ESA/STAT/AC. 202/S5.3

International Seminar on Early Warning and Business Cycle Indicators

14 to 16 December 2009
Scheveningen, The Netherlands

The Euro-area recession and nowcasting GDP growth using statistical models

Gian Luigi Mazzi (Eurostat)
James Mitchell (NIESR)
Gaetana Montana (Eurostat)
Kostas Mouratidis (Sheffield University) and
Martin Weale (NIESR)

The Euro-area recession and nowcasting GDP growth using statistical models[1]

Gian Luigi Mazzi (Eurostat), James Mitchell (NIESR), Gaetana Montana (Eurostat), Kostas Mouratidis (Sheffield University) and Martin Weale (NIESR)

7 December 2009

Abstract

This paper assesses the performance of nowcasts for Euro-area quarterly GDP growth, constructed ahead of Eurostat's *Flash* estimate which is at 45 days, over a period which includes the recent recession. The exercise uses real-time data and allows for the staggered release of monthly information on indicator variables throughout the quarter. The results indicate that the recent recession, due to the global financial crisis of 2007-8, led to a dramatic deterioration in the performance of nowcasts at 0 and 15 days, but a clear improvement relative to autoregressive benchmarks. The utility of constructing nowcasts using indicator variables increased over the recessionary period. But the performance of the different statistical nowcasting models varies considerably according to which statistical model is used. The relative performance of different nowcasting models - and different indicator variables - switched suddenly in the recession with "soft" data from qualitative surveys becoming more important relative to "hard" data on industrial production, even when the "hard" data are known. Despite the fact that prior to the recession it paid to wait for release of two months of within quarter industrial production data, over the recessionary period one would have been better off ignoring this statistical evidence and constructing nowcasts zero-weighting the industrial production data and focusing on the qualitative surveys alone. But our results suggest that now the recession appears to over the utility of indicator-based nowcasts, whether used in a regression or factor-based model, will diminish and the autoregressive benchmark will once again become a competitive, albeit perhaps not the best, nowcasting model. They also suggest that, as before the recession, waiting for two months of industrial production data will improve accuracy.

[1] Address for correspondence: James Mitchell, National Institute of Economic and Social Research, 2 Dean Trench Street, Smith Square, London SW1P 3HE, U.K. Email: j.mitchell@niesr.ac.uk.

Introduction

GDP data are published after a lag. Eurostat publishes its *Flash* quarterly GDP estimates for the Euro-area about 45 days after the end of the quarter. Inevitably, this means that economists and policymakers do not know where we are now, yet alone where we might be in the future. This has impeded economists' ability to track, in real-time, the course of the recent recession in the Euro-area.

There is, therefore, a demand for forecasts of where the economy is now - this quarter or even this month. One option is to construct a composite coincident indictor, which summarises the information in many series all believed to tell us something about the economy's prospects. Such an indicator is often interpreted as representing the "underlying state of the economy"; e.g., see Carriero and Marcellino (2007). Another option, which has the attraction of being directly observable and verifiable, is to nowcast GDP growth itself, since GDP represents an overall measure of the status of the economy. Here we consider nowcasts of quarterly GDP growth. (An alternative strategy, which we do not consider, is to construct monthly GDP estimates; see Mitchell *et al.* (2005), Frale *et al.* (2008) and Angelini *et al.* (2008b)). Consistent with the focus of statistical offices (as opposed to econometricians), we attempt to minimise dependence on forecasts and therefore construct nowcasts only at the end of the quarter, rather than within-quarter or one quarter ahead as in, for example, Angelini *et al.* (2008a,b). Specifically, we construct nowcasts of quarterly GDP growth, ahead of Eurostat's *Flash* estimate, at 0 and 15 days after the end of the quarter. To achieve this improvement in timeliness does nevertheless require recourse to forecasting models, although to a lesser extent than if the nowcasts were produced before the end of the quarter of interest.

Nowcasts are typically produced by statistical forecasting models. These statistical models by construction, and in contradistinction to structural or economic models, are reduced-form. They seek to explain and then nowcast GDP growth by exploiting information on *indicator* variables. These are variables which are meant to have a close relationship with GDP but are made available more promptly than the data for which they stand as a proxy. In practice there is a large number of potential indicator variables, both quantitative ("hard") and qualitative ("soft"). But the advantage of the "soft" data is that they tend to be published ahead of "hard" data. The set increases further, when as possible indicator variables, we consider variables not directly related to GDP but presumed to have some indirect relationship. For example, interest rates or other financial data might be considered as they might help explain/predict GDP growth.

In this paper we assess the ability of some widely used statistical models to anticipate the recession in the Euro-area (EA12), and then adapt to it. This involves comparing and contrasting both regression and factor-based approaches of using "soft" and "hard" data to nowcast quarterly GDP growth. Nowcasts are produced at 0 and 15 days after the end of the quarter. It is important to distinguish between the timeliness of different nowcasts, since there is a trade-off between the timeliness and accuracy of nowcasts. Nowcasts can always be produced more quickly by exploiting less *hard* information, but we might expect the quality of the nowcasts to deteriorate as a result. Importantly, we identify what if any indicator variables were most helpful in anticipating the recession. We study the changing performance of models' relative performance and establish that there are clear benefits, over an autoregressive benchmark, to constructing indicator-based nowcasts – but principally during the recessionary period. During the stable pre-recessionary period we find that the models' relative performance is different. There is little to choose between not just indicator-based nowcasts and autoregressive forecasts, but different means of using the indicator variables. This finding is consistent with the view that in the stable ("Great Moderation") period from the mid 1980s to 2007, while it became "easier" to forecast (in the sense that the root mean squared error of forecasts declined), it became "harder" to beat an autoregressive benchmark (e.g., see Stock and Watson, 2007).

Indicator variables

There is always a question about the set of indicator variables to consider. Giannone *et al.* (2008) provide a formal analysis, stressing the distinction between the timeliness and informational content of *hard* and *soft* data. Industrial production is an obvious (*hard*) indicator for GDP growth, and we will therefore consider it. (We also experimented with retail sales data but found them to have no additional informational content.) Industrial production data, unlike GDP data, are available monthly and are published by Eurostat 45 days or so after the end of the month to which they relate. This means that if we want a nowcast half way through the month, after the end of the quarter, only the third month of the quarter (for industrial production) needs to be forecast. Inaccurate forecasts of one month may nevertheless deliver reasonably accurate projections of the quarter to which they belong.

However, we might expect (or hope) that a broader range of indicator variables for GDP growth are available. A popular source of information on the current and future direction of the economy is the findings from qualitative business surveys, of the sort collected in the European Commission's Euro-wide database of business and consumer surveys. A perceived attraction of these surveys is that they are forward, as well as backward, looking. Given the size of the German economy, we also consider the IFO surveys separately and in addition; but we do note that the source of DG-ECFIN's data for Germany is the IFO so, in principle, the two data sets should be, at the minimum, highly correlated.

