NEW EUROCOIN: TRACKING ECONOMIC GROWTH IN REAL TIME

Filippo Altissimo, Riccardo Cristadoro, Mario Forni, Marco Lippi, and Giovanni Veronese*

Abstract—Removal of short-run dynamics from a stationary time series to isolate the medium- to long-run component can be obtained by a bandpass filter. However, bandpass filters are infinite moving averages and can therefore deteriorate at the end of the sample. This is a well-known result in the literature isolating the business cycle in integrated series. We show that the same problem arises with our application to stationary time series. In this paper, we develop a method to obtain smoothing of a stationary time series by using only contemporaneous values of a large data set, so that no end-of-sample deterioration occurs. Our method is applied to the construction of New Eurocoin, an indicator of economic activity for the euro area, which is an estimate, in real time, of the medium- to long-run component of GDP growth. As our data set is monthly and most of the series are updated with a short delay, we are able to produce a monthly real-time indicator. As an estimate of the medium- to long-run GDP growth, Eurocoin performs better than the bandpass filter at the end of the sample in terms of both fitting and turning-point signaling.

I. Introduction

THIS paper presents a method to estimate in real time the current state of the economy, with an application to the euro area. The resulting indicator, New Eurocoin (NE), is intended to replace the Eurocoin indicator proposed by Altissimo et al. (2001) and published monthly by the Centre for Economic Policy Research (see www.cepr.org).

The main objective of NE is to make an assessment of economic activity that is (a) comprehensive and nonsubjective, (b) timely, and (c) free from short-run fluctuations. Requirement a is obvious. Regarding requirements b and c, both private agents and economic policymakers require for their decisions a clear distinction, in real time, between transitory and long-lasting changes in the state of the economy. For example, if an upward change occurs, it is crucial to decide whether it is the beginning of a long, positive swing or a short-lived phenomenon. In particular, a countercyclical policy should target medium- rather than short-term waves. The latter are both less detrimental and more difficult to fight, owing to the delays of policy reactions and the effects of intervention on economic activity.

None of the available macroeconomic series provides a measure of the state of the economy that fulfills criteria a, b, and c. GDP, the most comprehensive indicator of real activity, fails to meet requirements b and c. Regarding timeliness, the GDP is available only quarterly and with a long delay. For instance, the preliminary estimate of euro-area GDP for the first quarter of the year does not become available until May. Moreover, the GDP is affected by a sizable short-run component.

NE is a real-time estimate of GDP growth, cleaned of short-run oscillations. More precisely, we focus, first, on the growth rate of the GDP and define the *medium- to long-run growth* (MLRG), as the component of the GDP growth rate obtained by removing the fluctuations of a period shorter than or equal to 1 year. This component, which is, of course, a smoothing of GDP growth, is our ideal target. Second, NE is a monthly and timely estimate of the MLRG for the euro area: around the twentieth of each month, we are able to produce a reliable estimate for the previous month.

To avoid possible misunderstandings, let us point out that "medium- to long-run growth" denotes only the smoothed component of the growth rate defined above; it bears no relationship to any definition of trend. In particular, integration of the medium- to long-run component will never be considered.

The MLRG, as defined above, is obtained by applying a bandpass filter, which is, however, an infinite, two-sided moving average. Empirical applications imply replacing missing with predicted data, and therefore a possible deterioration at the end of the sample. In particular, poor end-of-sample estimation and serious revisions as new data become available have been consistently stated in the literature trying to isolate business cycle fluctuations in macroeconomic integrated time series (see Baxter & King, 1999; Christiano & Fitzgerald, 2003). The same end-of-sample problem arises applying bandpass filters to any stationary time series. However, this paper concentrates on a particular stationary series, the euro-area GDP growth, and the bandpass filter that removes one year or shorter fluctuations. The end-of-sample deterioration, for this case, is discussed in the paper and assessed in a real-time exercise in section VI.

A substantial mitigation of this conflict between timeliness and removal of the short-run fluctuations is the main contribution of this paper. Our indicator NE, an estimate of the MLRG, is based on a large data set, including 145 euro-area macroeconomic variables. We construct a small number of smooth factors, which are generalized principal components of current values of the variables in the data set, specifically designed to remove short-run and variable-specific sources of fluctuation. NE is obtained as a linear combination of the smooth factors. Because only current values of the variables are used, no end-of-sample deterioration occurs. Moreover, although NE cannot compete with the truncated bandpass filter within the sample, we show that NE outperforms the bandpass filter at the end of the sample in terms of both fitting and turning-point signaling.

This result can be explained by observing that the data set contains variables that are leading with respect to current GDP, and the smoothness of our factors is obtained by

Received for publication June 6, 2007. Revision accepted for publication October 29, 2008.

* Altissimo: Brevan Howard Asset Management; Cristadoro and Veronese: Bank of Italy; Forni: Università di Modena e Reggio Emilia, CEPR, and RECent; Lippi: Università di Roma 'La Sapienza.'

We gratefully acknowledge encouragement and support from the Bank of Italy and the Centre for Economic Policy Research. We thank Lucrezia Reichlin who started the Eurocoin project some years ago, Marc Hallin for his suggestions at an early stage of the project, and Antonio Bassanetti for his ideas and unique expertise in the realization of the previously published version of Eurocoin. We also thank Andreas Fisher, Domenico Giannone, James Stock, Mark Watson, and three anonymous referees for constructive criticism and suggestions. The views expressed here do not necessarily reflect those of the Bank of Italy, the Brevan Howard Asset Management, or any other institutions with which the authors are affiliated.

TABLE 1.—THE CALENDAR OF SOME MACROECONOMIC SERIES

Time	December 2004	January 2005	February 2005	March 2005	April 2005	May 2005	June 2005	Delay
GDP	Q3 - 2004	Q3 - 2004	Q4 - 2004	Q4 - 2004	Q4 - 2004	Q1-2005	Q1-2005	45–90
Industrial production	Oct. 04	Nov. 04	Dec. 04	Jan. 05	Feb. 05	Mar. 05	Apr. 05	45–50 days
Surveys	Dec. 04	Jan. 05	Feb. 05	Mar. 05	Apr. 05	May 05	Jun. 05	0–25 days
Retail sales	Oct. 04	Nov. 04	Dec. 04	Jan. 05	Feb. 05	Mar. 05	Apr. 05	45–50 days
Financial markets	Dec. 04	Jan. 05	Feb. 05	Mar. 05	Apr. 05	May 05	Jun. 05	0 days
CPI	Nov. 04	Dec. 04	Jan. 05	Feb. 05	Mar. 05	Apr. 05	May 05	15 days
Car registrations	Nov. 04	Dec. 04	Jan. 05	Feb. 05	Mar. 05	Apr. 05	May 05	2–30 days
Industrial orders	Oct. 04	Nov. 04	Dec. 04	Jan. 05	Feb. 05	Mar. 05	Apr. 05	50 days

The table body lists the last available update period.

linearly combining current values of variables that are lagging, coincident, and leading with respect to the GDP. Therefore, the information contained in future values of the GDP, which are unavailable at the end of the sample, can be partially recovered using the smooth factors.

The method we use is based on the large-scale generalized dynamic factor model (GDFM) proposed by Forni et al. (2000, 2005) and Forni and Lippi (2001) (see also the literature cited in section V). Valle e Azevedo, Koopman, and Rua (2006) propose a multivariate method with bandpass filter properties, which exploits information from a relatively small number of variables. We are not far in spirit from their work, the main difference being that our procedure is designed to extract information from a large panel of time series.

Our ideal target, being an infinite moving average, is, strictly speaking, unobservable. However, as we show in Appendix A, our finite-sample version of the bandpass filter provides a good approximation to the ideal target until we are 1 year away from the end and beginning of the sample. (The appendixes are available online at http://www.mitpressjournals.org/doi/suppl/10.1162/REST_a_00045.) This is the basis for our adopted empirical target and measure of performance for NE. The performance of NE at time t, with $t \leq T - 12$, is measured as the difference between NE at time t, and the empirical target at t, obtained using the data up to T.

