

Extracting a better signal from uncertain data

By Alastair Cunningham and Christopher Jeffery of the Bank's Conjunctural Assessment and Projections Division.

Most macroeconomic data are uncertain — they are estimates rather than perfect measures. One symptom of that uncertainty is the propensity of statistical agencies to revise their estimates in light of new information or methodological advances. While revisions should move estimates closer to the 'truth', the potential for early estimates to be revised poses challenges for forecasting and economic analysis. Over the past few years, Bank staff have undertaken a range of research into how best to deal with the ensuing uncertainty. The results of that research have been used for some time as part of the toolkit available to staff when briefing the Monetary Policy Committee. This article describes some further developments in that research effort aimed at refining the staff's toolkit.

Introduction

Most macroeconomic data are uncertain — they are estimates rather than perfect measures. Measurement errors arise because data are often based on samples. And they also arise because many variables — for example, in-house software investment — are not easily observable at all, necessitating the use of proxies. Such uncertainty poses challenges for both forecasting and economic analysis. As such, according to Lomax (2004), 'few subjects consume more of [the Monetary Policy Committee's] time and energy'.

But how can the extent of the data uncertainty problem be judged and what can be done about it?

One symptom of data uncertainty is the propensity of statistical agencies to revise their estimates. In order to provide a timely indication of economic developments, the Office for National Statistics (ONS) publishes early estimates based on the survey responses available at the time. These estimates are inevitably revised as more information is received. Additionally, the ONS periodically reviews its statistical methods. To ensure comparability of the National Accounts through time, the ONS reconsiders the back data in the light of any methodological changes — leading to further revisions. The scale of the ensuing revisions gives one indication of the extent of data uncertainty in the past. And to the extent that past revisions give a good guide to the likely scale of revisions in the future, they can also be used to gauge the uncertainty associated with the latest data.

Recognition of this uncertainty leads naturally to a probabilistic view of the past. Estimation of a confidence interval around the published data is a first step, and gives an indication of the potential scale of revisions. Going further,

economists can make use of additional evidence about the current economic conjuncture and the past patterns in revisions to assess the likely direction of future revisions.

Treating uncertain data in this way is neither new nor unique to the Bank. A study by the Statistics Commission (2004) concluded that 'the main users of the [ONS] statistics knew that revisions should be expected, understood the reasons for them, and were able to make some allowance for them when taking important decisions'. However, most attempts to allow for potential revisions are informal — recognising that revisions might occur but not offering any quantification of how large they could be.

Recognising the potential for revisions to macroeconomic data, Bank staff have undertaken a range of research into how best to deal with data uncertainty. Some of that research has focused on the potential implications of data uncertainty for forecasting and policy formation — see, for example Jääskelä and Yates (2005). Other work has aimed to enhance the interpretation of uncertain data. Lomax (2004) describes the array of evidence — such as business surveys and reports from the Bank's regional Agents — deployed by staff in interpreting the recent conjuncture. And Ashley *et al* (2005) set out a first-pass method for formalising — and hence making more rigorous — the Bank staff's approach to combining the evidence from such publicly available sources. The statistical methods outlined in Ashley *et al* (2005) have been used by Bank staff for some time when briefing the Monetary Policy Committee (MPC) on developments in output growth. And this method was used in the August 2007 *Inflation Report*.

This article describes further developments in this research which aim to exploit a richer array of evidence. The next section describes the scale of revisions to early National

Accounts estimates. Subsequent sections describe how the uncertainty caused by prospective revisions can be mitigated. As mentioned above, the aim of the exercise is to make the best use of publicly available evidence when interpreting the picture painted by the latest ONS estimates.

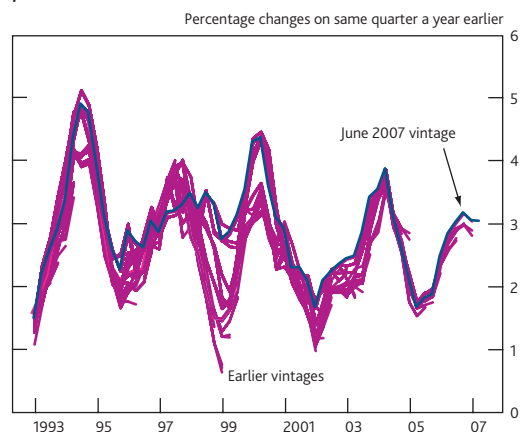
The cornerstone of the approach described in this article is the use of the experience of past revisions to proxy current data uncertainty. This raises two important caveats. First, data uncertainty may not be fully captured in revisions — even where data are not subject to revision, they may be based on samples and proxy measures and hence offer an uncertain measure. Second, past revisions may not always be a good indicator of prospective revisions. The statistical methods described in this article should not therefore be used in isolation. They need to be complemented with a careful understanding of the way in which macroeconomic aggregates are compiled and revised. The box on page 366 describes the revisions process applied in the production of the United Kingdom’s National Accounts and introduces some issues in mapping from the scale of past revisions to a view of current data uncertainty.

The scale of past revisions

Reviewing the scale of past revisions is a natural first step in interpreting data that are subject to revision. And in recent years, a number of ‘real-time’ data sets — describing the evolution of estimates through successive data releases (or vintages) — have been developed to facilitate this sort of exercise. The Bank first published a limited real-time data set in 2002. This database has subsequently been updated and materially extended and now covers around 100 macroeconomic time series.⁽¹⁾

To illustrate the potential scale of the uncertainty in National Accounts data, **Chart 1** compares the latest estimate of GDP growth with earlier vintages released since January 1993. Revisions to GDP growth have often been large.

