

The output gap and its role in monetary policy decision-making

Florin Citu and James Twaddle, Economics Department¹

Understanding the link between the real economy and inflation is essential to monetary policy formulation. In conventional macroeconomics, the concept of the output gap is an important component of that link. Used in this context, the output gap provides a useful way of thinking about inflationary pressure in the economy. This article discusses the output gap concept, its strengths and weaknesses, and how it fits into the monetary policy process at the Reserve Bank.

While the output gap is a useful device in assisting the understanding and forecasting of inflation developments, it does have weaknesses. The output gap is unobservable, and developing reliable estimates of it is a key weakness. A second is the possibility that the linkage between the real economy and inflation is not well represented by models or frameworks that use the output gap concept. Being alert to these weaknesses, the Reserve Bank looks at a range of other indicators of inflationary pressure when formulating monetary policy and uses judgement where appropriate.

1 Introduction

Monetary policy directed at maintaining low and stable inflation uses interest rates to lean against persistent inflationary and deflationary pressures. The changes that a central bank makes to interest rates take several quarters before they have their largest impact on inflation. This is because it takes time for people to react to the interest rate changes, and for peoples' reactions to affect inflation. These lags between interest rate changes and changes in inflation mean that a central bank needs to have a view on the inflationary pressures that are likely to prevail in the future so that they can decide the appropriate setting for interest rates today.

There are a number of factors that will influence inflation in the future, such as people's expectations of future inflation, exchange rate movements, wage developments, and import price changes. Empirically, one of the most important of these influences seems to be the state of the business cycle. When the economy is going through a period in which resources are underutilised there tends to be more disinflation pressure than inflation pressure. Conversely, when the

economy is going through a period in which resources are heavily used, the balance of pressure tends to push inflation up.

Accordingly, to formulate the appropriate monetary policy today we need a measure of the state of the business cycle that shows the strain that current economic activity is exerting on resources. There are a number of indicators of such strain. For example, the number of job advertisements appearing in newspapers or the difficulty that businesses are having in finding appropriately qualified workers tell us about the strains in the labour market, and the amount of time machines are idle tells us about strains on equipment.

The output gap is an aggregate measure of resource strain in the economy. It has long been a device used at central banks to represent how "hot" or "cold" the economy is at any particular time and to forecast likely inflationary pressures. It can also be a useful indicator about the volatility of activity in the economy, whether resources are alternating quickly between periods of substantial resource strain to periods of substantial resource slack, or whether the economy is moving smoothly between moderate levels of resource strain and slack.

¹ We are grateful for many helpful comments from our colleagues at the Reserve Bank. Particular thanks go to David Archer and Geof Mortlock.

At the Reserve Bank, the output gap forms an important part of our framework for thinking about inflation developments. However, the output gap is not a perfect indicator of future inflationary pressures and we do not treat it as such. It cannot be directly measured, it only tells us about aggregate resource use, and its link with inflation is not always stable. Thus, we use the output gap with the appropriate degree of caution about the signals provided.

This article discusses a range of issues relating to the output gap and the role it plays in the Reserve Bank's approach to monetary policy formulation.² The following two sections discuss the output gap concept and the methods available for estimating the output gap. We then place the output gap in the context of policy-making at the Bank before reviewing our method for estimating the output gap. Section five covers weaknesses with the output gap for forecasting inflation. Section six discusses how judgement forms an important part of our application of the output gap concept.

2 The output gap concept

A matter of "hot or "cold?

Most observers of the economy can readily relate to the notion that, when an economy is operating at cyclical lows, prices are more likely to drift down, and vice versa. This is the economy-wide counterpart of the familiar story of tradespeople being prepared to shave prices when the next job is weeks away, but seeking a higher price to attend to your job when the waiting list is long.

One way to represent whether the economy is running "hot" or "cold" is to use the output gap. The output gap is the difference between actual output produced in the economy over a given period and the "normal" or "trend" level of output produced in the economy (often called "potential" output). Normal or trend output in this conception is the level of output that is consistent with stable inflation – i.e. the level of output produced when the economy is neither

too hot nor too cold.³ Trend output is not constant; it evolves with the economy's structural capacity to produce goods and services. However, there is no observable direct measure of trend output, nor do we know how trend output will develop in the future.