The paper is organized as follows. Section II collects some preliminary observations. Section III defines our target—the medium- to long-run component of GDP—and discusses its interpretation. Sections IV and V describe and motivate our estimation procedure. Section VI constructs the NE indicator and analyzes its real-time performance in comparison with alternative indicators. Section VII concludes. The appendix contains a detailed discussion of the ideal target, the empirical target and their distance, a description of the data set, and a short comparison between New and Old Eurocoin.

II. Preliminary Observations

To gauge the current state of the economy given the delay with which GDP is released, market analysts and forecasters resort to more timely and higher-frequency information and on this basis obtain early estimates of GDP. However, two problems immediately arise: (a) looking at the typical release calendar for the euro area, one can see that timeliness varies greatly even among monthly statistics (end-of-sample unbalance), and (b) since GDP is quarterly, we have to handle monthly and quarterly data simultaneously.

In what follows, we show how to combine the comprehensive and nonsubjective information provided by GDP with the early information provided by surveys and other monthly series to obtain a reliable and timely picture of current economic activity.

Our data set includes monthly series of consumer and production prices, wages, share prices, money, unemployment rates, job vacancies, interest rates, exchange rates, industrial production, orders, retail sales, imports, exports, and consumer and business surveys for the euro-area countries and the euro area as a whole (see Appendix B for details). The data set has been organized taking into account the calendar of data releases that is typical in real situations, with the aim of reproducing the staggered flow of information available through time to policymakers and market forecasters. This lack of synchronism, though little considered in the literature, is crucial for assessing realistically the performance of alternative real-time indicators.[1]

As illustrated in table 1, financial variables and surveys are the most timely data, while industrial production and other "real variables" are usually available with longer delays. Around the twentieth of month $T+1$, when we calculate the indicator for month T, surveys and financial variables are usually available up to time T (thus with no delay), car registrations and industrial orders up to $T-1$, and industrial production indexes up to $T-2$ or $T-3$. The GDP series is observed quarterly, so that its delay varies with time. For example, on April 20, only data up to the fourth quarter of the previous year are available; thus there is a three-month delay with respect to T, which is March. On May 20, the delay with respect to T is reduced to one month, as a first-quarter preliminary estimate is released, and will be two months when T is May—hence, an average delay of two months.

The most timely variables (such as purchasing managers indexes, consumer surveys, and business climate indexes) are far from being comprehensive and smooth. Other standard series, such as industrial production and exports, complement the information content of the surveys but are less

[1] Important exceptions are Bernanke and Boivin (2003) and Giannone, Reichlin, and Sala (2002).

FIGURE 1.—SOME ECONOMIC INDICATORS FOR THE EURO AREA (NORMALIZED SCALE)

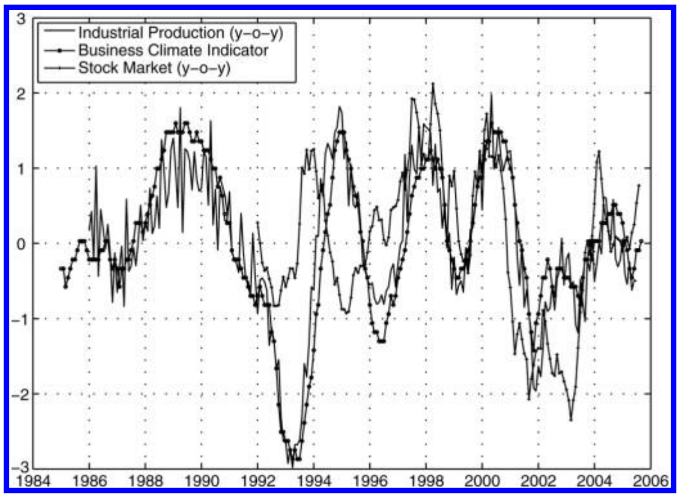

timely. Furthermore, all monthly series exhibit heavy short-run fluctuations and might provide contradictory signals (see figure 1). As a result, none of them is fully satisfactory, and "there is much diversity and uncertainty about which indicators are to be used" (Zarnowitz & Ozyildirim, 2006).

We tackle the end-of-sample unbalance in the following way. Let x^*_{it}, $i = 1, \ldots, n$, be the series after outliers and seasonality have been removed and stationarity achieved by a suitable transformation (see Appendix B). Let k_i be the delivery delay (in months) for variable x^*_{it}, so that when we are at the end of the sample, its last available observation is $x^*_{i,T-k_i}$. We define the panel x_{it}, $i = 1, \ldots, n$, by setting

$$x_{it} = x^*_{i,t-k_i}, \tag{1}$$

so that the last available observation of x_{it} is at T for all i. Of course, this realignment implies cutting some observations at the beginning of the sample for several variables. As a result, after transforming and realigning, the data set goes from June 1987 to June 2005—hence, $T = 217$. The same realignment is used both when we consider the whole sample up to T and when we consider subsamples $[1\ \tau]$, as in the pseudo-real-time exercises carried out in section VI.

To use our monthly data set to obtain a timely GDP indicator, it is convenient to think of GDP as a monthly series of quarterly aggregates with missing observations. The figure for month t, denoted by z_t, is defined as the aggregate of GDP for months t, $t - 1$, and $t - 2$, so that there is a two-month overlap between two subsequent elements of the series. Obviously the monthly series is observable only for March, June, September, and December.

The monthly GDP growth rate is defined as

$$y_t = \log z_t - \log z_{t-3}.$$

Thus, y_t is the usual quarter-on-quarter growth rate, except that it is defined for all months.

How to deal with the missing observations in GDP is discussed in detail in section III and in Appendix A.1.

III. The MLRG and Its Interpretation

A natural way to define the medium- to long-run fluctuations of a time series is by considering its spectral representation. Assuming stationarity, y_t can be represented as an integral of sine and cosine waves with frequency ranging between $-\pi$ and π, with respect to a stochastic measure (see Brockwell & Davis, 1991). Based on the spectral representation, we define the medium- to long-run component of y_t by taking the integral over the interval $[-\pi/6\ \pi/6]$ instead of $[-\pi\ \pi]$. The frequency $\pi/6$ for a monthly series corresponds to a 1-year period; thus, we cut off seasonal and other higher-frequency waves.

This frequency-domain construction has a time-domain counterpart, which is known as bandpass filter. Here we do not delve into the details of this correspondence and go directly to the result (Baxter & King, 1999; Christiano & Fitzgerald, 2003). Our medium- to long-run component—call it c_t—is the following infinite, symmetric, two-sided linear combination of the GDP growth series:

$$c_t = \beta(L)y_t = \sum_{k=-\infty}^{\infty} \beta_k y_{t-k}, \quad \beta_k = \begin{cases} \frac{\sin(k\pi/6)}{k\pi} & \text{for } k \neq 0 \\ 1/6 & \text{for } k = 0. \end{cases} \tag{2}$$

The time series y_t therefore has the decomposition

$$y_t = c_t + s_t = \beta(L)y_t + [1 - \beta(L)]y_t, \tag{3}$$

where s_t includes all the waves of period shorter than 1 year. Since $\beta(1) = 1$, the mean of the GDP growth series, denoted by μ, is retained in c_t, while the mean of s_t is 0. Because c_t and s_t are orthogonal, the variance of y_t is broken down into the sum of a short-run variance and a medium- to long-run variance. The medium- to long-run component c_t is our ideal target MLRG.

Application of equation (2) to the GDP growth rate requires some elaboration. First, as we know, y_t is not observed monthly. Several solutions are possible, including linear interpolation of the missing values or the more sophisticated techniques introduced in Chow and Lin (1971). However, as far as the variable c_t is concerned, the particular interpolation of the missing values in y_t makes no significant difference (see Appendix A.1 for details).

