Citation for published version:

Deschamps, B & Ioannidis, C 2014, 'The efficiency of multivariate macroeconomic forecasts', *Manchester School*, vol. 82, no. 5, pp. 509-523. <https://doi.org/10.1111/manc.12016>

DOI:

[10.1111/manc.12016](https://doi.org/10.1111/manc.12016)

Publication date:

2014

Document Version

Early version, also known as pre-print

[Link to publication](#)

University of Bath

Alternative formats

If you require this document in an alternative format, please contact:
openaccess@bath.ac.uk

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

THE EFFICIENCY OF MULTIVARIATE MACROECONOMIC FORECASTS*

by

BRUNO DESCHAMPS

Nottingham University Business School China

and

CHRISTOS IOANNIDIS†

University of Bath

We examine the efficiency of multivariate macroeconomic forecasts by estimating a vector autoregressive model on the forecast revisions of four variables (GDP, inflation, unemployment and wages). Using a data set of professional forecasts for the G7 countries, we find evidence of cross-series revision dynamics. Specifically, forecasts revisions are conditionally correlated to the lagged forecast revisions of other macroeconomic variables, and the sign of the correlation is as predicted by conventional economic theory. This indicates that forecasters are slow to incorporate news across variables. We show that this finding can be explained by forecast underreaction.

1 INTRODUCTION

Following Nordhaus (1987), numerous studies have tested the efficiency of macroeconomic forecasts by regressing forecast revisions on their lagged values. Forecast efficiency has often been rejected on the basis that forecast revisions are autocorrelated, indicating that forecasters incorporate new information too slowly. Forecast efficiency has been rejected for various economic series and countries using consensus forecasts (Loungani, 2001; Isiklar *et al.*, 2006; Ager *et al.*, 2009), as well as individual forecasts (Davies and Lahiri, 1995; Lahiri and Sheng, 2008; Dovern and Weisser, 2011).

Previous studies have analyzed forecast efficiency for single variables, such as GDP growth or inflation. However, professional forecasters often forecast multiple macroeconomic series. For instance, banks typically issue monthly reports providing forecasts for output growth, inflation, unemployment, budget deficit, etc. The paper contributes to the forecast efficiency literature by extending the analysis of forecast efficiency to the case of multivariate forecasts, i.e. forecasts made on multiple macroeconomic series rather than a single series. Our objective is to provide a more comprehensive analysis of forecast efficiency which considers both

* Manuscript received 30.9.11; final version received 26.3.13.

† We are grateful to anonymous referees. We would like to thank *Inquire UK* for their financial support.

within-series and cross-series forecast revisions dynamics. For multivariate forecasts, forecast efficiency requires forecast revisions to be both uncorrelated with their lagged values, and also conditionally uncorrelated with the lagged forecast revisions of the other variables. Using the *Consensus Economics Inc.* data set, we measure the degree of inefficiency of multivariate forecasts by examining the cross-temporal correlations between the consensus forecast revisions of the different series. We focus on four variables (GDP growth, inflation, unemployment rate and wages growth) for the G7 countries for the years 1992–2009.

We follow the vector autoregressive (VAR) approach introduced by Isiklar *et al.* (2006), and later used by Lahiri and Isiklar (2009). They estimate a VAR on the GDP growth consensus forecast revisions of the G7 countries, and find that forecasts are slow to incorporate foreign news. We extend their methodology to the case of multivariate forecasts and estimate, for each country, a VAR on the consensus forecast revisions of the four variables. The main finding of this paper is that forecast revisions, in addition to being autocorrelated, are also conditionally correlated to the lagged forecast revisions of other variables. Consensus forecasts typically take several months to fully incorporate news emanating from other variables. Moreover, for any pair of variables, the sign of the cross-series dynamics is typically the one predicted by conventional economic theory. These findings are robust across countries and variables.

We provide two alternative interpretations to the slow reaction to news coming from other variables. First, the cross-series dynamics could reflect forecasters' misunderstanding of the relationships that exist between the variables. Alternatively, we argue that the cross-series dynamics could also derive directly from underreaction to new information. When the degree of underreaction varies across variables, we show that the lagged forecast revisions of a variable x help predict the future forecast revisions of variable y (controlling for the lagged forecast revisions of y). This applies to both consensus and individual forecasts. For instance, when most forecasters are conservative on output but bullish on unemployment, one can attribute the output forecast conservatism to a high degree of underreaction, rather than to the absence of new information. Hence, unemployment forecasts are informative of true output expectations, which may explain the cross-series dynamics. This article is organized as follows. Section 2 presents the data. In Section 3 we develop the empirical prediction and estimate the VAR model. Section 4 develops the theoretical model to interpret the econometric evidence. Finally Section 5 concludes.

