JENA ECONOMIC RESEARCH PAPERS

2019 – 006

Predicting Ordinary and Severe Recessions with a Three-State Markov-Switching Dynamic Factor Model

by

**Kai Carstensen
Markus Heinrich
Magnus Reif
Maik H. Wolters**

www.jenecon.de

ISSN 1864-7057

The JENA ECONOMIC RESEARCH PAPERS is a publication of the Friedrich Schiller University Jena, Germany. For editorial correspondence please contact markus.pasche@uni-jena.de.

Impressum:

Friedrich Schiller University Jena
Carl-Zeiss-Str. 3
D-07743 Jena
www.uni-jena.de

© by the author.

Predicting Ordinary and Severe Recessions with a Three-State Markov-Switching Dynamic Factor Model*

Kai Carstensen[†]

University of Kiel, Ifo Institute, CESifo

Markus Heinrich[‡]

University of Kiel

Magnus Reif[§]

University of Kiel, Ifo Institute

Maik H. Wolters[¶]

University of Jena, Kiel Institute for the World Economy, IMFS at Goethe University Frankfurt

13th September 2019

Abstract

We estimate a Markov-switching dynamic factor model with three states based on six leading business cycle indicators for Germany preselected from a broader set using the Elastic Net soft-thresholding rule. The three states represent expansions, normal recessions and severe recessions. We show that a two-state model is not sensitive enough to reliably detect relatively mild recessions when the Great Recession of 2008/2009 is included in the sample. Adding a third state helps to clearly distinguish normal and severe recessions, so that the model identifies reliably all business cycle turning points in our sample. In a real-time exercise the model detects recessions timely. Combining the estimated factor and the recession probabilities with a simple GDP forecasting model yields an accurate nowcast for the steepest decline in GDP in 2009Q1 and a correct prediction of the timing of the Great Recession and its recovery one quarter in advance.

Keywords: Markov-Switching Dynamic Factor Model, Great Recession, Turning Points, GDP Nowcasting, GDP Forecasting

JEL-Codes: C53, E32, E37

*We thank Rüdiger Bachmann, Christian Schumacher, and seminar participants at the University of Kiel, the University of Nuremberg, Deutsche Bundesbank, Ifo Munich, DIW Berlin, IWH Halle and RWI Essen for helpful comments.

[†]carstensen@stat-econ.uni-kiel.de

[‡]m.heinrich@stat-econ.uni-kiel.de

[§]Correspondence to: Magnus Reif, Ifo Institute, 81679 Munich, Germany. Email: reif@ifo.de

[¶]maik.wolters@uni-jena.de

1 Introduction

The failure of macroeconomists to predict the Great Recession of 2008 and 2009 has evoked much public criticism. While the debate mostly focuses on the state of macroeconomic modeling, it has also raised the question why professional forecasters even at the onset of the Great Recession did not foresee the steep output contraction that loomed around the corner. The case of Germany illustrates this failure. It was not until November 2008 that professional forecasters started predicting a recession despite clear warning signals accumulating throughout the year 2008.¹ For example, the expectation component of the Ifo business climate index—viewed by professional forecasters as one of the most important early indicators for German GDP—began its descent already in June 2007 and plunged heavily in July 2008, well before GDP plummeted in the fourth quarter of 2008 and the first quarter of 2009.

In this paper we take up the debate and ask whether it is possible to reliably predict in real time both business cycle turning points and GDP growth rates around these turning points, particularly during the Great Recession episode. We focus on Germany as a representative of the group of countries that show little persistence in GDP growth (other countries with this characteristic are, *inter alia*, Italy, Japan, Australia, and Norway). The lack of persistence is important because the usual approach to predict GDP growth by augmenting an autoregressive distributed lag model with a business cycle measure derived from coincident indicators (see, e.g., Chauvet and Potter, 2013) works well only for countries like the US that exhibit significant sample autocorrelations.² As a more promising approach for low-persistence countries we suggest to directly exploit the information of leading indicators for future GDP. For Germany we show that this yields very competitive one-quarter ahead forecasts of business cycle turning points and GDP growth.

To extract information from leading indicators of the German business cycle, we use the Markov-switching dynamic factor model (MS-DFM) proposed by Diebold and Rudebusch (1996) and Kim and Yoo (1995) because it has been shown to be a valuable device for assessing the state of an economy (Chauvet, 1998; Kim and Nelson, 1998; Camacho et al., 2014) and its results are much more timely available than those of simple benchmark approaches such as the Bry-Boschan algorithm. However, unlike the previous literature we specify the MS-DFM with three states. Specifically, we add to the conventional expansion and (ordinary) recession states a third state which reflects a severe recession.³ This is motivated both by the general perception that the Great Recession was different from previous post-war recessions and may thus require a special

¹See Drechsel and Scheufele (2012) for an analysis of the performance of leading indicators during the financial crisis and Heilemann and Schnorr-Bäcker (2017) for a detailed documentation of the chronological sequence of data releases and publications of professional forecasts in 2008.

²The cross-country difference in the persistence of GDP growth and its implications for forecasting are hardly discussed in the literature. One exception is Stock and Watson (2005) who document that, among the G7 countries, Germany, Italy and Japan have negligible persistence in the post-1984 period.

³Three-state Markov-switching models have been applied mainly to the US (Boldin, 1996; Layton and Smith, 2000; Krolzig and Toro, 2001; Ferrara, 2003; Nalewaik, 2011; Ho and Yetman, 2012) but also to the euro area (McAdam, 2007; Artis et al., 2004; Anas et al., 2008). However, they have been implemented in univariate and vectorautoregressive contexts but not in a dynamic factor model. In addition, these papers intend to identify a recession, a normal growth regime, and a high growth regime, the latter being typically interpreted as a recovery in the line of Sichel (1994) and Morley and Piger (2012). The only exception is Hamilton (2005) who identifies a severe recession regime in a univariate model of the US.

econometric treatment, and by our empirical finding documented below that an MS-DFM with two states becomes instable in 2008.⁴

We also address the question of how to determine the number of states in real time. This is highly relevant as the severe recession state is only weakly identified before the Great Recession which is probably why studies analyzing pre-2008 data report that the German business cycle can well be represented with two states (Bandholz and Funke, 2003; Artis et al., 2004; Kholodilin and Siliverstovs, 2006). We propose to choose—at each point in real time—the number of states that optimizes the quadratic probability score which measures how well the MS-DFM fits the Bry-Boschan algorithm. Thereby, we effectively train the MS-DFM to yield results close to a simple benchmark but at the same time exploit its advantage to detect turning points instantaneously at the sample end.

Another methodological contribution to the literature is to prepend a flexible indicator selection procedure to the MS-DFM. This is important because there are many potentially useful business cycle indicators available for an economy to be fed into the MS-DFM, while the nonlinear one-step estimation approach by Kim and Yoo (1995), which simultaneously determines the factor and the state probabilities, is subject to numerical problems if the number of parameters is large.⁵ We use a soft thresholding procedure that accounts for multivariate correlations among the variables to extract a small number of variables from a medium-sized set of pre-selected indicators because Bai and Ng (2008) show that hard thresholding, i.e., using statistical tests to ensure that a predictor is significant irrespective of other predictors, might be inadequate in such situations. Specifically, we use the elastic net (EN) algorithm of Zou and Hastie (2005), which is a convex combination of a ridge regression and a Least Absolute Shrinkage and Selection Operator (LASSO). It is suited particularly for data sets with highly correlated variables like business cycle indicators.

We structure our empirical analysis in two parts. We first study whether the MS-DFM reasonably describes the German business cycle ex post using revised data for the period January 1991 to June 2016. Subsequently, we examine how well the MS-DFM is suited to timely detect, and predict, business cycle turning points in real time. In both parts, we compare the properties of models with two and three states, emphasizing the Great Recession period.

In the ex post analysis presented in Section 4, we apply the EN algorithm to select three out of 16 hard indicators such as new orders and three out of 19 survey indicators, all of which have been considered as early indicators in the literature on German business cycle dynamics.

⁴Another way to approach this problem is to stay with a two-regime model but make the regime-dependent growth rates follow a random walk as in Eo and Kim (2016). However, their setting differs in several important respects from ours. First, they analyze US GDP in a univariate approach. Extending it to a factor model is computationally very demanding. Second, they model the full postwar sample which is characterized by a secular decline in US growth rates while we model only the most recent 25 years of German data for which a similar decline is much less obvious. Third, their focus is on extracting in-sample features of the business cycle while we are mainly concerned with real-time forecasting for which too much parameter flexibility typically reduces forecast accuracy.

⁵While this problem can be circumvented by a two-step approach which first extracts a linear factor from the data set and subsequently uses this factor to estimate a univariate Markov-switching model, Camacho et al. (2015) argue that the one-step method is—although it involves a higher computational burden—more robust against misspecification. Furthermore, Doz and Petronevich (2016) compare the performance of both methods on dating French business cycle turning points and find that one-step estimation is more precise in indicating the beginning and end of recessions.

Using six indicators has been proven to capture business cycle dynamics quite well for several countries (see, e.g., Chauvet, 2001; Camacho and Martinez-Martin, 2015; Aastveit et al., 2016).⁶ We then feed these six indicators in one-factor MS-DFMs with two and three states, estimate the parameters, and smooth out the factors, which can be interpreted as composite leading indicators, and the conditional state probabilities. It turns out that the three-state model is superior in several dimensions. Its factor correlates more strongly with GDP growth (if aggregated to the quarterly frequency) and its states can be interpreted nicely as expansion, ordinary recession, and severe recession, while the two-state model seems to identify a low-growth regime and a medium-severe recession regime that is too fierce for any pre-2008 downturn and too mild for the Great Recession. The three-state model also dates recessions in general, and the Great Recession in particular, much more in line with conventional wisdom and the Bry-Boschan algorithm.⁷ In contrast, the two-state model is less sensitive and thus typically comes a bit late because the business cycle needs to deteriorate considerably before it is classified as medium-severe recession.

In Section 5 we present the second part of our empirical analysis. We ask whether the superiority of the three-state model carries over to a forecasting situation in real time in which the data exhibit a ragged-edge structure and the Bry-Boschan algorithm is not suited because its standard version requires a lag of at least 5 months until it is able to signal a turning point. To this end, we set up a recursive nowcasting exercise from January 2001 to June 2016 that in each month selects six indicators by means of the EN algorithm and estimates one-factor MS-DFMs with two and three states. We find that the two-state model signals turning points fairly well but becomes unstable during the Great Recession, while the three-state model appears poorly identified before the Great Recession but works properly thereafter. These results suggest that a forecaster would have dismissed the two-state model after the Great Recession and moved towards the three-state model. To operationalize this, we use real-time model selection based on the quadratic probability score and the BIC which yields a combined two-state/three-state model. It produces precise and timely nowcasts of business cycle turning points.

Using a recursive out-of-sample forecasting exercise we even demonstrate that the combined model is able to provide excellent 3-month ahead turning point predictions that would have been extremely useful for policy makers during the Great Recession. In particular, it predicts an upcoming recession with almost 100 percent probability already in July 2008 and thus four months ahead of most professional forecasters. Moreover, in March 2009 it correctly predicts that the recession comes to an end soon, one month before the German public started to discuss a third stimulus package.

⁶The results are similar when we only select four indicators or when we increase the number of selected indicators up to ten. Using more than ten indicators leads to noisy and more and more unreliable recession signals.

⁷We apply the Bry-Boschan algorithm because there is no widely accepted monthly business cycle chronology for the German economy available against which we can assess the results of our MS-DFM. The chronology published by the Economic Cycle Research Institute (ECRI) is based on both an unknown data set and an unknown method and is provided with a lag of approximately one year. The business cycle dates published by the OECD are determined by applying the Bry-Boschan algorithm on the OECD's composite leading indicator on a quarterly basis. A useful proposal is made by Schirwitz (2009) who suggests a consensus business cycle chronology based on the results of different methods. However, it is again on the quarterly frequency. Hence, we use the Bry-Boschan algorithm applied to monthly industrial production as a benchmark.

Finally, we assess whether point forecasts of German GDP growth rates benefit from including the information provided by the MS-DFMs. Specifically, we augment an autoregressive forecasting model with the dynamic factor and the recession probabilities extracted from the early indicators. Specifically, augmenting an autoregressive forecasting model with the dynamic factor and the recession probabilities extracted from the early indicators considerably improves nowcasts and short-term forecasts, especially during recessions. In particular, it yields an accurate nowcast for the steepest decline in GDP in 2009Q1.

This paper adds to the literature that applies Markov-switching models to the German business cycle. Ivanova et al. (2000) estimate univariate Markov-switching models for various interest rate spreads and examine their predictive power for business cycle turning points. Bandholz and Funke (2003) use an MS-DFM model with a bivariate data set to construct a leading indicator for the German business cycle. Kholodilin (2005) augment that model with a second factor and interpret it as a coincidence indicator. Abberger and Nierhaus (2010) demonstrate the predictive power of the Ifo business climate index with regard to business cycle turning points in a univariate framework. Proaño and Theobald (2014) use Probit models rather than a Markov-switching approach to predict German recessions. None of these contributions considers a flexible data selection approach based on a large data set or a distinction between severe and ordinary recessions. Moreover, they are based on revised data, while we analyse the predictive ability of the model in a real-time setting.

The remainder of the paper is structured as follows: Section 2 outlines our baseline MS-DFM model and the estimation method. In Section 3 we describe our data set and the variable selection procedure. Section 4 and 5 present our estimation results as described above. Finally, Section 6 concludes.

