OECD Economics Department Working Papers No. 1428

Designing fan charts for GDP growth forecasts to better reflect downturn risks

David Turner

https://dx.doi.org/10.1787/e86f1bfc-en

OECD

Unclassified **ECO/WKP(2017)60**

Organisation de Coopération et de Développement Économiques
Organisation for Economic Co-operation and Development **10-Nov-2017**

English - Or. English

ECONOMICS DEPARTMENT

ECO/WKP(2017)60
Unclassified

DESIGNING FAN CHARTS FOR GDP GROWTH FORECASTS TO BETTER REFLECT DOWNTURN RISKS

ECONOMICS DEPARTMENT WORKING PAPERS No. 1428

By David Turner

OECD Working Papers should not be reported as representing the official views of the OECD or of its member countries. The opinions expressed and arguments employed are those of the author(s).

Authorised for publication by Christian Kastrop, Director, Policy Studies Branch, Economics Department.

All Economics Department Working Papers are available at www.oecd.org/eco/workingpapers.

JT03422664

This document, as well as any data and map included herein, are without prejudice to the status of or sovereignty over any territory, to the delimitation of international frontiers and boundaries and to the name of any territory, city or area.

English - Or. English

OECD Working Papers should not be reported as representing the official views of the OECD or of its member countries. The opinions expressed and arguments employed are those of the author(s).

Working Papers describe preliminary results or research in progress by the author(s) and are published to stimulate discussion on a broad range of issues on which the OECD works.

Comments on Working Papers are welcomed, and may be sent to OECD Economics Department, 2 rue André Pascal, 75775 Paris Cedex 16, France, or by e-mail to eco.contact@oecd.org.

All Economics Department Working Papers are available at www.oecd.org/eco/workingpapers.

© OECD (2017)

You can copy, download or print OECD content for your own use, and you can include excerpts from OECD publications, databases and multimedia products in your own documents, presentations, blogs, websites and teaching materials, provided that suitable acknowledgment of OECD as source and copyright owner is given. All requests for commercial use and translation rights should be submitted to **rights@oecd.org**

ABSTRACT/RÉSUMÉ

Designing Fan Charts for GDP Growth Forecasts to Better Reflect Downturn Risks

Forecasts of GDP growth are typically over-optimistic for horizons beyond the current year, particularly because they fail to predict the occurrence or severity of future downturns. Macroeconomic forecasters have also long been under pressure to convey the uncertainty surrounding their forecasts, particularly since the financial crisis. The current paper proposes a method to address both these issues simultaneously by constructing fan charts which are parameterised on the basis of the historical forecasting track record, but distinguish between a "safe" regime and a "downturn-risk" regime. To identify the two regimes, use is made of recent OECD work on early warning indicators of a prospective downturn, relating to housing market or credit developments. Thus, when an early warning indicator is “flashing", the associated fan chart is not only wider to reflect increased uncertainty, but is also skewed to reflect greater downside risks using a two-piece normal distribution of the form used by central banks to provide fan charts around inflation forecasts. Conversely, in a safe regime, when the early warning indicators are not flashing, as well as being symmetric, the fan chart is narrower both relative to the downturn-risk regime and relative to what the fan chart would be if the dispersion was calculated with respect to the entire forecast track record with no distinction between regimes. The method is illustrated by reference to OECD GDP forecasts for the major seven economies made just prior to the global financial crisis, with fan charts calibrated using the track record of forecasts published in the OECD *Economic Outlook*. Fan charts which take account of early warning indicators in this way are much better at encapsulating the outturns associated with a downturn than a symmetrical fan chart calibrated indiscriminately on all forecast errors.

JEL codes: E58, E17, E65, E66, GO1
Keywords: Fan charts, economic forecasts, uncertainty, risk, downturn

*****

Concevoir des graphiques en éventail des prévisions de croissance du PIB pour mieux refléter les risques à la baisse

Les prévisions de croissance du PIB sont typiquement trop optimistes à des horizons dépassant l’année courante, particulièrement car elles échouent à prédire l’arrivée ou la sévérité des crises futures. Les macro-économistes sont depuis longtemps sous pression pour communiquer l’incertitude entourant leurs prévisions, surtout depuis la crise financière. Ce papier propose une méthode pour traiter simultanément ces deux problèmes, en construisant des graphiques en éventail, paramétrés par l’historique des prévisions passées, qui font la différence entre un régime « sûr » et « risqué ». Pour identifier ces deux régimes, les travaux récents de l’OCDE sur les indicateurs avancés des crises potentielles, liés au marché du logement ou à l’évolution du crédit, sont utilisés. Ainsi, lorsqu’un indicateur « clignote », le graphique en éventail associé est non seulement plus large pour refléter l’incertitude, mais aussi déformé pour montrer les risques plus importants à la baisse, en utilisant une distribution normale à deux-pièces similaire à celle utilisée par les banques centrales pour produire les graphiques en éventail des prévisions d’inflation. Inversement, dans un régime sûr, lorsque les indicateurs avancés ne clignotent pas, les graphiques en éventail sont non seulement symétriques mais également plus étroits comparés à ceux du régime risqué et à ce que seraient les graphiques si la dispersion était calculée sur tout l’historique des prévisions sans distinguer les régimes. Cette méthode est illustrée par les prévisions du PIB de l’OCDE pour les sept économies majeures effectuées juste avant la crise financière, avec des graphiques en éventail calibrés par l’historique des prévisions publiées dans les Perspectives économiques de l’OCDE. Les graphiques en éventail qui prennent en compte les indicateurs avancés sont beaucoup plus performants pour contenir les résultats associés à une crise que ceux qui sont symétriques et calibrés sur toutes les erreurs de prévisions.

Codes JEL: E58, E17, E65, E66, GO1
Mots clés : Graphique en éventail, prévisions économiques, incertitude, risque, crise

TABLE OF CONTENTS

1. Introduction and summary 6
2. An examination of OECD GDP growth forecast errors 7
2.1 Characteristics of current-year forecast errors 7
2.2 Characteristics of year-ahead forecast errors 8
2.3. Implications of past forecasting errors for formulating fan charts 10
3. Identifying downturn-risk regimes using early warning indicators 10
4. Parameterising fan charts for different regimes 12
4.1 Central bank approaches to constructing fan charts 12
4.2 Fan charts under a safe regime 12
4.3 Constructing fan charts to reflect downturn risks 13
5. Assessing risk in forecasts made just prior to the global financial crisis 17
5.1 Risk assessment in the May 2008 Economic Outlook 17
5.2 Fan charts for forecasts made just prior to the global financial crisis 17
5.3 Sensitivity testing 27
6. False alarms and missed downturns 28
7. Further possible developments 30

REFERENCES 32

ANNEX A. OECD FORECAST ERRORS OF G7 GDP GROWTH 34

ANNEX B. THE TWO-PIECE NORMAL DISTRIBUTION AND PREDICTION INTERVALS 37

Tables

Table 1. Definition of preferred domestic early warning indicators 11
Table 2. RMSE of year-ahead GDP forecast growth, full sample and safe regime 13
Table 3. Preferred equations to explain forecast errors with domestic early warning indicators 15
Table 4. Preferred equations to explain forecast errors with domestic and international indicators 16
Table 5. Position of the 2009 GDP growth outturn in alternative fan charts 19
Table 6. Sensitivity of preferred equations to inclusion of a global financial crisis dummy 27
Table 7. Sensitivity of the fan chart position of the 2009 GDP outturn to a crisis dummy 28
Table 8. Summary table of early warning alarms and downturns for the G7, 1981-2015 29

Figures

Figure 1. Histogram of pooled normalised year-ahead forecast errors for the G7 9
Figure 2. Q-Q plot of pooled normalised year-ahead forecast errors for the G7 9
Figure 3. Fan charts for the Spring 2008 GDP growth forecast for the United States 20
Figure 4. Fan charts for the Spring 2008 GDP growth forecast for Japan 21
Figure 5. Fan charts for the Spring 2008 GDP growth forecast for Germany 22
Figure 6. Fan charts for the Spring 2008 GDP growth forecast for France 23

Figure 7. Fan charts for the Spring 2008 GDP growth forecast for the United Kingdom 24
Figure 8. Fan charts for the Spring 2008 GDP growth forecast for Italy.....25
Figure 9. Fan charts for the Spring 2008 GDP growth forecast for Canada 26
Figure 10. The combined skew across all G7 fan charts.....29

DESIGNING FAN CHARTS FOR GDP GROWTH FORECASTS TO BETTER REFLECT DOWNTURN RISKS

By David Turner[1]

1. Introduction and summary

1. Forecasts of GDP growth are typically over-optimistic for horizons beyond the current year, particularly because they fail to predict the occurrence or severity of future downturns.[2] Macroeconomic forecasters have also long been under pressure to convey the uncertainty surrounding their forecasts (Tay and Wallis, 2000), particularly since the financial crisis (Pain and Lewis, 2014). The current paper proposes a method to address both these issues simultaneously by constructing fan charts which are parameterised on the basis of the historical forecasting track record, but distinguish between a "safe" regime and a "downturn-risk" regime using so-called early warning indicators of prospective downturns.

2. The main findings of the paper are as follows:

- The largest errors in year-ahead OECD forecasts of GDP growth for the G7 countries are overwhelmingly over-predictions associated with a failure to forecast the occurrence of downturns. Regression analysis for G7 countries suggests that early warning indicators, relating to credit and housing market developments, are strongly correlated with these over-predictions. These early warning indicators then form the basis of distinguishing a "safe" regime and a "downturn-risk" regime, with each regime having its own risk distribution and associated fan chart.

- The risk distribution associated with the downturn-risk regime is characterised by a skewed two-piece normal distribution where the downside risk distribution is parameterised from an estimate of the mean forecast error during downturn-risk periods. In addition to domestic early warning indicators, a sum of similar indicators in other countries above a minimum threshold, is also useful in assessing the international spillover risk of a downturn and can further increase the negative skew in the fan chart.

- Conversely, in a safe regime, when the early warning indicators are not flashing, the fan chart is symmetric and narrower both relative to the downturn-risk regime and relative to what it would be if the dispersion was calculated with respect to the entire forecast track record with no distinction between regimes.

---

1. The author is Head of the Macroeconomic Analysis Division in the Economics Department at the OECD. My thanks to Professor Ken Wallis (University of Warwick), Andrej Sokol (Bank of England), Marie Hesselman (Sveriges Riksbank), Nigel Pain, Douglas Sutherland, Mauro Pisu and Peter Hoeller (all from the OECD Economics Department) as well as other participants at an internal OECD Economics Department seminar and at the Project LINK meeting for comments on earlier versions of this paper, with the usual disclaimer. Particular thanks also to Mikkel Hermansen for advice in constructing early warning indicators, Jeroen Meyer for statistical analysis of forecast errors, Thomas Chalaux for generating the fan chart graphics and Veronica Humi for help in document preparation (all from the OECD Economics Department).
2. For example, Lougani (2001) analyses private sector forecasts of growth and concludes *"the record of failure to predict recessions is virtually untarnished"* as well as that *"there is a high degree of similarity between private forecasts and those of international organisations"*.

- Following the suggested methodology, the degree of negative skew when early warning indicators are flashing can be substantial. When a domestic early warning indicator is flashing the extent of negative skew typically ranges between 1 and 2 percentage points, but when such indicators are flashing in many countries simultaneously, as they were in the prelude to the global financial crisis, the degree of negative skew can be many times greater. Fan charts which take account of early warning indicators in this way are much better at encapsulating the outturns associated with a downturn than a symmetrical fan chart calibrated indiscriminately on all forecast errors.

- There will be "false alarms" when an early warning indicator is flashing but no downturn occurs in the immediate forecast horizon. However, on many such occasions there will be a series of further alarms as tensions in the housing market or credit growth continue to build until the bubble bursts and there is an eventual downturn. Indeed, such "false alarms" may provide a more useful and timely warning to which policy can react than a "genuine" alarm just prior to a downturn. In any case, the difficulty of precisely predicting the timing of any downturn, underlines that the early warning indicators are better employed in quantifying risks to a central forecast than in adjusting the forecast itself.

3. The remainder of the paper is organised as follows. Section 2 examines the OECD's forecast track record, distinguishing between current-year and year-ahead forecasts, as motivation for the form of fan charts which are advocated. In Section 3, recent OECD work to construct early warning indicators is summarised, as well as how it has been adapted for the current paper. In Section 4, the method of parameterising fan charts for different regimes is described and it is illustrated with fan charts for past OECD forecasts published just prior to the global financial crisis in Section 5. In Section 6, the problems in using early warning indicators in terms of false alarms and missed downturns are addressed. Finally further possible developments of the methodology are discussed in Section 7.

2. An examination of OECD GDP growth forecast errors

4. This section analyses OECD forecasts of GDP growth for each of the G7 economies published in the May or June (hereafter "Spring") edition of the *OECD Economic Outlook*, distinguishing between forecasts for the current year and the year-ahead (summary tabulations are provided in Annex A). Together with past evaluations of forecasting performance -- in particular by Vogel (2007) and Pain and Lewis (2014), summarised in Turner (2016) -- this provides the motivation for the form of the fan charts advocated in later sections. Forecast errors are defined throughout as the outturn less the forecast (so that a negative error corresponds to an over-prediction), and the outturn for a particular year is defined as that published in the Spring *Economic Outlook* of the following year.

