BANK OF CANADA
BANQUE DU CANADA

Discussion Paper/Document d'analyse
2015-1

Measuring Potential Output at the Bank of Canada: The Extended Multivariate Filter and the Integrated Framework

by Lise Pichette, Pierre St-Amant, Ben Tomlin and Karine Anoma

www.bank-banque-canada.ca

Bank of Canada Discussion Paper 2015-1

January 2015

Measuring Potential Output at the Bank of Canada: The Extended Multivariate Filter and the Integrated Framework

by

Lise Pichette, Pierre St-Amant, Ben Tomlin and Karine Anoma

Canadian Economic Analysis Department
Bank of Canada
Ottawa, Ontario, Canada K1A 0G9

Bank of Canada discussion papers are completed research studies on a wide variety of technical subjects relevant to central bank policy. The views expressed in this paper are those of the authors.
No responsibility for them should be attributed to the Bank of Canada.

ISSN 1914-0568

© 2015 Bank of Canada

Acknowledgements

The authors would like to thank Mikael Khan, for his contribution to early versions of this work, and Dany Brouillette, Agathe Côté, Glen Keenleyside, Sharon Kozicki, Natalia Kyui and Stephen Murchison for their comments and suggestions. They also thank participants in Bank of Canada internal seminars for useful feedback and Cassandra Cao and Mohanad Salameh for their research assistance.

Abstract

Estimating potential output and the output gap – the difference between actual output and its potential – is important for the proper conduct of monetary policy. However, the measurement and interpretation of potential output, and hence the output gap, is fraught with uncertainty, since it is unobservable. It is therefore important that we continually expand and improve upon existing models, and innovate by testing new approaches and incorporating them into the analysis of potential output and the output gap. Within this context, this paper first provides an assessment of the extended multivariate filter (EMVF), which the Bank has used since the late 1990s to come up with a baseline measure of the output gap. It is determined that the EMVF has several limitations that need to be addressed. Consequently, a modified version of the EMVF incorporating revised conditioning information is presented. In addition, a newly developed methodology, the integrated framework (IF), provides a separate analysis of trend labour input and trend labour productivity, and in doing so accounts for more long-term structural changes in the economy. While neither of these approaches is perfect, and both have limitations, they represent improvements over the conventional method. The paper also outlines how the modified EMVF, the IF, and information from the Bank's *Business Outlook Survey* and other sources are used to come up with an estimate of the current output gap and the future growth rate of potential output.

JEL classification: E0, E31, E5, E52
Bank classification: Economic models; Inflation and prices; Labour markets; Productivity

Résumé

L'estimation de la production potentielle et de l'écart de production – soit la différence entre la production effective et potentielle – est importante pour la conduite appropriée de la politique monétaire. Cependant, la production potentielle n'étant pas observable, la mesure et l'interprétation de cette variable, et par conséquent de l'écart de production, sont entourées d'incertitude. Il est donc essentiel de continuellement développer et améliorer les modèles existants, d'innover en mettant à l'essai de nouvelles méthodes et d'incorporer celles-ci dans l'analyse de la production potentielle et de l'écart de production. C'est dans cette optique qu'est d'abord proposée une évaluation du filtre multivarié élargi (EMVF) dont la Banque se sert depuis la fin des années 1990 pour établir une mesure de référence de l'écart de production. Ce filtre comporte plusieurs lacunes qu'il faut combler. Par conséquent, une version modifiée du filtre multivarié élargi comprenant une révision de l'information exploitée pour estimer la production

potentielle est exposée. De plus, une méthode nouvelle procure une analyse distincte de la croissance tendancielle du facteur travail et de la croissance tendancielle de la productivité du travail, et prend en compte certains changements structurels à long terme de l'économie. Même si aucune de ces approches n'est parfaite, et bien qu'elles présentent toutes les deux des lacunes, elles constituent un progrès par rapport à la méthode de mesure habituelle de l'écart de production. L'étude décrit aussi de quelle manière le filtre élargi, le cadre intégré et l'information de l'enquête sur les perspectives des entreprises menée par la Banque sont utilisés avec d'autres sources pour estimer l'écart de production actuel et le taux de croissance future de la production potentielle.

Classification JEL : E0, E31, E5, E52
Classification de la Banque : Modèles économiques; Inflation et prix; Marchés du travail; Productivité

1. Introduction

Estimating potential output and the output gap is important for the conduct of monetary policy, since the output gap – the difference between actual output of the economy and its potential – is a key indicator of inflation pressures in the economy. When demand is strong, actual output can move above potential, therefore pushing against the economy's capacity to produce and putting upward pressure on inflation. Potential output can thus be defined as the level of output that can be sustained in an economy without adding to inflationary pressures. Potential is determined by structural factors such as demographic developments, education, innovation and the stock of capital. However, estimating potential output is a difficult task because it is not observable and its determinants can be difficult to measure or estimate.

Since the late 1990s, the extended multivariate filter (EMVF), developed by Butler (1996) and augmented by other Bank of Canada economists, has been the main tool used at the Bank of Canada to measure past and present potential output. The EMVF's success has been such that it has been called the "conventional measure" of the output gap in the Bank's *Monetary Policy Report*. This success reflects the fact that the EMVF has several advantages. In particular, it combines mechanical filtering with additional information, including economic relationships, in a way that was very innovative when it was first developed. While the mechanical filtering part ensures that estimates do not deviate too much from actual data, which helps mitigate the risk of failing to account for structural breaks in the economy, the economic relationships can help to fine-tune the potential output estimate. In addition, the EMVF is easy to use, since it is almost fully automatized. The current paper provides a re-examination of the EMVF and discusses one alternative and one complement: a modified version of EMVF and the integrated framework (IF).

We consider four criteria in evaluating methods used to measure potential output. Because potential output developments need to be analyzed and communicated, a first criterion is that an approach should provide economic interpretation. Consistent with the Bank's inflation-targeting mandate, a second criterion is that estimated potential output should be helpful in identifying inflationary and disinflationary pressures in the economy. The third criterion is that the assumptions and statistical relationships conditioning potential output estimates should be consistent with the data. Fourth, given that the Bank needs to estimate historical and future growth rates of potential, the tools should be helpful for both estimating and projecting potential. The approaches presented in this paper – the EMVF, the modified EMVF and the IF – meet these criteria to different extents, and in this paper we discuss the pros and cons of each of them.

More specifically, we show that the EMVF does not fully meet the four criteria because: (i) it relies on a mix of conditioning economic relationships, mechanical filtering, end-of-sample

constraints and other prior information, which greatly complicates economic interpretation; (ii) some of the conditioning information is not supported by the most recent data and as a result the mechanical filter on which the EMVF is based could give distorted output gap estimates; and (iii) it is difficult to use it to project potential.

Given these limitations, in this paper we revise some features of the EMVF. The modified version of the filter uses updated conditioning information, partly addressing point (ii). Also, the determination of trend labour productivity is revised in a way that makes it easier to interpret and less dependent than the current EMVF on assumptions about the value of the trend labour share. However, limitations remain with the modified EMVF. In particular, it is difficult to interpret (despite the changes to trend labour productivity). Still, this revised version has its advantages and it makes sense to use it in an approach aimed at managing the risk of large errors. But it is a work in progress. Much conditioning information has been removed and some will need to be added.[1]

A recently developed method, the IF, allows for richer economic analysis. For instance, it can be used to estimate and project trend hours worked by cohorts or age groups. The IF is also easier to use for projections, and indeed a preliminary version of it has been used by Bank staff to project potential in recent years. In addition, it is flexible and allows for output gaps of different sizes and degrees of persistence, and the IF can be used to analyze the implications of certain structural changes in the economy; for instance, the implications of aging for trend labour input. It is important to note that the IF is a living approach that will continue to be refined and improved through time. It is also important to note that the IF should not be thought of as a single model – rather, it is a set of tools that are brought together to shed light on the movements in potential output. Specialists' judgment is involved in using it; therefore, the estimates it provides cannot always be replicated with data alone.[2]

Unfortunately, the perfect approach does not exist and the IF has limitations. For instance, it does not directly take into account all relevant pieces of information, and, being more structural, it could potentially lead to larger errors than methods relying more on mechanical filters. Also, in its current form, the IF relies on a mechanical filter to identify trend total factor productivity growth, which complicates its use for projections, and on a crucial assumption that is required to pin down the level of potential. The latter assumption is required because the growth accounting framework, used to measure trend labour productivity, can only be applied to determine the growth rate of that trend, not the level.