2 The Markov-switching dynamic factor model

We use a Markov-switching dynamic factor model (MS-DFM) to extract common nonlinear business cycle dynamics from a set of leading indicators. We distinguish between n_h hard indicators, $y_{it}^{(h)}$, such as new orders, interest rates, and oil prices, which typically account for rather short-term fluctuations, and n_s survey indicators, $y_{it}^{(s)}$, such as the Ifo business climate index and the ISM purchasing managers index which capture primarily medium-term business cycle dynamics. The distinction is important because quarterly growth rates of hard indicators generally correlate well with quarter-on-quarter GDP growth (and monthly growth rates of hard indicators correlate with monthly business cycle indicators like industrial production), while business surveys typically rather fit year-on-year GDP growth. We model these differences along the lines of Camacho et al. (2014): For the hard indicators we assume a standard factor structure,

$$y_{i,t-l_{h,i}}^{(h)} = \gamma_i^{(h)} f_t + z_{it}^{(h)}, \quad i = 1, \dots, n_h, \quad (1)$$

where $y_{i,t-l_{h,i}}^{(h)}$ is a hard indicator in monthly growth rates, $z_{it}^{(h)}$ is an idiosyncratic component, f_t is a scalar dynamic factor that leads the month-on-month business cycle dynamics by three months, and $l_{h,i}$ is the lag with which the hard indicator i enters the model. For the survey

indicators we assume a slightly different specification,

$$y_{i,t-l_{s,i}}^{(s)} = \gamma_i^{(s)} \sum_{k=0}^{11} f_{t-k} + z_{it}^{(s)}, \quad i = 1, \dots, n_s, \quad (2)$$

where $y_{i,t-l_{s,i}}^{(s)}$ is a soft indicator in levels, $z_{it}^{(s)}$ is an idiosyncratic component and $l_{s,i}$ is the lag with which the soft indicator i enters the model. We include the sum of lags 0 to 11 of the factor as a parsimonious way to incorporate the phase shift associated with a year-on-year growth cycle that correlates with the survey indicators.⁸

For all indicators, we take into account that they lead the cycle to different extents and thus should enter the factor model with different lags $l_{h,i}$ and $l_{s,i}$. To make the factor lead the business cycle by 3 months, we include indicators that lead GDP by 1, 2, and 3 quarters with a lag of 0, 3, and 6 months, respectively (in Section 3 below we describe in detail how we choose the indicators and their lags).

Following Doz and Petronevich (2016), we model the vector of idiosyncratic components, $z_t = [z_{1t}^{(h)}, \dots, z_{n_h t}^{(h)}, z_{1t}^{(s)}, \dots, z_{n_s t}^{(s)}]'$, as a diagonal VAR process of lag order q ,

$$z_t = \psi_1 z_{t-1} + \dots + \psi_q z_{t-q} + \varepsilon_t, \quad \varepsilon_t \sim i.i.d. N(0, \Sigma_z), \quad (3)$$

where $\psi_1, \dots, \psi_q$ and Σ_z are diagonal matrices, and ε_t is a vector of independent Gaussian shocks. We specify the common factor as an autoregressive process of lag order p with regime-dependent intercept,

$$f_t = \beta_{S_t} + \phi_1 f_{t-1} + \dots + \phi_p f_{t-p} + \eta_t \quad \eta_t \sim i.i.d. N(0, 1), \quad (4)$$

where η_t is an independent Gaussian shock. The intercept, β_{S_t} , depends on the state variable $S_t \in \{1, \dots, m\}$ as follows:

$$\beta_{S_t} = \beta_1 S_{1,t} + \beta_2 S_{2,t} + \dots + \beta_m S_{m,t},$$

where $S_{m,t}$ is equal to unity if S_t equals m and zero otherwise. We assume that S_t follows a first-order ergodic Markov chain. The corresponding $m \times m$ transition matrix, Π , has elements p_{ij} defining the probability to switch from regime i to regime j , with $\sum_{j=1}^m p_{ij} = 1$ for every $i = 1, \dots, m$. We do not impose restrictions on the duration of any regime. We consider models with two regimes ($m = 2$) that represent expansions and recessions and with three regimes ($m = 3$) with the aim to distinguish in addition between ordinary and severe recessions.

Defining the vector $y_t = [y_{1,t-l_{h,1}}^{(h)}, \dots, y_{n_h,t-l_{h,n_h}}^{(h)}, y_{1,t-l_{s,1}}^{(s)}, \dots, y_{n_s,t-l_{s,n_s}}^{(s)}]'$ of dimension

⁸As a robustness check we apply an Almon lag structure as a more flexible weighting scheme. Specifically, we model the survey indicators as $y_{i,t-l_{s,i}}^{(s)} = \gamma_i^{(s)} g(\delta, L) f_t + z_{it}^{(s)}$ where $g(\delta, L) = \sum_{k=0}^{11} c(\delta, k) L^k$ is a lag polynomial, L denotes the lag operator, and $\delta = [\delta_0, \delta_1]$. We specify $c(\delta, k)$ as an exponential Almon lag $c(\delta, k) = \frac{\exp(\delta_0 k + \delta_1 k^2)}{\sum_{k=0}^{11} \exp(\delta_0 k + \delta_1 k^2)}$. Since results do not improve, we keep the parsimonious specification. Results are available upon request.

$n = n_h + n_s$, we cast the model into state-space form,

$$y_t = Ba_t \quad (5)$$

$$a_t = \mu_{S_t} + Fa_{t-1} + R\omega_t, \quad (6)$$

where a_t is the state vector, ω_t is a vector of independent Gaussian shocks with mean zero and covariance matrix Q , B , F and R are coefficient matrices, and μ_{S_t} is a state-dependent intercept. For details, we refer to Appendix A.

We estimate the MS-DFM by numerically maximizing the highly nonlinear likelihood function.⁹ To this end, we employ the filter proposed by Kim (1994), see Appendix B for details. It yields the latent dynamic regime dependent factor as well as the Markov-switching probabilities.

We use the following starting values. In a first step, we approximate f_t by a static principal components analysis and plug it into (4) with invariant intercept to estimate starting values for ϕ_1 to ϕ_p by OLS. We also plug f_t into (1) and (2), and run OLS regressions to obtain starting values for $\gamma^{(h)}$ and $\gamma^{(s)}$. The residuals of these regressions approximate the idiosyncratic components z_t . We use them to estimate a diagonal VAR model of lag order q to find starting values for ψ_1 to ψ_q , and Σ_z . In the next step, we take all these values to initialize and estimate a dynamic factor model with a single regime. This yields starting values for $\gamma^{(h)}$, $\gamma^{(s)}$, $\phi_1, \dots, \phi_p$, $\psi_1, \dots, \psi_q$ and Σ_z . Finally, combining the results of the single-regime model with starting values for the transition matrix and the regime dependent means completes the set of required parameters. Specifically, we initialize the transition matrix by assuming persistent regimes (high values on the main diagonal and small values on the off-diagonal). We construct starting values for the regime dependent means as follows. In case of the two-state model we take the average over all positive factor values and the average of all negative factor values for the expansion and recession regime, respectively. For the three state model we use the same approach and take in addition the smallest factor value in the sample as starting value for the mean of the severe recession regime.

3 Indicator selection

While there are many business cycle indicators available for the German economy, the challenge is to reduce their number such that they carry all necessary cyclical information without overburdening the nonlinear maximum likelihood technique described above with estimating too many parameters. Boivin and Ng (2006) demonstrate that even linear factor models do not always benefit from adding more and more variables in particular in the context of forecasting. Camacho et al. (2015) focus specifically on MS-DFMs and show that, once a small number of high quality indicators is included, adding more indicators yields only minor improvements in terms of the identification of business cycle turning points. Finally, Schumacher (2010) shows that feeding only a set of targeted predictors into an otherwise standard factor model can improve prediction accuracy of German GDP. Hence, we first pre-select a medium-sized set of potentially useful indicators based on previous results in the literature and then apply to it a

⁹We use the Matlab `globalsearch` class based on the routine `fmincon` to obtain a global maximum.

variable selection algorithm that chooses only a few final indicators to be fed into the MS-DFM.

Our pre-selection is primarily based on previous results of the literature (Fritsche and Stephan, 2002; Kholodilin and Siliverstovs, 2006; Drechsel and Scheufele, 2012; Lehmann and Wohlrabe, 2016) on the German business cycle. As hard indicators we choose 6 industrial order inflow series, 2 commodity prices, 3 interest rates, the German contribution to the EMU M2, and the DAX index which have all been found to give early business cycle signals. To take into account Germany’s dependence on foreign markets, we also include US industrial production as a simple indicator for world market fluctuations. We finally add German consumer prices and employment as important economic state variables even though they are typically thought to lag the business cycle. We leave it to the indicator selection algorithm below to decide whether they are promising candidates. As survey indicators we pre-select 9 series published by the European commission and 7 series published by the Ifo institute. These series cover a broad range of economic activity, with a specific focus on expectations. We add the purchasing manager index for the US, the Belgium business confidence indicator—which is sometimes found to lead the EU cycle, see Vanhaelen et al. (2000)—and the Euro-coin index to reflect the importance of major foreign markets.¹⁰ Altogether, we pre-select a set of 35 monthly business cycle indicators, of which 16 are categorized as hard and 19 as survey indicators. To ensure stationarity, we apply log differencing to all hard indicators—except for interest rates and spreads where we compute differences without taking logs—while the survey indicators are stationary by construction. The indicators are then standardized to mean zero and variance one. A complete description is provided in Appendix E. Our sample starts in January 1991 in order to avoid any issues associated with the German reunification break, and runs until June 2016.

Based on the pre-selected data set, we employ an automatic indicator selection algorithm. As our goal is to provide early signals for business cycle turning points, the algorithm should select only those hard indicators that exhibit a strong lead correlation with quarter-on-quarter GDP growth rates, $\Delta \log(GDP_t)$, and only those survey indicators that exhibit a strong lead correlation with year-on-year GDP growth rates, $\Delta_4 \log(GDP_t)$. To this end, we transform our monthly indicators to quarterly frequency by averaging over the respective quarter and estimate the predictive regressions

$$\Delta \log(GDP_t) = \sum_{i=1}^{16} \sum_{l=1}^3 b_{i,l}^{(h)} y_{i,t-l}^{(h)} + u_t^{(h)} \quad (7)$$

for hard indicators and

$$\Delta_4 \log(GDP_t) = \sum_{i=1}^{19} \sum_{l=1}^3 b_{i,l}^{(s)} y_{i,t-l}^{(s)} + u_t^{(s)} \quad (8)$$

for survey indicators, where l denotes quarterly lags. The parameters $b^{(h)}$ and $b^{(s)}$ are estimated using the elastic net (EN) proposed by Zou and Hastie (2005) and successfully used by Bai and Ng (2008) for indicator selection.¹¹ The elastic net is a convex combination of a

¹⁰Although it consists of both hard and survey indicators, the Euro-coin index is assigned to the survey category because it exhibits, as the other survey indicators, the highest correlation with year-on-year GDP growth (Altissimo et al., 2010).

¹¹Another method to identify the relevant indicators in the context of predicting recessions are boosted regres-

ridge regression and a LASSO and yields nonzero parameter estimates only for a few important indicators. It solves the following optimization problem:

$$L = (\lambda_1, \lambda_2, \mathbf{b}) = |\mathbf{y} - \mathbf{X}\mathbf{b}|^2 + \lambda_1 |\mathbf{b}|_1 + \lambda_2 |\mathbf{b}|^2, \quad (9)$$

where $\mathbf{b} = (\mathbf{b}_1, \dots, \mathbf{b}_N)'$ is a $N \times 1$ dimensional coefficient vector, and

$$|\mathbf{b}|_1 = \sum_j |\mathbf{b}_j| \quad \text{and} \quad |\mathbf{b}|^2 = \sum_j \mathbf{b}_j^2.$$

$\mathbf{y} = (y_1, \dots, y_T)'$ denotes a centered response variable—in our setting either $\Delta \log(GDP_t)$ or $\Delta_4 \log(GDP_t)$ —and $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_N)$ is a set of N standardized predictors $\mathbf{x}_i = (\mathbf{x}_{1i}, \dots, \mathbf{x}_{Ti})'$ —in our setting either the hard indicators $y_{i,t-l}^{(h)}$, $i = 1, \dots, 16$, $l = 1, \dots, 3$, or the survey indicators $y_{i,t-l}^{(s)}$, $i = 1, \dots, 19$, $l = 1, \dots, 3$.

The tuning parameters λ_1 and λ_2 control the weight on the L_1 and L_2 -norm penalty, respectively. For increasing relative weight λ_1 the EN approaches the LASSO which is known to shrink coefficients to zero due to the non-smoothness of its objective function, while for increasing relative weight λ_2 the EN approaches the ridge regression which is capable of handling highly correlated predictors. Zou and Hastie (2005) show that the EN inherits both properties and is thus particularly suited for our purpose. They also demonstrate that the EN can be transformed into a LASSO problem which can be estimated by the Least Angle Regression (LARS) of Efron et al. (2004). This algorithm, called LARS-EN, is a forward stepwise additive fitting procedure. The number of steps, k , equal the number of included variables and corresponds, for given λ_2 , to a specific value of λ_1 . Hence, instead of choosing λ_1 and λ_2 , one may equivalently choose λ_2 and k which is what we do in the following.¹² For a detailed description of the LARS-EN algorithm along with the estimated coefficients, we refer to Appendix C.

We apply the LARS-EN algorithm to both (7) and (8) and choose in both cases $\lambda_2 = 100$ which is a fairly large value and allows high correlation between the selected indicators.¹³ We select $n_h = 3$ hard indicators and $n_s = 3$ survey indicators in order to avoid predominance of one category and so to balance their relative merits: Hard indicators are often thought to give more reliable signals ex post but suffer from publication lags and strong revisions in real time, while soft indicators are timely available and remain largely unrevised but might be more loosely connected to the “hard” outcome variables such as GDP we are interested in.¹⁴ In both cases,

sion trees (see Ng, 2014; Döpke et al., 2017), which complement the probit approach and thus are not applicable in our case.

¹²For given λ_2 this works as follows. Since LASSO shrinks coefficients to zero, start with a sufficiently large λ_1 (which yields zero estimates of all coefficients) and iteratively lower λ_1 until the prespecified number, k , of nonzero coefficient estimates is obtained.