An alternative approach for gauging the temperature of the economy is to focus on the growth rate of output, comparing it with its trend growth rate. Sustained periods of above-trend growth can be associated with above-average inflation pressures, while sustained periods of below-trend growth can be associated with unusual disinflation pressures. Here also we have the problem of not knowing what the trend growth of the economy is, and how it will develop in the future. But there is a major difference between the two approaches, which is the source of both the relative strength and weakness of the output gap concept.

The output gap concept works in *levels* (of actual output and trend output), and so is influenced much more by the history of the components than is the growth approach. With the output gap, one might be saying: "the economy has been growing rapidly for a while, but it is coming out of a slack period and there are still underutilised resources, so there is no inflation threat at present". The validity of this statement depends crucially on whether previously-estimated normal rates of resource utilisation – or equivalently, normal or trend output – remain the valid reference point. By contrast, a focus on the *growth rate* would in the same circumstances suggest that, if the economy was coming out of a recession and growing above-trend, inflationary pressures would be intensifying, even if there were still machines standing idle and large numbers of employable people searching for jobs.

Consider the following scenario. Suppose that the economy has been through a "W" shaped growth period (a recession, a short burst of growth, a secondary lull, and then a recovery). Assessing the output gap in the recovery phase depends on estimates of the cumulative amount of slack in resource utilisation through the whole growth period, not just the current growth phase. In contrast, with a growth-focused


² An earlier Bulletin article looked at the output gap concept, how we estimate it, and issues relating to projecting future trend output (Conway and Hunt (1998)). Here, we focus more on how the output gap fits into the policy-making process. We will look at the use of the output gap in forecasting inflation but will not address issues around the optimal monetary policy response to the output gap.

³ Some commentators draw a distinction between potential and trend output. For simplicity, we use the terms interchangeably here to mean the level of output consistent with stable inflation.

approach it does not matter how big or small was the short burst of growth that formed the middle of the W – the focus is purely on the current growth rate relative to trend. Therefore, when comparing the actual growth rate of the economy with its trend growth rate, there is a need to explicitly take into account the possibility that above-trend growth may be occurring as a result of spare capacity in the economy and therefore might not necessarily pose a threat to inflation.

The strength of the output gap concept is that the intensity of resource utilisation does seem to matter for inflation pressures, so the history of resource use is likely to be important. One can see this from Figure 1, which shows the Bank's estimate of the output gap and variations in the non-tradables component of CPI inflation (that component being closer to the idea of domestically-sourced inflation than the total CPI, and hence more closely linked to the output gap concept). There appears to be a relatively strong and broadly stable lagged relationship between the estimated output gap and inflation. Given this apparent connection, using the output gap as an indicator of future inflation is conceptually a better choice for monetary policy formulation than using growth rates.

Figure 1
Output gap and non-tradable CPI inflation



That's the strength. What about the weakness? The output gap approach relies on the ability to identify the normal productive capacity of the economy (through all stages of a business cycle and taking into account productivity changes) and to assess the extent to which, at any given point in time, there is currently spare capacity in the economy. As we go on to explain, reliably assessing trend output and spare capacity in the economy is easier said than done.

3 Approaches to estimating the output gap

We can categorise the estimation methods into three groups. The first group tries to infer trend output by looking at the path of actual output. The second group attempts to isolate the effects of structural and cyclical influences on output, using economic theory. And finally, survey data can be used to infer the extent of excess demand or supply in the economy.

An approach used in the early days of output gap estimation was to fit a linear trend to the output data, following from the thought that actual output cycles around a trend level that would be unlikely to vary significantly. This approach relies on the level of trend output, and the economy's capacity to produce goods and services, growing at a constant rate – ie it assumes away changes in typical productivity growth and technological developments, or changes to the composition of the economy, that could bring about significant and long-term changes to the trend growth rate in the economy. The danger of relying on this assumption was shown by the experience of the 1970s, when unexpected and unrecognised declines occurred in trend growth rates in industrial countries. Partly because of the linear trend used to extrapolate normal output, policy-makers believed that slower output growth rates implied that they were in a disinflationary period, and that loosening monetary policy would reduce unemployment and generate higher output growth without increasing in inflation. With the benefit of hindsight, we know that during this period resources were not as underutilised as had been thought, and that the easier monetary conditions lead to higher inflation.