Second, we choose linear interpolation. Precisely, consider the months from 1 to $\tau \leq T$. We assume that 1 is the first publication date of the GDP and denote by T_P the last publication date within $[1\ \tau]$. Moreover, denote by $\hat{\mu}$ the mean of y_t, estimated using its quarterly observations within $[1\ \tau]$. Then define $\check{y}_t$ by setting $\check{y}_t = y_t$ for $t = 1,\ 4,\ \ldots,\ T_P$ and

$$y_t = \hat{\mu} \text{ for } t = -2,\ -5,\ -7, \ldots \text{ and}$$
$$t = T_P + 3,\ T_P + 6,\ T_P + 9.\ \ldots$$

The series $\check{y}_t$ is infinite, with two missing observations for each quarter.

Third, call $y_t(\tau)$ the result of the linear interpolation of $\check{y}_t(\tau)$ and define, for $t = 1, 2, \ldots, \tau$,

$$c_t^*(\tau) = \beta(L) y_t(\tau). \tag{4}$$

Thus, we use the notation $c_t^*(T)$ when the whole interval $[1\ T]$ is considered, or simply c_t^* if no confusion arises.

In Appendix A we show that $c_t^*(T)$ provides a very good approximation of c_t for $13 \leq t \leq T - 12$, where $T = 217$, the size of our sample. Our argument is based on both a simulation exercise and theoretical calculations.

The simulation design is as follows. Mimicking the dynamic structure of y_t, we generate 2,000 time series of length $M = 2N + T + 2N$:

$$y_{j,t}, \quad j = 1, \ldots, 2000;$$
$$t = -2N + 1, \ldots, 0, \quad \underbrace{1, \ldots, T,} \quad T + 1, \ldots, T + 2N.$$

A preliminary simulation determines N as such that the revision

$$c_{j,t}^*(2N + T + 2N) - c_{j,t}^*(N + T + N)$$

is negligible for all $1 \leq t \leq T$ and all $j = 1, \ldots, 2{,}000$ (see Appendix A.2). As a consequence, setting $\hat{M} = N + T + N$, for t belonging to the central subsample of length T, we set $c_{j,t} = c_{j,t}^*(\hat{M})$. Then for $1 \leq q \leq T$, we consider the ratio

$$v_{j,t,T} = \frac{\left(c_{j,t}^*(T) - c_{j,t}\right)^2}{\sigma_j^2},$$

where σ_j^2 is the estimated variance of $c_{j,t}^*(\hat{M})$. Denoting by $V_{t,T}$ its average over 2,000 replications, we find, for $T = 217$,

$$V_{T,T} = 0.14, \quad V_{T-12,T} = 0.008, \quad V_{T-108,T} = 0.0013,$$

thus a close approximation up to $T - 12$, followed by a rapid deterioration (further details on the deterioration are in Appendix A.2).

A similar pattern, as shown in Appendix A.3, results from the theoretical frequency-domain calculation of the ratio

$$\frac{\operatorname{var}\left(c_t^*(T) - c_t\right)}{\operatorname{var}(c_t)},$$

where c_t and $c_t^*(T)$ are obtained by applying, respectively, $\beta(L)$ and the truncated version equivalent to equation (4) (see Appendix A.3), to several monthly stationary processes.

We henceforth take $c_t^*(T)$ as our empirical target for $13 \leq t \leq T - 12$. In section VI, we use $c_t^*(T)$, within $[13\ \ T - 12]$, to compare the performance of NE and other indicators in terms of both fit and ability to signal turning points.

Figure 2 presents the approximation $c_t^*(T)$ for the euro zone GDP, $13 \leq t \leq T - 12$, along with quarterly GDP growth, y_t, where T is June 2005. We see that c_t^* closely tracks the GDP growth (it captures about 70% of the variance of y_t). Figure 2 provides a clear illustration of the smoothing effect of the bandpass filter. Short-run waves are removed, so that observers can distinguish longer oscillations and their turning points. The main task of this paper is a good estimate of c_t at the end of the sample, so that turning points can be detected in real time (see section VI).

FIGURE 2.—$c_t^*(T)$ AND THE MONTHLY QUARTER-ON-QUARTER GDP GROWTH RATE

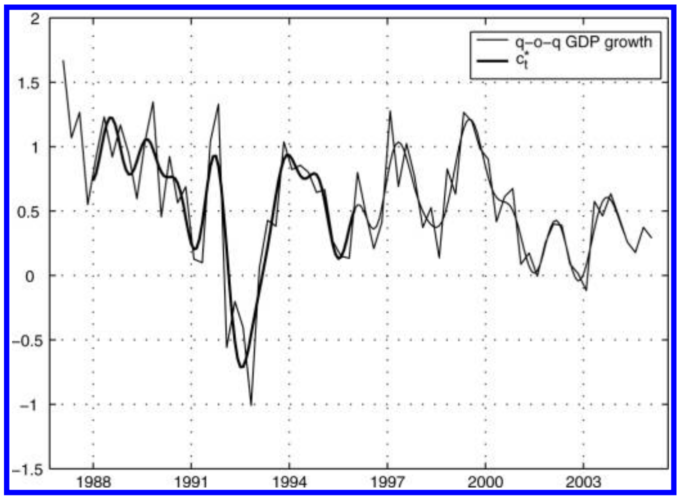

We conclude this section with a few observations about the relationship between MLRG and the year-on-year change of GDP, which is often reported as a measure of medium- to long-run growth. Indicating by $\tilde{y}_t$ the year-on-year change of GDP—the difference between the quarter ending at t and the quarter ending at $t - 12$ (divided by 4 to obtain quarterly rates)—we have

$$\tilde{y}_t = \frac{y_t + y_{t-3} + y_{t-6} + y_{t-9}}{4}.$$

Hence $\tilde{y}_t$ is a moving average of the y series which, unlike MLRG, is one-sided toward the past and hence not centered at t. As a result, $\tilde{y}_t$ is lagging with respect to both y_t and MLRG by several months (precisely four and a half), as is apparent from figure 3.

The phase shift is reduced if we compare MLRG with the future of $\tilde{y}_t$. In section VI D, we show that our indicator, which tracks MLRG, is a good predictor of future year-on-year growth.

IV. Estimation I: Projecting the MLRG on Monthly Regressors

We now begin the construction of NE, our alternative estimate of c_t. A brief outline of our procedure will be helpful to the reader:

- NE is the projection of c_t on a set of regressors, which are linear combinations of the variables in the data set.

FIGURE 3.—$c_t^*(T)$ AND THE MONTHLY YEAR-ON-YEAR GDP GROWTH RATE

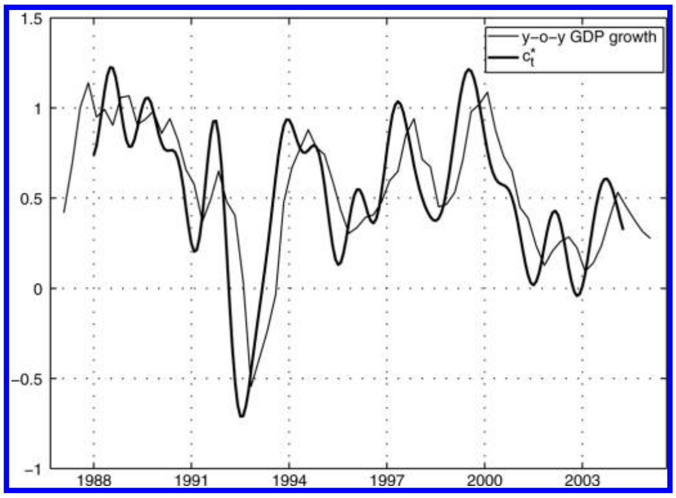

In this section, we give a detailed description of the way we compute the projection once the regressors have been constructed.

- In section V, we construct the regressors. Assuming that our data set can be modeled as a dynamic factor model, we determine the dimension of the factor space—call it r—and project c_t on the first r principal components of the series in the data set, which is a basis of the factor space—call κ_t the projection. Our regressors are generalized principal components, specifically designed to minimize the short-run component. For this reason we call them smooth factors. In section VI, we determine the number of smooth factors as the integer $\bar{s}$ such that the residual of the projection κ_t and that of the projection of c_t on the first $\bar{s}$ smooth factors are approximately of the same size. We show that $\bar{s}$ is smaller than r and that the projection on the $\bar{s}$ smooth factors is substantially smoother than κ_t.
- This projection of c_t on the first $\bar{s}$ smooth factors is NE. In section VI, we provide a detailed assessment of the real-time performance of NE.

