDP share a common trend, I include in the specification (cast in the form of an error correction model, with GDP changes as the left hand side variable) lagged levels of both.
- (2) An increasing amount of emphasis is being placed in recent times on indicators coming from survey data. The experts’ commentary on the economy’s near-term outlook changes from day to day in response to the release of such additional (typically monthly) data. However, such business cycle survey indicators are often highly correlated with each other, thus creating multicollinearity problems. One may wish to eliminate their noise while still exploiting the wealth of information contained in them either by forming a principal component of them or by simply averaging


⁷ That is, the Root Mean Square Error (RMSE) of the supply model is lower than the corresponding RMSE of the demand side model. See footnote 11 for a definition.

⁸ Whenever Industrial Production is not available for the entire quarter, the monthly series is extrapolated with an $ARMA(p,q)$ model of the form $(p=\|1,2,3,6,12\|,q=\|1\|)$.

across indicators (or using some other form of aggregation).⁹ The latter strategy has been adopted here, leading to the construction of a **coincident survey indicator (CSI)**, constructed as the sum of:

- Production expectations three months ahead index, *PEMA*
- Selling price expectations three months ahead, *SPEMA*
- Building construction index, *BCI*
- Assessment of stock of finished products, expressed with inverse sign, *ASFP*.

Figure 3: Cross-Correlation Between GDP Growth and Survey Variables



Note: The figure shows the quarterly cross-correlation between Italian GDP growth and the (first differenced) coincident survey indicator (DCSI) and leading survey indicator (DCLI). Positive (negative) entries on the horizontal axis correspond to the survey indicators leading (lagging) GDP.

- (3) In addition, I also use a **Leading Survey Indicator (CLI)** from the survey variable **Expected Business Situation**. The cross-correlogram of the first difference of this variable with GDP growth
- (3) Figure 3 shows that it has a strong positive correlation with GDP one period ahead. Similar results do not hold for the survey variable *PEMA* (production expectations months ahead), which has a strong correlation with GDP not only one quarter ahead

⁹ Forni, Hallin, Lippi and Reichlin (2000) reconcile dynamic principal components analysis with dynamic factor analysis in order to extract indicators from a large panel of economic variables (many variables for many countries). The procedure is used to estimate coincident and leading indicators for the euro area.

but also in the current quarter, and does not appear to improve fit and forecasting ability of the model when included in the leading indicator.

After the usual specification searches, the proposed model is therefore:

$$\Delta Y_t = a_1 + a_2 Y_{t-1} + a_3 \Delta IP_t + a_4 IP_{t-1} + a_5 TREND + a_6 \Delta CSI_t + a_7 \Delta CLI_{t-1} \quad (I1)$$

The equation has been estimated over the period 1985Q2–2000Q2 (the initial period reflects constrained data availability for the EU survey indicators) and the results are shown in Table 1. The coefficient estimates appear quite satisfactory and in line with expectations. The standard error of the regression is 0.4 percent and the coefficients are all statistically significant (the trend variable is the only exception: I explore this issue further below). The regression residuals pass the usual tests for stationarity and normality.

Table 1: Estimation Results of the Bridge Model for the Italian Economy

DEPENDENT VARIABLE: GDP Quarterly Growth, $\Delta Y(t)$				
Quarterly Data From 1985Q2 To 2000Q2				
Degrees of freedom	54			
R-squared	0.5967			
Adjusted R-squared	0.5519			
Standard Error of Estimate	0.0040			
Durbin-Watson Statistic	2.4072			
	VARIABLE	COEFFICIENT	STD ERROR	T-STAT
1	CONSTANT	1.5954	0.5791	2.7547
2	Y(t-1)	-0.1559	0.0526	-2.9658
3	$\Delta IP(t)$	0.1946	0.0412	4.7218
4	IP(t-1)	0.0896	0.0267	3.3537
5	TREND	0.0235	0.0162	1.4518
6	$\Delta CSI(t)$	0.0016	0.0006	2.7958
7	$\Delta CLI(t-1)$	0.0012	0.0005	2.3817
Analysis of the Residuals				
Ljung-Box Chi-Squared Test for Serial Correlation				
LB(24)	Test statistic (p-value)	34.4232 (.0593)		
F-test for Autoregressive Conditional Heteroskedasticity				
ARCH(24)	Test statistic (p-value)	0.7182 (.764)		
Jarque-Bera Normality Test				
Chi-Squ(2)	Test statistic (p-value)	0.7552 (.6855)		


Note: Specification (I1) with a linear trend and survey indicators.

