BANK OF ENGLAND

# Quarterly Bulletin

2017 Q2

## Topical article

Peering into the present: the Bank's approach to GDP nowcasting




ISSN 2399-4568

# Peering into the present: the Bank's approach to GDP nowcasting

By Nikoleta Anesti, Simon Hayes, Andre Moreira and James Tasker of the Bank's Monetary Analysis Directorate.(1)

- The Bank's GDP nowcast represents the Monetary Policy Committee's (MPC's) estimate of economic growth in the current quarter, before official data become available. The nowcast is informed by statistical models, but is ultimately judgemental, reflecting all available information.
- Users of nowcasts must be aware of the degree of accuracy that can be expected, as this varies across models and time. Models based on survey information tend to be more accurate early in the quarter, whereas high-frequency output data published by the ONS become more useful later.
- The MPC's *Inflation Report* nowcasts have been relatively accurate, with a root mean squared error of 0.3 percentage points over the past ten years — lower than a mechanical use of the models could have attained. GDP growth estimates have fallen within 0.1 percentage points of the MPC's expectation about half the time, although much larger surprises have occasionally occurred.

## Overview

The Monetary Policy Committee (MPC) has a keen interest in the current cyclical state of the economy as this has an important bearing on the appropriate stance of monetary policy. Because official data are published with a lag, the MPC produces estimates of the current rate of economic growth using a range of models and indicators in a process called 'nowcasting'. This article describes the Bank's approach to nowcasting and discusses some of the practicalities involved.

Statistical models that translate the range of available indicators into numerical predictions for GDP growth are key inputs to the MPC's nowcast. The Bank's toolkit currently includes three models, detailed in the article. Although such models are important, the MPC's actual nowcasts are not derived mechanically from the model outputs, but are ultimately a matter of judgement. That judgement is informed by knowledge of the models' strengths and weaknesses, and also reflects other information not captured by the models. Evidence suggests that this approach helps to improve accuracy.

Overall, the Bank's nowcasts have tended to track estimates of GDP growth relatively closely (**summary chart**). However, uncertainty is inherent in nowcasting and surprises are inevitable. We show that this uncertainty varies both across models and time, and so it is important to understand not only the degree of accuracy that can reasonably be expected in general, but also the confidence bands surrounding any given estimate.

**Summary chart** *Inflation Report* nowcasts against ONS preliminary GDP growth estimates

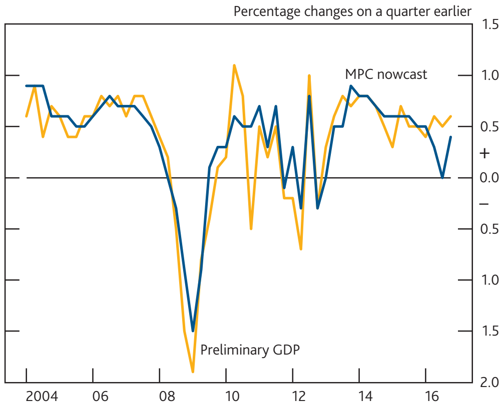


Sources: ONS and Bank calculations.

Occasionally large nowcast errors occur, and such events can contain important lessons. The sizable nowcast error in the third quarter of 2016, following the EU referendum, highlighted the extent to which indicators such as business surveys can overreact in some circumstances, rendering nowcasts based on them more uncertain than usual. The unprecedented nature of the event, however, means that the practical value of this lesson may be limited.

(1) The authors would like to thank Michal Stelmach for his help in producing this article.

## Introduction

Seeing into the future will always be difficult. But the nature of economic statistics is such that there is considerable uncertainty even about the present, as official data on the state of the economy are published with a lag. This has given rise to the industry of 'nowcasting' which, as the name suggests, involves estimating variables for the current period in advance of the publication of official data. This article focuses on the Bank's approach to nowcasting quarterly GDP growth, the first estimate of which is currently published around 25 days after the end of each quarter.(1)

GDP nowcasts are a key input into the Monetary Policy Committee's (MPC's) decisions, as the current cyclical state of the economy has an important influence on the appropriate stance of monetary policy. This article aims to shed light on the process by which the MPC arrives at its GDP nowcasts — published regularly in the *Inflation Report* — and also provides an update on the statistical toolkit developed by Bank staff to support that process.

The Bank's nowcasting toolkit comprises three core models: an industry model, a mixed-data sampling (MIDAS) model, and a dynamic factor model (DFM). These are detailed later in the article. It is useful to have multiple models because every approach is likely to be imperfect to some extent. An eclectic approach can help to avoid the particular pitfalls of any single model, making forecasts more robust.

Bank staff regularly brief the MPC on the latest economic data and what those mean for the model predictions. However, the MPC's actual nowcast does not flow mechanically from the model outputs. Rather, it involves the judgemental combination of both model nowcasts and other available information that the models might be unable to capture. Experience shows that this approach tends to produce more accurate nowcasts than could otherwise be achieved.

Even so, it is important to recognise that uncertainty cannot be eliminated altogether. The mapping from indicators to official GDP estimates is imperfect, and all nowcasts come with confidence bands. It is therefore important that users are aware of the degree of accuracy that can be expected. Moreover, we show that the uncertainty surrounding nowcasts varies both across models and over time, and so it is important to pay attention to the confidence bands surrounding any given estimate.

The remainder of the article is structured as follows. The first section describes the practicalities of GDP nowcasting at the Bank. The second section provides details of the Bank's three main GDP nowcasting models. The third section discusses the models' performance and how that information feeds into the MPC's judgemental nowcast. The fourth and final section focuses on measures of uncertainty around nowcasts, and sets up the box titled 'The 2016 Q3 nowcast', which examines in detail the nowcast made in the August *Inflation Report* following the referendum on the United Kingdom's membership of the European Union.

## How is the nowcast produced in practice?

The Bank's approach to GDP nowcasting is summarised in **Figure 1**. It shows how the MPC's nowcast is not the mechanical output of a single model, but a judgemental estimate drawing on a range of data, models, and other information.

