

Schreiber, Sven

**Conference Paper**

## Forecasting business-cycle turning points with (relatively large) linear systems in real time

Beiträge zur Jahrestagung des Vereins für Socialpolitik 2013: Wettbewerbspolitik und Regulierung in einer globalen Wirtschaftsordnung - Session: Forecasting, No. E20-V2

**Provided in Cooperation with:**

Verein für Socialpolitik / German Economic Association

*Suggested Citation:* Schreiber, Sven (2013) : Forecasting business-cycle turning points with (relatively large) linear systems in real time, Beiträge zur Jahrestagung des Vereins für Socialpolitik 2013: Wettbewerbspolitik und Regulierung in einer globalen Wirtschaftsordnung - Session: Forecasting, No. E20-V2, ZBW - Deutsche Zentralbibliothek für Wirtschaftswissenschaften, Leibniz-Informationszentrum Wirtschaft, Kiel und Hamburg

This Version is available at:

<https://hdl.handle.net/10419/79709>

**Standard-Nutzungsbedingungen:**

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

**Terms of use:**

*Documents in EconStor may be saved and copied for your personal and scholarly purposes.*

*You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.*

*If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.*

# Forecasting business-cycle turning points with (relatively large) linear systems in real time\*

by Sven Schreiber<sup>†</sup>

this preliminary version: March 2013<sup>‡</sup>

## Abstract

The detection of business-cycle turning points is usually performed with non-linear discrete-regime models such as binary dependent variable (e.g., probit or logit) or Markov-switching methods. The probit model has the drawback that the continuous underlying target variable is discretized, with a considerable loss of information. The Markov-switching approach in general presupposes a non-linear data-generating process, and the numerical likelihood maximization becomes increasingly difficult when more covariates are used. To avoid these problems we suggest to first use standard linear systems (subset VARs with zero restrictions) to forecast the relevant underlying variable(s), and in a second step to derive the probability of a suitably defined turning point from the forecast probability density function. This approach

---

\*This paper grew out of a project funded by the German federal ministry of finance (BMF), and we thank Daniel Detzer for excellent research assistance. For helpful comments we thank Jörg Breitung, Dieter Nautz, Thomas Theobald, and Jürgen Wolters.

The entire estimation and evaluation algorithm has been implemented as a program using the *gretl* econometrics software, see Cottrell and Lucchetti (2010).

<sup>†</sup>University of Hamburg (Institute for Growth and Fluctuations), and Macroeconomic Policy Institute (IMK) Düsseldorf. Email: svetosch@gmx.net, postal address: Von-Melle-Park 5, D-20146 Hamburg, Germany.

<sup>‡</sup>The empirical analysis of this paper is currently being updated with more recent data for Germany, and an analogous analysis for US data. This update could not be finished for the submission deadline, but is expected to be completed by the end of March.

will never fail numerically. We also discuss and show how this approach can be used in real time in the presence of publication lags and to capture features of the data revision process, and we apply the method to German data; the event of the recent “Great Recession” is first signalled in June 2008, several months before the official published data confirms it (but due to publication and recognition lags it is found after it already began in reality).

Keywords: density forecasts, business-cycle turning points, real-time data, now-casting, great recession, Germany

JEL codes: C53 (forecasting models), E37 (cycle forecasting)

## 1 Introduction

In this paper we suggest a linear system approach to the old problem of business-cycle turning-point prediction, taking into account the data availability and revision problems in real time. This approach differs from the usual methods used to detect business-cycle turning points, which is usually done with non-linear discrete-regime models such as binary dependent variable (e.g., probit) or Markov-switching methods.<sup>1</sup> The general idea is that we use a linear (system) model to predict the continuous output variable several ( $h$ ) steps ahead. Instead of using the point forecast we then use the estimated probability density functions (pdf) of the forecast to calculate the probability of a realization below the previously defined recession threshold (or above a certain boom threshold). As described in more detail below, using monthly available data we employ a threshold of a negative cumulated growth rate of -1% over a time span of five months to call a recession (we ). This threshold can be naturally interpreted according to the “triangle” approach by Harding and Pagan (2002), where episodes can be very short and intense or more drawn out and gradual to qualify as recessions. Our choice obviously requires an average output

---

<sup>1</sup>While a subclass of Markov-switching models can actually be regarded as linear, the general case –for example with regime-dependent dynamics– yields non-linear models (see Krolzig, 1997).

decline of at least 0.2% (approximately) per month.<sup>2</sup>

Our approach has the following advantages: first, because of the linearity of the estimator there exist analytical solutions for every step of the estimation algorithm, which guarantees the existence of a turning point forecast for every conceivable data constellation. In contrast, for the typically applied non-linear model classes it is always possible that the maximization of the likelihood function could fail numerically. In our method the only non-linear step appears at the end of the procedure, namely to find the p-value of a normal distribution, and this standard problem obviously cannot cause any numerical problems. Secondly, we define the real-time variables in our multivariate system such that we also model the revision process of consecutive data publications, by keeping (some of) the superseded data publications in the econometric system. This means for example that we do not take any new data publications at face value, but our forecasts will also cover the yet unreleased revised values.<sup>3</sup>

Finally, in relation to the estimation of probit models there is reason to hope that the direct forecast of the continuous output variable (industrial production) is better able to exploit the information contained in the data. After all, in order to fit the probit model it is necessary to reduce the target variable to a binary regime variable, which discards quite a bit of information.

There are also drawbacks of our approach which have to be acknowledged: In order to make use of a broad information set, we use relatively large VARs as the starting point for our forecasting model. These initial models are then reduced with automated coefficient restrictions following the general-to-specific method, but the initial models suffer from the curse of dimensionality, i.e. the combination of too many variables and lags may exceed the available degrees of freedom. In a scenario with only quarterly data and only short

---

<sup>2</sup>Often two consecutive quarters (i.e., six months) are implicitly considered as a minimum recession duration. However, not all “official” recessions fulfill this criterion, and thus we chose five months instead to shorten our turning point recognition lag.

<sup>3</sup>This is especially helpful if early data publications are systematically biased, see below for the evidence. But even if no revisions are expected on average, the variance of the revisions also plays a role for the forecast pdf; our approach accounts for that.

available revision data histories our approach may therefore not be the most suitable – although in a situation without long revision data histories it is not entirely clear why we would want to use real-time methods in the first place.<sup>4</sup>

A more fundamental restriction is given by the fact that our model presupposes linear time series processes. Thus if the DGP were actually non-linear, our forecasting models would only be justified as an approximation. On the other hand, the same variables that are used in the forecasting model are often analyzed with linear models in other macroeconomic contexts, and thus linear models seem to be perfectly reasonable. The fact that we find no significant deviations from normally distributed innovations (see below) also indicates that nonlinearities do not seem very important.

**Related literature** An early example that linear prediction models can be applied to the problem of turning point determination with continuous target variables is given in Stock and Watson (1993). For examples of non-linear turning point applications see Ng (2012), or Laytona and Katsuura (2001), and especially for Markov-switching approaches dealing with real-time data Hamilton (2011), or Nalewaik (2012).

Österholm (2012) uses a similar approach in the sense of applying a linear model (Bayesian VAR), and working with the predictive densities. However, our conclusions are quite different from his pessimistic outlook, since our model seems to work well. The main difference appears to be the size of the model, as Österholm focuses on trivariate systems. Our findings highlight the importance of a broad information set. For a further discussion of Bayesian approaches see section 3.4.