[1] For instance, while we removed the Phillips curve that was used in the EMVF (because it is no longer supported by the data), we have not yet added back a different Phillips curve. We are working on doing this.

[2] The term "integrated framework" is derived from the fact that the analysis is integrated with the Bank projection of variables such as real GDP and real business investment, and it is a tool that can be used to estimate historical as well as future potential output.

The level and the growth rate of potential output remain highly uncertain. This is why the Bank relies on various models, indicators and judgment to arrive at an estimate for the profile of potential output. It is important that the Bank continually expand and improve upon existing models, and innovate by testing new models and incorporating them into the analysis of potential output and the output gap. The revision of the EMVF and the introduction of the IF represent necessary and important steps in this process, and greatly improve the assessment of potential output. In particular, the methods used need to be sufficiently different if they are to protect against the risk of having the wrong model, and the modified EMVF and the IF satisfy this requirement.

In the current approach, Bank staff use the IF and the modified EMVF to build a base-case estimation, and use the IF to project potential. They then consider other sources of information, including various labour market indicators and information coming from the Bank's *Business Outlook Survey* (BOS), to apply judgment and arrive at final estimates and projections. The Bank's Governing Council may add judgment to form its own estimate and projection. When all of this information and the new methodologies are taken into consideration, it is believed that the Bank now has more balanced and reliable estimates of potential output and the output gap. The rest of this paper is organized as follows. Section 2 provides a brief reminder about the EMVF's original structure and of some of the changes made in subsequent years. Section 3 presents and discusses the economic relationships used to condition the filter. In section 4, we present the modified version of the EMVF and the IF, and compare these methods with the conventional measure. We then conclude in section 5.

2. The EMVF

This section explains the basic structure of the EMVF proposed by Butler (1996) and discusses subsequent changes. It ends with some discussion of limitations associated with this structure.

2.1 The EMVF's basic structure[3]

The EMVF's first building block is the Hodrick-Prescott (HP) filter (Hodrick and Prescott 1997). The HP filter chooses the trend of a time series as the solution to

$$\{y_t^g\}_{t=0}^{T+1} = argmin \sum_{t=1}^{T}(y_t - y_t^g)^2 + \lambda \cdot (\Delta^2 y_{t+1}^g)^2, \qquad (1)$$

where $\Delta^2 y_{t+1}^g = \Delta\Delta y_{t+1}^g$, the second difference of the trend, and λ is the smoothness parameter. Laxton and Tetlow (1992) propose adding a term to the equation:

$$\{y_t^g\}_{t=0}^{T+1} = argmin \sum_{t=1}^{T}(y_t - y_t^g)^2 + \lambda_g \cdot \left(\Delta^2 y_{t+1}^g\right)^2 + \sum_{i=1}^{n} \lambda_{\varepsilon_i} \cdot \varepsilon_{it}^2, \quad (2)$$

[3] Our presentation in this section is largely based on Butler (1996) and St-Amant and van Norden (1997).

where $\varepsilon_t = z_t - f\left(y_t^g, x_t\right)$, with z_t some economic variable of interest, and $f(.)$ modelling z_t as a function of both some explanatory variables x_t and the unobserved trend y_t^g. There can be i economic relations. Since equation (2) includes ε_{it}^2, the trend is chosen to simultaneously minimize deviations of output from trend, minimize changes in the trend's growth rate, and maximize the ability of the trend to fit some economic relationships. λ_g is the smoothness parameter and $\lambda_{\varepsilon_i}$are the relative weights put on these objectives.

Butler (1996) introduces two more terms: one limiting deviations of the growth rate of the trend from some assumed steady-state value at the end of the sample, and the second limiting the extent to which the trend could change from the previous period.

But the main distinguishing feature of Butler's EMVF is that, rather than filtering output directly, as was done with previous multivariate filters, it filters components of output identified on the basis of the Cobb-Douglas production function:

$$Y = Q \cdot N^{\alpha} \cdot K^{1-\alpha}, \quad (3)$$

where Q is total factor productivity, N is labour, K is the capital stock, and α is the labour-output elasticity (or labour's share of income). The marginal product of labour can be derived from this production function and then output can be written as the marginal product of labour (μ) times labour input (N) divided by the labour share (α):$Y = \mu N/\alpha$.

Labour input can be expressed as the product of population (pop), the employment rate (1-rate of unemployment (u)), and the participation rate ($part$):

$$n = pop(1-u)part.$$

Therefore, real output can be written as

$$Y = \mu n/\alpha = \mu(pop(1-u)part)/\alpha.$$

Potential is defined by the same equation, but using the filtered instead of actual series:

$$\hat{Y} = \hat{\mu}\hat{n}/\hat{\alpha} = \hat{\mu}(pop(1-\hat{u})\hat{part})/\hat{\alpha}.$$

An important assumption is that the labour share of output (or the labour/capital elasticity) is always equal to the trend labour share of output,[4] which simplifies the calculation of potential to

$$\hat{Y} = Y\frac{\hat{\mu}\hat{n}}{\mu n}.$$

[4] The trend labour share of output is estimated with a HP filter augmented with an end-of-sample restriction assuming a gradual return to a steady-state value (more details are provided in section 2.2.2).

Because the trend in average hours worked (AHW) can vary even when that of employment does not (and, indeed, a substantial change in AHW had taken place from the early 1970s to the early 1990s), and because such variation would affect trend output, Bank staff replaced employment as a measure of labour input in the EMVF with the product of employment and AHW (h). As a result, labour input was then measured as total hours worked rather than employment in the production function.[5]

2.2 Issues with the EMVF's basic structure

This section discusses some issues related to the HP filter, and the production function approach used by Butler (1996).

2.2.1 HP filter

The HP filter was chosen because it is easy to implement and because it is frequently used by macroeconomists. In addition, it has the advantage that it tends to revert back to the observed data. However, there are problems with this filter. First, there is an end-of-sample limitation resulting from the fact that the HP filter is a two-sided moving average (Baxter and King 1995). In addition, the HP filter can generate distorted business cycle measures when it is applied to persistent time series, such as the level of real GDP (Guay and St-Amant 2005). A third limitation is the fact that HP filtering the data does not provide an explanation as to why the trend is changing. This is perhaps the most important limitation for policy-makers, since they have to understand and explain changes in trends that affect their decisions.

2.2.2 The Assumed Decomposition of Output

As shown in section 2.1, the decomposition of output in the EMVF relies on a Cobb-Douglas production function. One feature of this function is that the labour share of income is constant. However, as shown in Chart 1, the actual share of labour in output decreased almost continuously from the peak reached in the 1990–91 recession to 2005, and has stabilized afterwards to a level below the previous historical average. Various factors have been identified as playing a role in this decline: globalization, technological progress and weakening bargaining power of unions. Using an error-correction model, Morel (2006) finds that increased labour productivity, decreasing union density and increasing openness to trade contributed to the decline observed in the labour share in Canada.

[5] Note that the EMVF does not use information about the capital stock. This can be an advantage given measurement issues.

In the EMVF, the trend labour share is measured using a very *stiff* HP trend (with a smoothness parameter λ of 10,000)[6] and with a steady-state restriction for the past 15 quarters. Instead of being constant, labour's share is thus allowed to change very gradually over history. The steady-state restriction implies that trend labour share is gradually moving back to its steady state at the end of sample. Assuming that labour's share of income is roughly constant, as assumed in the Cobb-Douglas production function, the steady state is set to an historical average. However, it has been well above its value over the past decade. Supported by factors such as globalization and technological developments, a lower value for the steady state provides a trend labour share that seems more reasonable (see Chart 1).[7] We calculate that such a change to the estimate of trend labour share increases excess supply at the end of sample by about 1 percentage point.