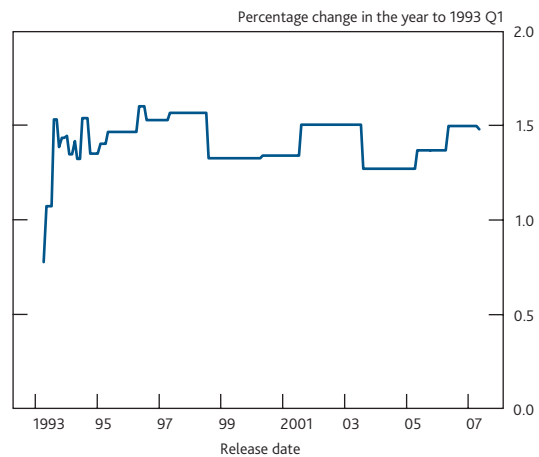
Chart 1 Real-time estimates of real GDP at market prices^(a)



(a) The blue line shows the profile of GDP growth published in the 2007 Quarterly National Accounts (QNA). Each pink line shows the profile of GDP growth published in an earlier vintage of the National Accounts.

Revisions have also occurred several years after the event, as shown in **Chart 2**, which plots successive estimates of annual GDP growth in 1993 Q1.

Chart 2 Successive estimates of annual growth in real GDP in 1993 Q1



The potential for revision means that early estimates can give a noisy signal of the underlying growth profile that will be revealed in more mature data. One metric for this uncertainty is the variance of revisions to the first estimates published by the ONS. **Table A** shows this ‘revisions variance’ for estimates of quarterly growth in real GDP and a selection of the output and expenditure components published in the first *Quarterly National Accounts (QNA)* release (column A).⁽²⁾

Table A Scale of revisions to first estimates of quarterly growth for select constant price National Accounts components^(a)

	Variance of revisions since first QNA release	Variance of growth shown in the latest vintage of data	Noise to signal ratio at first QNA release ^(b)
	A	B	C = A/B
GDP	0.10	0.07	1.38
Household consumption	0.22	0.25	0.91
Whole-economy investment	2.40	2.99	0.80
Government consumption	1.09	0.51	2.14
‘Economic’ exports	1.77	3.52	0.50
‘Economic’ imports	0.96	2.17	0.44
Service sector output	0.13	0.10	1.30
Production sector output	0.35	0.55	0.63

(a) Figures have been estimated over data released between 1993 and the latest (June 2007) QNA. While revisions could be estimated over a longer time horizon, there is some evidence of a structural break in the scale of revisions to National Accounts variables in the early 1990s (see Garratt and Vahey (2006)). The same estimation window is used throughout this article.
 (b) Note that figures are rounded so columns A and B may not map to column C.

(1) The data are available at www.bankofengland.co.uk/statistics/gdpdatabase.
 (2) **Table A** is based on revisions over the five years since the first QNA release. So, revisions to the first QNA estimate of each quarter’s growth are evaluated over a fixed five-year window. The calculations exclude the impact of any revisions made during the 1998 *Blue Book*, which saw — among other things — the introduction of a new system of national accounts (ie ESA 95) because those revisions were associated with a change in economic classification and are not judged to be indicative of ongoing data uncertainty. This treatment is retained throughout.

Past revisions to the United Kingdom's National Accounts as an indicator of current uncertainty

The Office for National Statistics (ONS) periodically reviews the causes and scale of past revisions to the United Kingdom's National Accounts — see, for example Obuwa and Robinson (2006). National Accounts estimates are revised for a number of reasons: to correct any processing errors; to incorporate additional information received; to re-reference and rebase; and to incorporate changes to either the accounting framework or the methods used to construct estimates. Revisions to correct processing errors have been infrequent in the past, but revisions for the other reasons have been material.

So how sure can one be that the scale of past revisions is representative of current uncertainty?

Revisions as information is received and processed

The evidence available to support the published National Accounts grows from the time of the first estimates. The body of evidence grows as surveys of firms' output are supplemented by increasing information on expenditure and income. Such quarterly information is then benchmarked to annual data from a variety of sources. This benchmarking exercise is typically completed three to fifteen months after the first estimates — the results being published in the annual *Blue Book*. Even then, the various sources available may not give a consistent impression of activity across the National Accounts. The ONS therefore applies a set of disaggregated coherence checks (known as input-output balancing), motivating further revisions. This process generates revisions up to ten quarters after the first estimates. And further revisions are possible as some evidence is received with a longer lag.

As long as early estimates continue to be based on incomplete information, the past experience of revisions in the first few quarters after data are released is likely to be informative about the magnitude of current data uncertainty. That said, the ONS has embarked on a major statistical modernisation programme (see Beadle (2007)), one aspect of which is to enhance input-output balancing. Successful delivery will accelerate the input-output balancing process and hence might increase the rate at which noise in early data estimates is reduced.

Revisions due to re-referencing and rebasing

The National Accounts measure activity in both real and nominal terms. Real measures adjust for any changes in the aggregate price level, and are currently referenced to 2003 prices. Since the *Blue Book* published in 2005, ONS policy has been to shift the reference year forwards by one year with the publication of each annual *Blue Book*.⁽¹⁾

Re-referencing leads naturally to revisions to the levels of both price deflators and real measures, but has no effect on growth rates. However, re-referencing is also accompanied by changes to the latest base year. The relative price of different goods and services is fixed at the base-year level when calculating real growth rates. Rebasing will therefore lead to revisions to growth profiles, as the spending basket is updated to reflect changing patterns of expenditure.

As long as disaggregated spending patterns remain hard to measure, the past experience of revisions associated with rebasing is likely to be informative about current data uncertainty.

Revisions as methods are changed

Data remain subject to revision for many years after the initial release. One reason for such late revisions is that the methods used to manipulate statistical returns are subject to continuous review by ONS staff. When methods are changed, the ONS work through any implications for back data and incorporate revisions in subsequent annual *Blue Books*. So, for example, in the 2006 *Blue Book* the method used to estimate capital depreciation was changed, leading to revisions to the profile of investment from 1948. Revising the back data in this way helps ensure comparability across the whole time series.

The degree to which past methodological revisions are informative about current uncertainty depends on the nature of the methodological change. Some methodological revisions reflect changes to economic classification or one-off improvements to data processing technologies. For example, in September 1998, the National Accounts moved to a new accounting framework known as ESA 95. Such changes do not reflect ongoing difficulties in measurement and hence are not informative about current data uncertainty. But others follow from continued attempts to improve the measures used to capture aspects of economic activity — for example, the revisions that followed the *Atkinson Review* of public sector output and productivity (Atkinson (2005)). These considerations show that careful thought about the factors driving past methodological revisions is required when constructing estimates of current uncertainty.