2.1 Characteristics of current-year forecast errors

5. OECD forecasts of GDP growth for the G7 countries made in the Spring for the current year are available from 1971. There is no obvious sign of systematic bias in the current-year forecasts; over the full sample, the mean forecast error across all G7 countries varies between -0.02 and -0.16% points and formal tests in previous evaluations of OECD forecast performance do not reject unbiasedness.

6. A striking feature of the current-year forecast is that there is an improvement in forecast performance over time; both the mean absolute error (MAE) and the root-mean-square error (RMSE) have declined for all G7 countries and on average by about half between the periods 1970-90 and 2000-16 (Annex A, Table A1). Likely reasons for this improvement in performance include a decline in the volatility of GDP growth as well as improved access and timeliness of national accounts and other related

data by national statistical authorities, as well as better use of short-term indicator models which use monthly data to predict the current and next quarter's GDP growth (Turner, 2016).

2.2 Characteristics of year-ahead forecast errors

7. The available sample of year-ahead Spring forecasts begins with forecasts made in 1981, as year-ahead forecasts were not made before that year. The year-ahead forecast errors for the G7 economies tend to confirm problems identified in past evaluations of forecast performance.

- As might be expected, forecast accuracy deteriorates as the forecast horizon is extended; the MAE and RMSE for year-ahead forecasts averaged across all G7 countries is 1.4 and 1.8 percentage points, respectively, more than double that for current-year forecasts over a similar sample period (Annex A, Table A1 and A2).

- Of greater concern is evidence of bias in the year-ahead forecasts. The mean forecast error is negative and economically meaningful for all countries, averaging -0.6% points across all G7 countries, implying forecasts typically over-predict growth in the year ahead (Annex A, Table A3). This is consistent with formal tests in previous evaluations of OECD forecasts, which have often rejected a null hypothesis of unbiasedness in one-year-ahead forecasts.

- The largest absolute forecast errors are all over-predictions. Across all G7 countries, there are 13 occurrences when the absolute forecast error exceeds twice the RMSE of all forecast errors for that country, but all these occurrences are over-predictions rather than under-predictions and all are associated with negative growth outturns (Annex A, Table A3). This confirms a finding of previous evaluations of both OECD and other macroeconomic forecasts, namely that they are invariably poor at predicting turning points, particularly downturns (Fildes and Steckler, 2002; Loungani, 2001; Abreu, 2011; Pain and Lewis, 2016).

- After standardising forecast errors on the respective country RMSEs, the joint assumption of unbiasedness and normality is overwhelmingly rejected when pooling errors across the G7.[3] This is confirmed visually by a histogram and Q-Q plot which serve to emphasis both the negative skewness of the errors and the number of large negative outliers (Figures 1 and 2).

- Unlike the current-year forecasts, there is little evidence of any improvement in year-ahead forecast performance over time; neither the MAE nor the RMSE is significantly or systematically lower over the period 2000-16 compared to the 1980s and 1990s (Annex A, Table A2).

- A further feature of the year-ahead forecast errors is that they are strongly positively correlated across countries (Annex A, Table A4), particularly between the United States and Canada, and between the major European economies.

3. Wilk-Shapiro and Kolmogorov-Smirnov tests of normality are both strongly rejected at the 0.1% significance level.

Figure 1. Histogram of pooled normalised year-ahead forecast errors for the G7
1981-2015, % points

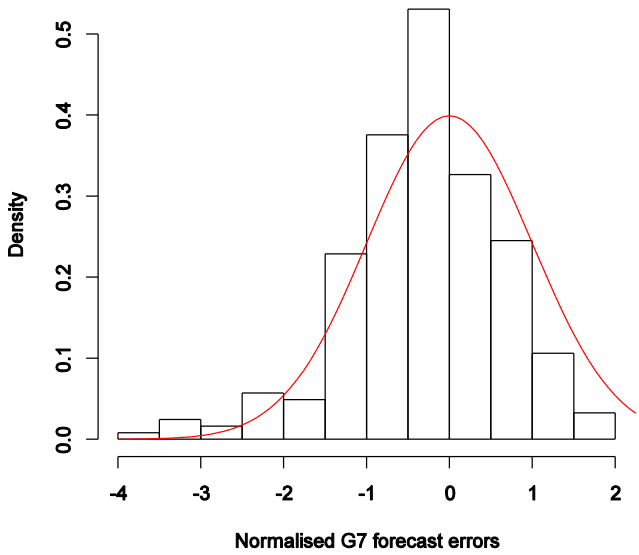

Note: The bars show the frequency distribution of year-ahead forecast errors for all G7 countries together after the errors have been normalised on each country's RMSE. The red curve shows the hypothetical frequency distribution of a standard normal variable.

Figure 2. Q-Q plot of pooled normalised year-ahead forecast errors for the G7
1981-2015, % points

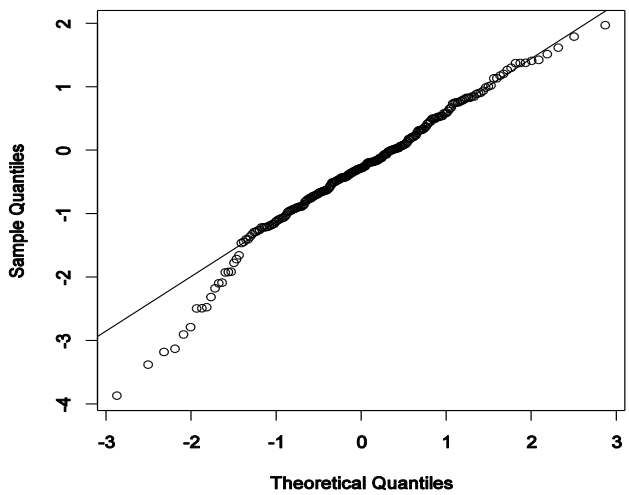

Note: A Q-Q plot is a scatterplot of two sets of quantiles; in this case, the quantiles on the y-axis are of the pooled year-ahead forecast errors (normalised by dividing by the country-specific RMSE) and on the x-axis the quantiles for the standard normal distribution. If both sets of quantiles came from the same distribution, the points would form a straight line.

2.3. Implications of past forecasting errors for formulating fan charts

8. Based on the preceding analysis, the distribution underlying the calculation of the forecast fan charts is based on the following assumptions:

- For the current year-forecast it will be assumed that errors are normally distributed with a zero mean around the published forecast. Given the improvement in forecast performance observed across all G7 countries, an estimate of the distribution's standard deviation is based on the forecasting track record since 2000.

- In order to gauge uncertainty surrounding the year-ahead forecast, there is a need to identify if there is risk of a downturn. If so, then uncertainty will be greater, with errors more likely to be skewed to the downside. The following sections describe the approach to identifying the downturn risks and parameterising the skewed-to-the-downside distribution.

3. Identifying downturn-risk regimes using early warning indicators

9. Recent OECD research by Hermansen and Röhn (2016), hereafter HR, provides empirical evidence on the usefulness of a set of early warning indicators to predict severe downturns in OECD economies over the period 1970-2014. The usefulness of the indicators is assessed on the basis of the signalling approach, whereby an indicator signals a vulnerable state of the economy if it crosses a threshold.[4] Thresholds levels are chosen so as to strike a balance between the risks of missing vulnerable states and issuing many false alarms. In particular, a loss function is used to determine the optimal thresholds, which explicitly takes into account preferences between these type I and type II errors. An indicator is judged to be 'useful' if its predictions result in a lower loss compared to a benchmark in which the indicator is ignored.

10. The current paper focuses only on the set of domestic indicators which were found by HR to consistently perform the best across all OECD, namely housing market variables (real house prices, the house-price-to-rent ratio and the house-price-to-disposable income ratio) and credit variables (total private credit and bank credit). However, an important difference with HR, is that the choice of indicators within this subset is country-specific, based on the same evaluation criteria described in HR but applied to each country individually. On this basis, it is possible to find a credit-related early warning indicator which is useful in predicting downturns for every G7 country and a housing-related indicator for every G7 country except Germany and Japan (Table 1).

4. For examples of the signalling approach see Kaminsky et al. (1998), Borio and Lowe (2002) and Behn et al. (2013).

Table 1. Definition of preferred domestic early warning indicators

Country	Early warning variable	Functional form	Threshold	Indicator variable name
United States	House-price-to-disposable-income ratio (% of GDP)	No transformation	75%	HPY
United States	Total private credit (% of GDP)	4 year growth rate	90%	PCR
Japan	Total private credit (% of GDP)	No transformation	80%	PCR
Germany	Private bank credit (% of GDP)	6 quarter growth rate	75%	BCR
France	Private bank credit (% of GDP)	6 quarter growth rate	75%	BCR
France	Real house prices	Deviations from a 20 quarter moving average	85%	RHP
United Kingdom	House-price-to-rent ratio	5 year growth rate	80%	HPR
United Kingdom	Total private credit (% of GDP)	3 year growth rate	85%	PCR
Italy	House-price-to-disposable-income ratio (% of GDP)	No transformation	80%	HPY
Italy	Private bank credit (% of GDP)	6 quarter growth rate	75%	BCR
Canada	Private bank credit (% of GDP)	6 quarter growth rate	75%	BCR
Canada	Real house prices	Deviations from a 20 quarter moving average	85%	RHP

Note: All variables correspond to those defined in the second column of Table 1 in Hermansen and Röhn (2016).

Source: Hermansen and Röhn (2016) and authors' calculations.

11. In most cases, the current paper adopts the same functional form of the indicator found to be optimal across all OECD countries by HR.[5]An exception is the United Kingdom, where the same housing market and credit variables are used as the basis for the early warning indicators, but where an alternative functional form and threshold was found to be more successful in predicting downturns.[6] In particular, the functional form found to be most successful when applied uniformly across all OECD countries, is simply the unadjusted ratio of the house-price-to-rent ratio, the house-price-to-disposable income ratio and the total-private-credit-to-GDP ratio, this is not the most appropriate functional form for all countries. For example, for the United Kingdom, there is a pronounced upwards trend in all of these ratios, which means that the variables only pass alert thresholds in the later stages of the sample, even though it is clear that rapid increases in these variables preceded downturns earlier in the sample. This suggests that a functional form capturing the rate-of-change might be more successful, and testing using the same methodology as HR finds that a growth-rate functional form is indeed more successful in predicting downturns for the United Kingdom. A similar argument applies to the United States in terms of the private credit variable expressed as a percentage of GDP. These examples raise the possibility that customising the form of the

5. The form of the indicator variables are those shown in the second column of Table 1 of Hermansen and Röhn (2016) corresponding to the preference parameter $\sigma = 0.8$.

6. A more minor change to the indicators used by HR was to the alert threshold for the US house-price-income ratio which was here lowered to 75% from 80%. Otherwise, the alert thresholds remain the same as those used by HR.

early warning indicators for individual countries may improve their predictive power and hence be more useful in constructing fan charts, an issue which is further discussed in Section 7.

4. Parameterising fan charts for different regimes

4.1 Central bank approaches to constructing fan charts

12. Forecast fan charts are one way for macroeconomic forecasters to convey the uncertainty surrounding their forecasts (Tay and Wallis, 2000). The term "fan chart" was coined by the Bank of England which has been using such charts since 1997 in its *Inflation Report* to describe the uncertainty around its forecast of inflation and GDP growth. For the Bank of England, the width of the fan chart is initially based on the dispersion of outturns around previous forecasts, but this is then modified by the Monetary Policy Committee based on their judgement as to whether uncertainty looking forward is likely to be greater or less than that past experience, and whether risks are skewed to the up- or down-side (Britton et al., 1997; Bank of England, 2005).

13. The Riksbank (the Swedish central bank) has also published an Inflation Report containing its inflation forecast with uncertainty bands around the forecast since 1997. In the case of the Riksbank, when the fan charts were first introduced, the underlying distribution was informed by model simulations of shocks that were judged likely to influence the forecast (Blix and Sellin, 1998). However, since 2007 fan charts are based on historical errors, with fan charts bands calculated using the historical RMSE and assuming that errors are normally distributed (Sveriges Riksbank, 2007).

14. While the Bank of England and Riksbank pioneered the use of fan charts in macroeconomic forecasting, this practice has since spread among many central banks as well as other major national and international institutions [see World Bank (2016) for a recent survey].

4.2 Fan charts under a safe regime

15. When no early warning indicators are flashing, the band widths of the fan chart are computed with respect to the RMSE of forecast errors which are typical for safe periods, rather than from the full sample of forecast errors. On average across G7 countries this will reduce the band widths of the fan chart by about one quarter (Table 2). For example, in the case of Italy, this reduces the dispersion of the fan chart for one-year ahead errors by more than one-third over any given prediction interval; thus (assuming errors are normally distributed) a 90% prediction interval for one-year ahead forecasts of Italian GDP growth based on the full sample of forecast errors would be more than 6 percentage points wide, whereas based on the reduced sample of safe periods it would be less than 4 percentage points.