In the rest of the paper, we first present the data, then the following section formally describes the statistical approach, including the notational conventions for the real-time data with publication lags and revisions, and (in subsection 3.5) we explain how we construct

---

<sup>4</sup>A different possibility to overcome this curse-of-dimensionality problem would be a Bayesian estimation. However, we do not have any priors about the coefficient values, and thus using Bayesian methods would then arguably be merely another device to reduce the model dimensions without the standard decision-theory foundation. Hence, as long as our approach works, we see it as a substitute.




Figure 1: German industrial production, five-month moving average growth. The time axis refers to the period of data publication, i.e. three months after the reference period for this second-release series. The straight blue line is at -1%.

and use the density forecasts to infer the recession probabilities. Afterwards the in-sample properties of the estimated models are presented and we report the predictors that turn out as the most relevant ones for the German industrial production forecasts. Then, in section 6, we present the pseudo out-of-sample results for turning-point / recession probability forecasts, which cover the period of the recent “great recession”, before we conclude.

## 2 Basic data properties

In Germany the distribution of the growth rates of industrial production is quite concentrated around zero, and thus there is a relatively fine line between recessions and expansions. In figure 1 we present the cumulated five-month growth rates of industrial production for each publication month (i.e. a datapoint refers to the time when it was published, which differs from the time of the corresponding economic activity, see below for further definitions and explanations), along with a horizontal line showing the -1% threshold. It can be seen that this threshold is crossed quite often and that the volatility of the series is quite large relative to the range.

Thus we expect that for German data it will be quite difficult to forecast the state of




Figure 2: Revisions of industrial production data. Difference of the second release published in a certain period and the first release published one period earlier (both in logs).

the business cycle at longer horizons. Nevertheless, even if in a certain period the ex-post perspective already indicated that the economy had experienced a turning point, the implied recession or boom probabilities from the forecasting (or nowcasting) exercise would still be interesting for economic policy as well as private agents, since that is an additional and different piece of information.

Since the central topic of this paper is the real-time perspective, we report in figure 2 the most important part of the publication history of the industrial production data, namely the realizations of the first revisions (differences of logs). It can be seen that the magnitude of revisions is sometimes substantial. Note also that the first release has been traditionally biased downward because in a simple autoregression the constant term turns out as significantly positive (with the coefficients of all autoregressive terms being negative). This may have changed very recently from about 2011 onwards as can be seen at the end of the graph.

### 3 The statistical real-time framework

#### 3.1 Notation

While the underlying econometrics is fairly common, it is not so common to use such VAR models for real-time data with publication lags and possible revisions, and thus we introduce some useful definitions and notation to capture these properties.

First of all, in the presence of publication lags and revisions we must distinguish two time axes: first the time period when a value is published, and secondly the time period when the measured activities or events actually took place. Throughout, we call the former “publication period” and the latter the “reference period”.<sup>5</sup> Since the publication time determines the available information set on which forecasts can be based, the econometrically relevant time index in our model will correspond to the publication periods.

We denote an individual observation for use in the statistical model below by

$$x_t(r, p), \quad r = 1, 2, \dots, rmax; p = 0, 1, 2, \dots; t = 1, 2, \dots \quad (1)$$

which means that this is the  $r$ -th publication ( $r$  for release) of the measured magnitude  $x$  which was published in period  $t$ . (Thus the  $r$ -th publication is the  $(r - 1)$ -th published revision.) The value  $rmax$  represents a cutoff point which is arbitrarily chosen to limit the number of regressors later in the model; in theory the revision process can go on forever.

The number  $p$  tells us the publication lag of this measurement, i.e. the number of periods it takes the statistical agencies to collect the input data and to publish the first result of their measurement. The publication lag is of course different across variables, but we assume that it is fixed over time for each variable. This is a potentially restrictive assumption, especially if we were using weekly or daily data, but for our monthly dataset it holds. This implies that the observation of the first publication  $x_t(r = 1, p)$  corresponds

---

<sup>5</sup>In the real-time literature the term “vintage” is also used and will typically correspond to our publication time. It seems a little ambiguous, however, because sometimes it may describe the revision number instead. Thus we avoid this term.



to the reference period (of economic activity)  $t - p$ . For later revisions of the data, the discrepancy between the underlying reference period and the time of publication also depends on the number of releases:  $t - p - (r - 1)$ .<sup>6</sup>

An example may be helpful, using as a final piece of notation the mapping  $t(\cdot)$  between the human-readable labels of time periods such as “2010m4” and the integer-valued time period index. Let  $x_t(r = 1, p = 2)$  be the first publication of the (log) industrial production, then  $x_t(2, 2)$  is the second publication (first revision), the element  $x_{t(2000m3)}(1, 2)$  contains the value of industrial production for the reference month of 2000m1, because the publication lag of the first publication is two months. The element  $x_{t(2000m4)}(2, 2)$  also contains a value for the same reference month 2000m1, namely the second publication (first revision).

The important feature of our real-time data representation is that each element – for example  $x_{t(2000m3)}(1, 2)$  – itself is fixed over time, even though the value of the economic variable for the implied reference period is revised later on, because the revised data are explicitly recorded in separate variables. It is this representation which enables us to use standard VAR tools; in terms of the included information set we retain the outdated data publications in our model in the form of lagged variables, which means that we implicitly also have a model of the revision process itself and exploit any systematic movements in there for our forecasts.

The publication lag  $p$  will actually be zero for many variables, especially for financial-market data which are immediately available. The maximum publication lag in our application will be 2 periods (months), for example concerning the industrial production index.

For certain model variants later on we collect all stationary variables in a (column) vector  $\mathbf{s}_t$  of dimension  $n_0 \times 1$  and the integrated variables in another vector  $\mathbf{d}_t$  of dimension

---

<sup>6</sup>We assume that revisions happen continuously in every period after the initial publication. This can be done without loss of generality, because in case no actual revision happens in a certain period we simply create a dummy observation with unchanged data:  $x_{t+1}(r + 1, p) = x_t(r, p)$ .

$n_1 \times 1$ .<sup>7</sup> Note that the sum of the vector lengths  $n_0 + n_1$  will in general exceed the number of “conceptual variables” in the model, because the different releases are included as separate statistical-technical variables.

Since we plausibly assume that the revision process of the data does not introduce any additional non-stationary behavior, the collection of the releases of a certain variable  $x_t(r, p)$ ,  $r = 1, 2, \dots, rmax$ , will either belong entirely into  $\mathbf{s}_t$  or exclusively into  $\mathbf{d}_t$ . A corollary of this assumption is that the different releases of a certain economic variable which is integrated will automatically be co-integrated with unit coefficients.

The union of all variables is denoted in two separate ways:

$$\begin{aligned} \mathbf{z}_t^0 &= (\mathbf{s}'_t, \Delta \mathbf{d}'_t)' \\ \mathbf{z}_t^1 &= (\mathbf{s}'_t, \mathbf{d}'_t)' \end{aligned} \quad (2)$$

In the model using  $\mathbf{z}_t^0$  the integrated variables are thus differenced to render all variables stationary, whereas the vector  $\mathbf{z}_t^1$  is a mixture of stationary and integrated variables in levels. Both vectors have dimension  $(n_0 + n_1) \times 1$ .

Because the topic of this paper is the detection of turning points of economic activity, we are ultimately interested only in the output variable (measured by log industrial production in our monthly data), which we denote with  $y_t(r, p = 2)$ , already using the knowledge about the publication lag of two months. Since (log) output is non-stationary, the collection  $y_t(r, p)$  (for all considered releases  $r$ ) is a part of  $\mathbf{d}_t$ .