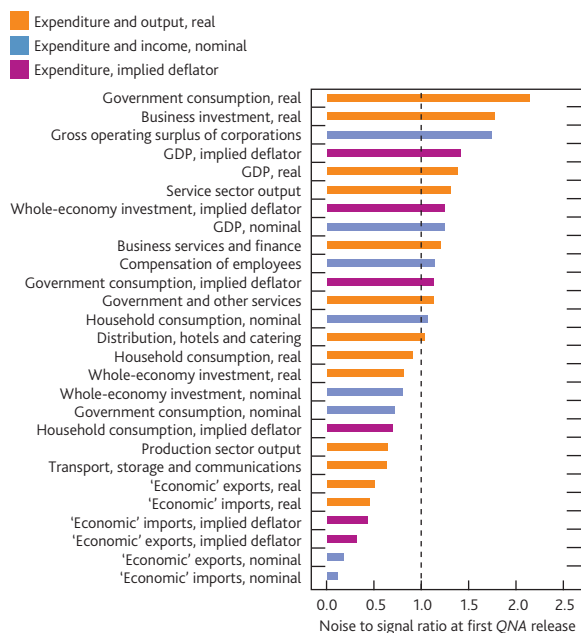
(1) The *Blue Book* published in 2007 was an exception. To free up resources necessary to allow delivery of modernised National Accounts, the ONS reduced the scope of that *Blue Book*. One element of the reduced scope was maintenance of 2003 as the reference price level.

Ranking variables by this measure may give a misleading impression of how far the potential for revision complicates economic analysis, because the measure does not control for differences in the volatility of their growth profiles over time. For example, a revision of 0.1 percentage points may be material for analysis of a variable with a relatively smooth growth profile (like GDP), but is unlikely to be material for a relatively volatile variable like whole-economy investment. Column B reports the variance of the quarterly growth rates shown in the latest vintage of data, and demonstrates a wide variation in the volatility of growth profiles.

The 'noise to signal' ratio (column C) provides a more natural metric for the scale of data uncertainty. This measure compares the variance of revisions (the noise — in column A) to the variance of the growth profile shown in the latest data (the signal — in column B)⁽¹⁾ and hence puts the scale of revisions in the context of the relative volatility of the underlying series (A/B). On this measure, revisions noise appears to have been most material for GDP, government consumption and service sector output.

Chart 3 plots the noise to signal ratio for the first QNA estimates of quarterly real GDP growth alongside 26 expenditure, output and income series. The chart shows considerable differences across variables — some early estimates providing a noisier signal than others.

Chart 3 Noise to signal ratio for first QNA estimates of quarterly growth^(a)



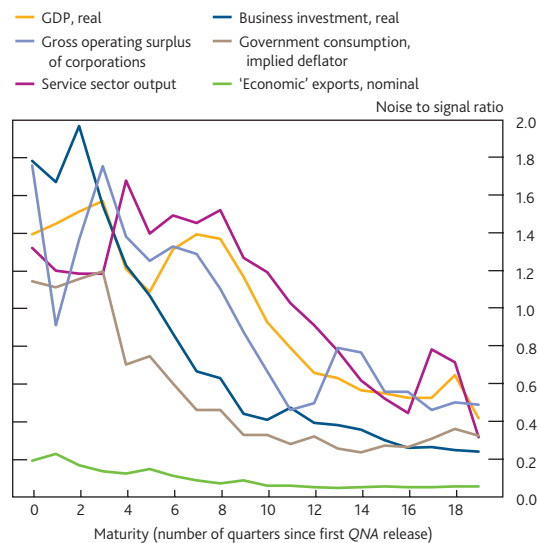
For around half of the variables in **Chart 3**, the noise to signal ratio is above one. Put another way, revisions have been more volatile than the growth profile shown in the latest vintage of the series in question. For these variables, revisions have led to large changes in the published growth rates. This is

particularly marked for estimates of real government consumption growth. This may be because early estimates of government consumption are based on only a small sample of the information that eventually becomes available. And it may also follow from the ongoing methodological changes made to measures of government consumption.

In contrast, the noise to signal ratio across all components of the trade accounts is relatively low. For trade data, revisions have not significantly altered the growth profile shown in early National Accounts releases. In other words, the challenge in interpreting trade data is the volatility of the growth profile rather than the propensity for revision.

Chart 3 shows the extent to which the potential for revisions can cloud the picture painted by the estimates published in the first QNA after each quarter. But as time elapses the ONS is able to incorporate more information so that more mature estimates might be expected to provide a less noisy signal. This fits the experience of past revisions, as shown in **Chart 4**, which plots the decrease in the noise to signal ratio in various published estimates over the five years since initial publication.⁽²⁾

Chart 4 Noise to signal ratio for estimates of quarterly growth of varying maturities^(a)



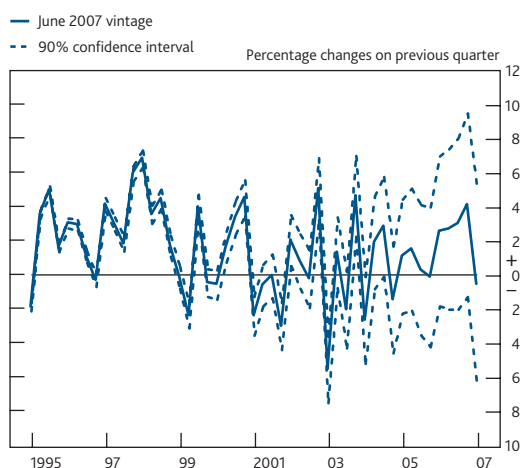
The chart also shows considerable differences across variables. Notably, revisions noise surrounding estimates of business investment has decreased more rapidly than has been the case for service sector output or the gross operating surplus of corporations.