Table 2. RMSE of year-ahead GDP forecast growth, full sample and safe regime

1982-2016, % points

	(1)	(2)	(3)
	RMSE		Ratio (2)/(1)
	Full sample	Safe regime	
United States	1.52	1.32	87%
Japan	2.29	1.67	73%
Germany	1.87	1.59	85%
France	1.42	1.24	87%
United Kingdom	1.64	0.98	60%
Italy	1.86	1.18	63%
Canada	2.13	1.65	77%

Note: The 'safe regime' is defined as periods in which early warning indicators are not flashing.

4.3 Constructing fan charts to reflect downturn risks

16. The approach adopted in this paper is different to that of the Bank of England and Riksbank, as the distribution underlying the fan chart is calculated solely from past forecast errors, but distinguishing between different regimes according to whether an early warning indicator is flashing.

- Under the safe regime it is assumed that possible outcomes are normally distributed around the published forecast, which is assumed to be the modal outcome. Given the symmetry of the normal distribution, the modal outcome is also the mean outcome. The dispersion of possible outcomes is based on the RMSE of forecast errors during periods in which early warning indicators are not flashing.

- Under the downside-risk regime, the possible outcomes are skewed to the downside, implying that if the published forecast is the modal outcome, then the mean outcome is below the published forecast.

17. While the approach to parameterising differs, the current paper follows the Bank of England and Riksbank in using the two-piece normal distribution to model the distribution of risks when they are skewed. The two-piece normal distribution can be summarised in terms of three parameters -- the mode (μ) and two standard deviations (σ_1 and σ_2) -- and is formed by taking two halves of normal distributions with parameters (μ_1 σ_1) and (μ, σ_2), and scaling them so that the probability density functions have a common value at the mode (see Annex B). With $\sigma_1 > \sigma_2$ the distribution will have negative skewness, and the relationship between the mean outcome, μ^*, and mode is given by (see John, 1982):

(1) $\mu - \mu^* = k\,(\sigma_1 - \sigma_2)$, where $k = (2/\pi)^{½}$

18. To estimate the three parameters of the two-piece normal distribution for the downturn-risk regime, the following assumptions are made:

i. The published forecast is the modal outcome, μ.

ii. The smaller standard deviation, σ_2, of the upper half of the distribution can be approximated by the RMSE of forecast errors in "safe" periods when there is deemed to be no risk of a

downturn. The reasoning is that if the outcome exceeds the published forecast so that there is no severe downturn, then there is little reason to expect the forecast error to be exceptional and hence it can be evaluated by the past forecast performance during safe periods.

iii. The larger standard deviation, σ_1, of the lower half of the distribution is computed by first estimating the mean forecast error during "downturn-risk" periods and then using equation (1), given the assumptions for μ and σ_2 above, to determine σ_1.

19. The mean forecast error during downturn-risk periods is estimated by regressing forecast errors over the full sample period for each individual G7 country separately on one, or sometimes two, early warning indicator(s), thus:

(2) $${}_t e_{t+1} = \beta_0 + \beta_1 w^1_t + \beta_2 w^2_t ,$$

where ${}_t e_{t+1}$ is the error in the forecast (outturn minus forecast) published in the Spring of year t for calendar year (t + 1) and w^i_t (i =1 or 2) is a binary variable representing the early warning indicator(s) taking the value 1 or 0, where 1 signals the risk of a downturn. Given that the original early warning indicators are quarterly variables and are intended to signal the onset of a downturn eight quarters in advance, w^i_t takes the value of 1 if the corresponding original quarterly early warning indicator was 1 in any of the quarters in year (t – 1) preceding the publication of the forecast or in the first two quarters of year t when the forecast was published. In addition, following the approach of HR, in the year following the occurrence of a severe downturn, the indicator is ignored by setting it to zero.[7]

4.3.1 Domestic downturn risks

20. Following this procedure, the estimated coefficients β_1 and (if present) β_2 then provide estimates of the mean error during downturn-risk periods.[8] Regressions of (1) for each of the G7 economies, using the preferred early warning indicators usually give well-determined statistically significant negative coefficients for β_1 and β_2, which take values of between -0.9 and -2.0 (Table 3). This suggests that the mean forecast error is significantly, in both statistical and economic terms, negative in “downturn-risk” periods. The magnitude of the effect of the early warning indicator on the mean forecast error is substantial and provides an estimate of the degree of skewness, the difference between the mean and modal outturns. Thus if a domestic early warning indicator is flashing, the degree of negative skew varies from -0.9 to -2.0 percentage points. On the other hand, the estimated intercept term, β_0, is usually small and statistically insignificant, consistent with the mean error being zero during “safe” periods.

7. HR explain that the first four quarters following the start of a severe recession are excluded from the evaluation sample since the behaviour of the vulnerability indicators is likely to be different during a severe recession compared to normal or pre-recession times as demonstrated by Bussiere and Frantzcher (2006).

8. If only w^1 is present, then the mean error is β_1 when the indicator is sounding an alarm. If both w^1 and w^2 are present, then the mean error depends on whether w^1, w^2 or both are sounding an alarm, in which cases the mean error is β_1, β_2 or (β_1+ β_1), respectively.

Table 3. Preferred equations to explain forecast errors with domestic early warning indicators

Dependent variable: year-ahead forecast error, sample period: 1981-2015

Variable	(1) USA	(2) USA	(3) USA	(4) JPN	(5) DEU	(6) FRA	(7) FRA	(8) FRA	(9) GBR	(10) GBR	(11) ITA	(12) ITA	(13) ITA	(14) CAN	(15) CAN
constant	0.00	-0.05	0.19	-0.01	-0.29	-0.28	-0.32	-0.20	0.13	0.26	-0.63	-0.50	-0.18	-0.321	-0.24
	-			-	-	-	-	-	-	-	*	-	-	-	-
HPY	-1.72		-1.68								-1.82		-1.67		
	**		**								***		***		
PCR		-1.16	-1.12	-1.88						-1.86					
		*	*	**						***					
BCR					-0.87	-0.90						-1.37	-1.21	-2.03	
					-	**						***	**	**	
RHP							-0.94								-1.46
							**								**
HPR									-1.60						
									***						
(RHP + BCR)								-0.62							
								**							

Notes:

1. Explanatory variables are all domestic early warning indicators, taking the value zero or one: HPY is related to the house-price-to disposable income ratio; PCR is related to the ratio of total private credit to GDP; BCR is related to the ratio of bank credit to GDP; RHP is related to real house prices. The definition of when each of these variables switch between zero and one is explained in Table 1.

2. Statistical significance of coefficients at the 1%, 5% and 10% levels are denoted by ***, ** and *, respectively. A coefficient which is not significant at the 10% level is denoted by "-".

4.3.2 International downturn risks

21. One of the main findings of HR, was that indicators of global risk, including global house prices and global measures of credit, often outperform domestic indicators, but subject to the caveat that they are suited to picking up recessions that affect a large number of countries simultaneously, such as the global financial crisis of 2008/9. The same indicators performed less well in the current framework. Moreover, constructing global indicators also raises difficult issues around country coverage and weighting schemes. Nevertheless, the strong cross-country correlation between year-ahead forecast errors, suggests international spillover effects are important.

22. To model such effects a simple alternative country-specific early warning variable has been constructed based on the number of 'other' G7 countries which are experiencing an early warning indicator alarm, but with some modification to give greater weight to other countries for which forecast errors are most closely correlated. The international early warning indicator is then judged to flashing if the number of other countries with an alarm exceeds a given threshold (which corresponds roughly with an 80th percentile threshold). Specifically, the country-specific international early warning indicator is determined as follows:

- For the United States and Japan, it is based on the number of other G7 countries above a threshold of three in which early warning indicators are flashing;
- For Canada, it is also based on the number of other G7 countries, but giving a triple weight to the United States and with a threshold of four;
- For the major European countries it is based on the number of other major European countries as well as the United States in which early warning indicators are flashing above a threshold of two.

23. Experimentation with the form of this variable suggested that it works better if instead of being a binary (0, 1) indicator, it is a variable taking a value of between 1 and (a hypothetical maximum of) 4, depending on the number of other G7 countries (exceeding the threshold) in which the early warning indicator is flashing.[9] This enables a distinction to be drawn between the severity of the international downturn pressures depending on the number of G7 countries in which early warning indicators are flashing: in the early 2000s, there were two to three G7 countries; in the early 1990s, three to four countries; and in the prelude to the global financial crisis of 2008/9 early warning indicators were flashing in five countries.

24. The addition of this international early warning variable to the regressions described by (2) is always highly significant, and for nearly all G7 countries (the only exception being France), it is possible to simultaneously estimate a mean forecast error related both to domestic and international early warning indicators (Table 4).

Table 4. Preferred equations to explain forecast errors with domestic and international indicators

Dependent variable: year-ahead forecast error, sample period: 1981-2015

Country	Constant[3]	Domestic indicator[1]		International indicator[2]		
		Coefficient[3]	Variable	Coefficient[3]	Other countries	Threshold
United States	0.35	-1.18	(HPY+PCR)	-1.76	Other G7	3
	-	***		**		
Japan	0.50	-2.18	PCR	-2.11	Other G7	3
	-	***		***		
Germany	0.24	-1.25	BCR	-1.23	FRA, GBR, ITA, USA	2
	-	**		***		
France	-0.25	-0.37	BCR	-1.50	DEU, GBR, ITA, USA	2
	-	-		**		
United Kingdom	0.44	-1.13	PCR	-2.50	DEU, FRA, ITA, USA	2
	*	**		***		
Italy	-0.14	-1.09	(HPY+BCR)	-1.12	DEU, FRA, GBR, USA	2
	-	***		**		
Canada	-0.01	-1.43	BCR	-1.16	Other G7, USA x 3	3
	-	**		***		

Notes:

1. Domestic indicator variables all take the value zero or one: HPY is related to the house-price-to disposable income ratio; PCR is related to the ratio of total private credit to GDP; BCR is related to the ratio of bank credit to GDP; RHP is related to real house prices. The definition of when each of these variables switch between zero and one is explained in Table 1.

2. The 'international indicator' is a variable measuring the number of other countries, defined in the penultimate column, which exceeds a threshold number of countries given in the final column.

3. Statistical significance of coefficients at the 1%, 5% and 10% levels are denoted by ***, ** and *, respectively. A coefficient which is not significant at the 10% level is denoted by "-".

9. In summing this indicator a complication is that two domestic early warning indicators are used in the preferred equations summarised in Table 4 for the United States and Italy. In the count of the international early warning indicator for other countries the United States and Italy are treated differently. In the case of the United States, the contribution of each indicator is scored separately, so if both credit and housing indicators are flashing, the United States adds two to the count of the international warning indicator. In the case of Italy, each indicator is assigned a weight of 0.5, so if both indicators are flashing it only adds one to the count of the international warning indicator. This differential treatment reflects the greater weight of the United States in the world economy and so in the transmission of shocks.

5. Assessing risk in forecasts made just prior to the global financial crisis

25. None of the early warning indicators used in the preceding empirical analysis are currently flashing, so instead of computing fan charts for current forecasts they are calculated for forecasts published in the period immediately preceding the global financial crisis.

5.1 Risk assessment in the May 2008 Economic Outlook

26. The projections published in the May 2008 *OECD Economic Outlook* did not include any explicit quantitative assessment of risk, but the discussion of individual country forecasts did provide a descriptive assessment. This discussion recognised that risks were unusually elevated, with risks from financial markets or from external demand or both being mentioned for all G7 countries. However, for the majority of G7 countries, risks were thought to be equally balanced to the up- and down-side, as demonstrated from the following excerpts from the May 2008 *Economic Outlook*:

- United States: *"Risks have widened. There is considerable uncertainty about the eventual scale of financial institution's losses… These risks could go either way."*

- Germany: "*(G)rowth could be weak if the financial turmoil turns out to be deeper and longer lasting than assumed… but… the risks could also go in the opposite direction."*

- France: *"The main risks to the projection remain focused on external developments. In particular, a longer-lasting downturn in the United States… however, the outcome could also turn out to be more favourable than projected, with positive spillover effects on French activity."*

- Italy: "*Risks to growth are symmetric. …although there may still be unpleasant surprises [from financial developments], the likelihood that Italy is vulnerable to domestic financial risk seems less than it did a few months ago. Growth could turn out to be lower than projected here, but a stronger pick-up in 2009… is also possible."*

- Canada: *"The risks to the outlook are for the credit crisis and adjustment in the housing market to be either sharper/longer or shallow/shorter, than currently expected, signifying respectively prolonged weakness or speedier-than-anticipated recovery."*

Only for Japan and the United Kingdom were risks unambiguously assessed as being to the downside:

- Japan: "(*E)xternal and domestic risks are mostly on the downside… including a sharper-than-expected decline in overseas demand and yen appreciation*."