2 DATA DESCRIPTION

Consensus Economics Inc. publishes monthly macroeconomic forecasts made by a panel of professional forecasting bodies (such as banks, securities com-

panies, research institutes, and large industrial corporations) for the G7 countries. The forecasts are made on a set of variables including GDP growth and its components, inflation, unemployment, corporate earnings, public deficit and wages. The forecasts are multivariate since each panelist forecasts an average of 10 variables each month. The number of forecasters fluctuates from month to month and between countries. There are for instance, on an average month, 25 forecasters for the USA and the UK, but only 12 for Italy. Forecasters are usually based in the country they forecast. For instance, US forecasts are primarily made by American companies.

The structure of the data is as follows. Every month, each panelist forecasts the macroeconomic variables for both the current and following year. Given that forecasts can be revised each month, each forecaster makes up to 24 repeated forecasts for each target year and each series. For instance, the first set of forecasts for the 2009 the US output growth was released in January 2008, and the final set was released in December 2009. The data set includes all the forecasts made for the target years 1992–2009. We focus on forecasts of GDP growth, inflation, unemployment rate and wages because these are the most significant variables that are available for all countries, with the exception of the wage series which is not available for Canada. Following the approach of Isiklar *et al.* (2006), we examine the efficiency of the consensus forecast, defined as the mean of the individual forecasts, rather than individual forecasts. The size and structure of the *Consensus Forecasts Inc.* data set makes it well suited to test forecast efficiency, both for consensus forecasts (Harvey *et al.*, 2001; Loungani, 2001; Isiklar *et al.*, 2006), and for individual forecasts (Gallo *et al.*, 2002; Batchelor, 2007; Lahiri and Sheng, 2008; Lahiri and Isiklar, 2009; Dovern and Weisser, 2011; Patton and Timmermann, 2011).

We denote by $f_{x,t,h}$ the consensus forecast of variable x made for target year t at horizon h (i.e. h months ahead). The consensus forecast revision at horizon h is denoted $r_{x,t,h} = f_{x,t,h} - f_{x,t,h+1}$. Table 1 shows the range of contemporaneous correlations of the consensus forecast revisions for any pair of variables, which provides preliminary evidence on the linkages between the variables. Each cell shows the range of correlation of consensus forecast

TABLE 1
RANGE OF CORRELATION IN CONSENSUS FORECAST REVISIONS FOR THE G7 COUNTRIES

	GDP	Inflation	Wages	Unemployment
GDP	1			
Inflation	0.18 to 0.45	1		
Wages	0.04 to 0.42	0.09 to 0.37	1	
Unemployment	−0.38 to −0.71	−0.08 to −0.32	−0.08 to −0.37	1

Each cell shows the range of correlation of consensus forecast revisions across countries. For instance, the correlation between inflation and GDP ranges from 0.18 and 0.45 depending on the country.

revisions across countries, and the signs are as predicted by conventional economic theory. Improvements in the economic conditions lead forecasters to revise upwards the forecasts of output, inflation, and wages and revise downward unemployment forecasts.¹ Additionally, the estimated correlation coefficients are significant and sometimes large. For instance the correlation between output and unemployment forecast revisions ranges from -0.38 to -0.71 depending on the country, suggesting that panelists apply Okun's law in their forecasts. This is reassuring given that the empirical support Okun's law is relatively strong (Blinder, 1997). The correlation between output and wages is somewhat weaker in comparison but still consistent with economic theory. Table 1 also reveals that the sign of the correlation coefficients for the same pair of variables is the same for all countries, which is indicative of robustness.