Today, instead of presuming that the trend of output growth evolves in an unchanging manner, the standard technique is to approximate the evolution of output growth using flexible trends derived using "filters".⁴ Like a sieve that lets fine grains through while retaining the larger ones, a filter aims to identify persistent movements in output while discarding the transitory or cyclical parts. A commonly used filter is the

⁴ Filters are statistical techniques that decompose variables into permanent, cyclical, and sometimes noise components.

Hodrick-Prescott (HP) filter.⁵ While some filters use only past historical output data to determine trend output, the HP filter looks at both past and future data.

Figure 2
Output, linear and flexible filters
 (Chain-linked 1995/96 prices)




Figure 2 illustrates the difference between a linear trend and a flexible filter. In 1995/1996 actual output was below the linear trend estimate of trend output, indicating a negative output gap. At this same point in time actual output was above the flexible filter estimate of trend output, indicating a positive output gap. The linear trend fails to recognise the slowing in trend output growth in the late-1980s, a period when there was a lot of scrapping of plant and equipment and low labour force growth. By assuming the economy's capacity to produce goods and services grew at a constant rate through the early to mid-90s, the linear trend would suggest that there were plenty of underutilised resources, and that the economy could grow quickly without generating inflationary pressures. In fact, during the mid-90s a number of measures indicated resources were under strain, and inflation from the domestically-focused sectors of the economy was rising (ie there was a positive output gap).

Simple flexible filters only look at actual output data in determining trend output, so they only pick up changes in trend output when they have started to appear in actual output. Other filtering approaches make use of both flexible filters and information from economic theory and other indicators of resource strain in the economy. These methods are often called semi-structural methods because they make use of information about the economy's underlying structure. Semi-structural methods use, for example, the economic

theory that suggests low utilisation of plant and machinery would coincide with a negative output gap to help estimate the output gap. The multivariate (MV) filter that we use for estimating trend output is a semi-structural method (the MV filter is covered in more depth in the next section).

Using even more economic theory are methods that seek to build up a view of how much the economy can produce without inflationary pressure by looking at the things that go into producing goods and services – labour, plant and machinery (capital), and the productivity of these inputs. A production function – an equation that links inputs like labour and plant and machinery to total output – is then used to translate the trend levels of these inputs into an estimate of trend output.

The production function approach has the benefit that it splits trend output developments into different components, such as labour force developments or additional investment in plant and machinery. However, although in theory this approach might pick up changes in the economy's ability to produce goods and services more readily than filtering approaches, this is not necessarily so. There are still key unobservable variables, such as businesses' target level of capacity utilisation, and future advances in technology. Moreover, the variables that are observable, such as plant and machinery, can also be difficult to measure. In practice, to estimate the unobservable variables, economists often extrapolate past trends as a gauge to possible future trends and often end up using the same filtering techniques discussed above for the different production function inputs.

The final approach to estimating potential output and the output gap is based on asking businesses how far away from "normal" is their utilisation of existing capacity. Survey responses are used to construct measures such as "capacity utilisation" that indicate the state of cyclical activity and how much strain resources are under. Commonly, these surveys ask businesses about their ability to expand production without additional cost, or how hard it is for firms to find the labour that they need.⁶

⁵ See Hodrick and Prescott (1997).

⁶ Examples in New Zealand include the NZIER's Quarterly Survey of Business Opinion series on capacity utilisation, and difficulty finding skilled and unskilled labour, and the ANZ job ads survey.

As with the previous methods of output gap estimation, survey-based measures of the output gap are also subject to a number of measurement difficulties. For example, different firms might interpret survey questions differently; only a small number of firms might respond to the surveys (making it difficult to reliably extrapolate the data for the economy as a whole); surveys are based largely on firms producing manufactured goods and do not necessarily capture other parts of the economy, etc. Thus, although survey measures provide a potentially useful source of information on the state of the economic cycle, they should not be interpreted as a definitive measure of the output gap.

