Kiel
Working Papers

Kiel Institute
for the World Economy

Systematic Errors in Growth
Expectations over the Business
Cycle

by Jonas Dovern and Nils Jannsen

1989| February 2015

Systematic Errors in Growth Expectations over the Business Cycle *

Jonas Dovern and Nils Jannsen

Abstract:
Using real-time data, we analyze how the systematic expectation errors of professional forecasters in 19 advanced economies depend on the state of the business cycle. Our results indicate that the general result that forecasters systematically overestimate output growth (across all countries) masks considerable differences across different business-cycle states. We show that forecasts for recessions are subject to a large negative systematic forecast error (forecasters overestimate growth), while forecasts for recoveries are subject to a positive systematic forecast error. Forecasts made for expansions have, if anything, a small systematic forecast error for large forecast horizons. When we link information about the business-cycle state in the target year with quarterly information about its state in the forecasting period, we find that forecasters realize business-cycle turning points somewhat late. Using cross-country evidence, we demonstrate that the positive relationship between a change in trend growth rates and forecast bias, as suggested in the literature, breaks down when only focusing on forecasts made for expansions.

Keywords: Macroeconomic expectations, forecasting, forecast bias, survey data

JEL classification: C5, E2, E3

Prof. Dr. Jonas Dovern
Heidelberg University
Bergheimer Str. 58
69115 Heidelberg, Germany
Telephone: +49-(0)6221 -54-2958
E-mail: jonas.dovern@awi.uni-heidelberg.de

Dr. Nils Jannsen**
Kiel Institute for the World Economy
Kiellinie 66
24105 Kiel, Germany
Telephone: +49 (0) 431-8814 298
E-mail: nils.jannsen@ifw-kiel.de

_____

1 Introduction

The literature on forecast evaluation has a long tradition of analyzing whether macroe-conomic forecasts are ex-ante unbiased. Since Mincer and Zarnowitz (1969) and Holden and Peel (1990), there have been numerous studies that analyze whether, for instance, growth or inflation expectations published by professional forecasters systematically under- or overestimate future growth or inflation outcomes, respectively.[1] This high interest in understanding the nature of macroeconomic expectations is due to the out-standing importance that expectations have for a wide range of macroeconomic matters, such as the implementation of monetary policy and fiscal rules, corporate investment decisions, and consumption-saving choices of households.[2]

However, with some exceptions (Sinclair et al., 2010; Messina et al., 2014; Loungani et al., 2013), the literature has been silent on whether such systematic forecast errors depend on the state of the business cycle.[3] Two issues in particular are not addressed at all in past studies. First, there is no broad evidence about the evolution of system-atic forecast errors around business-cycle turning points.[4] Second, past studies do not explicitly look at systematic forecast errors for recovery periods but instead assume a simple two-regime model of the business cycle that distinguishes only between re-cessions and expansions (see the discussion in Fildes and Stekler, 2002).[5] Time-series evidence suggests, however, that the early years of recoveries are distinct from more mature economic expansions (see, e. g., Kim et al., 2005; Boysen-Hogrefe et al., 2015); and anecdotal evidence (see, e. g., European Central Bank, 2014, Box 6) suggests there might be a tendency of forecasters to underestimate the strength of recoveries.

We address these issues by making a threefold contribution to the literature. First, using growth forecasts for a large panel of advanced economies provided by *Consensus*

---

[1]For some recent examples of studies that look at issues related to forecast bias, see inter alia Batchelor (2007), Ager et al. (2009), Ashiya (2009), Dovern and Weisser (2011), and Deschamps and Ioannidis (2013).

[2]For a detailed discussion of the relevance of expectations for macroeconomic policy, see, e. g., Wieland and Wolters (2013).

[3]We believe that use of the term 'bias' should be reserved for the unconditional expectation of a forecast error. Therefore, we prefer to speak of 'systematic (forecast) errors' whenever we make statements that are conditional on the state of the business cycle.

[4]Loungani et al. (2013) use an annual classification scheme to identify recession *years* that does not allow tracking of the evolution of forecasts (relative to actual outcomes) as a function of the temporal distance to recession starts.

[5]A notable exception is the earlier evidence provided in Zarnowitz and Braun (1993). Furthermore, Loungani (2002) analyzes the accuracy of forecasts made for recovery years but does not, however, investigate the associated systematic forecast error.

Economics and real-time data on GDP outcomes, we provide broad-based evidence on whether professional forecasts exhibit state-dependent systematic forecast errors. Second, we identify recovery periods and treat them as a distinct business-cycle phase. Third, we provide new evidence on the pattern of systematic forecast errors around business-cycle turning points by combining annual information about the state of the business cycle in the target year of the forecasts with quarterly information about its state in the forecasting period. Our approach allows a more detailed analysis of this process than an approach that uses only annual information, as the latter give only a very rough indication of when business-cycle turning points occur.[6]

Existing research has documented significant business-cycle effects on systematic forecast errors for the U.S. Using a modified Mincer-Zarnowitz regression, Sinclair et al. (2010) find that the Federal Reserve's Greenbook forecasts tend to overestimate the growth rate of real gross domestic product (GDP) during recessions and to underestimate it during expansions. These state-dependent systematic forecast errors offset each other when the analysis does not control for the state of the economy. Messina et al. (2014) confirm these results and show that the speed at which new information is incorporated into the forecasts does not change significantly during recessions. Loungani et al. (2013) adopt a broader perspective and analyze a large panel of growth forecasts for different countries and document that forecasters only slowly realize the occurrence of recessions leading to severe overestimation of growth rates prior to and during the early phases of recessions.[7]

There is broad consensus that professional forecasters typically target the first release of official GDP data and, hence, real-time data should be used when evaluating macroeconomic forecasts (Croushore and Stark, 2001; Croushore, 2011). This has not always been followed in studies that analyze international panels due to, presumable, the absence of 'ready-made' real-time data sets. Notable exceptions are Loungani (2001), who analyzes the accuracy of average growth forecasts in 63 economies, Isiklar and Lahiri (2007), who analyze for a panel of 18 advanced economies up to which forecast horizon survey-based forecasts provide valuable information, Batchelor (2007), who analyzes systematic bias in growth and inflation forecasts in the G7 economies,

---

[6]In particular, the quarterly business cycle dating scheme gives a better indication of whether the professional forecasters are aware that the state of the economy might have changed when they make their forecasts compared with an annual business-cycle scheme.

[7]In a related paper, Loungani (2001) provides some evidence that forecasts for recessions are subject to a large systematic forecast error. However, he does not formally test the significance of his findings and does not provide similar evidence for recovery and expansion years.

and Guisinger and Sinclair (2015), who analyze whether growth and unemployment forecasts reflect Okun's law in a panel of nine advanced economies.

Our main findings in this paper are as follows. First, the slightly negative forecast bias (across all countries and forecast horizons) masks tremendous differences in systematic forecast errors across business-cycle states, forecast horizons, and countries. We show that forecasts that are made for recession years are, on average, subject to a very large negative systematic forecast error (growth expectations are too optimistic), while forecasts that are made for recoveries, on average, underestimate the strength of the recovery. In contrast, forecast errors during expansions are, on average, not significantly different from zero for small horizons and are small in the case of longer-term forecasts. Second, forecasters anticipate business-cycle turning points rather late. Third, the adjustment of forecasts (i.e. the reduction of the systematic forecast error) around business-cycle turning points is a gradual process consistent with theories of imperfect information (Sims, 2003; Coibion and Gorodnichenko, 2012). Finally, we document a large heterogeneity of systematic forecast errors across countries and show that the positive relationship between a change in trend growth rates and forecast bias, as suggested by Batchelor (2007), cannot explain the cross-country variation of systematic forecast errors during expansions.

The remainder of this paper is structured as follows. Section 2 presents a stylized model of state-dependent expectation errors. Section 3 describes our data set and the approach we use to date recessions and recoveries. Section 4 describes the econometric approach we use. Section 5 presents a description of our empirical results. Finally, Section 6 presents our conclusions.

2 A Stylized Model of State-Dependent Expectation Errors

To organize thoughts about the state dependency of expectation errors it is useful to consider the following stylized model.[8] Assume that forecasters try to predict $y_t = x_t + z_t\, i_t$, where the components are given by

$$x_t = \mu + \rho x_{t-1} + \varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0, \sigma^2) \tag{1}$$

$$z_t \sim \mathcal{N}(\Delta\mu_r, \sigma_r^2) \tag{2}$$

$$i_t \sim \mathcal{B}(p_r) \tag{3}$$

---

[8]For expositional ease, the model is specified having only two different business-cycle states. However, it could easily be generalized to include more different states.

x_t is the process that governs growth dynamics in the first business-cycle regime. i_t follows a Bernoulli distribution and determines whether the economy is in the first business-cycle regime ($i_t = 0$) or in the second one ($i_t = 1$); the second regime occurs with probability p_r, where this probability is quite low.[9] z_t is a process that (with $\Delta\mu_r < 0$) determines how much lower growth is in the second regime. ε_t, z_t, and i_t are assumed to be independent of each other and across time. This model abstracts from most dynamic effects but captures three important features: First, growth is somewhat persistent and, hence, predictable. Second, the occurrence of recessions is unpredictable. Third, the depth of each recession is of stochastic nature and not predictable.

Now, at each point in time the optimal one-step ahead forecast under quadratic loss is given by the conditional expectation for y_{t+1}, $E[y_{t+1|t}] = E[x_{t+1|t}] + E[z_{t+1|t}]E[i_{t+1|t}] = \mu + \rho x_t + p_r\Delta\mu_r$. If forecasters followed this strategy, their forecasts would be unbiased on average, but conditional on the economy being in regime 1 or 2 we would observe systematic forecast errors of $-p_r\Delta\mu_r$ or $(1 - p_r)\Delta\mu_r$, respectively.

In contrast, forecasters could also try to forecast the economy conditional of being in the first (more likely) regime. Why would they want to do this? It might be that loss functions are such that issues other than the quadratic expectation error affect the forecaster's payoff (see, e. g., Ehrbeck and Waldmann, 1996; Laster et al., 1999). In our context, one hypothesis could be, for instance, that it is generally accepted that recessions are unpredictable and forecasters are not judged based on their recession forecasts but rather based on their performance during expansions.[10] In such an environment, it is optimal for a forecaster to form for ecasts based on the conditional expectation $E[y_{t+1,t}|i_{t+1} = 0] = \mu + \rho x_t$. It follows, that the systematic forecast error is equal to zero if $i_{t+1} = 0$ and equal to $\Delta\mu_r$ if $i_{t+1} = 1$. The unconditional bias is given by $p_r\Delta\mu_r$ in that case. This could explain findings of a negative bias of growth forecasts as reported for instance in Batchelor (2007) and Dovern and Weisser (2011). If indeed an average negative bias goes along with the absence of systematic forecast errors during business-cycle expansions will be one of the issues analyzed in the empirical part of this paper below.

---

[9]Looking at the U. S., e. g., the economy was only around 14 % of all quarters between 1950 and 2014 in recession.

[10]It is very long known that forecasters fail to predict recessions (Stekler, 1972). See also Hamilton (2011) for a recent account of the difficulties of forecasting recessions in real time.