**Figure 1** The nowcast process

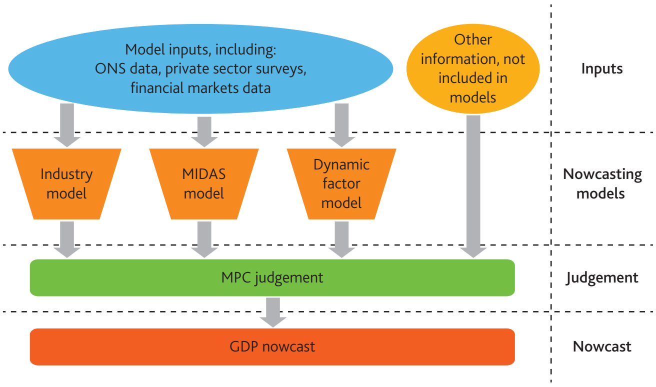


The starting point is the array of economic data available in advance of the GDP release, including high-frequency official estimates of output, private sector surveys, and other data, including from financial markets. The set of available information changes over time: at the start of the quarter relatively few data are available, but by the time GDP is released useful indicators are plentiful. This process of data publication is known as the 'data cycle' and has a bearing on which models are most useful, as well as on the overall uncertainty around the MPC's nowcast.

Statistical models are used to translate this unstructured set of information into numerical predictions for GDP growth. The econometric literature offers a number of possible modelling approaches. Three are used at the Bank, employing different statistical techniques and information sets.

### Why use multiple models?

It is a well-known empirical finding that combining forecasts from different models — for example by taking a simple or weighted average — can help to improve forecast

(1) Even the initial official estimates are not the final word, as they are subsequently revised as more information is collected and processed. The focus in this article is on nowcasting the preliminary (first) estimate of GDP growth. For information on how the Bank translates this into a prediction of the 'mature' estimate of GDP growth see Cunningham and Jeffery (2007), or Cunningham *et al* (2012) for further technical details.

performance.[1] This may be true in two senses. First, combination can provide a degree of insurance against large errors from any one model, thereby reducing the error variance. Second, there is also scope to increase accuracy in absolute terms, provided the errors of the individual predictions that feed into the combination tend to cancel out.

Forecast combinations can outperform because all models are imperfect to some extent. One reason is that the 'true' model specification is unobservable: in other words, the forecaster cannot be certain about which variables should optimally be included in the model, and in what form. A further challenge is that, for any given model, it can be difficult to estimate the relevant parameters with precision, in part due to the limited samples available.

In practice, the different perspectives that often arise from using different modelling approaches and information sets can contribute to making forecasts more robust, and may sometimes help to alert the MPC to particular factors, or signal relevant risks.

### The role of judgement

As a result of the different approaches employed, the uncertainties surrounding the nowcasts from each of the models will generally differ from one another. Moreover, those uncertainties tend to change through the quarter, with some models being relatively more useful towards the start of the data cycle and others towards the end, close to the publication of GDP. This raises an additional challenge, in that appropriate weights must be assigned to the different model predictions at different times.

It would be possible to mechanise such a weighting scheme — and later in the article we explore one possible mechanical combination approach. But the nowcast is ultimately a matter of judgement, reflecting the MPC's best estimate of GDP growth in light of all available information, including that not captured by the models.

A common scenario where judgement is applied is when the models suggest different central nowcasts from one another. Knowledge of the models' past performance can help to inform how much weight to place on each. But past performance need not carry over into the future, so a detailed understanding of the models is important in choosing appropriate weights at specific points in time. For instance, there might be an outlier in one of the input variables, affecting some models but not others, which might warrant less weight than usual; or if one model has shown persistent errors in a particular direction in recent quarters, it might suggest making an adjustment for this.

Judgements of another kind are occasionally incorporated, when the models are unlikely to pick up specific features of the economic conjuncture. Examples include events such as the Diamond Jubilee and Olympics in 2012, which shifted the pattern of quarterly growth that year. Evidence from other sources such as the Bank's Agency network may also have an influence on the nowcast.

## The Bank's nowcasting models

There are three core models in the Bank's nowcasting toolkit: an industry model, a combined-MIDAS model and a dynamic factor model. The industry model is the longest-standing and has already been introduced (see Bell *et al* (2014)). The combined-MIDAS model and the dynamic factor model (DFM) have been developed more recently and bring the Bank's toolkit closer to the technological frontier.

The three models have different strengths. The industry model mirrors the construction of the first estimate of GDP, using the high-frequency output data published by the ONS for the main sectors. The combined-MIDAS model extracts signals from other high-frequency indicators such as business surveys, which are timelier than official data and highly correlated with GDP growth. Finally, the DFM exploits the largest data set of all three models.

The remainder of this section briefly recaps the main features of the industry model and introduces the combined-MIDAS and DFM models. Further details can be found in a technical annex at the end of the article. It should be noted that the Bank's toolkit undergoes frequent maintenance and development, so what follows is a snapshot at this point in time, and is subject to change.

### The industry model

The industry model is a 'bottom-up' representation of the economy, based on the output side of the national accounts. A key feature of the model is that it incorporates official monthly output estimates as they are published by the ONS, replicating the process by which they feed into the first estimate of GDP.

Total output in the model is disaggregated into seven industries. In descending order of their share in the economy, the three main sectors are services, production, and construction.[2] Of those, we split services output into distribution services, private non-distribution services, and government services; and production output into manufacturing, utilities, and extraction. The ONS publishes monthly output estimates covering the listed industries with varying lags. In order to form a view of growth in the quarter,

(1) See eg Bates and Granger (1969), Timmermann (2006), Pesaran and Timmermann (2007), Tian and Anderson (2014).
(2) Services output accounts for a little over three quarters of the economy, production about 15%, and construction roughly 5%. Agriculture output, which accounts for less than 1% of output, is assumed to not contribute to quarterly growth.

the industry model approach involves forecasting those monthly outturns that are not yet available at each point in time.