This analysis of historical revisions can be used to estimate a 'confidence interval' surrounding the latest vintage of data. As an example, **Chart 5** plots the resulting confidence interval

(1) At the time of writing, the latest QNA data were those published on 29 June 2007.
 (2) Revisions in **Chart 4** are evaluated over a five-year window from each maturity.

around the June 2007 vintage of quarterly growth in real business investment. Were future revisions to be of a similar magnitude to those observed in the past, there is a 90% probability that any mature data point would fall within the dotted lines. To draw out the decrease in the noise to signal ratio as data become more mature, the chart smoothes through any bumps in the profile shown in **Chart 4**. In doing so, the data are assumed to get better over time until, eventually, they are no longer revised.

Chart 5 Confidence interval around the June 2007 vintage of real business investment



Reflecting the decline in the noise to signal ratio as time elapses, the confidence interval is narrower for estimates of growth some years ago than it is for estimates in the most recent past. And because revisions noise decreases relatively rapidly for business investment, the confidence interval narrows quickly — the estimated variance of revisions declines by just under 50% every five quarters. There is, however, substantial uncertainty surrounding the most recent past. As a result, undue emphasis should not be placed on small changes in the quarterly growth profile shown by the early estimates.

Assessing the likely direction of future revisions

Confidence intervals of this form are helpful in forming an initial impression of the significance of small changes in published growth rates. But they do not give any indication of whether apparent 'news' is more likely to be revised away than it is to be amplified through subsequent revision.

Making fuller use of the available evidence may help shed light on the probable direction of future revisions. In particular economists can appeal to:

(a) Any patterns in past revisions — such as any tendency to revise weak early estimates up; or for revisions to growth

in one quarter to correlate with revisions to growth in the adjacent quarters.

- (b) The indications offered by other measures of activity — such as business surveys or the scores produced by the Bank of England's regional Agents.
- (c) The time-series properties of the data — recognising, for example, that if quarterly growth rates have not been volatile in the past, one should be wary of early estimates that show large quarterly swings.

(a) Patterns in past revisions

The ONS periodically reviews any large revisions to specific components to check whether they were due to predictable factors; and hence, whether data processing procedures should be improved. While it is hard to identify obvious and recurring factors driving specific revisions, there are patterns in revisions to the National Accounts aggregates. Recognising these patterns can help data users (such as Bank staff) to interpret the picture painted by early data releases.

Early estimates have tended to be revised upwards

For most components of the National Accounts, early releases have tended to be revised up more often than down. And upward revisions have tended to be larger than have downward revisions.

Table B shows the mean revision to early estimates of quarterly growth for real GDP and a selection of the output and expenditure components published in the first QNA release (column A). As in **Table A**, it is important to control for the relative volatility in the different variables (column B) to understand how materially any average revisions might affect economic analysis. So, column C shows the mean revision normalised to take account of the variation in the time profile of growth for each variable. On this basis, the tendency to revise up appears to have been most marked for GDP and service sector output.

Table B Direction of revisions to first estimates of quarterly growth for select constant price National Accounts components^(a)

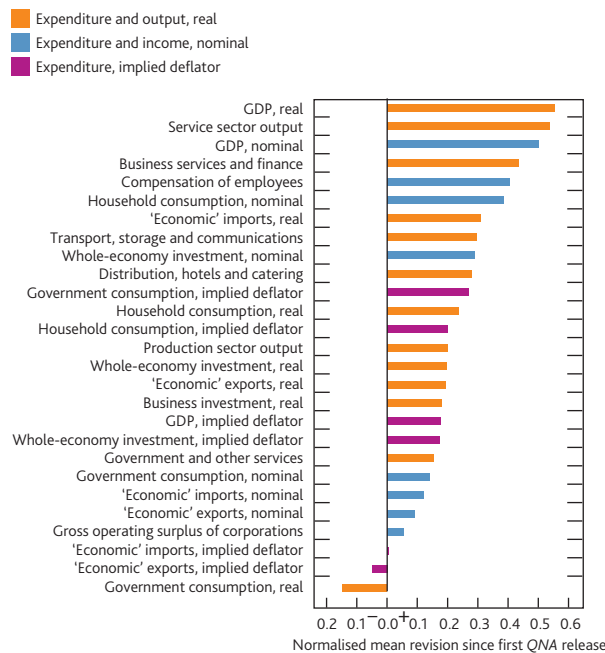
	Mean revision since first QNA release	Variance of growth shown in the latest vintage of data	Normalised mean revision since first QNA release ^(b)
	A	B	C = A/√B
GDP, real	0.15	0.07	0.55
Household consumption	0.12	0.25	0.24
Whole-economy investment	0.34	2.99	0.20
Government consumption	-0.11	0.51	-0.15
'Economic' exports	0.37	3.52	0.19
'Economic' imports	0.45	2.17	0.31
Service sector output	0.17	0.10	0.54
Production sector output	0.15	0.55	0.20

(a) The sample used is as in **Table A**.

(b) Note that figures are rounded so columns A and B may not map to column C.

The pervasive tendency for revisions to move estimates upwards is revealed in **Chart 6**, which shows the 'normalised mean revision' to early estimates of quarterly growth for a wider range of National Accounts aggregates.⁽¹⁾

Chart 6 Direction of revisions to first QNA estimates of quarterly growth^(a)

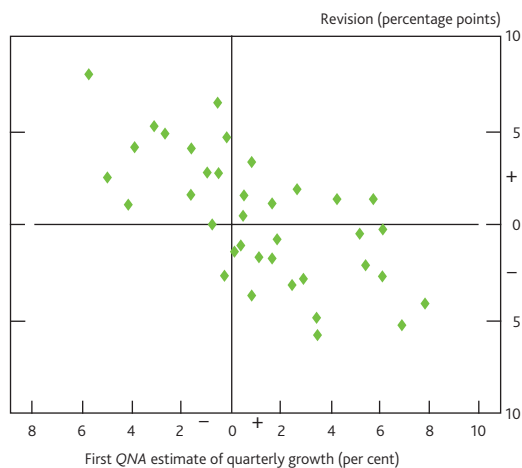


(a) The sample used is as in Table A.

Weak early estimates have tended to be revised up by more than strong early estimates

In general, revisions have been inversely related to the strength of the early estimates. As an example, **Chart 7** plots the first QNA estimates of quarterly growth of real business investment against the revisions to those estimates.