3 Data and Business Cycle Dating

3.1 Forecast Data and Real-Time Realizations

We use a panel of forecasts for the growth rate of real GDP for 19 advanced economies over the period 1990 to 2013 provided by Consensus Economics. Each month, Consensus Economics asks a panel of professional forecasters to report their forecasts for the annual GDP growth rates of the current and next calendar year. We base our analysis on the average (so-called consensus) forecasts from each forecast period. Thus, the data set contains a sequence of 24 forecasts for each target year and country. These forecasts are made between January of the year prior to the target year and December of the target year. The data, therefore, have a three-dimensional panel structure. We use $F_{it,h}$ to denote a forecast for the growth rate in country i for a target year t made with a forecast horizon of h months. As an example, the consensus forecast in country i for the annual growth rate of 2014 from September 2013 is denoted by $F_{i2014,16}$.

We denote the corresponding actual GDP growth rate by A_{it}. We use real-time data to compute the actual growth rates. More precisely, we use the first available data vintage that allows the computation of the annual growth for a particular year. We follow this real-time approach, because there is a broad consensus in the literature that professional forecasters typically target the first data release when forming their macro forecasts.[11]

We use real-time data from three different sources. First, we use the *Original Release Data and Revisions Database* (ORDRD) provided by the *OECD*, which contains real-time vintages from the *Main Economic Indicators*, starting with the vintage from January 1999. Second, we extend this database back to 1990 with data from the *Real-Time Historical Database for the OECD* provided by the *Federal Reserve Bank of Dallas*.[12] This source is consistent with the ORDRD. Finally, we use data vintages from the *OECD Economic Outlook* to fill gaps of missing data that result from gaps in the first two data sets.[13]

Given these quantities, we compute a sequence of 24 forecast errors for each country and target year as $FE_{it,h} = A_{it} - F_{it,h}$. The sample covers data for Austria, Belgium, Canada, Switzerland, Germany, Denmark, Spain, Finland, France, Greece, Ireland, Italy, Japan, the Netherlands, Norway, Portugal, Sweden, the United Kingdom, and

---

[11]We demonstrate below how using current-vintage data can lead to misguided conclusions.

[12]For a detailed description of this source, see Fernandez et al. (2011).

[13]A detailed account of the real-time data sources is given in Appendix A.2.

the United States. The sample contains a total of 10,677 (consensus) forecasts made for 24 target years between 1990 and 2013 (Table 1). Due to data limitations, the sample for Greece starts with forecasts for the target year 1993.

[Table 1 about here.]

3.2 Business Cycle Dating

Due to the specific structure of the survey data, it is appropriate to use two different methods for dating recessions and recoveries. Thereby, we can account for the fact that forecasts of our sample are made for *annual* growth rates and simultaneously analyze how the average forecast errors evolve around business-cycle turning points identified based on *quarterly* data. Throughout the paper, we identify the different states of the business cycle (at annual and quarterly frequencies) based on the most recent data vintage. We believe this data status provides the most reliable proxy for the state of the business cycle at each point in time because it is based on a broader information set compared with earlier data vintages.

We use a very simple and transparent rule to identify recession years at an *annual* frequency: We identify a particular year as a recession year (in a particular country) if the annual growth rate is negative ($A_{it} < 0$). This rule closely follows Bry and Boschan (1971) and has recently been used, e.g., in Jordà et al. (2013). We use an indicator function $\Lambda_{it} = \mathcal{I}(A_{it} < 0)$ that is equal to 1 if year t is a recession year in country i and 0 otherwise. In addition, we use an indicator function $\Lambda_{it}^* = \mathcal{I}(A_{it} < 0, A_{it-1} \geq 0)$ that is equal to 1 if year t is the first year of a sequence of recession years in country i and 0 otherwise. Based on this classification, we identify a particular year as a recovery year if the annual growth rate is non-negative and the previous year is identified as a recession year. We use an indicator function $\Upsilon_{it} = \mathcal{I}(A_{it} \geq 0, A_{it-1} < 0)$ that is equal to 1 if year t is a recovery year in country i and 0 otherwise.

In addition, we make use of a quarterly business cycle dating scheme based on quarterly GDP data.[14] Here, we resort to the algorithm proposed by Harding and Pagan (2002), who adapt the method of Bry and Boschan (1971) for use with quarterly data. This algorithm identifies local peaks and troughs in real GDP. We use the convention that a recession starts in the first month of the quarter that follows on the peak quarter

---

[14]Given the monthly frequency of the survey data, a monthly dating of recessions would be ideal. GDP data are, however, available only at a quarterly frequency for most countries.

and that a recovery starts in the first month of the quarter that follows on the trough quarter.

In our empirical analysis, we will distinguish in a first step whether a forecast is made for a recession year, a recovery year, or an expansion year (defined as a year that is neither a recession nor a recovery year). In a second step, we also keep track of whether a recession (or recovery) had already begun when a forecast was made. To this end, we define variables $\lambda_{it,h}$ and $\lambda_{it,h}$ that are 1 if, according to the quarterly dating scheme, a recession or recovery, respectively, has already begun at the corresponding forecast period, and 0 otherwise. Note that once these variables switch to 1 for a particular target year, they remain at this value until the end of the target year, i.e. if $\lambda_{it',h'} = 1$ for a particular target year t' and forecast horizon h', then $\lambda_{it',h'-k} = 1$ for all $k = 1, \ldots, h' - 1$ (and, analogously, for $\upsilon_{it,h}$). In a third step, we keep track of how long before (or how long after) the start of a recession or recovery a forecast is made. Formally, we define $\Delta_{it,h}^{\lambda}$ to measure the distance (in months) of a forecast period to the start of a recession and $\Delta_{it,h}^{\upsilon}$ to measure the distance (in months) of a forecast period to the start of a recovery.

[Figure 1 about here.]

Figure 1 demonstrates our dating convention using two examples.[15] The left plot refers to the most recent recession in the US. The annual GDP was lower than in the previous year in 2008 and in 2009, and hence, we classify these two years as recession years. 2010 is the first year with a positive growth rate following the recession, and hence, we classify it as a recovery year. The analysis of the quarterly data suggest that the recession started in the first quarter of 2008 and, hence, we define January 2008 as the start of the recession. To take two examples, $\Delta_{US\,2008,16}^{\lambda} = -6$ because the forecast for 2008 made in July 2007 was made 6 months before the recession started, and $\Delta_{US\,2009,16}^{\lambda} = 6$ because the forecast for 2009 made in July 2008 was made 6 months after the start of the recession. The recession ended in the second quarter of 2009, and thus, we treat July 2009 as the first month of the recovery. The right plot refers to the recession of 2001/02 in Japan. These examples highlight that when there is a mild recession, which does not result in a negative annual growth rate, we do not classify any of the years as a recession or recovery year. In the given example, the mild recession

---

[15]The results of our quarterly business-cycle dating method are broadly in line with business-cycle datings provided by other sources, such as the *National Bureau of Economic Research* (NBER) or the *Economic Cycle Research Institute* (ECRI).

starts in the second quarter of 2001, and the recovery starts in the second quarter of 2002. Thus, $\Delta^\lambda_{JP\,2001,9} = \Delta^\lambda_{JP\,2002,21} = 0$ and $\Delta^v_{JP\,2002,9} = \Delta^v_{JP\,2003,21} = 0$.

Table 1 contains an account of the frequency of recessions and recoveries identified by our different identification methods. The number of recessions identified by the annual dating algorithm ranges from one in Austria, Ireland, and Norway to four in Italy and Portugal. Several recessions last more than one year, so the number of recession years is considerably higher than the number of recessions for many countries. In some cases, the number of recovery years is lower than the number of recession years because six economies were in a recession in 2013 when our sample ends. The number of recessions identified by the quarterly dating algorithm is typically larger and ranges from two in Canada, Ireland, the Netherlands, and the United States to six in Germany and Japan. In total, our sample covers 45 (74) recessions based on the annual (quarterly) dating scheme.

4 Econometric Methodology

We follow the standard approach to analyze systematic forecast errors in fixed-event forecasts by using (modified) Holden-Peel (1990) regressions. The basic version of the Holden-Peel regression is given by

$$FE_{it,h} = \alpha + \varepsilon_{it,h}. \tag{4}$$

We add several terms to this basic version to account for the state of the economy in the target year (target-year effects), for the state of the economy when the forecast was made (forecast-period effects), and for the forecast horizon (horizon effects). In addition, we add country fixed effects (α_i) to the model because we want to control for country-specific factors, which influence the average forecast error but are omitted in our specification. Batchelor (2007) and Dovern and Weisser (2011) argue, for instance, that changes in the trend growth rate of a particular country lead to systematic forecast errors because forecasters have to learn about the structural change. In the end, our

most general specification is:

$$
\begin{aligned}
FE_{it,h} =& \alpha + \alpha_i \underbrace{+\beta_1 \Lambda_{it} + \beta_2 \Lambda_{it}^* + \beta_3 \Upsilon_{it}}_{\text{target-year effects}} \\
& \underbrace{+\gamma_1 h + \gamma_2 h \times \Lambda_{it} + \gamma_3 h \times \Lambda_{it}^* + \gamma_4 h \times \Upsilon_{it} + \gamma_5 h \times \lambda_{it,h} + \gamma_6 h \times \upsilon_{it,h}}_{\text{horizon effects}} \\
& \underbrace{+\delta_1 \lambda_{it,h} + \delta_1 \upsilon_{it,h}}_{\text{forecast-period effects}} + \varepsilon_{it,h}.
\end{aligned}
\tag{5}
$$

In this specification, β_1, β_2, and β_3 measure how much information about forecast errors is contained in the business-cycle-regime dummies. As Sinclair et al. (2010) argue, if these coefficients are different from zero, forecasters do not systematically anticipate the state of the economy in their forecasts. γ_1, ..., and γ_6 measure how strongly the systematic forecast errors vary with the forecast horizon. γ_1 measures the baseline effect, while the other coefficients correspond to five terms that are constructed by interacting the forecast horizon and our business-cycle-dating variables. These coefficients measure how the 'horizon effect' changes if forecasts are made for a recession or recovery (γ_2, γ_3, γ_4) or when they are made around a business-cycle turning point (γ_5, γ_6). Finally, $\alpha + \alpha_i$ measures the country-specific systematic forecast error during 'normal' expansions after controlling for potential horizon effects. α is identified by imposing $\sum_{i=1}^{N} \alpha_i = 0$, and it measures the average systematic forecast error after controlling for other effects in the model.[16]

Testing for a systematic forecast error, or bias, in equation (4) boils down to a t-test of H_0: $\alpha = 0$ versus H_1: $\alpha \neq 0$. In the adjusted Holden-Peel regression (5), the matter is complicated by the fact that the model implies situation-specific systematic forecast errors that depend on the forecast horizon, the nature of the target year, and the business-cycle state that prevails during the forecast period. Two examples illustrate this point: Checking whether forecasters, on average, make systematic forecast errors when $h = 13$ and forecasts are made for an expansion year can be performed by using an F-test of H_0: $\alpha + 13\gamma_1 = 0$ versus H_1: $\alpha + 12\gamma_1 \neq 0$. Likewise, checking whether forecasters, on average, make systematic forecast errors when $h = 1$, forecasts are made

---

[16]Note that we use a standard fixed-effects model that allows only the constant to differ across countries, because the number of recessions and recoveries for each country is not large enough to sensibly estimate country-specific versions of the other parameters in the model. In other words, due to the limited amount of information for each country, we have to restrict the effects the business cycle has on systematic forecast errors to be equal across countries.

for a recession year and the recession has already begun can be performed by using an F-test of H_0: $\alpha + \beta_1 + \delta_1 + 1 \cdot (\gamma_1 + \gamma_2 + \gamma_5) = 0$ versus the corresponding alternative hypothesis.[17]

Because the time-varying forecast horizon is such a pronounced feature of our data set, we also look at the forecast bias as a function solely of the forecast horizon at some points below. A priori, it is sensible to allow for such differences because longer-term forecasts might, for instance, be biased due to misperceptions about trend growth rates, while short-term forecasts are more likely to be unbiased because much of the relevant information set is already revealed when forecasters make these forecasts. For this analysis, we do not impose any restriction on the relationship between the forecast horizon and forecast bias, but we regress the latter on a set of 24 dummies that correspond to the different forecast horizons

$$FE_{it,h} = \sum_{r=1}^{24} \alpha_r D_h^r, \tag{6}$$

where $D_h^r = 1$ if $h = r$ and zero otherwise. To take differences across business-cycle states into account, we estimate equation (6) for three different sub-samples that correspond to forecasts made for recession years, those made for recovery years, and those made for expansion years. However, due to the large number of regressors, it is not possible in this specification to add interaction terms involving the business-cycle state in the forecast period.