These qualitative surveys ask many questions about recent and expected (future) experiences. Accordingly, since respondents typically reply "up", "same" or "down", we consider the balance statistics (the proportion of optimists less pessimists) first as compiled and aggregated in the Economic Sentiment survey. This summarises evidence across many questions; specifically, it involves us considering 5 indicators: (i) the industrial confidence indictor; (ii) the services confidence indicator; (iii) the consumer confidence indicator; (iv) the retail trade confidence indicator and (v) the Economic Sentiment indicator which is a composite of the previous 4 indicators. These survey data, as in Charpin, Mathieu and Mazzi (2008), are considered both in levels and first differences. We note that the Flash results were much improved when the survey data were entered in first differences, as well as levels.

Secondly, we consider a far larger set of balance statistics. This includes consideration of not just the 5 questions above - for the EA - but the corresponding national data. Euro-area GDP is an aggregate variable. Therefore, rather than simply forecast the aggregate (EA growth) using aggregate indicators, we follow the suggestion of Hendry and Hubrich (2006) and include disaggregate indicators in the aggregate model. As mentioned above, we also separately consider as indicators the business climate, expectations and situation indices and balance statistics from IFO. In addition, following Charpin, Mathieu and Mazzi (2008), we consider households' opinion of major purchases over the next 12 months, the construction confidence indicator, households' financial position over the next 12 months and employment expectations in construction.

In summary, we consider two blocks of qualitative survey data, the first being far smaller and the second far larger. When producing nowcasts via the regression methodology (considered below) we consider just the small set; while in the factor analysis we consider both sets separately. Boivin and Ng (2006) show that there can be benefits (less noise) to limiting the number of variables from which the factors used in nowcasting/forecasting are extracted.

These qualitative survey data have both backward and forward looking aspects. The statistical models we consider below let the data decide which questions are most helpful when nowcasting, and thereby automate the choice so that judgment plays no role.

We also test the informational content of some financial data. We consider them despite the fact that one might argue that indicator variables like interest rates, whatever their influence on demand, are not so suitable for the production of (pre-) first estimates of data (as opposed to forecasts) since they do not have such clear links to the data they are supposed to represent. Specifically, we considered quarterly averages of monthly ECB data for the 10 year Euro bond yields, the 1 year bond yield, the difference between these (the interest-rate spread) and the growth rate of M3. These data are published at the end of the month to which they relate. Financial data, in particular the interest-rate spread (the yield curve), have been found in other work to be effective at forecasting recessions (e.g., see Estrella and Mishkin, 1998).

Just as disaggregate (national-level) "soft" indicators are considered we could add in "hard" national GDP data if available in-time. But national GDP data for the EA countries are not been published much, it at all, ahead of Eurostat's *Flash* estimate at 45 days. Over the last few months Belgian GDP data have been available 10 days ahead of the EA data (at about 35 days), but this has not been the case historically and therefore we cannot yet conduct a historical examination of their use. But we should expect the increased timeliness of national GDP data to deliver more accurate EA nowcasts.

The modelling approach

Since current and lagged values of these indicator variables, and lags of GDP itself, can plausibly help explain GDP one must consider carefully how one selects the indicator variables for the model used to explain GDP growth. The number of possible indicator variables can easily get very large. The problem is then to either “select”, in some sense, the “best-fitting” indicators (these best-fitting indicators can be chosen on the basis of both *a priori* or objective in-sample performance criteria) or “reduce” the set of indicator variables “automatically”. Once this is done the models that can be used to nowcast GDP growth are estimable using classical statistical methods- there are no degrees of freedom constraints. Different models involve different ways of linking the indicator variables to GDP. This can be done at a quarterly, monthly or mixed frequency. It is an empirical question which is most sensible.

In this paper we consider forecasting quarterly GDP growth directly using both regression and factor based models. An alternative is to nowcast and forecast monthly GDP, using monthly indicator variables, subject to aggregation constraints; e.g., see Mitchell *et al.* (2005), Frale *et al.* (2008) and Angelini *et al.* (2008b).

Regression-based nowcasts (with a small number of indicator variables)

The regression-based approach focuses on a small set of indicator variables that bear a close relationship to GDP. Its attraction is that it is simple and easy-to-interpret.
The approach adopted is designed to comply with the criterion that the models used to produce nowcasts should be credible to policy-makers and other non-statisticians. We see this as ruling out processes with lengthy lags in exogenous variables, since it is difficult to defend a situation where an indicator is sharply influenced by some other variable up to six months or so ago.

The desire to produce clear models with short lags is reinforced by the fact that in many cases the data series we have available are although monthly, generally short in duration. This makes it more difficult to explore cointegration satisfactorily and our models will be regression equations constructed with the dependent variable entering only in log differences.

The modelling framework requires only a one period ahead forecast. Regression-based nowcasts are produced as special cases of the following general regression equation expressed at the quarterly frequency ($t=1,\ldots,T$):

$$\Delta y_t = c + \sum_{i=1}^{p} \alpha_i \Delta y_{t-i} + \sum_{i=0}^{p} \sum_{j=1}^{k} \beta_{ij} x_{t-i,j} + u_t; \; (t = 1, \ldots, T) \qquad (1)$$

where Δy_t is the log of the dependent variable (quarterly GDP growth in our application), $x_{t,j}$ is the *j*-th (quarterly) indicator variable (*j=1,…,k*) in logs when appropriate, *c* is an intercept, *p* the number of lags and u_t a disturbance. All indicator variables that enter (1), if necessary, are differenced until stationary. The use of (log) first differences is deemed sufficient to render all series stationary. We also note that for short horizons the forecasting performance from univariate nonlinear models is typically worse or not much better; see Stock and Watson (1999) and Marcellino (2008). We therefore confine attention to simpler linear models.

It should be noted that contemporaneous quarterly values of the indicator variables are included in (1). This reflects the fact that these indicators by their nature are published ahead of the variables to which they are assumed to relate, even though they may relate to the same time period.

Various specifications of (1) might be considered based on different methodological approaches to selecting the specification to use for regression-based estimation. Given *k* indicator variables and a given number of lags (we consider *p=4*), for *t=1,...,T*, we consider all possible combinations of (1) of the *(p+1).k* exogenous and *p* lagged endogenous variables thus generated. Since, however, this creates a very large possible number of regressions and bearing in mind the well-known benefits of parsimony in forecasting models, we limit ourselves to those equations containing no more than three explanatory variables. There are ${}^{[(p+1)k+p]}C_3 + {}^{[(p+1)k+p]}C_2 + [(p+1)k+p]+1$ such equations. We then "automatically" select the preferred model using the Bayesian Information Criterion (BIC). We then use this model, and its estimated coefficients from the sample *t=1,…,T*, and the quarter T+1 values of the explanatory variables in the preferred model to nowcast Δy_{T+1}. Recall that the T+1 values of the indicator variables are published ahead of the T+1 values for Δy_{T+1} and can therefore be exploited when nowcasting.

Monthly Bridge Equations

Various methods are available to nowcast a quarterly variable like GDP exploiting monthly indicator variables. As indicated above one option is to consider the problem as one of constructing monthly GDP estimates. The estimation of monthly GDP amounts to a temporal disaggregation or a distribution problem. A popular alternative is to use bridge equations, since this framework sits naturally with our focus on (small *k*) regression-based nowcasts.