- United Kingdom: "*There are downside risks to growth… For example, GDP growth could slow more markedly if financial sector health continues to deteriorate or if the housing market falls into a more significant slump."*

5.2 Fan charts for forecasts made just prior to the global financial crisis

27. At the time of the May 2008 *Economic Outlook* domestic early warning indicators were flashing in all G7 countries Except Japan and Germany. Thus, warning thresholds had been surpassed for: growth of private credit for the United States and the United Kingdom; growth of bank credit for France, Italy and Canada; house-price income ratios in the United States and Italy; growth in the house-to-rent ratio in the United Kingdom; and growth in real house prices in France and Canada.

28. For comparative purposes fan charts are calculated in two or three ways for each G7 country: firstly, to serve as a benchmark or 'straw man', a symmetrical dispersion is calculated assuming normality of forecast errors using the RMSE over the full historical sample; secondly, assuming a two-piece normal distribution with the skew to the fan chart computed using domestic warning indicators only, using the coefficient estimates in Table 3;[10] finally, the skew to the two-piece normal distribution is computed using both domestic and international early warning indicators, using the coefficient estimates in Table 4. These fan charts are not computed with real-time information, rather all of them use information on the full sample of errors and early warning indicators. Instead, the interest here is in comparing the fan charts to show the difference from utilising information on the early warning indicators.[11]

29. Even once the parameters of the two-piece normal distribution are determined, there is still a separate technical issue as to how to determine the prediction intervals on which the fan chart is based. The illustrative fan charts shown later in this section follow the practice of the Bank of England so that the baseline forecast is treated as the mode of the forecast distribution and the use of *highest probability density* intervals ensure that the baseline forecast is always within the central (darkest shaded) interval, although this choice is not without controversy (Wallis, 1999; and see the discussion in Appendix B).

30. A further issue relates to the number of prediction intervals shown in the fan chart. When the Bank of England began publishing fan charts they displayed intervals for each decile -- cumulatively 10%, 20% and 30% up to 90% -- although they have also more recently begun publishing an additional version of the fan chart with (only) three wider bands, covering the cumulative intervals 30%, 60% and 90%. Arguably, the latter version is visually easier to interpret. Moreover, when surveyed, macroeconomic forecasters often resort to very rounded numbers in attaching probabilities to possible outcomes and in the light of this "*uncertain uncertainty*" (Boero et al., 2015), showing fewer prediction intervals seems more appropriate. Accordingly, the fan charts in the current paper only distinguish three prediction intervals at 50%, 75% and 90%.[12]

31. On this basis, a series of fan charts is generated for each G7 economy for the May 2008 *Economic Outlook* forecasts (Figures 3 to 9) and the position of the outturn for GDP growth in 2009 in each fan chart is also summarised in Table 5:

- Calibrating the fan chart on the basis of historical errors, assuming symmetry (the 'straw man'), implies the outturn for 2009 is below, and usually far below, the 90% prediction interval for all G7 countries (as denoted by an "X" in in Table 5).

- Taking into account domestic early warning indicators brings the outturn for 2009 (as denoted by "W" in Table 5): within the 50-70% prediction interval for the United States; within the 70-90% interval for Canada and Italy; and for France and the United Kingdom the 2009 outturn is closer to, but still below, the lower limit of the 90% prediction interval. However, the position

10. For some countries there is more than one regression in Table 3 that can be used to parameterise the fan chart. In these cases the early warning indicators corresponding to these selected in Table 4 are chosen, although the coefficient estimates are still taken from Table 3.

11. HR perform an out-of-sample exercise on the early warning indicators and find that domestic credit and asset market indicators "perform particularly well, similar to the full sample results" as do the global early warning indicators.

12. An analogy with the appropriate disaggregation of prediction intervals might be with the reporting of the number of decimal places; quoting a number with many decimal places suggests a degree of confidence or certainty in a number, which in macroeconomic forecasting is rarely the case. Moreover, it is less easy to quickly discern the main messages from a table of many numbers, if each number is displayed with a large number of spurious decimal places.

of the outturn for Japan and Germany remains unchanged and well outside the lower limits of the fan chart, simply because domestic early warning indicators were not flashing in these countries just prior to the crisis.

- Finally, allowing for the international dimension of the crisis (as denoted by "G" in Table 5), by taking into account the early warning indicators flashing in other G7 countries as well domestic indicators, brings the outturn for 2009 within the 50% prediction interval for the United States, Japan and Canada, and within the 50-70% prediction interval for all other G7 countries.

Table 5. Position of the 2009 GDP growth outturn in alternative fan charts

Fan chart interval	USA	JPN	DEU	FRA	GBR	ITA	CAN
0-50%	G	G					G
50-70%	W		G	G	G	G	
70-90%						W	W
90-99%	X			W	W		X
99%+		XW	XW	X	X	X	

Note: The table shows the prediction interval in which the outturn for GDP growth in 2009 is located in fan charts constructed around the forecast made in the May 2008 *Economic Outlook,* according to alternative parametrisations of the fan chart: "X" denotes the position of the outturn in a fan chart constructed from past historical errors, assuming that they are symmetrical and ignoring the early warning indicators; "W" denotes the position in a fan chart constructed using information on the size of forecast errors when domestic early warning indicators are flashing ; and "G" denotes the position in a fan chart constructed when the global dimension of the crisis is taken into account by using information on both domestic and international early warning indicators are flashing.

Figure 3. Fan charts for the Spring 2008 GDP growth forecast for the United States

(A) Based on historical forecast errors, ignoring early warning indicators

(B) Allowing for alarm from domestic early warning indicator

(C) Allowing for alarms from domestic and international early warning indicators

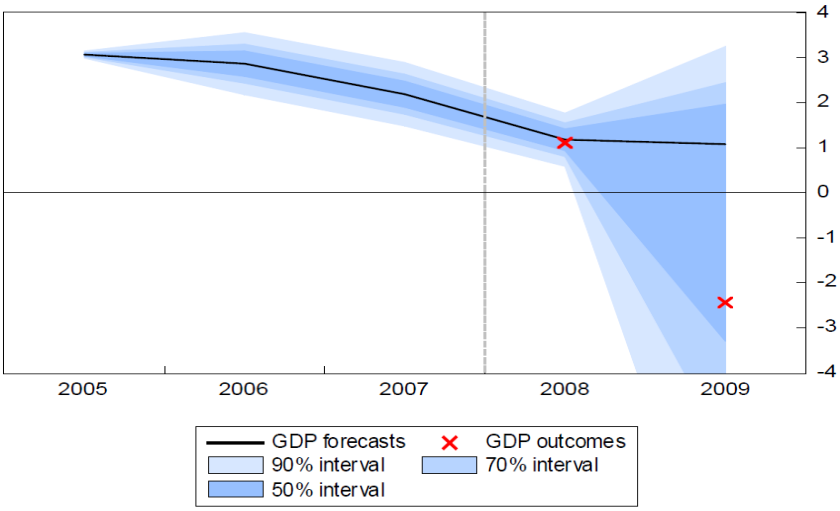

Note: Shaded blue areas show successively the 50%, 70% and 90% prediction intervals. The solid black line is the outturn up to 2007 and the projection for 2008 and 2009, as reported in the May 2008 *Economic Outlook*. The red crosses show the outturn according to the *Economic Outlook* published in the year following the first outturn data. The prediction intervals around the historical growth path reflect the extent to which historical estimates of GDP growth are subsequently revised.

Figure 4. Fan charts for the Spring 2008 GDP growth forecast for Japan

(A) Based on historical forecast errors, ignoring early warning indicators

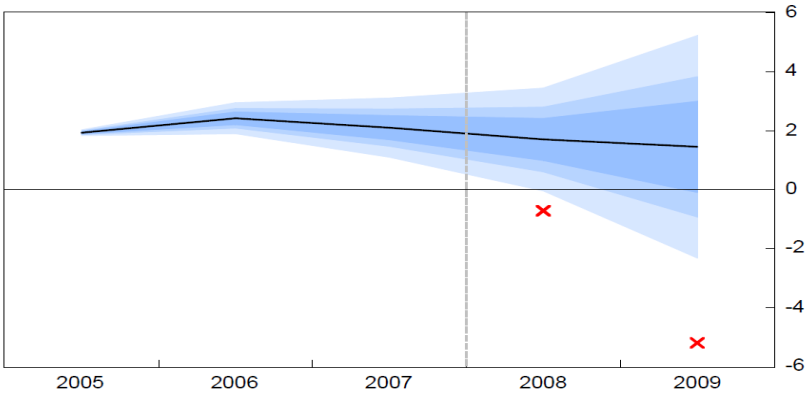

(B) Allowing for alarms from domestic and international early warning indicators

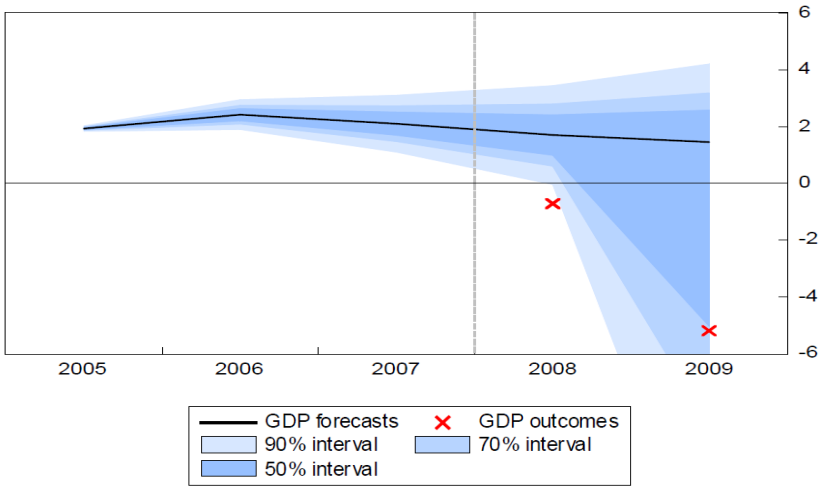

Note: Shaded blue areas show successively the 50%, 70% and 90% prediction intervals. The solid black line is the outturn up to 2007 and the projection for 2008 and 2009, as reported in the May 2008 *Economic Outlook*. The red crosses show the outturn according to the *Economic Outlook* published in the year following the first outturn data. The prediction intervals around the historical growth path reflect the extent to which historical estimates of GDP growth are subsequently revised.

Figure 5. Fan charts for the Spring 2008 GDP growth forecast for Germany

(A) Based on historical forecast errors, ignoring early warning indicators

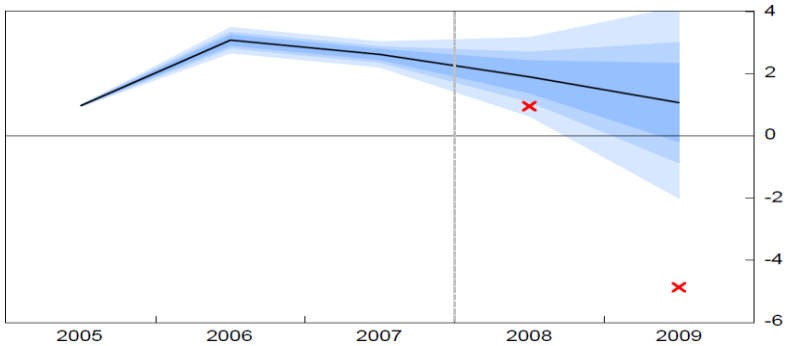

(B) Allowing for alarms from domestic and international early warning indicators

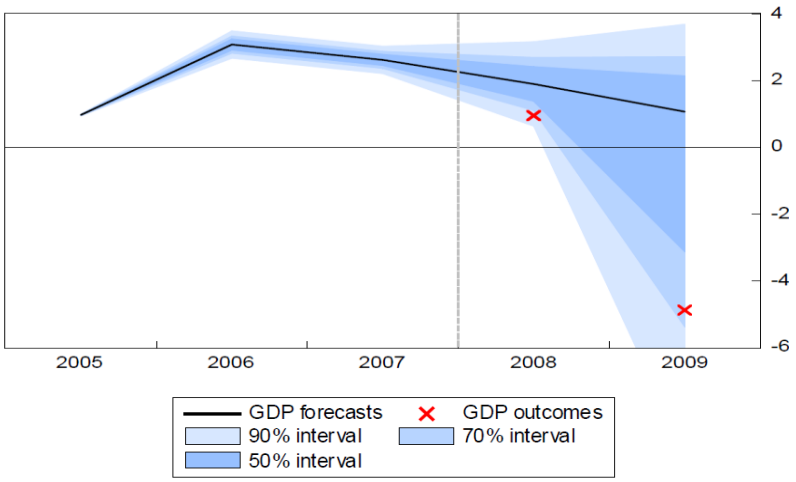

Note: Shaded blue areas show successively the 50%, 70% and 90% prediction intervals. The solid black line is the outturn up to 2007 and the projection for 2008 and 2009, as reported in the May 2008 *Economic Outlook*. The red crosses show the outturn according to the *Economic Outlook* published in the year following the first outturn data. The prediction intervals around the historical growth path reflect the extent to which historical estimates of GDP growth are subsequently revised.