5 Empirical Results

5.1 Are Systematic Forecast Errors State-Dependent?

The evidence from Sinclair et al. (2010), Loungani et al. (2013), and Messina et al. (2014) suggest that forecasters make systematic errors when forecasts are made for recessions because they fail to anticipate the change in the business-cycle regime well in advance. In this section, we provide much broader international evidence regarding this question. In addition, we present analogous results with respect to forecasts that are made for recovery years.

---

[17]In Table 2 below, we present test statistics for a wider range of cases. It is straightforward to apply this idea of testing for situation-specific systematic forecast errors to any case.

Table 2 summarizes the results from different specifications of the panel version of the modified Holden-Peel regression. The first model corresponds to a (fixed-effects) panel version of equation (4).[18] The estimate for α indicates that forecasters, on average, overestimate future growth by 0.4 percentage points. The null hypothesis of unbiased forecasts is clearly rejected. This confirms earlier evidence in Batchelor (2007) and Dovern and Weisser (2011) that indicates that GDP forecasts often have a negative bias.

[Table 2 about here.]

However, how much heterogeneity across business-cycle states does this result mask? Models M2 to M4 sequentially add the target-year dummies Λ_{it}, Λ_{it}^*, and Υ_{it} to the specification. The first of these specifications closely corresponds to the regressions run by Sinclair et al. (2010) and suggest that forecasters do not systematically anticipate recession years. For these years the growth rate is overpredicted by a wide margin ($\hat{\beta}_1 = -2.1$). Second, it is interesting that this effect is not significantly smaller when forecasts are made for later years in a sequence of recession years (M3); this finding is in line with the evidence provided in Loungani (2002). Finally, forecasters seem, on average, to underestimate the strength of recoveries. Including Υ_{it}^* in M4 yields an estimate for β_3 of 0.6, which is highly significantly different from 0. Turning to the new estimates for α shows that the result of model M1 is mainly driven by recessions. Once we control for recessions, the (negative) estimate for α becomes much smaller in absolute values and is no longer significantly different from 0 (M2-M3). However, once we control for recoveries, it becomes obvious that this result masks a significant positive systematic forecast error in recoveries with a small but significant negative systematic forecast error in expansions ($\hat{\alpha} = -0.1$ in M4).

The evidence in Sinclair et al. (2010) suggests that forecasters are unable to systematically anticipate recessions but that they recognize business-cycle troughs once they occur. To see whether this also holds true in our data set and whether the same also applies to the anticipation of recoveries, we include information about whether a recession or recovery had already started when a forecast was made in M5. This information is captured by the variables $\lambda_{it,h}$ and $\upsilon_{it,h}$.[19] The estimate for δ_1 is equal to

---

[18]Note that we identify α by restricting the average country-fixed-effect to 0.

[19]Note that we cannot know exactly when forecasters became aware of business-cycle turning points due to different publication lags in different countries. Thus, δ_1 and δ_2 should, strictly speaking, only be interpreted as an approximation of the reduction of systematic error due to the gain of knowledge about the occurrence of a business-cycle turning point.

1.4, which indicates that the systematic (negative) forecast error for recession years is indeed smaller in situations in which a forecast for a recession year is made after the start of the recession. (Think, for instance, of a forecast made in October of the target year when the recession has started in the second quarter of that particular year.) At the same time, the estimate for β_1 declines to -3.4. Thus, the systematic bias before the start of a recession is even larger than the estimates for models M2-M4 suggest. The issue is different and somewhat puzzling on first sight in the case of recoveries. The estimate for β_2, the overall 'recovery effect' on systematic forecast errors is approximately cut in half in this specification, and $\hat{\delta}_2 = 0.4$ suggests that the excessive pessimism of forecasters with respect to the strength of the recovery becomes even larger (and significantly so) after the recovery has already begun. We can explain this result by the fact that some of the forecasts for recovery years (those with a high forecast horizon) are made when the economy is still in a recession. With the recession becoming longer, forecasters apparently become more pessimistic and tend to revise their forecasts downward. Consequently, the systematic forecast error for a recovery year can increase for some months even when the forecast horizon is becoming smaller. If the forecasters, in addition, become aware of the recovery rather late or if they revise their forecasts upward only very cautiously after the recovery has started, the systematic forecast error is, on average, higher after a recovery has started than before.[20]

Thus far, all specifications implicitly make the assumption that the systematic forecast errors are the same across all forecast horizons. This is clearly an unrealistic assumption, and we relax it for the last model (M6). To retain a certain level of parsimony, we impose that the systematic forecast error be linearly related to the forecast horizon. We also include various interaction terms of the forecast horizon with the different variables we use to measure the state of the business cycle. This allows us to investigate whether the 'horizon effect' differs across forecasts made for expansions, recession years, or recovery years.

Looking at the estimate for γ_1 in model M6 shows that the average systematic forecast error during expansions decreases by 0.015 percentage points with each increase of the forecast horizon by 1 month. Taken together with the constant term that is estimated to be 0.08, the systematic forecast error for expansions is given by $\hat{\alpha} + \hat{\gamma}_1 = 0.07$ for a forecast horizon of $h = 1$. This is not significantly different from 0 according to the test result, which can be found in the lower part of Table 2 (H_0^F). For $h = 13$,

---

[20]We return to this issue in section 5.2 below.

13

the systematic error is $(\hat{\alpha} + 13 \cdot \hat{\gamma}_1 = -0.1)$, and for $h = 24$, it reaches its largest value in absolute terms $(\hat{\alpha} + 24 \cdot \hat{\gamma}_1 = -0.3)$. F-tests indicate that the systematic forecast error is not significantly different from zero (at the 5% level) for forecast horizons from 1 to 12 months. For $h > 12$, professional forecasters, on average, significantly overestimate GDP growth during expansions.

Concerning the 'horizon' effects for forecasts made for either recession or recovery years, we obtain the following results. First, the systematic (negative) forecast error corresponding to forecasts for recession years is strongly decreasing with a shrinking forecast horizon $(\hat{\gamma}_1 + \hat{\gamma}_2 = -0.15)$. Second, this effect is even more pronounced once a recession has started already: the parameter estimate for $\hat{\gamma}_5 = -0.071$ is significantly different from zero. Thus, forecasters reduce the systematic error of their forecasts faster once the recession has actually started. Finally, there is no evidence of an additional 'horizon' effect corresponding to forecasts made for recovery years. Both $\hat{\gamma}_4$ and $\hat{\gamma}_6$ are small and not significantly different from zero. In addition, the results for M6 reveal that the coefficients on the various business-cycle dummies are no longer significantly different from zero (or only marginally so) after we include the different horizon terms in the model. This suggests that systematic forecast errors are gradually changing with the horizon rather than being subject to abrupt changes.

Although the results for the 'horizon' effects overall seem to be reasonable, they are based on the restrictive assumption of linear 'horizon' effects. Using the regression framework described in equation (6), we now analyze whether this assumption is acceptable or too restrictive. Figure 2 shows horizon-specific estimates for the systematic forecast errors corresponding to recession years, recovery years, and expansions. The results suggest that the relationship between the forecast horizon and the systematic forecast error shows some signs of non-linearities. For recession forecasts, the systematic error decreases monotonically with a shrinking forecast horizon. The reduction, however, is not linear but is more pronounced between $h = 18$ and $h = 6$, i.e. around the turn of the year. For recovery forecasts, the systematic forecast error initially increases between $h = 24$ and $h = 18$ and is then gradually lowered toward the end of the target year. (This pattern also partly explains why δ_2 is estimated to be positive in M5 and M6 above.) For expansion forecasts, the systematic forecast error is small for all forecast horizons. There seems to be a break around $h = 15$, and forecasts with a lower forecast horizon do not exhibit a significant systematic forecast error on average. Overall, the non-linearities are not too pronounced, and using a linear approximation in the panel model seems reasonable.

[Figure 2 about here.]

The results show that the expansion forecasts exhibit no systematic forecast errors for most of the forecast horizons available in our sample and that the recovery forecasts do not fail systematically at small forecast horizons.[21] Only forecasts made for recession years exhibit systematic error, albeit small, even in December of the target year. Hence, for long-term forecasts the existence of an average forecast bias cannot be explained by a desire of forecasters to produce predictions that have no systematic forecast error during expansions (see the stylized model presented in section 2). In contrast, the state-dependent systematic errors observed for smaller forecast horizons can support such behavior of forecasters.

If we had used current-vintage data and/or had neglected differences across business-cycle states, we would have found systematic forecast errors for almost all forecast horizons. Appendix B contains a table and a figure analogous to those of this section but based on current-vintage data. These results clearly show that using current-vintage data leads to an overestimation of systematic forecast errors—especially for expansion forecasts.

5.2 How Early Are Recessions and Recoveries Anticipated?

Thus far, we have analyzed the systematic forecast errors in terms of the characteristics of the target year and the forecast horizon. We have made use of the quarterly (and more precise) business-cycle dating only via $\lambda_{it,h}$ and $\upsilon_{it,h}$, but we have not yet analyzed how the pattern of systematic forecast errors evolves around the start of recessions or recoveries. We now turn to this issue.

