Technical Report No. 79 / Rapport technique n° 79

Measurement of the Output Gap: A Discussion of Recent Research at the Bank of Canada

by Pierre St-Amant and Simon van Norden

Bank of Canada Banque du Canada

August 1997

Measurement of the Output Gap: A Discussion of Recent Research at the Bank of Canada

Pierre St-Amant and Simon van Norden

ISSN 0713-7931
ISBN 0-662-26019-8

Printed in Canada on recycled paper

ACKNOWLEDGMENTS

The authors wish to thank Pierre Duguay, Chantal Dupasquier, Paul Fenton, Gabriele Galati, Alain Guay, Seamus Hogan, Irene Ip, Robert Lafrance, René Lalonde, David Longworth, Tiff Macklem, John Murray, and Brian O'Reilly for useful comments and discussions. They also thank Jennifer Page and Rebecca Szetto for their excellent research assistance and Patricia Buchanan and David R. Miller for the final editing. Of course, since the authors are solely responsible for the paper's content, none of the above mentioned are responsible for any remaining errors.

ABSTRACT

In this paper, we discuss some methodologies for estimating potential output and the output gap that have recently been studied at the Bank of Canada. The assumptions and econometric techniques used by the different methodologies are discussed in turn, and applications to Canadian data are presented.

The use of the Hodrick-Prescott (HP) filter to measure the output gap has been justified on the basis that this filter extracts business-cycle frequencies from the data and that it can estimate an unobserved cyclical component. We note that the HP filter is unlikely to do well in achieving these objectives for series whose spectra have the typical Granger shape, such as real output, and that it will often fail to measure cyclical components adequately. The problems of the HP filter are accentuated at the end of samples, which is the place most relevant for policymakers. Finally, we note that univariate filters will only be able to give us information about the current output gap if the gap is Granger-caused by output growth; this is not the case if we believe that potential output is exogenous.

Extensions to the HP filter, such as those proposed by Laxton and Tetlow (1992) and Butler (1996), have focussed on incorporating additional information derived from assumed or estimated economic relationships. The motivation behind these "hybrid" methods is a desire to obtain estimates of the output gap that are conditioned by structural information but that remain "smooth." However, existing hybrid methods have proved hard to estimate. In addition, they may not be robust to alternative reasonable calibrations, and they do not allow for easy calculation of confidence intervals. We also find that Butler's method does not perform as well as the simple HP filter in terms of isolating fluctuations of output originating from business-cycle frequencies. We also discuss the "TOFU" approach (a Trivial Optimal Filter that may be Useful), which replaces the smoothness assumptions of the hybrid methods with an unrestricted but linear filter.

We then turn to multivariate filtering methods based on VARs that incorporate long-run restrictions. Unlike univariate filters, VAR-based

methods do not suffer from obvious end-of-sample problems, and they can provide projected values for the output gap. Relative to other multivariate methods (such as the multivariate Beveridge-Nelson method), one advantage of the VAR method using long-run restrictions is that it does not restrict the dynamics of potential output a priori. We investigate the implications of long-run restrictions on real output only and on real output *and* inflation. We argue that the latter approach should be of interest for policymakers focussing on movements of real output associated with movements in the trend of inflation. Unfortunately, the VAR applications that we consider display wide confidence intervals, similar to those reported on the basis of other methods. Using VARMAs or constrained VARs instead of unconstrained VARs may reduce that uncertainty.

RÉSUMÉ

Les auteurs de l'étude se penchent sur les méthodes d'estimation de la production potentielle et de l'écart de production récemment étudiées à la Banque du Canada, passent en revue les hypothèses et les techniques économétriques sur lesquelles repose chacune d'elles et fournissent des exemples d'applications aux données canadiennes.

Certains économistes justifient l'emploi du filtre de Hodrick-Prescott (HP) pour mesurer l'écart de production par le fait que ce filtre permet d'extraire des données les fréquences correspondant au cycle économique et d'estimer une composante cyclique non observée. Les auteurs de l'étude constatent quant à eux que le filtre HP réussit mal à extraire ces fréquences dans le cas de séries telles que la production réelle, dont la forme spectrale ressemble à celle que Granger a mise en lumière, et qu'il parvient rarement à mesurer correctement la composante cyclique. Les lacunes du filtre HP s'accentuent également en fin d'échantillon, là où les valeurs présentent le plus d'intérêt pour les décideurs publics. Enfin, les auteurs font remarquer que les filtres univariés ne peuvent fournir des indications sur l'écart de production observé que si cet écart est déterminé, au sens de Granger, par la croissance de la production. Si la production potentielle est considérée comme exogène, ces filtres ne sont d'aucune utilité.

Les tentatives visant à améliorer la tenue du filtre HP, comme celles de Laxton et Tetlow (1992) et de Butler (1996), ont surtout consisté à lui incorporer des éléments d'information additionnels tirés de relations économiques hypothétiques ou estimées. Ces méthodes hybrides sont motivées par le désir d'obtenir des estimations de l'écart de production qui tirent parti de renseignements de nature structurelle mais dont le profil reste lisse. Toutefois, il s'est avéré difficile d'estimer l'écart de production à l'aide de ces méthodes. En outre, celles-ci sont sensibles au choix de l'étalonnage, et le calcul des intervalles de confiance est malaisé. Les auteurs de l'étude constatent également que la méthode proposée par Butler ne réussit pas aussi bien qu'un simple filtre HP à isoler les fluctuations de la production qui sont d'origine conjoncturelle. Ils examinent aussi l'approche « TOFU », qui repose sur l'emploi d'un filtre ne comportant aucune restric-

tion mais ayant une forme linéaire, et non sur l'imposition de contraintes de lissage, propre aux méthodes hybrides.

Les auteurs analysent ensuite les méthodes de filtrage à plusieurs variables qui font appel à des vecteurs autorégressifs (VAR) assortis de restrictions de long terme. À la différence des filtres univariés, ces méthodes ne présentent pas de lacunes manifestes en fin d'échantillon et permettent de prévoir les valeurs de l'écart de production. Comparativement aux autres méthodes multivariées (p. ex. la méthode de Beveridge-Nelson), les méthodes reposant sur l'emploi de VAR assortis de contraintes de long terme ont l'avantage de ne pas restreindre a priori la dynamique de la production potentielle. Les auteurs étudient les conséquences de l'imposition de restrictions de long terme à la production réelle seule, puis à la fois à la production réelle et à l'inflation. Ils font valoir que les résultats de cette deuxième approche devraient intéresser les décideurs publics qui s'attachent aux mouvements de la production réelle associés aux variations du taux de l'inflation tendancielle. Malheureusement, les résultats qu'ils obtiennent au moyen de VAR sont assortis d'intervalles de confiance aussi larges que ceux que produisent les autres méthodes. Il est possible que l'utilisation de VARMA au lieu de simples VAR réduise l'incertitude relative aux estimations.

CONTENTS

1 INTRODUCTION

Most macroeconomic models used for forecasting and policy analysis require an estimate of potential output. At the Bank of Canada, for example, estimates of potential output are important inputs in various "Phillips-curve" models and in the staff's Quarterly Projection Model, where the gap between actual and potential output is a key variable determining the evolution of prices and wages. A level of real GDP above potential (a positive output gap) will often be seen as a source of inflationary pressures and a signal that monetary authorities interested in avoiding an acceleration of inflation should tighten monetary conditions. A level of real GDP below potential (a negative output gap) will have the opposite implication.

The output gap can thus be defined as the component of real output that is associated with changes in inflation.[1] Note that gaps could be calculated in markets other than goods and services. For example, gaps in the labour market have frequently been calculated, and authors such as Hendry (1995) have presented "money gaps."

Unfortunately, measuring the output gap is not an easy task. Different sets of assumptions can be used in combination with various econometric techniques to provide different measures of the output gap. One common assumption is that the output gap is some part of the transitory (cyclical) component of real output. The methods discussed in this paper make that assumption.

The purpose of this paper is to examine in detail the properties of the various methods themselves rather than the economic reasonableness of the measures produced by those methods.

The first group of methods we consider are those that simply use some implicit or explicit assumptions about the dynamics of real output to

1. To be more precise, we should take into account expected inflation, and therefore define the output gap with respect to *unexpected* changes in inflation. Some models also imply a relationship between the change in the gap and inflation.

identify the output gap. For example, if real output is believed to be composed of a stationary component and a simple log-linear trend, the output gap could be measured as the residuals of a regression of log output on a linear time trend. Unfortunately, such a simple model does not adequately describe the behaviour of output, and measuring the temporary component in more complex models is problematic.

In this paper, we will assume that real output is I(1); that is, that the level of output is subject to permanent shocks so there is no deterministic trend towards which output tends to revert.[2] Many approaches have been proposed to identify the permanent and cyclical components of real output in such models, such as those proposed by Hodrick and Prescott (1997), Watson (1986) or Beveridge and Nelson (1981). The problem is that the measured cyclical component may differ considerably from one method to another. Quah (1992) argues that this is an intrinsic problem and that "...without additional ad hoc restrictions those [univariate] characterizations are completely uninformative."