3 EMPIRICAL ANALYSIS

3.1 Methodology

The concept of forecast efficiency, as introduced by Nordhaus (1987), proposes that forecasts are efficient if they incorporate all the available information. Under this definition, forecast revisions should be unpredictable and therefore uncorrelated with their lagged values. It follows that forecast efficiency can be rejected if forecast revisions are autocorrelated. In that context, the standard forecast efficiency test for a variable x (output, inflation etc.) consists in regressing forecast revisions on their k th lagged value:

$$r_{x,t,h} = \beta_x r_{x,t,h+k} + \varepsilon_{t,h} \quad (1)$$

where $k \geq 1$ is the order of the lagged forecast revision. k is usually set equal to one or two, depending on the assumptions made on the structure of forecast errors. Forecast efficiency requires $\beta_x = 0$, but most studies find that the hypothesis of efficiency is rejected in favor of the alternative $\beta_x > 0$. Theoretically, the rejection of efficiency should open the possibility of improving forecasts by adjusting them for the persistence shown. This strategy tends however to be inaccurate because the degree of autocorrelation can vary considerably across forecasters and time periods.²

We extend the literature by considering the case of multivariate forecasts. For multivariate forecasts, forecast efficiency requires forecast revisions uncorrelated both with their lagged values and with the lagged forecast revisions of other series. Our objective is not to test the efficiency of individual

¹Note that the relationship between GDP growth and inflation can be negative during periods of stagflation. The positive correlation suggests that forecasters did not expect stagflation to occur in the G7 countries for the years 1992–2009.

²For a discussion on the limitations of the forecast debiasing strategy, see Goodwin and Lawton (2003).

forecasts, but rather to determine how long it takes for consensus forecasts to fully incorporate news from other macroeconomic variables. We adopt the approach proposed by Isiklar *et al.* (2006). They estimate a VAR model on the consensus GDP forecast revisions for the G7 countries, which is appropriate to examine the linkages of GDP forecast revisions across countries. They find significant off-diagonal impulse responses, which shows that GDP forecasts revisions in one country respond to the lagged GDP forecast revisions of other countries. We apply this methodology to the case of multivariate forecasts. Specifically, we estimate a VAR for each country on the consensus forecast revisions of GDP, inflation, unemployment and wages; which is appropriate to examine the linkages of forecast revisions across series within a country. For each country, the VAR(p) model takes the form

$$r_{t,h} = c + B_1 r_{t,h+1} + B_2 r_{t,h+2} + \dots + B_p r_{t,h+p} + \varepsilon_{t,h} \quad (2)$$

where c is a vector of constants, $r_{t,h} = (rgdp_{t,h}, rinflation_{t,h}, rwage_{t,h}, runemployment_{t,h})'$ is the vector of forecast revisions for target year t and horizon h , $rgdp_{t,h}$ denotes the consensus forecast revision for GDP growth, $rinflation_{t,h}$ the inflation forecast revision, $rwage_{t,h}$ the wages growth forecast revision, and $runemployment_{t,h}$ the unemployment rate forecast revision. B_k is a (4×4) matrix of coefficients for forecast revisions $r_{t,h+k}$.³ The optimum lag length is chosen by Schwartz Bayesian information criterion (SBIC). Since the orthogonalized impulse response functions depend on the ordering of variables in the VAR, we prefer to use generalized impulse response functions which do not depend on the ordering, as proposed by Pesaran and Shin (1998). Because of the well-documented autocorrelation of consensus forecast revisions, we expect the diagonal elements of B_k to be positive, at least for the first lag. The off-diagonal elements of B_k will determine whether the cross-series revision dynamics are significant. If forecasts efficiently incorporate news emanating from other variables, the off-diagonal elements of B_k should be insignificantly different from zero. The off-diagonal elements show how the forecast revisions of one variable respond to the lagged forecast revisions of the other variables. Evidence that some off-diagonal elements of B_k are significantly different from zero would indicate that forecast revisions are predictable using other variables.

3.2 Results

We estimate equation (2) for each country, and plot the cumulative impulse response functions for the USA in Fig. 1. We do not report here the graphs for the other countries since they all look very similar to the USA.⁴ The cumulative impulse response functions (IRFs) show the long-run impact of forecast

³For Canada B_k is 3×3 because the wages variables is not available.

⁴The graphs for the other countries are available from the authors upon request.