We analyze how the systematic forecast errors change during the months before and after recessions or the start of recoveries. To this end, we make use of our quarterly business-cycle dating and define, based on $\Delta_{it,h}^{\lambda}$ and $\Delta_{it,h}^{\upsilon}$, dummy variables for each month before and after the start of recessions and recoveries, respectively. For example, when looking at the recession case, the dummy D_{-2}^{Λ}, which captures the typical systematic forecast error made two months before the start of a recession, is equal to 1 for all observations that correspond to a forecast made two months before a recession starts in a particular country and 0 otherwise. We regress the forecast error on a set

---

[21]In contrast, average forecasts, i.e. without controlling for the state of the business cycle, are significantly biased for all but the very low forecast horizons. The point estimates for these biases are also much larger for all forecast horizons compared to the size of the systematic forecast errors during expansion years.

of dummy variables of this type separately for recessions and recoveries.[22] In these regressions, we differentiate between forecasts that are made for recession years (or for recoveries following recession years) and those forecasts that are not made for recession or recovery years. The latter can be interpreted as forecasts that are made before, during, and after mild recessions.[23]

By using the dummy-variable approach, we allow the systematic forecast error to be a very flexible function of the temporal distance to the two types of business-cycle turning points. This is appropriate because we do not have any information a priori about when exactly forecasters learn about business-cycle turning points and adjust their forecasts accordingly. On the one hand, forecasters might notice the start of recessions and recoveries only some months after the event because publication lags and data revisions make identification in real time very difficult. On the other hand, survey data or financial market data that are available much earlier and that are not subject to revisions are frequently used to assess the state of the business cycle in real time. Overall, we expect that forecasters become aware of business-cycle turning points a few months after the event, i.e. at $\Delta_{it,h}^{\lambda} > 0$ or $\Delta_{it,h}^{v} > 0$. In the end, our empirical findings can shed light on the question of how long it takes until the occurrence of a recession/recovery is reflected in the average forecast.

We make pragmatic choices about the upper and lower bounds of the window around the business-cycle turning points we look at. We trace forecasts up to 1 year after the start of a recession or recovery, i.e. $\Delta_{it,h}^{\lambda,\max} = \Delta_{it,h}^{v,\max} = 12$. While it would also be interesting to trace the evolution of systematic forecast errors during an entire year before the start of recessions or recoveries, we face data limitations that force us to adopt an asymmetric window. Figure 3 shows the number of available observations (of forecasts made for recession years or recovery years) for each value of $\Delta_{it,h}^{\lambda}$ and $\Delta_{it,h}^{v}$, respectively. The number of observations is sufficiently large and does not change substantially for the time after the turning points. However, the farther away we move before the turning points, the lower the number of observations, becoming very small for low values of $\Delta_{it,h}^{\lambda}$ and $\Delta_{it,h}^{v}$. Consequently, comparisons of results across different estimates become increasingly dominated by the changing composition of the sample ('composition effect'). Therefore, we only look at the 6 months prior to all business-cycle turning points of

---

[22]As a robustness check, we include the forecast horizon h in these regressions. The results are very similar to those presented below and are available upon request.

[23]In these cases, the recession identified by the quarterly business-cycle dating approach does not cause a negative annual growth rate of GDP.

our sample, i.e. $\Delta_{it,h}^{\lambda,\min} = \Delta_{it,h}^{v,\min} = -6$. Thus, the sets of dummy variables we use are given by $D_{-6}^{\Lambda}, D_{-5}^{\Lambda}, \ldots, D_0^{\Lambda}, \ldots, D_{12}^{\Lambda}$ and by $D_{-6}^{\Upsilon}, D_{-5}^{\Upsilon}, \ldots, D_0^{\Upsilon}, \ldots, D_{12}^{\Upsilon}$, respectively.

[Figure 3 about here.]

Figure 4 shows the estimated systematic forecast errors as a function of the temporal distance to the business-cycle turning points. In other words, it shows how the systematic forecast error evolves around the start of recessions and recoveries. Looking at the recoveries, the systematic forecast error corresponding to forecasts for recovery years, i.e. recoveries from strong recessions, increases in the run-up to the recovery, i.e. during the recession. In other words, during the recession, forecasters seem to become more pessimistic about the strength of the subsequent recovery. Interestingly, the positive systematic forecast error is not significantly reduced during the year following the start of the recovery. Twelve months after the start of a recovery, forecasters typically still underestimate the strength of the economic catch-up. For recoveries from mild recessions, forecasters, on average, significantly overestimate the strength of the recovery 6 months before the start of the recovery. However, this systematic forecast error is quickly reduced and becomes insignificant for $\Delta_{it,h}^{v} \geq 3$.

[Figure 4 about here.]

Turning to the dynamics of systematic forecast errors around the start of recessions, we observe the following features. First, forecasters initially overestimate growth in the case of mild recessions (by approximately 1 percentage point). This systematic error is very persistent and starts to be corrected approximately 4 months after the start of the recession. Approximately 9 months after the recession start, there is no longer a significant systematic forecast error. Second, the systematic overestimation of growth rates is far more severe in those cases we define as recession years. Initially, forecasters overestimate growth by 3–4 percentage points; and even 1 year after the start of the recession, growth is significantly overestimated (by more than 1 percentage point). The pattern over time, however, is similar to the case of mild recessions: the size of the systematic error starts to be reduced approximately 3 months after the recession start while being fairly constant during the months before the start of the recession (if the wide confidence bands are taken into consideration). The largest reductions in the systematic forecast error can be observed 6 to 12 months after the recession starts. Finally, and related to the last point, we observe a puzzling increase in the systematic

17

forecast error (in terms of point estimates) during the 6 months prior to the start of recessions (from -3 to -4 percentage points).

Taking a closer look at the data reveals that the latter feature is an artifact of the 'composition effect' mentioned above. Figure 5 compares the point estimates from the previous figure with the averages of *synthetic* forecast errors that we obtain in three steps. First, we compute the change of the forecast error from one forecast period to the next for all observations. Second, we estimate the average of these changes for each value of $\Delta_{it,h}^{\lambda}$ and $\Delta_{it,h}^{v}$. Finally, we normalize the level of the average synthetic forecast error by setting it equal to the original estimate for the period of the business-cycle turning points, and we use the estimated average changes to compute a synthetic level for other values of $\Delta_{it,h}^{\lambda}$ and $\Delta_{it,h}^{v}$. By using the change in forecast errors to compute the synthetic forecast error, we control, to some extent, for jumps that are due to the 'composition effect'. The new proxy is more in line with our prior in case of strong recessions: the smaller the $\Delta_{it,h}^{\lambda}$, the larger the systematic forecast error. Interestingly, forecasters seem to anticipate lower growth well before the start of the recession, even though they do not reduce their forecasts nearly sufficiently.[24]

[Figure 5 about here.]

5.3   What Explains Cross-Country Differences in Systematic Errors?

Up to this point, we have focused on average results and have not discussed important differences across countries, which are surely masked (as in any cross-country analysis) by the average results. The relevant data are displayed in Figure 6. It shows for each country the average forecast errors made for expansion, recession, and recovery years as a function of the forecast horizon.[25] It is evident that forecasts are not subject to systematic errors for expansions for small forecast horizons in most countries. For some countries, even longer-term forecast errors are not significantly different from 0. For other countries, negative (e. g. Spain) or positive (e. g. Greece or Ireland) systematic forecast errors are sizable.

---

[24]Note that for all other cases, there are no major differences between the original estimates and the synthetic proxy, suggesting that the 'composition effect' does not play a substantial role in those cases.

[25]Note that we do not consider the country-specific estimates for systematic forecast errors for recession and recovery years to be reliable because they are based on very few observations (as indicated by the very wide standard errors). We still show them as giving a tentative impression of the performance of forecasters during recessions and recoveries in each of the countries.

[Figure 6 about here.]

To analyze the cross-country differences more formally, a look at the size of the country-specific fixed effects of the panel versions of the Holden-Peel regressions in section 5.1 is a good starting point and reveals that the systematic forecast errors are indeed very different across countries. In model M1, not controlling for any business-cycle effects, the systematic forecast errors range from −0.7 percentage points in Italy to 0.3 percentage points in Ireland. Overall, the standard deviation of fixed effects is equal to 0.24. The bias is significantly different from 0 for 18 countries (out of 19).

Because we showed above that systematic forecast errors are most pronounced before and during recessions, one hypothesis is that it might be possible to explain differences across countries by their different business-cycle experiences over the sample period. One would expect, for instance, that forecasters tend to overestimate GDP growth more in those countries that experienced more and/or longer recessions during the period covered by our sample. However, our results reveal that this does not seem to be the primary reason behind the differences across countries. First, we also observe a high variation with respect to the fixed effects in specification M5. The standard deviation of the fixed effects is 0.21—not much smaller than for M1—and the implied systematic forecast errors for expansions range from −0.3 (Italy) to 0.6 (Ireland). Here, the systematic forecast error during expansions is significantly different from 0 for 14 countries—significantly less than when not controlling for business-cycle states. Second, the correlation between the systematic forecast errors for expansions implied by M5 and the country-specific biases implied by M1 is very high (Figure 7). Including the observations for Ireland, the correlation between the two sets of estimates is 0.87 (and highly significant); excluding the outlier, the correlation is 0.59 (p-value: 0.054). Because we control for business-cycle effects in M5, the different business-cycle experiences over our sample (e. g. differences in the number of recessions that occurred between 1990 and 2013) explains only a minor part of the differences in the systematic forecast errors across countries.

[Figure 7 about here.]

Batchelor (2007) claims that one factor behind systematic forecast errors is changing trend growth rates because forecasters have to learn about these structural changes. This argument is mostly valid for forecasts made with a large forecast horizon, as little actual information about the outcome is revealed by the time forecasters make these

forecasts. Therefore, we concentrate on forecasts with $h \geq 13$ to determine whether we find this effect for our sample.

As a rough measure for the change of trend growth rates over our sample, we use the difference between the average growth rate of potential output between 2009 and 2013 and that between 1985 and 1989. We take estimates for these growth rates from the OECD (see Appendix A.3).

Figure 8 shows the relation of the changes of the growth rate of potential output over our sample to the country-specific systematic forecast errors based on specification M1 and the country-specific systematic forecast errors for expansions based on M5.[26] The blue line indicates the slope of a regression of the growth rate change on the overall country-specific bias of longer-term forecasts. This slope is equal to 0.17 and is significantly different from 0 at a 5% confidence level. This confirms the results in Batchelor (2007) based on a much larger data sample.

In contrast, the red line, which represents the slope of a regression of the growth rate change on the country-specific systematic (longer-term) forecast errors during expansions, has a much lower slope (0.08) that is not significantly different from 0. In other words, the effect identified by Batchelor does not seem to be independent of the business-cycle state. If anything, it is confined to the forecasts for recessions and/or recoveries. This, however, is puzzling because the argument involves a misjudgment of the *trend growth rate*, which should be independent of the state of the business cycle and should, hence, lead to systematic forecast errors during all states of the business cycle. In summary, neither the country-specific business-cycle experience nor differences in the change of trend growth rates seem to be convincing determinants of the variation (across countries) of the overall forecast bias or the systematic forecast error in expansions, respectively.

[Figure 8 about here.]

6   Concluding Remarks

In this paper, we use average GDP growth forecasts for 19 advanced economies for the years 1990 to 2013 to analyze whether and how systematic forecast errors depend on the state of the business cycle.

---

[26]Note that both M1 and M5 have been estimated based only on those observations with $h \geq 13$ for this exercise.