### 3.2 The statistical VAR model

The underlying two systems have the following general form,

$$\mathbf{z}_t^i = \sum_{k=1}^K A_k \mathbf{z}_{t-k}^i + \boldsymbol{\mu} + \mathbf{u}_t, \quad i \in \{0, 1\}, t = 1..T(\tau), \quad (3)$$

---

<sup>7</sup>The letter d stands for difference stationary. Other types of non-stationarity (trend stationarity, deterministic breaks) are not considered in our setup.

which is a completely standard VAR model in terms of econometrics, with multivariate white-noise residuals that are assumed to be normally distributed:  $\mathbf{u}_t \sim N(0, \Omega)$ . In order not to overload the notation we do not explicitly differentiate the model parameters according to the model variant  $i$ , because the context should make that sufficiently clear.

The sample end point  $T(\tau)$  is not fixed here because we conduct a pseudo-out-of-sample evaluation with the evaluation sample running from  $\tau : 2007m1$  to  $\tau : 2010m11$ . This means that the systems are re-estimated for every publication period in the evaluation sample; an alternative approach could be to hold the parameter estimates from an earlier sample as fixed and make use of newer data only for the forecasts. We can also easily consider a rolling window for estimation by letting the starting period of the sample move in parallel to the end period; this may be desirable in practical on-going applications of our method if structural breaks are suspected in the beginning of the sample. In our formal presentation in this paper, however, we assume a time-invariant data-generating process (DGP).

We set  $rmax = 2$  and thus only consider the first two publications. For the given data it turned out that this captures most of the revision process, see Schreiber, Theobald, Proaño, Stephan, Rietzler, and Detzer (2012).

The vectors  $\mathbf{z}_t^i$  contain 21 elements: first and second publications of the variables industrial production, domestic orders, foreign orders, the (consumer) price index, as well as the variables for which no revision history is available: oil prices, CDAX stock index, REX bond market value index, vacancies, the Euribor interest rate, the yield spread for corporate bonds, five different interest rate term spreads, a business climate (IFO) and expectations index (ZEW). Variables are appropriately transformed; depending on the model variant, the non-stationary variables are either included in log-levels or transformed by log-differencing to growth rates. In any case, some variables are treated as stationary, namely the interest rate spreads and the Ifo and ZEW indicators.

The log-level variant using  $\mathbf{z}_t^1$  implicitly takes into account the existence of multiple coin-

tegrating relations between variables. E.g. the first and second publications of the same concept should be expected to be cointegrated, and also between industrial production and new incoming orders a cointegration relationship is normally expected. The second model variant with  $\mathbf{z}_t^0$  by construction includes only stationary variables, because such specifications are often used for short-term forecasts of inflation in the literature. Here the levels information is neglected, i.e. the implicit error correction terms of the level model are not included, but this could be negligible for relatively short time horizons. The differencing transformation of the variables may provide a certain stability of the predictions if there are shifts in the level relations (Clements and Hendry, 1999). Ultimately, the choice between the model variants is an empirical question.

The available sample in this model framework is limited by the requirement that the data source actually must provide first and second publications even for the earlier reference periods. In our case the dataset (mainly from the Bundesbank real-time database) supports a starting date of 1995.

### 3.3 Sequential model reduction

Our estimation approach for a given  $\tau$  begins with the usual OLS estimates of each equation of this system. Due to the initial VAR form this is also the efficient system estimator, but this is irrelevant here because of the following model reduction. In our case the maximum number of lags would be  $K = 6$ . Since large systems typically require fewer lags than small systems, this half-year period should be sufficient. For the choice of  $K$  there are various possibilities, for example we chose  $K = 4$  based on the best  $\bar{R}^2$  in the levels equation for the first publication of industrial production. For the specification with growth rates this translates to  $K = 3$  because obviously there the fourth lagged level appears implicitly in the differences.

The next steps are the automatic sequential elimination of insignificant regressors to im-

prove the overall precision of the estimation.<sup>8</sup> The regressor with the highest p-value of the usual two-sided t-statistic is removed and the equation is re-estimated (again with OLS). This procedure is repeated until no t-statistic has a p-value above the selected threshold  $p_{cutoff}$ , where we set  $p_{cutoff} = 0.05$ . Overall, and roughly speaking, in our case this method removes at least half of the regressors.

Once all the equations have been reduced in this way, the entire system is re-estimated with the SUR method, i.e. in an efficient manner, taking into account the cross-equation covariance matrix of the residuals. This feasible-GLS estimator of course also has analytical solutions and thus never fails. The final reduced system then provides the basis for the real-time forecasts based on the current margin  $\tau$ .

We denote the final SUR-estimated coefficients with the imposed zero restrictions through the described sequential model reduction by  $\hat{B}_k$ ; in principle it could happen that some of these matrices are actually zero. The resulting estimated subset-VAR model is thus:

$$\mathbf{z}_t^i = \sum_{k=1}^K \hat{B}_k \mathbf{z}_{t-k}^i + \hat{\boldsymbol{\mu}} + \hat{\mathbf{u}}_t, \quad i \in \{0, 1\}, t = 1..T(\tau), \quad (4)$$

with a corresponding standard estimate of the residual covariance matrix  $\hat{\Omega}$ .

### 3.4 Discussion of alternative solutions to the curse of dimensionality

The problem of scarcity of degrees of freedom in large dynamic models can be solved in different ways, all of which are a kind of "shrinkage" estimation methods. Firstly, the dimension of the data space can be reduced by the extraction of common factors or principal components. These factor models are an active field of research and application, for a recent application to German real-time data see Schumacher and Breitung (2008). However, there it is not easy to determine the relevance of certain predictors. Alternatively

---

<sup>8</sup>For simplicity we perform this elimination separately for each equation; it would be straightforward in principle to conduct it on the system level with the seemingly unrelated regressions (SUR) method, but there is no indication in the literature that an OLS-based procedure would actually be inferior.

one can reduce the parameter space: For example, if there is a priori information on the coefficients, Bayesian methods can be applied. Recently Banbura, Giannone, and Reichlin (2010) showed that big BVARs can have competitive prediction qualities, and with a suitable choice of priors there exists no initial degree-of-freedom problem. However, it is often dubious whether there actually is a priori information about the coefficients; rather it seems that BVARs are simply used in the literature as a convenient shrinkage tool. In the latter case, a more classical approach in the spirit of the general-to-specific methodology is also possible by successively removing the insignificant predictors in order to increase the efficiency of the forecast. This "Subset VAR" approach with automated zero restrictions is what we use in this paper.

### 3.5 Density forecasts and estimating the recession probability

Having estimated the parameters of the subset VAR with the sample ending in  $T(\tau)$ , we can produce  $h$ -step point forecasts in the usual way, which we denote by  $\hat{\mathbf{z}}_{T+h}^i$ . Note that an  $h$ -step forecast corresponds to the reference period  $T + h - p - (r - 1)$ ; if  $h < p + r - 1$  this is sometimes called a backcast since the reference period precedes the current publication period, and the term nowcast is used for  $h = p + r - 1$ .