**Figure 2** is a stylised representation of the various steps involved in producing the industry model nowcast.(1)

Figure 2 The industry model nowcast

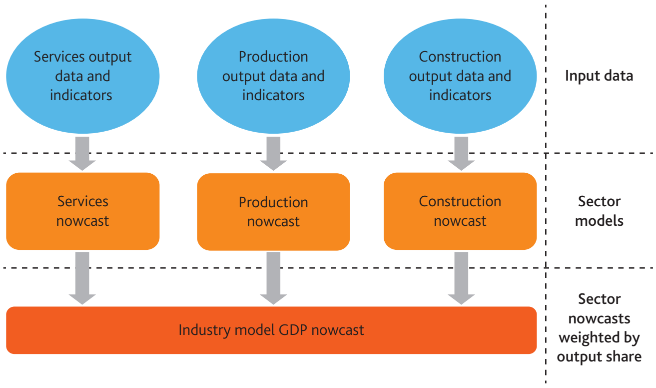


Starting from the latest ONS estimates, monthly output in each of the seven industries is projected forward based on its time series properties — ie its typical historical behaviour — supplemented by survey indicators.(2) The forecast monthly paths are used together with the available output data to estimate quarterly growth in each of the sectors. Finally, the industry model GDP nowcast is obtained by weighting together the resulting industry estimates according to their respective shares in the economy.

## The combined-MIDAS model

The term MIDAS is an acronym for mixed-data sampling. That captures the defining feature of MIDAS regressions: they deal with data sampled at different frequencies directly, without the need to convert them to a common frequency. More specifically, MIDAS models use higher frequency indicators — in our case monthly — to forecast a lower frequency variable — quarterly GDP growth.

Preserving the original frequency of the indicators is a desirable feature of this class of models, but the number of coefficients to be estimated can become prohibitively large as the difference in sampling frequency increases — for example, if one wishes to use daily observations to predict an annual variable — in which case more complicated techniques must be used. In macroeconomic applications, though, the number of parameters tends to be relatively small. In those circumstances it is often appropriate to follow the 'unrestricted MIDAS' approach, in which parameters can be estimated freely using standard techniques. That is the approach we take.

Including explanatory variables at their original monthly frequency, instead of converting to quarterly frequency by common methods such as averaging or summing, has the advantage of preserving information that might otherwise be lost: in this way, appropriate weights for the different months can be estimated; whereas averaging or summing would amount to imposing an equal-weight restriction.

The Bank's combined-MIDAS model consists of a collection of unrestricted MIDAS equations, each of which regresses the first estimate of GDP growth on monthly lags of a particular indicator of economic activity. The majority of indicators used in this model come from business surveys, such as the Markit/CIPS PMIs and those from the Confederation of British Industry (CBI) and the British Chambers of Commerce (BCC).(3)(4) Such survey indicators are usually timelier than official data, and are highly correlated with economic activity, so they can provide useful early signals. One reason to have a separate regression for each of the indicators is that it is often helpful to consider the range of nowcasts from the different models: for example, an unusually high dispersion of survey predictions might signal a more uncertain outlook for GDP than usual.

**Figure 3** summarises the combined-MIDAS model nowcast.

Figure 3 The combined-MIDAS nowcast

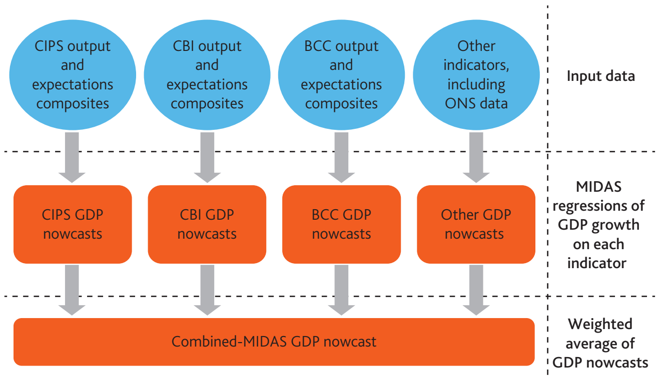


At each point in the data cycle there are a number of different GDP nowcasts, one from each indicator. The combined-MIDAS nowcast is obtained by taking a weighted average of the different model predictions, with those weights an increasing function of their relative past accuracy. At present, the Markit/CIPS and BCC surveys tend to receive the highest weights through most of the data cycle.

(1) The figure shows the high-level services, production and construction sectors, but forecasts are produced at a more disaggregated level, as described in this section.
(2) See Bell *et al* (2014) for details of the indicators that feed into the industry model.
(3) The current version of the combined-MIDAS model includes regressions based on the output and expectations composites from the Markit/CIPS, CBI and BCC surveys, the expected business activity balance from the Lloyds Business Barometer, ONS monthly output data, and the first principal component of a broader data set.
(4) The BCC survey is of quarterly frequency, so standard regressions are used in that case.

### The dynamic factor model (DFM)

Dynamic factor models have become standard in central banks' toolkits in recent years. Their principal advantage is that they enable forecasters to extract signals from larger data sets than is normally feasible using other techniques. In addition, DFMs provide a convenient framework to deal with some of the typical challenges faced by forecasters, such as the different frequencies of data releases, or the 'ragged ends' of the data set owing to staggered indicator publication dates as well as different starting points for the corresponding time series.

The Bank's DFM uses by far the largest information set among the three models: around 60 series spanning output, prices, and labour market data published by the ONS, business and consumer confidence indicators from private survey providers, as well as financial markets data. The DFM is able to deal with such a large data set by summarising the input data into a smaller number of variables called 'factors', which exploit the common movement of the series. Those factors are then used to nowcast the preliminary estimate of GDP growth.

One drawback of dynamic factor models in general is that it can be difficult to provide an economic interpretation of the factors. In trying to address this issue, the Bank's DFM follows a recent approach which involves grouping the data series into various blocks, according to an economic classification, and estimating a factor for each of these. Specifically, in addition to estimating one factor common to all the variables, we categorise each of our input series as real, nominal, 'soft' (business and consumer surveys) or financial, and estimate factors associated with each of these categories.