These problems have not prevented the widespread use of Hodrick and Prescott's filter to identify the cyclical component of output.[3] Arguments commonly made to justify its use are that

- it extracts the relevant business-cycle frequencies of output
- it closely approximates the cyclical component implied by reasonable time-series models of output

2. This is the most common assumption in modern applied macroeconomics and is consistent with the view that real output can be permanently affected by shocks such as technological innovations. An alternative view is that output is stationary around a time trend, but that this time trend is subject to occasional random changes in its slope and intercept. Evidence for such a view is discussed by Perron (1989) and Weber (1995). Since detecting changes in the slope or intercept near the end of a sample is quite difficult, such models imply that one cannot reliably measure the current deviations from trend. Since we argue below that this is what policymakers wish to measure, adoption of the "breaking-trend" model per se is not a solution to the problems of measuring output gaps.

3. We will henceforth refer to this method as the HP filter, although Hodrick and Prescott note that their method owes much to Whittaker (1923) and Henderson (1924). Note that, although the Hodrick and Prescott article was published in 1997, their working paper dates from 1981.

We examine these arguments in Sections 2.2 and 2.3. We also note that, unlike much of the literature on “detrending,” the main problem facing policymakers is to estimate the deviation from trend at the end rather than in the middle of a data sample.[4] We conclude that such methods are unlikely to be suitable for use in a policy context, and we discuss economic factors that limit our ability to estimate the current output gap.

An important class of alternatives to these univariate dynamic methods are approaches that combine their assumptions with information from assumed or “structural” relationships between the output gap and other economic variables, such as a Phillips curve or Okun’s law. We examine some of these in Section 3. Among them are the multivariate HP filters (MHPF) proposed by Laxton and Tetlow (1992) and Butler (1996), representing the general approach currently used in the staff economic projection of the Canadian economy at the Bank of Canada. In Section 3, we note that calibration of the MHPF methods has been problematic and that their estimates of the output gap have wide confidence intervals despite the inclusion of structural information. Spectral analysis of the Canadian output gap resulting from the application of the MHPF method also gives the disturbing result that the gap includes a very large proportion of cycles that are much longer than what is usually defined as a typical business cycle. A reaction to these methods is the “Trivial Optimal Filter that may be Useful” (TOFU) approach suggested by van Norden (1995), which replaces the HP smoothing problem with the simpler restriction of a constant linear filter. The TOFU approach has yet to be shown to be workable.

The third and final class of methods that we consider uses multivariate rather than univariate dynamic relationships, often in combination with structural relationships from economic theory, to estimate output gaps as a particular transitory component of real output. Some of these are

4. This is an oversimplification. More accurately, policymakers will usually be most interested in expected *future* values of the output gap, particularly when these expectations are conditioned by specific policy actions. This is more demanding than simply estimating the output gap at the end of sample, so our discussion of the additional difficulties introduced by end-of-sample problems underestimates the true difficulty of the policy problem. For that reason, we think good end-of-sample performance is a necessary rather than a sufficient condition for reliable estimation of the deviation from trend.

examined in Section 4. One example is the decomposition method suggested by Cochrane (1994) (henceforth CO). This method is based on the permanent-income hypothesis and uses consumption to define the permanent component of output, which can then be used as a measure of potential output. Multivariate extensions of the Beveridge-Nelson decomposition method (MBN) have also been proposed to identify the permanent component of output (Evans and Reichlin 1994). A major restriction, used by both the CO and the MBN methods, is that the permanent component of real output is a random walk.

Section 4.1 of this paper, which draws heavily on Dupasquier, Guay, and St-Amant (1997), discusses the CO and MBN methodologies and compares them with a structural vector autoregression methodology based on long-run restrictions imposed on output (LRRO). This method was proposed by Blanchard and Quah (1989), Shapiro and Watson (1988), and King et al. (1991). One characteristic of the LRRO approach is that it does not impose restrictions on the dynamics of the permanent component of output. Instead, it allows for a permanent component comprising an estimated diffusion process for permanent shocks that can differ from a random walk. The output gap then corresponds to the cyclical component of output excluding the diffusion process of permanent shocks, which is instead assigned to potential output. Section 4.2 presents an application of the LRRO method to Canadian data.

In Section 4.3 (which draws from Lalonde, Page, and St-Amant (forthcoming)), we present another methodology based on long-run restrictions imposed on a VAR that associates restrictions imposed on real output *and* inflation. The output gap is then a part of the cyclical component of real output that is consistent with changes in the trend of inflation.[5]

The final section concludes with some directions for future research.

5. Lalonde, Page, and St-Amant also present a method that associates the output gap with changes in the trend of inflation but does not impose that the output gap is stationary.

2 THE HP FILTER

In recent years, mechanical filters have frequently been used to identify the permanent and cyclical components of time series. The most popular of these mechanical filters is that proposed by Hodrick and Prescott (1997). This section evaluates the basic Hodrick-Prescott (HP) filter's ability to provide a useful estimate of the output gap. Section 3 then discusses some recently proposed extensions and alternatives to the basic HP filter.

Guay and St-Amant (1996) show that the HP filter does a poor job of extracting business-cycle frequencies from macroeconomic time series. Consequently, it does not constitute an adequate approach for estimating an output gap constrained to correspond to the business-cycle frequencies of real GDP. This is discussed in Section 2.2, where we further argue that constraining the output gap in that way is not very attractive in any case. Guay and St-Amant also show that the HP filter is likely to do a poor job of extracting an output gap assumed to correspond to the unobserved cyclical component of real GDP. This is discussed in Section 2.3. In Section 2.4, we focus explicitly on the HP filter's end-of-sample problems and conclude that these raise further doubts about the appropriateness of using the HP filter to estimate the output gap. Finally, Section 2.5 investigates what economic theory has to say about the possible usefulness of filters for estimating output gaps at the end of sample.

Most of the arguments in this section of the paper are drawn from Guay and St-Amant (1996) and van Norden (1995). Note that Guay and St-Amant show that the main conclusions they reach concerning the HP filter also apply to the band-pass filter proposed by Baxter and King (1995).

2.1 The optimization problem

The HP filter decomposes a time series y_t into an additive cyclical component y_t^c and a growth component y_t^g:

$$y_t = y_t^g + y_t^c . \tag{1}$$

Applying the HP filter involves minimizing the variance of the cyclical component y_t^c subject to a penalty for the variation in the second difference of the growth component y_t^g. This is expressed in the following equation:

$$\{y_t^g\}_{t=0}^{T+1} = \operatorname{argmin} \sum_{t=1}^{T} \left[\left(y_t - y_t^g\right)^2 + \lambda[(y_{t+1}^g - y_t^g) - (y_t^g - y_{t-1}^g)]^2\right], \qquad (2)$$

where λ, the smoothness parameter, penalizes the variability in the growth component. The larger the value of λ, the smoother the growth component. As λ approaches infinity, the growth component corresponds to a linear time trend. For quarterly data, Hodrick and Prescott propose setting λ equal to 1,600. King and Rebelo (1993) show that the HP filter can render stationary any integrated process of up to the fourth order.

2.2 How well does the HP filter extract business-cycle frequencies?

Authors such as Singleton (1988) have shown that the HP filter can provide an adequate approximation of a high-pass filter when it is applied to stationary time series. Here we need to introduce some elements of spectral analysis. A zero-mean stationary process has a Cramer representation like:

$$y_t = \int_{-\pi}^{\pi} \varepsilon^{i\omega t} dz(\omega), \qquad (3)$$

where $dz(\omega)$ is a complex value of orthogonal increments, i is the imaginary number ($\sqrt{(-1)}$) and ω is frequency measured in radians, i.e., $-\pi \le \omega \le \pi$ (see Priestley 1981, Chapter 4). In turn, filtered time series can be expressed as

$$y_t^f = \int_{-\pi}^{\pi} \alpha(\omega) e^{i\omega t} dz(\omega), \qquad (4)$$

with

$$\alpha(\omega) = \sum_{h=-k}^{k} a_h e^{-i\omega h}. \qquad (5)$$

Equation (5) is the frequency response (Fourier transform) of the filter. That is, $\alpha(\omega)$ indicates the extent to which y_t^f responds to y_t at frequency ω

and can be seen as the weight attached to the periodic component $e^{i\omega t}dz(\omega)$. In the case of symmetric filters, the Fourier transform is also called the gain of the filter.

An ideal high-pass filter would remove low-frequency, or long-cycle, components and allow high-frequency, or short-cycle, components to pass through, so that $\alpha(\omega) = 0$ for $|\omega| \leq \omega^p$, where ω^p has some predetermined value and $\alpha(\omega) = 1$ for $|\omega| > \omega^p$. Figure 1 shows the squared gain of the HP filter. Very high frequencies are left out because the focus is business-cycle frequencies as defined by NBER researchers since Burns and Mitchell (1946), i.e., cycles lasting no fewer than six and no more than thirty-two quarters. We see that the squared gain is 0 at zero frequency and is close to 1 from around frequency $\pi/8$ (cycles of sixteen quarters) and up. On the basis of Figure 1, the HP filter would appear to be an adequate approximation of a high-pass filter in that it removes most low frequencies and passes through most higher frequencies, including business-cycle frequencies.

The output gap could be associated with business-cycle frequencies plus higher frequency volatility in the data. Figure 1 would then suggest that the HP filter is an adequate measure of the output gap. One problem with this approach is that most macroeconomic time series are either integrated or highly persistent processes. In their study, Guay and St-Amant (1996) conduct a systematic investigation of the HP filter's ability to capture business-cycle frequencies, i.e., the area delimited by the spectrum of an original series at frequencies between six and thirty-two quarters. Their main finding is that, when the peak of a series is at zero frequency and the bulk of the variance is located in low frequencies, which is the shape described by Granger as typical for macroeconomic time series, the HP filter cannot adequately capture business-cycle frequencies. This is illustrated in Figure 2, which shows the spectrum of an autoregressive process having its peak at zero frequency and the spectrum of the cyclical component resulting from the application of the HP filter.