Figure 6. Fan charts for the Spring 2008 GDP growth forecast for France

(A) Based on historical forecast errors, ignoring early warning indicators

(B) Allowing for alarm from domestic early warning indicator

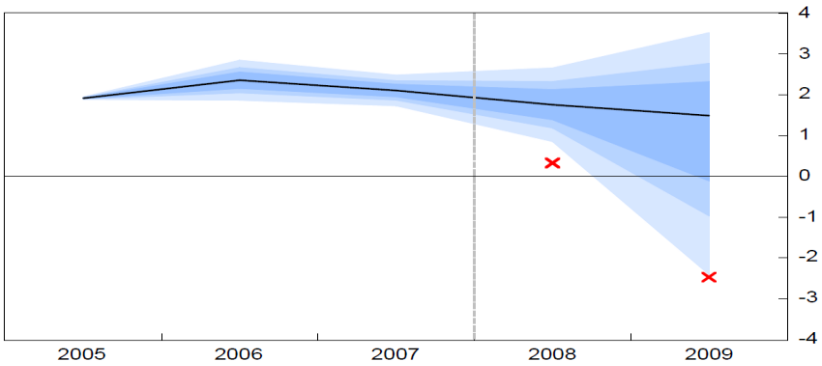

(C) Allowing for alarms from domestic and international early warning indicators

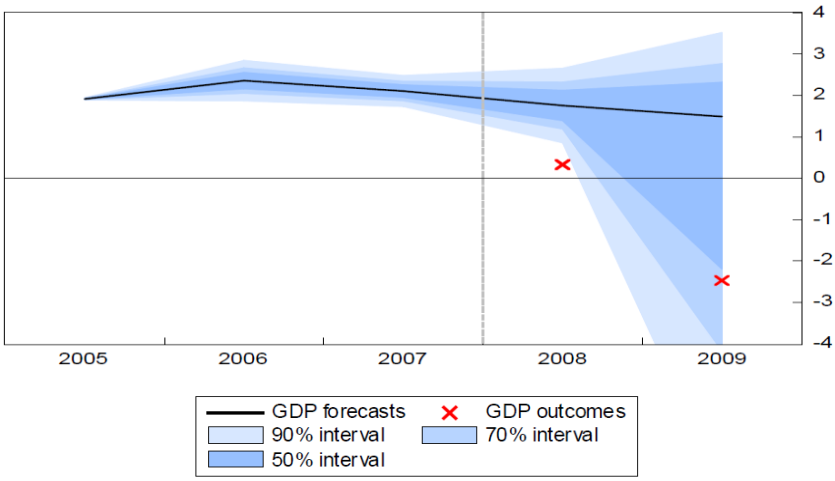

Note: Shaded blue areas show successively the 50%, 70% and 90% prediction intervals. The solid black line is the outturn up to 2007 and the projection for 2008 and 2009, as reported in the May 2008 *Economic Outlook*. The red crosses show the outturn according to the *Economic Outlook* published in the year following the first outturn data. The prediction intervals around the historical growth path reflect the extent to which historical estimates of GDP growth are subsequently revised.

Figure 7. Fan charts for the Spring 2008 GDP growth forecast for the United Kingdom

(A) Based on historical forecast errors, ignoring early warning indicators

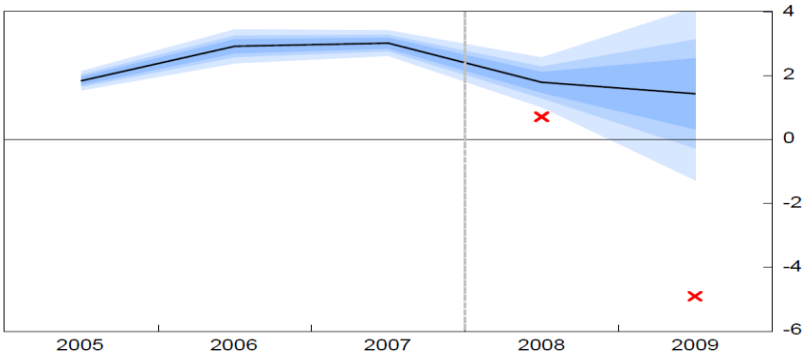

(B) Allowing for alarm from domestic early warning indicator

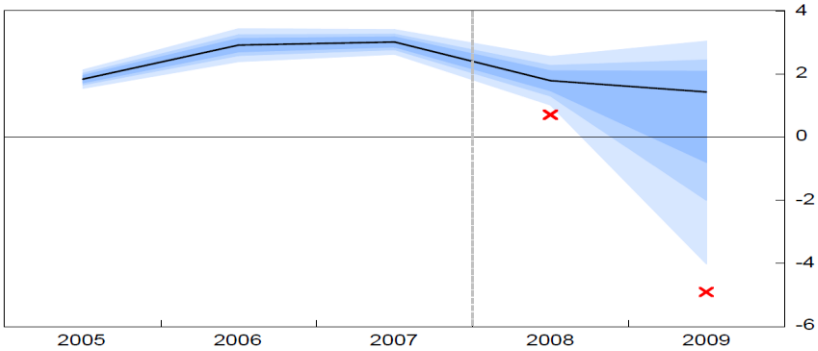

(C) Allowing for alarms from domestic and international early warning indicators

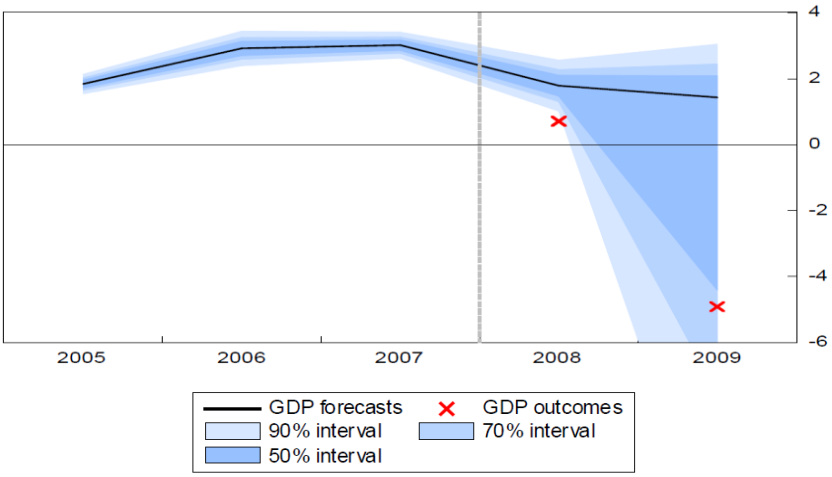

Note: Shaded blue areas show successively the 50%, 70% and 90% prediction intervals. The solid black line is the outturn up to 2007 and the projection for 2008 and 2009, as reported in the May 2008 *Economic Outlook*. The red crosses show the outturn according to the *Economic Outlook* published in the year following the first outturn data. The prediction intervals around the historical growth path reflect the extent to which historical estimates of GDP growth are subsequently revised.

Figure 8. Fan charts for the Spring 2008 GDP growth forecast for Italy

(A) Based on historical forecast errors, ignoring early warning indicators

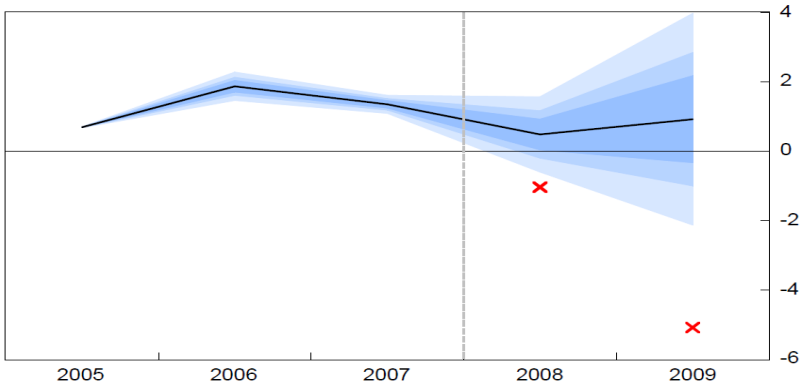

(B) Allowing for alarm from domestic early warning indicator

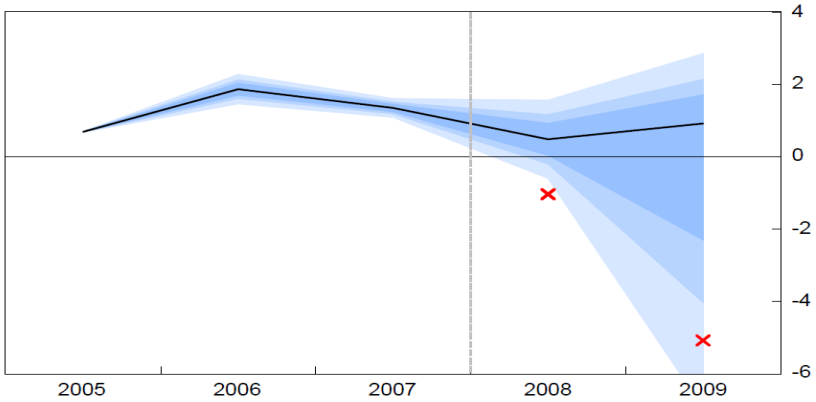

(C) Allowing for alarms from domestic and international early warning indicators

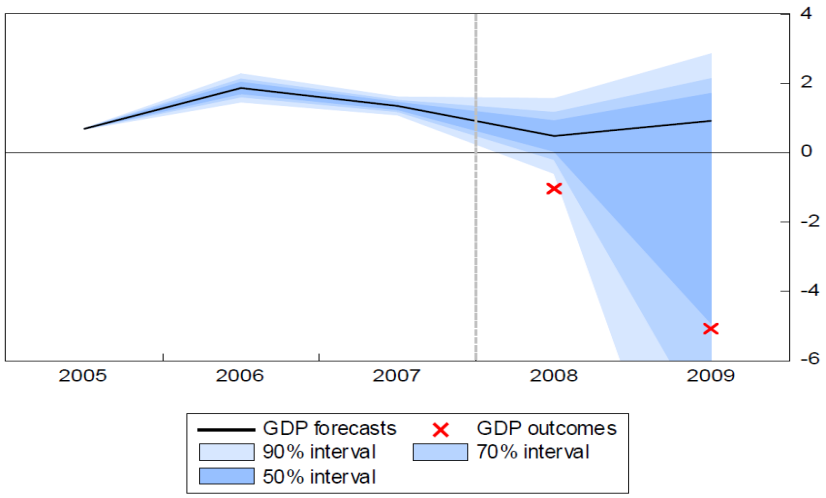

Note: Shaded blue areas show successively the 50%, 70% and 90% prediction intervals. The solid black line is the outturn up to 2007 and the projection for 2008 and 2009, as reported in the May 2008 *Economic Outlook*. The red crosses show the outturn according to the *Economic Outlook* published in the year following the first outturn data. The prediction intervals around the historical growth path reflect the extent to which historical estimates of GDP growth are subsequently revised.

Figure 9. Fan charts for the Spring 2008 GDP growth forecast for Canada

(A) Based on historical forecast errors, ignoring early warning indicators

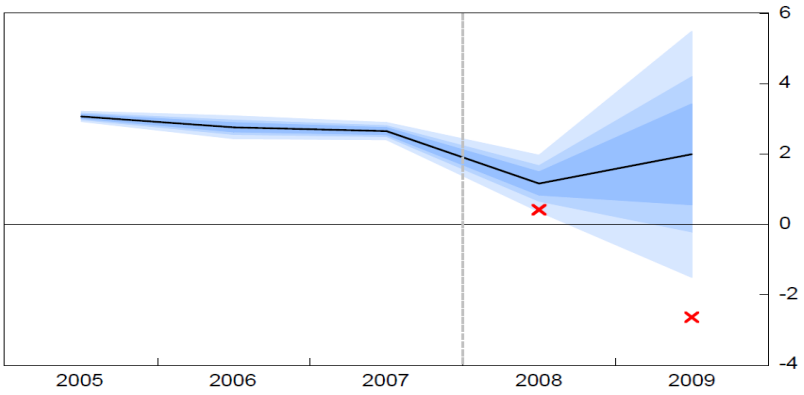

(B) Allowing for alarm from domestic early warning indicator

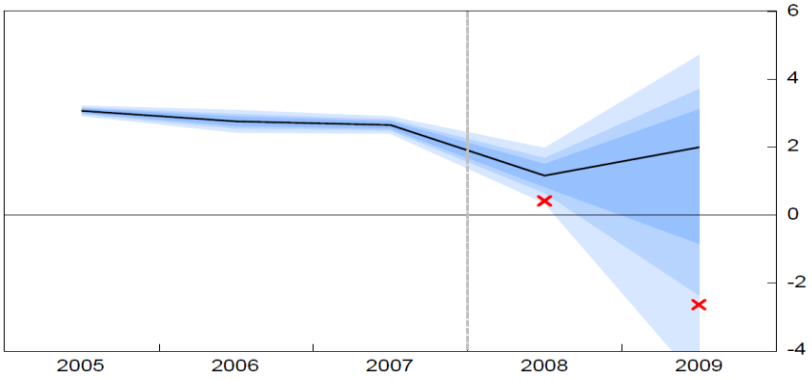

(C) Allowing for alarms from domestic and international early warning indicators

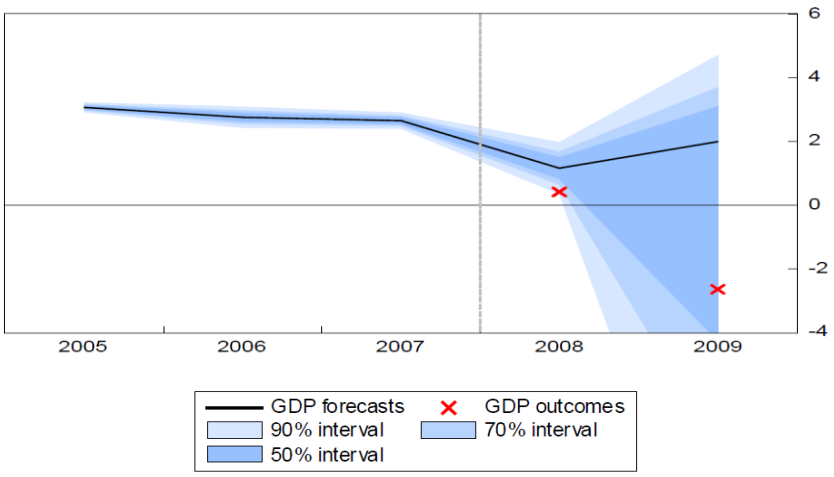

Note: Shaded blue areas show successively the 50%, 70% and 90% prediction intervals. The solid black line is the outturn up to 2007 and the projection for 2008 and 2009, as reported in the May 2008 *Economic Outlook*. The red crosses show the outturn according to the *Economic Outlook* published in the year following the first outturn data. The prediction intervals around the historical growth path reflect the extent to which historical estimates of GDP growth are subsequently revised.