Chart 1: Labour share

Quarterly data

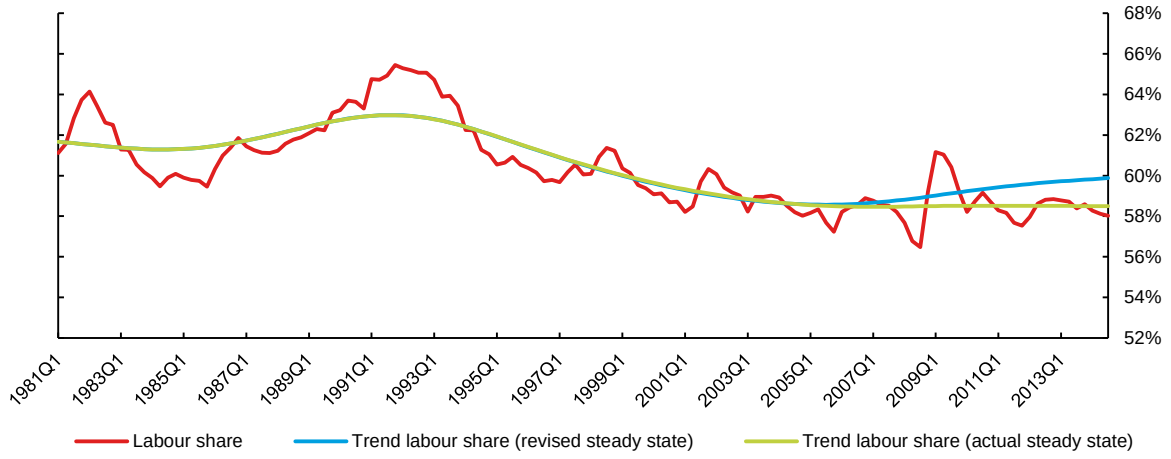

Sources: Statistics Canada and Bank of Canada calculations

Last observation: 2014Q3

Because labour's share of output is set to be always equal to its trend, revising down its steady-state assumption implies a lower actual marginal product of labour. It also affects the estimate of trend marginal productivity, but to a lesser extent, because the filter is smoothing the impact of this change. Therefore, the downward revision to trend labour share has a negative effect on the productivity gap, and thus on the output gap.

[6] Labour share is the only component for which a value of the smoothness parameter is different from the standard 1,600.

[7] The average labour's share of income over the past decade is used for the green line in Chart 1.

3. The Conditioning Information in the EMVF

As discussed in section 2.1, various economic relationships and other elements of information are used to condition the estimate of potential output within the EMVF. This section presents the conditioning economic relationships, discusses problems that some of them pose, and describes their influence, as well as that of other elements of information, on potential output estimates.

3.1 Economic relationships used in the EMVF

The economic relationships introduced by Butler (1996) that are still in use include a Phillips curve, a modified Okun's law and a cointegration relation between the marginal product of labour and real wages.

The **Phillips curve** links inflation with the output gap and expected inflation. This Phillips curve is non-linear, since excess demand gaps have been found to have more effects on inflation than excess supply gaps (for more details see Laxton, Rose and Tetlow 1993). When re-estimating this relation with an updated sample, the results are that (i) the non-linearity is no longer supported by the data, and (ii) for a sample starting with the adoption of the inflation target, the coefficient on the output gap is also insignificant.

The **modified Okun's law** links the current change in the unemployment rate gap with the lagged marginal product of labour gap. This is intended to capture the quantity adjustment process that firms undertake when their marginal product deviates from its short-run equilibrium value. This relationship is defined as follows:

$$u_t - \hat{u}_t = \theta(\mu_{t-1} - \hat{\mu}_{t-1}) + 0.9(u_{t-i} - \hat{u}_{t-i}),$$

where u and $\hat{u}$ are the observed and filtered unemployment rates, respectively, and θ is the "Okun coefficient." For the whole sample, the estimated parameters are not very different from what is used in the current EMVF. However, when the sample period is split to account for the inflation-target regime, there is evidence of a structural break (the coefficient on the productivity gap is not significant anymore).

The **real producer wage** is used as external information for the trend of the marginal product of labour. Neoclassical theory suggests that the equilibrium marginal product of labour and real producer wage should be cointegrated.

Additional economic information is introduced subsequently: measures of trend average hours worked and trend unemployment, and an estimate of the trend participation rate. The HP trends of **average hours worked** and **unemployment** are assumed to be cointegrated with trends identified on the basis of structural vector autoregressions (SVARs) based on long-run restrictions à la Blanchard and Quah (1989). The average hours worked SVAR model has three variables (growth rates of average hours worked, U.S. GDP and CPIX), and the trend unemployment SVAR model uses six variables (employment insurance generosity, degree of unionization, unemployment rate, inflation rate, real overnight rate and real long-term interest rate).

Updating the SVAR used to estimate trend average hours worked provides counterintuitive results (the average-hours gap persists below zero since the early 1990s, while we would expect a zero mean). This is probably because inflation is affected by a structural break in 1992, as a consequence of the adoption of the inflation targets. Moreover, some variables that might be important to determine Canadian hours worked, such as demographic variables, are not included.

Updating the measure of trend unemployment using the original SVAR model was not possible because of data limitations.[8] A more recent estimate of trend unemployment obtained from an updated version of the model developed by Côté and Hostland (1996) is used in section 4.2.

The **trend participation rate** is estimated and projected using a demographic approach based on 24 age and gender groups of the working-age population.[9]

3.2 Impact of the conditioning information

If one removes the effects of the conditioning information, the EMVF simply consists of applying the Hodrick-Prescott filter on each component and then summing these to obtain potential output. The impact of the conditioning information on the estimated EMVF output gap can therefore be summarized by comparing it with an HP-filter-based gap (see Chart 2).

[8] The unionization rate series that was provided by Human Resources and Skills Development Canada (HRSDC) was terminated in 2004. Data on unionization are also available at Statistics Canada, but they start only in 1997. These two series look very different over the overlapping period, which has prevented us from splicing them together.
[9] This model was estimated years ago, but its projections (used in estimating trend participation) are still of merit. Nevertheless, this model will need to be updated. We leave this to future work.

Chart 2: HP filter and EMVF output gaps

Quarterly data

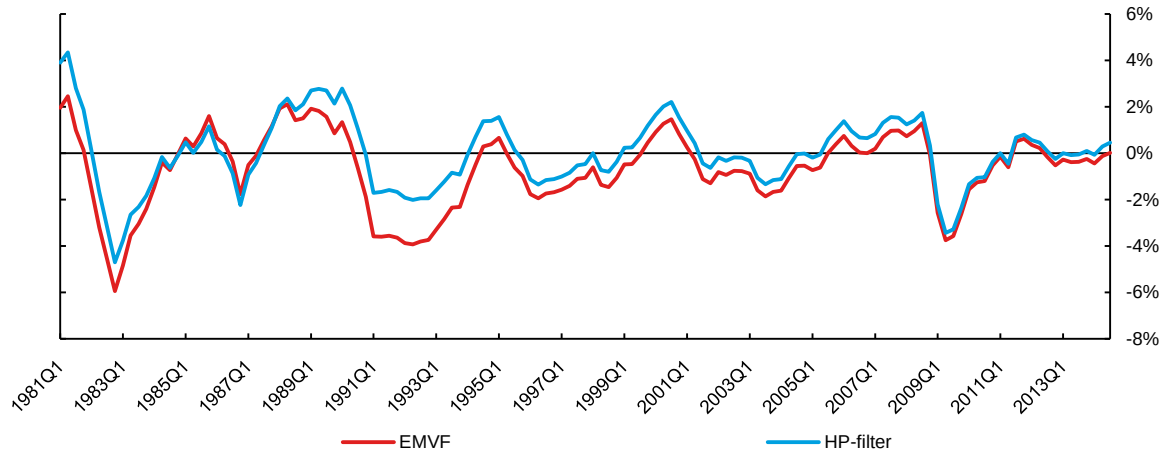

Source: Bank of Canada calculations

Last observation: 2014Q3

The EMVF output gap sometimes differs substantially from the HP output gap. This is the case, for instance, at the time of the early 1990s recession.[10] At that time, the EMVF output gap is substantially more negative, which seems consistent with the significant disinflation that then took place. In fact, the HP filter is simply a moving average applied to the data. By construction the resulting output gap will be zero, or approximately zero, on average. Conversely, on average, the EMVF output gap is negative over the historic sample examined (about -0.7 per cent).