¹³Higher values for λ_2 do not change the selection. Smaller values for λ_2 cause LARS-EN to select only one of a set of correlated indicators which is problematic in our setting because we rather select similar indicators with high correlation and good forecasting power for GDP than very different indicators of which some are only loosely related to GDP. We also tried to choose λ_2 according to cross validation based on the MSE but this method leads to inferior results which is why we do not report them here.

¹⁴A series of robustness checks showed that our specification is in fact optimal to produce reliable real-time recession signals. It clearly dominates the alternative specifications $n_h = 1$ and $n_s = 5$, $n_h = 2$ and $n_s = 4$, and $n_h = 5$ and $n_s = 1$ and is slightly better than the specification $n_h = 4$ and $n_s = 2$. The results are available upon request.

we thus set the elastic net parameter k to 3.¹⁵

4 Ex post business cycle dating for Germany

In the following, we apply our dynamic factor model combined with the LARS-EN indicator selection to identify the German business cycle turning points in the full sample. Such an ex post business cycle dating based on revised data is of its own interest as it complements simple but purely univariate dating algorithms like Bry-Boschan and undisclosed multivariate procedures like the one published by the Economic Cycle Research Institute (ECRI). Our main interest is, however, to show that our empirical approach produces reasonable results in-sample before we subsequently use it to predict turning points out-of-sample in a real-time forecasting setting.

4.1 Selected indicators

We first apply the LARS-EN algorithm with the aforementioned settings to the pre-selected set of indicators. We obtain the following results. The selected hard indicators comprise—in the order of selection—foreign orders of capital goods, domestic orders of intermediate goods, and domestic orders of capital goods. The selected survey indicators include—again in the order of selection—overall production expectations, Ifo business expectations, and Ifo export expectations. All six indicators are selected with a lag of one quarter implying that they lead the business cycle by three months. To obtain a factor with the same lead property, we include the indicators contemporaneously in the monthly factor model, i.e., set $l_{h,i} = l_{s,i} = 0$ in equations (1) and (2) for all i . Altogether the selection reflects common knowledge that orders of production inputs and business expectations are valuable early indicators. It also highlights the openness of the German economy as foreign trade plays a role in both indicator sets.

4.2 Factor estimate for MS(2)-DFM

Based on the selected indicators, we first estimate a “classical” two-state model, MS(2)-DFM, that distinguishes between expansions and recessions. Before estimation, we have to determine the lag orders of the factor and the idiosyncratic components. Camacho and Perez-Quiros (2007) and Aastveit et al. (2016) argue that the main dynamics of a business cycle can be captured solely by a switching intercept, and Boldin (1996) shows for univariate Markov-switching models that overparameterization can lead to severe problems. Therefore, we set the lag order, p , of the factor to zero.¹⁶ This allows us to treat our intercept as a switching mean. The autocorrelation functions of the idiosyncratic components indicate a lag order of $q = 2$.

The estimated means, probabilities and factor loadings of the MS(2)-DFM are reported in Table 1, while the autoregressive parameters of the idiosyncratic components are presented in

¹⁵In some instances, the LARS-EN algorithm selects two different lags of the same indicator. In such a case, we include in our factor model the lag selected first and increase k by one to select another indicator to be included.

¹⁶We also estimated models with $p = 1$ and $p = 2$ but obtained inferior results for the in-sample fit, thereby confirming the results of Camacho and Perez-Quiros (2007) and Aastveit et al. (2016).

Table 8 in Appendix D. State 1 features a positive mean of $\beta_1 = 0.32$, a high persistence probability, and occurs 87 percent of the time unconditionally. It can thus be interpreted as an expansionary regime. State 2 exhibits a negative mean of $\beta_2 = -2.12$, is less persistent, and takes place 13 percent of the time which is why it appears like a standard recession regime. However, the estimated means have a strong implication. To see this, recall that the factor is constructed from standardized indicators and has a sample mean of approximately zero. Therefore, the expansionary (recessionary) mean describes the average positive (negative) deviation from “normal times”. While the scale is arbitrary, the relative sizes are not. Hence, the estimates imply that a recession is, in absolute terms, about 6.5 times stronger than an expansion. This appears very large and is a consequence of effectively treating the Great Recession as a normal recession.

Table 1: Estimated parameters of the MS(2)-DFM

Parameter	β_1	β_2	p_{11}	p_{22}	P_1	P_2	$\gamma_1^{(h)}$	$\gamma_2^{(h)}$	$\gamma_3^{(h)}$	$\gamma_1^{(s)}$	$\gamma_2^{(s)}$	$\gamma_3^{(s)}$
Estimate	0.32 (0.10)	-2.12 (0.34)	0.97 (0.02)	0.79 (0.13)	0.87	0.13	0.23 (0.03)	0.38 (0.05)	0.20 (0.03)	0.12 (0.01)	0.11 (0.01)	0.11 (0.01)

Notes: Estimated standard errors are reported in parentheses below the estimates. $P_1 = P(S_t = 1)$ and $P_2 = P(S_t = 2)$ are the unconditional probabilities of being in the expansionary and recessionary states, respectively. $\gamma_1^{(h)}$ to $\gamma_3^{(h)}$ and $\gamma_1^{(s)}$ to $\gamma_3^{(s)}$ refer to the factor loadings of the hard indicators (new foreign orders of capital goods, new domestic orders of intermediate goods, and new domestic orders of capital goods) and soft indicators (overall production expectations, overall business expectations, and export expectations), respectively.

Nevertheless, the factor corresponds closely to GDP growth, see Figure 1 where we display quarterly averages of the filtered factor along with quarterly German GDP growth rates. Even though the factor solely summarizes the fluctuations of the six leading indicators identified above, it tracks GDP growth remarkably well. In several instances it appears to lead GDP growth as intended by construction. In fact, it exhibits the strongest correlation of 0.64 to GDP growth with a lead of one quarter which suggests that already the MS(2)-DFM may be well suited to forecast business cycle turning points.

Finally, note that the estimated factor loadings are all positive and significantly different from zero, implying procyclicality of the selected indicators. As in previous studies (Camacho et al., 2014; Camacho and Perez-Quiros, 2010), the soft indicators load more weakly on the factor than the hard indicators.

4.3 Factor estimate for MS(3)-DFM

Now we introduce a third state. The idea is to account for, and predict, extraordinary strong output contractions like the Great Recession. The majority of the literature only considers two regimes. The few exceptions that consider three regimes rather aim at identifying weak growth phases (sometimes called stall phases) in addition to recessions and expansions (Boldin, 1996; Ferrara, 2003; Artis et al., 2004; Nalewaik, 2011). Instead, we aim at identifying regime 1 as expansionary, regime 2 as ordinary recession and regime 3 as severe recession as in Hamilton

Figure 1: Filtered factor of the MS(2)-DFM and GDP growth

Notes: The figure displays quarterly averages of the filtered factor estimated from an MS(2)-DFM, and quarterly German GDP growth rates, 1991Q1-2016Q2. The factor is re-scaled to fit mean and variance of GDP.

(2005).¹⁷

To identify the three regimes and obtain numerically stable results of the numerical estimation procedure, we impose two economically sensible restrictions on the 3×3 transition matrix. Specifically, as in Hamilton (2005) we do not allow to directly switch from regime 1 (expansion) to regime 3 (severe recession) or vice versa. This is motivated by the observation that the Great Recession started off like an ordinary recession at the beginning of 2008, became severe after the Lehman collapse (industrial production dropped by more than 3 percent in each of the four months between November 2008 and February 2009), and phased out in the subsequent months.¹⁸ The restricted transition matrix reads as follows:

$$\Pi_r = \begin{bmatrix} p_{11} & (1 - p_{11}) & 0 \\ p_{21} & p_{22} & (1 - p_{21} - p_{22}) \\ 0 & (1 - p_{33}) & p_{33} \end{bmatrix}. \quad (10)$$

Except for adding a third state, we apply the same specification choices as before. In particular, we include the same six indicators as in the two-state model, set the lag order, p , of the factor to zero and the lag order, q , of the idiosyncratic components to two. The estimated parameters of the MS(3)-DFM are reported in Table 2. They compare favorably to the results of the two-state approach because the relative size of the means is more in line with what one would expect. The first regime has a positive mean implying that an expansion is characterized by a positive deviation from average times. The second regime has a negative mean of

¹⁷Proaño (2017) also identifies three business cycle states for Germany. He distinguishes above trend growth, trend growth and recessions, rather than the expansionary, recessionary and severe recessionary regime that we are interested in.

¹⁸A likelihood ratio test of the two restrictions was not rejected with a p -value of almost 1. In a model without the two zero restrictions the point estimates of the two transition probabilities are virtually zero with large standard errors which suggests that they are not well identified by the data.

an absolute size that is 2.5 times the mean of the first regime. Hence, a normal recession is characterized by a negative deviation from average times, and it is 2.5 times as strong as an expansion. The third regime has a much lower mean and can thus safely be interpreted as a severe recession. The estimate implies that a severe recession is more than five times worse than a normal recession. Not much surprisingly given the development of the Great Recession, a severe recession is estimated to be much less persistent than normal recessions and expansions. In addition, the probability to switch from the ordinary recession to the severe recession is much lower ($1 - \hat{p}_{21} - \hat{p}_{22} = 0.01$) than to switch back ($1 - \hat{p}_{33} = 0.34$), and the unconditional probability of being in a severe recession is much lower than that of being in an ordinary recession. The factor loadings are significantly positive and also very similar in magnitude compared to the ones of the MS(2)-DFM.

Table 2: Estimated parameters of the MS(3)-DFM

Parameter	β_1	β_2	β_3	p_{11}	p_{22}	p_{33}	p_{21}	P_1	P_2	P_3
Estimate	0.61 (0.12)	-1.42 (0.29)	-7.93 (1.08)	0.94 (0.03)	0.83 (0.10)	0.66 (0.51)	0.16 (0.09)	0.73	0.26	0.01
Parameter	$\gamma_1^{(h)}$	$\gamma_2^{(h)}$	$\gamma_3^{(h)}$	$\gamma_1^{(s)}$	$\gamma_2^{(s)}$	$\gamma_3^{(s)}$				
Estimate	0.20 (0.02)	0.24 (0.04)	0.18 (0.02)	0.09 (0.01)	0.09 (0.01)	0.09 (0.01)				

Notes: Estimated standard errors are reported in parentheses below the estimates. $P_1 = P(S_t = 1)$, $P_2 = P(S_t = 2)$, and $P_3 = P(S_t = 3)$ are the unconditional probabilities of being in the states of expansion, recession, and severe recession, respectively. $\gamma_1^{(h)}$ to $\gamma_3^{(h)}$ and $\gamma_1^{(s)}$ to $\gamma_3^{(s)}$ refer to the factor loadings of the hard indicators (new foreign orders of capital goods, new domestic orders of intermediate goods, and new domestic orders of capital goods) and soft indicators (overall production expectations, overall business expectations, and export expectations), respectively.

The factor of the MS(3)-DFM corresponds closely to GDP growth, see Figure 2. Again it appears to lead GDP growth. It exhibits the strongest correlation of 0.68 to GDP growth with a lead of one quarter. This correlation is slightly larger than the one of the two-state factor which indicates that the three-state model might be better suited to predict German business cycle turning points.

4.4 Which model gives a more realistic characterization of the German business cycle?

In the following we present the smoothed recession probabilities of the two-state and three-state models and assess whether they give a realistic picture of the German business cycle phases. As a benchmark we would ideally use a generally accepted monthly business cycle chronology for Germany comparably to the one of the NBER for the US. Since this is not available, we construct our own benchmark. To this end, we apply the Bry-Boschan business cycle dating algorithm because it is an often-used method and easily replicable. Given a monthly benchmark series, x_t , the algorithm defines peaks by

$$\wedge_t = \{(x_{t-d}, \dots, x_{t-1}) < x_t > (x_{t+1}, \dots, x_{t+d})\},$$

Figure 2: Filtered factor of the MS(3)-DFM and GDP growth

Notes: The figure displays quarterly averages of the filtered factor estimated from an MS(3)-DFM, and quarterly German GDP growth rates, 1991Q1-2016Q2. The factor is re-scaled to fit mean and variance of GDP.

and troughs by

$$\forall t = \{(x_{t-d}, \dots, x_{t-1}) > x_t < (x_{t+1}, \dots, x_{t+d})\},$$

where d is the minimum duration which also implies that peak and trough must be at least d periods apart. The definition reveals the major drawback of the algorithm. To identify a turning point it requires at least d subsequent observations. Throughout the literature it has become standard to assume $d = 5$ months (and additionally a minimum length of a full cycle of 15 months). We follow this convention. Thus the algorithm exhibits a lag of at least 5 months until it signals that the state of the business cycle has changed, while the MS-DFM is—if it is applied in a real-time situation—designed to identify turning points instantaneously.

As benchmark series, x_t , to be fed into the Bry-Boschan algorithm we choose industrial production excluding construction.¹⁹ This is motivated by the stylized fact that German industrial and overall activity are so strongly correlated that the industry sector, which exhibits a much more pronounced cyclical behavior than GDP, is generally thought of as the driver of the German business cycle. As industrial production is available at a monthly frequency, it enables us to determine the state of the economy on a monthly basis.

Figure 3 presents the smoothed recession probabilities of the MS(2)-DFM (panel a) and the MS(3)-DFM (panel b). Generally, they match the Bry-Boschan classification (indicated by shaded areas) quite well. In particular, they start rising slightly before, or at the beginning of, all benchmark recessions. Further, the MS(3)-DFM model identifies the steepest contraction of GDP during the Great Recession as a severe recession regime, while the probability of a severe recession is close to zero for the rest of the sample.

There are some important differences between the recession probabilities of the two models.

¹⁹We exclude construction because particularly in the 1990s after German reunification the construction cycle was decoupled from the overall business cycle in Germany.

The two-state model detects the Bry-Boschan recessions starting in January 1995 and September 2002 with probabilities of less than 0.4, while the three-state model identifies them with probabilities of more than 0.9.²⁰ This finding indicates that the three-state model is much more sensitive than its two-state counterpart because the distinction between ordinary and severe recessions allows it to assign already mildly weak times as (ordinary) recessions. The increased sensitivity can also be inferred from panel (c) of Figure 3 which displays the recession probabilities of the two-state model and the joint probabilities of a normal or severe recession of the three-state model. Clearly, the latter are always higher than the former.