For multistep forecasts ( $h > 1$ ) the forecasts for lower horizons are needed, e.g.  $\hat{\mathbf{z}}_{T+2}^i = \hat{\mathbf{B}}_1 \hat{\mathbf{z}}_{T+1}^i + \sum_{k=2}^K \hat{\mathbf{B}}_k \mathbf{z}_{T+2-k}^i + \hat{\boldsymbol{\mu}}$ . Therefore in this approach we implicitly forecast all the variables contained in the model, even though we are ultimately only interested in the forecast of the output variable (industrial production), and its probability density. Such an iterative multistep forecast is well known to be optimal if the true model is a VAR.

But of course the future realizations of the variables are random due to the innovations  $\mathbf{u}_{T+1}, \dots, \mathbf{u}_{T+h}$ . We denote the (estimated) multivariate probability distribution function (pdf) corresponding to the forecast errors by  $\hat{f}_{T+h}^i$ . Due to the normality of the innovations and the linearity of the model this pdf is also Gaussian, centered around the point forecast  $\hat{\mathbf{z}}_{T+h}^i$ , and the forecast error is in general a moving average of the future innovations which

yields standard textbook formulae for its covariance matrix  $\hat{\Psi}_h^i$ , see for example Lütkepohl (2007, e.g. section 3.5). Therefore the following expression is directly operational:

$$\hat{f}_{T+h}^i : N(\hat{\mathbf{z}}_{T+h}^i, \hat{\Psi}_h^i). \quad (5)$$

In the following sections we must distinguish between the subset VAR in (log) levels and growth rates, because while for the log-level specification with  $\mathbf{z}_t^1$  we “only” need the corresponding diagonal element of  $\hat{\Psi}_h^i$  to get the forecast error variance for future output, for the specification in growth rates with  $\mathbf{z}_t^0$  (and thus with  $\Delta y_t(r, p)$ ) we need to re-cumulate the predicted growth rates to arrive at the predicted output level.

For the estimation of the forecast density and in particular for its covariance  $\hat{\Psi}_h^i$  we consider only the residual variance  $\Omega$  and ignore the additional estimation uncertainty of the parameters. However, this should be important only in the tails of the distributions, and these are the areas which are not very relevant for our purposes. After all, we are not (very) interested whether a recession probability is really 6% instead of 2%, but we care more about whether the probability crosses the 50% threshold.

### 3.5.1 Level specification

Let us start with the somewhat simpler case of the log-level specification working with  $\mathbf{z}_t^1$ , where we simply need the distribution of the  $h$ -step forecast of log industrial production, which we denote by  $\hat{f}_{y(r,p),T+h}^1$ , with mean value (point forecast) of  $\hat{y}_{T+h}(r, p)$ . Its forecast error variance is the corresponding element on the diagonal of  $\hat{\Psi}_h^1$ , let us call the square root of this element and thus the standard deviation of this output forecast error  $\hat{\psi}_h^1$ .

We limit the forecast horizon to some maximum  $hmax$  which is related to our underlying implicit recession duration: for our applications  $hmax = 5$ . Now we subtract from the distribution of a  $h$ -step forecast the (log) value of the corresponding variable five reference periods (months) earlier in order to calculate the cumulative growth rate over this five-month period. This benchmark reference period is  $T + h - p - (r - 1) - 5$ , where  $r$  and  $p$

are still given from the variable  $y(r, p)$  which is forecast. For this reference period the  $i$ -th publication happens in period  $T + h - (r - 1) - 5 + (i - 1)$ , and we always use the latest available release in real time published in  $T$ , which is therefore given by  $i^* = \min(5 + r - h, rmax)$ . Thus we get the following expression for the recentered distribution:

$$\hat{f}_{y(r,p),T+h}^1 - y_{T-(5+r-h-i^*)}(i^*, 2) \quad (6)$$

The complicated time and release indexing reflects the fact that we need to pick data that refers to an actual five-month lag with respect to the economic activity, not with respect to the date when the latest available release was published. For example, if we forecast the value of the second release ( $r = 2$ ) three publication months ahead ( $h = 3$ ), the reference period is  $T + h - p - (r - 1) = T$  (a nowcast), the reference period for subtraction is  $T - 5$ , and in principle the latest available release in  $T$  for this reference period would be  $i = 5 + r - h = 4$ . However,  $i^* = rmax < 4$  if the higher releases are not available in the dataset. Another example: forecasting the initial publication ( $r = 1$ ) five publication months ahead ( $h = 5$ ) implies that we must also use the initial publication ( $i^* = 1$ ) for subtraction which was published in period  $T$ , because the later releases would not have been published yet from a real-time perspective.

The intermediate result that we obtain is the density forecast of the predicted cumulative five-month growth rate for this  $h$ -step horizon, and the variance is still given by the already estimated forecast error variance  $\hat{\psi}_h^1$ , since the subtraction of a known value only re-centers the distribution.<sup>9</sup>

Finally, we determine how much probability mass falls below the selected threshold which defines a recession. Our threshold is based on the so-called “triangle approach” by Harding and Pagan (2002), where for simplicity we use a fixed 5-month period for evaluation,

<sup>9</sup>Since we treat the published data as fixed here, we are implicitly modelling the revision process as non-random. In principle it would be possible to generalize our approach. However, note that for a low forecast horizon  $h$  we subtract the final ( $rmax$ -th) release of the data which represents the end of the revision process (at least within our chosen information set), so no further random changes in those values would be possible anyway.



and our critical value is  $-0.025 * 2/5 = -1\%$  which is an arbitrary but commonly used value. That is, we want to determine the CDF value of the normal distribution with mean  $\Delta_5 \hat{y}_{t+h}(r, p) = \hat{y}_{T+h}(r, p) - y_{T-(5+r-h-i^*)}(i^*, 2)$  and standard deviation  $\hat{\psi}_h^1$  at the given value of  $-1\%$ . This of course is equivalent to finding the CDF value of the standard normal distribution denoted with  $\Phi()$ , at the corresponding standardized value:

$$\Phi \left( (-0.01 - \Delta_5 \hat{y}_{t+h}(r, p)) / \hat{\psi}_h^1 \right), \quad (7)$$

This standard numerical problem is the only non-linear step in our approach. This value gives the probability that the publication in  $T + h$  will meet the definition of a recession.

Apart from the publication lag, this method is also affected by a turning-point recognition lag. If for example we get in some publication period a turning-point signal based on the first publication, this means that the beginning of the recession (turning point) actually happened in some reference period up to  $-(h - p - 5)$  months ago (for  $h < 5 + p$ ), and so in reality the recession would likely be already underway. Although this is unfortunate, it is the logical consequence of the definition that a decline in economic activity must have a certain minimum duration to call it a recession.<sup>10</sup>

### 3.5.2 Specification in growth rates

If the system is specified in growth rates, an explicit cumulation of forecast errors from step 1 to step  $h$  is necessary, in parallel to the fact that we are interested in the cumulated (or average) growth over this period. Therefore in this case the distribution of the  $h$ -th step alone is not enough, and the determination of the distribution of the cumulative growth rate is complicated by the fact that the forecast errors are not independent. Formally, the vector of output growth forecasts  $(\Delta \hat{y}_{T+1}(r, p), \dots, \Delta \hat{y}_{T+h}(r, p))'$  from this model has

---

<sup>10</sup>In probit or Markov-switching models a formal minimum-duration requirement is typically missing for the forecasts. Instead at estimation time a specification is chosen that somehow delivers reasonable regime classifications in sample, where “reasonable” usually also means that the regime episodes should not be too short.

a covariance matrix which is not diagonal, and these correlations have to be taken into account when analyzing the sum of the forecasts.