One reason why the EMVF output gap is on average negative is the presence of a "bias corrector." The corrector has a positive value in the hours worked and unemployment components (pushing the trend of these variables upward), but a negative value in the marginal product of labour component. Overall, the net impact on potential output is positive, and thus the output gap is more negative than if no adjustment was applied. This is shown in Chart 3. These bias correctors reflect judgment about the level of the gap at the time the EMVF was made operational.

[10] Of course, this is the ex post output gap. The real-time output gap could be, and has been, different.

Chart 3: Output gap - sensitivity to bias correction

Quarterly data

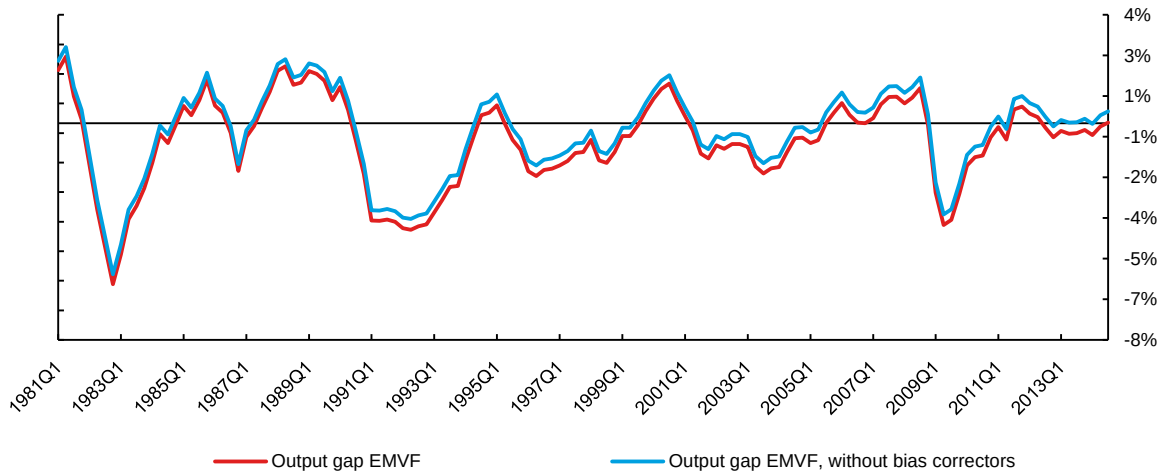

Source: Bank of Canada calculations

Last observation: 2014Q3

To assess the effect of the conditioning economic information, we have performed a sensitivity analysis that shows that very few economic relations and end-of-sample restrictions have a significant impact on the overall assessment of the output gap. As discussed in section 2.2.2, the steady-state value assumed for the labour share of income has a significant impact on the output gap. Similarly, imposing a weight of zero on this end-of-sample constraint, when calculating trend labour share, implies larger excess supply. Another end-of-sample restriction that has some impact is the weight applied on the deviations of the marginal product of labour from its steady state. In times where productivity growth is persistently below the assumed steady state, the restriction poses upward pressures on trend productivity at the end of sample. Recently, however, productivity growth has been approaching the steady state imposed in the EMVF, leaving little material impact for this end-of-sample restriction (see Charts 4 and 5, red lines). The only other conditioning information that seems to have an effect on the estimate of the output gap is the cointegration relationship linking the marginal product of labour with the real producer wage. Removing this relationship from the EMVF increases excess supply by more than one percentage point at the end of the sample (see Charts 4 and 5, green lines). The effect of this conditioning information is more important at the end of the sample, even though it is taken into account over the whole sample, because the marginal product of labour is overestimated relative to wages since it is calculated using trend labour share.

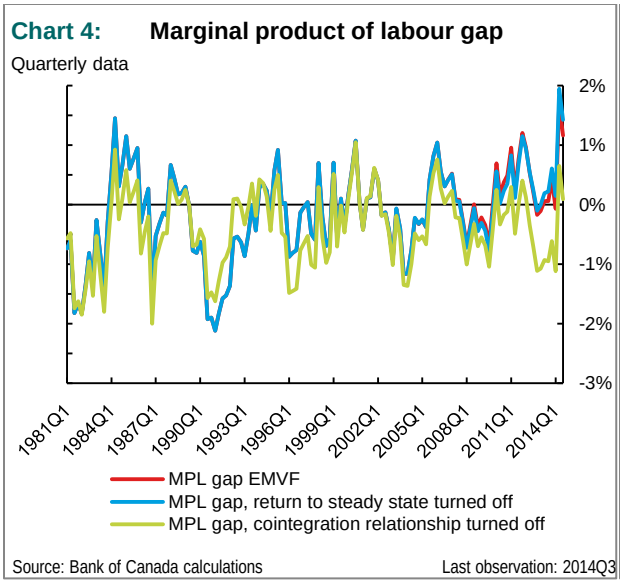

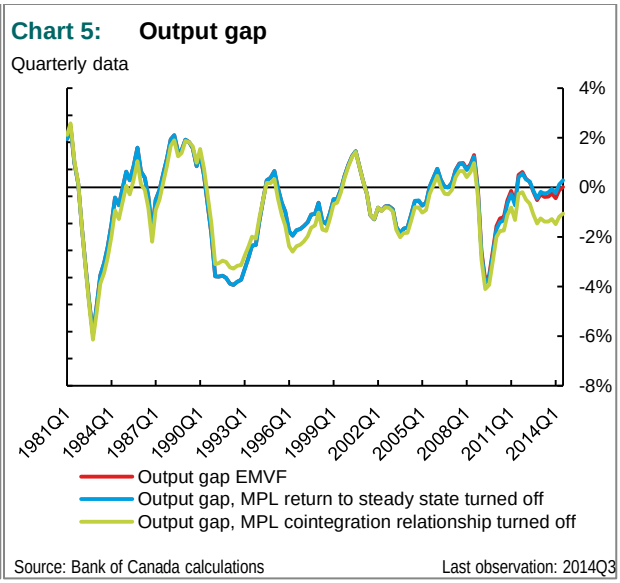

4. The EMVF versus Alternative Methods

This section compares the EMVF with two alternatives: a modified version of the EMVF and a different approach called the integrated framework. In the introduction, four criteria were proposed for assessing methods designed to assess potential output. Section 3 showed that the EMVF does not meet some of these criteria, since some of the economic relationships conditioning the filter are not supported by recent data. In addition, the EMVF is very difficult to interpret. Sections 4.1 and 4.2 present two alternatives and assess them against our criteria. Section 4.3 compares the approaches. Section 4.4 provides a brief discussion of other approaches, including some that may be worth exploring further in future work.

4.1 The modified EMVF

The modified version of the EMVF takes into account the conclusions drawn from the re-examination presented in the previous sections. More specifically, based on the analysis presented in section 3, the information provided by the modified Okun's law, the VAR-based trend average hours worked and the trend unemployment based on a structural VAR is excluded. In the latter case, it is replaced by a more recent estimate of the trend unemployment rate based on an updated version of the model developed by Côté and Hostland (1996).