As a potential drawback, an increased sensitivity may go hand in hand with a higher risk of false alarms. In fact, the three-state model indicates the existence of a recession in a few cases when both the two-state model and the Bry-Boschan algorithm do not. It is instructive to examine these additional signals in 1998, 2005 and at the end of 2009 in more detail. In September 1998 the recession probability of the three-state model exceeds 0.5 for 7 months in a row while the two-state probability remains slightly below 0.5 and the benchmark does not indicate a recession at all. At that time the German business cycle was temporarily fragile as indicated by a majority of the selected indicators. After a peak in July 1998, industrial production exhibited a weak period of more than 6 months before it picked up again. However, the trough was already in November 1998 which is not more than five months away from the peak. Hence, the Bry-Boschan algorithm neglects this episode.

In 2005 the recession probability of the three-state model rises to a value of just below 0.4. While it thus gives only a weak signal, it does so for good reasons. In mid-2004 the selected soft indicators started a gradual decline that continued until April 2005, and the selected hard indicators (domestic and foreign orders of capital goods and domestic orders of intermediate goods) exhibited two weak months in February and March 2005. As a consequence, the recession probability increases. The model result coincides with the assessment of professional forecasters at that time. For example, according to the Ifo business cycle forecast of June 2005 the German economy “started stuttering” (Flaig et al., 2005). Today we know that the German economy in 2005 was rather stagnating. Industrial production decreased in February, May, August and November 2005 but not in two or more months in a row. Hence, there is no local minimum to be identified by the Bry-Boschan algorithm, which therefore neglects this episode.

From November 2009 to January 2010 the recession probability of the three-state model exhibits a brief hike with a maximum of nearly 0.7. It reflects a decline in domestic orders of capital goods during September 2009 to January 2010 and drops in foreign orders of capital goods in November 2009 and January 2010, indicating a weakening business cycle. In addition, the soft indicators increase only very moderately in these months. In fact, the recovery of the German economy from the Great Recession paused during the winter 2009/10. After industrial production had increased by 11% in the first five months after its trough in April 2009, it stagnated until February 2010, but again there was no clear minimum which is why the Bry-Boschan algorithm does not indicate a recession.

²⁰A similar episode is the Bry-Boschan recession of February to June 2016. We do not take it too seriously, however, because the data are still relatively preliminary which may induce divergences in the information content of the early indicators and industrial production that may vanish after future data revisions.

Figure 3: Recession probabilities of MS(2)-DFM and MS(3)-DFM

Notes: Panel (a) displays smoothed recession probabilities of the MS(2)-DFM. Panel (b) displays smoothed probabilities of an ordinary recession (red line) and severe recession (blue line) of the MS(3)-DFM. Panel (c) compares the probability of a recession from the MS(2)-DFM (red line) with the joint probability of an ordinary or severe recession from the MS(3)-DFM (black line). Shaded areas correspond to the recessions dated by the Bry-Boschan algorithm.

These examples demonstrate that it is ultimately a matter of definition whether an episode should be classified as a recession and that it is important to combine information from hard and soft indicators. It also shows that what might appear as oversensitivity at first sight, may carry useful information that is more nuanced than a 0-1 rule.

To illustrate the leading properties of the two models, Figure 4 takes a closer look at the Great Recession. In panel (a) the red line represents the smoothed recession probabilities of the two-state model. Since the factor is designed to lead GDP by one quarter, a recession probability measured in month t refers to month $t + 3$. Specifically, the recession probability first exceeds 0.5 in June 2008 and thus predicts that a recession starts in September 2008, the month of the Lehman collapse. While this appears like a sensible result, it is by now conventional wisdom that the Great Recession in Germany started earlier that year²¹ while the most severe production declines came a few months later. The root of the problem is again the missing distinction between ordinary and severe recessions. As the two-state model identifies a single “average” recession, it comes late when a recession is mild.

In contrast, the three-state model almost perfectly matches the Great Recession. Panel (b) of Figure 4 displays the smoothed probabilities of an ordinary recession (red line) and a severe recession (blue line). The probability of an ordinary recession first rises above 0.5 in January 2008 indicating a recession start three months later in April which compares well with the development of output: the second quarter of 2008 saw the first (small) decline in GDP. The probability of a severe recession exceeds 0.5 during October to December 2008 implying that January to March 2009 are the core recession months. In fact, GDP loss in the first quarter of 2009 was by a large margin the steepest of the Great Recession. Also, industrial production fell maximally in January 2009. Altogether, the three-state model indicates that the Great Recession occurred between April 2008 and May 2009 while the Bry-Boschan algorithm identifies May 2008 to April 2009.

To more formally evaluate the two-state and three-state models against the Bry-Boschan benchmark, we employ the quadratic probability score

$$QPS = \frac{1}{T} \sum_{t=1}^T [B_{t+k} - P_t(\text{recession})]^2, \quad (11)$$

where B_{t+k} denotes the binary Bry-Boschan benchmark series with lead equal to $k \in \{0, 1, 2, 3\}$ and $P_t(\text{recession})$ is the smoothed probability to be in a recession (two-state model) or in an ordinary or severe recession (three-state model). QPS takes an optimal value of zero if the smoothed probabilities calculated by a model coincide with the benchmark.

²¹Using a simple rule-of-thumb that defines a recession as at least two consecutive quarters of negative real GDP growth, one would date the start of the recession in the second quarter of 2008. Official business cycle dates from the CEPR Euro Area Business Cycle Dating Committee are only available for the euro area as a whole. According to those the business cycle peak occurred in the first quarter of 2008. Business cycle dates for Germany are released by the Economic Cycle Research Institute (ECRI) which dates the peak of the previous expansion in April 2008.

Figure 4: Recession probabilities of MS(2)-DFM and MS(3)-DFM during the Great Recession

Notes: Panel (a) displays smoothed recession probabilities of the MS(2)-DFM during the Great Recession. Panel (b) displays smoothed recession probabilities of the MS(3)-DFM during the Great Recession. The solid lines depict the model-based recession probabilities which lead the business cycle by three months. Shaded areas correspond to the recessions dated by the Bry-Boschan algorithm.

In addition, we compute the false positives measure

$$FPS = \frac{1}{T} \sum_{t=1}^T [B_{t+k} - I\{P_t(\text{recession}) > 0.5\}]^2, \quad (12)$$

where $I\{P_t(\text{recession}) > 0.5\}$ is an indicator function taking the value of 1 if the smoothed probability of being in a recession is higher than 0.5 and 0 otherwise. Hence, this measure counts the number of false signals, i.e. incorrectly predicted periods, of the model. The lower

the FPS is, the better is the model's ability to reliably predict recessions.

Table 3 reports the QPS and FPS measures for the two-state and three-state models. According to both quality measures the three-state approach provides a superior in-sample fit for all measures. This suggests that using an MS(3)-DFM gives a more realistic characterization of the German business cycle than using a more classical two-state model. Since it provides detailed information in terms of regime probabilities we also prefer it over a simple 0-1 classification scheme like the Bry-Boschan algorithm that in addition can only classify downturns that last at least five months as recessions.

Table 3: QPS and FPS measures

k	QPS				FPS			
	0	1	2	3	0	1	2	3
MS(2)-DFM	0.1830	0.1702	0.1661	0.1725	0.2164	0.2131	0.2164	0.2262
MS(3)-DFM	0.1491	0.1240	0.1089	0.1121	0.2164	0.1803	0.1574	0.1541

Notes: QPS is the quadratic probability measure defined in (11). FPS is the false positives measure defined in (12). $k \in \{0, 1, 2, 3\}$ refers to the lead of of the Markov-switching models compared to the Bry-Boschan benchmark.

Additionally, the QPS and FPS measures corroborate that the Markov-switching models exhibit a lead compared to the Bry-Boschan benchmark. Specifically, the QPS measure is minimal at $k = 2$ suggesting that both models have a lead of two months, while the FPS measure is lowest at $k = 1$ month for the two-state model and $k = 3$ for the three-state model. Taken together, these results indicate that it is possible to achieve a leading property of almost one quarter by carefully selecting a set of leading indicators and integrating them into a Markov-switching dynamic factor model.

4.5 Monthly business cycle chronology for Germany

In some situations it may be valuable to have a dichotomous monthly business cycle chronology (even though recession probabilities are much more informative). Characterizing months with a recession probability greater than 0.5 as recessionary and assuming a lead of three months, we derive such a chronology from our preferred three-state model, see Table 4. We also report the chronologies based on the two-state model and the Bry-Boschan algorithm.

According to our three-state model Germany has experienced eight recessionary phases since January 1991. Particularly pronounced episodes are the post-reunification recession (May 1992 to July 1993), the “dot com” recession (March 2001 to January 2002), the Great Recession (April 2008 to May 2009), and the European sovereign debt crisis which consists of two phases (September 2011 to February 2012 and September to December 2012) summarized in columns 8a and 8b of Table 4.

The two-state model identifies solely those four pronounced recessions. However, the timing is always a little late and the recession lengths appear a bit underestimated. For example, according to the two-state model the Great Recession lasted only nine months and the European debt crisis as little as four months. In contrast, the Bry-Boschan benchmark indicates eight recessionary phases which in most cases coincide well with the three-state model. Exceptions are the two

Table 4: Benchmark recession dates for Germany

MS(3)-DFM	start	–	05.92	04.95	08.98	03.01	10.02	04.08	09.11	09.12	–
	end	–	07.93	09.95	02.99	01.02	04.03	05.09	02.12	12.12	–
MS(2)-DFM	start	–	07.92	–	–	06.01	–	09.08	–	09.12	–
	end	–	02.93	–	–	01.02	–	05.09	–	12.12	–
Bry-Boschan	start	01.91	03.92	01.95	–	03.01	09.02	05.08	–	08.11	08.15
	end	09.91	07.93	10.95	–	11.01	09.03	04.09	–	01.13	12.15

Notes: Recessions are defined as $PR[S_t = 2|\Psi_T] \geq 0.5$ (MS(2)-DFM) and $PR[S_t = 2|\Psi_T] + PR[S_t = 3|\Psi_T] \geq 0.5$ (MS(3)-DFM), where Ψ_T is the information set available at the sample end. Episodes that last less than 4 months are excluded.

episodes at the sample beginning and the sample end which may be the result of a sample edge problem (in particular, potential data revisions render the 2015 recession tentative), and the episode between August 1998 and February 1999 already discussed in the previous subsection.

5 Real-time business cycle assessment and forecasting

In this section, we apply the Markov-switching dynamic factor models to nowcast and forecast business cycle turning points, as well as GDP growth rates, in real time. In doing so, we exploit the advantage of these models to indicate turning points instantaneously and thereby circumvent the endpoint problem inherent to the Bry-Boschan algorithm which leads to delayed signals.

5.1 Nowcasting German business cycle turning points

To assess the nowcasting ability of the two-state and three-state models, we perform a nowcasting experiment over the evaluation period January 2001 until June 2016 using real-time data. We choose this evaluation period because it includes five recessions which allows us to judge the results with some confidence, while the initialization sample of ten years (1991M01-2000M12) is still sufficient to estimate an MS-DFM. In addition, we include equally long periods before and after the Lehman bankruptcy which helps us to understand whether adding a third state—which is hardly identifiable before the Great Recession—would have made a difference in real time.

We construct a real-time data set consisting of the same pre-selected set of 35 indicators as in the previous section. To this end, we take the series of new orders, employed persons, and inflation from the real-time database of Deutsche Bundesbank,²² and US industrial production from the real-time database of the OECD. The remaining hard indicators are determined on financial markets and are not revised.²³ The survey indicators are revised only very marginally, hence we neglect these revisions.

²²Some releases miss some observations at the beginning of the sample. In such cases, we use growth rates from previous releases to fill the gaps by means of backward chaining.

²³The only exception is the German contribution to EMU M2. However, it is so rarely and slightly revised that we can safely take it as being unrevised.

In each step of the nowcasting experiment, we go through the selection and estimation stages described in previous sections. To obtain a nowcast for month $\tau \in \{2001M01, \dots, 2016M06\}$, we first apply the LARS-EN algorithm to the sample available at the end of this month and select a set of three hard and three survey indicators. Subsequently, we feed these indicators into the Markov-switching dynamic factor models with zero lags for the factor and two lags for idiosyncratic component, estimate the parameters, and smooth out the state probabilities. As a result, we not only obtain a series of real-time probabilities but also a time-varying selection of indicators for the period January 2001 until June 2016. Figures 12 and 13 in the Appendix depict the recursive selection of the hard and survey indicators, respectively.

It turns out that the real-time indicator selection is stable in the sense that changes in the chosen indicator sets occur infrequently. The selection reflects the traditional dependence of the German business cycle on global developments. Of the six indicators, the LARS-EN algorithm always picks two hard indicators (foreign orders of capital goods and, with very few exceptions, one of the two commodity prices) and one survey indicator (the Euro-coin indicator until April 2013 and the Ifo export expectations thereafter) that summarize external information while only two survey indicators (the Ifo business expectations and another Ifo indicator) are more closely related to the domestic situation. Interestingly, the Great Recession does not seem to affect the selection with one exception which may signal an increased relevance of the domestic economy: the sixth indicator is foreign orders of intermediate goods until February 2009 but domestic orders of intermediate goods thereafter.

The real-time nowcasts of the recession probabilities are constructed using all available information at a certain point of time. Since we only select indicators that lead the business cycle by at least 3 months, it would be sufficient to include indicators of period $\tau - 3$ and earlier in order to compute filtered probabilities of period τ . However, such an approach would neglect important information as, at the end of period τ , the realizations of, say, new orders for period $\tau - 2$ and survey indicators for period τ are already known. Therefore, we compute the real-time probabilities by means of backward smoothing taking all observations into account that are known in period τ .²⁴ Like Chauvet and Hamilton (2006) and Hamilton (2011), we find that these smoothed probabilities are much more stable and reliable than their filtered counterparts.