We find that the negative systematic forecast error (across all countries), which is obtained when different states of the business-cycle are not taken explicitly into account, masks enormous differences across business-cycle states. We show that forecasts that are made for recession years are subject to a very large negative systematic forecast error. Interestingly, the magnitude of this systematic error does not depend on whether the forecast is made for the first year in a sequence of recession years or later recession years. The systematic forecast error becomes significantly smaller after the start of a recession. This, however, is a gradual process (which slowly starts some months before a recession starts) rather than a shift in expectations at a specific point in time. The largest revisions are made, on average, 6 to 12 months after the recession starts. Our results show that the negative systematic forecast error is still approximately 1.5 percentage points even 12 months after the start of a recession. In the case of mild recessions (recessions identified based on a quarterly identification scheme that do not lead to a decline in GDP on an annual frequency), forecasters also tend to lower their forecasts even before a recession starts and to gradually adapt to the recession regime during the following months. The forecasts do not exhibit a significant systematic forecast error approximately 9 months after the start of such mild recessions.

With regard to recoveries, we find that forecasters make significant positive systematic forecast errors. These errors, however, are considerably smaller in absolute magnitude than the errors made for recession years. It appears that forecasters have more problems adjusting their forecasts properly before and after the start of recoveries relative to the observed behavior around recession starts: on average, the systematic forecast error becomes even larger after the start of a recovery.[27] Forecasts for recoveries following mild recessions do not exhibit significant systematic forecast errors.

When controlling for these particularities of recession and recovery forecasts, we find that forecasts for expansions basically do not exhibit systematic forecast errors. We only find evidence for very small systematic forecast errors for forecast horizons beyond 14 months. Apparently forecasters focus on providing predictions that do not have a systematic forecast error during expansions. It is worth noting that this result crucially depends on the use of real-time data to compute the forecast errors.

In the last part of the paper, we document that systematic forecast errors are noticeably different across countries. We show that this variation can be explained only to

---

[27]This is due to two issues. First, professional forecasters tend to become more pessimistic about the growth outlook for recoveries during the recessions that precede the recoveries. Second, forecasters do not revise their forecasts much once a recovery has started.

some extent by the different business-cycle experiences of the countries over our sample. Supporting the evidence in Batchelor (2007), we also show that there is a positive relationship between the change of the trend growth rate over the sample period and the country-specific bias of longer-term forecasts. We find no such relationship, however, when correlating the change of the trend growth rate with the country-specific systematic forecast errors during expansions, which raises doubts about the causal mechanisms in Bachelor's theory. Likewise, also the business-cycle experience of a country is not strongly related to the observed forecast bias for this country. We leave the question of what drives cross-country variations in forecast bias and state-specific systematic forecast errors for future research.

In summary, we find convincing and broad-based evidence that professional forecasts for expansions are not subject to any sizable systematic forecast errors, while average forecast errors for recessions and recoveries are large and significantly different from 0. Once again, the use of real-time data vintages has proven to be critical for obtaining an accurate assessment of forecast errors. Moreover, our results indicate that forecasters anticipate the state of the economy rather late when business-cycle turning points occur. This confirms evidence confined to the U.S. that has been presented by Sinclair et al. (2010) and Messina et al. (2014). The rather gradual adjustment of forecasts around business-cycle turning points that we document is in line with theories of imperfect information (e. g. Sims, 2003), which imply a sluggish adjustment of forecasts (Coibion and Gorodnichenko, 2012; Dovern et al., 2015).

References

Ager, Philipp, Marcus Kappler, and Steffen Osterloh, "The accuracy and efficiency of the Consensus Forecasts: A further application and extension of the pooled approach," *International Journal of Forecasting*, 2009, *25* (1), 167–181.

Ashiya, Masahiro, "Strategic bias and professional affiliations of macroeconomic forecasters," *Journal of Forecasting*, 2009, *28* (2), 120–130.

Batchelor, Roy, "Bias in macroeconomic forecasts," *International Journal of Forecasting*, 2007, *23* (2), 189–203.

Boysen-Hogrefe, Jens, Nils Jannsen, and Carsten-Patrick Meier, "A Note on Banking and Housing Crises and the Strength of Recoveries," *Macroeconomic Dynamics*, 2015, *forthcoming.*

Bry, Gerhard and Charlotte Boschan, *Cyclical Analysis of Time Series: Selected Procedures and Computer Programs* NBER Books, National Bureau of Economic Research, Inc, 1971.

Coibion, Olivier and Yuriy Gorodnichenko, "What Can Survey Forecasts Tell Us about Information Rigidities?," *Journal of Political Economy*, 2012, *120* (1), 116 – 159.

Croushore, Dean, "Frontiers of Real-Time Data Analysis," *Journal of Economic Literature*, 2011, *49* (1), 72–100.

__ **and Tom Stark**, "A real-time data set for macroeconomists," *Journal of Econometrics*, 2001, *105* (1), 111–130.

Deschamps, Bruno and Christos Ioannidis, "Can rational stubbornness explain forecast biases?," *Journal of Economic Behavior & Organization*, 2013, *92* (C), 141–151.

Dovern, Jonas and Johannes Weisser, "Accuracy, unbiasedness and efficiency of professional macroeconomic forecasts: An empirical comparison for the G7," *International Journal of Forecasting*, 2011, *27* (2), 452–465.

__ , **Ulrich Fritsche, Prakash Loungani, and Natalia T. Tamirisa**, "Information Rigidities: Comparing Average and Individual Forecasts for a Large International Panel," *International Journal of Forecasting*, 2015, *31* (1), 144–154.

Ehrbeck, Tilman and Robert Waldmann, "Why are Professional Forecasters Biased? Agency Versus Behavioral Explanations," *The Quarterly Journal of Economics*, 1996, *111* (1), 21–40.

European Central Bank, "Predicting the Strength of Recoveries," Monethly Bulletin 07-2014, July 2014.

Fernandez, Adriana Z., Evan F. Koenig, and Alex Nikolsko-Rzhevskyy, "A real-time historical database for the OECD," Globalization and Monetary Policy Institute Working Paper 96, Federal Reserve Bank of Dallas 2011.

Fildes, Robert and Herman O. Stekler, "The state of macroeconomic forecasting," *Journal of Macroeconomics*, 2002, *24* (4), 435–468.

Guisinger, Amy Y. and Tara M. Sinclair, "Okuns Law in real time," *International Journal of Forecasting*, 2015, *31* (1), 185 – 187.

Hamilton, James D., "Calling recessions in real time," *International Journal of Forecasting*, 2011, *27* (4), 1006 – 1026.

Harding, Don and Adrian Pagan, "Dissecting the cycle: a methodological investigation," *Journal of Monetary Economics*, 2002, *49* (2), 365–381.

Holden, Ken and David A. Peel, "On Testing for Unbiasedness and Efficiency of Forecasts," *The Manchester School of Economic & Social Studies, University of Manchester*, 1990, *58* (2), 120–27.

Isiklar, Gultekin and Kajal Lahiri, "How far ahead can we forecast? Evidence from cross-country surveys," *International Journal of Forecasting*, 2007, *23* (2), 167–187.

Jordà, Òscar, Moritz Schularick, and Alan M. Taylor, "When Credit Bites Back," *Journal of Money, Credit and Banking*, 2013, *45* (s2), 3–28.

Kim, Chang-Jin, James Morley, and Jeremy Piger, "Nonlinearity and the permanent effects of recessions," *Journal of Applied Econometrics*, 2005, *20* (2), 291–309.

Laster, David, Paul Bennett, and In Sun Geoum, "Rational Bias in Macroeconomic Forecasts," *The Quarterly Journal of Economics*, 1999, *114* (1), 293–318.

Loungani, Prakash, "How accurate are private sector forecasts? Cross-country evidence from consensus forecasts of output growth," *International Journal of Forecasting*, 2001, *17* (3), 419–432.

— , "There Will Be Growth in the Spring," *World Economics*, 2002, *3* (1), 1–6.

— **, Herman O. Stekler, and Natalia Tamirisa**, "Information rigidity in growth forecasts: Some cross-country evidence," *International Journal of Forecasting*, 2013, *29* (4), 605–621.

Messina, Jeff, Tara M. Sinclair, and Herman O. Stekler, "What Can We Learn From Revisions to the Greenbook Forecasts?," Working Papers 2014-14, The George Washington University, Institute for International Economic Policy 2014.

Mincer, Jacob A. and Victor Zarnowitz, "The Evaluation of Economic Forecasts," in "Economic Forecasts and Expectations: Analysis of Forecasting Behavior and Performance" NBER Chapters, National Bureau of Economic Research, Inc, August 1969, pp. 1–46.

Sims, Christopher, "Implications of Rational Inattention," *Journal of Monetary Economics*, 2003, *50* (3), 665–690.

Sinclair, Tara M., Herman O. Stekler, and Fred Joutz, "Can the Fed Predict the State of the Economy?," *Economics Letters*, 2010, *108*, 28–32.

Stekler, Herman O., "An Analysis of Turning Point Forecasts," *American Economic Review*, 1972, *62* (4), 724–729.

Wieland, Volker and Maik Wolters, "Forecasting and Policy Making," in "Handbook of Economic Forecasting, Volume 2," Elsevier, August 2013, pp. 239–325.

Zarnowitz, Victor and Phillip Braun, "Twenty-two Years of the NBER-ASA Quarterly Economic Outlook Surveys: Aspects and Comparisons of Forecasting Performance," in "Business Cycles, Indicators and Forecasting" NBER Chapters, National Bureau of Economic Research, Inc, August 1993, pp. 11–94.

Table 1: Descriptive Statistics

	Forecasts					Business Cycle Dating			
	First TY	#TY	#FC	Mean	SD	#R	#RY	#U	#HPR
AT	1990	24	566	2.0	1.1	1	1	1	4
BE	1990	24	566	1.8	1.1	3	3	3	4
CA	1990	24	567	2.5	1.2	2	2	2	2
CH	1990	24	566	1.6	0.8	2	4	2	5
DE	1990	24	567	1.6	1.3	3	3	3	6
DK	1990	24	566	1.9	1.0	3	4	3	5
ES	1990	24	566	2.1	1.7	3	5	2	3
FI	1990	24	566	2.2	1.7	3	6	2	4
FR	1990	24	567	1.9	1.1	2	3	2	3
GR	1993	21	482	1.8	2.6	2	7	1	5
IE	1990	24	566	3.5	2.7	1	3	1	2
IT	1990	24	567	1.4	1.3	4	6	3	5
JP	1990	24	567	1.5	1.7	3	5	3	6
NL	1990	24	566	1.8	1.4	2	3	1	2
NO	1990	24	566	2.5	1.0	1	1	1	5
PT	1990	24	566	1.6	1.8	4	7	3	5
SE	1990	24	566	2.0	1.4	2	5	2	3
UK	1990	24	567	1.9	1.3	2	3	2	3
US	1990	24	567	2.5	1.2	2	3	2	2
Total	-	-	10677	2.0	1.6	45	74	39	74

Notes: *First TY* refers to the first target year for which forecasts are available; *#TY* refers to the number of target years; *#FC* refers to the total number of available forecasts; *Mean* and *SD* refer to the mean and standard deviation of the sample of consensus forecasts; *#R* refers to the number of distinct recessions identified by the annual dating scheme; *#RY* refers to the total number of recession years; *#U* refers to the number of recoveries corresponding to the annual dating scheme; and *#HPR* refers to the number of recessions identified based on quarterly data using the algorithm by Harding and Pagan (2002).