Fig. 1. US Cumulative Generalized Impulse Responses

Note: USG is the US GDP forecast revisions, USI is the US inflation forecast revisions, USU is the US unemployment forecast revisions, USW is the US wages forecast revisions

revisions shocks, either from the same variable (diagonal), or from a different variable (off-diagonal). Forecast efficiency is strongly rejected for all seven countries and all four variables. Consider first the cumulative IRF of each variable to its own shocks. In all 26 cases, forecast revisions respond to lagged shocks, and the cumulative IRFs are highly significant. In other words, there is evidence of positive autocorrelation, which is in line with prior studies, possibly suggesting that forecasters underreact to new information.

The novelty of this paper is to examine how forecast revisions respond to shocks emanating from other variables. We find that the off-diagonal elements of the cumulative IRFs are in most cases significantly different from zero. In Fig. 1 for instance, the response of US GDP forecast revisions to the unemployment forecast revisions is significant and negative. Across all seven countries, 67 of the 78 off-diagonal cumulative IRFs are significant at the 5 per cent level. This demonstrates that the cross-series revision dynamics is a widespread phenomenon for all countries and variables, indicating that forecasts do not adjust rapidly enough to the news content of other variables. This constitutes a second level of forecast inefficiency, i.e. inefficiency across series rather than within a single economic series.

We provide additional significance test for the cross-series dynamics by running a series of Granger causality test, and report the results in Table 2. The causality test informs on whether the forecast revisions of one of the variables included in this study can be explained by the joint lagged forecast revisions of the other three. The results clearly support the existence of cross-series correlation. For instance, in six countries the GDP forecast revisions can be predicted more accurately by considering the joint forecast revisions of inflation, wages and unemployment. In most cases, the Granger causality test is significant at 1 per cent, a statistical result formally establishing the existence of cross-series dynamics.

We proceed by examining the signs of the cumulative IRFs in Table 3. It is striking that for the 67 significant off-diagonal elements, the sign of the cumulative IRF is always the identical to the sign of the correlation coefficient between the two variables. This can be seen by comparing the sign of the cumulative IRFs to the signs of the coefficient shown in Table 1. For most pairs of variables, the cross-series revision dynamics are highly significant. For instance, in all seven countries the response of unemployment forecast revisions to GDP forecast revisions is negative and significant. Forecast efficiency can be rejected irrespective of the sign of the cumulative IRFs. Nonetheless, the fact that the signs of the cumulative IRF and the relationship always match is an interesting finding because it facilitates the interpretation of the cross-series revision dynamics, as we argue in Section 4.

It can be argued that the inefficient forecasts can be improved by adjusting the revisions for the persistence shown. In the context of multivariate forecasts, the linkages across variables offer greater possibilities to predict future revisions. Table 4 displays the adjusted R^2 of forecast revisions using

TABLE 2
GRANGER CAUSALITY WALD TESTS THAT THE FORECASTS REVISIONS OF A VARIABLE ARE GRANGER CAUSED BY THE LAGGED FORECASTS REVISIONS OF THE OTHER VARIABLES

	<i>USA</i>	<i>UK</i>	<i>Japan</i>	<i>Germany</i>	<i>France</i>	<i>Italy</i>	<i>Canada</i>
GDP	27.33***	28.26***	22.26***	27.83***	29.46***	25.07***	5.31
Inflation	28.08***	34.57***	9.94	21.47***	25.24***	37.93***	16.51***
Wages	30.54***	61.54***	37.81***	30.05***	34.28***	22.41***	n.a.
Unemployment	22.07***	26.54***	31.92***	8.84	42.78***	20.94***	12.52

Each number shows the Chi-squared from the Granger causality test.
***/**/* denote 1/5/10 per cent significance.

TABLE 3
NUMBER OF COUNTRIES WITH POSITIVE AND NEGATIVE CUMULATIVE GENERALIZED
IMPULSE RESPONSE

Forecast revision	Cumulative response to a shock			
	GDP	Inflation	Wages	Unemployment
GDP	7+	6+	4+	7–
Inflation	7+	7+	5+	6–
Wages	5+	5+	7+	6–
Unemployment	7–	6–	4–	7+

Number of countries in which the sign of the IRFs are significantly (5 per cent) positive or negative. For instance, in five countries wages forecast revisions respond positively to GDP forecast revisions.