Bridging involves linking monthly data, typically released early in the quarter, with quarterly data like GDP; e.g., see Salazar and Weale (1999) and Baffigi *et al.* (2004). In effect a two-equation system is now used to nowcast Δy_{T+1}, with the second equation comprising the forecasting model for the monthly variable $x_{t,j}$. The errors between the two equations, at the underlying monthly frequency, are assumed orthogonal so that the equations are estimated separately. In common with much previous work, see Diron (2008), we consider simple AR models for $x_{t',j}$:

$$x_{t',j} = \sum_{i=1}^{p} \beta_i x_{t'-i,j} + e_{t',j} \tag{2}$$

where t'=1,…,Tm denotes the monthly data with *m=3* months in the quarter.

The Flash model for Δy_t, (1), is therefore estimated (in-sample, *t=1,...,T*) using hard quarterly data on $x_{t,j}$. However, when wanting to nowcast Δy_{T+1} since we may only have partial information on $x_{T+1,j}$ (for some indicator variables, *j*) the predicted values $\widehat{x}_{T+1,j}$ from the AR model are used instead when nowcasting from (1). For example, when wanting to nowcast GDP growth with only two months of *hard* (Eurostat) data on industrial production available, the final month in the quarter is forecast using the AR model. This forecasted value is then combined with the two months of hard data to obtain $\widehat{x}_{T+1,j}$.

Combination nowcasts

We also consider variants of the above regression-based methods which involve combining nowcasts across different models.

Combination offers a means of integrating out model uncertainty, in other words of insuring ourselves against having picked the wrong (regression) model. While the BIC selects the “best” regression model, this model may not be selected with probability one. There is a considerable body of work that has found forecast combination to often work well; see Timmermann (2006) for a recent survey. Equal weighting is often found to work as well as more complex (optimal - variance weighted) alternatives (see Smith and Wallis, 2009).

Therefore, we consider the benefits of combining the nowcasts from the many regressions. That is, rather than use the BIC to select the best-fitting model we combine all of the nowcasts. This is achieved by taking an equal-weighted average and also a BMA weighted average where each forecast is weighted in-line with its BIC value using a non-informative prior.

Factor-based nowcasts

The factor-based approach can consider a large(r) set of indicator variables than these regression based approaches and summarises their information in a small number of (unobserved) common factors. These factors are then used to help predict the variable of interest. While, as Eklund and Kapetanios (2008) review, there are many factor based forecasting approaches, we focus on the most popular. This is the static (principal components or diffusion index) approach popularised by Stock and Watson (2002) but building on Rhodes (1937). This involves extracting up to *k* principal components from the set of indictors, possibly stacked over time also, and then relating these to GDP growth via a linear regression. Giannone *et al.* call this "bridging with factors", although their focus is on use of within-quarter (monthly) information on the indicator variables.

Given our focus is constructing nowcasts at the end of the quarter of interest the "soft" data, although available monthly, are known for the whole quarter and we therefore do not consider as in Giannone *et al.* (2008, 2009) and Angelini *et al.* (2008a,b) extracting factors from a set of monthly and quarterly indicator variables where there is a "ragged-edge" such that not all observations are known for the whole quarter. We confine attention to quarterly indicator variables.

Specifically, we take the first 8 principal components from both the small and large set of survey balances mentioned above; this kind of method has been found to be an effective means of modelling a large number of noisy survey variables in a parsimonious manner (see Hansson *et al.*, 2005). We use the BIC to select up to 3 preferred factors to use when nowcasting.

Benchmark models

To evaluate the performance of the nowcasting models, considered above, it is important to have a benchmark. Ability to beat the benchmark, systematically over time, suggests that the model is of use. In economic forecasting the most popular benchmark, which proves surprisingly difficult to beat, is the autoregressive forecast which involves setting $\beta_{ij} = 0$ in equation (1). We set *p=1*. We also consider the random walk model which also sets $\alpha_1 = 1$. This model is robust to structural breaks (see Clements and Hendry, 1998). Both of these benchmark models condition on the previous quarter's GDP estimate but, unlike the regression and factor models, do not exploit any within quarter information. Given the potential predictability of data revisions (to GDP growth), we also consider AR forecasts constructed not only from the latest vintage of GDP data but previous monthly vintages. We consider both using the BIC to select the preferred data vintage to use when nowcasting and also the merits of equal-weighted combinations across the different vintages. Clements and Galvão (2009) show that forecasts computed from an AR estimated using the latest vintage of data only need not minimise RMSE when data revisions are not mean zero.

The accuracy of nowcasts: a real-time exercise

We compare the accuracy of nowcasts of Euro-area (EA12) GDP growth in recursive out-of-sample experiments using real-time data. Specifically, we use the real-time data triangles for real GDP and industrial production available from Hendyplan's EuroIND database. This database provides a daily snapshot of the data held by Eurostat. The qualitative survey data are not revised (in a significant manner at least) and we use "final" vintage data only. Models are estimated on data vintages back to 2001 with data back to 1991q1. Seasonally adjusted data are used.

It is important to use real-time data, namely data available to the forecaster at the time they actually made their forecast rather than the latest release from Eurostat, given that data are revised (see Mitchell, 2004 and Garratt and Vahey, 2006). Use of the latest vintage of data may give a misleading impression of the accuracy of a given forecasting model/strategy, since in reality the forecaster used an earlier vintage of the data to make their forecasts. We therefore use the latest vintage of data available to the forecaster when they made their forecast.[2]

Nowcasts for GDP growth are computed recursively from 2003q2-2009q2 using the various models and data vintages. The importance of conducting such a nowcasting competition follows from the view that it is always possible to explain past growth reasonably well, using a relatively small number of carefully chosen variables with carefully chosen lags. But there is no reason to expect that such equations will necessarily be good nowcasting tools.

The trade-off between the timeliness and accuracy of nowcasts

We focus on producing nowcasts of quarterly GDP growth to two timescales. Both assume that we know the value of GDP in the previous quarter, although they use different releases of the national accounts and take into account the staggered release of data on the monthly indicator variables.

The first nowcast is produced at 0 days after the end of the quarter. At this point in time the qualitative survey data are known for the calendar quarter. But monthly industrial production data are available only for the first month in the quarter, and we therefore forecast the remaining two months using equation (2). At 0 days we know the first release of the national accounts (which contains the second estimate of GDP - the *Flash* estimate is Eurostat's first estimate) and therefore use these data to construct the nowcasts.

The second nowcast is produced at 15 days when we have two month's hard data for industrial production (IP), and only have to forecast the one remaining month in the quarter using equation (2), and we also know qualitative survey data for all three months in the quarter. We have the third GDP release at this point, with the publication of the second release of the national accounts.

[2] As indicated above, there is an issue when using real-time data of whether, in the jargon of this literature, it is best to estimate models on data from the latest column, or the diagonal, from the real-time data triangle; see Corradi *et al.* (2009). Here we stick to the convention of going down the column and using the latest release of data; but as discussed when constructing benchmark AR nowcasts we do consider the merits of exploiting more than one vintage of data at the same time.

We did experiment with a third nowcast, produced at 30 days, when we also know retail trade (RT) data for two months of the calendar quarter. The final month in the quarter for RT was nowcast as in equation (2). But these retail trade data did not deliver improved nowcasts, however exploited, and we do not consider them further.