Table 2: Modified Holden-Peel Panel Regressions (Real-Time Data)

Regression Results	M1	M2	M3	M4	M5	M6
β_1		-2.108^{***}	-2.087^{***}	-2.022^{***}	-3.354^{***}	0.008
		(-7.86)	(-5.56)	(-5.31)	(-8.68)	(0.01)
β_2			-0.034	-0.026	0.361	-0.257
			(-0.10)	(-0.08)	(0.97)	(-0.76)
β_3				0.623^{***}	0.322	-0.637^*
				(3.43)	(1.47)	(-2.03)
δ_1					1.375^{***}	0.553
					(4.85)	(0.96)
δ_2					0.418^{**}	0.646^*
					(2.26)	(2.02)
γ_1						-0.015^{***}
						(-3.53)
γ_2						-0.139^{***}
						(-4.30)
γ_3						0.000
						(0.01)
γ_4						0.056^{***}
						(2.90)
γ_5						-0.071^{**}
						(-2.49)
γ_6						0.015
						(0.93)
α	-0.380^{***}	-0.033	-0.033	-0.098^{**}	-0.098^{**}	0.081^*
	(-25.62)	(-0.74)	(-0.74)	(-2.15)	(-2.17)	(1.92)
Obs.	10677	10677	10677	10677	10677	10677
R^2	0.00	0.25	0.25	0.26	0.28	0.40

Test for Systematic Errors (F-tests)

Average errors (no horizon control)

$H_0^A: \alpha = 0$	656.20^{***}	0.55	0.54	4.62^{**}	4.70^{**}	
$H_0^B: \alpha + \beta_1 = 0$		91.30^{***}	39.28^{***}	37.98^{***}	89.37^{***}	
$H_0^C: \alpha + \beta_1 + \delta_1 = 0$					35.35^{***}	
$H_0^D: \alpha + \beta_3 = 0$				8.84^{***}	1.17	
$H_0^E: \alpha + \beta_3 + \delta_2 = 0$					12.18^{***}	

Horizon-specific errors ($h = 1$)

$H_0^F: \alpha + 1 \cdot \gamma_1 = 0$						2.76
$H_0^G: \alpha + \beta_1 + \delta_1 + 1 \cdot \gamma_1 + 1 \cdot \gamma_2 + 1 \cdot \gamma_5 = 0$						0.04
$H_0^H: \alpha + \beta_3 + \delta_2 + 1 \cdot \gamma_1 + 1 \cdot \gamma_4 + 1 \cdot \gamma_6 = 0$						2.90

Horizon-specific errors ($h = 13$)

$H_0^I: \alpha + 13 \cdot \gamma_1 = 0$						5.25^{**}
$H_0^J: \alpha + \beta_1 + 13 \cdot \gamma_1 + 13 \cdot \gamma_2 = 0$						10.67^{***}
$H_0^K: \alpha + \beta_3 + 13 \cdot \gamma_1 + 13 \cdot \gamma_4 = 0$						0.02

Horizon-specific errors ($h = 24$)

|						
$H_0^L: \alpha + 24 \cdot \gamma_1 = 0$						10.13^{***}

Notes: The estimated equation is given by (5). The constant α is identified by requiring that the sum of the country-fixed effects (not shown in the table) is equal to 0. The numbers in parentheses are t-statistics based on Huber-White robust standard errors. The numbers in the lower panel of the table are F-statistics corresponding to the different hypotheses.
* denotes significance at a 10% confidence level.
** denotes significance at a 5% confidence level.
*** denotes significance at a 1% confidence level.

Figure 1: Recession and Recovery Dating Scheme
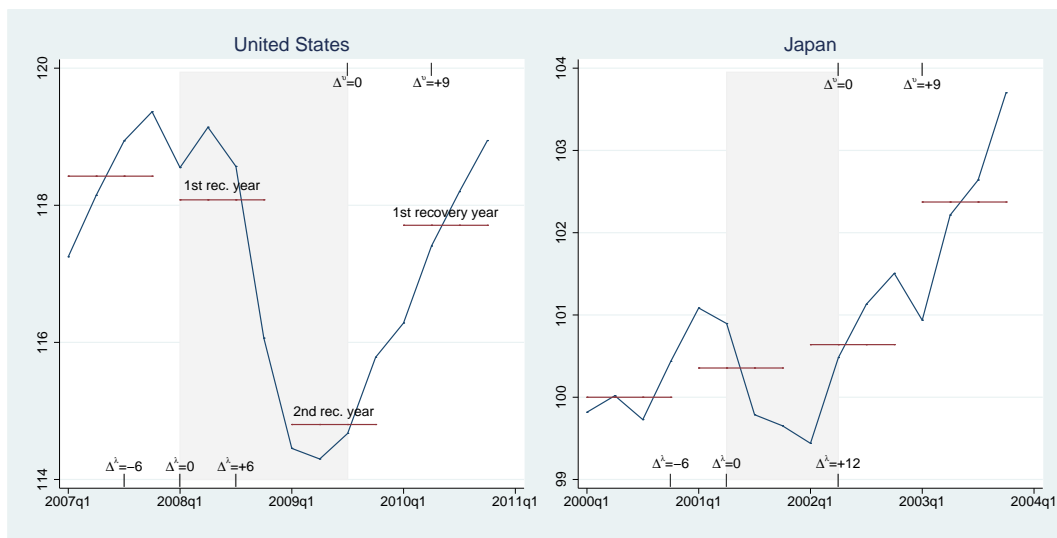
Notes: Quarterly data on the level of real GDP (blue line, index: 2000=100). The red lines indicate the annual average level of GDP for each year. The shaded areas refer to recessions as identified by the quarterly business-cycle dating scheme. The variables Δ^λ and Δ^υ are defined in the text and indicate the distance to the start of recessions or recoveries, respectively.
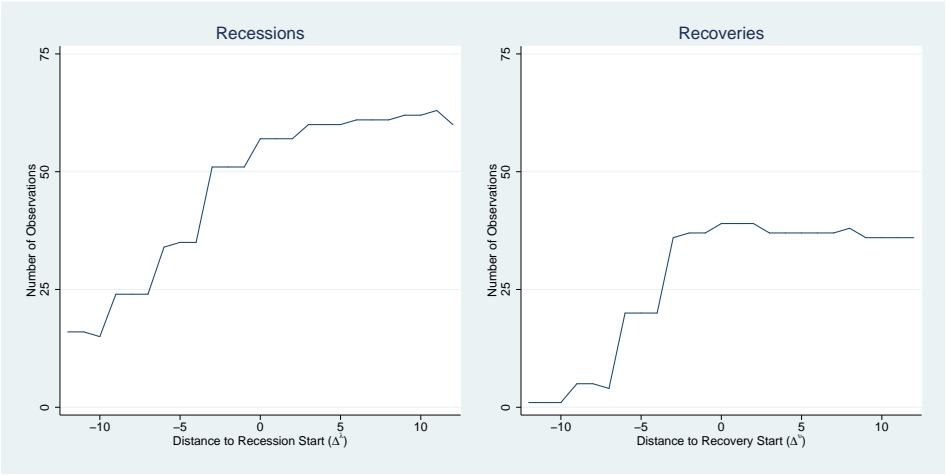
Figure 2: Systematic Forecast Errors by Horizon
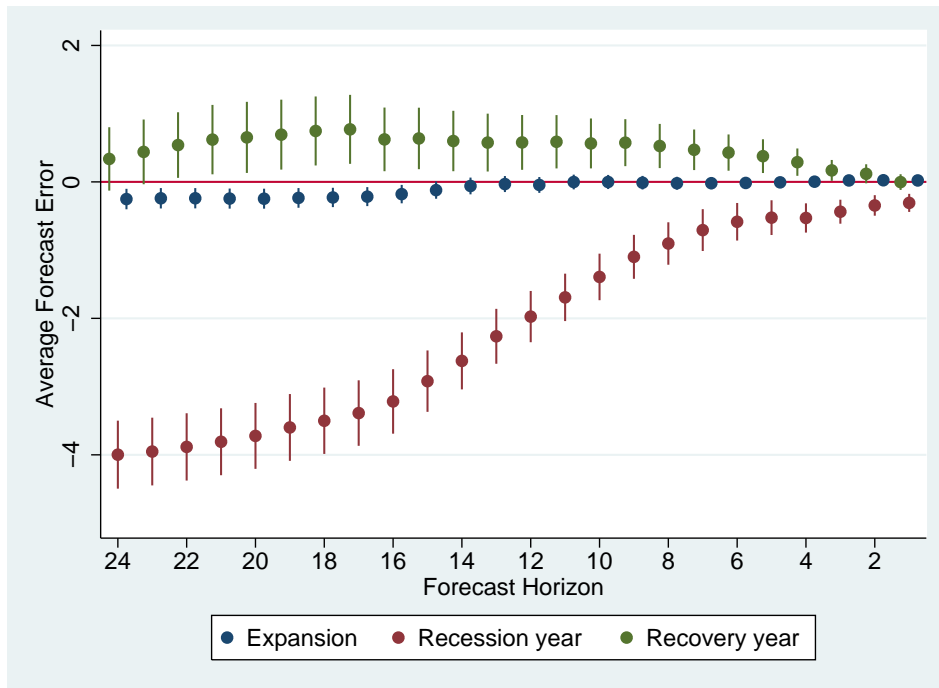
Notes: This figure shows estimates for the systematic forecast errors (in percentage points) as a function of the forecast horizon. The dots represent point estimates from a regression of the forecast error on a set of 24 dummy variables (one for each of the different forecast horizons). The vertical lines indicate 95% confidence bands surrounding the point estimates.
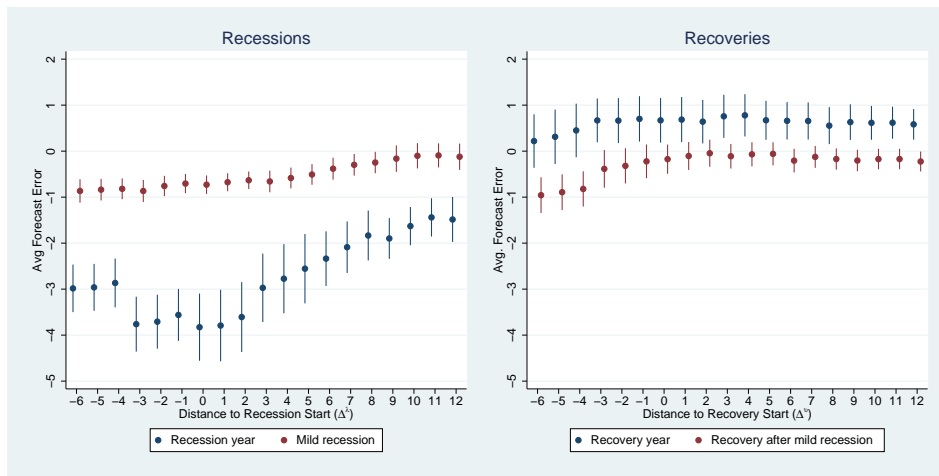
**Figure 3:** Number of Observations Around Business-Cycle Turning Points

Notes: Number of available forecast for recession and recovery years as a function of the temporal distance to the start of recessions or recoveries, respectively.