The variables in the data set are observed monthly. The regressors, denoted by w_{kt}, $k = 1, \dots, r$, are contemporaneous linear combinations of such variables and are therefore monthly variables. The projection of c_t on the regressors requires some discussion.

The population projection of c_t on the linear space spanned by $w_t = (w_{1t}, \dots, w_{rt})'$ and the constant is

$$P(c_t|w_t) = \mu + \Sigma_{cw}\Sigma_w^{-1}w_t, \tag{5}$$

where Σ_{cw} is the row vector whose kth entry is $\mathrm{cov}(c_t, w_{kt})$ and Σ_w is the covariance matrix of w_t. NE is obtained by replacing the above population moments with estimators:

$$\hat{c}_t = \hat{\mu} + \hat{\Sigma}_{cw}\hat{\Sigma}_w^{-1}w_t. \tag{6}$$

Estimation of $\hat{\Sigma}_w$ is standard once the regressors w_t have been defined. Estimation of $\hat{\Sigma}_{cw}$ is less obvious. The covariances between c_t and w_t can be estimated using w_t and the approximation c_t^*, leaving aside end- and beginning-of-sample data. Alternatively, we can start by estimating the cross-covariances between the quarterly series y_t and w_t. Note that this is possible for any monthly lead and lag,[2] while it is not possible to estimate a monthly autocovariance for y_t. Using such cross-covariances we obtain an estimate of the cross-spectrum between c_t and w_t—call it $\hat{S}_{cw}(\theta)$. Finally, $\hat{\Sigma}_{cw}$ is obtained by integrating $\hat{S}_{cw}(\theta)$ over the band $[-\pi/6\ \pi/6]$.

The results obtained with the two techniques do not differ substantially. The latter is more elegant and therefore has been selected.[3]

V. Estimation II: Constructing the Regressors

A. *Dynamic Factor Models*

The regressors w_{kt} are constructed using techniques from large-dimensional dynamic factor models. We assume that each of the variables x_{it} in the data set is driven by a small number of common shocks, plus a variable-specific (usually called idiosyncratic) component. The idea that this common-idiosyncratic decomposition provides a useful description of macroeconomic variables goes back to the seminal work of Burns and Mitchell (1946) and has been recently developed in the literature on large-dimensional dynamic factor models (see Bai, 2003; Bai & Ng, 2002; Forni, Hallin, Lippi, & Reichlin, 2000, 2001, 2004, 2005—henceforth FHLR; Forni & Lippi, 2001; Stock & Watson, 2002a, 2002b; Kapetanios & Marcellino, 2009).

Large-dimensional factor models estimate a small (relative to the size of the data set) number of "common factors," obtained as linear combinations of the x_{it}'s, which remove the idiosyncratic components and retain the common sources of variation. The innovation of this paper with respect to this literature is a procedure to remove both the idiosyncratic and the short-run components, so that the resulting factors are both common and smooth.

Let us briefly recall the main features of the dynamic factor model we are referring to. Each series x_{it} in the data set is the sum of a common component—call it χ_{it}—that is driven by a small number of common shocks, and an idiosyncratic component, ξ_{it}:

$$\begin{aligned} x_{it} &= \chi_{it} + \xi_{it} \\ &= b_{i1}(L)u_{1t} + b_{i2}(L)u_{2t} + \cdots + b_{iq}(L)u_{qt} + \xi_{it}. \end{aligned} \tag{7}$$

[2] Our method is closely related to the mixed-data sampling (MIDAS) approach. See Ghysels, Sinko, and Valkanov (2007).

[3] We should keep in mind that the series $y_t(T)$, which used construct $c_t^*(T)$, has been obtained by linear interpolation, so that $c_t^*(T)$, strictly speaking, is not covariance stationary, or costationary with the variables in the data set.

FIGURE 4.—COMPARING c_t^* (T) WITH VARIOUS INDICATORS

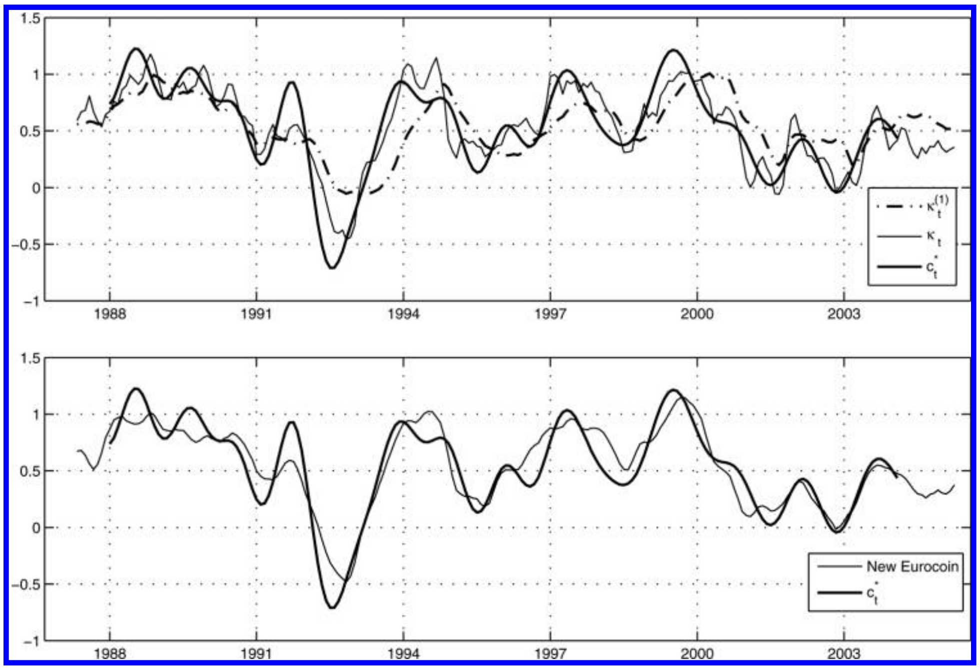

Note: Upper panel: $c_t^*(T)$, thick line, κ_t, solid line, $\kappa_t^{(1)}$, dash-dotted line. Lower panel: $c_t^*(T)$, thick line, New Eurocoin, solid line.

Common and idiosyncratic components are orthogonal at all leads and lags. Moreover, the idiosyncratic components ξ_{it} and ξ_{jt} are mutually orthogonal at all leads and lags for $i \neq j$.[4]

Model (7) is further specified by assuming that the common components χ_{it} can be given the static representation

$$\chi_{it} = c_{i1}F_{1t} + c_{i2}F_{2t} + \cdots + c_{ir}F_{rt}. \tag{8}$$

Under equation (8), different estimators, which are consistent as both the number of observations in each series (T) and the number of series in the data set (n) tend to infinity, have been proposed for the space G_F spanned by the factors F_{jt} (see Stock & Watson, 2002a, 2002b; FHLR, 2005). In particular, Stock and Watson use the first r principal components of the variables x_{it}. Consistent estimates of the common components χ_{it} are obtained by projecting the variables x_{it} on the estimated factors.

Our assumption is that the variables x_{it}, as well as the GDP, are driven by the factors F_{kt}. On the other hand, y_t is a quarter-on-quarter rate of change, whereas the x's, that is, the variables used to construct the factors, are month-on-month rates of change, so that, as we argue in Appendix A.4, representing y_t in terms of the factors transformed by $(1+L+L^2)^2$, that is,

$$y_t = c_{y1}[(1 + L + L^2)^2 F_{1t}] + c_{y2}[(1 + L + L^2)^2 F_{2t}] + \cdots + c_{yr}[(1 + L + L^2)^2 F_{rt}] + \xi_{yt},$$

is parsimonious and fairly reasonable. Thus, the projection of c_t on the factor space will always be estimated by using the transformed regressors $(1 + L + L^2)^2 F_{kt}$, $k = 1, \ldots, r$ (the same transformation will be applied to the smooth factors; see section VB).