TABLE 4
ADJR-SQUARED OF THE CONSENSUS FORECAST REVISIONS USING THE VAR AND
A SIMPLE REGRESSION

	VAR	Simple regression		VAR	Simple regression
USA			France		
GDP	0.29	0.22	GDP	0.39	0.33
Inflation	0.50	0.45	Inflation	0.37	0.31
Wages	0.10	0.02	Wages	0.20	0.11
Unemployment	0.35	0.30	Unemployment	0.34	0.25
UK			Italy		
GDP	0.47	0.41	GDP	0.40	0.33
Inflation	0.31	0.23	Inflation	0.42	0.12
Wages	0.36	0.24	Wages	0.11	0.04
Unemployment	0.45	0.41	Unemployment	0.20	0.09
Japan			Canada		
GDP	0.33	0.28	GDP	0.26	0.24
Inflation	0.24	0.19	Inflation	0.21	0.16
Wages	0.16	0.06	Unemployment	0.18	0.13
Unemployment	0.22	0.13			
Germany					
GDP	0.53	0.48			
Inflation	0.36	0.30			
Wages	0.24	0.16			
Unemployment	0.50	0.42			

either the VAR or a linear regression model (i.e. regress the forecast revisions of each variable on its lagged values, with four lags). On average, the adjusted R^2 increases from 0.24 to 0.32 by estimating the VAR rather than univariate regressions. The finding that additional information can be extracted from multivariate forecasts also shows that multivariate forecasts are in general more informative of a specific variable than univariate ones. Among all the four series, wage forecasts revisions are the least predictable, indicating that wage forecasts exhibit the lowest degree of autocorrelation. For the USA for instance, the simple regression R^2 is 0.22 for GDP and 0.45 for inflation, but

only 0.02 for wages. A possible explanation is that wages forecasts do not receive as much attention as output and inflation, and the incentive to smooth forecast revisions for strategic purposes may be weaker.

4 THEORETICAL INTERPRETATIONS

In this section we seek a theoretical explanation to our findings, and we first argue that the cross-series dynamics could reflect forecasters' tendency to systematically underestimate the relationships between the different series. We subsequently develop an alternative interpretation and show that cross-series correlation could also derive directly from forecast underreaction. Both interpretations apply both to individual and to consensus forecasts.

4.1 *Misunderstanding of Cross-series Relationships*

A possible explanation for the cross-series revision dynamics is that forecasters misunderstand the relationships between different series. For instance, the debates on the exact shape of the Philips curve or Okun's law are still broadly unresolved. In that context, it is not implausible that forecasters underestimate the relationship between output and unemployment (Okun's law), between inflation and unemployment (Philips curve), and more generally between any two series. If forecasters underestimate the strength of the relationship between series x and y , then new information on x will not be immediately incorporated in forecasts of y . Rather, forecasts of y will take several months to fully adjust, resulting in the observed cross-series dynamics. This reasoning applies both to individual forecasters and to consensus forecasts. In sum, our findings are consistent with forecasters systematically underestimating the correlation between variables.

4.2 *Underreaction*

We propose an alternative explanation and show that underreaction can be sufficient to generate information stickiness between series. We consider the case of an individual forecaster, but this model can also be directly applied to consensus forecasts. Suppose that there exists a well-established structural relationship between the expected values of variables x and y at horizon h which is given by $E_h(y) = g(E_h(x))$.⁵ E_h refers to the forecaster's expected value at horizon h . We denote by $f_{x,h}$ and $f_{y,h}$ the forecasts for forecaster i at horizon h (the horizon is the time variable). Forecasts are optimal when equal to the expected value of the variable, i.e. $f_{x,h} = E_h(x)$ and $f_{y,h} = E_h(y)$. Under forecast optimality, $f_{y,h} = g(f_{x,h})$ holds at all time.

⁵The function $g(\cdot)$ could for instance indicate a causality running from x to y . Alternatively, the structural relationship can also capture a situation in which variables x and y are both endogenous and driven by exogenous common factors.