Figure 4: Systematic Forecast Errors around Business-Cycle Turning Points

Notes: This figure shows estimates for the systematic forecast errors (in percentage points) as a function of the temporal distance (in months) to the start of a recession or recovery. The dots represent point estimates from a regression of the forecast error on a set of 19 dummy variables (one for each month around the start of a recession or recovery). The vertical lines indicate 95% confidence bands surrounding the point estimates. 'Mild recession' refers to those cases in which a recession did not lead to a negative annual growth rate.

Figure 5: Synthetic Forecast Errors around Business-Cycle Turning Points
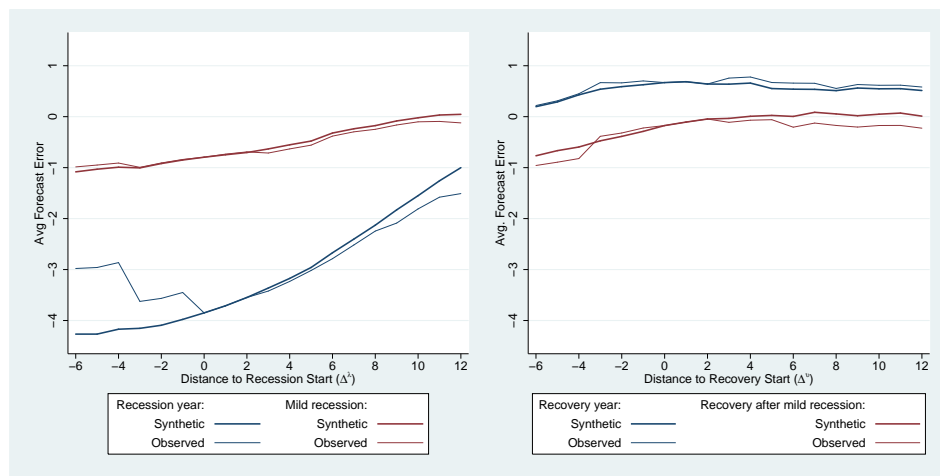
Notes: This figure shows the evolution of observed systematic forecast error as a function of the temporal distance to the start of recessions or recoveries together with the synthetic measure of these forecast errors as described in the text. 'Mild recession' refers to those cases in which a recession did not lead to a negative annual growth rate.
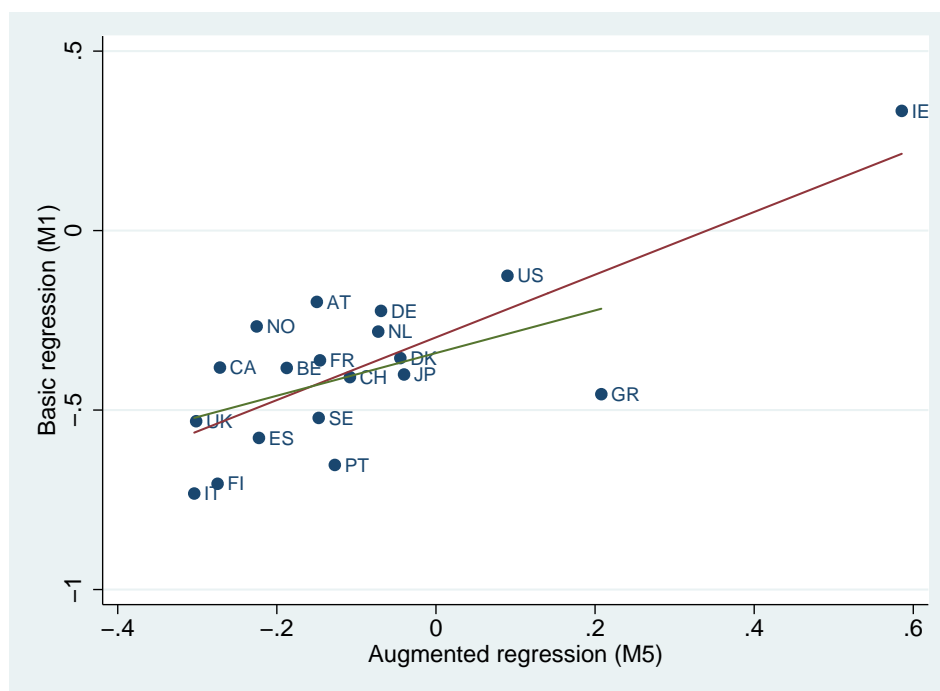
Figure 6: Systematic Forecast Errors by Horizon and by Country
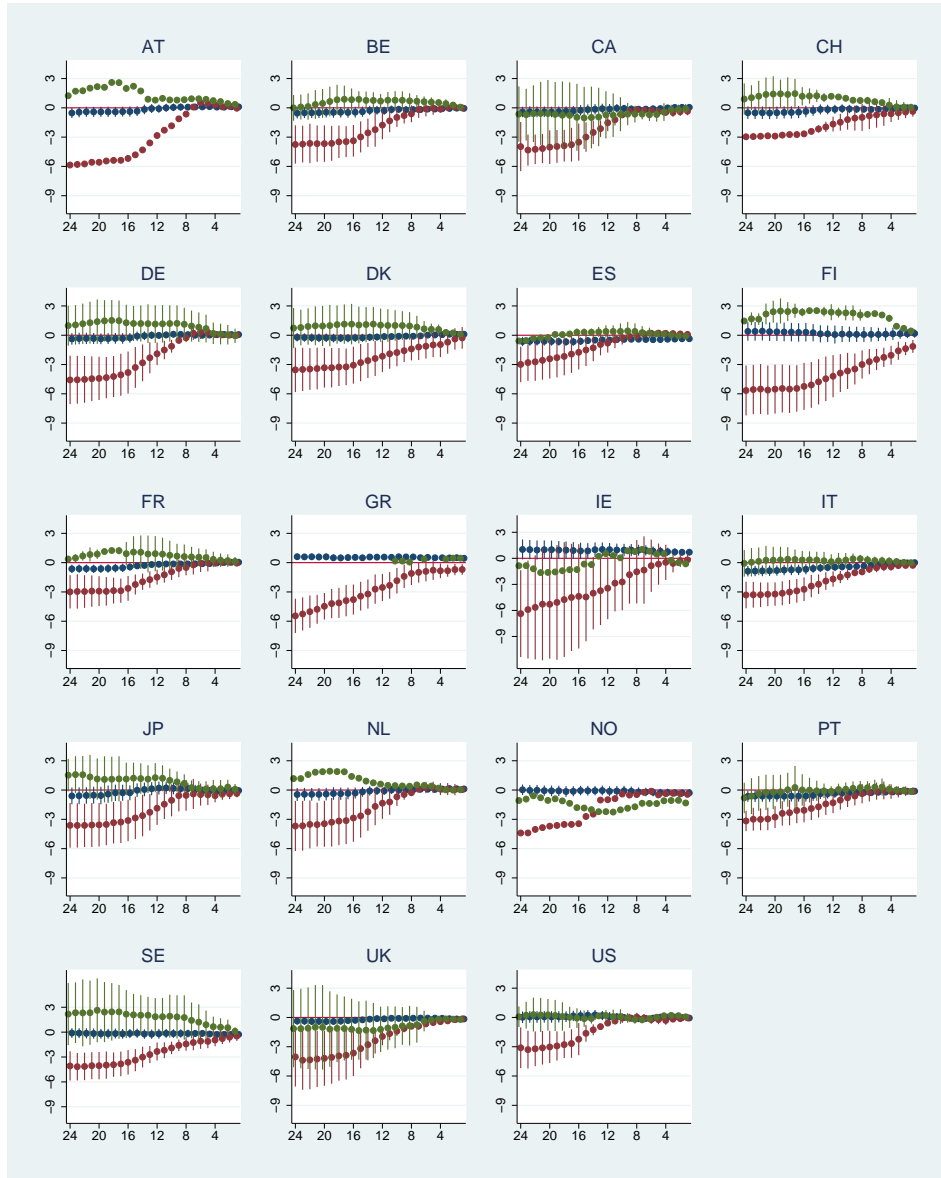
Notes: All graphs show systematic forecast errors in percentage points. For further notes, see table 2.

Figure 7: Relating Forecast Bias and Systematic Forecast Errors during Expansions

Notes: The dots represent the estimates of the country-specific systematic forecast errors ($\alpha + \alpha_i$, in percentage points) from specifications M1 (basic regression) and M5 (augmented regression). The two lines visualize correlation coefficients estimated by OLS with (red) and without (green) the observation for Ireland.

Figure 8: Systematic Forecast Errors and Changes in Trend Growth Rates
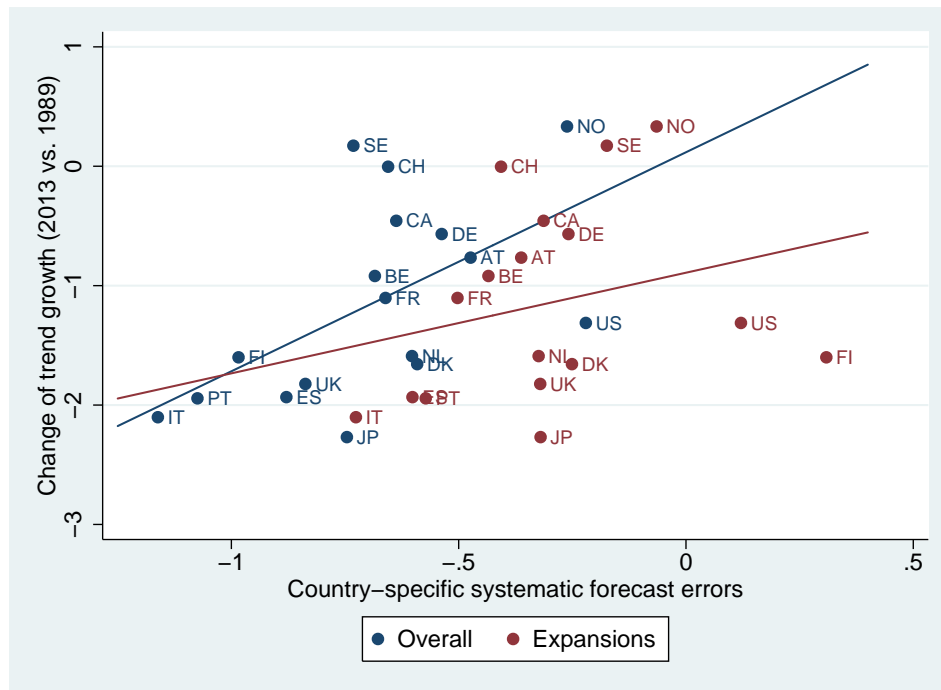
Notes: Each dot represents one country in our sample. The change in trend growth is calculated as the difference (in percentage points) between the average growth rate of potential output between 2009 and 2013 and that between 1985 and 1989 (as estimated by the OECD). The country-specific systematic forecast errors ($\alpha + \alpha_i$, in percentage points) are based on specifications M1 (basic regression) and M5 (augmented regression).