The upper panel of Figure 5 depicts the smoothed recession probability generated in real time by the MS(2)-DFM. It shows a roughly similar evolution as the one based on full sample estimates discussed in Section 4 but deviates from it in two episodes, 2004 and 2005, when it falsely signals recessions. A difference between real-time and ex post analysis can be caused by two factors. First, the ex post model is applied to revised data which is relevant in many cases because revisions of some hard indicators can be huge. Second, the real-time model suffers from the usual sample-end problem while the ex post model knows how the indicators evolve over the whole sample. This affects not only the smoothed probabilities but also the variable selection

²⁴Note that this leads to ragged edges in the data structure. We deal with that complication by using the method of Mariano and Murasawa (2003) which is extended to the nonlinear Markov-switching framework by Camacho et al. (2018). It consists of replacing the missing observations at the end of the sample by random numbers distributed independently of the model's parameters. These random numbers are in turn eliminated by an appropriately defined Kalman filter. As shown by Camacho et al. (2018), neither the maximum of the likelihood function nor the estimated filtered probabilities depend on these random numbers.

Figure 5: Real-time nowcasts of recession probability

(a) MS(2)-DFM

(b) MS(3)-DFM

Note: Smoothed recession probabilities of (a) MS(2)-DFM and (b) MS(3)-DFM recursively estimated with real-time data. Ψ_t denotes the information set as of period t . Shaded areas correspond to the recessions of the benchmark business cycle chronology from Section 4.4.

algorithm. For example, it may take a while until the real-time model replaces an indicator with deteriorating information content by another one that is better suited.

During the first episode, the real-time recession probability rose to slightly below 0.5 in June 2004, mainly because the selected three soft indicators—Ifo business expectations, intermediate goods production expectations, and the EuroCoin index—started to ease off at the beginning of 2004. The ex post recession probability does not react because based on the full sample, the latter two indicators are replaced by the overall production expectations and the Ifo export

expectations which evolved more positively. In particular, export expectations tended to increase in the first three quarters of 2004. Nevertheless, industrial production stagnated—in March 2004 it was on the same level as in November 2003—but without a clear local minimum which is why the (ex post) Bry-Boschan algorithm does not indicate a recession.

In the second episode of June to August 2005, the real-time recession probability exceeded 0.8. Again, this was primarily due to a temporary decline in the selected soft indicators at the beginning of 2005. In addition, foreign orders of capital goods and foreign orders of intermediate goods exhibited a few weak months. While there are differences to the ex post analysis both due to indicator selection and data revisions, it is remarkable that the three-state model based on revised full-sample data gives a (weak) recession signal at the same time, see the discussion in Section 4.4 above. As argued there, the German economy stagnated at that time but did not slip into a recession. Hence, the Bry-Boschan algorithm does not react but there is good reason for an increased recession probability.

The real-time recession probabilities estimated by the MS(3)-DFM are shown in the lower panel of Figure 5. They differ substantially from those based on the full sample. In particular, before the Great Recession the third state is not well identified in real time and the probabilities of being in the second or third state exhibit erratic fluctuations. Immediately after both orders and the early indicators have plummeted in the end of 2008, the third state starts to identify a severe recession. Hence, the advantage of having a third state kicks in at this point of time.

After the Great Recession, the real-time MS(3)-DFM raises two false alarms which do not show up in the ex post analysis. In August 2013 the real-time recession probabilities increased to slightly below 0.5, caused by a temporary weakness of both the selected hard indicators—foreign orders of capital goods, domestic orders of intermediate goods, and the HWWA index—and the selected survey indicators. In particular, the Ifo business expectations declines from February until May. In the ex post analysis the recession probabilities do not exceed 0.2 because of data revisions and differences in the selection of the hard indicators. Most notably, a real-time stagnation of foreign orders of capital goods in May is revised into a strong increase by roughly 2.4%. Moreover, the HWWA index, which is selected in real time, is ex post replaced by domestic orders of capital goods, which evolve less negatively. Industrial production in turn shows a very erratic behavior with alternating months of positive and negative growth between June and November 2013. Therefore, the Bry-Boschan algorithm cannot detect a local minimum and does not signal a recession in autumn 2013.

In November and December 2014 there was another false alarm with real-time recession probabilities exceeding 0.5, primarily due to a marked decline in Ifo business expectations and overall production expectations. In addition, domestic orders of intermediate goods and the HWWA commodity price index, which reflects the demand situation on world markets, exhibited weak or even negative growth rates over most of the year. The ex post analysis does not signal a recession mainly for two reasons. First, the indicator selection differs. In particular, instead of the HWWA commodity price index the ex post model selects domestic orders of capital goods which evolve less negatively. Second, the downswing of the domestic order inflow is much more pronounced in real time than using revised data. For instance, in June 2015 domestic orders of capital goods drop by 3.5% in real-time, while the revised decline is only 2.7%. The real-

time results also differ from the ex post results of the Bry-Boschan algorithm because industrial production sharply decreased only in August 2014 and in January 2015 while it also saw a few positive months such that a clear local minimum is missing.

To further understand what happens inside the two models, Figure 6 takes a closer look at their recursively estimated state-specific means. The two-state model (panel a) exhibits a break at the beginning of the Great Recession. Before, the model is remarkably stable with a first state that has a positive mean and a second state that has a negative mean of similar absolute magnitude. Since the factor is extracted from standardized indicators and thus has a sample mean of approximately zero, the first state can be interpreted as expansion, while the second state represents a recession. During the Great Recession, however, the expansion mean is estimated as approximately zero whereas the recession mean falls dramatically. At that time, a user of this model would have found the model's result unconvincing, both because of its instability and—perhaps more importantly—because of its interpretation: neither an upswing with a growth rate that merely equals the sample average nor an extreme contraction could have been easily reconciled with what was observed as expansions and recessions before the onset of the Great Recession. These findings probably would have been interpreted as a signal that “this time is different” and that a third state is necessary to characterize the German business cycle properly.

In contrast, the three-state model is instable before the Great Recession because the third state is only weakly identified during this time. Until 2005 the first two states would have been interpreted as expansion and ordinary recession, while the third state having a mean considerably smaller than the second state would have been labeled a severe recession. However, during the boom of 2006 to mid 2008 which preceded the Great Recession, the first state signals a strong boom and the third state a recession of similar absolute magnitude whereas the second state indicates “average times” with mean zero and thus average growth—an interpretation difficult to reconcile with prior experience. This changes again with the beginning of the Great Recession. As more and more bad news come in, the model starts to extract a severe recession regime with a very negative mean that fluctuates—after a few months of undershooting—in a range that is considerably below the pre-crisis level. In addition, the means of the first and second state stabilize at levels that lend to the interpretation of expansion and mild recession, respectively. This stabilization is also visible in Figure 5 where the smoothed real-time recession probabilities largely coincide with those based on the full sample shown in Figure 3.

5.2 Model selection in real time

The results of the nowcasting experiment directly raise the issue of model selection in real time. We suggest to use either of the following two criteria to compare the two-state and three-state models. The first criterion is the QPS which measures how closely the Markov-switching models match the business cycle turning points identified by the Bry-Boschan algorithm applied to German industrial production in real time. We select the model with the better fit.²⁵ The

²⁵We take industrial production from the real-time database of the Bundesbank and run the Bry-Boschan algorithm on the information set available at each point in time. This implies that the real-time Bry-Boschan algorithm gives different results at sample ends than the ex post Bry-Boschan algorithm because identification of

Figure 6: Recursively estimated means for MS(2)-DFM and MS(3)-DFM

Note: Means of an MS(2)-DFM and MS(3)-DFM, estimated recursively with real-time data. Blue line: first state, red line: second state, yellow line: third state.

second criterion is the BIC which may have the advantage over the QPS that it balances fit against parsimony.²⁶ Both criteria are applied exclusively to the information sets available at each point in time to make sure this is in fact a real-time model selection without any benefit of hindsight.

a turning point requires a lag of $d = 5$ months. While this means that model selection may react with a delay in real time, it is probably exactly the way an applied researcher would proceed who does not benefit from hindsight. For the Great Recession we therefore find that the real-time Bry-Boschan algorithm detects the recession start of May 2008 not before using the industrial production data vintage released on 7 November 2008 which includes the first five recession months May to September 2008.

²⁶Smith et al. (2006) propose a specific Markov-switching specific criterion. However, it is designed for models in which all parameters switch and thus does not work with our model. They also show that the Akaike information criterion (AIC) always selects the model with more states. Hence, we prefer the BIC over the AIC.

Figure 7: Recursive differences in QPS and BIC between MS(2)-DFM and MS(3)-DFM

Notes: Panel (a) shows the recursively computed $dQPS = QPS(2) - QPS(3)$ with 95% confidence bands. Panel (b) shows the recursively computed $dBIC = BIC(2) - BIC(3)$. In panel (b) we trim the observation of December 2005 (-72.56) to -40 to make the graph better readable.

Figure 7 plots the differences between the two-state and three-state models in terms of the QPS, $dQPS = QPS(2) - QPS(3)$, and in terms of the BIC, $dBIC = BIC(2) - BIC(3)$, based on the real-time estimates over the period January 2001 until July 2016. In both cases, a positive value indicates an advantage of the MS(3)-DFM. We find that the two-state model is superior up to the end of 2008, while the three-state model is favored thereafter. The exact change dates are very similar: the $dBIC$ selects November 2008 as the first month with an advantage of the three-state model, while the $dQPS$ identifies December 2008. Note that this date coincides with the aforementioned break in the recursively estimated state-specific means of the two-state model, see panel (a) of Figure 6. Hence, a user of these models would have noticed by December 2008 that introducing a third state is necessary to obtain a well-specified model.

Panel (a) of Figure 8 shows the smoothed nowcast probabilities of a combination of the

MS(2)/MS(3)-DFM, with the shift implemented in December 2008, for the whole sample, while panel (b) zooms in on the Great Recession period. The shift occurs when the probability for a severe recession reaches one in December 2008. The economy gets back to an ordinary recession in April 2009. This information about the magnitude of the recession might have been extremely helpful at this point in time as it perfectly matches the steepest part of the Great Recession: industrial production dropped by -7.2% in January 2009 and GDP dropped by -4.6% in the first quarter of 2009. Further taking the publication lag of two months for industrial production and one quarter for GDP into account, the model could have given timely information about the economic situation at that time and thus provided background for policy-makers to counteract the situation before knowing how deep the recession really was.

5.3 Forecasting German business cycle turning points

Markov-switching models can also be used to forecast future turning points. While nowcasting business cycle turning points in real time is generally difficult enough and accuracy deteriorates quickly with the forecast horizon (Hamilton, 2011), our selection of early indicators that lead GDP by up to three months enables us to directly filter the probabilities $Pr[S_t = i | \Psi_{t-3}]$ from the data. It turns out that the probability forecasts of both the two-state and three-state models are somewhat more volatile than the corresponding nowcasts. This is not surprising because less information is available. Technically, this is reflected in the fact that the nowcasts are smoothed probabilities while the 3-month ahead forecasts are only filtered.

To save space, we solely report the predicted recession probabilities of the combined MS(2)/MS(3)-DFM with the shift taking place in December 2008 as discussed above.²⁷ Figure 9 shows that they contain very useful information. For example, in July 2008 the model forecasts a recession with almost 100 percent probability for October which is remarkable as most forecasters identified the recession not before November (Heilemann and Schnorr-Bäcker, 2017). It also predicts the recovery very timely. The forecast made in January 2009 already predicts for April 2009 that the severe recession ends and the economy is back in a normal recession. And in March 2009 the model first predicts that the recession ends three months later in June 2009. We believe that this information would have been valuable at that time. For example, the German parliament passed a large stimulus package known as “Konjunkturpaket II” in February 2009, and in April the German public started to discuss another stimulus package because the end of the recession seemed far away.²⁸

5.4 Point forecasts of German GDP

Chauvet and Potter (2013) compare a large number of GDP-forecasting models including linear univariate and multivariate time series models, DSGE models and Markov-switching models. They find that MS-DFMs are by a large difference the most successful models in predicting GDP during US recessions in real time and even outperform expert forecasts from the Blue

²⁷Results for the single MS(2)-DFM and MS(3)-DFM models are available upon request.

²⁸The combined MS(2)/MS(3)-DFM forecasts include the same false alarms, and for the same reasons, as the respective nowcasts. Therefore, we do not discuss them here but refer the interested reader to the detailed analysis presented in Section 5.1.

Figure 8: Real-time nowcast of recession probabilities using an MS(2) and an MS(3)-DFM

(a) Full evaluation sample 2001M01 to 2016M06

(b) Great Recession sample

Note: Smoothed recession probabilities of an MS(2) (left panel) and an MS(3)-DFM (right panel) during the Great Recession recursively estimated with real-time data. The split indicates the shift from MS(2)-DFM to MS(3)-DFM. Shaded areas correspond to the recessions from the benchmark business cycle chronology from Section 4.4.

Chip Survey. To check whether they are also useful for predicting German GDP, we conduct an out-of-sample forecast experiment using real-time data. Our MS-DFMs do not include GDP and thus do not provide directly a GDP forecast. Therefore, we augment an autoregressive distributed lag (ADL) model for quarterly GDP growth with the estimated factor and the

Figure 9: Real-time forecast recession probabilities of MS(2)/MS(3)-DFM

smoothed recession probabilities,

$$\Delta \log(\text{GDP}_{t+h}) = c + \sum_{j=1}^p \alpha_j \Delta \log(\text{GDP}_{t-j}) + \sum_{j=0}^r \gamma_j f_{t-j} + \sum_{j=0}^s \delta_j \Pi_{t-j} + \varepsilon_t, \quad (13)$$

where h denotes the forecast horizon, f_t denotes the quarterly average of the monthly factor, and Π_{t-j} is the quarterly average of the smoothed probability that period $t-j$ experiences a recession. Note that we use a direct rather than an iterative forecasting procedure. We compare the performance of the following forecasting models including a nested benchmark AR-model:

- AR: Our benchmark is a purely autoregressive model with p lags ($\gamma_j = \delta_j = 0$).
- ADL-DFM(1): This is a one-state, i.e. linear, dynamic factor model including p lags of GDP growth and r lags of the factor and $\delta_j = 0$ as there are no switches between states. We consider this model in order to check whether including additional information via a linear factor is already sufficient to improve upon the benchmark AR forecasts or whether a Markov-switching framework is essential.
- ADL-DFM(2) and ADL-DFM(3): These are ADL models which include p lags of GDP growth, r lags of a state-dependent factor and s lags of the recession probabilities generated by the MS-DFM(2) and the MS-DFM(3) model, respectively. For the latter it turned out that distinguishing between mild and severe recessions did not improve forecasting power which is why we only report results based on the joint probability $Pr[S_t = 2|\Psi_T] + Pr[S_t = 3|\Psi_T]$.