Accuracy of the nowcasts: results

The nowcasts are evaluated in Tables 1 and 2 by reporting their root mean squared error (RMSE) against the subsequent 'outturn', defined as either the first (*Flash*) or latest (as of 13 August 2009) Eurostat estimate. Both are defined as quarterly GDP growth in percentage points measured as 100 times the first-difference of the logarithm of real GDP. Figure 1 plots the real-time nowcasts from selected models since 2003q2.

Importantly, rather than as is traditional evaluate the accuracy of the competing nowcasts "on average" over some 'arbitrary' evaluation period, we follow Giacomini and Rossi (2010) and consider how models' relative performance has changed over time. This is important, since a model may nowcast well relative to its competitors during one part of the evaluation period, but there may be a reversal during another part of the evaluation period. We focus on isolating the effect the recent recession, due to the global financial crisis of 2007, has on models' relative performance. Therefore Table 1 summarises the performance of the models over the 2003q2-2009q2 evaluation period, while Table 2 consider the pre-recessionary period 2003q2-2007q4.

Comparison of Tables 1 and 2 is therefore instructive in helping us learn about how the behaviour of the nowcasting models has changed with the onset of the recession. The recession led to a dramatic increase in the models' RMSE, with most models' RMSE statistics doubling. These nowcasting models were unable to keep up with the rapid fall in GDP growth which started in 2008q2. But, as we will discuss, some models adjusted more quickly than others.

As well as a decline in absolute accuracy, the relative performance of the models changed substantially with the recession. First, in the pre-recessionary period, Table 2 shows that the regression, factor and combination nowcasts all perform, certainly relative to Table 1, similarly relative to the autoregressive benchmark. But there are, nevertheless, some minor differences between the alternative nowcasting models.

The best nowcasts are produced (with a RMSE of 0.201) at 15 days when industrial production data are used as the sole indicator variable. But importantly this model has to be chosen judgmentally, rather than selected automatically using statistical evidence. This is seen by the fact that when the BIC is used to select the preferred indicator(s) to use in the regression model it does not, as shown in Table 3, select contemporaneous industrial production. The data suggest that lagged industrial production, rather than contemporaneous industrial production, is the preferred indicator. But worse nowcasts are obtained if one uses lagged industrial production data.

Comparing the RMSE statistics in Table 2 for the AR(1) against both Eurostat's First (*Flash*) estimate and the latest release, indicates that data revisions do matter; the RMSE rises from 0.206 to 0.266 when the nowcasts are evaluated against the latest rather than the first GDP release.

Turning to Table 1 which includes the recession in the evaluation period but focusing on models' relative performance, we see that, in contrast to Table 2, it is easier to beat the AR benchmark with both regression and factor-based models delivering RMSE statistics, for some variants, less than 0.5. The AR(1) delivers a RMSE of 0.712. So within quarter information is more clearly improving accuracy, once the evaluation period includes the 2008-9 recession caused by the global financial crisis.

Also in contrast to Table 2 we see from Table 1 that waiting 15 days for the release of two months' "hard" data on industrial production no longer appears to be worth it. As accurate, or more accurate, nowcasts are produced without reference to the industrial production data. This means better nowcasts are produced at 0 days than 15 days. But since the data (the BIC) if left to decide on the preferred indicators to use in the regression model sometimes select industrial production, as shown in Table 3, it again pays to exclude the industrial production data from the set of indicator variables considered. If excluded, more accurate nowcasts are produced. For example, with 2 indicators selected by the BIC, the regression model delivers nowcasts with a RMSE of 0.433 if IP is excluded but 0.52 if included. The BIC when left to decide on the preferred indicators, as shown in Table 3, does select different indicators over time. Table 3 shows considerable instability in terms of the preferred indicators, whether the nowcasts are produced at 0 or 15 days. Moreover, Table 3 indicates the clear role "soft" data play, since qualitative survey data like the ESI are often selected as the preferred indicator(s). Thanks to their forward-looking nature, these soft data detected the recession more quickly than the hard data. But comparison with Table 1, when nowcasts are produced from a regression where the ESI is the sole indicator variable, indicates there are clear gains to using the BIC to select multiple indicators from a larger set of indicators. The ESI alone, which delivers a RMSE of 0.528, is not the best way of summarising the informational content of the wide array of soft data available, since a regression with 2 contemporaneous soft indicators (or lagged hard indicators), selected by the BIC, delivers a RMSE of 0.433 at 0 days.

Table 1 also indicates that selection is better than combination. While combined nowcasts worked well in pre-recessionary times (see Table 2), during the recession it is better to select one particular model rather than combine across many. The recession was marked by such a dramatic fall in GDP growth that combination nowcasts, which tend to be very smooth (as Figure 1 confirms), cannot adapt to the change quickly enough. This is so even when a weighted, rather than equal-weighted, average is taken.

But Table 1 indicates that the factor methods perform competitively with the regression models where the BIC selects the preferred indicators (from a set of indicator which excludes IP). But importantly, the factor methods are much more robust to whether IP is included in the set of indicator variables or not. There is also not much between the factor-based nowcasts when a small or large information set is used. This result contrasts that in Giannone *et al.* (2009) who find that in an application nowcasting Euro-area GDP growth that disaggregated information on surveys does not increase forecast accuracy; although, as discussed, their focus is within-quarter rather than end-of-quarter nowcasting.