In Figure 2, the spectrum of the cyclical component resulting from the application of the HP filter is very different from that of the original

FIGURE 1. Squared gain of the HP filter

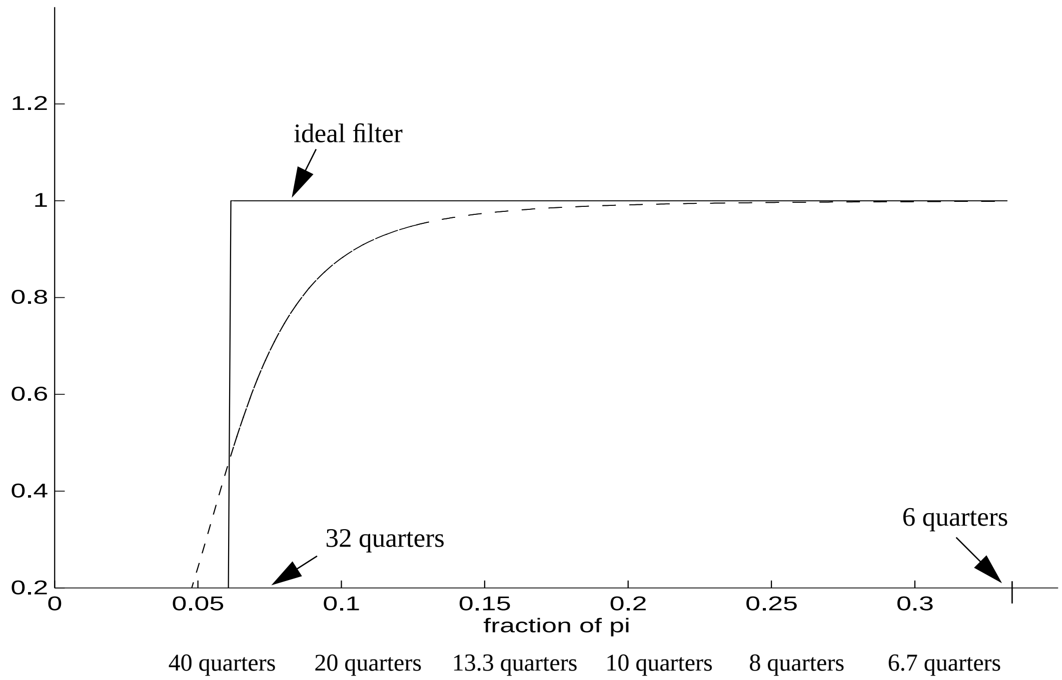

FIGURE 2. Spectrum of a series with Granger's typical shape and of that series HP-filtered (AR(3) coefficients: 1.2, -0.11, -0.16)

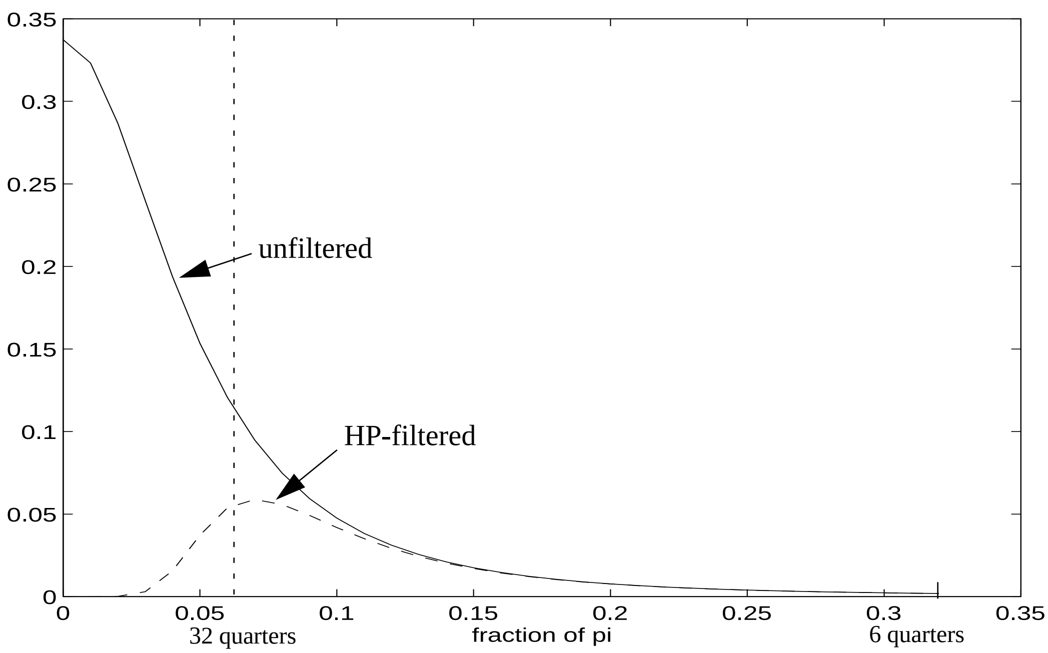

series. This comes as no surprise, since the filter is designed to extract low frequencies from the data. However, we can see that business-cycle frequencies are not left intact. In particular, the HP filter induces a peak inside business-cycle frequencies even though such a peak is absent from the original series, and it fails to capture a significant fraction of the variance contained in business-cycle frequencies but captures some variance originating outside these frequencies. Guay and St-Amant (1996) show that this is typical of time series having the typical Granger shape, i.e., most macroeconomic series. Indeed, the unfiltered spectrum shown in Figure 2 is a parametric estimate of the spectrum of U.S. real GDP.

The intuition behind this result is simple. Figure 1 shows that the gain of the HP filter at low business-cycle frequencies is smaller than that of the ideal filter. Indeed, the squared gain of the HP filter is 0.49 at frequencies corresponding to 32-quarter cycles and does not reach 0.95 before frequency $\pi/8$. Note also that the squared gain does not fall immediately to zero at lower frequencies. The problem is that a large fraction of the power of typical macroeconomic time series is concentrated in the band where the squared gain of the HP filter differs from that of an ideal filter. Also, the shape of the squared gain of the HP filter is such that a peak in the spectrum of the cyclical component is induced when it is applied to typical macroeconomic time series. In short, applying the HP filter to series dominated by low frequencies results in the extraction of a cyclical component that does not capture an important fraction of the variance contained in the business-cycle frequencies of the original series, but does capture an important part of variance situated at lower frequencies than business-cycle frequencies and induces spurious dynamic properties.

An additional problem is that associating the output gap with the business-cycle frequencies in the data might not be a good idea in the first place. Note in particular that part of the variance associated with business-cycle frequencies could reflect the dynamics of shocks to potential output. As noted by King et al. (1991), “productivity shock[s] [set] off transitional dynamics, as capital is accumulated and the economy moves toward a new steady state.” To the extent that such dynamics reflect the evolution of potential output itself, one might prefer to use a different approach to identify potential output and the output gap. Section 4 of the paper provides a more detailed discussion of this point.

2.3 How well does the HP filter extract the cyclical component?

In the previous section, we have seen that the HP filter does not have sufficiently good spectral properties to be able to isolate accurately the component of a series associated with fluctuations at business-cycle frequencies. As discussed by King and Rebelo (1993), another justification for the use of the HP filter is that in some cases it will be the optimal filter for identifying the cyclical component of a series. However, as shown by King and Rebelo and by Harvey and Jaeger (1993), these are cases when, in particular, the series is I(2), there are identical propagation mechanisms for innovations in the growth rate and in the cycle (or the transitory component is white noise), and the smoothing parameter λ is known. These conditions are rarely met in practice.

Of course, the fact that the HP filter is not an optimal filter does not necessarily mean that it will not be a good approximation of an optimal filter. We therefore consider whether the HP filter can reliably isolate the cyclical component of a variety of time series.

It is often argued that macroeconomic time series actually comprise a permanent component and a cyclical component. The permanent component could be driven by an I(1) technological process with drift, while monetary shocks, among others, could generate the cyclical component. In order to assess the HP filter's ability to extract such a cyclical component, consider the following data-generating process (DGP):

$$y_t = \mu_t + c_t \,, \tag{6}$$

where

$$\mu_t = \mu_{t-1} + \varepsilon_t \tag{7}$$

$$c_t = \phi_1 c_{t-1} + \phi_2 c_{t-2} + \eta_t \tag{8}$$

and

$$\varepsilon_t \sim NID\left(0, \sigma_\varepsilon^2\right),\ \ \eta_t \sim NID\left(0, \sigma_\eta^2\right). \tag{9}$$

Equation (6) defines y_t as the sum of a permanent component, μ_t, which in this case corresponds to a random walk, and a cyclical component, c_t.[6] The dynamics of the cyclical component are specified as a second-order autoregressive process, so that the peak of the spectrum could be at zero frequency or at business-cycle frequencies. We assume that ε_t and η_t are uncorrelated.