Using our data set over the whole sample period [1 T], the dimension of the factor space G_F has been estimated using the Bai-Ng criteria PC_{P1} and PC_{P2} (see Bai & Ng, 2002; we set $r_{\max} = 25$), the result being $r = 12$. Second, c_t has been projected on the first twelve principal components, filtered with $(1 + L + L^2)^2$, the projection being based on equation (6). This projection, denoted by κ_t, is shown in Figure 4, together with $c_t^*(T)$.[5]

We find that κ_t is a fairly good approximation to $c_t^*(T)$. Indeed the R^2 of the regression of $c_t^*(T)$ on κ_t, over the period [13 $T - 12$], is as high as 0.77. However, as Figure 4 shows, κ_t (upper panel, solid line) contains a sizable short-run component.

Smoother versions of κ_t can be obtained by reducing the number of principal components. In fact, the first principal component is quite smooth, but all the others, starting with the second, exhibit substantial short-run oscillations. The projection of c_t on the first principal component (filtered with $(1 + L + L^2)^2$)—call it $\kappa_t^{(1)}$—is plotted together with $c_t^*(T)$ in figure 4 (upper panel, dash-dotted line). A considerable improvement in smoothness is obtained, but first, the R^2 falls to 0.45, and, second, $\kappa_t^{(1)}$ has a systematic phase delay with respect to $c_t^*(T)$. As soon as we project on two principal components, the gain in smoothness almost disappears.[6] The next

[4] This assumption can be relaxed. See for example, FHLR (2000) and Stock and Watson (2002a, 2002b).

[5] We compute κ_t only for the whole sample. Therefore, we do not need the notation $\kappa_t(T)$.

[6] The plots of the projection of c_t on principal components are available on request.

section shows how smoothness can be obtained by a different definition of principal components.

B. *Smooth Factors*

We claim that by conveniently choosing a basis in G_F (different from the twelve principal components used to construct κ_t), we can obtain a projection with approximately the same fit but with a considerably reduced short-run component. Our construction is as follows. Let x_t, χ_t, and ξ_t be the vectors of the variables x_{it}—their common components and their idiosyncratic components, respectively. Let ϕ_{it} be the medium- to long-run component of χ_{it}, precisely $\phi_{it} = \beta(L)\chi_{it}$, and $\psi_{it} = \chi_{it} - \phi_{it}$. For the spectral density matrices, we have

$$S_x(\theta) = S_\chi(\theta) + S_\xi(\theta) = S_\phi(\theta) + S_\psi(\theta) + S_\xi(\theta).$$

Integrating over the interval $[-\pi\ \pi]$, we obtain the following decompositions of the variance-covariance matrix of the x's:

$$\Sigma_x = \Sigma_\chi + \Sigma_\xi = \Sigma_\phi + \Sigma_\psi + \Sigma_\xi.$$

Consistent estimates $\hat{\Sigma}_\chi$, $\hat{\Sigma}_\phi$, and $\hat{\Sigma}_\xi$ can be obtained from the estimates of the spectral density $S_x(\theta)$. (See Forni et al., 2000, for estimates of $S_\chi(\theta)$ and $S_\xi(\theta)$.) $\hat{\Sigma}_\chi$ and $\hat{\Sigma}_\xi$ are obtained by integrating $\hat{S}_\chi(\theta)$ and $\hat{S}_\xi(\theta)$ over $[-\pi\ \pi]$ (see Forni et al., 2005), and $\hat{\Sigma}\phi$ by integrating $S_\chi(\theta)$ over $[-\pi/6\ \pi/6]$.

The matrices $\hat{\Sigma}_\chi$, $\hat{\Sigma}_\phi$, and $\hat{\Sigma}_\xi$ are all we need to construct our smooth regressors. We start by determining the linear combination of the variables in the panel that maximizes the variance of the common component in the low-frequency band, that is, the smoothest linear combination. Then we determine another linear combination with the same property under the constraint of orthogonality to the first, and so on. These generalized principal components (GPC), denoted by W_{kt}, can be obtained by means of the generalized eigenvectors $v_1, \ldots, v_n$ associated with the generalized eigenvalues $\lambda_1, \ldots, \lambda_n$, ordered from the largest to the smallest, of the pair of matrices $(\hat{\Sigma}_\phi,\ \hat{\Sigma}_\chi + \hat{\Sigma}_\xi)$, that is, the vectors satisfying

$$\hat{\Sigma}_\phi v_k = \lambda_k(\hat{\Sigma}_\chi + \hat{\Sigma}_\xi)v_k, \tag{9}$$

with the normalization constraints $v_k'(\hat{\Sigma}_\chi + \hat{\Sigma}_\xi)v_k = 1$ (see Anderson, 1984, theorem A.2.4). The eigenvalue λ_k is equal to the ratio of common-low-frequency to total variance explained by the kth generalized principal component W_{kt}.[7] Of course, this ratio is decreasing with k, so that the greater is k, the less smooth and more idiosyncratic is W_{kt}.

Regarding the projection of c_t on G_F, observe first that since our model has been specified by equation (8), the first r GPC's span the same space G_F spanned by the first r ordinary principal components (see FHLR, 2005), so that projecting c_t on the first r GPCs would give the same result as projecting on the first r PCs—namely, κ_t. However, the variable c_t is by construction very smooth. Therefore, its projection on G_F is likely to be well approximated using only the first, and smoother, GPCs. In other words, a fit almost as good as that obtained by the first r ordinary principal components should be obtained by a substantially smoother approximation.

TABLE 2.—DETERMINING THE NUMBER OF GENERALIZED PRINCIPAL COMPONENTS

Number of GPCs	1	3	5	6	
R^2	0.34	0.50	0.75	0.79	
Number of PCs	1	3	5	6	12
R^2	0.45	0.61	0.71	0.71	0.77

In the second (fourth) row we report the R^2's of the regressions of $c_t^*(T)$ on s generalized (standard) principal components.

VI. Results

A. *The Number of Smooth Factors and the Definition of NE*

Based on the definition of smooth factors and the discussion above, the number of smooth factors is determined by the following procedure. First, we estimate q, the dimension of the white noise u_t (see equation (7)). Applying the criterion proposed in Hallin and Liška (2007), we set $q = 2$. Based on the determination of q, we estimate $\hat{S}_\chi(\theta)$ and $\hat{S}_\xi(\theta)$ as in FHLR (2000, 2005) and compute the covariance matrices $\hat{\Sigma}_\chi$, $\hat{\Sigma}_\phi$, and $\hat{\Sigma}_\xi$ as indicated above. Second, we estimate r, the dimension of G_F, using Bai and Ng's criterion. The result (see section V), is $r = 12$. Then, using $\hat{\Sigma}_\chi$, $\hat{\Sigma}_\phi$, and $\hat{\Sigma}_\xi$, we compute the generalized eigenvectors v_k, $k = 1, \ldots, r$ satisfying equation (9), and the associated GPCs $W_{kt} = v_k'x_t$. Finally, let $\kappa_t^{(s)}$ be the projection of c_t on the first s GPCs, while κ_t, as defined in section V, is the projection of c_t on the first r principal components. In both cases, the principal components are filtered with $(1 + L + L^2)^2$ (see section V), and the projection is based on equation (6). Then let ρ and ρ_s be the R^2's obtained by projecting c_t^* on κ_t and $\kappa_t^{(s)}$, respectively. Starting with $s = 1$, the number of GPCs is increased. We stop when the difference between ρ and ρ_s becomes negligible. Call $\bar{s}$ the number of GPCs so determined.

The fit of the indicator κ_t, that is, the R^2 of the regression of $c_t^*(T)$ on κ_t, over the period $[13\quad T - 12]$, is 0.77 (see section V and table 2). The R^2s of the regression of $c_t^*(T)$ on $\kappa_t^{(s)}$, with s equal to 1, 3, 5, 6, are reported in table 2. Visual inspection shows that the tiny improvement of the fit between five and six factors is not offset by reduced smoothness; thus, we set $\bar{s} = 6$.