One could have of course estimated a cointegration vector between GDP and Industrial Production and included the lagged cointegration vector in the regressions without entering separately the two variables. Cointegration tests between the two variables point to one cointegration vector of the simple form $[-1 \ 1]$. Working with such a cointegration vector does not significantly affect the results of the specification or its forecasting performance, to which I now turn.

C. Rolling Regression Estimates

In order to verify the model performance under “real-time” conditions, I compute one-step ahead GDP forecasts using rolling estimation technique (starting in 1996Q2). The results from this exercise are shown in Figure 4.

Figure 4: One-Step Ahead Forecasts, Model (I1)



Note: Comparison of one-step ahead forecasts and original data for Italian annualized GDP quarter-on-quarter changes. The specification is (I1) and the forecast are generated via rolling regression estimates.

The model performs reasonably well in the years 1999 and 2000, and captures in sign, if not in magnitude, some of the recent turning points of the Italian business cycle. Of course, it must be stressed that the one-step ahead forecast crucially depends on the latest data releases, which are frequently revised up to two years after the initial publication. However, given the


often erratic changes in GDP growth over the last five years, the performance of the model is remarkably good, with a Root Mean Square Error¹⁰ (RMSE) of 0.44 percent.

D. A Graphical Illustration of the Bridge Model

It is possible to illustrate how the model in (II) works by means of a simple graph. At each point in time, the GDP forecast depends on two sets of indicators: (1) past history of GDP, Industrial Production Index and Survey Leading Indicator, a linear trend term; (2) two current quarter indicators, Coincident Survey Indicator on the one hand and industrial production changes on the other. As forecasts are being prepared over the quarter, the GDP estimate varies depending on the value that the indicators in (2) assume.

Figure 5 plots the implied GDP quarter-on-quarter growth forecast as a function of current quarter industrial production growth, given the coefficient estimates reported in Table 1 (short-run elasticity of GDP to industrial production of 0.1946 and semielasticity of GDP to the coincident survey indicator of 0.16).

Figure 5: GDP Growth Forecast for Different Values of Industrial Production Growth



¹⁰ See footnote 11 for a definition of the RMSE.

Note: For given values of the lagged variables, the thick line shows the implied change in GDP growth (ΔY) forecasts for different values of industrial production index changes (ΔIP), given constant survey indicator values ($\Delta CSI=0$ in the quarter). The estimated model is (I1). The thinner lines above and below correspond to ΔY forecast as a function of ΔIP if the survey indicator variable is respectively 1 unit up or down on the previous quarter.

If the coincident survey indicator stays constant on the previous quarter, the estimated linear relationship between ΔIP and ΔGDP corresponds to the continuous boldfaced line in the figure (its slope is 0.1946). If the survey indicator increases (decreases) by 1 unit (1 standard error), the GDP forecast is revised upwards (downwards) by 0.16 percent, as the two dashed lines indicate.


E. The Sensitivity of the Results to the Inclusion of a Trend

The inclusion of the trend terms in the Parigi and Schlitzer (1995) specification—equation (PS)—is justified with the need to proxy the trend in the output of the service sector, which is not well accounted for by changes in the industrial production index. It might be the case that the trend plays an important role in driving the forecast results. This section further explores this issue.