With the “true” output gap unobservable, we must look at other properties of different output gap estimation methods to assess whether one output gap estimation method is better than others. Typically there are three ways to assess different estimation methods⁷:

- assessing the inflation forecasting ability of output gaps based on different estimation methods⁸;
- assessing whether particular output gap estimates match business cycle regularities (eg length and amplitude of cycle); and
- assessing the extent to which output gap estimates derived from different methodologies are revised as more data become available.

None of these properties should be looked at to the exclusion of the others in deciding which output gap estimation method to use, given that they sometimes need to be traded off against each other. For example, an output gap from one method may prove to be good at forecasting inflation with the ex-post revised data, but may perform poorly if it were estimated on unrevised data available at the time the forecast had to be made.

The diversity of methods being used by central banks and others illustrates that there is far from total agreement about which method is best, and also that using a number of

indicators of resource strain can help in understanding the state of activity in the economy. Some institutions, such as the OECD and European Union, use a production function as their main method for estimating the output. Other central banks, such as the Bank of Canada and the Reserve Bank of New Zealand, use semi-structural methods that rely on both filters and structural information about the economy.

4 How the output gap fits into the Reserve Bank policy-making process

The output gap in context

The output gap has its largest role in the Bank’s policy-making process as an input into the Bank’s economic projections, which form an important part of each *Monetary Policy Statement* and inform our Official Cash Rate (OCR) decision-making.

We use the Forecasting and Policy System (FPS) model to help develop our projections of future developments in the economy.⁹ The output gap is an important input into FPS because in FPS it is the key way that domestic activity influences inflation. The output gap is linked to inflation by a Phillips curve. A Phillips curve is an economic relationship that relates inflation developments to how “hot” or “cold” the economy is. When the output gap is positive, the Phillips curve will translate that strain on resources into inflationary pressure, and when the output gap is negative into pressure drawing inflation down. The inflationary pressures coming from domestic activity are then combined with inflation coming from other sources captured in FPS, such as the exchange rate and inflation expectations, to generate an outlook for inflation.

Producing projections is more complicated than simply putting numbers into a computer-based model and allowing the computer to generate final results. Models are simplifications of an inherently complex economy, and only capture in an aggregate form some of the main relationships we think are at work, making judgement an integral feature

⁷ Claus, Conway, and Scott (2000) use some of these techniques.

⁸ If, as we discuss later in the article, the link from the output gap to inflation were to weaken, then looking at the inflation forecasting ability of different output gap estimation techniques may not be particularly useful.

⁹ See Reserve Bank of New Zealand (1997) for a summary of FPS, Drew and Frith (1998) for how FPS is used to prepare projections.

of their use. There is a lot of information that cannot be captured by models, and the policy-making process recognises this by looking at a range of other information, such as inflation forecasts from alternative ways of thinking about the causes of inflation, information learnt from our business visits, and views about the balance of risks around a particular OCR decision.

How the Reserve Bank estimates the output gap¹⁰

Our main method for estimating the output gap uses the multivariate (MV) filter, a filter that is based around the common HP filter that we mentioned earlier. The HP filter determines trend output at a particular point by looking at past and future output in order to determine the persistent component. For example, if past and future output are low, then the filter will conclude that current trend output is low.

Because the filter uses both past and future data, there is a problem at the end of a sample period because there is no future data – something often referred to as the “end-point problem”. In this situation the filter only uses past data. When the future data become available, the filter may have a different estimate of trend output for that point in time. Frustratingly, the end-of-sample estimates are the most important for monetary policy because these are the ones we use for forecasting inflation.

The MV filter differs from the HP filter because it makes use of information from other indicators of resource strain (it is ‘semi-structural’). For example, we include information on firms’ capacity utilisation, because when capacity utilisation is high it generally indicates that resources are under strain and that there is a positive output gap. Using similar intuition, we also include information from the labour market – when unemployment is low then firms’ ability to expand production may be limited and existing employees may have to be induced by higher pay rates to work longer hours. The final piece of information we use to supplement the HP filter is inflation data. As we noted earlier, a positive output gap generally leads to inflation drifting up, so when inflation is high it can often indicate that we have been experiencing a

positive output gap. If the link from these other indicators to the output gap is stable then the indicators can reduce the severity of the end-point problem because the information from these indicators does not tend to be revised.¹¹

The last component of the MV filter is a constraint that is imposed so that estimates at the end of the sample are less likely to be revised because of the end-point problem. The constraint takes the form of a trend output growth assumption that the filter uses to help decide what trend output is. Aside from purely mitigating the end-point problem, the constraint is a tool for adding information not included in our filter into our trend output estimates. For example, if we were witnessing a large surge in immigrants coming to New Zealand, it is likely that they would add to our capacity to produce goods and services and reduce the strain on resources. In this example, we could change the constraint so that our trend output estimates better took account of the migration surge.