5.3 Sensitivity testing

32. The previous exercise is open to criticism that the estimations on which the fan charts over the global financial crisis are calibrated already include the experience of the large forecast error made for 2009. This approach was deliberately taken because excluding the experience of the global financial crisis would imply discarding much information, particularly regarding synchronised severe downturns in the major OECD economies, for which there is otherwise limited experience over the available sample period.

33. It is, however, possible to conduct a sensitivity exercise to effectively exclude the crisis from the parameterisation of the fan charts, by re-running the regressions of forecast errors on early warning indicators, but this time adding a dummy which takes the value unity for the one-year ahead forecast error of GDP growth in 2009 and zero elsewhere. The coefficient on this dummy variable is mostly large and negative and tends to reduce the explanatory power of the early warning indicators (Table 6). Nevertheless, for the most part the coefficients on the domestic early warning indicators remain negative, large and significant, and whereas the magnitude of the coefficient on the international indicators are more noticeably reduced and are less significant, the coefficients for most countries remain quite large, so potentially adding to any skew in the fan charts.

Table 6. Sensitivity of preferred equations to inclusion of a global financial crisis dummy

Dependent variable: year-ahead forecast error, sample period: 1981-2015

Country	Constant[3]	Domestic indicator[1]		International indicator[2]			Financial crisis dummy[3,4]
		Coefficient[3]	Variable	Coefficient[3]	Other countries	Threshold	
United States	0.37	-1.25	(HPY+PCR)	-1.85	Other G7	3	-0.46
	-	**		**			-
Japan	0.39	-2.18	PCR	-0.92	Other G7	3	-4.27
	-	***		-			-
Germany	0.01	-1.16	BCR	-0.16	FRA, GBR, ITA, USA	2	-5.50
	-	**		-			**
France	-0.26	-0.48	BCR	-0.83	DEU, GBR, ITA, USA	2	-1.56
	-	-		-			-
United Kingdom	0.36	-1.19	PCR	-1.21	DEU, FRA, ITA, USA	2	-3.10
		**		*			*
Italy	-0.22	-1.03	(HPY+BCR)	-0.51	DEU, FRA, GBR, USA	2	-2.70
	-	***		-			-
Canada	0.13	-1.68	BCR	-1.67	Other G7, USA x 3	3	3.61
	-	**		***			-

Notes:

1. Domestic indicator variables all take the value zero or one: HPY is related to the house-price-to disposable income ratio; PCR is related to the ratio of total private credit to GDP; BCR is related to the ratio of bank credit to GDP; RHP is related to real house prices. The definition of when each of these variables switch between zero and one is explained in Table 1.

2. The 'international indicator' is a variable measuring the number of other countries, defined in the penultimate column, which exceeds a threshold number of countries shown in the final column.

3. Statistical significance of coefficients at the 1%, 5% and 10% levels are denoted by ***, ** and *, respectively. A coefficient which is not significant at the 10% level is denoted by "-".

4. The financial crisis dummy takes the value of 1 for the forecast error made in Spring 2008 for growth in calendar year 2009 and zero elsewhere.

34. For the purpose of the sensitivity test, the revised coefficients from this additional regression, ignoring the coefficient on the dummy variable, can then be used to recalculate the fan charts for the May 2008 *Economic Outlook* projections of GDP growth for the G7. By ignoring the coefficient on the dummy variable, any exceptional information from the financial crisis is essentially being discarded in the

parameterisation of the fan chart. This alternative parameterisation of the fan chart tends to push the 2009 outturn more towards the periphery of the fan chart as summarised in Table 7, although for the United States, Canada and France the outturn remains within the same prediction interval. Moreover, the 2009 outturn for all countries, with the exception of Germany, remains at least within the 90% prediction interval. The reason the 2009 outturn is well outside the fan chart for Germany is likely a reflection of the limited experience of severe downturns (and hence large forecast errors) for Germany in the sample period once the global financial crisis is excluded.

Table 7. Sensitivity of the fan chart position of the 2009 GDP outturn to a crisis dummy

Fan chart interval	USA	JPN	DEU	FRA	GBR	ITA	CAN
0-50%	G d	G					G d
50-70%			G	G d	G	G	
70-90%		d			d	d	
90-99%							
99%+			d				

Note: The table shows the prediction interval in which the outturn for GDP growth in 2009 is located in fan charts constructed around the forecast made in the May 2008 *Economic Outlook,* according to alternative parametrisations of the fan chart: "G" denotes the position in a fan chart constructed from domestic and international early warning indicators; "d" denotes the position in the fan chart when it is parameterised from the same regression, but with a dummy variable included.

6. False alarms and missed downturns

35. While the examples in the previous section illustrate the potential usefulness of using the early warning indicators to calculate fan charts which correctly identify future downturn risks, there will inevitably be false alarms when an early warning indicator flashes, but a downturn does not occur within the immediate forecast horizon. The pattern suggested by the early warning indicators used in the present study (see Table 8), suggests that on many of these occasions there will be a series of further alarms from the indicators as tensions in the housing market or credit growth continue to build until the bubble bursts and there is an eventual downturn. Most notably, in the run-up to the global financial crisis, early warning alarms were sounding in 2005, but they became steadily louder -- as measured by the combined skew across the fan charts for all G7 countries -- until 2008, by which time when they were deafening (Figure 10). This bears out Dornbusch's observation that *"The crisis takes a much longer time coming than you think, and then it happens much faster than you would have thought"*. It also raises the issue as to whether the "false alarms" in, say, 2006 might have been more useful from a policy perspective than a "genuine" alarm in 2008, by which time taking remedial policy action would presumably have been too late. In any case, the difficulty of precisely predicting the timing of any downturn, suggests that the early warning indicators are better employed in designing fan charts that identify potential risks rather than in adjusting the baseline forecast.

Table 8. Summary table of early warning alarms and downturns for the G7, 1981-2015

Country	Indicator	1981	1982	1983	1984	1985	1986	1987	1988	1989	1990	1991	1992	1993	1994	1995	1996	1997	1998	1999	2000	2001	2002	2003	2004	2005	2006	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016
United States	HPY+PCR		X									X																	X	X							
Japan	PCR													X					X	X									X	X							
Germany	BCR													X									X	X						X							
France	BCR													X															X	X							
UK	PCR	X										X																	X	X							
Italy	HPY+BCR													X										X					X	X			X	X			
Canada	PCR		X								X	X																		X							

X = fall in GDP per capita = Early warning indicator flashing yyyy = Warnings in 4+ countries

Note: The red cells indicate years in which early warning indicators for particular countries are flashing signalling the risk of a downturn in the following year. Crosses "X's" indicate years in which there is an absolute fall in GDP per capita. Years highlighted in pink/mauve indicate years in which early warning indicators are flashing in more than four countries.

Figure 10. The combined skew across all G7 fan charts

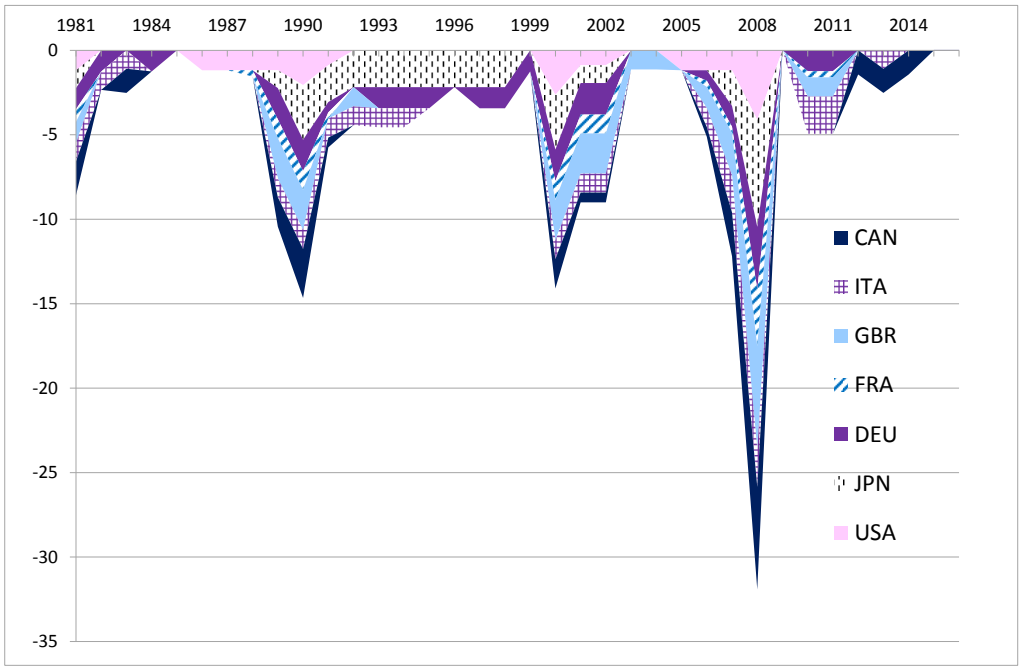

Note: The chart shows the skew across all G7 countries from both domestic and international early warning indicators calculated on the basis of the coefficients in Table 4, using equation (1) and the method described in the text.

36. There will also be severe downturns which occur without prior warnings from domestic housing or credit indicators. However, in many of these cases -- for example, Japan and Germany during the global financial crisis of 2008/9 -- the downturn is predicted by simultaneous warning indicators in other G7 countries and so captured by the international early warning indicator. Nevertheless, there will be other occasions which are unrelated to both domestic and international housing and credit indicators; for example, the severe downturn in the UK in 1980/81 was mainly driven by decisions to target the money supply in order to reduce inflation, which led to high interest rates and a marked appreciation of sterling (Buiter and Miller, 1981).

7. Further possible developments

37. There are a number of directions in which the current work could be developed.

1. *To test the robustness of the results, the method could be extended to other forecasts*. In the first instance, it could most easily be applied to the corresponding *OECD Economic Outlook* forecasts of G7 GDP growth made in the Autumn, rather than the Spring. Given the different time horizons, the parameterisation would be different, but the expectation would be that the same early warning indicators that are used in parameterising the fan charts for the Spring forecasts ought to be useful for the Autumn forecasts. The method could also be extended to non-G7 OECD countries for which there is a long series of past forecast vintages, although applying it to non-OECD countries would be difficult because the number of published forecasts is more limited. Finally, given that all forecasters are poor at identifying downturns, it would be interesting to apply the same approach to forecasts from different institutions.[13]

2. *The reliability of the early warning indicators could be improved by customising them to make them more country-specific*. The work by HR, from which the current paper draws inspiration, aimed at ranking early warning indicators according to their performance across *all* OECD countries, adopting the same functional form and alert thresholds across all countries. The application here is different in trying to identify the one or two early warning indicators which work best for individual countries. However, it is clear that some variables that work well for a majority of countries are likely to work less well for others; for example, housing market variables are likely to work well for the United Kingdom, but less well for Germany. Moreover, as was discussed in Section 3 in relation to UK housing and credit variables, the optimal functional form of any early warning indicator may differ from country to country; for example, if a particular ratio has a long-standing historical upward trend, it is likely that a rate-of-change functional form will be more successful than an adjusted level formulation in predicting past downturns.[14] In any case, if such early warning indicators are to be used "live" to parameterise any forecast fan chart, it would be essential that the forecasters with immediate responsibility for the forecast had confidence in the reliability of the early warning indicators.[15]

3. *Alternative forms of the international early warning indicator could be investigated*. It is apparent from the results described in this paper, that in some periods the international warning indicator plays a dominant role, which suggest there should be scope to investigate other forms of such an indicator. For example, by only taking some combination of early warning alarms in

13. Indeed, there is a striking synchronicity of year-ahead forecast errors both between different international institutions and with private sector forecasts [see for example Turner (2016), especially Figure 3; Abreu (2011), especially Figures 1 and 2; and Loungani (2000), especially Figures 8 to 10], so that typically the difference between forecasts for year-ahead GDP growth is relatively small compared to the average absolute error that all forecasters make

14. Another possible modification would be to relax the threshold for what HR define as a "severe downturn", namely a decline in quarterly GDP per capita from peak-to-trough of 3½ per cent, since it is apparent that this excludes some downturns which have been associated with large forecast errors. This would likely imply an increase in the number of downturns identified over the period since 1970, from an average of 2-3 downturns to 3, 4 or 5 downturns, per country. Increasing the number of downturn episodes might help to better identify the early warning indicators which are associated with large historical forecast errors.