Chart 7 Relationship between first QNA estimates of real business investment and revisions^(a)



(a) The sample used is as in Table A.

The chart reveals a negative relationship between the strength of the first QNA estimates and subsequent revisions to them.

Such negative relationships appear pervasive across the National Accounts aggregates — albeit typically less pronounced than has been the case for real business investment.

Revisions to quarterly growth rates have tended to be partially offsetting from one quarter to the next

Upward revisions to data in one quarter have typically been partially offset — in terms of their impact on the level of the series in question — by downward revisions in adjacent quarters. In other words, negative serial correlation is another pervasive feature of the experience of past revisions to the National Accounts. One corollary of this feature is that early estimates of annual growth have tended to provide a less noisy signal than early estimates of quarterly growth.

(b) Other measures of activity

Although the ONS is the primary source of macroeconomic data for the United Kingdom, it is by no means the only one. The Bank's regional Agents report on the experience of their contacts across the country. And several business organisations publish surveys that provide indications of, for example, output growth and costs for particular industries.

A number of caveats should be borne in mind when interpreting such business surveys: they rely on substantially smaller samples than the official data; and they are typically based on an aggregation of qualitative responses by individual firms. Nevertheless, such alternative indicators can be used to provide a cross-check on early National Accounts estimates — identifying where the early estimates appear most surprising in the light of other available evidence.

The usefulness of the cross-check depends on how closely the indicator has correlated with mature National Accounts data in the past, and on whether there is any doubt that past correlations might break down.

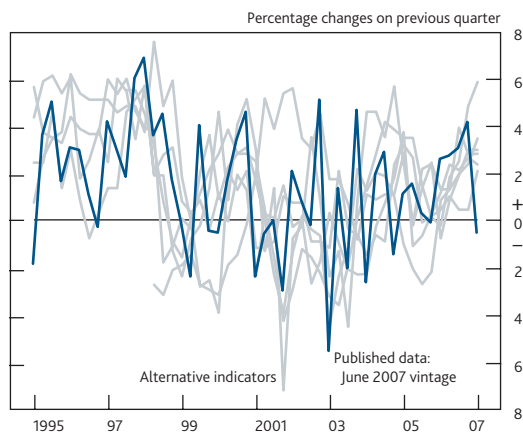
As an illustration, **Chart 8** shows the range of indicators available to help interpret the picture painted by early ONS estimates of business investment growth.⁽²⁾ Each grey line shows the profile of one alternative indicator. The blue line is the June 2007 vintage of the National Accounts. The past experience of revisions suggest that ONS estimates of the recent past are quite uncertain. So the alternative indicators provide a cross-check on the picture of the recent past painted by the blue line. Earlier in the sample, where the ONS

(1) The chart does not provide any statistical test of the significance of average revisions. Garratt and Vahey (2006) find that — over the period 1961–99 — the tendency to revise real GDP growth upwards was statistically significant at the 5% level. There is, however, some evidence that statistical quality improved during the early 1990s.

(2) The alternative indicators are two balances from the *CBI Quarterly Industrial Trends Survey* (the capital expenditure balance, and the proportion of respondents viewing uncertainty about demand as a constraint on investment), quarterly profit warnings, and sectorally weighted investment intentions balances from the *British Chambers of Commerce Quarterly Economic Survey* and the Bank of England's *Agents' Summary of Business Conditions*. The alternative indicators have been rescaled to have the same mean and standard deviation as the published data over the longest available common subsample.

estimates are more mature, the blue line provides a guide to the information content of the alternative indicators.

Chart 8 Range of alternative indicators of real business investment



The chart shows that most of the alternative indicators would be consistent with some recovery in investment growth through 2006 but that the ONS estimates were towards the top end of the range through the year. It also reveals that the survey indicators have not correlated particularly strongly with the profile of business investment growth in the past.

(c) Time-series properties of the data

In gauging where early estimates appear most surprising, economists can also appeal to what they know of the time-series properties of the data. For most macroeconomic variables, growth outturns have tended to cluster around their average. In other words, episodes of extreme growth have been rare. Taken together with the tendency for weak early estimates to be revised up by more than strong early estimates, this might suggest some caution before taking extreme early estimates at face value.

More generally, even before receiving any estimates, economists can draw on past patterns in the data to form a 'prior' view of how they expect the economy to evolve. And, given the uncertainty surrounding early releases it is unlikely to make sense to discard this prior as soon as the first estimates of National Accounts data become available.

Combining evidence from different sources

The discussion above suggests that there are a range of factors to consider when assessing the likely direction of revisions. Cross-checking the official data along these lines is neither new nor unique to the Bank. But approaching this issue formally can add rigour to the exercise of combining such diverse sources of information: helping economists to challenge evidence about different variables in a consistent way.

This sort of exercise is known as a 'signal extraction problem' and its output is a prediction of the profile that will be revealed once the early ONS estimates have matured. Research into this sort of problem is not new: one early example is Howrey (1978), who used a Kalman filter to predict revisions to US disposable income.⁽¹⁾

Following this example, Bank staff have developed a signal extraction model to help predict how far, and in what direction, the latest National Accounts data might be revised. Given the focus on the profile of growth in the past, the exercise might be described as 'backcasting' — as opposed to forecasting — economic activity. Intuitively, the model proceeds in two stages:

- Early ONS estimates are adjusted for any past tendency to be revised up or down.
- The official estimates and survey indicators are used to update a prior view of how the data should evolve.

The degree to which the resulting backcast 'aims off' the early ONS estimates depends on the noise surrounding the early ONS estimates, and the degree to which early estimates are 'surprising'. The 'surprise' in the early estimates is quantified in the light of: (i) past patterns in revisions; (ii) the profile of the survey indicators; and (iii) the time-series properties of the data — that is, the prior view of how the data would evolve. Importantly, the model recognises that revisions to adjacent quarters are unlikely to be reinforcing when there has been significant negative serial correlation in past revisions.

The noise in early ONS estimates reflects both the scale of past revisions to early estimates and the rate at which that noise has dissipated with maturity. In other words, it maps directly from the revisions experience used to estimate the confidence interval in **Chart 5**. The annex to this article explains some further details of the model set-up and its estimation.