In light of the discussion in section 2.2.2, the modified EMVF also improves the conventional measure by changing its structure on the marginal product of labour side. It modifies the decomposition of potential output and focuses on trend labour productivity (output per hour). This is consistent with Bank analysis of potential presented in *Monetary Policy Reports*. This

change enables us to avoid making assumptions about the production function and to move away from the concept of the marginal product of labour, which is defined using the labour share of income. As shown in section 2.2.2, this variable was found to have a significant impact at the end of the sample, producing a larger excess of supply when its steady state was revised down. But the explanation of this effect is not straightforward in terms of economic interpretation. We also revised down the steady state for labour productivity growth to be in line with the most recent assessment of the Bank,[11] and eliminated the bias corrector discussed in section 3.2, since it does not seem relevant anymore.

A pseudo real-time analysis presented in Appendix A suggests that the estimate of the output gap could be subject to slightly larger revisions with the use of the modified EMVF. However, these revisions appear less biased, on average, than the ones affecting the current EMVF.

4.2 The integrated framework

This section presents an approach called the integrated framework (IF). The IF is an extension of the approach that has been used by the Bank in recent years to project potential output growth. It produces annual estimates of potential output and quarterly estimates are obtained by interpolation.

The IF defines potential output as the product of trend labour input and trend labour productivity. We briefly discuss the estimation of these two trends.

Trend labour input is estimated using an updated version of the cohort-based model developed by Barnett (2007). In this model, trend labour input growth is decomposed as the sum of growth of the working-age population, trend employment rate and trend average weekly hours worked per employee. The growth rates in the working-age population are based on actual Labour Force Survey data (Statistics Canada) and on Statistics Canada's medium-growth population projections (Bank staff typically run sensitivity analyses using alternative population scenarios). The trend employment rate is estimated using a reduced-form model factoring in cohort (capturing, in particular, the increasing labour market participation of female cohorts) and age-group fixed effects, as well as the effects of other explanatory variables such as wealth over disposable income, an employment insurance disincentive index (Sargent 1995), a measure of the job offer rate and a real after-tax interest rate. The reduced-form model used to estimate trend average hours is similar, but includes a variable to capture the effects of school enrolment for youth and does not include cohort fixed effects. Small but important improvements have been made to the model presented in Barnett (2007), and work is ongoing to incorporate other important factors affecting trend labour input, such as changes in the level of education in Canada over history. The trend employment rate and trend average hours

[11] See Box 3 in the October 2014 *Monetary Policy Report*.

worked are first projected for different age groups and then aggregated using the projected age composition of the working-age population (provided by Statistics Canada). Therefore, this approach allows for explicitly taking into account demographic changes in the Canadian population (for example, population aging).

Chart 6 shows estimated trend labour input over recent history (1992–2014) as well as actual data for average hours worked and the gap between the two. IF estimations suggest a slowdown of trend labour input since around 2000. This reflects a combination of two main factors: a decline in the growth rates of the working-age population, owing to relatively low fertility rates over earlier decades, and a falling trend employment rate and trend average hours worked, largely due to an increasing proportion of older workers in the Canadian population.

Chart 6: Total hours worked (labour input), trend labour input and labour input gap

Source: Bank of Canada calculations

Last observation: 2014Q3

An advantage of this approach to measuring trend labour input is that it minimizes the usage of mechanical filters. It instead relies on a few structural factors and thus allows for some economic interpretation, partly addressing our first criterion. The story that could be extracted remains, however, somewhat limited, because this is not a fully structural model.

An additional advantage of the trend labour input model is that it allows for more detailed labour market analysis. For instance, it allows for estimating the effects of population aging and changing population structure, as well as projection of trend labour input by age group. Chart 7

shows the decomposition of the labour input gap by age group: youth (15–24), prime-age workers (25–54) and older workers (55+). Prime-age and youth workers were the main contributors to the labour input gap in the wake of the recent recession. Even though we observe a steadily closing labour input gap among prime-age workers, the labour input gap among youth shows a significantly slower path toward the recovery. The latter is in line with other economic indicators, such as elevated involuntary part-time work among youth.

Chart 7: Labour input gap decomposition by age groups, in %

Quarterly data

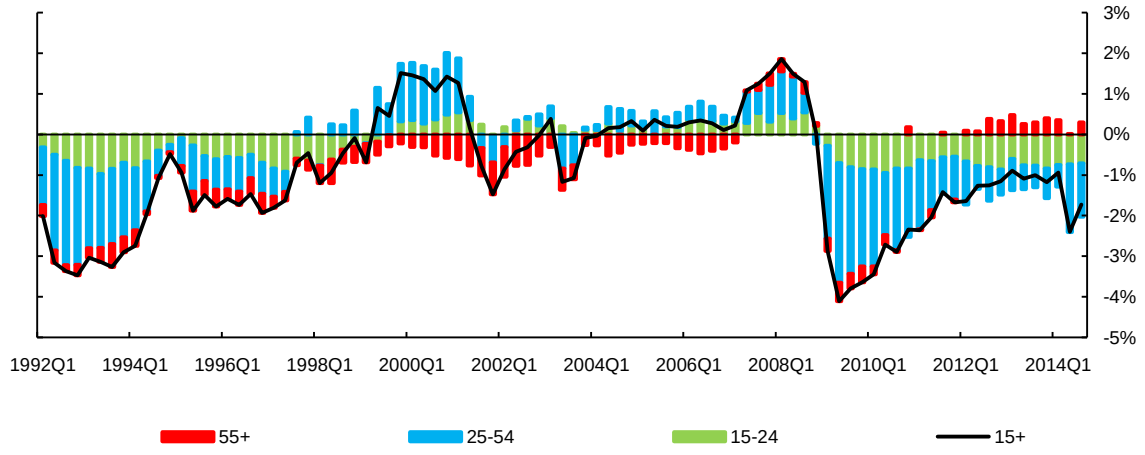

Source: Bank of Canada calculations

Last observation: 2014Q3

Trend labour productivity is measured using the growth accounting framework. The growth rate of trend labour productivity is decomposed as the sum of the growth rates of capital deepening (capital stock per hour worked) and total factor productivity (TFP). The trends of these two components are identified separately using a mix of economic data, a simple approach linking investment with capital deepening, mechanical filters, and judgment. This approach to identifying trend labour productivity is new, and for this reason, a more detailed description of the steps involved is provided in Appendix B.

Moving to this framework includes a number of advantages relative to the current EMVF. First, by using the growth accounting framework, which is a generalization of a broad set of constant-returns-to-scale production functions, we move away from the strong assumptions underlying the Cobb-Douglas production function. Second, we also reduce the usage of mechanical filters, and whenever it is used, we opt for the band-pass filter detailed by Christiano and Fitzgerald (2003), which has been found to perform better than the HP filter to identify trends at the end

of samples. Finally, this method incorporates more economic structure, which enables us to tell a story explaining changes in the trend of labour productivity. On the one hand, this more structural approach can generate large and persistent output gaps, which in some circumstances is helpful in explaining persistent weak inflation; on the other hand, if these large gaps are inaccurately measured by the IF, this method could lead to large errors. Therefore, one must exercise caution in interpreting the output gap estimates generated by the IF.

While the labour productivity model is useful for economic intuition, there are two important limitations worth noting. First, there is a significant amount of judgment used to derive estimates of trend labour productivity (especially in the identification of trend total factor productivity growth). Second, and this is related to the first limitation, the labour productivity model generates growth rates of the trend, and therefore the level of trend labour productivity is determined by assuming a point in the past where trend and actual labour productivity were equal, which relies on judgment.[12] Furthermore, the IF does not incorporate all information that could be helpful for estimating output gaps (for instance, the Bank's BOS). It therefore needs to be supplemented with other information.