15. At the OECD, each country forecast is under the immediate responsibility of individual Country Desk economists, although there is considerable degree of centralised oversight and input [see Turner (2016) for a description of the forecast process].

other G7 countries, the international indicator for Japan used in the current paper would miss the effect of the Asia crisis in 1998 which led to a downturn (and a large forecast error) in Japan.

4. *Finally, the procedure described here could be used as an initial starting point to parameterise a fan chart, which could then be subsequently modified according to the forecaster's judgement*. Thus, while the process described in this paper is mechanical, it could provide an initial estimate for the degree and asymmetry of uncertainty associated with a forecast, which could then be subsequently judgementally revised according to how current circumstances differ from past historical experience or a previous forecast fan chart, similar to the way described by Blix and Sellin (1998).

REFERENCES

Abreu, I. (2011), "International Organisations' vs Private Analysts' Forecasts: An Evaluation", Banco de Portugal Working Papers, 20/2011, July.

Bank of England (2008), "Inflation Report", May.

Bank of England (2005), "Assessing the MPC's Fan Charts", *Bank of England Quarterly Bulletin*. http://www.bankofengland.co.uk/publications/Documents/quarterlybulletin/qb050301.pdf

Behn, M. C. Detken, T., Peltonen and W. Schnudel (2013), "Setting Countercyclical Capital Buffers Based on Early Warning Models: Would it Work?", *ECB Working Paper Series*, No. 1604, https://www.ecb.europa.eu/pub/pdf/scpwps/ecbwp1604.pdf.

Blix, M. and P. Sellin (1998), "Uncertainty Bands for Inflation Forecasts", *Sveriges Riksbank Working Paper Series*, No. 65.

Boero, G., J. Smith and K. F. Wallis (2015), "The Measurement and Characteristics of Professional Forecasters' Uncertainty", *Journal of Applied Econometrics*, Vol. 30: 1029–1046. doi: 10.1002/jae.2400.

Borio, C. and P. Lowe (2002), Assessing the Risk of Banking Crisis", *BIS Quarterly Review*, December 2002, pp. 43-52, http://www.bis.org/publ/qtrpdf/r_qt0903e.pdf.

Britton, E., P. Fisher and J. Whitley (1997), "Asymmetry, Risks and a Probability Distribution of Inflation", *Bank of England Quarterly Bulletin*, http://www.bankofengland.co.uk/archive/Documents/historicpubs/qb/1998/qb980101.pdf

Buiter, W. and M. Miller (1981), "The Thatcher Experiment: The First Two Years", *Brookings Papers on Economic Activity*, Vol. 2.

Bussiere, M. and M. Fratzscher (2006), "Towards a New Early Warning System of Financial Crises", *Journal of International Money and Finance*, Vol. 25/6, pp. 953-973. DOI: http://dx.doi.org/10.1016/10.1016/j.jimonfin.2006.07.007.

Casella, G. and R. Berger (2002), *Statistical Inference*, Second Edition, Duxbury Advanced Series.

Fildes, R. and H. Stekler, (2002), "The State of Macroeconomic Forecasting", *Journal of Macroeconomics*, 24(2), pp. 435-468.

Hermansen, M and O. Röhn (2016), "Economic Resilience: The Usefulness of Early Warning Indicators in OECD Countries", *OECD Journal: Economic Studies*, No. 1, Vol. 2016, Issue, 1, pp. 9-35. OECD Publishing, Paris. DOI: http://dx.doi.org/10.1787/eco_studies-2016-5jg2ppjrd6r3

John, S. (1982), "The Three-Parameter Two-Piece Normal Family of Distributions and its Fitting", *Communications in Statistics - Theory and Methods*, 11(8), pp. 879-885. DOI: http://dx.doi.org/10.1080/03610928208828279

Kaminsky, G. S., S. Lizondo and C. Reinhart (1998), "Leading Indicators of Currency Crises", *IMF Staff Paper*, Vol. , 45, No. 1, http://www.imf.org/external/pubs/ft/staffp/1998/03-98/kaminsky.htm.

Loungani, P. (2001), "How Accurate are Private Sector Forecasts? Cross-Country Evidence from Consensus Forecasts of Output Growth", *International Journal of Forecasting*, 17(3), pp. 410-432.

OECD (1990), *OECD Economic Outlook*, No. 47, Vol. 1990/1, June, OECD Publishing, Paris. DOI: http://dx.doi.org/10.1787/eco_outlook-v1990-1-en

OECD (2008), *OECD Economic Outlook*, No. 83, Vol. 2008/1, June, OECD Publishing, Paris. DOI: http://dx.doi.org/10.1787/eco_outlook-v2008-1-en

Pain, N. and C. Lewis (2014), "Lessons Learned from OECD Forecasts During and after the Financial Crisis", *OECD Journal: Economic Studies*, No. 5, Vol. 2104, Issue, 1, pp. 9-39, OECD Publishing, Paris. DOI: http://dx.doi.org/10.1787/eco_studies-2014-5jxrcm2glc7j

Stark, T. (2010), "Realistic Evaluation of Real-Time Forecasts in the Survey of Professional Forecasters", Research RAP Special Report, https://ideas.repec.org/a/fip/fedprr/y2010imayx1.html.

Sveriges Riksbank (2007), "Calculation Method for Uncertainty Bands", Box in Chapter 1 of *Monetary Policy Report* 2007/1.

Tay, A. S. and K. F. Wallis (2000), "Density Forecasting: A Survey", *Journal of Forecasting*, 19, pp. 235-254.

Wallis, K. F. (2014), "The Two-Piece Normal, Binormal, or Double Gaussian Distribution: Its Origin and Rediscoveries", *Statistical Science*, Vol. 29, No. 1, pp. 106-112, https://projecteuclid.org/euclid.ss/1399645739

Wallis, K. F. (2004), "An Assessment of Bank of England and National Institute Forecast Uncertainties", *National Institute Economic Review*, No. 189, July.

Wallis K F. (1999), "Asymmetric Density Forecasts of Inflation and the Bank of England's Fan Chart", *National Institute Economic Review*, January.

World Bank (2016), "Quantifying Uncertainties in Global Growth Forecasts", *Special Focus 2* in *Global Economic Prospects*, June.

ANNEX A. OECD FORECAST ERRORS OF G7 GDP GROWTH

This Annex provides summary statistics on the performance of OECD forecasts of GDP growth for each of the G7 economies published in the May or June ("Spring") edition of the OECD *Economic Outlook*, distinguishing between forecasts for the current year and the year-ahead. Forecast errors are defined throughout as the outturn less the forecast (so that a negative error corresponds to an over-prediction), and the outturn for a particular year is defined as that published in the Spring *Economic Outlook* in the following year.

Table A1. Current-year forecast performance over time, 1971-2016

	USA	JPN	DEU	FRA	UK	ITA	CAN	Average
A. Mean absolute error								
Full sample	0.55	1.11	0.69	0.55	0.58	0.66	0.70	0.69
1971-79	1.10	1.65	0.95	0.73	0.99	1.12	0.58	1.02
1980-89	0.52	1.23	0.83	0.70	0.59	0.58	1.14	0.80
1990-99	0.57	1.03	0.56	0.44	0.58	0.55	0.89	0.66
2000-16	0.28	0.79	0.56	0.44	0.36	0.52	0.40	0.48
B. RMSE								
Full sample	0.78	1.57	0.87	0.78	0.79	0.95	0.94	0.95
1971-79	1.35	2.41	1.04	1.24	1.17	1.60	0.85	1.38
1980-89	0.63	1.68	1.01	0.79	0.85	0.80	1.43	1.03
1990-99	0.74	1.21	0.66	0.52	0.73	0.66	0.99	0.78
2000-16	0.36	1.06	0.77	0.55	0.47	0.66	0.49	0.62

Table A2. One-year-ahead forecast performance over time, 1982-2016

	USA	JPN	DEU	FRA	GBR	ITA	CAN	Average
A) Mean absolute error								
(a) Full sample	1.24	1.78	1.44	1.06	1.11	1.48	1.54	1.38
(b) 1982-1999	1.37	1.89	1.21	1.00	1.09	1.26	1.74	1.37
(c) 2000-16	1.08	1.66	1.71	1.13	1.14	1.74	1.30	1.39
Ratio (c)/(b)	79%	88%	141%	113%	104%	139%	75%	105%
B) RMSE								
(a) Full sample	1.52	2.29	1.87	1.42	1.64	1.92	2.13	1.83
(b) 1982-1999	1.65	2.23	1.56	1.33	1.45	1.55	2.42	1.74
(c) 2000-16	1.36	2.36	2.19	1.52	1.85	2.27	1.74	1.90
Ratio (c)/(b)	83%	106%	141%	114%	128%	146%	72%	113%

Table A3. Year-ahead forecast errors and outturns of GDP growth

Error equals outturn less forecast, percentage points

Year (t)	USA		JPN		DEU		FRA		GBR		ITA		CAN	
	Error	Outturn	Error	Outturn	Error	Outturn	Error	Outturn	Error	Outturn	Error	Outturn	Error	Outturn
1982	-2.70	-1.70	-1.00	3.00	-3.10	-1.10	-0.05	1.70	0.95	1.20	-2.05	-0.30	-7.20	-4.80
1983	1.15	3.40	-1.00	3.00	-1.95	1.30	-1.80	0.70	1.35	3.10	-3.70	-1.20	3.00	3.00
1984	2.30	6.80	2.30	5.80	0.85	2.60	1.20	1.70	0.15	2.40	0.60	2.60	-0.04	4.70
1985	-0.30	2.20	0.85	4.60	-0.35	2.40	-0.55	1.20	0.95	3.20	0.05	2.30	2.50	4.50
1986	-0.25	2.50	-2.00	2.50	-0.35	2.40	0.00	2.00	0.20	2.70	0.95	2.70	-0.15	3.10
1987	-0.85	2.90	1.20	4.20	-1.30	1.70	-0.60	1.90	2.25	4.50	0.35	3.10	0.65	3.90
1988	1.15	3.90	3.70	5.70	1.40	3.40	1.40	3.40	1.45	3.70	1.40	3.90	1.76	4.50
1989	0.50	3.00	1.15	4.90	2.25	4.00	1.95	3.70	0.05	2.30	0.95	3.20	-0.34	2.90
1990	-1.35	0.90	1.35	5.60	1.75	4.50	0.05	2.80	-1.40	0.60	-1.25	2.00	-1.34	0.90
1991	-3.19	-0.71	0.47	4.46	-0.24	3.14	-1.61	1.25	-4.09	-2.17	-1.71	1.44	-4.09	-1.53
1992	-1.02	2.11	-2.17	1.33	-0.24	1.97	-1.37	1.35	-2.23	-0.62	-1.77	0.93	-2.17	0.89
1993	-0.58	3.01	-2.95	0.10	-3.58	-1.25	-3.54	-0.92	-0.71	1.90	-2.78	-0.66	-1.92	2.39
1994	0.94	4.08	-2.72	0.55	1.50	2.87	1.20	2.67	0.85	3.79	0.44	2.18	0.02	4.51
1995	-0.92	2.04	-1.84	0.85	-0.61	1.95	-0.73	2.21	-0.75	2.41	0.40	3.02	-2.08	2.24
1996	0.10	2.43	1.24	3.56	-1.34	1.36	-1.72	1.46	-0.82	2.14	-2.24	0.69	-1.92	1.47
1997	1.72	3.76	-1.47	0.90	-0.21	2.18	0.08	2.44	0.22	3.25	-0.84	1.51	0.37	3.82
1998	1.92	3.88	-5.67	-2.82	0.00	2.75	0.39	3.20	-0.59	2.06	-0.37	1.43	-0.36	2.99
1999	2.09	4.15	-1.07	0.26	-1.42	1.47	0.10	2.93	0.30	2.06	-1.31	1.43	1.14	4.19
2000	2.99	4.99	1.70	1.70	0.62	2.95	0.60	3.25	1.47	3.04	0.70	2.91	1.93	4.68
2001	-1.80	1.19	-2.67	-0.43	-2.48	0.56	-0.91	2.03	-0.10	2.21	-1.35	1.79	-1.51	1.49
2002	-0.64	2.44	-0.79	0.34	-2.26	0.18	-1.51	1.16	-0.75	1.80	-2.10	0.37	0.21	3.37
2003	-0.42	3.12	2.34	2.68	-2.63	-0.10	-2.44	0.52	-0.54	2.22	-2.41	0.36	-2.32	1.71
2004	0.48	4.44	1.53	2.64	-0.72	1.03	-0.28	2.33	0.52	3.14	-1.46	0.97	-0.64	2.76
2005	-0.19	3.52	-0.03	2.72	-1.05	1.09	-1.11	1.45	-0.83	1.82	-1.75	0.11	-0.35	2.93
2006	-0.01	3.32	0.54	2.21	1.22	2.98	0.02	2.06	0.35	2.76	0.79	1.94	-0.39	2.75
2007	-0.96	2.19	-0.15	2.10	0.99	2.62	-0.12	2.11	0.13	3.02	0.03	1.35	-0.61	2.65
2008	-1.41	1.11	-2.79	-0.72	-1.23	0.95	-1.83	0.33	-1.77	0.71	-2.70	-1.04	-2.60	0.41
2009	-3.52	-2.44	-6.65	-5.20	-5.95	-4.88	-3.96	-2.47	-6.35	-4.92	-6.01	-5.09	-4.64	-2.65
2010	1.97	2.85	3.26	3.97	3.35	3.50	1.17	1.38	1.27	1.25	0.86	1.25	2.41	3.08
2011	-1.42	1.74	-2.79	-0.75	0.93	3.06	-0.37	1.70	-1.82	0.65	-0.98	0.54	-0.75	2.46
2012	-0.91	2.21	-0.15	2.00	-1.66	0.87	-2.08	0.02	-1.56	0.27	-4.01	-2.39	-1.01	1.84
2013	-0.75	1.88	0.07	1.54	-1.41	0.54	-0.94	0.30	-0.24	1.66	-1.49	-1.84	-0.60	2.01
2014	-0.39	2.39	-1.44	-0.09	-0.34	1.60	-0.61	0.17	1.29	2.81	-0.79	-0.35	0.16	2.44
2015	-1.10	2.43	-0.69	0.55	-0.61	1.45	-0.32	1.22	-0.34	2.33	-0.48	0.64	-1.55	1.17
2016	-1.23	1.53	-0.65	0.75	-0.58	1.74	-0.42	1.23	-0.31	2.03	-0.68	0.83	-1.04	1.22
Mean	-0.25		-0.54		-0.59		-0.59		-0.33		-1.05		-0.73	
RMSE	1.52		2.29		1.87		1.42		1.64		1.92		2.13	