- ADL-DFM(2&3): This is an ADL model which includes p lags of GDP growth and r lags of the factor and s lags of the recession probabilities generated by the MS-DFM(2) or MS-DFM(3) depending on which one is preferred by the BIC. The switch from the MS-DFM(2) to the MS-DFM(3) occurs in the fourth quarter of 2008.

We recursively construct real-time nowcasts ($h = 0$) and h -step forecasts for $h = 1, \dots, 4$ quarters based on an expanding window of vintage data.²⁹ Since we apply direct-step forecasting, for each model we consider one lag order specification per forecast horizon h . We proceed as follows. It is a well-known feature of German GDP growth that it has almost no autocorrelation (see, e.g., Pirschel and Wolters, 2018, for a comparison of autocorrelation functions of German and US GDP). Therefore, we include only one lag of GDP ($p = 1$) in all specifications. The recession probability Π_t is a first-order Markov process and includes by construction all relevant information which is why, at least theoretically, it is not necessary to include distributed lags. Since, in addition, Π_t leads GDP by one quarter, we include solely its first lag in the nowcast specifications ($s = 1, \gamma_0 = 0$) and its contemporaneous value in the forecast specifications ($s = 0$).³⁰ The factor f_t also leads GDP by one quarter. Therefore, we again exclude its current value from the nowcast specifications ($\delta_0 = 0$), but include it in the forecast specifications. At each recursion of our out-of-sample forecasting experiment we then choose the maximum lag order r as the one that minimizes the BIC.

We evaluate nowcasts and forecasts over the sample 2001Q1 to 2016Q2. Since GDP figures are subject to data revisions, we compare each forecast with the realisation published two quarters later. For example, a nowcast of 2001Q1 is compared with the value released by the end of 2001Q3. Exceptions are the major revisions of the German national account in 2005, 2011 and 2014. Here we use the last release before the revision to ensure that we take into account early data revisions but abstract from benchmark revisions which are difficult to forecast.

Before evaluating nowcasts and forecasts systematically based on RMSEs we graphically inspect the main characteristics of the different forecasting models. Figure 10 shows nowcasts and one-step ahead forecasts for the AR, ADL-DFM(1), ADL-DFM(2) and the ADL-DFM(3) model. The simple AR model captures the mean of GDP growth well, but mostly misses expansions and recessions. Both nowcast and forecast are basically flat until the Great Recession when the AR model reacts too late and too moderate. Afterwards, the nowcast becomes somewhat more accurate, while the forecast remains always close to the sample mean. Given the weak autocorrelation of German GDP, the result is not surprising.

The ADL-DFM(1) includes additional information via the factor estimated from a one-state, and thus linear, dynamic factor model. It turns out that this information and in particular the leading property of the factor is extremely valuable in generating accurate predictions. The model detects most turning points and misses only some episodes like the strong expansion in

²⁹A real-time nowcast of, say, 2001Q1 uses the data vintage available at the end of this quarter which includes—due to its one-quarter publication lag—GDP until 2000Q4. Correspondingly, an h -step ahead forecast is based on the data vintage available at the end of 2001Q1– h which includes GDP until 2000Q4– h . Note that some indicators also have publication lags but the dynamic factor model has not because it filters out the last observation of the vintage based on the information contained in surveys and commodity prices that are published without delay.

³⁰We checked specifications that allowed higher choices for s but got worse RMSEs which supports our argument.

Figure 10: Real-time nowcasts and one-quarter ahead forecasts of GDP growth

Note: Nowcasts and forecasts are based on real-time GDP and quarterly averages of monthly recession probabilities and dynamic factors. Shaded areas correspond to recessions according to the Bry-Boschan algorithm.

2006 or the spike of GDP growth in 2010. It gets the timing of the largest drop in GDP during the Great Recession right, even one quarter in advance, though by far not its actual depth. After the Great Recession the model considerably overestimates the strength of the recovery.

Turning to the Markov-switching models, both the ADL-DFM(2) and ADL-DFM(3) models improve during expansions and normal recessions only slightly upon the ADL-DFM(1) model. This changes, however, during the Great Recession and the subsequent recovery when their nowcasts for 2009Q1 are almost exactly correct and their one-quarter ahead forecasts for 2009Q1 outperform the ADL-DFM(1) model by a noticeable amount, even though they still underpredict

the actual depth of the recession. During the recovery from the Great Recession, they again make more accurate predictions than the ADL-DFM(1) model. It is further noticeable that except for the Great Recession the differences between the nowcasts and the 1-quarter ahead forecast are surprisingly small for all versions of the ADL-DFM framework. The 1-quarter forecasts are almost as accurate as the nowcast.

Based on the graphical analysis we conclude that in normal times it is sufficient to use the leading information extracted by a linear factor model. In contrast, during highly volatile times like the Great Recession and the subsequent recovery when disagreement among forecasters is usually high (see e.g. (Dovern, 2015)), predictions improve substantially when applying the MS-DFM(2) and MS-DFM(3) to account for the potential nonlinearity induced by those extraordinary business cycle movements.

As to the question whether to specify two or three states, we find that the predictions of the ADL-DFM(2) and ADL-DFM(3) are extremely close to each other, both before and after the Great Recession. Hence, using the information provided by the three-state Markov-switching model throughout the entire sample does not worsen GDP forecast accuracy despite the erratic switches between states before the Great Recession documented, *inter alia*, in Figure 5. Consequently, it does not make a difference here when the real-time model selection approach discussed above is applied. Since the shift from two to three states is detected in 2008Q4, the predictions of the ADL-DFM(2&3) model, which uses the information provided by the combined MS(2)/MS(3)-DFM, equal the predictions of the ADL-DFM(2) until the 2008Q4 and the ADL-DFM(3) thereafter. This is why we do not include these predictions in Figure 10.

Table 5 reports RMSEs relative to the AR model for forecasts up to $h = 4$. To test whether the forecast are significantly different from the benchmark AR-model, we employ the test proposed by Clark and West (2007). We find that all factor models provide significantly better predictions than the AR benchmark up to forecast horizon $h = 2$, with decreasing margin as the forecast horizon h increases. For a horizon of $h = 3$, only the ADL-DFM(1) outperforms the benchmark, and for $h = 4$ the AR model dominates even if not significantly so. These results are not surprising as by construction the factor leads GDP by only one quarter. Hence, for higher forecast horizons, the information provided by the factor models is much less relevant while the additional parameter estimation uncertainty remains unchanged. However, using the information from the MS-DFM is beneficial for forecasting GDP growth in spite of the low persistence of German GDP. In line with the graphical inspection, we also find that differences between the ADL-DFM(2), ADL-DFM(3), and ADL-DFM(2&3) models are rather small, especially for forecast horizons of up to two quarters.

The graphical analysis showed that the Markov-switching models perform particularly well during recessions. Hence, it is of interest to analyze differences in forecast precision between recessions and expansions systematically. To this end, we employ the quarterly version of the Bry-Boschan algorithm (Harding and Pagan, 2002) to GDP. The recession subsample includes 11 quarters (2002Q4-2003Q1, 2004Q3-2005Q1, 2008Q2-2009Q1, and 2012Q4-2013Q1), while the expansion subsample covers the remaining 55 quarters. Table 6 reports the corresponding RMSEs relative to the AR-benchmark model. These results confirm the finding by Chauvet and Potter (2013) that the advantage of Markov-switching models is largest during recessions. Interestingly,

Table 5: Relative RMSEs

Model	$h=0$	$h=1$	$h=2$	$h=3$	$h=4$
ADL-DFM(1)	0.7669***	0.8815***	0.9205***	0.9431***	1.2943
ADL-DFM(2)	0.6565***	0.8215***	0.8839***	1.3114	1.4523
ADL-DFM(3)	0.6472***	0.8616***	0.8983***	1.2471	1.5145
ADL-DFM(2&3)	0.6426***	0.8602***	0.9072***	1.2653	1.5201

Notes: Root mean squared errors relative to an AR-benchmark model. *, ** and *** denotes significance on the 10% level, 5% level and 1% level, respectively, according to the Clark-West test with Newey-West standard errors.

these models also improve upon the linear factor model during expansions, albeit to a smaller extent.

Table 6: Relative RMSEs for recessions

Model	$h=0$	$h=1$	$h=2$	$h=3$	$h=4$
Recessions					
ADL-DFM(1)	0.6186**	0.8154**	0.8567**	0.9100**	1.3229
ADL-DFM(2)	0.4910**	0.7274***	0.7634**	0.9923	1.5005
ADL-DFM(3)	0.5111**	0.7920**	0.8611**	1.1746	1.4960
ADL-DFM(2&3)	0.5019**	0.7831**	0.8643**	1.1715	1.5006
Expansions					
ADL-DFM(1)	0.9812***	1.0186***	1.0665***	1.0184	1.2160
ADL-DFM(2)	0.8816***	1.0066***	1.1353***	1.8687	1.3172
ADL-DFM(3)	0.8403***	1.0048***	0.9871**	1.4066	1.5621
ADL-DFM(2&3)	0.8406***	1.0171***	1.0087***	1.4669	1.5705

Notes: Relative root mean squared errors during recessions and expansions. *, ** and *** denotes significance on the 10% level, 5% level and 1% level, respectively, according to the Clark-West test with Newey-West standard errors.

6 Conclusion

We provide evidence that Markov-switching dynamic factor models together with a flexible variable pre-selection algorithm are an appropriate device to predict and date business cycle turning points for the German economy. It turns out that a three-state model is more sensitive than a two-state model and provides a better ex post characterization of the German business cycle, especially because it identifies the Great Recession as a severe recession. Using real-time data we show that nowcasts and one-quarter ahead forecasts capture business cycle dynamics in Germany well even though German GDP growth is characterized by very low persistence.

During the Great Recession the model predicts the timing of events one quarter in advance starting with the initially mild downturn, the severe recessionary phase afterwards, and finally the recovery. Further, a comparison of the two- and three-state model clearly signals that the three-state model would have been preferable in December 2008 right before the biggest downturn of the German economy. Hence, for professional forecasters using this framework

during the Great Recession would have been valuable to predict events systematically based on leading indicators. Moreover, the framework would have been highly useful for policymakers in order to plan the timing of policies to mitigate the crisis without the danger of stimulating the economy when the recovery was already on the way.

References

- Aastveit, K. A., Jore, A. S., and Ravazzolo, F. (2016). Identification and real-time forecasting of Norwegian business cycles. International Journal of Forecasting, 32(2):283–292.
- Abberger, K. and Nierhaus, W. (2010). Markov-Switching and the Ifo Business Climate: the Ifo Business Cycle Traffic Lights. OECD Journal: Journal of Business Cycle Measurement & Analysis, 2010(2).
- Altissimo, F., Cristadoro, R., Forni, M., Lippi, M., and Veronese, G. (2010). New Eurocoin: tracking economic growth in real time. The Review of Economics and Statistics, 92(4):1024–1034.
- Anas, J., Billio, M., Ferrara, L., and Mazzi, G. L. (2008). A system for dating and detecting turning points in the euro area. Manchester School, 76(5):549–577.
- Artis, M., Krolzig, H.-M., and Toro, J. (2004). The European business cycle. Oxford Economic Papers, 56(1):1–44.
- Bai, J. and Ng, S. (2008). Forecasting economic time series using targeted predictors. Journal of Econometrics, 146(2):304–317.
- Bandholz, H. and Funke, M. (2003). In search of leading indicators of economic activity in Germany. Journal of Forecasting, 22(4):277–297.
- Boivin, J. and Ng, S. (2006). Are more data always better for factor analysis? Journal of Econometrics, 132:169–194.
- Boldin, M. D. (1996). A check on the robustness of Hamilton’s Markov Switching model approach to the economic analysis of the business cycle. Studies in Nonlinear Dynamics & Econometrics, 1(1).
- Camacho, M. and Martinez-Martin, J. (2015). Monitoring the world business cycle. Economic Modelling, 51:617–625.
- Camacho, M. and Perez-Quiros, G. (2007). Jump-and-rest effect of US business cycles. Studies in Nonlinear Dynamics & Econometrics, 11(4).
- Camacho, M. and Perez-Quiros, G. (2010). Introducing the Euro-STING: short term indicator of Euro area growth. Journal of Applied Econometrics, 25(4):663–694.
- Camacho, M., Perez-Quiros, G., and Poncela, P. (2014). Green shoots and double dips in the euro area: A real time measure. International Journal of Forecasting, 30(3):520–535.