Unfortunately, readily available formulae for this cross-horizon forecast covariance do not seem to exist. For this reason we use here a parametric Monte Carlo simulation (bootstrap) approach. We use the estimated distribution of the system residuals (their estimated covariance matrix  $\hat{\Omega}$ ) and repeatedly draw random realizations of this multivariate (normal) distribution to simulate possible paths of the endogenous variables, based on the estimated subset VAR of the current sample. Since we are holding the estimated coefficients fixed, we again abstract from the parameter uncertainty, as in the levels case.

As before we focus especially on output (log industrial production) up to forecasting horizon  $h$ . For each simulated path the growth rates are summed to approximate the cumulative growth. For horizons  $h < 5$  the known  $5 - h$  growth rates of the past are added in order to always consider a uniform 5-month time span.<sup>11</sup> With these simulated frequency distributions of the cumulative growth rate we can then work as before. Since in our application the center of the distribution is more relevant than the margins, we can work with a relatively low number of repetitions (500). Of course, this simulation must be re-run for each new observation that is added to the sample in the evaluation period. It therefore takes significantly longer if a pseudo out-of-sample evaluation is performed over a longer evaluation sample, but for a practical real-time application at the current sample margin the computational burden of such a one-off bootstrap is quite small.

## 4 In-sample properties

Due to the large dimension of the underlying system, it is hardly possible to characterize the properties of all the estimated equations in detail. We therefore limit ourselves here to the two equations of industrial production (first and second publications). In addition,

---

<sup>11</sup> Again, for the past values we abstract from further revisions here, see our previous footnote about this issue.




Figure 3: One-step forecast errors (in-sample residuals) for the log level specification

Table 1: Test of normality of forecast errors

	w.r.t. first publ.	w.r.t. second publ.
log level specification	2.60 (0.273)	0.70 (0.706)
growth rate specification	4.26 (0.119)	1.52 (0.467)

**Notes:** Jarque-Bera-type test with  $\chi^2(2)$  distribution, p-values in parentheses. The forecast errors are displayed in figures 3 and 4.

we present only the results for the shortest sample, i.e. with the final publication period 2007m1.

First, the figures 3 and 4 show the in-sample forecast errors (one step) of the equations for the industrial production of the two model variants. The variance of the forecast error for the second publication (first revision) is much lower, which is of course due to the fact that the existing information of the first publication can be used to predict the second publication for the same reference period, which eliminates a significant portion of the uncertainty.

For the method used, it is sometimes important whether the residuals (i.e., forecast errors) can be regarded as normally distributed. In table 1 we report the formal tests of this




Figure 4: One-step forecast errors (in-sample residuals) for the growth rate specification null hypothesis, with the encouraging result that this hypothesis cannot be rejected. The growth rate specification comes closest to rejection with respect to the first publication, obviously due to some outliers at the end of 1995, which can also be seen in figure 4.

## 5 Relevant predictors

It is of course of major interest which indicators are relevant for the forecast equations. The tables 2 and 3 give an overview of the non-excluded regressors in the equations for the industrial production, where we do not distinguish between different lags. Due to the prior automatic variable selection procedures all remaining terms are obviously highly significant by construction, so we abstain from reporting p-values. When interpreting this table it should be noted that the removal of the predictor in an equation is only directly relevant for the one-step forecast; for multi-step forecasts of a variable  $x$  the predictor  $z$  can nevertheless be relevant due to the cross-connections in the system, even if in the equation of  $x$  it is not included directly.

The bottom line in the log-level specification table 2 is that in principle all variables

contribute to the overall information. This includes the various revisions of the relevant variables, and so it appears that the revisions contain systematic information. Apart from this it is striking that the Ifo business climate index is only significant in the larger sample and for the first publication.

In the following table 3 with respect to the specification in growth rates it is apparent then, however, that some variables drop out completely from the equations for the industrial production. The inflation rate for example turns out as an obviously lagging variable that has no predictive content for production. The same applies to the oil price, the REX bond index and the yield spread for corporate bonds (corp. spread). But as a result it is interesting to mention that both the IFO as well as the ZEW index are classified as apparently irrelevant for the first publication of production.

## **6 Out-of-sample evaluation of the turning point forecasts**

We now turn to the forecasting performance of our approach within the available period of evaluation. As already mentioned, the time axis in the following diagrams reflects the real-time information set, i.e. the date of publication; for example at date 2007m1 all data published up to that date are used for estimation. For the case of a 2-step forecast, because of the publication lag of two months this implies a prediction of the first publication for the current reference month and thus an effective nowcast. It is crucial that the recession signals (or lack thereof) can be interpreted as real-time signals, the implicitly affected reference months must always be derived using the meta knowledge of the publication lags.

In figure 5 the predicted recession probabilities are presented based on the forecasts of the first publications; the upper graph refers to the specification in growth rates, and the lower graph is with respect to the levels specification. The shaded areas are used to distinguish between “trivial” and informative signals: in the shaded area the definition of a recession is already fulfilled by the published data (based on the latest available five-month period)

Table 2: Regressors in the industrial production equations, log-level specification

	$\tau = 2007m1$	$\tau = 2010m11$	$\tau = 2007m1$	$\tau = 2010m11$
	First publ.	First	Second	Second publ.
	(rel. 1)	publ.(rel. 1)	publ.(rel. 2)	(rel. 2)
Ind.prod. rel. 1	x	x	x	x
Ind.prod. rel. 2	x	x	x	x
CPI rel. 1		x	x	
CPI rel. 2	x	x	x	
AA rel. 1	x	x	x	x
AA rel. 2	x	x	x	
AI rel. 1	x	x		x
AI rel. 2		x	x	x
Spread 10Y		x	x	
Spread 1Y		x		x
Spread 2Y	x	x		x
Spread 3Y		x	x	x
Spread 5Y		x		x
Ifo index		x		
ZEW index	x	x		x
oil price	x			x
CDAX stock index	x	x		x
REX bond index			x	x
vacancies	x	x		x
Euribor int.rate	x		x	
Corp. spread	x		x	

**Notes:** AA = Foreign new orders, AI = Domestic new orders

Table 3: Regressors in the production equations, growth rate specifications

	$\tau = 2007m1$ First publ. (rel. 1)	$\tau = 2010m11$ First publ. (rel. 1)	$\tau = 2007m1$ Second publ.(rel. 2)	$\tau = 2010m11$ Second publ.(rel. 2)
Ind.prod. rel. 1	x	x	x	x
Ind.prod. rel. 2	x	x	x	x
CPI rel. 1				
CPI rel. 2				
AA rel. 1	x			
AA rel. 2	x			
AI rel. 1	x			
AI rel. 2		x	x	x
Spread 10Y			x	x
Spread 1Y		x	x	x
Spread 2Y	x	x		x
Spread 3Y		x	x	
Spread 5Y		x		
Ifo index				
ZEW index			x	
oil price				
CDAX stock index	x	x		
REX bond index				
vacancies	x	x	x	x
Euribor int.rate		x	x	
Corp. spread				

**Notes:** AA = Foreign new orders, AI = Domestic new orders

without requiring a forecast. (But keep in mind that the definition of the shaded area purely relates to the past, while the signals of the forecasts may relate to other periods depending on the forecasting horizon.) This “trivial” ex-post recession signal first applies in the (publication) period 2008m8. It is also apparent that at the end of 2008 there is an occasion where the cumulative 5-month growth lies above the defined recession threshold so that this shaded area includes a gap. The path of the published industrial production data around the end of 2008 thus resembles a double-dip recession pattern.