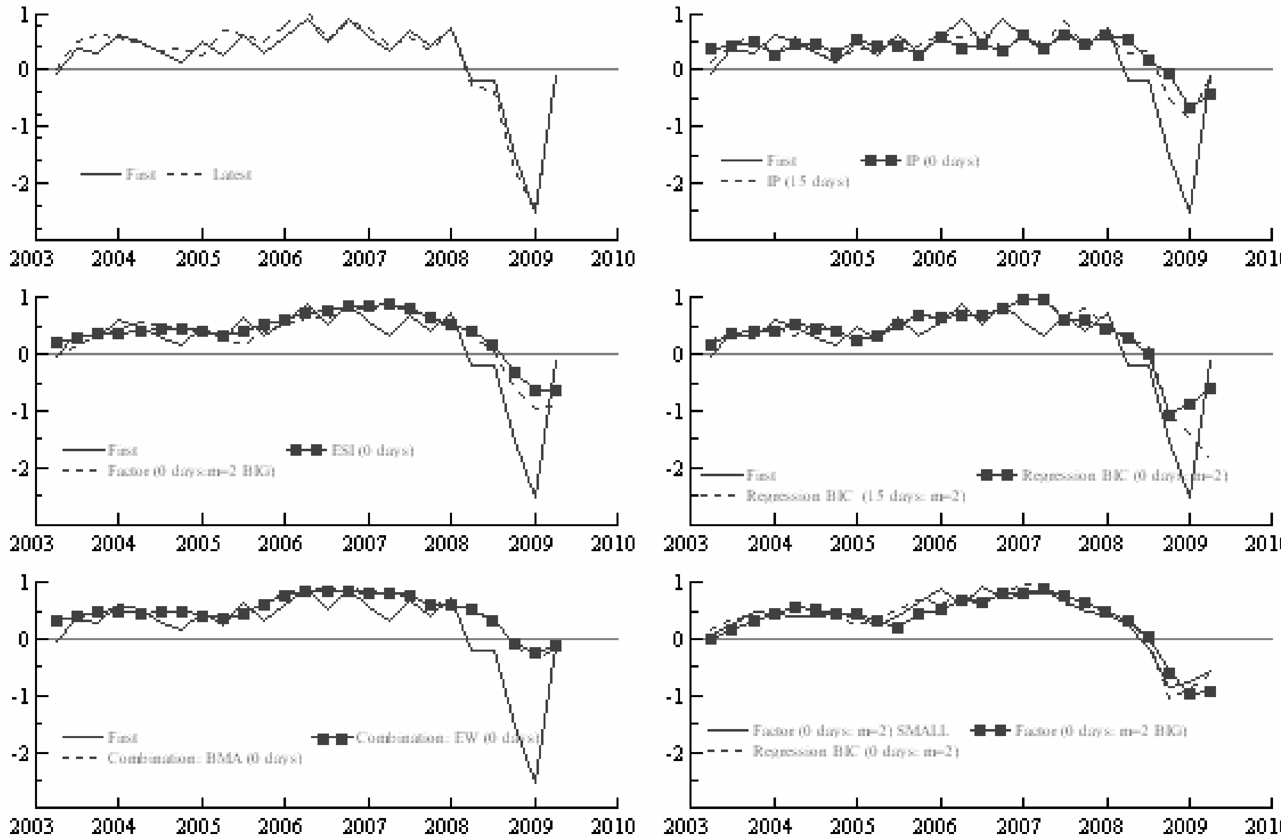

Figure 1: Eurostat's Flash quarterly GDP growth estimates for the Euro-area (100 times by the log first difference of GDP) plus Flash and Factor nowcasts at t+0 and t+15 days

Instability

Table 3 has shown that the preferred indicators selected by the BIC change over time. Comparison of Tables 1 and 2 has also shown that models' relative performance, in particular relative to the autoregressive benchmark, changed dramatically with the onset of the recession. It is therefore instructive to draw out further the instability over time in models' performance, and in particular consider how the role of different indicator variables changes over time. A given indicator may be helpful in one period but not in another.

Therefore, to examine further how models' performance changes over time we follow Giacomini and Rossi (2010) and consider the evolution of the models' performance relative to an autoregression. This involves measuring the local relative forecasting performance of the models. Given our short sample, we eschew the statistical tests proposed by Giacomini and Rossi (2010), which work off the average of the difference between the models' squared forecast errors, and simply plot in Figure 2, as a function of time, the squared forecast error of the nowcast of interest against the squared forecast error from an AR(1).

Figure 2 considers nowcasts constructed from the ESI indicator and from IP alone at both 0 days (when we forecast 2 months in the quarter) and at 15 days (when we forecast the final month in the quarter).

Figure 2 indicates that coincident with quarterly GDP growth becoming negative in 2008q2 the indicator-based nowcasts start to perform better than the autoregression. The gain over the AR is most marked for the nowcasts using IP at 15 days when only the final month in the quarter (of IP) has to be forecast. The superiority of the indicator-based nowcasts increases as the recession deepened but begins to wear off after the trough of the recession in 2009q1.

Conclusion

Previous studies (e.g., Banbura and Rünstler, 2007; Angelini *et al.*, 2008b) have shown that improved nowcasts/forecasts result from increasing the weight on *hard* data relative to *soft* data as the publication lag increases. Our findings, with a focus on backcasting rather than forecasting, complement these in showing that the relative informational content of soft and hard data also depends on the "regime"; in a recessionary regime the utility of the soft data remains high even when the hard data are known. This contrasts Banbura and Rünstler (2007), who find that the *optimal* weight on the soft data declines as soon as the hard data are known.

Perhaps unsurprisingly we also show that the recent recession in the Euro-area has seen a marked deterioration in the performance of some commonly used nowcasting methods, including regression and factor based approaches. But as well as a dramatic rise in models' RMSE when evaluated over the recessionary period our exercise, using real-time data, also indicated that there have been important and rapid switches in models' relative performance. In particular, the utility of constructing nowcasts using indicator variables increased over the recessionary period. Over the relatively stable period up to 2007q4 the improvement over autoregressive benchmark nowcasts is pretty limited. (This finding is shared by others; e.g. see Rünstler *et al.* (2009)). But once the evaluation period is extended to include the recession the autoregressive nowcasts are clearly beaten by the

Figure 2: Forecasting GDP growth over time using real-time data: relative performance of indicator-based nowcasts against an
AR(1)

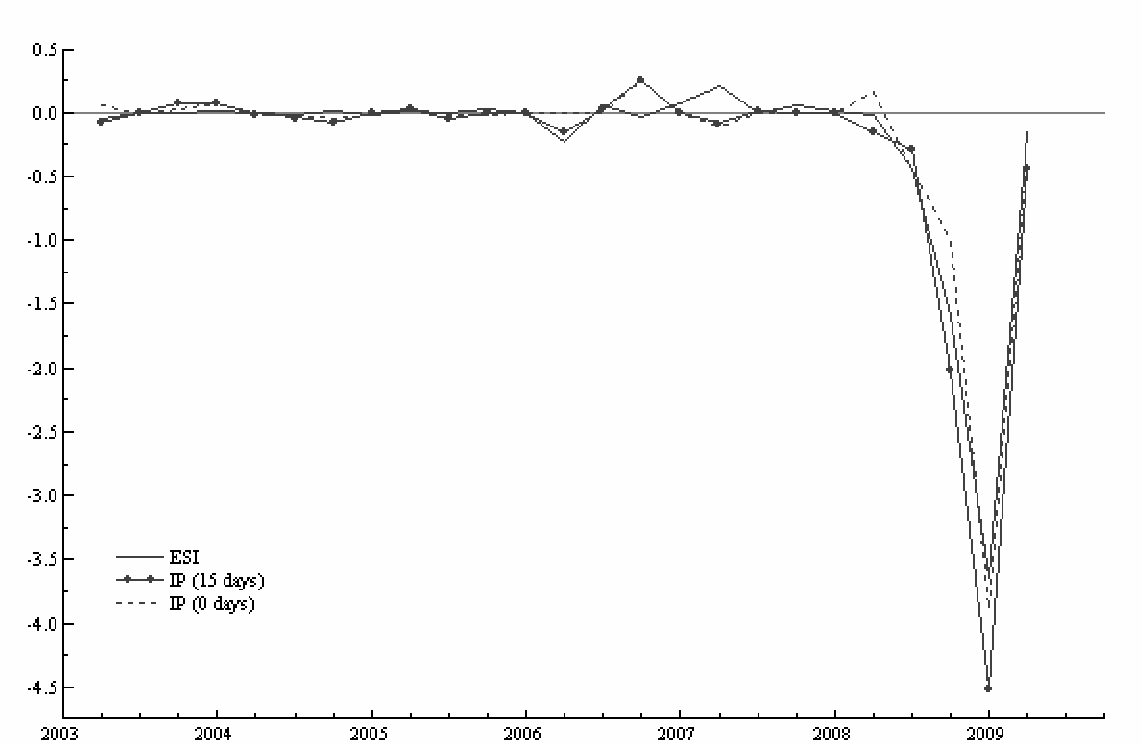

Notes: The lines represent the squared forecast error of the nowcast of interest (ESI, IP at 15 days or IP at 0 days) against the squared forecast error from an AR(1) computed at 15 days

indicator-based nowcasts, which adapt more quickly to the recessionary "regime".