- Camacho, M., Perez-Quiros, G., and Poncela, P. (2015). Extracting nonlinear signals from several economic indicators. Journal of Applied Econometrics, 30(7):1073–1089.
- Camacho, M., Perez-Quiros, G., and Poncela, P. (2018). Markov-switching dynamic factor models in real time. International Journal of Forecasting, 34(4):598–611.
- Chauvet, M. (1998). An econometric characterization of business cycle dynamics with factor structure and regime switching. International Economic Review, 39(4):969–996.
- Chauvet, M. (2001). A monthly indicator of Brazilian GDP. Brazilian Review of Econometrics, 21(1):1–15.
- Chauvet, M. and Hamilton, J. D. (2006). Dating business cycle turning points. Contributions to Economic Analysis, 276:1–54.
- Chauvet, M. and Potter, S. (2013). Forecasting output. Handbook of Economic Forecasting, 2(Part A):141–194.
- Clark, T. E. and West, K. D. (2007). Approximately normal tests for equal predictive accuracy in nested models. Journal of Econometrics, 138(1):291–311.
- Diebold, F. X. and Rudebusch, G. D. (1996). Measuring business cycles: A modern perspective. The Review of Economics and Statistics, 78(1):67–77.
- Döpke, J., Fritsche, U., and Pierdzioch, C. (2017). Predicting recessions with boosted regression trees. International Journal of Forecasting, 33(4):745–759.
- Dovern, J. (2015). A multivariate analysis of forecast disagreement: Confronting models of disagreement with survey data. European Economic Review, 80:16–35.
- Doz, C. and Petronevich, A. (2016). Dating Business Cycle Turning Points for the French Economy: An MS-DFM approach. Dynamic Factor Models (Advances in Econometrics, Volume 35) Emerald Group Publishing
- International Journal of Forecasting,
28(2):428–445.
- Efron, B., Hastie, T., Johnstone, I., and Tibshirani, R. (2004). Least angle regression. The Annals of Statistics, 32(2):407–499.
- Eo, Y. and Kim, C.-J. (2016). Markov-switching models with evolving regime-specific parameters: Are postwar booms or recessions all alike? Review of Economics and Statistics, 98(5):940–949.
- Ferrara, L. (2003). A three-regime real-time indicator for the US economy. Economics Letters, 81(3):373–378.

- Flaig, G., Nierhaus, W., Kuntze, O.-E., Gebauer, A., Henzel, S., Hülsewig, O., Dehne, A., Langmantel, E., Meister, W., Ruschinski, M., and Schimpfermann, B. (2005). ifo Konjunkturprognose 2005/2006: Nur zögerliche Erholung. ifo Schnelldienst, 58(12):29–63.
- Fritsche, U. and Stephan, S. (2002). Leading Indicators of German Business Cycles. An Assessment of Properties/Frühindikatoren der deutschen Konjunktur. Eine Beurteilung ihrer Eigenschaften. Journal of Economics and Statistics, 222(3):289–315.
- Hamilton, J. D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. Econometrica, 57(2):357–384.
- Hamilton, J. D. (2005). What’s real about the business cycle? Federal Reserve Bank of St. Louis Review, 87(4):435–452.
- Hamilton, J. D. (2011). Calling recessions in real time. International Journal of Forecasting, 27(4):1006–1026.
- Harding, D. and Pagan, A. (2002). Dissecting the cycle: a methodological investigation. Journal of Monetary Economics, 49(2):365–381.
- Hastie, T., Tibshirani, R., and Friedman, J. (2017). Elements of statistical learning. Springer.
- Heilemann, U. and Schnorr-Bäcker, S. (2017). Could the start of the German Recession 2008-2009 have been foreseen? Evidence from real-time data. Journal of Economics and Statistics, 237(1):29–62.
- Ho, W.-Y. A. and Yetman, J. (2012). Does US GDP stall? BIS Working Papers 387, Bank for International Settlements.
- Hoerl, A. E. and Kennard, R. W. (1970). Ridge Regression: Biased Estimation for Nonorthogonal Problems. Technometrics, 12(1):55–67.
- Ivanova, D., Lahiri, K., and Seitz, F. (2000). Interest rate spreads as predictors of German inflation and business cycles. International Journal of Forecasting, 16(1):39–58.
- Kholodilin, K. A. (2005). Forecasting the German Cyclical Turning Points: Dynamic Bi-Factor Model with Markov Switching. Journal of Economics and Statistics, 225(6):653–674.
- Kholodilin, K. A. and Siliverstovs, B. (2006). On the forecasting properties of the alternative leading indicators for the German GDP: Recent evidence. Journal of Economics and Statistics, 226(3):234–259.
- Kim, C.-J. (1994). Dynamic linear models with Markov-switching. Journal of Econometrics, 60(1-2):1–22.
- Kim, C.-J. and Nelson, C. R. (1998). Business cycle turning points, a new coincident index, and tests of duration dependence based on a dynamic factor model with regime switching. Review of Economics and Statistics, 80(2):188–201.

- Kim, M.-J. and Yoo, J.-S. (1995). New index of coincident indicators: A multivariate Markov switching factor model approach. Journal of Monetary Economics, 36(3):607–630.
- Krolzig, H.-M. and Toro, J. (2001). A New Approach to the Analysis of Business Cycle Transitions in a Model of Output and Employment. Economics Series Working Papers 59, University of Oxford, Department of Economics.
- Layton, A. P. and Smith, D. (2000). A further note on the three phases of the US business cycle. Applied Economics, 32(9):1133–1143.
- Lehmann, R. and Wohlrabe, K. (2016). Looking into the black box of boosting: the case of Germany. Applied Economics Letters, 23(17):1229–1233.
- Mariano, R. S. and Murasawa, Y. (2003). A new coincident index of business cycles based on monthly and quarterly series. Journal Applied Econometrics, 18(4):427–443.
- McAdam, P. (2007). USA, Japan and the Euro Area: Comparing Business-Cycle Features. International Review of Applied Economics, 21(1):135–156.
- Morley, J. and Piger, J. (2012). The asymmetric business cycle. Review of Economics and Statistics, 94(1):208–221.
- Nalewaik, J. J. (2011). Forecasting recessions using stall speeds. Finance and Economics Discussion Series 2011-24, Federal Reserve Board.
- Ng, S. (2014). Boosting recessions. Canadian Journal of Economics, 47(1):1–34.
- Pirschel, I. and Wolters, M. H. (2018). Forecasting with large datasets: compressing information before, during or after the estimation? Empirical Economics, 55(2):573–596.
- Proaño, C. R. (2017). Detecting and predicting economic accelerations, recessions, and normal growth periods in real-time. Journal of Forecasting, 36:26–42.
- Proaño, C. R. and Theobald, T. (2014). Predicting german recessions with a composite real-time dynamic probit indicator. International Journal of Forecasting, 30:898–917.
- Schirwitz, B. (2009). A comprehensive German business cycle chronology. Empirical Economics, 37(2):287–301.
- Schumacher, C. (2010). Factor forecasting using international targeted predictors: The case of German GDP. Economics Letters, 107(2):95–98.
- Sichel, D. E. (1994). Inventories and the three phases of the business cycle. Journal of Business & Economic Statistics, 12(3):269–277.
- Smith, A., Naik, P. A., and Tsai, C.-L. (2006). Markov-switching model selection using Kullback–Leibler divergence. Journal of Econometrics, 134(2):553–577.
- Stock, J. H. and Watson, M. W. (2005). Understanding changes in international business cycle dynamics. Journal of the European Economic Association, 3(5):968–1006.

Vanhaelen, J.-J., Dresse, L., and De Mulder, J. (2000). The Belgian industrial confidence indicator: leading indicator of economic activity in the euro area? Working Paper 12, National Bank of Belgium.

Zou, H. and Hastie, T. (2005). Regularization and variable selection via the elastic net. Journal of the Royal Statistical Society: Series B (Statistical Methodology), 67(2):301–320.

Appendix

A Construction of the state space form

We start defining the $(12 + nq)$ -dimensional state vector

$$a_t = [f_t, \dots, f_{t-11}, z'_t, \dots, z'_{t-q+1}]'.$$

Now the measurement equations (1) and (2) can be jointly written as

$$y_t = Ba_t,$$

where

$$B = \begin{bmatrix} \gamma^{(h)} & 0_{n_h \times 1} & \cdots & 0_{n_h \times 1} & I_{n_h} & 0_{n_h \times n_s} & 0_{n_h \times (nq-q)} \\ \gamma^{(s)} & \gamma^{(s)} & \cdots & \gamma^{(s)} & 0_{n_s \times n_h} & I_{n_s} & 0_{n_s \times (nq-q)} \end{bmatrix}$$

and $\gamma^{(h)} = [\gamma_1^{(h)}, \dots, \gamma_{n_h}^{(h)}]'$ and $\gamma^{(s)} = [\gamma_1^{(s)}, \dots, \gamma_{n_s}^{(s)}]'$.

The transition equation can be written as

$$a_t = \mu_{S_t} + Fa_{t-1} + R\omega_t$$

using the following definitions. The system matrix is

$$F = \begin{bmatrix} F_{11} & 0_{12 \times nq} \\ 0_{nq \times 12} & F_{22} \end{bmatrix}$$

where F_{11} is the (12×12) -dimensional companion matrix of an AR(12) process with lag coefficients ϕ_1 to ϕ_{12} of which coefficients 3 to 12 restricted to zero because we only allow a maximum lag order of $p = 2$ for f_t , and F_{22} is the $(nq \times nq)$ -dimensional companion matrix of an n -dimensional VAR process with q lags and coefficient matrices ψ_1 to ψ_q . The intercept vector is nonzero only for f_t and thus is

$$\mu_{S_t} = [\beta_{S_t}, 0_{1 \times (11+nq)}]'$$

The vector of iid shocks, $\omega_t = [\eta_t, \varepsilon'_t]'$, is iid normally distributed with mean zero and diagonal covariance matrix

$$Q \equiv E(\omega_t \omega'_t) = \begin{bmatrix} 1 & 0_{1 \times n} \\ 0_{n \times 1} & \Sigma_z \end{bmatrix}.$$

Finally, we define the coefficient matrix

$$R = \begin{bmatrix} R_{11} & 0_{12 \times n} \\ 0_{nq \times 1} & R_{22} \end{bmatrix},$$

where $R_{11} = [1, 0_{1 \times 11}]'$ and $R_{22} = [I_n, 0_{n \times (nq-n)}]'$.

B Estimation of the MS-DFM

We employ the filter proposed by Kim (1994) to estimate the MS-DFM. Based on the initialization $a_{0|0} = (I - F)^{-1}\mu_{S_t}$ and $P_{0|0} = (I - F \otimes F)^{-1}\text{vec}(Q)$, the recursion consists of the usual prediction and updating steps. To this end, let us define $P_{t|t-1}^{(j,i)}$ as the variance of z_t conditional on Ψ_{t-1} , the information available in $t-1$, and on $S_t = j$ and $S_{t-1} = i$, $P_{t|t}^{(i)}$ as the variance of z_{t-1} conditional on Ψ_t and $S_{t-1} = i$, and equivalently $a_{t|t-1}^{(j,i)}$ and $a_{t-1|t-1}^{(i)}$. Then the prediction step is

$$a_{t|t-1}^{(j,i)} = F a_{t-1|t-1}^{(i,k)} + \mu_{S_t}^{(j)}, \quad (14)$$

$$P_{t|t-1}^{(j,i)} = F P_{t-1|t-1}^{(i,k)} F' + R Q R', \quad (15)$$

and the updating step is

$$a_{t|t}^{(j,i)} = a_{t|t-1}^{(j,i)} + K_t^{(j,i)}(y_t - B a_{t|t-1}^{(j,i)}), \quad (16)$$

$$P_{t|t}^{(j,i)} = (I_{2n+p} - K_t^{(j,i)} B) P_{t|t-1}^{(j,i)}, \quad (17)$$

where the Kalman gain is defined by $K_t^{(j,i)} = P_{t|t-1}^{(j,i)} B' (B P_{t|t-1}^{(j,i)} B')^{-1}$. However, each recursion generates an m -fold increase in the number of states to be considered. Therefore, we apply the approximation by Kim (1994),

$$a_{t|t}^{(j)} = \frac{\sum_{i=1}^m \Pr[S_{t-1} = i, S_t = j | \Psi_t] a_{t|t}^{(j,i)}}{\Pr[S_t = j | \Psi_t]} \quad (18)$$

$$P_{t|t}^{(j)} = \frac{\sum_{i=1}^m \Pr[S_{t-1} = i, S_t = j | \Psi_t] (P_{t|t}^{(j,i)} + (a_{t|t}^{(j)} - a_{t|t}^{(j,i)})(a_{t|t}^{(j)} - a_{t|t}^{(j,i)})')}{\Pr[S_t = j | \Psi_t]}, \quad (19)$$

which reduces the number of possible states of $a_{t|t}$ and $P_{t|t}$ to m per period by taking weighted averages over the states and feeding them into the prediction steps (14) and (15).

The corresponding log likelihood function is obtained by Hamilton (1989):

$$\ln L = \sum_{t=1}^T \ln \left(\sum_{S_t=1}^m \sum_{S_{t-1}=1}^m f(y_t | S_t, S_{t-1}, \Psi_{t-1}) \Pr(S_t = j, S_{t-1} = i | \Psi_{t-1}) \right). \quad (20)$$

Evaluating it requires calculating the weights $\Pr(S_t = j, S_{t-1} = i | \Psi_{t-1})$, which can be expressed as the product of the probability of being in a certain regime at period $t-1$ and the corresponding transition probability:

$$\Pr(S_t = j, S_{t-1} = i | \Psi_{t-1}) = p_{ij} \Pr(S_{t-1} = i | \Psi_{t-1}). \quad (21)$$

Updating this probability with information up to period t yields the filtered probabilities:

$$\begin{aligned} Pr(S_t=j, S_{t-1}=i|\Psi_t) &= \frac{f(S_t=j, S_{t-1}=i, y_t|\Psi_{t-1})}{f(y_t|I_{t-1})} \\ &= \frac{f(y_t|S_t=j, S_{t-1}=i, \Psi_{t-1})Pr(S_t=j, S_{t-1}=i|\Psi_{t-1})}{\sum_{S_t=1}^m \sum_{S_{t-1}=1}^m f(y_t|S_t=j, S_{t-1}=i, \Psi_{t-1})Pr(S_t=j, S_{t-1}=i|\Psi_{t-1})} \end{aligned}$$

and

$$Pr(S_t=j|\Psi_t) = \sum_{i=1}^m Pr[S_{t-1}=i, S_t=j|\Psi_t].$$

Based on an initialization—we employ the unconditional probabilities as derived by Hamilton (1989)—the steps can be iterated forward over the sample to obtain the filtered probabilities for each period. Along with the filter recursions, this yields all the information we need to estimate the latent dynamic regime dependent factor as well as the Markov-switching probabilities.