Figure 6 shows one-step ahead forecasts for our model (I) depending on which specification for the trend one assumes. Four different combinations were tried: one of them represents the benchmark model (I1) described in the previous paragraphs:

- (0) No trend (I0 model);
- (1) Linear trend (I1 model);
- (2) Linear and quadratic trend (I2);
- (3) Linear and cubic trend (I3)

Figure 6: GDP Growth Forecasts for Different Trend Specifications



Note: GDP growth (solid line) and GDP growth forecasts (dashed lines) for the same quarter generated one quarter before using the model in (I1) with alternative trend specifications.

Over the period 1997Q2–2000Q2, it appears that the linear trend and the no trend specification provide higher GDP forecasts than the specifications including a quadratic trend and a cubic trend.

Table 2 presents the values of two traditional measures of goodness of forecast with respect to the final value for the rolling regression period from 1997Q3 to 2000Q2, mean forecast error and the root mean square error.¹¹ The evidence from the third column is that the forecasts from the four different models are not biased.

¹¹ Let R be the actual realisation of the GDP, F be its one step ahead forecast, and T the time horizon for the forecasts. The mean forecast error (MFE) is $(1/T) \sum (F-R)$, whereas the root mean square error (RMSE) is the square root of $(1/T) \sum (F-R)^2$. The mean forecast error is a measure of unbiasedness, whereas the RMSE is a measure of efficiency. Unbiasedness is a necessary condition for efficiency.

Table 2: Tests of the Accuracy of Forecast Models

UNBIASEDNESS TEST			GOODNESS OF FORECAST TEST	
Series	Mean forecast error	Significance level of zero mean forecast error test	Root mean square error	Obs.
RESID_I0	-0.01%	93.86%	0.29%	12
RESID_I1	0.07%	42.83%	0.30%	12
RESID_I2	-0.14%	14.05%	0.30%	12
RESID_I3	-0.15%	12.31%	0.31%	12

Note: Tests of the accuracy of forecast models, for different trend specifications. The mean forecast error is the average of the difference between forecast and realised events. The root mean square error RMSE is the square root of the variance of the forecast error.

Table 3: One-Step Ahead Forecasts Given the Different Trend Specifications

ENTRY	ΔY_{I0}	ΔY_{I1}	ΔY_{I2}	ΔY_{I3}	Average
2000Q1	0.62%	0.66%	0.52%	0.52%	0.58%
2000Q2	0.66%	0.67%	0.57%	0.58%	0.62%
2000Q3	0.41%	0.46%	0.33%	0.33%	0.38%
Average	0.56%	0.60%	0.47%	0.48%	0.53%

Note: Specification (I0) refers to GDP growth forecast made using the model (I) with no trend, (I1) to the model (I) with linear trend, (I2) with quadratic trend, (I3) with cubic trend.

Overall, it seems that the specification adopted is not very sensible to the inclusion of the trend terms. For instance, in 2000Q1, the GDP growth forecast would have ranged from 0.52 percent to 0.66 percent (see first numbered row in Table 3).

F. Two Quarters Ahead Forecasts

Looking at equations such as (I) reveals that such a model is capable of forecasting in principle only one step ahead. Still, one can push the bridge model to make forecasts two quarters ahead. One way to do so is to specify an equation that helps forecasting each exogenous variable. A straightforward approach is to fit an ARIMA for each of the exogenous contemporaneous regressors and then to project the estimated equation two

quarters ahead. In our case, the variables that one needs to project are IP and CSI ¹² only, given that it is the lagged level of the leading indicator that enters the specification.

Forecasts from bridge models beyond the two quarters horizon tend to build up error and wander away from observed data series quite quickly. If one wants to make such longer-horizon forecasts, one might want to turn toward other approaches.