**Figure 4** summarises the DFM nowcast process, in which each series feeds into the common factor as well as into one of the category factors, and those factors are used to produce a GDP growth estimate.

Figure 4 The dynamic factor model nowcast

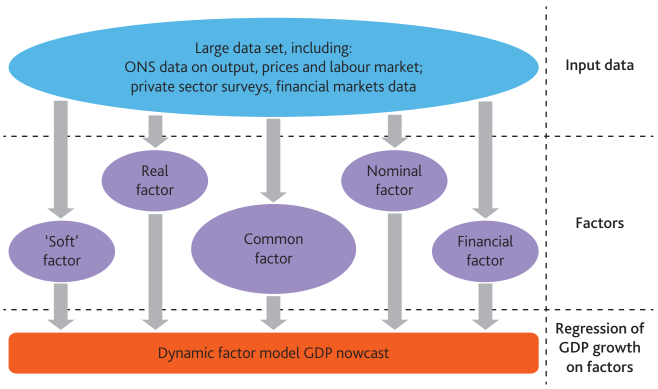


In practice, every time there is a data release the DFM nowcast is revised, and the impact of the release is calculated along with the contribution of each block to the nowcast. This is an appealing feature of the model, as it allows us to answer more particular questions, for example about the news specifically from the soft (survey) block or the financial block to the GDP nowcast at different points in time.

## Model performance and the MPC's judgement

This section evaluates the performance of the three models by means of an 'out-of-sample' forecast exercise over the ten-year period from 2007 to 2016.(1)(2)

We show that the relative performance of the models can vary, both across subperiods and within the nowcast quarter. For that reason, continued monitoring and evaluation of the statistical toolkit play an important role in informing the weights assigned to each of the models. Moreover, we present evidence in support of the Bank's eclectic approach from a simple mechanical forecast combination scheme, and illustrate how the judgemental incorporation of information not included in the models has improved performance.

### Model nowcasts across time

The results of the out-of-sample exercise show that the different models would have broadly tracked the first estimate of GDP growth over 2007–16. This is illustrated in **Chart 1**, which plots the evolution of the three models' nowcasts against preliminary GDP growth.(3)

Chart 1 Out-of-sample nowcasts

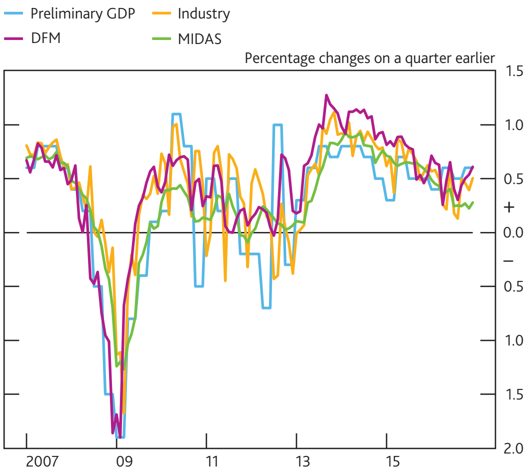


Sources: ONS and Bank calculations.

(1) By out-of-sample it is meant that we estimate the models on the information available up to each point and forecast one-step ahead. For the industry and combined-MIDAS models the exercise uses 'real-time' data, ie the actual data available at each point in time (this may differ from current vintages due to revisions). For the DFM a subset of the input data used is latest vintage.
(2) The start of the evaluation period is determined by data availability.
(3) The chart shows three nowcasts per quarter for each model, at the beginning of each month. The cut-off point is after the Markit/CIPS PMI releases.

From the chart it is clear that no single model is closest to the preliminary GDP line at all times: instead, it suggests that different models may be more useful in different subperiods. For example, the dynamic factor model would have been the most accurate going into the financial crisis, benefiting from its use of financial market data, although it would have systematically overestimated growth between 2013 and 2015; whereas the industry model is often better able to capture volatile moves in GDP growth, but can sometimes make very large errors, particularly early in the data cycle. Such examples help to illustrate the robustness rationale for maintaining a suite of models.

## Model performance within the nowcast quarter

Although relative model performance may vary across subperiods, we can take the 2007–16 period as a whole to gauge which models tend to be more accurate at different stages in the data cycle. The remainder of this section comments on this based on a common metric of forecast performance: the root mean squared error (RMSE).(1)

**Chart 2** shows the evolution of the RMSEs from the three models through the nowcast quarter — ie the 90 days or so between consecutive preliminary GDP releases. In addition to the models' RMSE lines, it also includes three RMSE diamonds from a mechanical nowcast combination scheme that will be discussed later in the section.

Chart 2 Model RMSEs through the data cycle (2007–16)

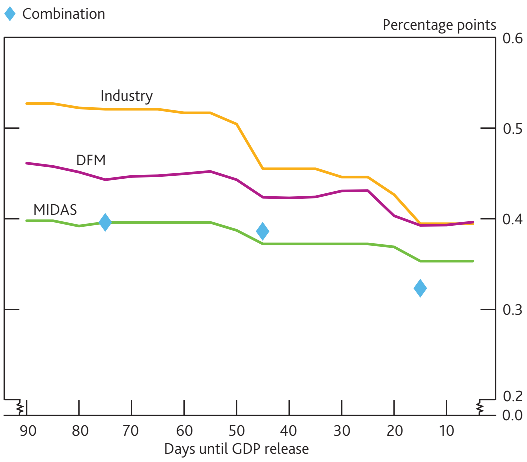


Source: Bank calculations.

In the chart, a lower RMSE indicates a higher degree of nowcast accuracy. Unsurprisingly, RMSEs tend to decline through the quarter as more data become available.

The combined-MIDAS model has the lowest RMSE through the entire data cycle. It starts off at around 0.4 percentage points, substantially lower than the DFM and the industry model, in part reflecting its heavier use of business surveys: such indicators are particularly helpful early in the quarter, as they are timelier than official data and highly correlated with GDP growth. The combined-MIDAS RMSE then progressively falls to about 0.35 percentage points as the preliminary GDP release draws near.