We have made some small changes to the MV filter since we began using it in 1997, the main one being an increase in the importance of the end-of-sample constraint. We have increased the importance of the constraint to further mitigate the end-point problem and to attach greater weight to our judgemental adjustments. However, actual output data remain the largest determinant of our trend output estimates through the HP filter component.

The MV filter produces the trend output estimate for the current period (ie the period prevailing at the time a set of forecasts is prepared). For the Bank’s economic projections it is necessary to project a measure of trend output into the future. The MV filter is not used for this purpose because there is no future data to filter. Instead, a production function approach is used. We take FPS projections of labour and capital, combine them with a view on how productively they will be used in the future, and use a production function to form our projections of trend output.

¹⁰ For a more detailed summary of the MV filter see Conway and Hunt (1998).


¹¹ Conway and Hunt (1998) show that the end-point problem is less pervasive in the MV filter relative to the HP filter because of the addition of the other indicators (the ‘structural’ information).

5 Weaknesses with using the output gap

Revisions to output gap estimates

With no “official” measure of the output gap, there is some uncertainty about how well our current output gap estimates represent the resource strains that the current level of activity is creating in the economy. Often in subsequent quarters, as we get more and better data, and as our understanding develops of how the economy is evolving, our output gap estimates are revised. Figure 3 illustrates the revisions that output gap estimates can undergo, plotting our output gap estimate from each *Monetary Policy Statement (MPS)* since we began using FPS, referred to as ‘real-time’ estimates, against our revised output gap estimates for that period from the November 2002 projection. There can be large differences between real-time and current best output gap estimates (eg in the second half of 1999 the estimates had different signs).

Figure 3
Output gap estimates



Output gap estimates can be revised for three reasons:

- Data revisions and lags. First-release output data are often revised in subsequent quarters as the statistical agency gets better estimates of what took place in the economy. In addition, first-release data are available only with a lag that forces us to make estimates of output for those quarters where we lack data.
- The end-point problem. Many filters use past and future output to determine trend output. At the end of the sample there is no future data, forcing them to look only at past data. Subsequently, when we get the future data and re-filter the output data, we may get a different

estimate of trend output from that calculated using only the past data.

- Structural change and judgement. Erroneous output gap estimates can occur if the economy is changing in a way we do not know about – eg such as productivity growth rate changes brought about by technical innovation. Equally, the judgement we introduce into our output gap estimates may be incorrect and subsequently be revised.

The extent and source of revisions to real-time output gap estimates will depend on the estimation method. For example, structural methods are more susceptible to revisions that are caused by changes in the structure of the economy, while simple filters will be more susceptible to revisions caused by the end-point problem.

By giving the MV filter different vintages of data we can learn which sources of revisions – data lags, data revisions, or the end-point problem – are most important under the MV filter method that we use. Revisions because of data lags show up as the difference in output gap estimates made using the data available at a particular time and subsequent estimates made with first-release (but not final) output data. Similarly, revisions caused by the end-point problem will be the difference between output gap estimates made using the latest revised or final data available up until a particular time and estimates made using the whole span of output data.

Figure 4 plots the real-time MV filter output gap estimates at each point in time if we had been using the current version of the filter since late 1994, against the latest estimates.¹² Figure 5 decomposes the difference in estimates into data lags, data revisions, and the end-point problem.

¹² The series plotted are not exactly the same as those in figure 3 because of simplifying assumptions, the most important of which are: (i) we ignore lags and revisions in the structural information, though revisions to this information are not likely to be a large cause of revisions to output gap estimates; (ii) no weight is given to the end-of-sample constraint because we do not know what we would have set it to if we had been using the MV filter in real-time over the whole of this period.