But counter-intuitively, nowcasts, when evaluated over the recessionary period, are in general more accurate at 0 days than 15 days. This is because in a regression-based approach the BIC will select IP data, if in the information set, since in-sample it is highly correlated with GDP growth. But over the recent (recessionary) past, one is better off not considering IP growth since the qualitative survey data provide more accurate impressions of the future path of the recession. That is, one is better off looking at the qualitative survey data even though historically they are not as well correlated with GDP growth: put simply, at times of recession their informational content increases relative to *hard* (IP) data. Nowcasting models which focus on the *hard* data proved slower at adapting to the end of the period economists have called the Great Moderation. As a result they were, in general, less successful at anticipating the recession.

Therefore it appears that we should distinguish between the performance of the nowcasts in the stable ("Great Moderation") period that existed until last year and their performance in the current recession. The qualitative survey data, with their forward-looking questions, appear better able to adapt to the rapidly changing circumstances than the models which use hard indicator data. The regression BIC method and the (quarterly) factor methods appear to offer the best means of handling these changes, although as discussed it is important to zero weight the IP data in the regression but not in the factor methods. It will be interesting to see in future research whether mixed-frequency factor models, of the sort used by Angelini *et al.* (2008b), are able to pick up the rapid switch in the utility of soft and hard indicators automatically (*ex ante*).

Throughout we have focused on the point nowcasts – or the 'central' (conditional mean) predictions from the statistical models. At times of uncertainty it is, in fact, particularly important to consider not just models' central estimates but their density forecasts. While in Figure 1 we have seen "disagreement" increase, between the models' nowcasts, uncertainty more generally has not been considered. An interesting exercise for future research is to assess the ability of these models to predict the probability of a recession and more generally to examine their density nowcasts.

Table 1: RMSE of nowcasts of quarterly Euro-area real GDP growth (at a quarterly rate) 0 and 15 days after the end of the quarter against Eurostat's first (*Flash*) estimate (at 45 days) and their "final" release (as of 13 August 2009): 2003q2-2009q2. Estimates at 0 days are presented both when industrial production (IP) is included and excluded from the information set

	First Estimate			Final Estimate		
Horizon: days	0	0	15	0	0	15
	-IP	+IP		-IP	+IP	
Statistical model						
Regression m=1	0.527	0.682	0.576	0.563	0.710	0.611
Regression m=2	0.433	0.521	0.496	0.444	0.543	0.513
Regression m=3	0.519	0.460	0.491	0.536	0.477	0.510
IP BIC	-	0.623	0.595	-	0.665	0.635
IP	-	0.543	0.457	-	0.591	0.499
Factor: m=1 SMALL	0.584	0.584	0.583	0.616	0.616	0.615
Factor: m=2 SMALL	0.457	0.457	0.457	0.477	0.477	0.477
Factor: m=3 SMALL	0.445	0.451	0.446	0.466	0.473	0.466
Factor m=1 BIG	0.466	0.466	0.465	0.508	0.508	0.507
Factor m=2 BIG	0.462	0.462	0.460	0.505	0.505	0.503
Factor m=3 BIG	0.455	0.455	0.455	0.493	0.492	0.493
Combination: Equal	0.605	0.608	0.592	0.631	0.640	0.620
Combination: BMA	0.595	0.592	0.580	0.620	0.623	0.607
ESI (in levels)	0.528	-	0.528	0.564	-	0.564
AR(1)	0.712	-	0.712	0.759	-	0.759
AR(1) revision: BIC	0.684	-	0.670	0.711	-	0.698
AR(1) revision: Comb	0.632	-	0.634	0.673	-	0.675
Random walk	0.701	-	0.705	0.753	-	0.757

Notes: -IP denotes without IP; +IP denotes IP is considered as a possible indicator variable. Regression denotes nowcasts from a regression model with multiple indicators, like equation (1), where we limit the computational burden to models with at most m explanatory variables from the set of *hard* and *soft* indicators - with the BIC used to select the preferred model. The set consists of the 5 survey balances, survey data as used in Charpin et al. (2008), IFO data, plus the finance data. The survey data are considered both in levels and first differences. Combination is an equal-weighted combination of all nowcasts considered when m=2 (i.e. 1540 models at 0 days and 1711 models at 15 days). AR(1) is a first-order autoregressive model in quarterly GDP growth. AR(1) revision uses not just latest vintage but previous vintage GDP data, with BIC selecting the best-fitting revisions model and Comb combining across the different revisions models. The random walk nowcasts use the latest vintage GDP estimates. The "small" information set used in the factor analysis consists of the five survey balances, survey data as used in Charpin et al. (2008), IFO survey data and the financial variables. The "large" information set adds disaggregate survey information also to create a 142 variable dataset. ESI (IP) denotes nowcasts from a model when the Economic Sentiment Indicator (industrial production data) is considered as the sole indicator variable.

Table 2. RMSE of nowcasts of quarterly Euro-area real GDP growth (at a quarterly rate) 0 and 15 days after the end of the quarter against Eurostat's first (*Flash*) estimate (at 45 days) and their "final" release (as of 13 August 2009): 2003q2-2007q4. Estimates at 0 days are presented both when industrial production (IP) is included and excluded from the information set

PRE-RECESSION VARIANT OF TABLE 1

	First Estimate			Final Estimate		
Horizon: days	0	0	15	0	0	15
	-IP	+IP		-IP	+IP	
Statistical model						
Regression m=1	0.214	0.247	0.239	0.232	0.261	0.246
Regression m=2	0.241	0.274	0.258	0.228	0.283	0.253
Regression m=3	0.284	0.260	0.295	0.276	0.260	0.282
IP BIC	-	0.242	0.236	-	0.270	0.252
IP	-	0.231	0.201	-	0.273	0.227
Factor: m=1 SMALL	0.287	0.287	0.287	0.293	0.293	0.293
Factor: m=2 SMALL	0.236	0.237	0.238	0.243	0.244	0.244
Factor: m=3 SMALL	0.211	0.211	0.213	0.221	0.221	0.221
Factor m=1 BIG	0.221	0.221	0.220	0.260	0.260	0.259
Factor m=2 BIG	0.218	0.218	0.205	0.261	0.261	0.248
Factor m=3 BIG	0.213	0.213	0.211	0.246	0.246	0.244
Combination: Equal	0.225	0.204	0.221	0.204	0.200	0.203
Combination: BMA	0.235	0.208	0.229	0.211	0.203	0.209
ESI (in levels)	0.220	-	0.220	0.237	-	0.237
AR(1)	0.206	-	0.207	0.266	-	0.267
AR(1) revision: BIC	0.294	-	0.292	0.299	-	0.301
AR(1) revision: Comb	0.224	-	0.270	0.256	-	0.292
Random walk	0.276	-	0.285	0.351	-	0.357