By contrast, the industry model has the highest RMSE early on at around 0.5 percentage points, since at that point no monthly output estimates covering the nowcast quarter are available. But its RMSE then drops the most steeply as those data get incorporated — for example, notice the decline after the first month of production data is published, around 45 days before GDP — eventually becoming marginally lower than that of the DFM, at 0.4 percentage points.(2) Although the industry model replicates the construction of preliminary GDP, there is still considerable uncertainty ahead of the release as the official data it uses are published with relatively long lags and so a forecast component always remains.(3)

Finally, the RMSE of the dynamic factor model tends to sit in the middle through most of the cycle, dropping from around 0.45 percentage points at the beginning to 0.4 percentage points just before the GDP release.

## The MPC's judgemental nowcast

Performance metrics of the kind discussed in this section help to inform the weights put on each of the models in a typical quarter. As such, the MPC's nowcast can normally be expected to be most influenced by the combined-MIDAS model early in the cycle, hence effectively putting more weight on survey information; whereas later in the quarter it will increasingly take a signal from official data, as processed by the industry model. However, this is not a hard and fast rule and will often be adjusted, for example in the light of the models' recent performance.

Although the MPC's nowcasts are not derived mechanically from the models, evidence from one possible mechanical combination scheme, based on a similar principle to that described in the previous paragraph, lends support to the MPC's eclectic approach. The diamonds in **Chart 2** represent the RMSE of a combination scheme that gives more weight to those models that have performed better in the recent past.(4) The combination performs roughly on par with the best model early in the data cycle, but outperforms all three models towards the end of the quarter.

(1) The RMSE can be thought of as a 'typical' error. It differs from the mean absolute error in that larger errors are disproportionally penalised.
(2) The industry model RMSE figures are likely somewhat overstated, owing to the fact that monthly construction output data only begun to be published by the ONS in 2010, and so could not be used for most of the evaluation period. In practice these data are routinely incorporated into the industry model nowcast.
(3) Under current release calendars, only one month of services output data and two months of production/construction data are published before the preliminary GDP estimate.
(4) More specifically, the weights are inversely proportional to each model's RMSE over a two-year rolling window.

## The 2016 Q3 nowcast

This box examines the most recent sizable nowcast error, in 2016 Q3, following the referendum on EU membership.

In the August 2016 *Inflation Report*, prepared around one month after the vote to leave the European Union, the Monetary Policy Committee (MPC) forecast that output would practically stagnate in Q3. However, in the event the ONS published a preliminary GDP growth estimate of 0.5%.

Although large by historical standards, that surprise was not unprecedented: there had been a number of similar or larger nowcast errors over the previous ten years (see **Chart 5**). In previous instances, significant nowcast errors had generally been associated with large swings in GDP growth. In 2016 Q3, however, growth remained steady and it was volatility in survey indicators that drove the divergence, as most surveys released soon after the vote pointed to a weaker near-term outlook than implied by current ONS estimates.

Producing a Q3 nowcast soon after the referendum required some difficult judgements to be made. For example: how much signal to take from one month of very weak survey data? How much weight to put on pre-referendum data? These judgements were made particularly challenging by the lack of historical precedent. In the event the MPC judged it appropriate to take some steer from the weakness in survey indicators but stopped short of taking a full signal.

### Data in the immediate aftermath of the referendum

The August *Inflation Report* was published six weeks after the EU referendum. At the time it was prepared no official data covering the post-referendum period had been released by the ONS. However, some of the surveys that the MPC usually monitors had already published their first post-referendum outturns, and were almost unanimous in pointing to a marked deterioration in the near-term outlook.

**Chart A** shows some of the important post-referendum indicators available ahead of the *Inflation Report*.

The Markit/CIPS PMI surveys — normally strongly correlated with economic activity — signalled outright falls in output: for example, the July level of the output composite, if sustained, was consistent with a 0.4% quarterly contraction in GDP. The