C LARS-EN algorithm

In the following, we explain in more detail how the elastic net works and present results for the full sample. Let us focus on the selection of hard indicators, $y_{it}^{(h)}$, since the selection of the soft indicators works equivalently. The aim is to choose those hard indicators that jointly predict quarterly GDP growth well. We start from the quarterly predictive regression (7),

$$\Delta \log(GDP_t) = \sum_{i=1}^{16} \sum_{l=1}^3 b_{i,l}^{(h)} y_{i,t-l}^{(h)} + u_t^{(h)}, \quad (22)$$

where $\Delta \log(GDP_t)$ is centered at zero and all regressors are standardized. Applying OLS would yield, in general, nonzero parameters for all three lags of all 16 indicators. To obtain a sparse solution, i.e., a solution that contains parameter estimates of zero, and thus really selects indicators, we estimate the parameters by means of the elastic net (EN) proposed by Zou and Hastie (2005). To this end, we define the $T \times 1$ vector $y = (\Delta \log(GDP_1), \dots, \Delta \log(GDP_T))'$ and the $T \times 48$ matrix X with rows

$$X_t = (y_{1,t-1}^{(h)}, y_{1,t-2}^{(h)}, y_{1,t-3}^{(h)}, \dots, y_{16,t-1}^{(h)}, y_{16,t-2}^{(h)}, y_{16,t-3}^{(h)})$$

and corresponding 48×1 vector of coefficients

$$b = (b_{1,1}^{(h)}, b_{1,2}^{(h)}, b_{1,3}^{(h)}, \dots, b_{16,1}^{(h)}, b_{16,2}^{(h)}, b_{16,3}^{(h)})'.$$

Then we solve the elastic net optimization problem

$$L = (\lambda_1, \lambda_2, b) = |y - Xb|^2 + \lambda_1 |b|_1 + \lambda_2 |b|^2, \quad (23)$$

where $|\cdot|_1$ and $|\cdot|^2$ denote the L_1 and L_2 norm, respectively. The specific shape of the L_1 norm induces, for sufficiently large λ_1 , a sparse solution that can be interpreted as regressor subset

selection, see Hastie et al. (2017). We follow Zou and Hastie (2005) who show that the elastic net optimization problem can be rewritten as a LASSO optimization problem which can be solved by an adaption of the least angle regression (LARS) originally proposed by Efron et al. (2004). The adaption to the elastic net, called LARS-EN, allows to transform the two tuning parameters λ_1 and λ_2 into the tuple (k, λ_2) , where k is the number of regressors to be selected. The intuition behind it is simple: the larger we choose λ_1 , the more dominates the L_1 norm which favors a sparse solution. One can think of the LARS-EN algorithm as starting, for fixed λ_2 , from a very large value of λ_1 such that $\mathbf{b}$ is estimated as a zero vector. By successively lowering λ_1 , more and more nonzero parameter estimates show up and thus k increases. Since we intend to select three hard indicators, we set $k = 3$.

We also need to choose a value for the other tuning parameter, λ_2 , which determines the weight of the L_2 norm in the optimization problem. To understand how λ_2 affects the estimation problem, note that the elastic net collapses to the LASSO if $\lambda_2 = 0$. The LASSO is known to select almost arbitrarily only one predictor from a subset of highly correlated regressors. This is the so-called grouping effect. By the very nature of our problem — we intend to extract common business cycle information from a set of selected leading indicators — our regressors are potentially strongly correlated and we deliberately want to choose correlated ones for the subsequent factor model to make sense. Further note that the elastic net reduces to the ridge regression if $\lambda_1 = 0$. Ridge regression is able to deal with highly correlated regressors. In fact, it was originally motivated for the extreme case that the cross-product $\mathbf{X}'\mathbf{X}$ is not even invertible (Hoerl and Kennard, 1970). In general, the elastic net is a kind of combination of the LASSO and ridge regression. The larger we choose λ_2 , the more dominates the L_2 norm which allows efficient handling of correlated regressors and avoids the grouping effect. We experimented with different choices for λ_2 and found the value of 100 to work well which is in the range of values considered by Zou and Hastie (2005). The results turned out to be robust to choosing higher values but smaller values gave rise to the grouping effect.

To get an idea of how the elastic net works with our data, let us consider the selection of hard indicators in the regression (7) for the ex post analysis. We set $\lambda_2 = 100$. The LARS-EN algorithm starts with a prohibitively large $\lambda_1 = 95.2$ so that all parameters are estimated as zero. Successively lowering λ_1 allows the inclusion of more and more regressors. The upper panel of Figure 11 shows how the parameter estimates evolve step by step. In step $k = 1$, λ_1 is lowered to 82.9 which allows to include the first regressor, foreign orders of capital goods (lag 1), with parameter 3.17. In step $k = 2$, λ_1 is lowered to 63.6. Now the first regressor has a larger parameter, 6.77, and a second regressor, domestic orders of intermediate good (lag 1), is added with parameter 3.60. In step $k = 3$, λ_1 is lowered to 59.6 which allows to add domestic orders of capital goods (lag 1) as third regressor. Hence, this choice of λ_1 corresponds to our objective of $k = 3$ and we use the selected three indicators in our Markov-switching models. Of course, it is possible to take more steps and thus add more variables. To illustrate this, step $k = 4$ is also shown.

The upper panel of Table 7 reports the estimated parameters of the third LARS-EN step applied to the selection regression (7). As a comparison we also show the OLS estimates of the same parameters. (Of course, OLS yields nonzero estimates of all parameters but for ease of

Note: The colored lines indicate how the parameter estimates for the regressors stated in the legend change step by step. Each step k corresponds to a specific value λ_1 that allows to include another regressor. The dashed line indicates step $k = 3$.

presentation we leave them out here.) Clearly, the elastic net estimates are absolutely smaller than the unconstrained OLS estimates. The lower panels of Figure 11 and Table 7 show the results of the analogous selection regression (8) for the survey indicators. While the parameter values obviously change, the general procedure remains the same.

Table 7: Parameters of the selection regressions estimated by LARS-EN and OLS

Indicator	b	β_{OLS}
<i>Regression of GDP on hard indicators</i>		
Foreign orders of capital goods	7.38	10.92
Domestic orders of intermediate goods	4.18	8.81
Domestic orders of capital goods	0.65	7.59
<i>Regression of GDP on soft indicators</i>		
Overall production expectations	1.13	10.40
Ifo business expectations	1.05	14.42
Ifo export expectations	0.91	16.41

Notes: **b** in the upper and lower panels denotes the estimated parameters of equations (7) and (8) that are nonzero based on LARS-EN with $\lambda_2 = 100$ and $k = 3$. β_{OLS} denotes the respective parameter estimates obtained by OLS.

D Detailed estimation results

In this section we report the estimated autoregressive parameters of the idiosyncratic components of both the MS(2)-DFM and the MS(3)-DFM. Recall that the vector of idiosyncratic components, $z_t = [z_{1t}^{(h)}, z_{2t}^{(h)}, z_{3t}^{(h)}, z_{1t}^{(s)}, z_{2t}^{(s)}, z_{3t}^{(s)}]'$ is modeled as a diagonal VAR process of lag order $q = 2$ with diagonal covariance matrix. Hence, each component $i = 1, \dots, 6$ follows an independent AR(2) process with AR parameters $\psi_{i,1}$ and $\psi_{i,2}$, where $\psi_{i,j}$ is the i th diagonal element of the parameter matrix ψ_j defined in (3). Table 8 shows these parameters estimated by maximum likelihood. It turns out that, while being stationary by assumption, most idiosyncratic components are fairly persistent.

Table 8: Autoregressive parameters of the idiosyncratic components

Indicator	MS(2)-DFM		MS(3)-DFM	
	$\psi_{i,1}$	$\psi_{i,2}$	$\psi_{i,1}$	$\psi_{i,2}$
$z_{1t}^{(h)}$ Foreign orders of capital goods	-0.63 (0.06)	-0.31 (0.06)	-0.66 (0.06)	-0.34 (0.06)
$z_{2t}^{(h)}$ Domestic orders of intermediate goods	-0.32 (0.07)	-0.09 (0.07)	-0.41 (0.07)	-0.16 (0.07)
$z_{3t}^{(h)}$ Domestic orders of capital goods	-0.65 (0.06)	-0.21 (0.06)	-0.68 (0.06)	-0.24 (0.06)
$z_{1t}^{(s)}$ Overall production expectations	0.63 (0.08)	0.14 (0.07)	0.73 (0.07)	0.06 (0.07)
$z_{2t}^{(s)}$ Ifo business expectations	1.03 (0.07)	-0.21 (0.07)	1.09 (0.07)	-0.27 (0.07)
$z_{3t}^{(s)}$ Ifo export expectations	0.83 (0.07)	0.07 (0.07)	0.88 (0.07)	0.01 (0.07)

Notes: $\psi_{i,j}$ denotes the autoregressive parameter of idiosyncratic component i for lag j . In terms of the notation of equation (3), it is the i th diagonal element of the (diagonal) parameter matrix ψ_j . Estimated standard errors are reported in parentheses below the estimates.

E Data: indicators, sources, and real-time selection

The majority of the series is downloaded from Thomson Reuters Datastream, while the remaining indicators are directly obtained from the German Bundesbank, the ECB and the OECD. Tables 9 and 10 list the hard and survey indicators, respectively, together with their sources and the transformations we applied. For the hard indicators we report the sources for both our ex post analysis and our real-time analysis. The survey indicators are stationary by construction and thus left untransformed. They are published without (noticeable) revisions, hence the use of a specific real-time data set is not necessary.

The hard and survey indicators selected by the LARS algorithm in each step of our real-time analysis are reported in Figures 12 and 13. Note that we exclude from the Figures all variables that are never selected.

Table 9: Hard Indicators

Name of Series	Datastream code (ex post data set)	Real-time code (real-time data set)	Transformation	
			diff	log
<i>New Orders</i>				
Domestic orders of capital goods, Germany	BDDCPORDG	BBKRT.M.DE.Y.I.IO2.ACM03.C.I	yes	yes
Domestic orders of consumer goods, Germany	BDDCNORDG	BBKRT.M.DE.Y.I.IO2.ACM04.C.I	yes	yes
Domestic orders of intermediate goods, Germany	BDDBPORDG	BBKRT.M.DE.Y.I.IO2.ACM02.C.I	yes	yes
Foreign orders of capital goods, Germany	BDOCPORDG	BBKRT.M.DE.Y.I.IO3.ACM03.C.I	yes	yes
Foreign orders of consumer goods, Germany	BDOCNORDG	BBKRT.M.DE.Y.I.IO3.ACM04.C.I	yes	yes
Foreign orders of intermediate goods, Germany	BDOBPORDG	BBKRT.M.DE.Y.I.IO3.ACM02.C.I	yes	yes
<i>Interest rates</i>				
Yield on German federal securities, residual maturity 9 to 10 years	BDT0557	no revisions	yes	no
Fibor – 3 month (monthly average)	BDINTER3	no revisions	yes	no
Term spread on German federal securities - (10y-3m)	BDT0557-BDINTER3	no revisions	yes	no
<i>Commodity prices</i>				
Brent crude oil price, US-Dollar	OILBREN	no revisions	yes	yes
HWWA commodity price index, Euro	BDHWWAINF	no revisions	yes	yes
<i>General economic indicators</i>				
Dax performance index	DAXINDX	no revisions	yes	yes
German contribution to EMU M2	BDTXI302A	no revisions	yes	yes
Employed persons – overall economy, Germany	last vintage	BBKRT.M.DE.S.L.BE1.CA010.P.A	yes	yes
Consumer prices – all categories, Germany	last vintage	BBKRT.M.DE.Y.P.PC1.PC100.R.I	yes	yes
<i>Foreign markets</i>				
US industrial production	last vintage	from OECD.Stat	yes	yes

Notes: “no revisions” indicates that the series is assumed to be published without revisions, “last vintage” indicates that we use the last vintage of the real-time data set as ex post data.

Table 10: Survey Indicators

Name of Series	Datastream-Code	Source
Overall production expectations, Industry, Germany	BDTTA5BSQ	European commission
Intermediate goods production expectations, Industry, Germany	BDITM5.BQ	European commission
Investment goods production expectations, Industry, Germany	BDIVE5.BQ	European commission
Overall employment expectations, Germany	BDTTA7BSQ	European commission
Overall order books, Industry, Germany	BDTTA2BSQ	European commission
Consumer confidence indicator, Germany	BDCNFCONQ	European commission
Consumer survey: economic situation next 12 months, Germany	BDEUSCEYQ	European commission
Economic sentiment indicator, Germany	BDEUSESIG	European commission
Economic sentiment indicator, Euro Zone	EKEUSESIG	European commission
Ifo business climate index, Germany	BDCNFBUSQ	Ifo institute
Ifo business climate manufacturing, Germany	BDIFDMTLQ	Ifo institute
Ifo business climate capital goods, Germany	BDIFDMPLQ	Ifo institute
Ifo business expectations index, Germany	BDCYLEADQ	Ifo institute
Ifo business expectations manufacturing, Germany	BDIFDMTKQ	Ifo institute
Ifo business expectations capital goods, Germany	BDIFDMPKQ	Ifo institute
Ifo export expectations, Germany	BDIFDMTJQ	Ifo institute
ISM purchasing managers index, USA	USCNFBUSQ	ISM institute
Belgium business indicator, Belgium	BGCNFBUSQ	National Bank of Belgium
Euro-coin, Euro-Zone	EMECOIN.Q	CEPR

Note: Recursive real time variable selection with LARS-EN. Variables selected at $\circ$ = lag 1, $\times$ = lag 2, $*$ = lag 3. Variables which do not appear are never selected.

Note: Recursive real time variable selection with LARS-EN. Variables selected at ○ = lag 1, × = lag 2, * = lag 3. Variables which do not appear are never selected.