The projection of c_t on the first six GPCs is the NE indicator. We use the notation $\hat{c}_t(T)$ for the indicator at time t obtained using the whole sample to estimate the necessary covariance matrices and $\hat{c}_t(\tau)$ when the subsample $[1\ \tau]$ is used.

New Eurocoin and $c_t^*(T)$ are plotted together in figure 4 (lower panel, solid line). The advantage of generalized over

[7] The generalized principal components used in FHLR (2005) are designed for a different purpose. They are obtained using the generalized eigenvectors of the couple $(\hat{\Sigma}_\chi, \hat{\Sigma}_\xi)$.

FIGURE 5.—PSEUDO-REAL-TIME ESTIMATES OF MLRG AT THE END OF THE SAMPLE, OBTAINED WITH NE (UPPER PANEL) AND BP (LOWER PANEL)

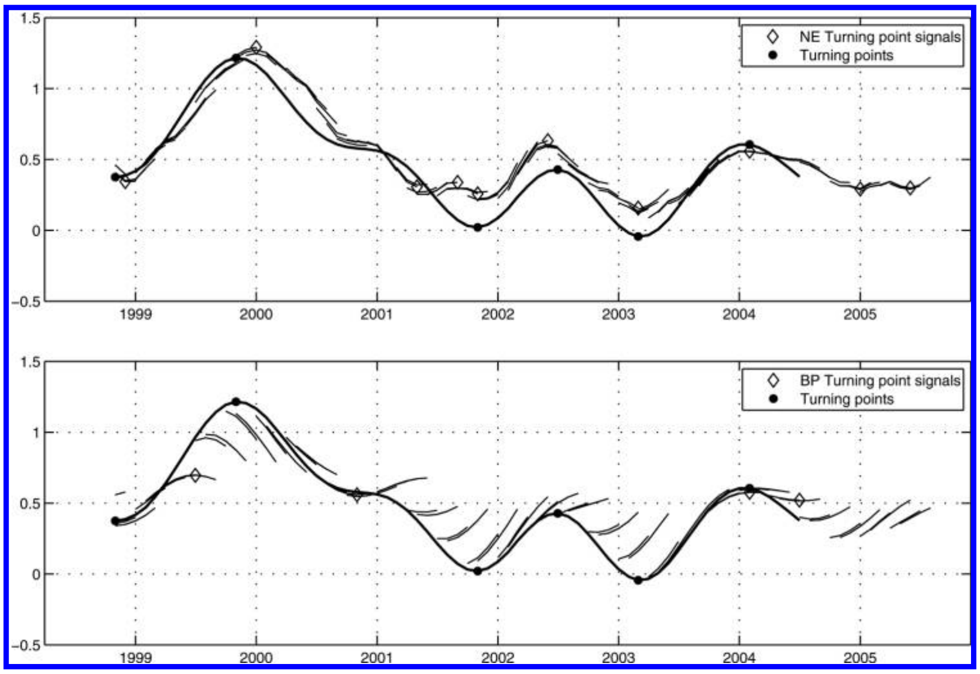

ordinary principal components, when fit and smoothness are jointly considered, is evident by comparing this to κ_t (figure 4, upper panel). The next section contains a systematic comparison, based on a real-time exercise, of NE to κ_t and $c_t^*(t)$ and the Christiano-Fitzgerald version of the bandpass filter.

B. The Real-Time Performance

In this section, we report a pseudo-real-time evaluation of NE. Here "pseudo" refers to the fact that we do not use the true real-time preliminary estimates of the GDP, but the final estimates as reported in GDP "vintage" available in September 2005. The same holds true for all other monthly variables.[8]

In figure 5 and table 5 we report quantities resulting from the estimation of $\hat{c}_t(\tau)$ and $c_t^*(\tau)$, for some values of τ and t running from November 1998 to June 2005 (the last part of the graph in figure 5 and the number of consistent signals in table 5). However, the results using the target $c_t^*(T)$ use data only up to June 2004.

Intuition for the results presented below is provided in figure 5. In the upper graph, the long, continuous line represents $c_t^*(T)$. The short line ending at t represents the three estimates: $\hat{c}_{t-2}(t)$, $\hat{c}_{t-1}(t)$, and $\hat{c}_t(t)$. Therefore, the three points on the short lines over a given t are the first estimate and two revisions of NE at t: $\hat{c}_t(t)$, $\hat{c}_t(t+1)$, and $\hat{c}_t(t+2)$. Revisions of NE at t are due to reestimation of the factors and the projection as new data arrive and are modest. The bullets indicate turning points, and the diamonds indicate turning point signals (see below for formal definitions).

The lower graph shows the corresponding estimates for BP. Each short line represents $c_{t-2}^*(t)$, $c_{t-1}^*(t)$, and $c_t^*(t)$. Clearly the bandpass filter estimates (BP), although very smooth, exhibit a large bias toward the sample mean. NE estimates are more accurate, and the revision errors are smaller.

Let us now establish the formal criteria used in our evaluation. We are interested in (a) the ability of $\hat{c}_t(t) - \hat{c}_{t-1}(t) = \Delta\hat{c}_t(t)$ to signal the correct sign of the change, that is, the sign of $\Delta c_t^*(T)$, as measured by the percentage of correct signs (see Pesaran & Timmermann, 2009); (b) the ability of $\hat{c}_t(t)$ to approximate (nowcast) $c_t^*(T)$, for the period $T - 81 \leq t \leq T - 12$, as measured by the ratio $\sum_{t=T-81}^{T-12}[\hat{c}_t(t) - c_t^*(T)]^2 / \sum_{t=13}^{T-12}[c_t^*(T) - \overline{c_t^*}(T)]^2$, where $\overline{c_t^*}(T) = \sum_{t=13}^{T-12} c_t^*(T)/(T-24)$; and (c) the size of the revision errors after 1 month, as measured by the ratio $\sum_{t=T-81}^{T-1}[\hat{c}_t(t+1) - \hat{c}_t(t)]^2 / \sum_{t=13}^{T-12}[c_t^*(T) - \overline{c_t^*}(T)]^2$.

Our indicator NE, at time t, is compared, using criteria a, b, and c, to three alternative approximations of $c_t^*(T)$, which use information up to time t:

TABLE 3.—END-OF-SAMPLE PERFORMANCE

Indicator	% Correct Prediction of Sign of Δc^*	MS of Nowcast Error/Variance of c^*	MS of Revision Error/Variance of c^*
NE	0.88[a]	0.13	0.005
BP	0.63	0.32	0.061
CFBP	0.66	0.27	0.133
PC	0.62	0.21	0.116

Sample: November 1998–June 2004. The first column reports the percentage of correct signs with respect to those of Δc^*.
[a]In this case, the null of no predictive power is rejected at the 1% significance level.
Source: Pesaran and Timmermann (2009).

[8] A true real-time exercise using the different vintages of the data would be preferable. Unfortunately vintages for most of the monthly series in the data set are not available. We prefer the pseudo-real-time exercise rather than resorting to a mixture of latest vintage and real-time vintage data, which could produce misleading results.