The strategic forecasting literature has repeatedly shown that reputational concerns may induce forecasters to underreact. Ottaviani and Sorensen (2006) find that it is optimal for careerist forecasters to underweight new information when private signals are normally distributed. Underreaction is also likely to occur when the private signals of able forecasters are correlated (Scharfstein and Stein, 1990), or when the consensus forecast is highly accurate (Trueman, 1994). Following the literature on strategic forecasting, we consider the possibility that forecasters underreact, and suppose that the forecasts are formed in the following manner:

$$f_{x,h} = \alpha_{x,h} E_h(x) + (1 - \alpha_{x,h}) f_{x,h+1} \quad (3)$$

$$f_{y,h} = \alpha_{y,h} E_h(y) + (1 - \alpha_{y,h}) f_{y,h+1} \quad (4)$$

where $\alpha_{x,h}$ and $\alpha_{y,h}$ are the weights placed by the forecaster on the optimal forecast. The weights placed on the lagged forecast, $1 - \alpha_{x,h}$ and $1 - \alpha_{y,h}$, capture the degree of underreaction, and are private information. The case of $\alpha = 1$ corresponds to forecast optimality, whereas $\alpha = 0$ corresponds to perfect stickiness ($f_{x,h} = f_{x,h+1}$). We cannot rule out the possibility that the underreaction incentives vary across variables, and hence allow $\alpha_{x,h}$ and $\alpha_{y,h}$ to take different values. To model the fluctuations of the degree of underreaction over time we consider the simplifying assumption that $\alpha_{x,h}$ and $\alpha_{y,h}$ are drawn every period from a uniform distribution $U[0, 1]$, depending for instance on the strength of reputational concerns at horizon h . This simplifying assumption captures the month-to-month variations in α that are unobservable by the market.

Allowing $\alpha_{x,h} \neq \alpha_{y,h}$ implies directly that forecasts may not be consistent with the structural relationship, i.e. $f_{y,h} \neq g(f_{x,h})$. For instance a forecaster with low α for GDP and high α for unemployment will make conservative GDP forecasts (excessively small forecast revisions), while at the same time make bolder revisions to the unemployment forecasts. Such inconsistencies seem to occur regularly in our data set. Figure 2 shows for the USA a negative relationship between GDP and unemployment forecast revisions. It is not infrequent for forecasters to revise both GDP and unemployment forecasts upwards during the same month (or both downwards), indicating possible forecast inconsistency given the negative relationship between the two series.

Proposition 1: The forecast $f_{y,h}$ is conditionally informative of x : $E_h(x|f_{x,h}) \neq E_h(x|f_{x,h}, f_{y,h})$. Also, $E_h(x|f_{x,h}, f_{y,h})$ increases/decreases with $f_{y,h}$ if $g(f_{x,h})$ increases/decreases with x .

Proof: See the Appendix. ■

Proposition 1 shows that the market can estimate more accurately the value of x if a forecaster produces forecasts on both x and y than if he only

FIG. 2. GDP Forecast Revisions versus Unemployment Forecast Revisions for the USA

forecasts x . It is indeed difficult for the market to form an accurate estimate based on a single forecast, because such a forecast may be biased due to underreaction. In the presence of underreaction, the forecaster could for instance be privately more optimistic/pessimistic about x than his published forecast indicates. In sum, the absence of information about the degree of underreaction makes it difficult to assess the true information content of economic forecasts. Now consider the case when the same forecaster forecasts both x and y . Suppose, for instance, that the forecasts are neutral about GDP, but much more optimistic about other variables related to economic activity, such as unemployment and wages, due to a different degree of underreaction. This inconsistency between forecasts would lend support to the view that underreaction (rather than the absence of news) is at the origin of the neutral GDP forecast. Thus, in this instance, the GDP forecast is probably biased, and it can be inferred that the forecaster is more optimistic about GDP than the published forecast indicates. In sum, multivariate forecasts help narrowing down forecasters' true expectations, which helps form a more accurate estimate of x . This intuitive result also applies directly to consensus forecasts.

Corollary 1: (i) Forecast revisions are conditionally correlated with the lagged forecast revisions of the other variables. (ii) The sign of the correlation between the cross-series revision is same as the sign of their underlying relationship.

Proof: See the Appendix. ■

Corollary 1, which automatically follows from Proposition 1, first shows that underreaction is sufficient to generate cross-series dynamics. It also shows that the sign of the cross-series conditional correlation of forecast revisions is identical to the sign of the relationship between the two series. For instance, a negative unemployment forecast revision at $h + 1$ indicates that the forecaster is more optimistic about GDP than his forecast indicates, which in turn implies that a positive GDP forecast revision at horizon h can be expected. In terms of the VAR estimation, this means that the sign of the off-diagonal IRFs will be the same as the sign of the structural relationship between the two economic series. As shown in Section 3, our findings strongly support this prediction. Thus, underreaction can explain the cross-series revision dynamics.