Notes: Regression denotes nowcasts from a regression model with multiple indicators, like equation (1), where we limit the computational burden to models with at most m explanatory variables from the set of *hard* and *soft* indicators **-** with the BIC used to select the preferred model. The set consists of the 5 survey balances, survey data as used in Charpin et al. (2008), IFO data, plus the finance data. The survey data are considered both in levels and first differences. IP BIC denotes a model when the BIC is again used to select the preferred indicator but only from IP data and their lags. IP denotes a regression where contemporaneous IP is the only indicator variable. Combination is an equal-weighted combination of all nowcasts considered when m=2 (i.e. 1540 models at 0 days and 1711 models at 15 days). AR(1) is a first-order autoregressive model in quarterly GDP growth. AR(1) revision uses not just latest vintage but previous vintage GDP data, with BIC selecting the best-fitting revisions model and Comb combining across the different revisions models. The random walk nowcasts use the latest vintage GDP estimates. The "small" information set used in the factor analysis consists of the five survey balances, survey data as used in Charpin et al. (2008), IFO survey data and the financial variables. The "large" information set adds disaggregate survey information also to create a 142 variable dataset. ESI (IP) denotes nowcasts from a model when the Economic Sentiment Indicator (industrial production data) is considered as the sole indicator variable.

Table 3: Indicators recursively selected by the BIC in the regressions when m=2

	0 days: without IP		0 days: with IP		15 days	
2003-2	ESI (L)	INDU (D)	Exp (L)	IP (-1)	ESI (L)	IP (-1)
2003-3	ESI (L)	INDU (D)	CONS (L)	IP (-1)	ESI (L)	INDU (D)
2003-4	ESI (L)	INDU (D)	CONS (L)	IP (-1)	CONS (L)	IP (-1)
2004-1	ESI (L)	INDU (D)	CONS (L)	IP (-1)	CONS (L)	IP (-1)
2004-2	INDU (D)	ESI (L: -1)	CONS (L)	IP (-1)	CONS (L)	IP (-1)
2004-3	ESI (L)	INDU (D)	ESI (L)	IP (-1)	ESI (L)	IP (-1)
2004-4	ESI (L)	INDU (D)	ESI (L)	IP (-1)	ESI (L)	IP (-1)
2005-1	ESI (L)	INDU (D)	ESI (L)	IP (-1)	ESI (L)	IP (-1)
2005-2	INDU (D)	ESI (L: -1)	ESI (L)	IP (-1)	ESI (L)	IP (-1)
2005-3	ESI (L)	INDU (D)	ESI (L)	IP (-1)	ESI (L)	IP (-1)
2005-4	ESI (L)	Exp (L)	ESI (L)	Exp (L)	ESI (L)	Exp (L)
2006-1	ESI (L)	Exp (L)	ESI (L)	Exp (L)	ESI (L)	Exp (L)
2006-2	ESI (L)	Exp (L)	ESI (L)	Exp (L)	ESI (L)	Exp (L)
2006-3	ESI (L)	Exp (L)	ESI (L)	Exp (L)	ESI (L)	Exp (L)
2006-4	ESI (L)	Exp (L)	ESI (L)	Exp (L)	ESI (L)	Exp (L)
2007-1	ESI (L)	Exp (L)	ESI (L)	Exp (L)	ESI (L)	Exp (L)
2007-2	ESI (L)	Exp (L)	ESI (L)	Exp (L)	ESI (L)	Exp (L)
2007-3	ESI (D)	ESI (L: -1)	ESI (L)	IP (-1)	ESI (L)	IP (-1)
2007-4	ESI (D)	ESI (L: -1)	ESI (L)	IP (-1)	ESI (L)	IP (-1)
2008-1	ESI (D)	ESI (L: -1)	ESI (L)	IP (-1)	ESI (L)	IP (-1)
2008-2	ESI (D)	ESI (L: -1)	ESI (L)	IP (-1)	ESI (D)	ESI (L: -1)
2008-3	ESI (D)	ESI (L: -1)	ESI (L)	IP (-1)	ESI (L)	IP (-1)
2008-4	ESI (D)	ESI (L: -1)	ESI (D)	ESI (L: -1)	ESI (D)	ESI (L: -1)
2009-1	ESI (L)	IFO (D: -1)	ESI (L)	IP (-1)	ESI (L)	IP (-1)
2009-2	ESI (L)	INDU (D: -1)	ESI (L)	IP (-1)	ESI (L)	IP (-1)

Notes: L denotes the survey balance in levels and D denotes in first differences. -1 denotes the previous quarter (lagged) value of the indicator is selected. ESI is the Economic Sentiment Indicator; IP is industrial production growth; Exp denotes consumers' expectations over the next 12 months; INDU is the Industrial Confidence Indicator ; IFO denotes German survey-data from IFO; CONS is the Consumer Confidence Indicator.

Table 4: Nowcasts of quarterly GDP growth: Eurostat estimates versus selected indicator-based regression, factor and combination nowcasts at 0 and 15 days after the end of the quarter

	Eurostat Flash	Eurostat Latest	IP	IP	ESI	Regress m=2 w/o IP	Regress m=2 w IP	Factor m=2 Small	Factor m=2 Big	Combin EW	Combin BMA
	45 days	2009	0 days	15 days	0 days	0 days	15 days	0 days	0 days	0 days	0 days
2003-2	-0.05	0.01	0.37	0.15	0.20	0.17	0.25	0.06	-0.01	0.32	0.30
2003-3	0.37	0.52	0.41	0.44	0.27	0.38	0.38	0.25	0.16	0.39	0.38
2003-4	0.30	0.62	0.49	0.58	0.37	0.40	0.40	0.47	0.32	0.50	0.51
2004-1	0.61	0.56	0.26	0.25	0.36	0.39	0.50	0.44	0.45	0.50	0.51
2004-2	0.51	0.46	0.48	0.58	0.42	0.55	0.32	0.40	0.55	0.45	0.46
2004-3	0.30	0.34	0.46	0.29	0.44	0.47	0.55	0.42	0.51	0.48	0.49
2004-4	0.15	0.36	0.31	0.19	0.45	0.41	0.41	0.44	0.45	0.49	0.50
2005-1	0.49	0.26	0.53	0.37	0.40	0.24	0.25	0.39	0.44	0.42	0.42
2005-2	0.26	0.71	0.42	0.43	0.35	0.35	0.31	0.25	0.31	0.36	0.36
2005-3	0.64	0.62	0.44	0.58	0.39	0.52	0.48	0.39	0.18	0.45	0.46
2005-4	0.31	0.49	0.24	0.41	0.52	0.68	0.67	0.64	0.43	0.63	0.63
2006-1	0.59	0.79	0.57	0.61	0.60	0.64	0.64	0.89	0.52	0.78	0.80
2006-2	0.88	1.12	0.38	0.58	0.74	0.71	0.71	0.60	0.67	0.87	0.89
2006-3	0.52	0.56	0.47	0.70	0.78	0.69	0.69	0.95	0.65	0.85	0.87
2006-4	0.90	0.86	0.35	0.36	0.84	0.82	0.82	0.73	0.81	0.86	0.89
2007-1	0.57	0.74	0.61	0.58	0.84	0.98	0.98	0.75	0.82	0.82	0.85
2007-2	0.34	0.41	0.38	0.47	0.89	0.97	0.97	0.86	0.87	0.82	0.85
2007-3	0.70	0.58	0.63	0.86	0.81	0.60	0.71	0.68	0.77	0.75	0.78
2007-4	0.41	0.36	0.47	0.42	0.66	0.63	0.80	0.49	0.63	0.61	0.63
2008-1	0.73	0.74	0.60	0.59	0.54	0.44	0.50	0.44	0.50	0.60	0.61
2008-2	-0.20	-0.28	0.55	0.31	0.43	0.29	0.27	0.25	0.33	0.52	0.53
2008-3	-0.20	-0.39	0.18	0.34	0.18	0.00	0.14	-0.17	0.05	0.32	0.31
2008-4	-1.53	-1.81	-0.08	-0.49	-0.30	-1.06	-1.06	-0.83	-0.58	-0.06	-0.11
2009-1	-2.54	-2.48	-0.68	-0.87	-0.62	-0.87	-1.42	-0.76	-0.97	-0.25	-0.31
2009-2	-0.11	-0.11	-0.41	-0.15	-0.64	-0.59	-1.85	-0.56	-0.93	-0.11	-0.20