4.3 Comparing the approaches

Chart 8 compares the estimated output gap for these three alternatives. As shown in this chart, the output gaps are highly correlated and broadly consistent across history. That is, excess supply is observed during recessions in all cases and most episodes of higher inflation coincide with periods of excess demand.[13] Some discrepancies between these measures might nevertheless appear in specific periods. For instance, during the 1990–91 recession, excess supply is more significant when measured with the conventional measure of the output gap potentially because of the Phillips curve that is conditioning this estimate. The disinflation that occurred with the introduction of the inflation target has been compatible with a larger excess supply incorporated through the Phillips curve. Another example is the most recent period, where the IF estimate of the output gap is lower (the output gap is more negative and more persistent) than the other measures. By construction, HP-filter-based methods such as the EMVF cannot generate such large and persistent gaps (the mechanical filter comes to interpret persistent weakness in actual data as a lower trend). The IF, on the other hand, can generate large and persistent output gaps following recessions. This is partly because it takes into

[12] We could make an assumption about the nature of the production function (such as assuming it is Cobb-Douglas) that would allow estimating the level of potential output directly. However, this would require further strong assumptions that would obfuscate the analysis of potential output. More details on this are provided in Appendix B.

[13] For a robust analysis, an assessment with real-time data should be done, but limited data prevent us from doing this exercise. This might be considered for future analysis.

account demographic variables that are little affected by business cycle fluctuations and are more likely to produce persistent deviations from trend.

Chart 8: Output gap

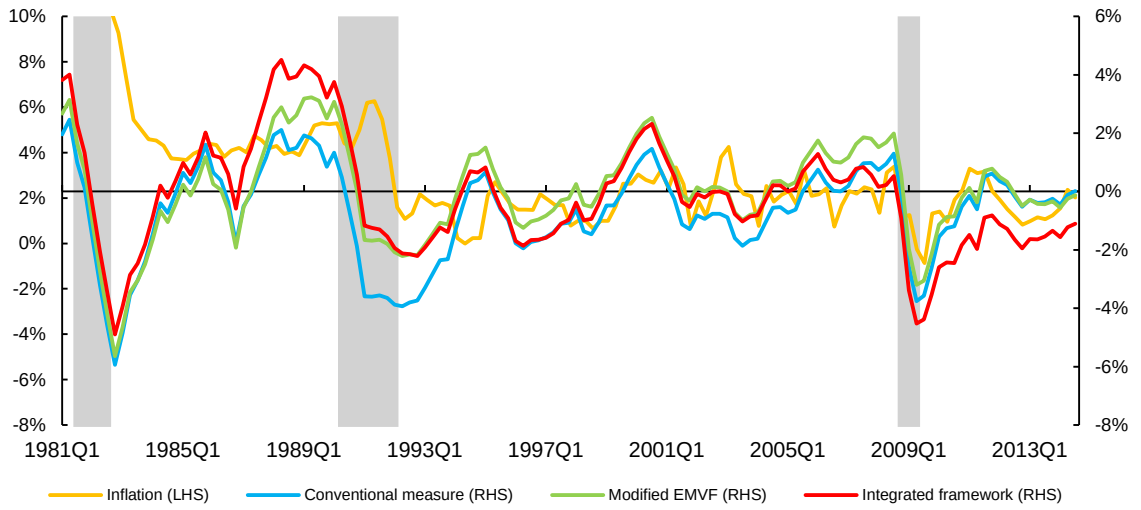

Sources: Statistics Canada and Bank of Canada calculations

Last observation: 2014Q3

4.4 Other methods

Other methods could be considered as alternatives to the EMVF. St-Amant and van Norden (1997) and Cayen and van Norden (2005) consider a few. It is, however, important to note that most of these methods provide very limited interpretability.

Borio et al. (2013) recently proposed measures of potential output that incorporate information about the financial cycle. We replicated their approach with Canadian data and found that estimates based on this approach are very sensitive to the choice of financial variables (OECD (2013) reaches the same conclusion).[14]

[14] A similar approach may nevertheless be worth pursuing in future work.

5. Conclusions

Since the late 1990s, the extended multivariate filter (EMVF) developed by Butler (1996) and augmented later by Bank staff has been the main tool used at the Bank of Canada to measure past and present potential output. This paper provides a re-examination and highlights some strengths and limitations. One limitation is that part of the information used to condition the EMVF needs to be reconsidered. We address this by updating part of the conditioning economic relationships and long-term assumptions, and by eliminating others that we find are no longer supported by the data. We also change the productivity side of the filter by refocusing it on labour productivity instead of marginal productivity of labour. This gives us a modified version of the EMVF. And while this modified EMVF is a useful framework for assessing potential output – an essential part of the potential output tool box – it has limitations. In particular, it remains difficult to interpret potential output developments.

This is why the Bank is now incorporating more structural approaches into the analysis of potential output. In the short to medium term, this means incorporating a set of tools combining cohort-type models to analyze labour input with a measure of trend labour productivity obtained by analyzing the contributions of capital deepening and total factor productivity, using the growth accounting framework. This approach, which we call the integrated framework, provides an estimate of potential output that allows for more economic interpretation while minimizing the use of mechanical filters. It has already been used at the Bank of Canada in the 2013 and 2014 October analysis of potential output. Further research is, however, required to refine the IF. In particular, the labour productivity side of that approach is only formalized in linking investment with the stock of capital. More research is required to link labour productivity with other factors that may affect it (for example, technological progress, economies of scale, firm creation and labour skills).

Although the IF allows for richer economic interpretation and is useful to project potential, it, too, has limitations. There is considerable uncertainty surrounding potential output estimations, an unobservable variable, and it would not be prudent to base potential output estimations and projections on a single method. This is why the EMVF is also used in the analysis of potential output. It is important that the Bank continually expand and improve upon existing models, and innovate by testing new models and incorporating them into the analysis of potential output and the output gap. The revision of the EMVF and the introduction of the IF represent necessary and important steps in this process, and greatly improve the assessment of potential output. In particular, the methods used need to be sufficiently different if they are to protect against the risk of having the wrong model, and the EMVF and the IF satisfy this requirement.

For the final assessment of potential output and the output gap, additional information is considered to further refine the estimates provided by the EMVF and IF. For instance, the assessment of the output gap estimate is informed by information based on the Bank's *Business Outlook Survey*, movements in wages and CPI inflation, labour market indicators, and movements in other indicators of capacity pressures. When all of this information and the new methodologies are taken into consideration, it is believed that the Bank now has more balanced and reliable estimates of potential output and the output gap.

References

Baldwin, J. R., W. Gu and Y. Beiling. 2007. “User Guide for Statistics Canada's Annual Multifactor Productivity Program.” The Canadian Productivity Review 2007, Statistics Canada, Economic Analysis Division.

Bank of Canada. 2014. *Monetary Policy Report*. October.

Barnett, R. 2007. “Trend Labour Supply in Canada: Implications of Demographic Shifts and the Increasing Labour Force Attachment of Women.” *Bank of Canada Review* (Summer): 5–18.

Baxter, M. and R. G. King. 1995. “Measuring Business Cycles: Approximate Band-Pass Filters for Economic Time Series.” Working Paper No. 5022. National Bureau of Economic Research.

Blanchard, O. and D. Quah. 1989. “The Dynamic Effects of Aggregate Demand and Supply Disturbances.” *American Economic Review* 79(4): 655–73.

Borio, C., F. P. Disyatat and M. Juselius. 2013. “Rethinking Potential Output: Embedding Information about the Financial Cycle.” BIS Working Paper No. 404, Bank for International Settlements.

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

Cayen, J.-P. and S. van Norden. 2005. “The Reliability of Canadian Output-gap Estimates.” *The North American Journal of Economics and Finance*, Elsevier, 16(3), 373–93, December.

Christiano, L. J. and T. J. Fitzgerald. 2003. “The Band Pass Filter.” *International Economic Review*, 44, 435–65, May.

Côté, D. and D. Hostland. 1996. “An Econometric Examination of the Trend Unemployment Rate in Canada.” Bank of Canada Working Paper No. 96-7.

Guay, A. and P. St-Amant. 2005. “Do the Hodrick-Prescott and Baxter-King Filters Provide a Good Approximation of Business Cycles?” *Annales d'Économie et de Statistique*, ENSAE, 77, 133–55.