III. VECTOR AUTOREGRESSIONS

The 1960s saw the large Keynesian macroeconomic models as the natural tool to meet the demand for macroeconomic forecasts. At the end of the 1960s there were several models of the U.S. economy each taking its roots in the original Klein (1950) 16 equations prototype. Since the 1970s, however, such models have been criticised on several grounds: the profession attacked the logical foundations of these models, in particular that they were based on too many exclusion restrictions to identify the parameters, that they were unsuitable for policy exercises, and that they were not based on individual optimising behaviour. As a consequence, the 1980s saw a shift toward joint modelling of the time series behaviour of all the variables in a system, in an unrestricted way. Sims (1980), in particular, advocated the use of vector autoregressions (VAR).

Vector autoregressions offer a simple way to generate forecasts. Consider the following VAR specification:

$$y_t = Ay_{t-1} + u_t \quad (2)$$

where y_t is a vector of variables and u_t is a multivariate error term with the usual properties. It follows that the n step ahead forecast for y_t is simply $A^n y_t$. Analogous arguments apply to the multi-lags case.

Because of the extreme simplicity, a VAR would seem unlikely to produce accurate forecasts. Litterman (1986), however, showed that a small VAR with six to eight variables produced better estimates (in the four to six quarters ahead range) than forecasts coming from small macroeconomic models or from commercial forecasting services. However, forecasts from VAR models easily suffer from overparametrization of the model. One way to overcome this problem would be to impose explicit restrictions by putting some groups of coefficients to zero. This would reduce the number of parameters to estimate, but would violate the approach which is largely atheoretical. An alternative is a Bayesian VAR (BVAR).

In a Bayesian VAR, one specifies loose restrictions on the coefficients, rather than hard shape or exclusion restrictions. The method assumes that coefficients on higher lags are more likely to be close to zero than coefficients on shorter lags. However, the data allow overriding the assumption, if evidence about a parameter is strong enough.

¹² Monthly variables are all projected with an ARMA of the form AR(1,2,3,6,12), MA(1). Details are available from the author upon request.

The Bayesian VAR requires one to specify means and standard deviations of the variables' prior distributions. In particular, the widely used Bayesian VAR uses the so-called Minnesota prior (see Litterman, 1986). The Minnesota prior is based on three main sets of restrictions (see also Bikker, 1998, for a clear summary):

- (1) *Overall restriction*, which is based on the assumption that the economic variable that one wishes to forecast follows a random walk around a deterministic component (plus constant, dummies and/or a time trend). That is,

$$Y_t = Y_{t-1} + u_t \quad (3)$$

This restriction is imposed with an overall tightness parameter γ (which is made proportional to the standard deviation of the lag coefficients higher than (1), which can also be varied for each individual equation in the VAR.

- (2) *Higher order lags* contain less information than lower order ones.
- (3) *Cross lags restrictions*: for each equation, own lags contain more information than lags of other variables.

Restrictions (2) and (3) are imposed by restricting the standard deviation function for the prior distribution to the form:

$$S(i, j, l) = \frac{\{\gamma g(l) f(i, j)\} s_i}{s_j} \quad (4)$$

where $S(i, j, l)$ denotes the standard deviation of the prior distribution for lag l of variable j in equation i . In this equation, s_i denotes the standard deviation of a univariate regression on equation i , and simply reflects different scales in the variables. The part in braces consists of the product of three terms, reflecting the three items above: γ is the overall tightness parameter in (4); $g(l)$ is the tightness of lag l relative to lag 1; $f(i, j)$ is the relative weight on variable j in the equation for i relative to i itself.

The meaning of this formula can be understood by looking first at the special case of the restriction (1). For the first lag, $l=1$, it is assumed that the prior distribution of the own variable lag is 1, that is $f(i, i)=1$. In addition higher order lags are in principle less and less informative, that is the tightness on lag l increases with l , at a harmonic rate dictated by a decay function g . Finally we specify the function $f(i, j)$, that is, the tightness on variable j in equation i relative to variable i .

In specifying the VAR, GDP was expressed in levels. This is in line with what was suggested by Sims, Stock, and Watson (1990) who recommend against differencing when one suspects cointegration among the variables. Following Litterman (1986), the prior distribution (i.e., the tightness of each of the restrictions above) was chosen in order to minimise the root mean square error of the model.