Figure 4
Real-time versus latest output gap estimates





Figure 5
Contributors to revisions to output gap estimates



Through the earlier part of the sample, the end-point problem is a larger source of revisions than in the latter half of the sample. For the output gap estimates in the second half of the sample there is less future data, so those estimates are based more on past data. When more future data become available, the estimates in the latter half of the sample will be revised by more. A muted version of the same pattern is seen from data revisions.

Taking account of the fact that revisions from the end-point problem are likely to be understated in the second half of the sample, the end-point problem is a slightly larger cause of output gap revisions than data revisions or data lags, a result common to a wide range of filter-based techniques (Orphanides and Van Norden (2002)).¹³

Output gap revisions due to output data lags and revisions do not tend to be as persistent as those due to the end-point problem, though they are often as large (eg late 1997/early 1998). At times, these revisions tend to be offsetting, meaning that the output gap revision due to data revisions and lags in aggregate are likely to be smaller than their individual parts.

The decomposition referred to above uses the latest version of the MV filter, and assumes that the economy has not undergone any major structural change over the period. It gives us no estimate of the impact of structural change as a source of revisions. While, over the relatively stable mid to late-1990s, structural change is unlikely to have been a big issue, over periods of major economic evolution and reform, as in the 1970s and 1980s, structural change can impede reasonable estimates of the output gap.

Work with US and UK output gap estimates has suggested that poor real-time output gap estimates during the 1970s, largely because of the use of the linear trends we discussed earlier failing to detect a trend growth slowdown, may have contributed to the higher inflation that occurred over this period.¹⁴ These results reinforce the point that, when there is large structural change, or major economic reforms, real-time output gap estimates are more susceptible to error.

The effects of structural change often take time to appear in official data, and therefore in output gap estimates. Sometimes in these situations policy-makers can detect the changes going on in the economy faster than output gap estimation techniques, and so by adding judgement they can improve output gap estimates. However, the judgemental adjustments of monetary policy-makers may also contribute to incorrect estimates of the output gap.

The link from output to inflation

Another potential risk in using the output gap for formulating monetary policy is that the relationship between the output gap and inflation probably changes over time. This problem has become more pronounced as central banks around the world have become more successful at controlling inflation.

¹³ If we had used the end-of-sample constraint and set it to the average trend growth rate over 1994-2002, the revisions attributable to the end-point problem would have been reduced and revisions from other sources would have increased slightly. However, doing this would have introduced information not known in 'real-time'.

¹⁴ See Orphanides (2001) for the US, and Nelson and Nikalov (2001) for the UK.


Throughout the 1990s, industrial countries have experienced very low variability in inflation rates, while the real economy – and hence the output gap – has continued to follow a conventional business cycle profile (albeit perhaps a little more damped than in earlier times).

We do not know whether such a change in the relative variability of inflation and the output gap would be associated with a fundamental breakdown in the relationship between the business cycle “temperature” and inflation, or instead the enhanced role of offsetting monetary policy actions. Another way of putting it is that we do not know whether the apparent change in the relationship reflects a decline in the inflation pressure associated with a given output gap, or whether the pressure remains in latent form.

Interestingly, there is mixed evidence in New Zealand of a weakened relationship between the output gap and inflation. Figure 1 suggested that over the 1990s there was a stable relationship between the output gap and inflation, and Claus (2000) finds that output gaps from different estimation techniques are still useful indicators of inflation.

The stable relationship between the output gap and inflation also holds if, instead of using our latest revised output gap estimates, we use our real-time, unrevised output gap estimates. Figure 6 shows that, while the link between real-time output gap estimates and inflation is not as strong as between our latest revised output gap estimates and inflation, there is a relatively strong relationship.

Figure 6
Output gap estimates and non-tradables CPI inflation




6 Uncertainty and the output gap

With the uncertainties that surround the output gap and the link from the output gap to inflation having the potential to give a misleading outlook for inflationary pressures, the Bank avoids using the output gap mechanically. We look at a range of indicators of resource strain, such as capacity utilisation and labour shortages, rather than rely solely on our estimates of the output gap. In addition to providing other estimates of the output gap, these indicators can also help give a disaggregated picture of resource strain, which the output gap, as an aggregate measure, does not show. For example, other indicators can tell us whether a period of resource slack is primarily falling on labour or plant and machinery use. A number of the alternative indicators of resource strain are not revised, which provides a cross-check on our real-time, unrevised output gap estimates.