The model detailed in the annex develops the work of Ashley *et al* (2005), who used regression analysis to combine information from the latest vintage of ONS data and alternative data sources such as business surveys. One important difference between the two models is that Ashley *et al* (2005) assume that ONS estimates accurately capture the underlying movements in the data once they have been fully balanced (usually around two years after publication of the initial estimate). In practice, ONS estimates remain subject to revision for several years after the initial release, and that is explicitly accounted for in the model used below.

(1) The Kalman filter is a tool for estimating the value of dynamic variables in light of a set of incomplete or noisy measurements. It has a wide range of applications across the physical and social sciences.

Example 1: real business investment

The recent profile of real business investment is considered as an example. As shown in **Chart 5**, the June 2007 National Accounts pointed to a sustained period of growth in real business investment throughout 2006, in contrast to the experience of the preceding five years. However, **Chart 5** also highlighted that early estimates of growth in business investment have been prone to significant revision — a point flagged in recent *Inflation Reports*. And while that noise has tended to decay with maturity, there has still been substantive uncertainty surrounding data a year after the initial release. So how well founded is the picture of recovery in business investment through 2006?

Casual inspection of the past profile of business investment suggests that the reported growth rates during 2006 were by no means unprecedented. And there have been two other episodes over the past fifteen years in which investment has grown for four or more consecutive quarters. But large swings from quarter to quarter have been more typical — in other words, the time-series properties of the data show little persistence in deviations of growth from its average. So, the sustained growth during 2006 that was reported in the June 2007 vintage of data appears unusual. In gauging how far to challenge this profile, the model appeals to the profile of alternative indicators and any patterns in past revisions.

As noted above, most of the alternative indicators would be consistent with some pickup in investment growth during 2006; but perhaps not to the full extent shown in the official data. However, neither the alternative indicators nor the early official estimates have correlated particularly strongly with mature estimates of business investment growth.

Based solely on this evidence, one might be cautious about taking the evidence from the early ONS estimates at face value, and draw only limited comfort from the recovery apparent in the alternative indicators. But simple inspection of correlations between early estimates, surveys and mature data misses an important feature of past revisions — that it has been rare for estimates of strong growth across successive quarters to be revised down.

Chart 9 shows a 'backcast' for real business investment that follows when all these factors are taken into account using the model summarised in the annex to this article.⁽¹⁾ At best, making use of this wide range of evidence can only reduce, not eliminate, any uncertainty. It therefore makes sense to view such backcasts in probabilistic terms, and a fan chart can be used to depict a distribution of possible values of the mature data. Such charts are constructed so that mature ONS estimates would be expected to lie within each pair of purple bands 10% of the time. Consequently, the mature data points are expected to lie somewhere within the entire fan chart 90% of the time.

Chart 9 'Backcast' for quarterly real business investment growth

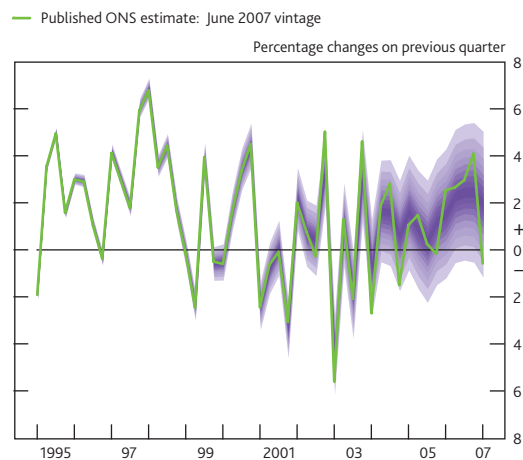
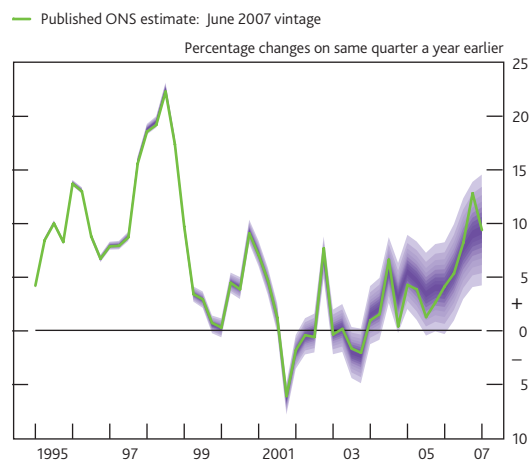


Chart 9 shows the estimated probability distribution for quarterly growth. The centre point of the fan chart is slightly below the published data through 2006, suggesting that downward revisions are more likely than upward revisions. The odds are not, however, extreme and the likelihood that data will be revised far enough to show a fall in investment during 2006 is low. Indeed, the chart also shows a reasonable probability that the profile will be revised to reveal even stronger growth during 2006.

The picture is clearer when the backcast is plotted for annual, as opposed to quarterly, growth rates (**Chart 10**). That chart makes it clear that the estimated recovery in business investment during 2006 is likely to be a robust feature of the data.

Chart 10 'Backcast' for annual real business investment growth



Looking at the experience before 2006, the charts suggest that uncertainty surrounding the backcast decays quite rapidly. This follows from the relatively short half-life of past revisions

(1) Full details are set out in Cunningham *et al* (2007 forthcoming).

to investment data — in common with the confidence interval estimated around the latest official estimates (**Chart 5**). The quarterly backcast fan chart (**Chart 9**) is narrower than that confidence interval, reflecting the in-sample gains from allowing for patterns in revisions, alternative indicators and past patterns in mature data. One important caveat in interpreting these results is, however, that the model behind the fan charts relies on past experience providing a good guide to the future. In practice, this may not always be the case.

Drawing on stories about other related variables

The model used above captures an array of patterns in revisions and dynamics in the uncertain data. But it retains one major simplifying assumption — namely that revisions to one variable are assumed to be independent of revisions to other variables within the National Accounts. However, while Bank staff may track separate indicators for a range of National Accounts components, those components are related by a lattice of accounting identities.