For the VAR specification, drawing on the various VAR and leading indicator papers on Italy (e.g., Gaiotti, 1999; Bikker, 1998; and Altissimo, Marchetti, and Oneto, 2000), the following variables were chosen:

- (1) Y (log of real GDP);

- (2) *CHOU* (log of real household consumption);
- (3) *REALRATE* (T-Bill rate minus annualised consumer price inflation);
- (4) ΔCSI (quarterly changes in coincident survey indicator);
- (5) *LREER* (log of real exchange rate);
- (6) ΔP (annualised quarterly changes in consumer price index);
- (7) *YGER* (log of German GDP).

Many of the choice variables are standard in any VAR model focused on estimating output growth. The inclusion of ΔCSI and *LREER* and *YGER* reflects, on the one hand, the desire to incorporate in the specification the EU survey variables, on the other, the importance of introducing into the specification foreign sector variables. Altissimo, Marchetti, and Oneto (1999), for instance, report that German output and its index of industrial production are strongly correlated with Italian GDP, and that they lead by approximately one quarter.

For the Bayesian VAR four lags were chosen. The model was initially estimated over 1985Q2–1995Q4. Simulated out of sample forecasts were obtained starting from that final date until 2000Q2, by extending the estimation period one quarter a time.¹³

The parameters of the BVAR model, including the coefficients of the function $f(i,j)$, are presented in Table 4. They were all chosen through a combination of grid search and application of prior economic theory.

Table 4: Summary of the Prior for the Bayesian VAR Model

Overall Tightness Parameter: 0.8 Harmonic Lag Decay with Parameter 2							
Standard Deviations as Fraction of Tightness and Prior Means Listed Under the Dependent Variable							
	Y	CHOU	REALRATE	ΔCSI	LREER	DP	YGER
Y	1	0.05	0.17	0.45	0.2	0.4	0.04
CHOU	1	0.4	0.02	0.35	0.01	0.05	0.04
REALRATE	1	0.05	0.4	0.05	0.01	0.25	0.04
ΔCSI	0.33	0.05	0	0.5	0.01	0.05	0.04
LREER	0.5	0.01	0.05	0.2	0.4	0.25	0.04
DP	0.4	0.05	0.25	0.05	0.01	0.6	0.04
YGER	0.5	0.01	0.05	0.05	0.01	0.05	0.4
Prior Means	1	1	1	1	1	1	1

¹³ The coefficient estimates were updated using a Kalman filter algorithm: that is, at each point in time the model was estimated through some period before the end of the data set and its performance was evaluated comparing actual and forecast data.

For instance, in the equation for *Y* (first numbered column) a relatively high weight is given to the coefficients on the real interest rate and on previous consumption; in the equation for *YGER* (last column) it is essentially assumed that German GDP follows a univariate autoregressive process (coefficients on variables other than *YGER* are forced to close to zero). It has already been shown that one of the assumptions of a BVAR is that own lags matter more and lags further away carry less information.

Table 4 also shows that a lower coefficient is assigned to inflation than to consumption in the GDP equation. For instance, the difference in the two GDP equations obtained via VAR and BVAR shows that while the coefficients on the first lag of *GDP* and *CHOU* are roughly of the same magnitude in the two models (Table 5), inflation and real interest rates carry much less weight in the BVAR compared to the simple VAR. That is why the BVAR model would have predicted less GDP growth for 1998 compared to the VAR in a period of falling real rates.