References

Angelini, E., G. Camba-Mendez, D. Giannone, L. Reichlin and G. Rünstler (2008a), "Short-term Forecasts of Euro Area GDP Growth". ECB Working Paper No. 949.

Angelini, E., M. Banbura and G. Rünstler (2008b), "Estimating and forecasting the Euro area monthly national accounts from a dynamic factor model". ECB Working Paper No. 953.

Baffigi, A., R. Golinelli and G. Parigi (2004), "Bridge models to forecast the euro area GDP", *International Journal of Forecasting*, 20, 447-460.

Banbura, M. and G. Rünstler (2007), "A look into the factor model black box: publication lags and the role of hard and soft data in forecasting GDP", ECB working paper 751.

Boivin, J. and S. Ng (2006), "Are More Data Always Better for Factor Analysis?", *Journal of Econometrics*, 132, 169-194.

Carriero, A. and M. Marcellino (2007), "A comparison of methods for the construction of composite coincident and leading indexes for the UK", *International Journal of Forecasting*, 23, 219-236.

Charpin, F., C. Mathieu and G. Mazzi (2008), "Construction of coincident indicators for the euro area". Eurostat working paper (September 2008).

Clements, M. and A. B. Galvão (2009), "Real-time forecasting with Autoregressive Models in the presence of data revisions", Warwick University Discussion Paper.

Corradi, V., A. Fernandez and N.R. Swanson (2009), "Real-time datasets really do make a difference: definitional change, data release and forecasting". Available at http://www.bundesbank.de/download/vfz/konferenzen/20090323_24_berlin/paper_corradi_fernandez.pdf

Clements, M.P. and D.F. Hendry, (1998), *Forecasting Economic Time Series*, Cambridge University Press: Cambridge.

Diron, M. (2008), "Short-term forecasts of euro area real GDP growth: an assessment of real-time performance based on vintage data", *Journal of Forecasting*, 27, 371-390.

Eklund, J. and G. Kapetanios (2008), "A review of forecasting techniques for large datasets", *National Institute Economic Review*, 203, 109-115.

Estrella, A. and F. Mishkin (1998), “Predicting U.S. recessions: Financial variables as lead-ing indicators”, *Review of Economics and Statistics,* 80, 45–61

Frale, C., M. Marcellino G-L. Mazzi and T. Proietti (2008), “A Monthly Indicator of the Euro Area GDP”, CEPR Discussion Paper 7007.

Garratt, A. and S.P. Vahey (2006), “UK Real-Time Macro Data Characteristics”, *Economic Journal*, 116, 119-135.

Giacomini, R. and B. Rossi (2010), “Forecast comparisons in unstable environments”, *Journal of Applied Econometrics*. Forthcoming.

Giannone, D., L. Reichlin and D. Small (2008), “Nowcasting: The real-time informational content of macroeconomic data,” *Journal of Monetary Economics,* 55, 665-676.

Giannone, D., L. Reichlin and S. Simonelli, (2009), “Nowcasting Euro-area economic activity in real time: the role of confidence indicators, *National Institute Economic Review,* 210, 90-97.

Hansson, J., P. Jansson and M. Löf (2005), “Business survey data: Do they help in forecasting GDP growth?”, *International Journal of Forecasting,* 21, 377-389.

Hendry, D.F. and K. Hubrich (2006), “Forecasting economic aggregates by disaggregates”, ECB Working Paper No. 589.

Marcellino, M. (2008), “A linear benchmark for forecasting GDP growth and inflation?,” *Journal of Forecasting*, 27, 305-340.

Mitchell, J. (2004), “Revisions to economic statistics”. *Statistics Commission Report No. 17 Vol. 2.* (http://www.statscom.org.uk/uploads/files/reports/Revisions_vol_2.pdf)

Mitchell, J., R.J. Smith, M.R. Weale, S. Wright and E.L. Salazar (2005), “An indicator of monthly GDP and an early estimate of quarterly GDP growth, *Economic Journal*, 115, 108-129.

Pesaran, M H. and A. Timmermann (2007), “Selection of estimation window in the presence of breaks”, *Journal of Econometrics*, 137, 134-161.

Rhodes, E.C. (1937), “The construction of an index of business activity” (with discussion), *Journal of the Royal Statistical Society: Series C*, 100, 18-66.

Rünstler, G., K. Barhoumi, S. Benk, R. Cristadoro, A. Den Reijer, A. Jakaitiene, P. Jelonek, A. Rua, K. Ruth and C. Van Nieuwenhuyze (2009), “Short-term forecasting of GDP using large monthly datasets - a pseudo real-time forecast evaluation exercise”, *Journal of Forecasting*, 28, 595-611.

Salazar, E. and M. Weale (1999), “Monthly data and short-term forecasting: an assessment of monthly data in a VAR model”, *Journal of Forecasting*, 18, 447-462.

Smith, J. and K.F. Wallis (2009), “A simple explanation of the forecast combination puzzle”, *Oxford Bulletin of Economics and Statistics*, 71, 331-355.

Stock, J.H. and M.W. Watson (1999), “A comparison of linear and nonlinear univariate models for forecasting macroeconomic time series”, *Cointegration, causality and forecasting: a festschrift in honor of Clive W.J. Granger* (editors Engle, R. and R. White). Oxford: OUP, pp. 1-44.

Stock, J.H. and M.W. Watson (2002), “Macroeconomic forecasting using diffusion indexes”, *Journal of Business and Economic Statistics*, 20, 147-162.

Stock, J.H. and M.W. Watson (2007), **“**Has Inflation Become Harder to Forecast?” *Journal of Money, Credit, and Banking*, 39, 3-34.

Timmermann, A. (2006), “Forecast combinations”. Handbook of Economic Forecasting Volume 1 (editors Elliott, G., C W J. Granger and A. Timmermann). North-Holland: Amsterdam, pp. 135-196.