These accounting identities can be used to challenge whether stories about one variable are consistent with stories about other variables. For example, the output and expenditure sides of the National Accounts should balance. So if economists expect upward revisions to household consumption, they must also expect either upward revisions to output components or downward revisions to other expenditure components. Alternatively, any top-down assessment of the likely direction of revisions to a National Accounts aggregate (for example overall service sector output) can be cross-checked with evidence of likely revisions to its components (bottom-up).

It is quite likely that ‘top-down’ and ‘bottom-up’ estimates will give a slightly different impression of the profile of growth. The models may well draw on different business surveys. And, in small samples, the time-series properties of the data and patterns in past revisions may differ slightly. In other words, there is likely to be a ‘residual’ between the bottom-up and top-down estimates. Closer inspection of this accounting residual can help cross-check the top-down estimates. The issue is how far to adjust those top-down estimates in the light of the residual.

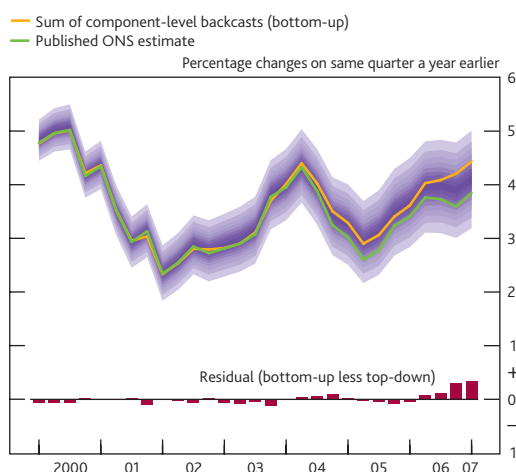
One approach to adjusting the top-down estimate in the light of evidence about prospective revisions to its components is to use a simple rule to allocate any ‘residual’ between backcasts on both sides of the identity. This follows a method first developed by Weale (1985) to produce reconciled National Accounts estimates. The rule used allocates any accounting ‘residual’ according to the degree of uncertainty surrounding the components — the larger the component and the more

uncertain the backcast for a variable, the greater the share of any residual attributed to it. So if the top-down estimates are much less uncertain than the bottom-up estimates then the cross-check will not add much value. But if both are equally uncertain, the bottom-up cross-check may help interpret the picture painted by the aggregate data.

Example 2: services output

As an example, **Chart 11** compares a backcast estimated for aggregate service sector output with the sum of backcasts for its constituent parts — in both cases using the new toolkit described above.⁽¹⁾ The green line shows the profile of the published data, the purple bands show the probability distribution derived from a top-down backcast, and the orange line shows the central (or point) estimate derived from the bottom-up sum of backcasts for the various components of service sector output.

Chart 11 Service sector output

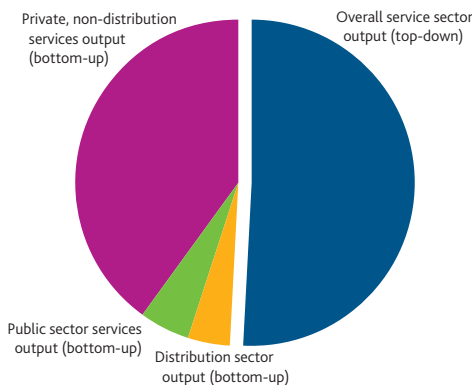


The orange (bottom-up) line is reasonably close to the centre of the (top-down) fan chart, suggesting that any differences between top-down and bottom-up estimates are small relative to the uncertainty surrounding those estimates. But there are periods of discrepancy — as shown by the red bars. For example, in late 2006/early 2007, the disaggregated picture is a little stronger than the top-down assessment — in part due to the strength of some surveys of non-distribution output. The bottom-up estimates suggest growth is increasing while the top-down estimates suggest more of a flattening off.

A consistent picture across aggregate service sector output and its components can only be derived by eradicating any residual. **Chart 12** shows the proportion of the residual that would be allocated to the backcasts of overall service sector output and its various components when their relative uncertainty is used to guide that process.

(1) The constituent parts used here are private non-distribution services output; distribution sector output; and public sector services output.

Chart 12 Allocation of accounting residual across 'top-down' and 'bottom-up' backcasts



If the top-down and bottom-up approaches generated equally uncertain estimates, any residual would be allocated 50–50. In practice, this appears to be the case. So the bottom-up view provides a cross-check on the top-down fan chart.

Conclusion

Bank staff have long recognised the potential for revisions to macroeconomic data and have undertaken a range of research into how best to deal with the ensuing data uncertainty. Early results of that research have been used for some time as part of the toolkit available to the staff when briefing the MPC on developments in output growth. This article describes some further developments in this research effort that were undertaken to refine the staff's toolkit. The aim of this exercise — and the earlier work — is to make the best use of publicly available evidence when interpreting the picture painted by the latest ONS estimates. The model described in this article uses the historical experience of revisions as a basis for estimating how confident one should be in early releases and predicting how far and in what direction those early releases might be revised.

Given the focus on uncertainty, the output of the model is a fan chart outlining the probability distribution across potential revisions. Such charts make clear that one should not place undue emphasis on small changes in growth rates shown in early estimates and that uncertainty may persist for some time.

The techniques described in this article add to the toolkit available to staff when briefing on data that are subject to revision. Bank staff can apply these modelling techniques

when briefing the MPC on recent developments. There is, however a substantial role for economic judgement in gauging how much weight to place on model results.

One natural caveat in interpreting model results is that the statistical methods rely on past revisions as a good indicator of current uncertainty. But this may not always be the case:

- Revisions may become less predictable in the future. In the past, some major changes to statistical practices appear to have led to changes in the patterns of revisions — for example, Garratt and Vahey (2006) found evidence of a structural break in revisions in the years following the *Pickford Report* (Pickford (1989)). Looking forward, successful delivery of the ONS's Statistical Modernisation Programme will enable more timely balancing of National Accounts data from differing sources and facilitate internal reviews of collation procedures.
- Significant methodological revisions in the past — such as the introduction of the ESA 95 accounting framework — may not be representative of current uncertainty. One important judgement in applying models of the type described in this article is, therefore, whether to exclude any past revisions from the analysis.