Since the early 1990s, a number of these indicators have generally followed similar patterns and so pointed to similar amounts of resource strain (figure 7), though on occasion the indicators have given conflicting messages (eg in 1997). When producing projections, these additional indicators of resource strain are compared to our output gap estimates and we can refine our output gap estimates using judgement if we feel the alternative indicators tell a more compelling story.

Figure 7
Indicators of resource strain



To more fully understand the implications of output gap uncertainty, we frequently run alternative scenarios that look at different output gap estimates. So, for example, if there was the risk that the output gap could be more positive than we had incorporated in our central scenario, we could put

the higher output gap estimate into our projections as an alternative scenario to see its implications for future output, inflation, interest rates, etc.

In the end, monetary policy decisions are far from being based solely on the estimated output gap; they take account of the balance of risks around the Bank's central view of the economy. For example, greater weight might be attached to the inflation signals provided by an estimate of the output gap if all other resource constraint indicators showed similar indications of strain, and less weight if these other indicators showed substantially different degrees of constraint. A very broad radar screen is used when assessing the many factors that can influence inflationary pressures.

7 Conclusions

The output gap is a conceptually useful way of thinking about the inflationary pressure coming from the domestic economy. It also appears to have broad empirical relevance. However, the output gap can be difficult to measure and reflects only aggregate resource pressures, which could readily lead to policy errors if the output gap were mechanically used as the sole indicator of inflation pressure. At the Bank, we use the output gap as only one of several inputs to the policy-making process. We look at a range of indicators, debate them, and are willing to exercise judgement when we feel it produces more plausible estimates of inflation pressure.

It is frequently noted that economics, and particularly monetary policy, is both an art and a science. Our use of the output gap shows both of these attributes; we estimate the output gap, but supplement our estimates with judgement where we feel the output gap has weaknesses.

References

Bank of Canada (2003), 'Indicators of Capacity and Inflation Pressures for Canada,' Bank of Canada Website: <http://www.bankofcanada.ca/en/indinf.htm>

Butler, L (1996), 'The Bank of Canada's New Quarterly Projection Model, Part 4. A Semi-Structural Method to Estimate Potential Output: Combining Economic Theory with a Time-Series Filter,' *Bank of Canada Technical Report* No. 77

Claus, I (2000), 'Is the output gap a useful indicator of inflation,' *Reserve Bank of New Zealand Discussion Paper* No. 05.

Conway, P, I Claus, and A Scott (2000), 'The output gap: measurement, comparisons, and assessment,' *Reserve Bank of New Zealand Research Paper* No. 44.

Conway, P and B Hunt (1998), 'Estimating the potential output of the New Zealand economy,' *Reserve Bank of New Zealand Bulletin* 61(3) (September) pp192-202.

Drew, A and M Frith (1998), "Forecasting at the Reserve Bank of New Zealand," *Reserve Bank of New Zealand Bulletin*, 61, 4, pp 317-27.

European Central Bank (2000), 'Potential output growth and output gaps: concept, uses and estimates', *European Central Bank Monthly Bulletin* (October) pp 24-31.

European Policy Committee (2001), 'Report on Potential Output and the Output Gap', *European Union* (October).

Hodrick, R and E Prescott, (1997), 'Post-war U.S. business cycles: An empirical investigation', *Journal of Money Credit and Banking*, Vol 29, no 1.

Nelson, E and K Nikolov (2002), 'UK inflation in the 1970s and 1980s: the role of output gap mismeasurement,' Bank of England Working Paper no.148.

Orphanides, A (2001), 'Monetary policy rules, macroeconomic stability and inflation: a view from the trenches,' Federal Reserve Board Finance and Economic Discussion Series, 2001-63.

Orphanides, A and S Van Norden (2002), 'The unreliability of output gap estimates in real time,' *Review of Economics and Statistics* 84(4) pp569-583.

Reserve Bank of New Zealand (1997), 'The Forecasting and Policy System: an introduction,' *Reserve Bank of New Zealand Bulletin* 60(3) pp 225-35.