Data are generated from equation (6) with ϕ_1 set at 1.2 and different values for ϕ_2 to control the location of the peak in the spectrum of the cyclical component. We also vary the standard-error ratio for the disturbances $\sigma_\varepsilon/\sigma_\eta$ to change the relative importance of each component. We follow the standard practice of assigning a value of 1,600 to λ, the HP filter smoothness parameter. We also follow Baxter and King's (1995) suggestion of dropping 12 observations at the beginning and at the end of the sample, which should favour the filter considerably by partly eliminating its end-of-sample problems (see Section 2.4). The resulting series contains 150 observations, a standard size for quarterly macroeconomic data. The number of replications is 500.

The performance of the HP filter is assessed by comparing the autocorrelation function of the cyclical component of the true process with that obtained from the filtered data. We also calculate the correlation between the true cyclical component and the filtered cyclical component and report their relative standard deviations ($\hat{\sigma}_c/\sigma_c$). Table 1 presents the results.

Table 1 shows that the HP filter performs particularly poorly when there is an important permanent component. Indeed, in most cases the correlation between the true and the filtered components is not significantly different from zero for high $\sigma_\varepsilon/\sigma_\eta$ ratios. The estimated autocorrelation function is invariant to the change in the cyclical component in these cases (the values of the true autocorrelation functions are given in brackets and

6. This is Watson's (1986) specification for U.S. real GDP.

TABLE 1: Simulation results for the HP filter

DGP			Estimated values				
			Autocorrelations				
$\sigma_\varepsilon/\sigma_\eta$	ϕ_1	ϕ_2	1	2	3	Correlation	$\hat{\sigma}_c/\sigma_c$
10	0	0	.71[0] (.59,.80)	.46[0] (.30,.60)	.26[0] (.08,.43)	.08 (-.07,.21)	12.96 (10.57,15.90)
10	1.2	-.25	.71[.96] (.61,.80)	.47[.90] (.31,.61)	.27[.84] (.08,.44)	.08 (-.11,.28)	4.19 (2.77,6.01)
10	1.2	-.40	.71[.86] (.60,.80)	.46[.63] (.30,.60)	.26[.41] (.08,.44)	.13 (-.12,.36)	6.34 (4.82,8.07)
10	1.2	-.55	.71[.77] (.60,.80)	.46[.38] (.29,.60)	.26[.03] (.06,.43)	.14 (-.08,.33)	6.93 (5.36,8.70)
10	1.2	-.75	.71[.69] (.60,.78)	.46[.27] (.30,.59)	.25[-.19] (.07,.41)	.15 (-.01,.31)	6.37 (4.79,7.95)
5	0	0	.69[0] (.58,.78)	.45[0] (.30,.58)	.26[0] (.09,.41)	.15 (.02,.27)	6.50 (5.28,7.85)
5	1.2	-.25	.71[.96] (.61,.80)	.46[.90] (.32,.61)	.26[.84] (.08,.43)	.16 (-.01,.36)	2.11 (1.43,3.04)
5	1.2	-.40	.72[.86] (.61,.80)	.46[.63] (.31,.60)	.25[.41] (.08,.42)	.23 (-.01,.45)	3.26 (2.47,4.15)
5	1.2	-.55	.71[.77] (.61,.80)	.46[.38] (.30,.59)	.24[.03] (.06,.41)	.24 (.01,.44)	3.60 (2.83,4.52)
5	1.2	-.75	.70[.69] (.61,.79)	.43[.27] (.26,.57)	.20[-.19] (.00,.38)	.29 (.11,.44)	3.30 (2.53,4.17)
1	0	0	.43[0] (.27,.57)	.28[0] (.11,.42)	.20[0] (-.02,.31)	.59 (.49,.70)	1.61 (1.41,1.85)
1	1.2	-.25	.76[.96] (.67,.83)	.51[.90] (.37,.62)	.29[.84] (.11,.44)	.51 (.33,.68)	.66 (.44,.91)
1	1.2	-.40	.75[.86] (.67,.81)	.44[.63] (.28,.55)	.16[.41] (-.03,.33)	.71 (.56,.82)	1.02 (.83,1.22)
1	1.2	-.55	.72[.77] (.66,.78)	.34[.38] (.21,.47)	.01[.03] (-.17,.19)	.76 (.56,.82)	1.15 (.83,1.22)
1	1.2	-.75	.68[.69] (.63,.72)	.15[.27] (.04,.27)	-.27[-.19] (-.44,.10)	.83 (.75,.89)	1.16 (1.04,1.29)
.5	0	0	.16[0] (.01,.32)	.10[0] (-.04,.24)	.04[0] (-.10,.18)	.82 (.75,.88)	1.16 (1.07,1.27)
.5	1.2	-.25	.79[.96] (.71,.85)	.53[.90] (.38,.65)	.30[.84] (.11,.46)	.61 (.41,.79)	.55 (.37,.76)
.5	1.2	-.40	.77[.86] (.69,.81)	.43[.63] (.29,.54)	.13[.41] (-.05,.29)	.84 (.73,.92)	.87 (.74,.99)
.5	1.2	-.55	.72[.77] (.67,.78)	.28[.38] (.17,.39)	-.10[.03] (-.25,.06)	.89 (.83,.94)	.98 (.89,1.07)
.5	1.2	-.75	.67[.69] (.63,.71)	.07[.27] (-.03,.18)	-.42[-.19] (-.57,-.27)	.94 (.90,.96)	1.02 (.97,1.08)
.01	0	0	-.08[0] (-.21,.06)	-.06[0] (-.21,.06)	-.06[0] (-.19,.06)	.98 (.96,.99)	.97 (.94,.99)
.01	1.2	-.25	.80[.96] (.72,.86)	.54[.90] (.38,.67)	.30[.84] (.11,.48)	.66 (.45,.83)	.51 (.34,.69)
.01	1.2	-.40	.78[.86] (.72,.83)	.43[.63] (.30,.55)	.12[.41] (-.05,.28)	.90 (.82,.96)	.81 (.71,.90)
.01	1.2	-.55	.73[.77] (.67,.77)	.26[.38] (.15,.37)	-.14[.03] (-.30,.01)	.96 (.91,.99)	.92 (.86,.96)
.01	1.2	-.75	.67[.69] (.62,.71)	.02[.27] (-.08,.13)	-.50[-.19] (-.61,-.35)	.99 (.97,1.0)	.97 (.95,.99)

the confidence band in parenthesis). When the ratio $\sigma_\varepsilon / \sigma_\eta$ is equal to 0.5 or 1 and the peak of the cyclical component is located at zero frequency ($\phi_2 > -0.43$), the dynamic properties of the true and the filtered cyclical components are significantly different, as indicated by the estimated parameter values. In general, the HP filter adequately characterizes the series' dynamics when the peak of the spectrum is at business-cycle frequencies and the ratio $\sigma_\varepsilon / \sigma_\eta$ is small. However, even when the ratio of standard deviations is equal to 0.01 (i.e., the permanent component is almost absent), the filter performs poorly when the peak of the spectrum of the cyclical component is at zero frequency. Indeed, for $\phi_2 = -0.25$, the dynamic properties of the filtered component differ significantly from those of the true cyclical component, the correlation is only equal to 0.66, and the standard deviation of the filtered cyclical component is half that of the true cyclical component.

It is interesting to note that the HP filter does relatively well when the ratio $\sigma_\varepsilon / \sigma_\eta$ is equal to 1, 0.5, or 0.01 and the spectrum of the original series has a peak at zero frequency and at business-cycle frequencies (i.e., the latter frequencies contain a significant part of the variance of the series). Consequently, the conditions required to identify adequately the cyclical component with the HP filter can be summarized as follows: the spectrum of the original series must have a peak located at business-cycle frequencies, and these frequencies must account for an important part of the variance of the series. If the variance of the series is dominated by low frequencies, which is the case for most macroeconomic series in levels, including real output, the HP filter does a poor job of extracting an output gap associated with the cyclical component of real output.

2.4 The HP filter at the end of samples

In examining the performance of the HP filter in the last two sections, we looked at how well it isolates particular business-cycle frequencies or the cyclical component of the series. Both cases implicitly looked at the performance of the HP filter over the available sample of data as a whole. However, it is useful to remember that the focus for policy advice is on estimating the *current* output gap. This is a more difficult task, since future information will presumably be useful in determining whether recent changes in output are persistent or transitory. We should therefore consider how the conclusions from the two previous sections might be affected by this added complication.[7]

To understand how the HP filter behaves at the end of sample, recall that the optimization problem it solves trades off the size of deviations from trend and the smoothness of that trend. In the face of a transitory shock, the filter is therefore “reluctant” to change the trend very much since this implies raising the trend before the shock and lowering it afterwards. At the end of the sample, however, the latter penalty is absent, implying that the optimal trend will be more responsive to transitory shocks at the end of sample than in mid-sample.

We can show this difference in several ways. Figure 3 shows the HP-filter trend expressed as a moving average of the unfiltered data. The weights in this moving average change as we move from the mid-sample towards the end of sample. The former gives us a smooth two-sided average in which no observation receives more than 6 per cent of the weight. The latter, however, gives a one-sided average where the last observation alone accounts for 20 per cent of the weight. Not surprisingly, this makes the HP trend more variable at the end of sample. Figure 4 shows that the deviations from the HP trend have different frequency responses as a result. In particular, the one-sided or end-of-sample filtered deviations from trend capture less of the variation at business-cycle frequencies (indicated by the dotted vertical lines).[8]

Figures 5 and 6 show how the deviations from the HP trend differ depending on whether we are at the end of sample or at mid-sample. The solid line in Figure 5 shows the usual deviation from the HP trend for Canadian GDP. The dashed line then shows the estimate obtained by using

7. This problem has been mentioned in other studies as well. Much of the analysis we present can also be found in Butler (1996).