Hodrick, R. and E. C. Prescott. 1997. “Postwar U.S. Business Cycles: An Empirical Investigation.” *Journal of Money, Credit, and Banking*, 29 (1), 1–16.

Laxton, D. and R. Tetlow. 1992. “A Simple Multivariate Filter for the Measurement of Potential Output.” Technical Report No. 59. Ottawa: Bank of Canada.

Laxton, D., D. Rose and B. Tetlow. 1993. “Is the Canadian Phillips Curve Non-linear?” Bank of Canada Working Paper No. 93-7.

Morel, L. 2006. “A Sectoral Analysis of Labour’s Share of Income in Canada.” Internal Bank of Canada note (presented at the 40th Annual Meeting of the CEA).

OECD. 2013. *Uncertainty about Potential Output and the Implications for Underlying Fiscal Balances.* Working Party No. 1 on Macroeconomic and Structural Policy Analysis.

Sargent, T. C. 1995. “An Index of Unemployment Insurance Disincentives.” Department of Finance Working Paper No. 95-10.

St-Amant, P. and S. van Norden. 1997. “Measurement of the Output Gap: A Discussion of Recent Research at the Bank of Canada.” Technical Report No. 79. Ottawa: Bank of Canada.

Appendix A: A Pseudo Real-Time Analysis for Assessing Potential Revisions to the Output Gap

This appendix presents the results of a pseudo real-time analysis aimed at assessing the extent of potential revisions to the estimate of the output gap that would originate from the chosen approach (the EMVF versus the modified version). The exercise always uses the latest vintage of data; therefore, these potential revisions cannot be attributed to data revisions. Chart A.1 shows the estimates of the output gap using the EMVF (left panel) and the modified EMVF (right panel) for samples adding one observation at a time (starting with a sample ending in 1990Q1).

Chart A.1: Output – All "pseudo" vintages since 1990

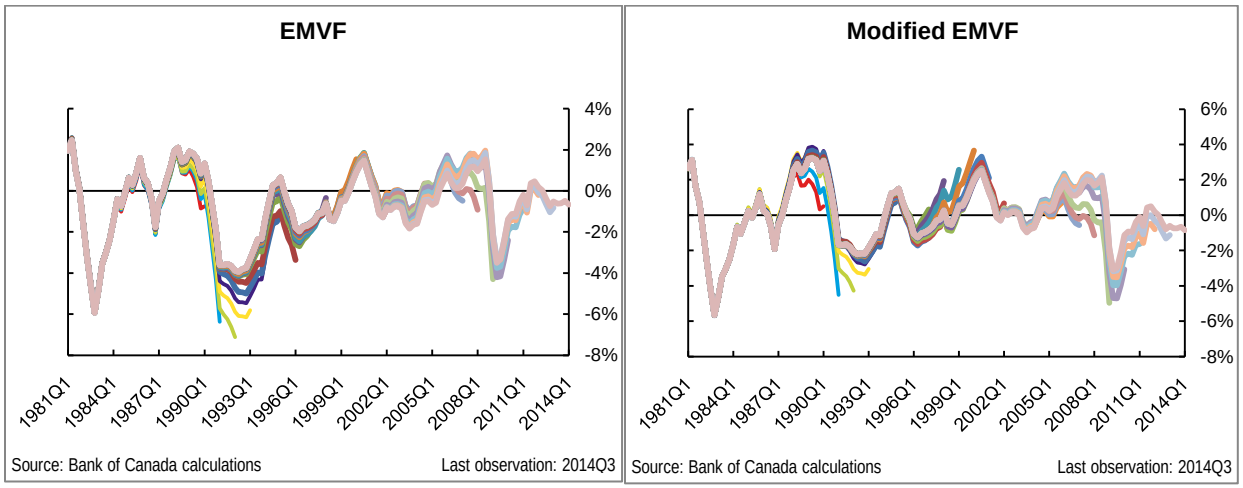

- The mean revision to the EMVF output gap since 1990 would have been 0.7 percentage points (pp), with a standard deviation of 1.0 if the raw data had never been revised. The corresponding mean absolute revision is 0.9 pp.[15]
- The mean revision to the modified EMVF output gap is 0.4 pp, with a standard deviation of 1.2. The mean absolute revision is 1.1 pp.

Chart A.2 reports the revisions to the output gap with the contributions from the labour input gap and the productivity gap for the EMVF in the left panel and the modified EMVF in the right panel. Revisions resulting from both approaches are shown in Chart A.3.

[15] The latest vintage used for those calculations is 2011Q1, to avoid including the end of sample, which is affected by the limitations of the HP filter and thus might still be subject to significant revisions.

Chart A.2: Revisions from labour input and productivity gaps

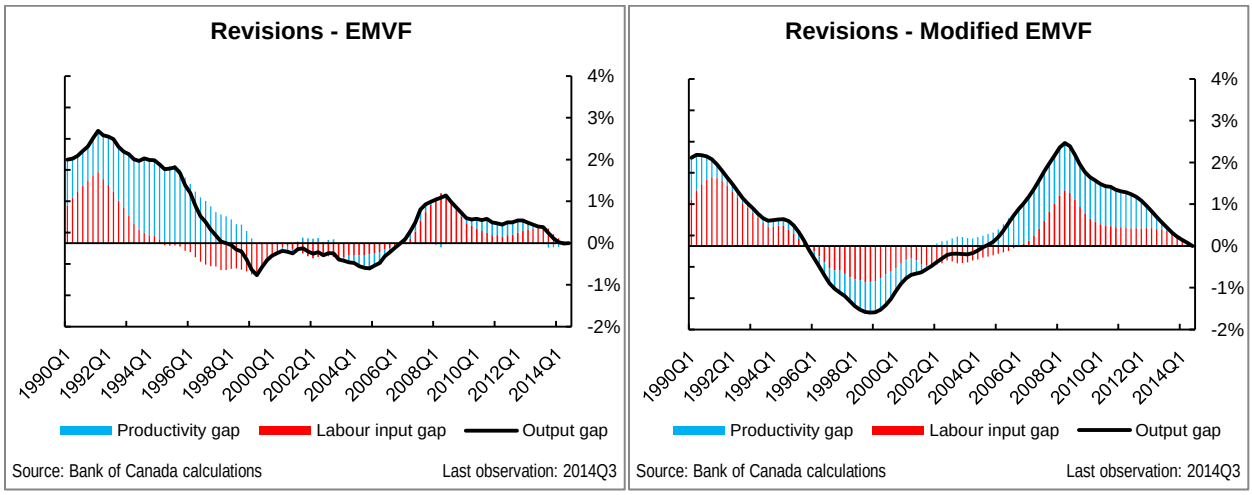

Chart A.3: Revisions – Output gaps

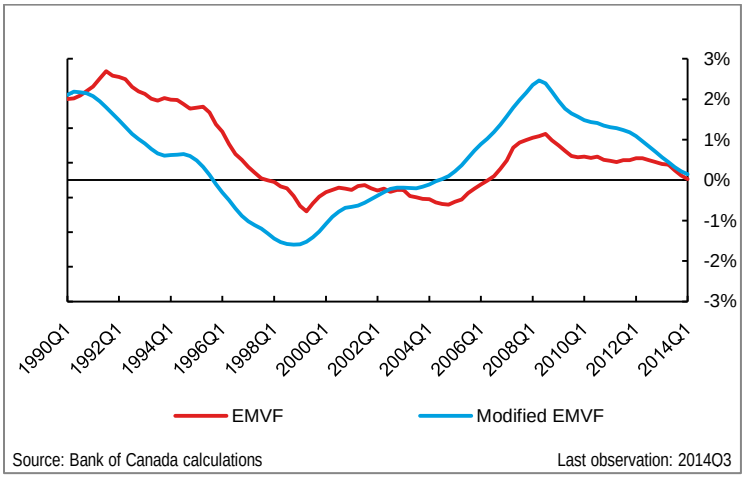

The results of this pseudo real-time analysis suggest that, on average, the modified EMVF output gap would be subject to slightly larger revisions in absolute terms than the current EMVF. This might be explained by the removal of a number of economic relationships that were used to condition the EMVF. As a consequence, the relative weights on the mechanical filter increase. However, the revisions to the EMVF output gap seem more biased on the upside than for the modified version, and Chart A.2 shows that this is mainly due to the productivity gap, which has almost never been revised downward over the sample examined. Yet, looking at the most recent decade, revisions on the productivity block are significantly more important with

the modified EMVF than the current version. One potential explanation is the link between productivity and wages that has been removed. Since real wages have persistently been below the marginal product of labour over the recent years, it has put downward pressures on trend productivity. There are also other factors that might contribute to this difference, since we have completely revisited the productivity block in the modified EMVF, using labour productivity instead of marginal productivity of labour.