**Chart A** Key indicators ahead of the August 2016 *Inflation Report*[a]

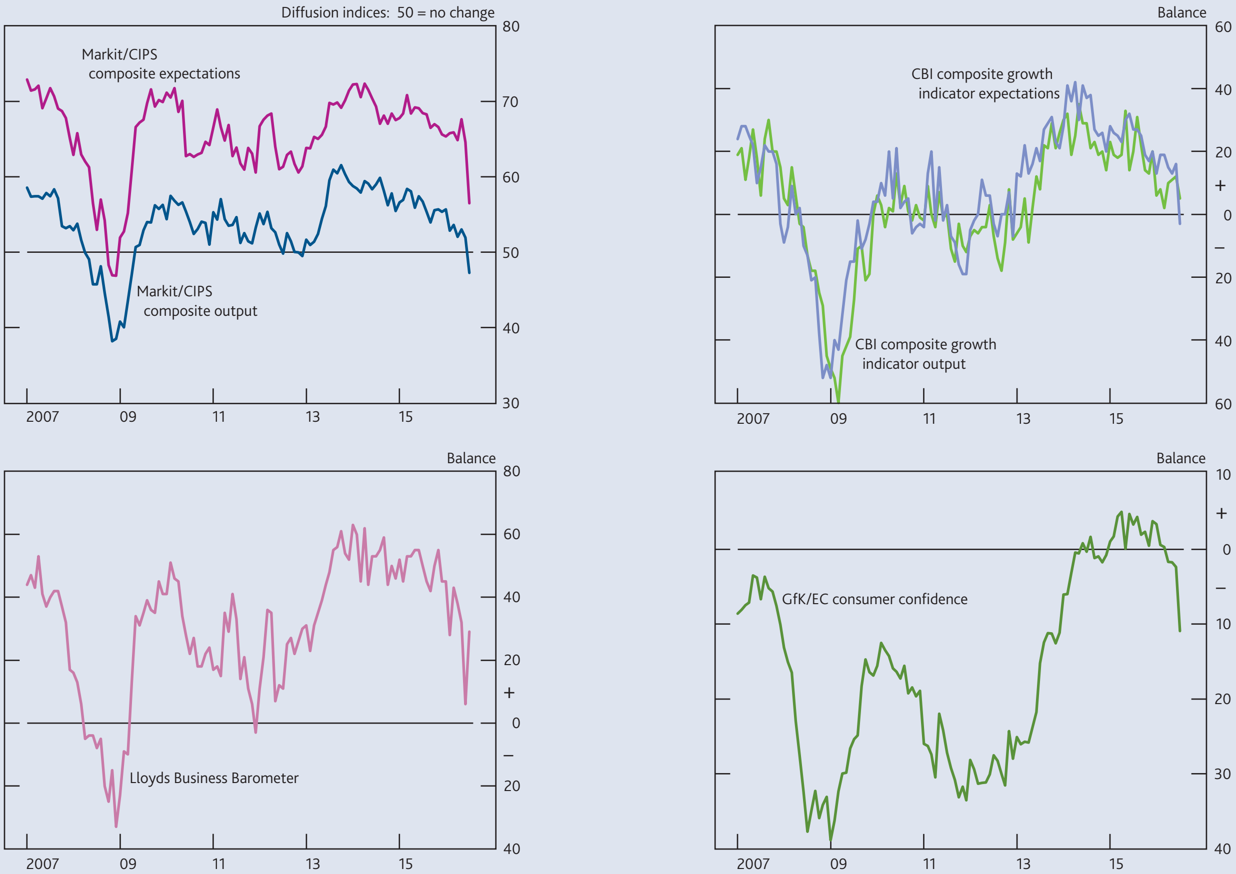


Sources: CBI, GfK (research carried out on behalf of the European Commission), IHS Markit, Lloyds Banking Group and Bank calculations.

(a) Two post-referendum observations were available for the Lloyds Business Barometer before the *Inflation Report* (as the June survey was conducted in the week after the EU referendum). For the remaining series July was the only post-referendum data point.

CBI's Composite Growth Indicator appeared to corroborate that weakness; and other indicators of business and consumer confidence had also deteriorated sharply.

### Model nowcasts in early August

Just ahead of the August *Inflation Report*, all three nowcasting models discussed in this article also signalled slowing, albeit slightly positive, growth in Q3:(1)

- The industry model nowcast was 0.1%, reflecting the collapse in the Markit/CIPS PMIs, which — in the absence of intraquarter output data — dominated the industry model nowcast at that point.
- The combined-MIDAS model nowcast, heavily based on business surveys, was also at 0.1%, reflecting the falls seen across qualitative indicators.(2)
- The dynamic factor model nowcast was still slightly stronger, at 0.3%, although that was in part because the available post-referendum surveys constituted a smaller subset of its input data, compared to the other models.

At that stage all the models remained influenced by pre-referendum data. But the generalised deterioration in post-referendum indicators suggested a potential departure from normal economic dynamics, and so it appeared plausible that pre-referendum data might have less predictive power with respect to Q3 developments than lagged data normally do. It was therefore judged appropriate to discount pre-referendum lags somewhat, which at face value suggested that the models might tend to overestimate growth.

Despite that judgement, the MPC took a limited signal from the post-referendum surveys, and its near-term outlook was not as pessimistic as the average of other forecasters: upon extensive analysis of a broader range of data, and being alert to the possibility of excessive volatility in the available data, the committee judged it appropriate to aim up from the negative growth rates suggested by prominent indicators such as the Markit/CIPS PMIs, and so it forecast output would broadly stagnate in Q3.

### What can be learned from the 2016 Q3 experience?

Statistical analysis by its very nature involves the study of imperfect relationships. Since the read-across from activity indicators to official GDP estimates is imperfect, nowcast errors are to be expected, and large surprises are likely to occur, albeit infrequently. As such, the occurrence of a large error does not, in and of itself, mean that the nowcast process is faulty and in need of repair.

Nevertheless, it is good practice to scrutinise sizable errors to see if any important lessons can be drawn from them. The lessons from the 2016 Q3 experience are somewhat limited. Even prior to this it was known that indicators such as business surveys could over-react in some circumstances — this is what led the MPC to produce a higher nowcast than some of those indicators would, taken at face value, have suggested. However, the lack of many such unusual events meant there was a great deal of uncertainty about the exact size of the adjustment that should be made. In raising its nowcast relative to the growth rates suggested by post-referendum business surveys the MPC made the right directional judgement, but in the event fell short on its magnitude. This latter fact will be taken into account if and when similar events occur in the future.

This last point is important, however. The vote to leave the European Union, which was widely unexpected, was an event without historical precedent, and it is quite possible that no such similar event will occur again. As such, although it acts as a warning for nowcasters to be alert to jumps in the uncertainty around nowcasts in the face of major events — more formally, it adds to our knowledge of the distribution of possible errors in unusual circumstances — whether it will provide much additional guidance on the extent to which the nowcaster should 'aim off' the model outputs in future is questionable.

(1) Note that the figures below may differ slightly from those underlying **Charts 1** and **3** because more information was incorporated into those nowcasts (eg CBI surveys, consumer confidence) than assumed available at the point in the data cycle represented in the charts.
(2) A judgement was made to exclude the quarterly BCC balances from the model at that time, as post-referendum data would only become available in September.

As discussed earlier in the article, in addition to the model outputs, the MPC takes into account a range of other relevant information, such as knowledge of specific events and evidence from the Bank's Agents. Experience demonstrates that this approach tends to produce more accurate nowcasts than could otherwise be achieved: the RMSE of the MPC's judgemental nowcasts, made roughly 80 days before the GDP release, is 0.3 percentage points — substantially lower than any of the pure model-based RMSEs at the corresponding stage in the data cycle.

**Chart 3** shows the MPC's judgemental nowcasts alongside the described mechanical combination scheme at a comparable time in the quarter.

It can be inferred from the chart that the lower RMSE of the MPC's nowcasts is due in part to successful judgements made around the Diamond Jubilee and the London Olympics in 2012 — although they also appear to better track the evolution of GDP growth more generally.