To sum up, we have proposed two possible explanations that can account for the cross-series revision dynamics, both interpretations have similar empirical implications and at this stage we cannot reject either of them.

5 CONCLUSION

In this paper, we have examined the efficiency of multivariate macroeconomic forecasts. Using a data set of professional forecasts made for the G7 countries for the years 1992 and 2009, we have established a number of empirical regularities. Additionally to the well-documented autocorrelation of forecast revisions, multivariate forecasts are characterized by cross-series revision dynamics. Using a VAR approach, we have found that the forecast revisions of a single variable are conditionally correlated to the lagged forecast revisions of other macroeconomic variables. Consensus forecasts are therefore slow to incorporate news emanating from other variables, which constitutes a second level of forecast inefficiency. We also found that the sign of the cross-series revision dynamics is systematically identical to the structural relationship between the two series. Our results underline the usefulness of considering the cross-series linkages to predict the direction and size of future forecast revisions.

We have proposed two possible interpretations to the existence of cross-series dynamics. The first is that forecasters simply misunderstand the strength of the structural relationships between the different macroeconomic variables. The alternative interpretation, instead, relies on the premise that the forecasters underreact to new information and that the degree of underreaction varies across variables. Multivariate forecasts help determine to what extent forecasts truly reflect forecasters' private belief. In other words multivariate forecasts can help the market form more accurate estimates of the forecasted variables.

APPENDIX

Proof of Proposition 1

We consider the $f_{x,h} > f_{x,h+1}$ case, but the proof is symmetric when $f_{x,h} < f_{x,h+1}$. It follows from (3) that $E_h(x|f_{x,h}, f_{y,h})$ increases with $f_{y,h}$ if $E_h(\alpha_{x,h}^{-1}|f_{x,h}, f_{y,h})$ does. Consider two forecasts $f_{y,h}$ and $f'_{y,h}$ such that $f'_{y,h} > f_{y,h}$. We have to show that $E_h(\alpha_{x,h}^{-1}|f_{x,h}, f'_{y,h}) > E_h(\alpha_{x,h}^{-1}|f_{x,h}, f_{y,h})$, which simplifies to $\int \alpha_{x,h}^{-1} q(\alpha_{x,h}^{-1}|f_{x,h}, f'_{y,h}) d\alpha_{x,h}^{-1} > \int \alpha_{x,h}^{-1} q(\alpha_{x,h}^{-1}|f_{x,h}, f_{y,h}) d\alpha_{x,h}^{-1}$ where $q(\cdot)$ is the density function of $\alpha_{x,h}^{-1}$. We use the monotone likelihood ratio property, i.e. show that

$$\frac{q(\alpha_{x,h,1}^{-1} | f_{x,h}, f_{y,h,1})}{q(\alpha_{x,h,0}^{-1} | f_{x,h}, f_{y,h,1})} \geq \frac{q(\alpha_{x,h,1}^{-1} | f_{x,h}, f_{y,h,0})}{q(\alpha_{x,h,0}^{-1} | f_{x,h}, f_{y,h,0})} \quad (A1)$$

for any $\alpha_{x,h,1}^{-1} \geq \alpha_{x,h,0}^{-1}$ and $f_{y,h,1} \geq f_{y,h,0}$, with at least one strict inequality. By Bayes rules, (A1) simplifies to

$$\frac{\Pr(f_{y,h,1} | E^1(x))}{\Pr(f_{y,h,0} | E^1(x))} \geq \frac{\Pr(f_{y,h,1} | E^0(x))}{\Pr(f_{y,h,0} | E^0(x))}$$

where $E^1(x)$ is forecaster's expected value of x that is implied by $f_{x,h}$ and $\alpha_{x,h,1}$. Note that $E^1(x) \geq E^0(x)$ since $\alpha_{x,h,1}^{-1} \geq \alpha_{x,h,0}^{-1}$. There are two possible cases.

Case 1: $f_{y,h,0} < f_{y,h,1} < g(E^0(x))$ or $g(E^0(x)) < f_{y,h,0} < f_{y,h,1}$. The condition (A1) holds, but no strict inequality.