TABLE 4.—CLASSIFICATION OF SIGNALS

	$\Delta\hat{c}_{t-2}(t-1)$	$\Delta\hat{c}_{t-1}(t-1)$	$\Delta\hat{c}_{t-1}(t)$	$\Delta\hat{c}_t(t)$	Consistency	Signal Type
1	−	−	−	+	Yes	**Upturn** at $t-1$
2	+	−	−	+	Yes	Uncertainty
3	−	−	−	−	Yes	Deceleration
4	+	−	−	−	Yes	Slowdown
5	+	+	+	−	Yes	**Downturn** at $t-1$
6	−	+	+	−	Yes	Uncertainty
7	+	+	+	+	Yes	Acceleration
8	−	+	+	+	Yes	Recovery
9	−	−	+	−	No	Trembling deceleration
10	+	−	+	−	No	Downturn at $t-2$ shifted
11	−	−	+	+	No	Missed upturn
12	+	−	+	+	No	Downturn at $t-2$ not confirmed
13	+	+	−	+	No	Trembling acceleration
14	−	+	−	+	No	Upturn at $t-2$ shifted
15	+	+	−	−	No	Missed downturn
16	−	+	−	−	No	Upturn at $t-2$ not confirmed

All possible sign patterns for $\Delta\hat{c}_{t-2}(t-1)$, $\Delta\hat{c}_{t-1}(t-1)$, $\Delta\hat{c}_{t-1}(t)$, $\Delta\hat{c}_t(t)$ are listed and classified into signal types.

- BP
- CFBP: the optimal approximation to the bandpass filter proposed in Christiano and Fitzgerald (2003). Their filter is applied to the interpolated series $y_t(\tau)$, for τ running from $T-81$ to $T-12$. We use the program recommended by Christiano and Fitzgerald[9] in the stationary version, with a long moving average whose coefficients are obtained by inverting an AR model estimated for the interpolated series $y_t(T-81)$, as defined in section III.[10]
- PC: κ_t, the estimate obtained using ordinary principal components.

All the comparisons reported below are fair, in that the same information set is available at any time t for each of the four competing indicators (though different indicators use different subsets).

Table 3 shows that as far as the criteria a, b, and c are concerned, NE scores better than BP and CFBP, the second outperforming the first as regards the nowcast error and the slope changes.[11] As expected, PC performs fairly well as far as b and c are concerned,[12] but is outperformed by NE by criterion a. Hence, NE dominates the other indicators for the criteria we selected.

As regards the nowcast error of BP and CFBP, we must keep in mind that in the pseudo-real-time exercise, the delay of the GDP with respect to t can be 1, 2, or 3 months (see section II). The figures 0.32 and 0.27 in the table can be referred to the average delay, which is 2 months. We should also observe that further publication delay for the GDP may occur in actual NE production.

[9] The code was downloaded from http://www.clevelandfed.org/research/models/bandpass/bpassm.txt.

[10] The observation in note 3 applies to the AR model estimated using the covariances of $y_t(T-81)$.

[11] The advantage of CFBP at the end of the sample vanishes at $T-12$. In other words, the target computed using BP and CFBP is almost identical.

[12] By construction, NE should nowcast as well as PC; hence, the better performance of NE in the second column is due to the particular sample chosen for the real-time exercise.

C. *The Behavior Around Turning Points*

The figures in the first column of table 3, concerning the percentage of correct signs, suggest that NE should perform well in signaling turning points in the target. In the remainder of this section, we explore this issue, but for that purpose, we need precise definitions of *turning point, turning point signal*, and *false signal*.

We define a turning point as a slope sign change in our target, $c_t^*(T)$. We have an *upturn (downturn)* at time t if $\Delta c_{t+1}^*(T) = c_{t+1}^*(T) - c_t^*(T)$ is positive (negative), whereas $\Delta c_t^*(T) = c_t^*(T) - c_{t-1}^*(T)$ is negative (positive). According to this definition, in the subsample involved in the pseudo-real-time exercise, the target exhibits three downturns and three upturns (see figure 5).

Next we define a rule to decide when a slope sign change of our indicator $\hat{c}$ can be interpreted as a reliable signal of a turning point in the target c^*. To this end, we focus on the sign of the last two changes of the current estimate of our indicator $\hat{c}_t(t)$ and the sign of the last two changes of the previous estimate $\hat{c}_t(t-1)$, that is,

$$\text{Current estimate:} \quad \cdots \quad \Delta\hat{c}_{t-1}(t) \quad \Delta\hat{c}_t(t) \tag{10}$$

$$\text{Previous estimate:} \quad \Delta\hat{c}_{t-2}(t-1) \quad \Delta\hat{c}_{t-1}(t-1) \quad \cdots \tag{11}$$

A sign change between $\Delta\hat{c}_{t-1}(t)$ and $\Delta\hat{c}_t(t)$ makes equation (10) a candidate as a *signal at t locating a turning point at* $t-1$. However, we accept the sign change in equation (10) as a turning point signal only if (a) the signal is consistent, that is, the signs of $\Delta\hat{c}_{t-1}(t-1)$ and $\Delta\hat{c}_{t-1}(t)$ coincide, and (b) there is no sign change in equation (11) between $t-2$ and $t-1$, that is, the signs of $\Delta\hat{c}_{t-2}(t-1)$ and $\Delta\hat{c}_{t-1}(t-1)$ coincide.

The reason for conditions a and b is that we want to be strict enough to rule out sign changes that may be caused by unstable estimates rather than by true turning points. Condition a is obvious. Condition b rules out a sign change between $t-1$ and t that follows the opposite change between $t-2$ and $t-1$ in the previous estimate.

Table 4 lists the eight possible consistent (rows 1–8) and the eight possible inconsistent signals (rows 9–16), which

TABLE 5.—REAL-TIME DETECTION OF TURNING POINTS (TP)

Target	Consistent Signals	Uncertainty Signals	TP Signals	TP Signals Excluding Last 12 Months	Correct TP	Correct over Signaled TP	Missed over All TP
NE	81	0	11	8	6	6/8	0
BP	80	0	4	4	1	1/4	5/6
CFBP	68	0	6	6	2	2/6	4/6
PC	81	9	22	20	6	6/20	0

Sample: November 1998–August 2005. The first column reports the number of consistent signals (over a total of 81 signals). The fourth column reports the number of turning point signals when excluding the last 12 signals. The fifth column counts the number of correct turning point signals—those matching the ones in the target. The last shows the percentage of turning points in the target missed by each indicator.

TABLE 6.—HOW TO RELATE THE MONTHLY INDICATOR TO ACTUAL GDP GROWTH

	RMSE with Respect to Different Growth Rates (%)			
Indicator	Quarter-on-Quarter Current Quarter	Year-on-Year Current Quarter	Year-on-Year One Quarter Ahead	Year-on-Year Two Quarters Ahead
NE	0.20	0.18	0.13	0.17
BP	0.32	0.22	0.22	0.25
CFBP	0.32	0.23	0.23	0.25
PC	0.21	0.17	0.15	0.21

Sample: December 1998–June 2005.

in principle could arise. Note that only two of the eight consistent sign changes in equation (10) are classified as turning point signals; those in the first and the fifth rows of table 4—an upturn and a downturn, respectively. Once we have established a rule to identify turning point signals in our indicator, we can compare them with turning points that actually occurred in the target.

We say that an upturn (downturn) signal at t locating a turning point at $t-1$ is false if c^* has no upturns (downturns) in the interval $[t-3, t+1]$, and correct otherwise. With this definition, an upturn signal in $\hat{c}_t$ leading or lagging the true upturn (an upturn in c_t^*) by a quarter or more is false, whereas a 2-month error is tolerated.

Table 5 shows results for the competing indicators in our real-time exercise. Signals are reported up to the last possible date within our data set, which is August 2005, although interestingly, across methods most signals in real time are consistent—all of them for NE and PC. PC also provides nine uncertain signals. NE signals eleven turning points (third column), eight of which are before the last 12 months, where c^* is reliable. The latter include all of the six turning points in the target. The PC indicator correctly signals all turning points but produces many false signals. By contrast, BP and CFBP produce only a few turning point signals, but most of them are false. Overall, the results on turning points are consistent with the figures in the first columns of table 3 and, as regards BP, with figure 5.

D. *Forecasting Properties of the Indicator*

In section III, we argued that we should expect a close match between NE and the GDP growth rate once the latter is smoothed with a moving average such as the one induced by the year-on-year transformation and adjusted for the phase shift.[13] This is confirmed by the results shown in the last two columns of table 6. While the root mean squared error of NE with respect to quarter-on-quarter GDP growth (first column) is 0.20, the same statistic with respect to year-on-year growth (divided by 4) is 0.18 (second column) and decreases to 0.13 and 0.17 when we adjust for the phase shift by considering future year-on-year growth (third and fourth column).[14] None of the competing indicators have similar forecasting properties.