8. The squared gains of the two HP trends also look quite different. For example, at the frequency corresponding to cycles of six quarters, the end-of-sample filter has a squared gain of about 1 while the mid-sample filter has a squared gain of about 0.75.

FIGURE 3. MA representation of the HP filter as a function of sample position (128 observations, $\lambda = 1,600$)

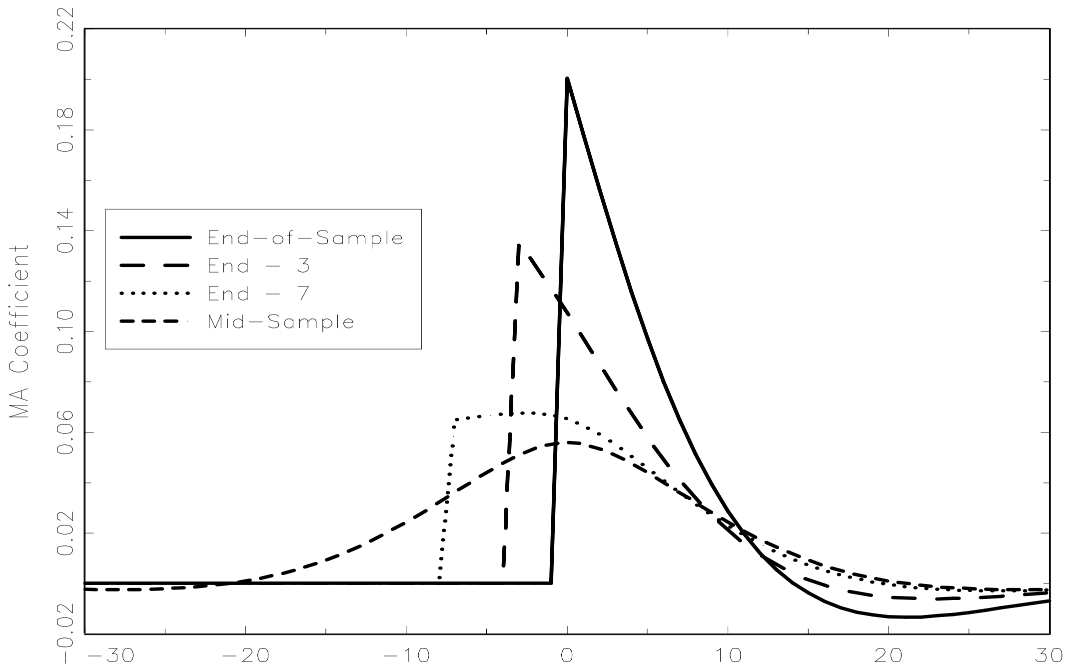

FIGURE 4. Squared gain of the HP Filter (128 observations, $\lambda = 1,600$)

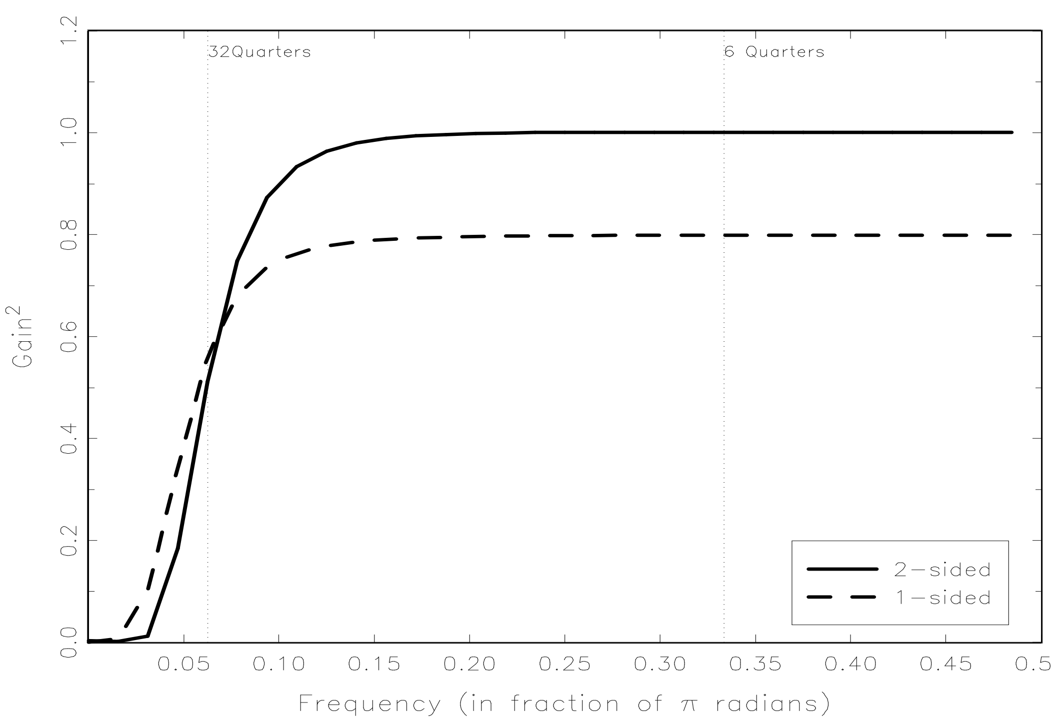

FIGURE 5. HP detrended real GDP (Canada, $\lambda = 1,600$)

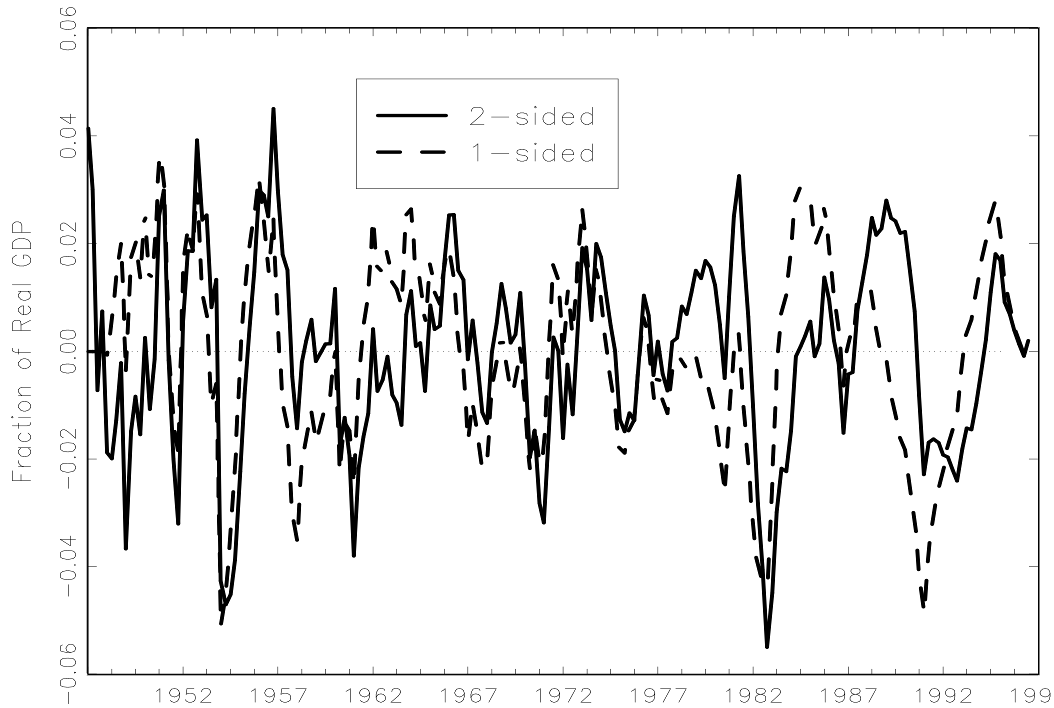

FIGURE 6. HP detrended real GDP, mid-sample to end of sample (Canada, $\lambda = 1,600$)