Table 5: Comparison Between Coefficient Estimates, Bayesian Versus Standard VAR

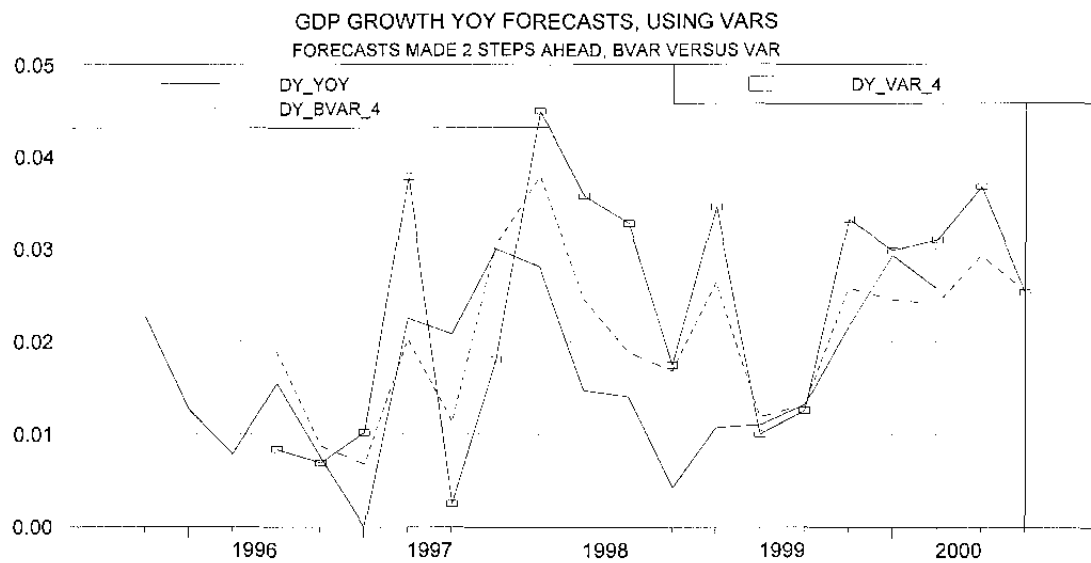
Dependent Variable: Y							
	Variable (lags in braces)	Coeff BVAR	Coeff VAR		Variable	Coeff BVAR	Coeff VAR
1	Y(1)	0.6095	0.5991	16	Δ CSI (4)	0.0000	-0.0003
2	Y(2)	0.1280	0.0725	17	LREER(1)	-0.0248	-0.0437
3	Y(3)	0.0343	0.3665	18	LREER(2)	0.0006	0.0483
4	Y(4)	0.0118	0.0440	19	LREER(3)	-0.0020	-0.0299
5	CHOU(1)	0.2854	0.3217	20	LREER(4)	-0.0008	0.0198
6	CHOU(2)	-0.0648	-0.2783	21	DP(1)	0.0096	-0.2661
7	CHOU(3)	-0.0206	0.0418	22	DP(2)	0.0274	0.2394
8	CHOU(4)	-0.0051	-0.1268	23	DP(3)	-0.0014	0.4876
9	REALRATE(1)	-0.0536	-0.1336	24	DP(4)	0.0078	0.0623
10	REALRATE(2)	-0.0195	-0.0106	25	YGER(1)	-0.0673	-0.1614
11	REALRATE(3)	0.0041	0.1098	26	YGER(2)	0.0092	0.0936
12	REALRATE(4)	-0.0075	-0.1008	27	YGER(3)	-0.0078	-0.1888
13	Δ CSI (1)	0.0011	0.0023	28	YGER(4)	0.0008	0.2104
14	Δ CSI (2)	0.0001	0.0012	29	Constant	0.5386	-0.4845
15	Δ CSI (3)	0.0000	0.0005	30	TREND	0.0160	-0.0015

The model forecasts statistics of the Bayesian VAR model turned out to be better than those generated by a similar unrestricted VAR model, which was outperformed in terms of Root mean square error, as Figure 7, showing two quarters ahead out of sample forecasts, shows (for ease of exposition, each point on the dashed line shows the forecast of GDP growth—

year-on-year—that one would have made estimating the model up to time t and projecting it two quarters ahead).


Finally, Figure 8 shows six quarters ahead GDP year-on-year growth forecasts generated from the BVAR model.