TABLE 7.—PSEUDO-REAL-TIME FORECAST PERFORMANCE

	Target Growth Rates (%)		
Model	Quarter-on-Quarter Current Quarter	Year-on-Year Current Quarter	Year-on-Year One Quarter Ahead
NE	0.20	0.18	0.13
AR (AIC)	0.29***	0.16	0.17
AR (BIC)	0.29***	0.16	0.17
ARMA (AIC)	0.31**	0.18	0.21**
ARMA (BIC)	0.30**	0.17	0.19*
Random walk	0.31**	0.18	0.19

Sample: December 1998–June 2005. The first column reports the root mean square forecast error with respect to current quarter-on-quarter GDP growth rate, the second with respect to current year-on-year GDP growth rate, the third with respect to next quarter year-on-year GDP growth rate. NE is the New Eurocoin forecast obtained using the monthly data set with information updated at most up to the last month of the current quarter. The AR and ARMA models are selected at each step according to their in-sample performance (in parentheses the selection criterion used) and are estimated on the quarterly GDP series. ***, **, *: Rejection of the null of equal forecast accuracy at 1%, 5% or 10%, respectively, according to Diebold and Mariano (1995) test.

To better gauge the forecasting ability of NE, we compare it with univariate ARMA models of quarterly GDP growth, selected by standard in-sample criteria. Such models are often used as benchmarks in forecasting studies (Stock & Watson, 2002b).

As shown in table 7, for quarter-on-quarter GDP growth (first column) and for the year-on-year growth rate one quarter ahead (third column), the forecast error of the indicator is far lower than those obtained either with the ARMA or with the random walk.

[13] A similar idea is exploited in Cristadoro et al. (2005) to motivate their result that a core inflation indicator obtained as a smoothed projection of CPI inflation on factors is a good forecaster of the CPI headline inflation.

[14] Obviously we can compare our monthly indicator with actual GDP growth rates only at the end of each quarter.

VII. Summary and Conclusion

Our coincident indicator NE is a timely estimate of the medium- to long-run component of the euro-area GDP growth. The latter, our target, has been defined as a centered, symmetrical moving average of GDP growth, whose weights are designed to remove all period fluctuations shorter than 1 year. As observed in section III, the target, which has a rigorous spectral definition, leads the "popular" measure of medium- to long-run change—namely, year-on-year GDP growth—by several months.

We avoid the end-of-sample bias typical of two-sided filters by projecting the target onto suitable linear combinations of a large set of monthly variables. Such linear combinations are designed to discard useless information (idiosyncratic and short-run noise) and retain relevant information (common, cyclical and long-run waves). Both the definition and the estimation of the common, medium- to long-run waves are based on recent factor model techniques. Embedding the smoothing into the construction of the regressors is, in our opinion, an important contribution of this paper.

The performance of NE as a real-time estimator of the target has been presented in detail in section VI. The indicator is smooth and easy to interpret. In terms of detection of turning points, it scores much better than the competitors that naturally arise as estimators of the medium- to long-run component of GDP growth in real time. The reliability of the signal is reinforced by the fact that the revision error of our indicator is fairly small as compared with the competitors. We have also shown that NE is a very good forecaster of year-on-year GDP growth one and two quarters ahead; it also scores well in forecasting quarter-on-quarter GDP growth, with an RMSE of 0.20, which ranks well even in comparison with best-practice results.

REFERENCES

Altissimo, F., A. Bassanetti, R. Cristadoro, M. Forni, M. Hallin, M. Lippi, L. Reichlin, and G. Veronese, "A Real Time Coincident Indicator for the Euro Area Business Cycle," CEPR discussion paper no. 3108 (2001).

Anderson, T. W., *An Introduction to Multivariate Statistical Analysis* (Hoboken, NJ: Wiley, 1984).

Bai, J., Inferential theory for factor models of large dimension. *Econometrica* 71 (2003), 135–171.

Bai, J., and S. Ng, "Determining the Number of Factors in Approximate Factor Models," *Econometrica* 70 (2002), 191–221.

Baxter, A., and R. G. King, "Measuring Business Cycle Approximate Band-Pass Filters for Economic Time Series," this REVIEW 81 (1999), 575–593.

Bernanke, B., and J. Boivin, "Monetary Policy in a Data-Rich Environment," *Journal of Monetary Economics* 50:3 (2003), 525–546.

Brockwell, P. J., and R. A. Davis, *Time Series: Theory and Methods*, 2nd ed. (New York: Springer-Verlag, 1991).

Burns, A. F., and W. C. Mitchell, *Measuring Business Cycles* (Cambridge, MA: NBER, 1946).

Chow, G. C., and A. Lin, "Best Linear Unbiased Interpolation, Distribution, and Extrapolation of Time Series by Related Time Series," this REVIEW 53 (1971), 372–375.

Christiano, L. J., and T. J. Fitzgerald, "The Band-Pass Filter," *International Economic Review* 84 (2003), 435–465.

Cristadoro, R., M. Forni, L. Reichlin, and G. Veronese, "A Core Inflation Indicator for the Euro Area," *Journal of Money Credit and Banking* 37:3 (2005), 539–560.

Diebold, F., and R. S. Mariano, "Comparing Predictive Accuracy," *Journal of Business and Economic Statistics* 13 (1995), 253–263.

Forni, M., M. Hallin, M. Lippi, and L. Reichlin, The Generalized Dynamic Factor Model: Identification and Estimation, this REVIEW 82 (2000), 540–554.

Forni, M., M. Hallin, M. Lippi, and L. Reichlin, "Coincident and Leading Indicators for the Euro Area," *Economic Journal* 111 (2001), 62–85.

——— "The Generalized Dynamic Factor Model: Consistency and Rates," *Journal of Econometrics* 119 (2004), 231–255.

——— "The Generalized Dynamic Factor Model: One-Sided Estimation and Forecasting," *Journal of the American Statistical Association* 100 (2005), 830–840.

Forni, M., and M. Lippi, "The Generalized Dynamic Factor Model: Representation Theory," *Econometric Theory* 17 (2001), 1113–1141.

Ghysels, G., A. Sinko, and R. Valkanov, "MIDAS Regressions: Further Results and New Directions," *Econometric Reviews* 26 (2007), 53–90.

Giannone, D., L. Reichlin, and L. Sala, "Tracking Greenspan: Systematic and Unsystematic Monetary Policy Revisited," CEPR discussion paper no. 3550 (2002).

Hallin, M., and R. Liška, "Determining the Number of Factors in the Generalized Factor Model," *Journal of the American Statistical Association* 102 (2007), 603–617.

Kapetanios, G., and M. Marcellino, "A Parametric Estimation Method for Dynamic Factor Models of Large Dimensions," *Journal of Time Series Analysis* 30 (2009), 208–238.

Pesaran, M. H., and A. Timmermann, "Testing Dependence among Serially Correlated Multi-Category Variables," *Journal of the American Statistical Association* 104 (2009), 325–337.

Stock, J. H., and M. W. Watson, "Forecasting Using Principal Components from a Large Number of Predictors," *Journal of the American Statistical Association* 97 (2002a), 1167–1179.

——— "Macroeconomic Forecasting Using Diffusion Indexes," *Journal of Business and Economic Statistics* 20 (2002b), 147–162.

Valle e Azevedo J., S. J. Koopman, and A. Rua, "Tracking the Business Cycle of the Euro Area: A Multivariate Model-Based Bandpass Filter," *Journal of Business and Economic Statistics* 24 (2006), 278–290.

Zarnowitz, V., and A. Ozyildirim, "Time Series Decomposition and Measurement of Business Cycles, Trends and Growth Cycles," *Journal of Monetary Economics* 53 (2006), 1717–1739.

This article has been cited by:

1. João Valle e Azevedo. 2010. A multivariate band-pass filter for economic time series. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* no-no. [CrossRef]