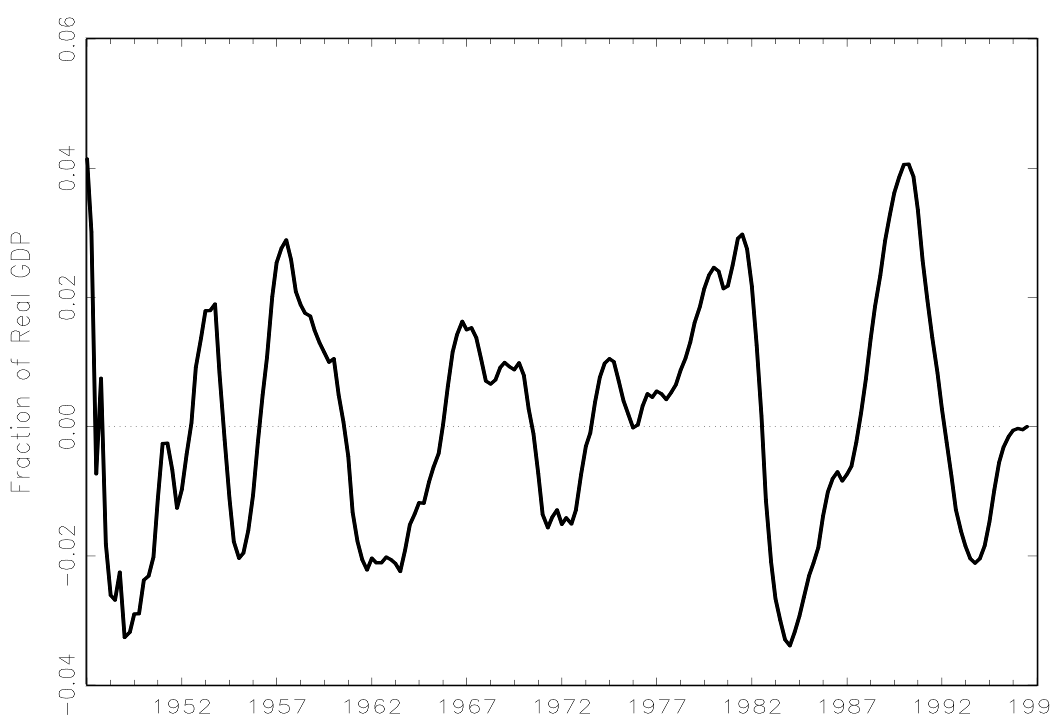

only data available up to that point in time (i.e., the corresponding end-of-sample estimates). Although the two series tend to move together, there are some important differences in size and timing. Comparing Figure 5 with Figure 6, we see that while deviations from trend are usually less than 3 per cent of GDP, the difference between its mid-sample and end-of-sample estimates is often as large as 2 per cent of GDP. Suppose we accept that the difference between these two measures is simply one component of the measurement error of end-of-sample estimates. This implies that the measurement errors of the latter must be roughly as large as the estimates themselves.[9] The implication is that end-of-sample estimates cannot be very reliable estimates of deviations from trend.

Figure 7 shows the implications of applying the HP filter to the "typical Granger-shape" series we considered previously. At the end of sample, even less of the variance of the deviations from the HP trend is due to variations at business-cycle frequencies and more is due to "leakage" from lower frequencies. This suggests that the results we obtained in Section 2.2 probably overstate the reliability of the HP filter for identifying an output gap associated with business-cycle frequencies. This is consistent with the results of Laxton and Tetlow (1992) and Butler (1996), who note that related filters also seem to perform worse at the end of samples. These related filters are discussed in Section 3.

9. We reach the same conclusion if we look at the range of the series, or at their standard deviations. The range (maximum - minimum) of the one-sided estimate is 8.7 per cent of GDP while the range of the difference between the one- and two-sided estimates is 7.5 per cent; the comparable standard errors are 1.8 per cent and 1.8 per cent. These comparisons are only approximate; small sample problems in the one-sided estimate at the beginning of the sample may make their difference appear excessively volatile, while the fact that the two estimates are constrained to be identical at the end of the sample will tend to understate the volatility of their difference.

FIGURE 7. Spectrum of series with typical Granger shape (128 observations, $\lambda = 1,600$ **)**

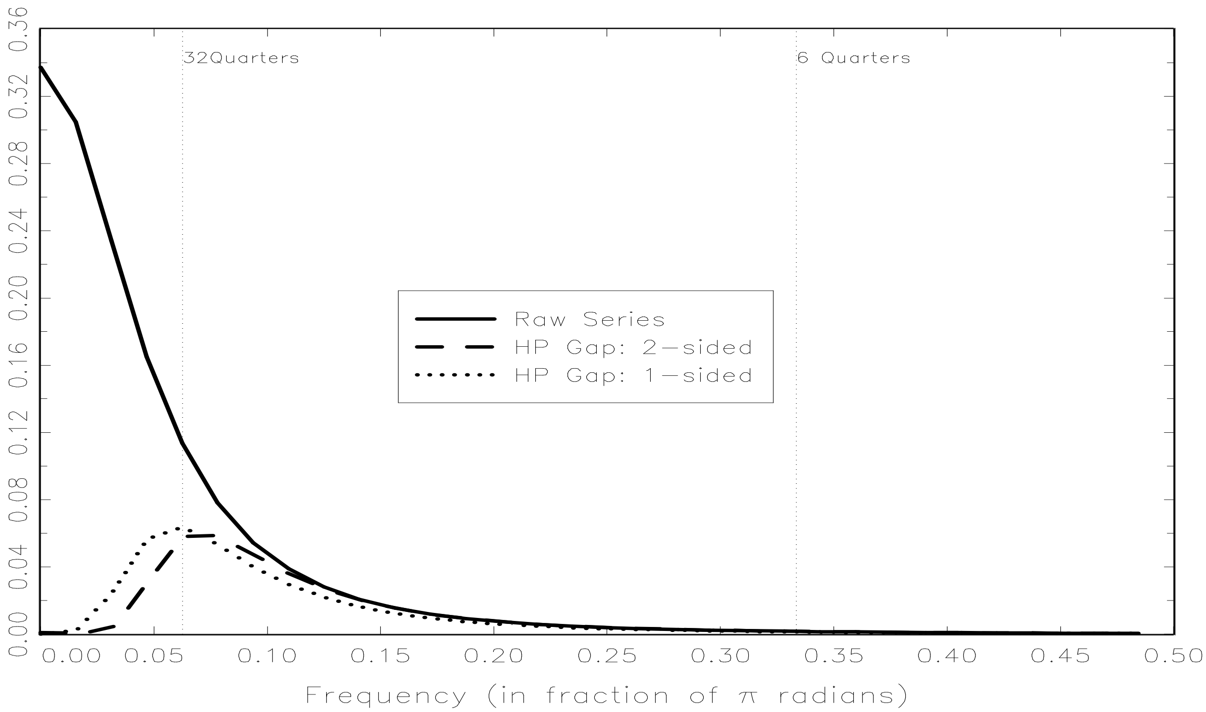

2.5 Limits to one-sided filtering

Part of the end-of-sample problem discussed in Section 2.4 reflects the fact that the HP filter behaves differently at the end of sample and at mid-sample, as shown in Figure 3. This suggests that other univariate filters might be able to measure output gaps more reliably. In this section, we consider one intrinsic limit to the ability of univariate filters to measure the current output gap, and show how this will in turn relate to beliefs about the economic relationships between actual and potential output. We show that models in which potential output is exogenous with respect to actual output and the output gap indicate that univariate filters will never be able to provide much information about contemporaneous output gaps.[10]

Suppose that potential output can be expressed as a linear filter of actual output, so that

10. This section draws heavily on van Norden (1995).

$$q_t = A(L) \cdot y_t + \varepsilon_t, \tag{10}$$

where q_t is (the log of) potential output, y_t is (the log of) actual output, ε_t is an innovations process that is uncorrelated with y_t at all leads and lags, and $A(L)$ is a two-sided polynomial in the lag operator (i.e., it takes a weighted sum of leads, lags and contemporaneous values of y_t). A sufficient but not necessary condition for such a representation to exist is that output y_t has a unit root and that the output gap $q_t - y_t$ is stationary.

We typically think of y_t as being non-stationary in mean, since it tends to drift upwards over time. To ensure that q_t and y_t move together in the long run (so that the gap is stationary), we will further assume that

$$A(1) = 1, \tag{11}$$

which simply means that the weights in $A(L)$ must sum to one. This in turn implies that we can write

$$A(L) - 1 = (1 - L) \cdot \tilde{A}(L), \tag{12}$$

and therefore that

$$q_t - y_t = \tilde{A}(L) \cdot \Delta y_t + \varepsilon_t. \tag{13}$$

Thus, we should be able to express the output gap as the weighted sum of past, present, and future output growth. The difference between equation (13) and the HP filter representation of the output gap is that the HP filter implies a particular set of restrictions on $\tilde{A}(L)$ that vary with the position in the sample. Let

$$E(q_t - y_t | H_y) = \tilde{A}(L) \cdot \Delta y_t \tag{14}$$

and

$$V((q_t - y_t) - E(q_t - y_t | H_y)) = \sigma^2, \tag{15}$$

where H_y is the set of all past, present, and future values of y_t.

So far, we have assumed that $\tilde{A}(L)$ is two-sided, whereas its use for policy purposes requires that it be one-sided. To understand how such a restriction on $\tilde{A}(L)$ will affect the accuracy of our estimate, note that the law of iterated expectations and equation (14) imply that

$$E(q_t - y_t | H_y^-) = E(E(q_t - y_t | H_y) | H_y^-) = E(\tilde{A}(L) \cdot \Delta y_t | H_y^-) , \tag{16}$$

where H_y^- is the set of all past values of y_t. If we define

$$\tilde{A}(L) = \tilde{A}^-(L) + \tilde{A}^+(L) , \tag{17}$$

where $\tilde{A}^-(L)$ has only positive powers of L and $\tilde{A}^+(L)$ has only non-positive powers of L, then equation (16) implies that

$$\begin{aligned} E(q_t - y_t | H_y^-) &= E(\tilde{A}^-(L) \cdot \Delta y_t | H_y^-) + E(\tilde{A}^+(L) \cdot \Delta y_t | H_y^-) \\ &= \tilde{A}^-(L) \cdot \Delta y_t + \sum_{j=1}^{m} a_j^+ \cdot E(\Delta y_{t+j} | H_y^-) = \bar{A}(L) \cdot \Delta y_t \end{aligned} \tag{18}$$

where a_j^+ is simply the coefficient on L^{-j} in $\tilde{A}^+(L)$. Similarly, we can show that

$$V(q_t - y_t | H_y^-) = V(q_t - y_t | H_y) + V(\tilde{A}(L) \cdot \Delta y_t | H_y^-) = \sigma^2 + V(\tilde{A}^+(L) \cdot \Delta y_t | H_y^-) , \tag{19}$$

where $V(X | \Omega)$ is the variance of the error in forecasting X given the information set Ω.