It is apparent that the predicted recession probabilities correctly signal the recession that (shortly) afterwards appears as a fact in the published data. Nevertheless, some forecasts also falsely indicate recessions that did not happen. Since this appears to be due to the volatility of various predictions it is also possible to consider the average probability over all forecast horizons. The general idea to combine forecasts from different models to improve the reliability is described in Timmermann (2006), the result for our application is shown in figure 6 and is quite encouraging. In particular, the level specification now no longer gets any recession wrong. At the same time the first recession signal happens in 2008m6, two months before the signal of the published data. And also the end of the recession appears from this series two months before the publication of relevant data.

The corresponding results for the forecasts regarding the second publication is shown in the figures 7 and 8. It is striking here that generally this seems to be less prone to false recession signals. Of course, it is easier to forecast the second publication (first revision) since the forecaster has one additional month’s information at hand. Thus it is crucial here whether this additional delay is reflected in the timing of the recession signal. Surprisingly, this is not necessarily the case; especially the (log-) levels specification gives a signal which occurs two months before the signal of the published data (with the exception of the 1-step horizon).

While this appears to be less urgent in the case of seconda publications, a simple averaging over the forecasting horizon is also possible, see figure 8. For the specification in levels the result is quite similar as before, whereas regarding the growth rate specification a






Figure 5: Out-of-sample recession probabilities (release 1) – i.e. according to first publications. Upper panel VAR specification in growth rates (for the  $I(1)$  variables) with a maximum of 3 lags, lower panel VAR specification in log-levels with at most 4 lags. — In the shaded areas the used recession definition is already met with the published data known in real time; a turning point forecasting signal would have to lie to the left of the shadow borders in order to provide an actual improvement.

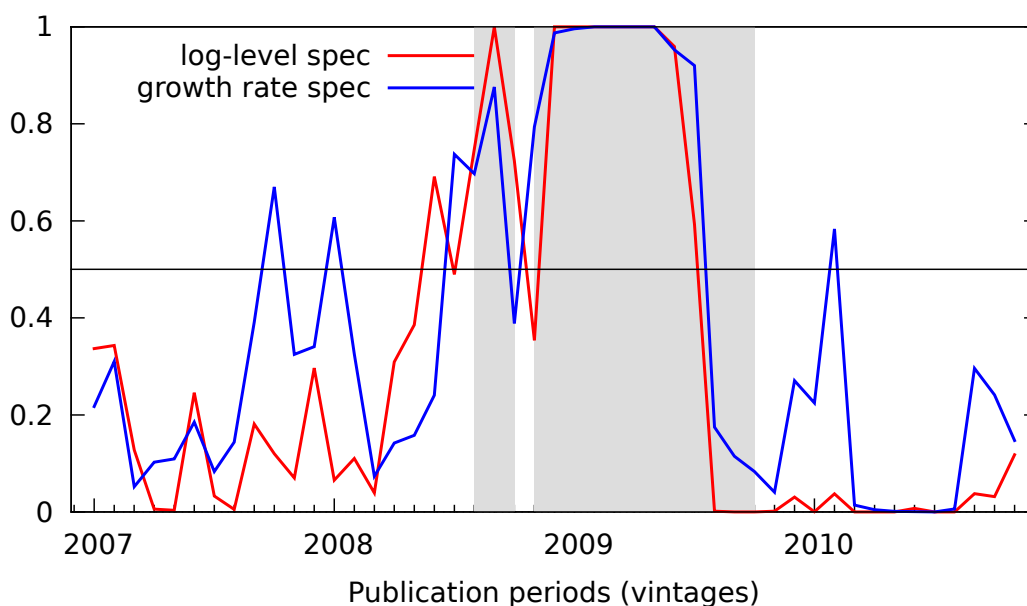
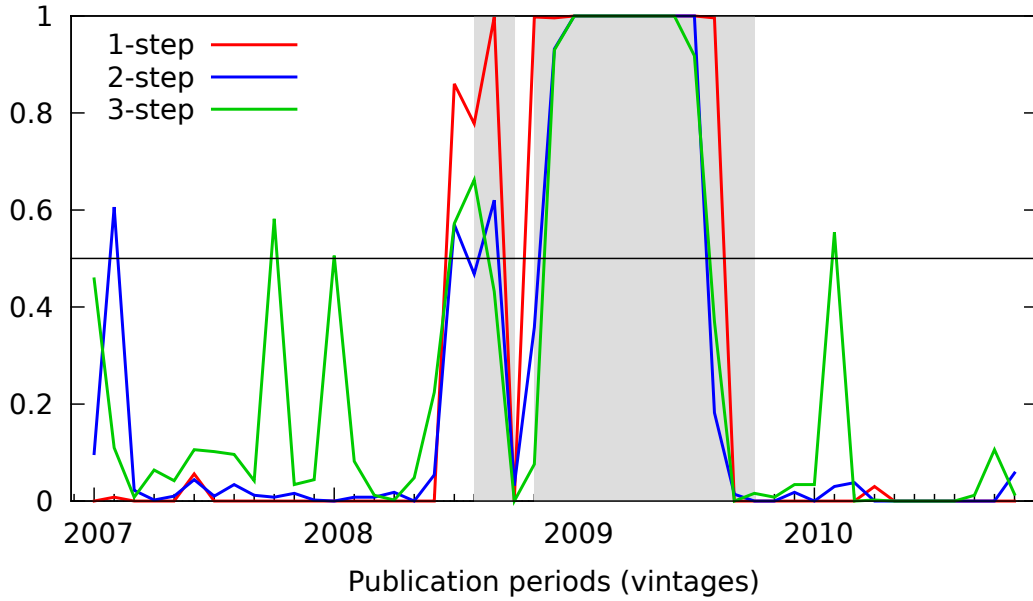


Figure 6: Average out-of-sample recession probabilities (release 1) – averages of the previously shown forecasting horizons, always referring to the first publication. Also see the notes to figure 5.

significant improvement is visible, as there are no longer any false recession signals.

Finally, for the assessment of the timing of the signals from the recession forecasts it is useful to refer to quantitative evaluation measures, which we report in table 4. For the interpretation of the evaluation statistics in table 4, however, the information asymmetry and the recession recognition lag must be taken into account: In comparison with the ex post determined benchmark series (see Schreiber, Theobald, Proaño, Stephan, Rietzler, and Detzer, 2012), which relates to reference months, we will shift the time axis of the predicted probabilities by the publication lag (backward) and by the forecast horizon (forward). As a result, we compare the recession signals always for the same reference month. However, it should be noted here that still different information sets for the forecasts sometimes remain. E.g. the 2-step forecast of the first publication based on publication period  $t$  refers to the reference period  $t$  of the same month, as does the 1-step forecast of the first publication from the forecast base period  $t + 1$ . But of course the forecast at time  $t + 1$  has an information advantage. This complication should be noted both when comparing different specifications as well as for comparisons with other models.

### Growth-rate specification



### Log-level specification




Figure 7: Out-of-sample recession probabilities (rel. 1) – i.e. referring to the second publication (first revision). Upper panel VAR specification in growth rates (for the  $I(1)$  variables) with at most 3 lags; lower panel specification in log-levels with at most 4 lags. Also see the notes to figure 5.