It is also quite possible that alternative indicators that provided a good mapping to mature ONS data in the past could offer a worse indication in the future — for example if the sample of respondents to a particular business survey becomes unrepresentative.

From all of this it should be clear why it is users of data (such as Bank staff) rather than data providers (such as the ONS) who set up this kind of signal extraction model. The degree to which past patterns in revisions are representative of current uncertainty is an economic judgement rather than a 'hard' statistical fact.

With each major methodological advance in published statistics, Bank staff will need to assess the extent to which past revisions provide a robust guide to ongoing data uncertainty. Close dialogue between users and providers of data is therefore vital to help ensure that use of statistical techniques to extract the signal from uncertain data is founded on a proper understanding of the way in which data are compiled.

Annex

Model of uncertain data

The model is set up to predict the cumulative impact of revisions to the profile shown in the latest National Accounts. It is founded on a representation of the patterns apparent in past revisions. This aspect of the model, termed a 'measurement' system, describes how the latest ONS estimate relates to the 'true' data — assumed here to be the profile that will be revealed once data are sufficiently mature that uncertainty has decayed completely.

This mapping draws on a number of features of the historical revisions experience. It treats the early estimates as equal to the truth plus a term describing the average revision and a measurement error.

$$y_t^{t+n} = y_t + c^n + v_t^{t+n} \quad (1)$$

where y denotes the 'true' data at time t

n describes the maturity of the data – the initial release having a maturity of 1

y_t^{t+n} is the ONS estimate of y at time t released at time $t+n$

c^n is the average revision at maturity n

v_t^{t+n} is the measurement error.

In order to capture the statistical properties of historical revisions, some structure is imposed on the measurement errors. First, serial correlation in revisions is accommodated by expressing errors in the measurement of growth in any period as a function of errors in measures of growth in the previous p quarters.

$$v_t^{t+n} = \sum_{i=1}^p b_i v_{t-i}^{t+n} + \varepsilon_t^{t+n} \quad (2)$$

The model allows for the tendency for measurement errors to tail off as data become more mature. The variance of measurement errors is assumed to decay as maturity increases — in line with the treatment used to estimate the confidence interval shown in **Chart 5**.

$$\sigma_{v_n}^2 = \sigma_{v_1}^2 (1 + \delta)^{n-1} \quad (3)$$

where the variance at maturity n (that is, $\sigma_{v_n}^2$) is a function of the variance at the initial release and δ , the rate of decay in the revisions variance [$-1 < \delta \leq 0$]. The average revision in equation (1) also decays with maturity in a similar way, so that the ONS estimates are assumed to converge on the 'true' data eventually.

A further measurement equation describes the relationship between any alternative indicators and the 'true' data. The measurement errors associated with alternative indicators are modelled far more crudely than the uncertainty surrounding the official estimates — assuming a constant mapping between the indicators and the mature data.

$$y_t^s = c^s + Z^s y_t + v_t^s \quad (4)$$

where y_t^s is the alternative indicator

c^s and Z^s describe the relationship between the rescaled indicator and the true data

v_t^s is the measurement error.

Under this representation, there is assumed to be no improvement in measurement as indicators become more mature — after all, surveys are not typically subject to revision. And, for simplicity, the model does not allow for any serial correlation in measurement errors for survey indicators. This simplifying assumption may not always be warranted, motivating careful thought about the relationship between indicators and mature data when interpreting model results.

The final leg of the model is a description of the time-series properties of the 'true' data — termed a 'transition equation'. The transition equation helps establish the degree to which early estimates are 'surprising' in light of past experience — for example, whether large swings in the data have been common in the past. A simple autoregressive model is used to describe the properties of the 'true' data:

$$y_t = \mu + \sum_{i=1}^q A_i y_{t-i} + e_t \quad (5)$$

The model is estimated in two steps:

- The first step is to estimate parameters driving the revisions process, using real-time data.
- The second step is to estimate the remaining parameters using the latest available vintage of ONS data and any alternative indicators. In doing so, the model allows for any past correlation between measurement errors and the growth rates revealed by the mature National Accounts data.

The model's output is a profile of the 'true' data — the backcast — that is consistent with these parameters and the latest profiles shown by the official estimates and survey indicators. Full details are set out in Cunningham *et al* (2007 forthcoming).

References

Ashley, J, Driver, R, Hayes, S and Jeffery, C (2005), 'Dealing with data uncertainty', *Bank of England Quarterly Bulletin*, Spring, pages 23–29.

Atkinson, T (2005), *Atkinson Review: Final Report: Measurement of government output and productivity for the National Accounts*, Palgrave Macmillan.

Beadle, J (2007), 'Modernising the UK's national accounts', *Economic & Labour Market Review*, Vol. 1(4), pages 32–38.

Cunningham, A, Eklund, J, Jeffery, C, Kapetanios, G and Labhard, V (2007), 'A state space approach to extracting the signal from uncertain data', *Bank of England Working Paper*, forthcoming.

Garratt, A and Vahey, S (2006), 'UK real-time macro data characteristics', *The Economic Journal*, Vol. 116, pages F119–F135.

Howrey, E (1978), 'The use of preliminary data in econometric forecasting', *The Review of Economics and Statistics*, Vol. 60(2), pages 193–200.

Jääskelä, J and Yates, A (2005), 'Monetary policy and data uncertainty', *Bank of England Working Paper no. 281*.

Lomax, R (2004), 'Stability and statistics', *Bank of England Quarterly Bulletin*, Winter, pages 495–501.

Obuwa, D and Robinson, H (2006), 'Revisions to quarterly GDP growth and its production (output), expenditure and income components', *Economic Trends*, Vol. 637, pages 28–39.

Pickford, S (1989), *Government economic statistics: a scrutiny report*, Cabinet Office, HMSO.

Statistics Commission (2004), 'Revisions to economic statistics: report by Statistics Commission', *Statistics Commission Report No. 17(1)*.

Weale, M (1985), 'Testing linear hypotheses on National Account data', *Review of Economics and Statistics*, Vol. 67(4), pages 685–89.