Note: Forecast errors for GDP growth in calendar year t are calculated as the difference of the forecast for that year published in the May/June *OECD Economic Outlook* of the preceding year (t-1), and the outturn published in the Economic Outlook of the following year (t+1). Shaded cells in the error columns highlight forecast errors for which the magnitude exceeds twice the RMSE of all forecast errors for that country, whereas shaded cells in the outturn column highlight outturns of less than -1.0%.

Table A4. Cross-country correlations in one-year ahead forecast errors, 1982-2016

	USA	JPN	DEU	FRA	GBR	ITA	CAN
USA	1.000						
JPN	0.347	1.000					
DEU	0.493	0.494	1.000				
FRA	0.575	0.422	0.826	1.000			
GBR	0.618	0.498	0.412	0.602	1.000		
ITA	0.494	0.449	0.794	0.856	0.634	1.000	
CAN	0.751	0.373	0.498	0.434	0.610	0.468	1.000

ANNEX B. THE TWO-PIECE NORMAL DISTRIBUTION AND PREDICTION INTERVALS

The two-piece normal distribution

The two-piece, or split, normal distribution has come to prominence in macroeconomics because of its use by central banks, especially the Bank of England (Britton et al, 1998; Bank of England, 2005) and Riksbank (Blix and Sellin, 1998), in constructing asymmetric fan charts, but it has a much longer history in statistical theory (Wallis, 2013).

The two-piece normal distribution can be thought of as the combination of two halves of two different normal distributions with the same mode, μ, but different standard deviations, σ_1 and σ_2, with the respective probability density functions being re-scaled to have the same value at the mode. If $\sigma_1>\sigma_2$ then the distribution is skewed to the left so that the mean is less than the mode, whereas if $\sigma_1<\sigma_2$ then the distribution is skewed to the right with the mean exceeding the mode. The resulting two-piece normal distribution can be summarised in terms of the three parameters (μ, σ_1, σ_2).

In this paper whenever the two-piece normal is used, namely when early warning indicators suggest there is a downturn risk, the distribution is skewed to the left, so that $\sigma_1>\sigma_2$ and the mean is less than the mode. In the example illustrated in the chart below -- which has been computed to be typical of the distribution in the year-ahead forecast of GDP growth for a G7 country when a domestic early warning indicator is flashing -- the two-piece normal distribution is constructed by combining normal distributions with common mode equal to 2 (which would correspond to the baseline year-ahead forecast of GDP growth) and standard deviations of σ_1=3.0 and σ_2=1.3.

The choice of prediction intervals

Once the parameters of the two-piece normal distribution are determined, given that the distribution is asymmetric, there is still an issue as to how to determine prediction intervals; i.e. how to determine the values (a, b) defining a prediction interval of p%, where 0% < p%< 100%, so that the outcome in question (here year-ahead GDP growth) will occur within the interval (a, b) with probability p%.

An intuitive choice would seem to be *central prediction intervals* so that the prediction interval covers the stated proportion in the centre of the distribution, with the remaining probability split equally between the upper- and lower-tails outside this interval. This is, however, not the practice of central banks such as the Bank of England and Riksbank, which instead adopt *highest probability density* (HPD) intervals that coincide with choosing the shortest interval (b-a) for any given p% (Casella and Berger, 2002). While these two approaches give the same intervals for symmetric distributions, they are different when the distribution is asymmetric. Wallis (1999) argues for replacing the Bank of England intervals with central intervals on the grounds that many readers of the *Inflation Report* may wrongly interpret the 90% interval as implying a 5% probability that the outcome (inflation or GDP growth) will exceed the upper limit of the fan chart and also that it is difficult to justify the Bank's choice of intervals from a loss-function perspective.

An important difference between the two approaches is that as the prediction interval p% is progressively reduced, in the case of central prediction intervals the distribution collapses to the *median*, whereas in the case of HPD intervals it collapses to the *mode*. This difference can be visually striking as the baseline forecast (the mode) may no longer appear in the darkest shaded intervals of the fan chart if central prediction intervals are adopted instead of HPD intervals.

In terms of the previous example illustrated in Figure B1, a 90% central prediction interval covers the interval (-3.40, 3.80) with upper- and lower-tail probabilities of exactly 5% (Figure B2), whereas the 90% HPD prediction interval covers the interval (-2.94, 4.14) with a lower-tail probability of 7.0% and an upper-tail probability of 3.0% (Figure B3). Moreover, whereas the mode is, by construction, always within the HPD p% prediction interval for all values of p%, the mode lies outside the central prediction intervals as p% is reduced below 40%, as illustrated for a 30% central prediction interval in Figure B2. In terms of the associated forecast fan charts, this implies that the baseline forecast (the mode) will always be within the darkest (i.e. tightest) prediction interval shown when HPD prediction intervals are used, but this is not guaranteed if central prediction intervals are used.

Figure B1. The probability density function of an illustrative two-piece normal distribution

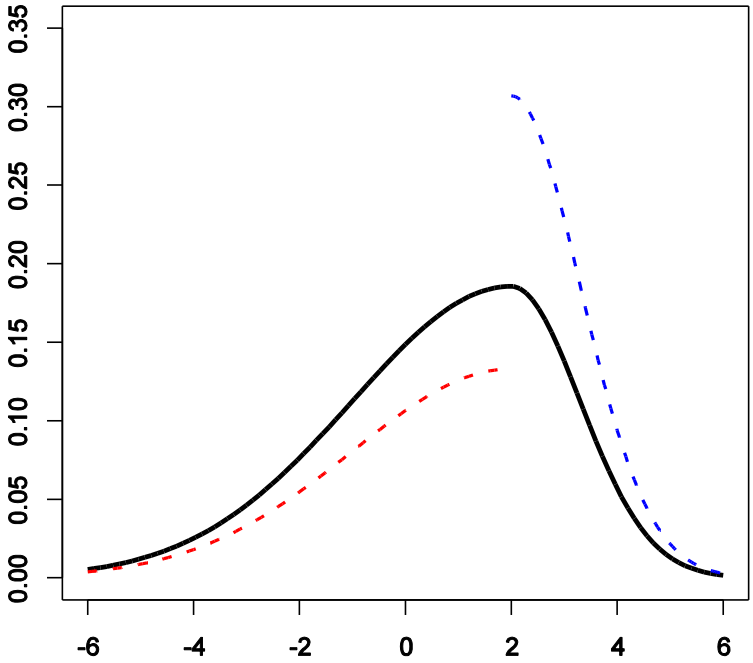

Note: The two-piece normal probability density function (pdf) represented by the solid black line is formed by re-scaling and combining two normal pdfs here represented by the dashed curves; the left-hand-side dashed curve is a normal pdf with mode of 2 and standard deviation of 3.0, the right-hand-side dashed curve has the same mode with a standard deviation of 1.3. The parameterisation has been chosen to roughly reflect the pdf of a year-ahead forecast of GDP growth for a typical G7 country when a domestic early warning indicator is flashing.

Figure B2. Central prediction intervals for an illustrative two-piece normal distribution

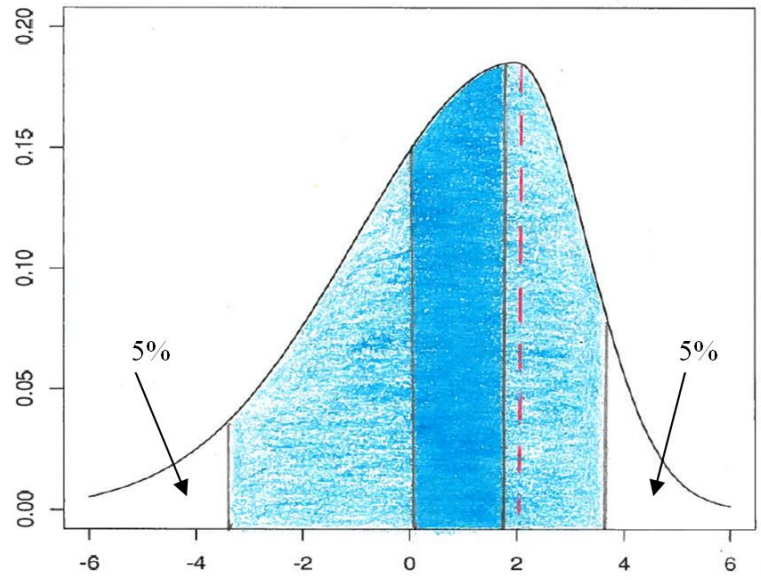

Note: The illustrative two-piece normal distribution has the same underlying parameters as that described in Figure B1. The lighter shaded area corresponds to a 90% central prediction interval, within which the darker shaded area corresponds to the 30% central prediction interval. The vertical red dashed line corresponds to the mode.

Figure B3. Highest probability density prediction intervals for an illustrative two-piece normal distribution

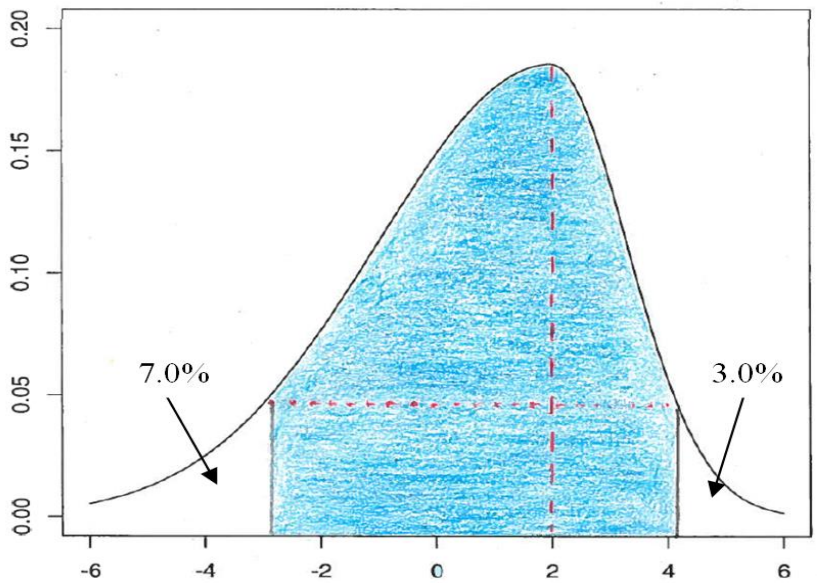

Note: The illustrative two-piece normal distribution has the same underlying parameters as that described in Figure B1. The shaded area corresponds to a 90% highest probability density prediction interval and the vertical red dashed line corresponds to the mode.