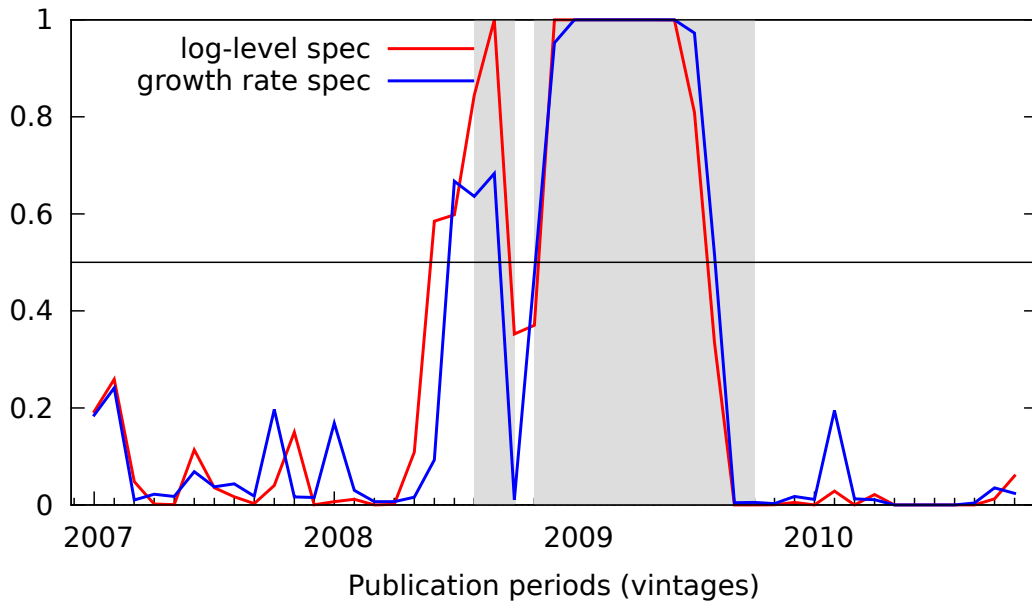


Figure 8: Average out-of-sample recession probabilities (release 2) – averages of the previously shown forecasting horizons, always referring to the second publication. Also see the notes to figure 5.

Therefore, the winning specification cannot be determined alone from the results in table 4.

Furthermore, in our approach a recession signal means that the implied reference period is about five months past the actual turning point, as we consider the cumulative growth rate over this period. In contrast, the benchmark series is generated entirely ex post.

Apart from these complications, the results from the evaluation point relatively clearly to the forecasts of the second publication as the best specification. Depending on the preferred measure either the 1-step forecast of growth rates or the average of horizons 1-3 in the level specification wins the competition.

## 7 Conclusions

In this paper we have proposed to forecast business-cycle turning points with linear systems that fully account for the publication lags and revisions of the data in real time. Our approach uses the whole forecast probability distribution (predictive density) to infer

Table 4: Out-of-sample forecast evaluation (comparison with ex-post benchmark)

specification	MAE	RMSE	Theil
<i>log-level specif.</i>			
1-step, release 1	0.1543	0.3498	0.3280
1-step, release 2	0.1203	0.3329	0.3104
2-step, release 1	0.2129	0.4035	0.3861
2-step, release 2	0.1563	0.3562	0.3319
3-step, release 1	0.2180	0.3947	0.3524
3-step, release 2	0.2064	0.3991	0.3788
average, release 1	0.1972	0.3475	0.3308
average, release 2	0.1569	<b>0.3256</b>	0.3156
<i>Growth-rate specif.</i>			
1-step, release 1	0.1851	0.3636	0.3489
1-step, release 2	<b>0.1187</b>	0.3318	<b>0.3103</b>
2-step, release 1	0.3134	0.4594	0.4186
2-step, release 2	0.2024	0.4012	0.3958
3-step, release 1	0.3586	0.4999	0.4270
3-step, release 2	0.3012	0.4794	0.4603
average, release 1	0.2991	0.4310	0.3970
average, release 2	0.2019	0.3844	0.3787

**Notes:** The comparison refers to the recession-regime benchmark series generated ex post with a non-parametric procedure, see Schreiber, Theobald, Proaño, Stephan, Rietzler, and Detzer (2012). All series are aligned to refer to the same reference periods (as opposed to publication periods). Release 2 is the second publication / first revision.

the probability of a recession. The main advantages over other popular methods are efficiency (if one is prepared to treat the data-generating process as approximately linear) and numerical stability. We included a relatively broad information set, which we view as essential for achieving a competitive forecasting model.

Using German monthly data and the sample 2007-2010 for forecast evaluation purposes, we show that the turning points can be predicted several months before official data publications confirm them. However, it should be pointed out again that longer forecast horizons provided unpleasant results, namely a clearly excessive volatility, such that the results for  $h = 4$  and  $h = 5$  were discarded and are not reported here. This problem is due to the features of the German data with relatively low growth rates even in expansion phases, and prior to the recent great recession comparably high growth rates in recession episodes, such that a distinction becomes difficult. A companion detailed comparison study (Schreiber, Theobald, Proaño, Stephan, Rietzler, and Detzer, 2012) also shows that our approach is generally competitive and often superior to other popular methods such as binary-regimes Probit models or Markov-switching setups; see appendix ?? below for the main results.

In the overall picture that takes into account the evaluation results with respect to the benchmark series as well as the timing of the recession signal, the log-level specification seems to be the best option for recession forecasting in real time. This holds especially when the three forecast horizons are combined into an average recession probability, and we recommend to use the second publication (first revision) of the output variable industrial production as the basis of the forecasts. With this specification we would have received the first signal for the recent great recession in Germany in the publication period 2008m6, which implies an effective turning point date in 2007m11 or shortly afterwards. Ex post and using only published data the turning point was recognized with the data publication in 2008m10. Therefore, despite the inherent recession recognition lag of our method a relatively early signal is apparently possible.

Regarding the question of the most useful indicators for the turning point detection, the

level and the growth-rate specifications yielded different results. For the level specification, in principle all considered variables appear to be relevant; while for the growth-rate specification some variables such as consumer prices, oil prices, REX bond index and corporate spread, but also the Ifo business climate index and ZEW expectations appear as relatively unimportant. In this sense the dynamics of the real revision-prone variables (such as lags and revisions to production and new orders), as well as timely available financial and labor market variables (spreads, CDAX stock index, reported vacancies) seem to be more important.

## **A Discussion of an alternative approach: forecasting equations for direct 1.. $h$ -step predictions of output (growth)**

In the main text of this paper we have described a VAR-type forecasting approach, where all included variables are also predicted, because these interim forecasts are then needed for the next iterative forecasting step. If the model specification correctly (or reasonably well) represents the data-generating process (DGP), this approach is indeed efficient.

Another possibility is to use equations that directly forecast output  $h$  steps ahead, by not including the first  $(h - 1)$  lags of regressors in the specification. Again we have to distinguish between an approach that uses the (log) level of the output variable, and one where the first difference (growth rate) is taken as the dependent variable. We sketch these approaches in turn; both have in common that the forecasts will not be efficient if the true model is a linear system (VAR).

### **A.1 Level specification**

The more straightforward case is given by the specification in levels. We would only need a single equation for (log) output  $y_t(r, p)$  where for an  $h$ -step forecast only regressors  $\mathbf{z}_{t-h}^1$  are used, i.e. the first  $h - 1$  lags are deliberately left out because they would not be

available when forecasting for period  $T + h$  in period  $T$ . For example, an equation for a direct two-step forecast can be written as follows:

$$y_{t+2}(r, p) = \sum_{k=0}^{K_2} \mathbf{a}'_{2,k} \mathbf{z}_{t-k}^1 + \mu_2 + u_{2,t}, \quad (8)$$

where the subscript “2” of various coefficients refers to the forecast horizon.