Figure 7: Year-on-Year GDP Growth Forecasts Two Quarters Ahead, BVAR Versus VAR Model



Note: Each point on the dashed line shows the forecast of GDP growth—year-on-year—that one would have made estimating the model up to time t and projecting it two quarters ahead.

Figure 8: BVAR Projections Six Quarters Ahead: Implied Year-on-Year GDP Growth Rates



IV. CONCLUSIONS

This study has briefly surveyed the main short-run forecasting methods with an application to the Italian GDP case.

In the first part, I have presented a bridge model that is applicable for one or two quarters ahead forecasting of the Italian GDP. The model relies on industrial production and a (promptly available) coincident survey indicator as the key variables that can help providing a first GDP estimate. It performs reasonably well in the years 1999 and 2000, and captures in sign, if not in magnitude, some of the recent turning points of the Italian business cycle.

For a horizon of one to two years, I have presented and estimated a Bayesian VAR model of the Italian economy, with particular attention to the GDP variable. The specification includes GDP, household consumption, real interest rate, survey indicator, exchange rate, inflation, and German GDP.

Both approaches appear to be useful as additional forecasting tools besides structural macroeconomic models, as their out-of-sample forecasting performance shows. Given their simplicity and their ease of use, they could be extended to other variables, including the external sector, and to other countries.

References

- Altissimo, F., D. Marchetti and G. Oneto, 1999, "The Italian Business Cycle: New Coincident and Leading Indicators and some Stylised Facts," ISAE, Documento di Lavoro n.8.
- Bikker, J.A., 1998, "Inflation Forecasting for Aggregates of the EU-7 and EU-14 with Bayesian VAR Models," *Journal of Forecasting*, 17, pp. 145–160.
- Cogley, T., and J.M. Nason, 1995, "Output Dynamics in Real-Business-Cycle Models," *American Economic Review*, vol. 85, pp. 492–511, June.
- Emerson, R., and D.F. Hendry, 1996, "An Evaluation of Forecasting Using Leading Indicators," *Journal of Forecasting*, 15, pp. 271–291.
- Forni, M., M. Hallin, M. Lippi, and L. Reichlin, 2000, "Coincident and Leading Indicators for the Euro Area," mimeo.
- Gaiotti, E., 1999, "The Transmission of Monetary Policy Shocks in Italy, 1967-97," Temi di Discussione del Servizio Studi, Banca d'Italia, n.363.
- Klein, L.R., 1950, *Economic Fluctuations in the United States 1921-41*, New York, Wiley.
- Klein, L.R., and J.Y. Park, 1993, "Economic Forecasting at High-Frequency Intervals," *Journal of Forecasting*, 12, pp. 301-319.
- Klein, L.R., and E. Sojo, 1989, "Combination of High and Low Frequency Data in Macroeconometric Models," in *Economics in Theory and Practice: An Eclectic Approach*, ed. by L.R. Klein and J. Marquez, Dordrecht: Kluwer, 1989, pp. 3-16.
- Litterman, R., 1986, "Forecasting with Bayesian Vector Autoregressions—Five Years of Experience," *Journal of Business and Economic Statistics*, 4, 1, pp. 25-38.
- Parigi, G., and G. Schlitzer, 1995, "Quarterly Forecasts of the Italian Business Cycle by Means of Monthly Economic Indicators," *Journal of Forecasting*, 14, pp. 117-141.
- Rathjens, P. and R. Robins, 1993, "Forecasting Quarterly Data Using Monthly Information," *Journal of Forecasting*, 12, pp. 321-330.
- Robertson, P., and E.W. Tallman, 1999, "Vector Autoregressions: Forecasting and Reality," *Federal Reserve Bank of Atlanta Economic Review*, 1st quarter, pp. 4-17.
- Sims, C.A., 1980, "Macroeconomics and Reality", *Econometrica*, Vol. 48, pp. 1-40.
- Sims, C.A., J. Stock, and M. Watson, 1990, "Inference in Linear Time Series Models with Some Unit Roots," *Econometrica*, 58, pp. 113-44.