Appendix A   Data Sources

A.1   Consensus Forecast Data

We use a panel of forecasts for the growth rate of real GDP for 19 advanced economies over the period 1990 to 2013 provided by Consensus Economics. Each month, Consensus Economics asks a panel of professional forecasters to report their forecasts for the annual GDP growth rates of the current and the next calendar year. We base our analysis on the average (so-called consensus) forecasts from each forecast period. Thus, the data set contains a sequence of 24 forecasts for each target year and country. These forecasts are made between January of the year before the target year and December of the target year.

A.2   Real-Time Data on GDP Growth

Table A.1 lists the sources for each real-time observation of real GDP growth in our sample. Figure A.1 shows the real-time growth rates, the current-vintage growth rates, and the average forecasts made in December of the previous year for all countries.

A.3   OECD Estimates of Potential GDP Growth

We use estimates of potential output ('GDPVTR') from the OECD Economic Outlook from spring 2014, whenever the data are available, to compute the average of the growth rate of potential output for the periods 1985–1989 and 2009–2013. In those cases for which these data vintages do not contain the required information, we resort to the most recent vintages of the Economic Outlook that contains the relevant information.

By this rule, we use data from the Economic Outlook from spring 2005 to compute the average of the growth rate of potential output in Germany for the period 1985–1989, and we use data from the Economic Outlook from spring 2011 to compute the average of the growth rate of potential output in Greece and Ireland for the period 1985–1989.

Figure A.1: Data by Country
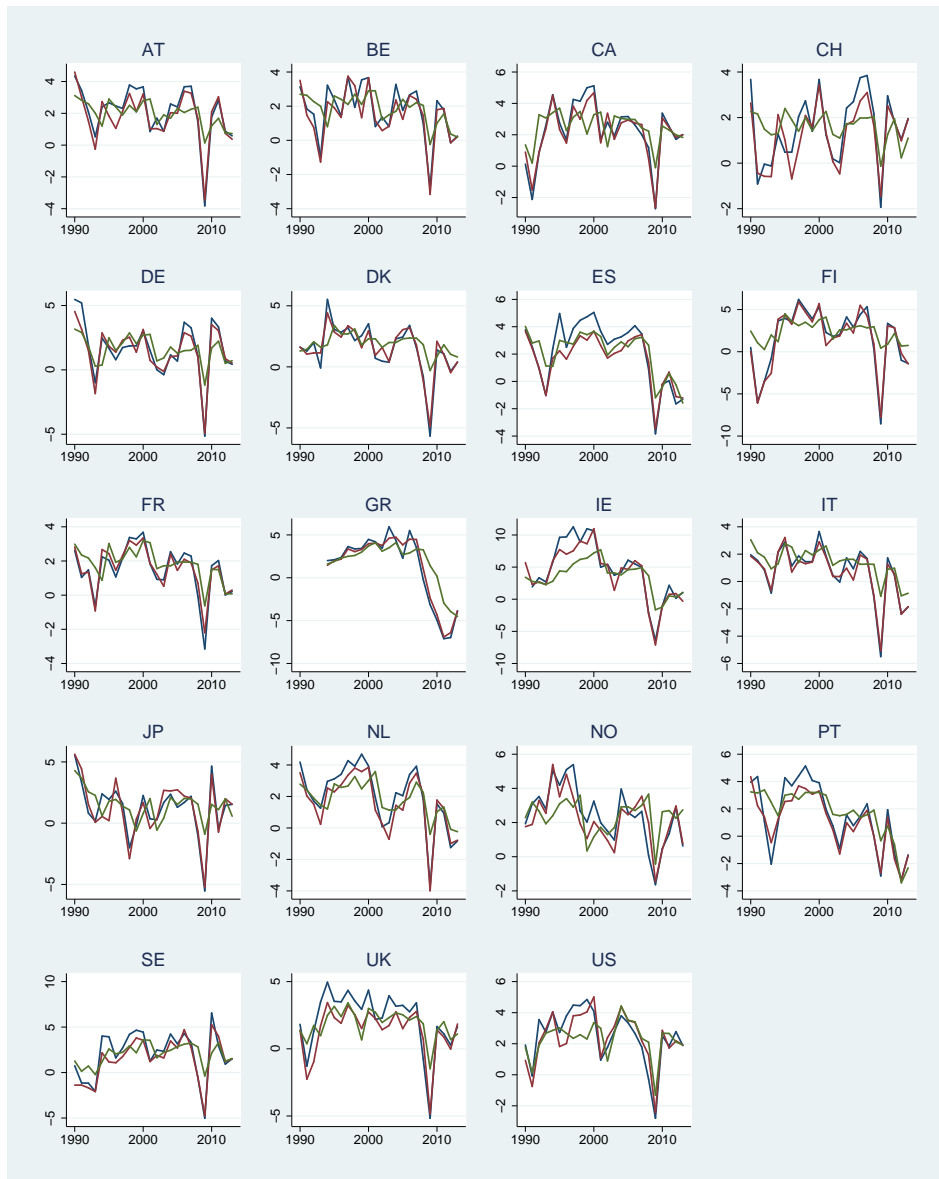
Notes: All graphs show annual real GDP growth and the consensus forecasts thereof in percentage points. The blue lines indicate the growth rates according to the current data vintage. The red lines indicate the growth rates in real-time (first available release). The green lines indicate the consensus forecasts for the annual growth rates made in December of the preceding year.

Table A.1: Sources of Real-Time Data

	ORDRD	RTHDS	EO
AT	1998-2013	1993-1995	1990-1992, 1996, 1997
BE	1998-2013	1997	1990-1996
CA	1998-2013	1990-1997	-
CH	1998-2013	1990-1997	-
DE	1998-2013	1990-1997	-
DK	1998-2013	1993-1997	1990-1992
ES	1998-2013	1994-1997	1990-1993
FI	1998-2013	1993-1997	1990-1992
FR	1998-2013	1990-1997	-
GR	2001-2010	-	1990-2000, 2011-2013
IE	2002-2013	-	1990-2001
IT	1998-2013	1990-1997	-
JP	1998-2013	1990-1997	-
NL	1998-2013	1993-1997	1990-1992
NO	1998-2013	1993-1997	1990-1992
PT	1998, 2000-2013	1997	1990-1996, 1999
SE	1998-2013	1990-1997	-
UK	1998-2013	1990-1997	-
US	1998-2013	1990-1997	-

Notes: The table shows the sources we used for the real-time annual growth rates for real GDP for each country and year in our sample. 'ORDRD' refers to the *Original Release Data and Revisions Database*, provided by the OECD. 'RTHDS' refers to the *Real-Time Historical Dataset for the OECD*, provided by the Federal Reserve Bank of Dallas (http://dallasfed.org/institute/oecd/index.cfm). 'EO' refers to the Economic Outlook published by the OECD. All values are taken from the earliest vintages of each source that provided the annual growth rate of the previous year or allowed the computation of that growth rate.

Appendix B   Results Based on Current Data Vintage

Table B.1: Modified Holden-Peel Panel Regressions (Current-Vintage Data)

	M1	M2	M3	M4	M5	M6
β_1		-2.607^{***}	-2.534^{***}	-2.467^{***}	-3.937^{***}	-0.605
		(-10.55)	(-7.28)	(-7.08)	(-8.81)	(-0.91)
β_2			-0.118	-0.113	0.300	-0.401
			(-0.31)	(-0.29)	(0.67)	(-1.43)
β_3				0.607^{***}	0.236	-1.172^{*}
				(2.92)	(0.99)	(-1.98)
δ_1					1.521^{***}	0.720
					(5.57)	(1.14)
δ_2					0.505^{**}	1.190^{*}
					(2.35)	(1.91)
γ_1						-0.014^{***}
						(-3.46)
γ_2						-0.138^{***}
						(-4.44)
γ_3						0.009
						(0.18)
γ_4						0.078^{**}
						(2.37)
γ_5						-0.070^{**}
						(-2.43)
γ_6						-0.008
						(-0.26)
α	-0.112^{***}	0.318^{***}	0.318^{***}	0.253^{***}	0.254^{***}	0.427^{***}
	(-6.69)	(7.81)	(7.70)	(5.73)	(5.68)	(10.35)
Obs.	10663	10663	10663	10663	10663	10663
R^2	0.00	0.30	0.30	0.31	0.33	0.42

Notes: The estimated equation is given by (5). The constant α is identified by requiring that the sum of the country-fixed effects (not shown in the table) be equal to 0. The numbers in parentheses are t-statistics based on Huber-White robust standard errors.
* denotes significance at a 10% confidence level.
** denotes significance at a 5% confidence level.
*** denotes significance at a 1% confidence level.

Figure B.1: Systematic Forecast Errors by Horizon (Current Vintage)
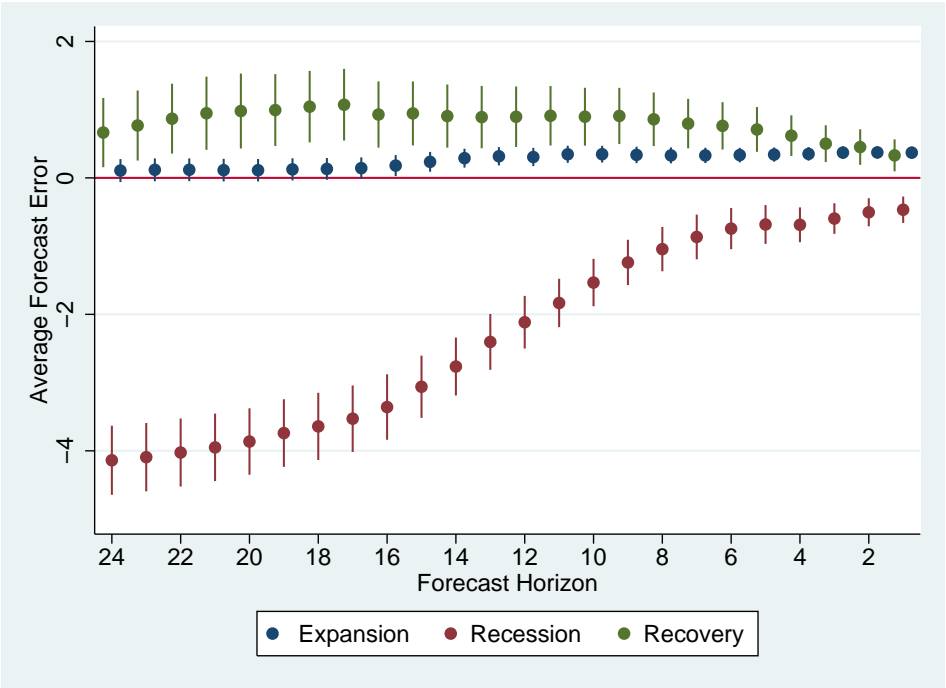
Notes: This figure shows estimates for the systematic forecast errors (in percentage points and computed based on current-vintage data) as a function of the forecast horizon. The dots refer to point estimates from a regression of the forecast error on a set of 24 dummy variables (one for each of the different forecast horizons). The vertical lines indicate 95% confidence bands surrounding the point estimates. All estimates are based on current-vintages data.