The advantage of this simple approach is that the distribution of the residuals directly corresponds to the relevant forecast error distribution. Again this distribution would have to be recentered (first by the point forecast and secondly by subtracting the known starting level in  $T + h - 5$ ). As before, insignificant regressors could be removed, although in this single-equation setup the need for this is less urgent.

The general disadvantage of this direct  $h$ -step forecast is that the omitted lagged variables will induce autocorrelation in the errors. One implication of this issue is that it may be useful to estimate the equation with GLS, e.g. by approximating the serial correlation with some simple AR- or MA- process and applying the correspondingly structured covariance matrix of the residuals.

A second problem is that in the presence of serial correlation in the residuals their conditional expectation does not vanish anymore; in the above two-step example with data available up to period  $T$ , this means  $E_T(y_{T+2}(r, p)) \neq \sum_{k=0}^{K_2} \mathbf{a}'_{2,k} \mathbf{z}_{T-k}^1 + \mu_2$  because the innovation  $u_{2,T}$  will be correlated with  $u_{2,T-2}$  and earlier lags. Since an estimate of the residual  $u_{2,T-2}$  is available, it would be possible to estimate also the residual  $u_{2,T}$  attached to the forecast target, by exploiting the (estimated) serial correlation.

## A.2 Growth rate specification

For the variant with growth rates  $\Delta y_t(r, p)$  we need an entire system of such forecasting equations for steps  $1..h$ , because we again need to cumulate the predicted output growth rates to infer the relevant 5-month total growth rate. For estimation this system would be



specified as follows:

$$\begin{aligned}
\Delta y_{t+1}(r, p) &= \sum_{k=0}^{K_1} \mathbf{a}'_{1,k} \mathbf{z}_{t-k}^0 + \mu_1 + u_{1,t} \\
\Delta y_{t+2}(r, p) &= \sum_{k=0}^{K_2} \mathbf{a}'_{2,k} \mathbf{z}_{t-k}^0 + \mu_2 + u_{2,t} \\
&\dots \\
\Delta y_{t+h}(r, p) &= \sum_{k=0}^{K_h} \mathbf{a}'_{h,k} \mathbf{z}_{t-k}^0 + \mu_h + u_{h,t}, \quad t = 1 + \max(K_1, \dots, K_h), \dots, T - h
\end{aligned} \tag{9}$$

The left-hand side of this system contains in fact different leads of output growth. This system can be estimated in various ways; however, if it is estimated equation-by-equation (e.g., OLS) care must be taken not to confuse the time indexing of the residuals, which is important for the calculation of their cross-equation covariance matrix for the forecast combination.

Of course, the equations of the system (except the first) again suffer from omitted variables because by construction some lags cannot be included in this approach. This implies that the respective residuals are likely to be autocorrelated. Again it may be fruitful to deal with this by using a GLS estimator, even when using an equation-by-equation approach.

And again it may be advisable to reduce the number of parameters in this model by automatically placing zero restrictions on the coefficient vectors.

After having estimated this system we want to use it for forecasting all the steps  $1..h$  based on the final available observation  $T$ . Here similar remarks as before apply, with respect to the impact of the serial correlation of the residuals on the way forecasts could be formed. Afterwards, the cross-equation sum of these forecasts  $(\sum_{i=1}^h \Delta \hat{y}_{T+i}(r, p))$  delivers the predicted (approximate) cumulated  $h$ -period growth rate. To infer the predictive density of this cumulation, i.e. the variance of the cumulated forecast error, we must again take the covariance between the forecast errors at different horizons into account. This simply means to estimate the cross-equation residual covariance matrix  $Cov(u_{1,t}, \dots, u_{2,t})$ , and use this estimate to compute the variance of the cross-equation sum of the different

forecast errors  $\sum_{i=1}^h u_{i,T}$  as the sum of all elements of that covariance matrix. That is the value one can use for the variance of the probability distribution of the cumulated growth rate.

Note that there are two different types of covariance matrices involved here: First each scalar residual  $u_i$  of each equation is a time series potentially with serial correlation. This serial correlation implies a covariance matrix  $E(\mathbf{u}_i \mathbf{u}_i') \neq \sigma_i^2 I$ , where  $\mathbf{u}_i$  denotes the vector of error terms in the  $i$ -th equation for all observation periods in the estimation sample. Secondly there is the cross-equation covariance matrix  $E(\mathbf{u}_t \mathbf{u}_t')$ , where  $\mathbf{u}_t = (u_{1,t}, \dots, u_{h,t})$ .

## References

- BANBURA, M., D. GIANNONE, AND L. REICHLIN (2010): "Large Bayesian vector auto regressions," *Journal of Applied Econometrics*, 25(1), 71–92.
- CLEMENTS, M., AND D. F. HENDRY (1999): *Forecasting non-stationary economic time series*. MIT Press.
- COTTRELL, A., AND R. LUCCHETTI (2010): *Gretl User's Guide Version 1.9*; available at <http://gretl.sourceforge.net/#man>.
- HAMILTON, J. D. (2011): "Calling recessions in real time," *International Journal of Forecasting*, 27, 1006–1026.
- HARDING, D., AND A. PAGAN (2002): "Dissecting the cycle: A methodological investigation," *Journal of Monetary Economics*, 49(2), 365–381.
- KROLZIG, H. M. (1997): *Markov switching vector autoregressions. Statistical inference and application to business cycle analysis*. Springer Science+Business Media.
- LAYTONA, A. P., AND M. KATSUURA (2001): "Comparison of regime switching, probit and logit models in dating and forecasting US business cycles," *International Journal of Forecasting*, 17(3), 403–417.
- LÜTKEPOHL, H. (2007): *New Introduction to Multiple Time Series Analysis*. Springer.
- NALEWAIK, J. J. (2012): "Estimating Probabilities of Recession in Real Time Using GDP and GDI," *Journal of Money, Credit and Banking*, 44(1), 235–253.
- NG, E. C. (2012): "Forecasting US recessions with various risk factors and dynamic probit models," *Journal of Macroeconomics*, 34(1), 11–125.

- ÖSTERHOLM, P. (2012): “The limited usefulness of macroeconomic Bayesian VARs when estimating the probability of a US recession,” *Journal of Macroeconomics*, 34, 76–86.
- SCHREIBER, S., T. THEOBALD, C. PROAÑO, S. STEPHAN, K. RIETZLER, AND D. DETZER (2012): “Verfahren der konjunkturellen Wendepunktbestimmung unter Berücksichtigung der Echtzeit-Problematik – Welche Indikatoren geben frühzeitig zuverlässige Signale?” IMK Study 27, IMK, available online at [www.imk-boeckler.de](http://www.imk-boeckler.de).
- SCHUMACHER, C., AND J. BREITUNG (2008): “Real-time forecasting of German GDP based on a large factor model with monthly and quarterly data,” *International Journal of Forecasting*, 24(3), 386–398.
- STOCK, J. H., AND M. W. WATSON (1993): “A procedure for predicting recessions with leading indicators: Econometric issues and recent experience,” in *Business Cycles, Indicators and Forecasting*. J. Stock and M. Watson, eds., NBER Book Series Studies in Business Cycles.
- TIMMERMANN, A. (2006): “Forecast combinations,” in *Handbook of Economic Forecasting*, Vol. 1, pp. 135–196. G. Elliot, C. Granger and A. Timmermann, eds., North-Holland, Amsterdam.