**Chart 3** MPC and combination nowcasts around 80 days ahead of the GDP release

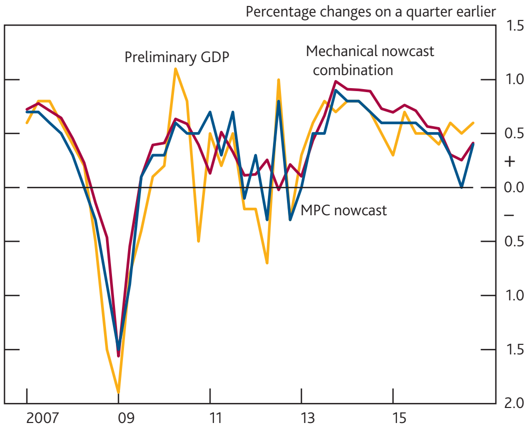


Sources: ONS and Bank calculations.

**Chart 4** Prediction intervals around hypothetical model nowcasts

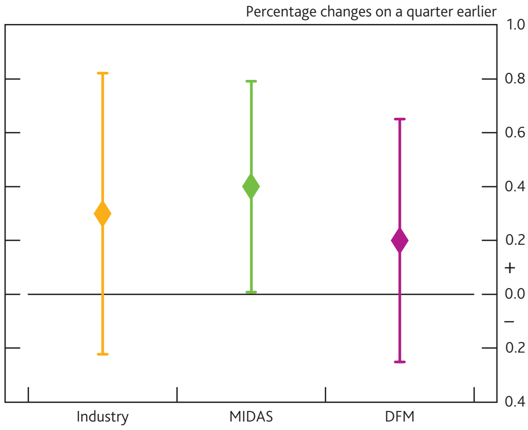


Source: Bank calculations.

## Uncertainty around nowcasts

Uncertainty is inherent in nowcasting, as no approach is able to predict GDP growth perfectly ahead of its publication by the ONS. In other words, nowcast errors are to be expected. That means that while a given nowcast represents the most likely outturn according to a particular model or judgement, a range of different outcomes is usually also plausible.

Past errors can offer a guide to the distribution of possible outcomes surrounding any given nowcast — although past surprises may not be fully representative of the true underlying uncertainty, particularly with limited samples. This section considers the uncertainty surrounding nowcasts from the Bank's three models, based on the errors from the out of sample exercise described above, and provides further details on the distribution of past surprises to the MPC's nowcasts.

### Uncertainty around model nowcasts

Nowcast uncertainty varies across models and through the data cycle. The root mean squared errors presented in the previous section provide a summary measure of that uncertainty: a higher RMSE normally indicates a wider dispersion of possible outcomes.

Under some simplifying assumptions, those RMSEs may be used to infer the distributions from which errors are drawn, and can be employed to construct prediction intervals around the central nowcasts from the respective models. Such intervals represent the range of most likely outturns, for a given level of confidence, and are commonly used to communicate forecast uncertainty — the greater the uncertainty, the wider the intervals will be.

To illustrate, **Chart 4** shows, based on the 2007–16 errors from each of the three models, the prediction intervals that would be expected to include the first estimate of GDP growth roughly two thirds of the time (+/-1 RMSE), given the respective — hypothetical — nowcasts made 80 days before the GDP release (around the time the *Inflation Report* forecast is produced).

For example, consider the MIDAS nowcast in the chart (the green diamond): given a central prediction of 0.4%, one would expect the published outturn to be between zero and 0.8% (those are the edges of the green line) about two thirds of the time, although it could still fall outside this range roughly a third of the time.

The width of the bands in **Chart 4** is based on the RMSEs from the previous section at an early point in the data cycle, and so would tend to become narrower as the publication of GDP draws closer.

### Uncertainty around the MPC's nowcast

Turning to the MPC's nowcast, it was shown in the previous section that its RMSE of 0.3 percentage points was lower than that of any of the models. Correspondingly, prediction intervals around the MPC's central nowcasts should be narrower than their model-based counterparts, as the evidence suggests the MPC's judgemental combination of model outputs and other information helps to mitigate uncertainty.

However, the RMSE is just a summary measure and it can be useful to consider the full error distribution. **Chart 5** shows the distribution of MPC nowcast errors (in absolute value) over 2007–16.

The MPC's nowcasts have often been accurate. The preliminary GDP estimate has fallen within 0.1 percentage points of the MPC's expectation about 50% of the time — and if one allows for a wider margin of error of 0.3 percentage points that figure rises to 80%. Still, it is important to be aware that larger errors can occasionally occur.

**Chart 5** Frequency of absolute nowcast errors (2007–16)

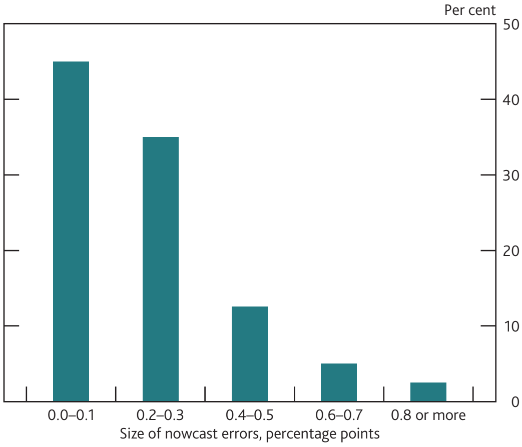


Source: Bank calculations.

The latest sizable surprise happened in the third quarter of 2016, following the EU referendum, and is explored in greater detail in the box on pages 128–29.

## Conclusion

The Bank's GDP nowcasts represent the MPC's estimates of growth in the current quarter and are an important input to monetary policy decisions.

The three models described in this article play an important part in forming those estimates. Within that, the MPC's nowcast will normally be most influenced by the combined-MIDAS model early in the data cycle, effectively putting more weight on survey information; whereas later in the quarter it increasingly takes a signal from official data, as processed by the industry model.

However, the MPC's nowcast does not flow mechanically from the models, and is ultimately a matter of judgement. Experience shows that this approach tends to produce more accurate nowcasts than could otherwise be achieved.