Appendix B: Updating the Measurement of Trend Labour Productivity Growth as Part of the Integrated Framework

This appendix outlines the steps taken to generate estimates of trend labour productivity in the integrated framework.

B.1 Data set and framework for the analysis of the growth of trend labour productivity

A data set combining data on real GDP and the capital stock from the National Accounts, with data on hours worked from the Labour Force Survey (which is consistent with the data used in the analysis of trend labour input and in the Bank of Canada staff projection) is used to create measures for total economy labour productivity and capital deepening (capital stock per hour worked). From this, a measure of total factor productivity (TFP) growth can be derived using the following growth accounting framework:[16]

$$\Delta\left(\frac{Y}{H}\right) = \alpha_k \cdot \Delta\left(\frac{K}{H}\right) + \Delta(TFP),$$

where Δ is the annual growth rate of the variable in parentheses, *Y* is output, *H* is total hours worked, *K* is the capital stock, and α_k is the capital share in output (which is allowed to vary over time).[17,18] It is important to note that the capital stock is for the business sector and is consistent with the investment series used in the staff projection.[19] The downside is that TFP growth will include changes in the public sector capital stock; the advantage is that we have a capital stock that is consistent with the staff forecast for investment, and therefore a capital stock series can be created through the end of the projection period.[20] This framework is used to establish long-term trends in capital deepening and TFP growth, and ultimately to estimate past, current and future levels of trend labour productivity.

B.2 The identification of trend labour productivity and the projection of the growth rate

The current and historical levels of trend labour productivity growth, as well as the projection for the growth rate of trend labour productivity growth, are identified in several steps using the

[16] See Baldwin, Gu and Yan (2007) for a description of the derivation of the growth accounting framework.
[17] α_k is assumed to be a two-year moving average of the actual annual share of capital in output.
[18] This framework differs from the growth accounting framework used in Statistics Canada's multifactor productivity (MFP) database, which includes measures of capital services and labour services in the calculation of TFP (or MFP) growth. We use the stock of capital and unweighted hours, to be consistent with the analysis of trend labour input and the staff projection.
[19] The data on the capital stock are collected from CANSIM Table 031-0003. The data include the end-of-year real net capital stock and annual real investment.
[20] In order to relate investment to the capital stock, the permanent inventory method ($K_t = (1-\delta)K_{t-1} + I_t$) is used to build the capital stock from the investment forecast through the end of the projection period (the staff forecast growth of investment is used to generate the level of investment). To get an estimate of the depreciation rate (δ), we back out the implied depreciation rate from the historical series and then use an average of the recent depreciation rate to determine the implied forecast for the capital stock.

growth accounting framework outlined above. The first step is to identify the growth rate of the trend level of capital stock per hour worked. The second step is to identify the growth rate of the trend level of TFP. In the final step, we combine the capital deepening and TFP growth estimates to get an estimate of the level of trend labour productivity (using judgment to identify the level).

Step 1: The growth rate of trend capital stock per hour worked

It is assumed that the trend level of capital stock per worker can be reasonably approximated by dividing the trend capital stock by trend hours worked. For trend capital stock, if it is assumed that there is no capital stock gap – that is, the capital stock represents all usable capital in the economy and is the maximum amount of capital that can be used for production – then the actual, observed capital stock equals the trend level of capital.[21] Therefore, using the methodology described above, the investment forecast can be linked to the capital stock to obtain the level of trend capital stock from 1980 through to the end of the projection period. This can then be combined with trend labour input estimates to obtain an estimate of trend (and the growth of the trend) capital deepening over the same time period.

Combining the growth rate of the trend capital deepening series with the capital share in output (α_k) provides the contribution of the growth rate of trend capital deepening to the growth rate of trend labour productivity.

Step 2: The growth rate of trend TFP

For TFP, we begin by estimating the growth rate of TFP over the sample period using the growth accounting framework. In this framework, the TFP variable picks up many things. TFP is generally thought to be a measure of technology (or technological progress in the growth accounting framework). However, TFP is a residual and as a result captures many things other than technology. As mentioned above, it will include growth in the public sector capital stock (since the capital measure is only for the business sector). It also captures capacity utilization, measurement error, returns to scale and a number of other factors. For this reason, it can be difficult to interpret movements in TFP over time.

Moreover, the growth accounting framework does not allow for the identification of the level of TFP, only the growth rate of TFP. Nevertheless, to establish the trend in the level of TFP, an an index number (normalized to 100 in 1990) is constructed using the growth rate of TFP, and

[21] Where a capital gap may exist is in capital utilization, since some of the capital stock may not be used for production in the current period. Capital utilization will therefore be captured in TFP in this model.

the band-pass filter, detailed in Christiano and Fitzgerald (2003), is used to estimate the trend in the level (index number) over time, and then to calculate the growth rate of this trend.[22]

In order to get a forecast for the growth rate of trend TFP, the data are analyzed to determine historical mean growth rates over different periods, and based on this as well as information on investment and future expected developments in intangible capital, the future growth rate of trend TFP is identified. The long-run growth rate of trend TFP is approximately 0.8 per cent. The band-pass filtered data suggest that there was a significant slowdown in the growth rate of trend TFP since 2000, but that the growth rate has been gradually returning to its long-run average since 2009. Using this information, judgment is used to increase the growth rate of TFP through the recession and to determine the forecasted values.

Step 3: Estimating the level of trend labour productivity

In this final step, the estimates of the growth rate of trend capital deepening and TFP are used to identify the growth rate of trend labour productivity, and then additional information is used to identify the level. Because we are using growth rates to identify a level over time, judgment is used to identify a period where it is believed that trend labour productivity was equal to actual labour productivity. From there, the combined growth rates of trend capital deepening and TFP are used to map out the level of trend labour productivity. It is determined that setting actual productivity equal to its trend in 1998 provides a long-run average annual labour productivity gap of -0.1 per cent (the smallest possible gap across the historical sample). Moreover, productivity growth was particularly strong in the second half of the 1990s, and 1998 represents the midpoint in the high growth period. For these reasons, it is believed that it is reasonable to use this as a starting point for determining the gap. Further judgment is used, if deemed necessary. With this, we can establish the level of trend labour productivity and identify the current labour productivity gap.

Alternatively, one could assume a Cobb-Douglas production function and estimate the level of potential output directly. However, it would be neccessary to make strong assumptions about the share of labour in output. Because it is assumed to be constant in the Cobb-Douglas framework, a constant level for the labour share would need to be chosen, which would, in part, determine the level of potential output. It is believed to be easier and more transparent to determine the growth rate of potential (using the growth accounting framework) and then back out the level, than to have the constant share of capital determine, in part, the level of potential output.

[22] The band-pass filter approach for identifying trend TFP is one of several methods that have been used to estimate trend TFP. Efforts have been made to develop a reduced-form model that would link TFP with investment and other variables, but so far this has not been helpful for identifying trend TFP. Work on this issue is ongoing.

It is important to note that up to this point, only annual data are being used. Quarterly estimates of the level of potential output are derived via interpolation. Once the quarterly level is attained, further judgment is used to refine the current level of potential, and hence the historical and future values of potential output.