Case 2: $f_{y,h,0} < g(E^0(x)) < f_{y,h,1}$, then (A1) holds with strict inequality. The full details of the proof are available from the authors. ■

Proof of Corollary 1

We can write forecast revisions as: $r_{x,h} = \alpha_{x,h}[E_h(x) - f_{x,h+1}]$. The expected future forecast revision is: $E_{h+1}(r_{x,h}) = E(\alpha_{x,h})E_h(x) - f_{x,h+1}$. Proposition 1 shows that $E_{h+1}(x)$ increases with $f_{y,h+1}$. Hence, $E_{h+1}(r_{x,h})$ also increases with $f_{y,h+1}$. ■

REFERENCES

- Ager, P., Kappler, M. and Osterloh, S. (2009). 'The Accuracy and Efficiency of the Consensus Forecasts: a Further Application and Extension of the Pooled Approach', *International Journal of Forecasting*, Vol. 25, pp. 167–181.
- Batchelor, R. (2007). 'Bias in Macroeconomic Forecasts', *International Journal of Forecasting*, Vol. 23, pp. 189–203.
- Blinder, A. (1997). 'Is There a Core of Practical Macroeconomics that we should All Believe?', *The American Economic Review*, Vol. 87, pp. 240–243.
- Davies, A. and Lahiri, K. (1995). 'A New Framework for Analyzing Survey Forecasts Using Three-dimensional Panel Data', *Journal of Econometrics*, Vol. 68, pp. 205–227.
- Dovern, J. and Weisser, J. (2011). 'Accuracy, Unbiasedness and Efficiency of Professional Macroeconomic Forecasts: an Empirical Comparison for the G7', *International Journal of Forecasting*, Vol. 27, pp. 452–465.

- Gallo, G. M., Granger, C. W. J. and Jeon, Y. (2002). 'Copycats and Common Swings: The Impact of the Use of Forecasts in Information Sets', *IMF Staff Papers*, Vol. 49, No. 1, pp. 4–21.
- Goodwin, P. and Lawton, R. (2003). 'Debiasing Forecasts: How Useful is the Unbiasedness Test?', *International Journal of Forecasting*, Vol. 19, pp. 467–475.
- Harvey, D., Leybourne, S. and Newbold, P. (2001). 'Analysis of a Panel of UK Macroeconomic Forecasts', *Econometrics Journal*, Vol. 4, pp. 37–55.
- Isiklar, G., Lahiri, K. and Loungani, P. (2006). 'How Quickly do Forecasters Incorporate News? Evidence from Cross-country Surveys', *Journal of Applied Econometrics*, Vol. 21, pp. 703–725.
- Lahiri, K. and Isiklar, G. (2009). 'Estimating International Transmission of Shocks Using GDP Forecasts: India and its Trading Partners', in S. Ghatak and P. Levine (eds), *Development Macroeconomics*, London, Routledge, pp. 123–162.
- Lahiri, K. and Sheng, X. (2008). 'Evolution of Forecast Disagreement in a Bayesian Learning Model', *Journal of Econometrics*, Vol. 144, pp. 325–340.
- Loungani, P. (2001). 'How Accurate are Private Sector Forecasts? Cross-country Evidence from Consensus Forecasts of Output Growth', *International Journal of Forecasting*, Vol. 17, pp. 419–432.
- Nordhaus, W. (1987). 'Forecasting Efficiency: Concepts and Applications', *Review of Economics and Statistics*, Vol. 69, pp. 667–674.
- Ottaviani, M. and Sorensen, P. (2006). 'The Strategy of Professional Forecasting', *Journal of Financial Economics*, Vol. 81, No. 2, pp. 441–466.
- Patton, A. and Timmermann, A. (2011). 'Predictability of Output Growth and Inflation: a Multi-horizon Survey Approach', *Journal of Business and Economic Statistics*, Vol. 29, No. 3, pp. 397–410.
- Pesaran, M. and Shin, Y. (1998). 'Generalized Impulse Response Analysis in Linear Multivariate Models', *Economics Letters*, Vol. 58, pp. 17–29.
- Scharfstein, D. and Stein, J. (1990). 'Herd Behavior and Investment'. *American Economic Review*, Vol. 80, No. 3, pp. 465–479.
- Trueman, B. (1994). 'Analyst Forecasts and Herding Behavior', *The Review of Financial Studies*, Vol. 7, pp. 97–124.