Equation (18) and equation (19) have an intuitive interpretation. The extent to which the one-sided filter $\bar{A}(L)$ is less informative than $\tilde{A}(L)$ will depend on the weight that $\tilde{A}(L)$ ascribes to current and future values of Δy_t and the extent to which those future values can be predicted from current and past values. The former will in turn depend on the Granger-causal relationship between $q_t - y_t$ and Δy_t, while the latter will depend on the degree to which output growth is serially correlated.

For most industrialized nations, 12 lags of quarterly output growth predict only 20 to 40 per cent of the variance of current output growth. Much of this explanatory power seems to come from the first few lags. This suggests that since predictability can be low, the extent of Granger-causality will play an important role in determining how accurately the one-sided univariate filter can estimate the current output gap. For simplicity, the role of Granger-causality will be discussed under the assumption that the past history of output growth is of no use in predicting present and future output growth.

From equation (18), we can see that $\bar{A}(L)$ will tell us as much about the output gap as $\tilde{A}(L)$ when $\tilde{A}^{+}(L) = 0$, which in turn implies that $q_t - y_t$ does not Granger-cause Δy_t. Van Norden (1995) shows that the latter condition in turn implies that q_t does not Granger-cause y_t. If y_t and q_t are cointegrated, this would imply that there is unidirectional causality from y_t to q_t. In other words, exogenous shocks to potential output would have no subsequent effect on actual output, but persistent shocks to actual output would eventually be followed by a similar change in potential. Such behaviour could describe a particularly severe form of hysteresis: one where output has no tendency to return to potential, but potential output is driven in the long run only by previous variations in actual output. In this kind of world, univariate filters can hope to be as effective in estimating the current output gap as they are in estimating past output gaps.

The latter will not be the case when $q_t - y_t$ Granger-causes Δy_t. Therefore, univariate methods will estimate the current output gap less accurately than past output gaps as long as output growth appears to respond in some degree to past changes in the output gap. The intuition behind this result should be clear. Since future output growth will reflect the influence of the current output gap, information about the current gap may be gained by observing future growth.

Univariate methods will be of no use in estimating the current output gap when Δy_t does not Granger-cause $q_t - y_t$.[11] In other words, if output growth that is faster or slower than normal tends to have no

11. Again, this conclusion assumes that past output growth is of no use in predicting future variations in output growth. As mentioned earlier, the data show that there is some serial correlation in output growth, so time-series methods would still have some explanatory power even in this case.

subsequent effect on the size of the output gap, then time-series estimates of the current gap will be as uninformative as possible. The intuition is similar to the explanation offered above. Past output growth is the only information about the gap that we have; if it tells us nothing about the current gap, then our estimates will not be enlightening.

It is more difficult to characterize the kind of economic model that could generate this kind of result, since Granger-causality from Δy_t to $q_t - y_t$ does not directly correspond to any statement about Granger-causality between y_t and q_t.[12] However, it is possible to give examples in which this result would hold. One simple case would be that where

$$\begin{aligned} \Delta y_t &= \alpha \cdot (y_{t-1} - q_{t-1}) + u_t \\ q_t &= q_{t-1} + v_t \end{aligned} . \tag{20}$$

Potential output follows a random walk that is independent of the behaviour of output. Actual output in turn is generated by a simple error-correction model, which ensures that actual and potential output move together in the long run. Such a model precisely satisfies the condition for no Granger-causality from Δy_t to $q_t - y_t$.

Clearly, there is a range of models in which univariate time-series methods will be of little use at the end of sample. Furthermore, it is the short-run dynamics of potential and actual output that are critical to determining whether models belong to this class. This is not an empirically testable question, since we cannot directly observe potential. However, we can try to ensure that our views on the determination of potential are consistent with the methods we use to measure it.

12. See van Norden (1995).

3 EXTENSIONS OF THE HP FILTER

The Bank of Canada has used various extensions of the HP filter to obtain measures of the output gap and to help guide policy. These "hybrid" methods were developed in the 1990s to try to balance the strengths and weaknesses of "structural" and "astructural" approaches to measuring the output gap for policymakers. The key papers explaining the justification for and implementation of this approach are Laxton and Tetlow (1992) and Butler (1996). Work in a similar vein has been pursued both at some of the Federal Reserve Banks (see Kuttner 1994) and at the OECD (see Giorno et al. 1995).

Understanding the contribution of these methods requires an appreciation of the problems these authors were trying to avoid. Laxton and Tetlow argue that there is insufficient knowledge about the true structural determinants of the supply side of the economy to make the purely structural approach practicable. At the same time, for policy purposes we need to distinguish between those movements in output caused by supply shocks and those caused by demand shocks, whereas most astructural (time-series) models attempt to distinguish between permanent and transitory components of output. They suggest as an alternative a way to combine the two approaches that we refer to as the multivariate HP filter.[13]

As is explained in Section 3.1, this methodology consists of adding the residuals of a structural economic relationship to the minimization problem that the HP filter is seeking to solve. Section 3.2 discusses the production function variant of this methodology. In Section 3.3, we examine additional modifications introduced to the filter to improve its performance at the end of the sample. Section 3.4 looks at these approaches from a

13. Laxton and Tetlow call their particular filter the "Multivariate Filter (MVF)" and Butler calls his the "Extended Multivariate Filter (EMVF)." In this paper, we broadly refer to all multivariate extensions of the univariate Hodrick-Prescott filter as multivariate HP filters (MHPF), which include the MVF and EMVF as special cases. The method currently used to estimate Canadian potential output for the Bank's staff projection will also be referred to as the EMVF. The latter differs somewhat from the implementation described in Butler (1996), but is conceptually the same.

different perspective and relates them to both the methods of Section 2 and other methods that use additional structural relationships.

3.1 A multivariate HP filter

As noted previously, the original Hodrick-Prescott filter chooses the trend as the solution to

$$\{y^g{}_t\}_{t=0}^{T+1} = \operatorname{argmin} \sum_{t=1}^{T} (y_t - y_t^g)^2 + \lambda \cdot (\Delta^2 y_{t+1}^g)^2 , \tag{21}$$

where $\Delta^2 y_{t+1}^g = \Delta \cdot \Delta \cdot y_{t+1}$ and $\Delta y_t = y_t - y_{t-1}$. The multivariate HP filter adds a term to the equation:

$$\{y^g{}_t\}_{t=0}^{T+1} = \operatorname{argmin} \sum_{t=1}^{T} (y_t - y_t^g)^2 + \lambda_g \cdot (\Delta^2 y_{t+1}^g)^2 + \lambda_\varepsilon \cdot \varepsilon_t^2 , \tag{22}$$

where $\varepsilon_t = z_t - f(y_t^g, x_t)$. Another economic variable of interest is z_t, and $f(.)$ models z_t as a function of both some explanatory variables x_t and the unobserved trend y_t^g. Since equation (22) includes a new term in ε_t^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 structural economic relationship $f(.)$. The relative weights put on these different objectives are reflected by λ_g and λ_ε.

The key to implementing the multivariate HP filter for the purpose of estimating potential output (or an output gap) is to specify $\{z_t, f(y_t^g, x_t)\}$ so as to capture some structural relationship that depends on either potential output or the output gap. For example, one could specify a Phillips-curve equation that relates observed inflation to a measure of inflation expectations, the output gap, and perhaps additional explanatory variables (such as oil prices). ε_t would then be the residual from this Phillips-curve equation, so the trend of output would be chosen in part to improve the explanatory power of the output gap for inflation. Alternatively, one could use an Okun's law relationship to link the rate of unemployment to the output gap and to various structural variables determining the non-accelerating inflation rate of unemployment (NAIRU). The trend of output would then be influenced by the evolution of the unemployment rate and its structural determinants.

Of course, there is no reason that we have to restrict ourselves to a single structural relationship. Equation (22) can be generalized to include an arbitrary number n of structural relationships with a common trend y_t^g, giving

$$\{y^g{}_t\}_{t=0}^{T+1} = \operatorname{argmin} \sum_{t=1}^{T} (y_t - y_t^g)^2 + \lambda_g \cdot (\Delta^2 y_{t+1}^g)^2 + \left(\sum_{i=1}^{n} \lambda_{\varepsilon_i} \cdot \varepsilon_{it}^2 \right). \qquad (23)$$

The original Laxton and Tetlow (1992) paper used information from both a Phillips-curve and an Okun's law relationship, while Butler (1996) also uses multiple structural relationships simultaneously.