Although the evidence suggests that the judgemental combination of model predictions and other information helps to mitigate nowcast uncertainty, sizable errors may still occasionally occur — as was the case in 2016 Q3 following the EU referendum — so it is important that nowcast users are aware of the degree of accuracy that can reasonably be expected.

## Annex
## The industry model

Most sector forecasts in the industry model are produced using some variant of the following basic relationship:

$$\widehat{output}_t = \hat{\beta}_0 + \hat{\beta}_1 output_{t-1} + \hat{\beta}_2 indicator_t, \qquad (1)$$

where $\widehat{output}_t$ is the output growth forecast in month $t$ for a given sector, $output_{t-1}$ is output growth in the previous month and $indicator_t$ is the month $t$ outturn of the relevant indicator.

Further details may be found in Bell *et al* (2014).

### The combined-MIDAS model

Mixed-data sampling (MIDAS) regressions explain a low-frequency variable by high-frequency variables and their lags. In the standard MIDAS approach only a small number of 'hyper-parameters' are estimated to find the best fit of a given function to the data as first introduced by Ghysels, Sinko and Valkanov (2007). This can have the advantage of greater parsimony — if a large number of parameters needed be estimated — otherwise it comes at the cost of flexibility. As the difference in frequencies in our application is small (the dependent variable is quarterly; the regressors are monthly), we can estimate coefficients for each of the months using standard ordinary least squares techniques, following Foroni and Marcellino (2014) and Foroni, Marcellino and Schumacher (2015) — the so-called unrestricted MIDAS approach.

As an example of the unrestricted MIDAS approach underpinning the combined-MIDAS model, suppose we aimed to forecast quarterly GDP growth in the first quarter of 2017 based on monthly Markit/CIPS output composite data up to March. We might then use the following regression model:

$$\widehat{GDP}_{Q1} = \hat{\beta}_0 + \hat{\beta}_1 CIPS_{Mar} + \hat{\beta}_2 CIPS_{Feb} + \hat{\beta}_3 CIPS_{Jan} + \hat{\beta}_4 CIPS_{Dec} \qquad (2)$$

where the forecast of first-quarter GDP growth, $\widehat{\text{GDP}}_{Q1}$, is a function of the four monthly outturns of the Markit/CIPS indicator between December and March.

A similar approach is followed for each of the indicators included in the model (the number of lags may vary across indicators and through the data cycle).

### The dynamic factor model

The dynamic factor model is a reduced form model using data in monthly and quarterly frequencies from an extensive array of indicators, following the framework described in Giannone, Reichlin and Small (2008) and Bańbura and Modugno (2014) for assessing the informational content of the real-time data flow. The model has the following general state space representation:

$$y_t = \Lambda f_t + \xi_t, \qquad (3)$$

$$f_t = A f_{t-1} + u_t, \qquad (4)$$

where $y_t$ is a vector of the target variable (GDP in our case), $f_t$ contains the factors extracted from the extensive data set of indicators, as explained in the main text, and $f_{t-1}$ includes the lags of the factors. $\Lambda$ is a matrix of factor loadings and A are matrices of autoregressive coefficients. $\xi_t$ and $u_t$ are the idiosyncratic disturbances and the residual errors, respectively.

The model is estimated using maximum likelihood techniques to easily account for missing data (eg because of the different starting and end points of the indicators in the data set).

## References

**Bańbura, M and Modugno, M (2014)**, 'Maximum likelihood estimation of factor models on datasets with arbitrary pattern of missing data', *Journal of Applied Econometrics*, Vol. 29, No. 1, pages 133–60.

**Bates, J M and Granger, C W J (1969)**, 'The combination of forecasts', *OR*, Vol. 20, No.4, pages 451–68.

**Bell, V, Co, L, Stone, S and Wallis, G (2014)**, 'Nowcasting UK GDP growth', *Bank of England Quarterly Bulletin*, Vol. 54, No. 1, pages 58–68; available at www.bankofengland.co.uk/publications/Documents/quarterlybulletin/2014/qb14q1.pdf.

**Cunningham, A, Eklund, J, Jeffery, C, Kapetanios, G and Labhard, V (2012)**, 'A state space approach to extracting the signal from uncertain data', *Journal of Business and Economic Statistics*, Vol. 30, No. 2, pages 173–80.

**Cunningham, A and Jeffery, C (2007)**, 'Extracting a better signal from uncertain data', *Bank of England Quarterly Bulletin*, Vol. 47, No. 3, pages 364–75, available at www.bankofengland.co.uk/publications/Documents/quarterlybulletin/qb070301.pdf.

**Foroni, C and Marcellino, M (2014)**, 'A comparison of mixed frequency approaches for nowcasting euro area macroeconomic aggregates', *International Journal of Forecasting*, Vol. 30, No. 3, pages 554–68.

**Foroni, C, Marcellino, M and Schumacher, C (2015)**, 'Unrestricted mixed data sampling (MIDAS): MIDAS regressions with unrestricted lag polynomials', *Statistics in Society*, Vol. 178, No. 1, pages 57–82.

**Ghysels, E, Sinko, A and Valkanov, R (2007)**, 'MIDAS regressions: further results and new directions', *Econometric Reviews*, Vol. 26, No. 1, pages 53–90.

**Giannone, D, Reichlin, L and Small, D (2008)**, 'Nowcasting: the real-time informational content of macroeconomic data', *Journal of Monetary Economics*, Vol. 55, No. 4, pages 665–76.

**Pesaran, M H and Timmermann, A (2007)**, 'Selection of estimation window in the presence of breaks', *Journal of Econometrics*, Vol. 137, No. 1, pages 134–61.

**Tian, J and Anderson, H (2014)**, 'Forecast combinations under structural break uncertainty', *International Journal of Forecasting*, Vol. 30, No. 1, pages 161–75.

**Timmermann, A (2006)**, 'Forecast combinations', *Handbook of Economic Forecasting*, Vol. 1, pages 135–96.