The usefulness of the multivariate HP filter depends on several factors. Obviously, the extent to which it improves upon the original HP filter will depend on the reliability and information content of the structural relationship(s) with which it is combined. These potentially offer a way of mitigating the problems of HP filters noted in Section 2. However, given the importance of obtaining good end-of-sample estimates of output gaps, we require structural relationships that can give good contemporaneous information.[14]

For the particular data-generating process they examine, Laxton and Tetlow find that the degree to which their filter does better than the univariate HP filter at estimating the output gap increases with the relative importance of demand shocks to supply shocks. While the MHPF can produce a large improvement, Laxton and Tetlow find that there is still substantial uncertainty in their point estimates of the output gap and that this uncertainty is larger at the end of sample. In their base case, they find that the 95 per cent confidence interval for the output gap at the end of sample is about 4 per cent on both sides, which implies that policymakers would rarely observe statistically significant output gaps.

Another factor key to the success of the multivariate HP filter is calibration. Instead of having a single λ parameter with a standard value of 1,600, we now have vectors of parameters $\vec{\lambda}_{g,\varepsilon}$ without a clear guide as to

14. In that respect, there may be limitations to the information we can expect to gain from Phillips-curve relationships if we believe that inflation responds to output gaps with a lag.

their appropriate values. In addition, we now also need to estimate the form of the structural relationships involving potential output. If we attempt to do this before calculating $\{y_t^g\}$, then we will be estimating a structural relationship that may be inconsistent with the values of $\{y_t^g\}$ produced by the MHPF. Furthermore, theory will often not be a sufficient guide to allow us to calibrate such a relationship tightly. The approach used by Laxton and Tetlow (1992) and Butler (1996) is to experiment with alternative weightings to see which produce reasonable results and how sensitive the outcomes are to these choices.

An alternative explored by Harvey and Jaeger (1993) and by Côté and Hostland (1994) is to estimate the structural relationship simultaneously with $\{y_t^g\}$ and $\{\lambda_g, \tilde{\lambda}_\varepsilon\}$ via maximum-likelihood methods.[15] Côté and Hostland found that the results can be sensitive to the specification of the structural relationships,[16] that the usefulness of the structural information vanishes when one considers only end-of-sample performance,[17] that the structural parameters cannot be estimated with much accuracy, and that maximization of the likelihood function was problematic.

To give some idea of how such filters perform in practice, Figure 8 compares three different estimates of the output (GDP) gap. The first is that produced by the Butler (1996) filter (labelled EMVF).[18] The others are those produced by a one-sided HP (1,600) filter and by the LRRO filter (described in Section 4). We can see from Figure 8 that the three methods produce gaps of roughly the same amplitude, and that there is a tendency

15. Butler (1996) mentions that a direct maximum-likelihood estimation was attempted but did not produce reasonable results for the λ s.

16. They find that specifying the dynamic relationship in levels or first differences has a large effect on the estimated values of λ_ε.

17. Côté and Hostland approximate the behaviour of the one-sided filter by using only the lags from the mid-sample representation of the HP filter. They also obtain much more useful results when they apply the two-sided filter at the end of sample using forecast values for the required leads in the filter.

18. The details of the filter used for the Bank's staff projection of the Canadian economy evolve over time in response to ongoing research and may therefore, as noted in footnote 13, differ slightly from the exposition in Butler (1996). The EMVF gaps shown in the figure reflect the specification used in the staff's December 1996 economic projection.

for the three series to rise and fall at similar times. While all three series show negative output gaps (i.e., excess supply) in the early 1990s, the LRRO and the HP show the economy returning to potential after a few years, while the EMVF shows large output gaps remaining through to the end of the sample (1996Q3). As we see in the next section, however, this last difference is more a reflection of the differences in the structural information used.

FIGURE 8. Comparison of three different measures of the Canadian output gap

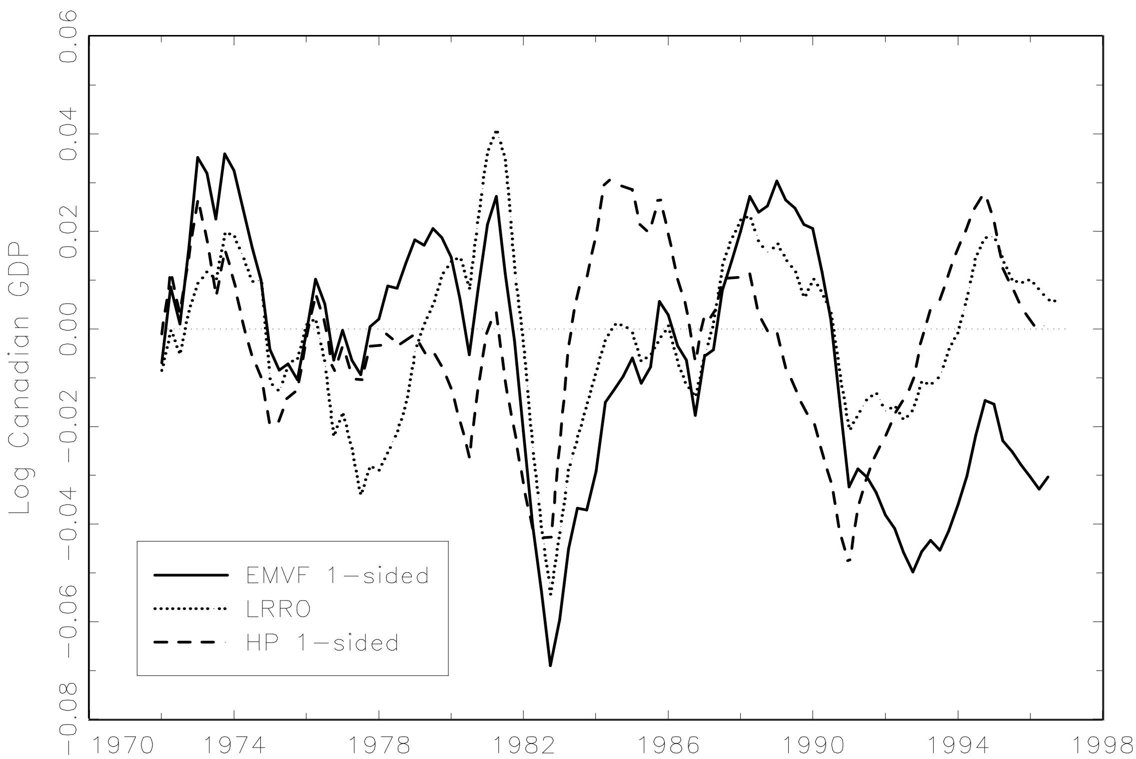

3.2 The production-function approach

Another important feature of the EMVF filter is that rather than filtering output directly, it decomposes output into a number of components that are then individually filtered. This allows both for a more direct link to sources of structural information and for an easier interpretation of the source of changes in the gap or potential.

The decomposition is based on a production function. Consider an aggregate Cobb-Douglas constant-returns-to-scale production function

$$Y = Q \cdot N^{\alpha} \cdot K^{1-\alpha}, \tag{24}$$

where Q is total factor productivity, N is labour, K is the capital stock, and α is the labour-output elasticity (as well as labour's share of income). With some algebra, we can show that

$$\mu \equiv \partial Y / \partial N = \alpha \cdot Y / N \Rightarrow y = n + \mu - \alpha, \tag{25}$$

where lower-case letters are the logs of their upper-case counterparts. This means that to estimate the trend in output, we estimate the trends in employment, the marginal product of labour, and the labour-output elasticity, and then sum them. One nice feature of the decomposition in equation (25) compared with equation (24) is that it avoids the problem of trying to estimate the capital stock reliably. We can then use the further decomposition that the log of total employment n is given by

$$n = Pop + p + \log(1-u), \tag{26}$$

where Pop is the log of the working-age population, p is the log of the participation rate, and u is the rate of unemployment.

Within this framework, the level of potential output is defined as the level of output consistent with existing population, trend rates of unemployment and participation, and trend levels of the marginal product of labour and the output share of labour. In practice, the trend levels of the participation rate and the output share of labour are determined by a combination of judgment, demographic factors, and univariate HP smoothing. Separate MHPF systems are then used to identify the trend level of the marginal product of labour and the trend unemployment rate. The determination of the trend unemployment rate is partly based on the structural method described in Côté and Hostland (1996) and the Phillips curve described in Laxton, Rose, and Tetlow (1993). A long-run relationship between the marginal product of labour and producer wages and a modified Okun's law relationship serve to identify the trend level of the marginal product of labour.

Analysis of the performance of the EMVF in Butler (1996) shows that this method has its own strengths and weaknesses. Butler notes, on the one hand, that the rolling and full-sample estimates of the trend unemployment rate and the equilibrium marginal product of labour are quite similar, and that the labour market gaps are highly correlated with inflation. On the other hand, he also notes that there is significant correlation in the errors across structural equations, suggesting that further efficiency gains may be possible.

To understand the source of the persistent output gap that the EMVF produces in the 1990s, we can decompose the output gap into its three components, as shown by the dotted lines in Figure 9. This shows that the aggregate output gap largely reflects a deviation of the participation rate from its trend level. Note that the participation rate is not filtered and that its trend level is essentially determined by judgment.

However, it would be wrong to attribute the aggregate gap entirely to structural information, as can be seen by comparing the filtered output gap (top-left graph, dotted line) with its unfiltered or “raw” counterpart (same graph, solid line). This shows that the effects of filtering over the most recent period have tended to increase the estimated size of the output gap by 1 to 2 per cent of GDP. The fact that the filtered unemployment rate gap is very close to its unfiltered counterpart implies that most of the difference between the filtered and unfiltered output gap is due to the effects of the filter on the marginal product of labour gap.

FIGURE 9. Decomposition of the EMVF output gap

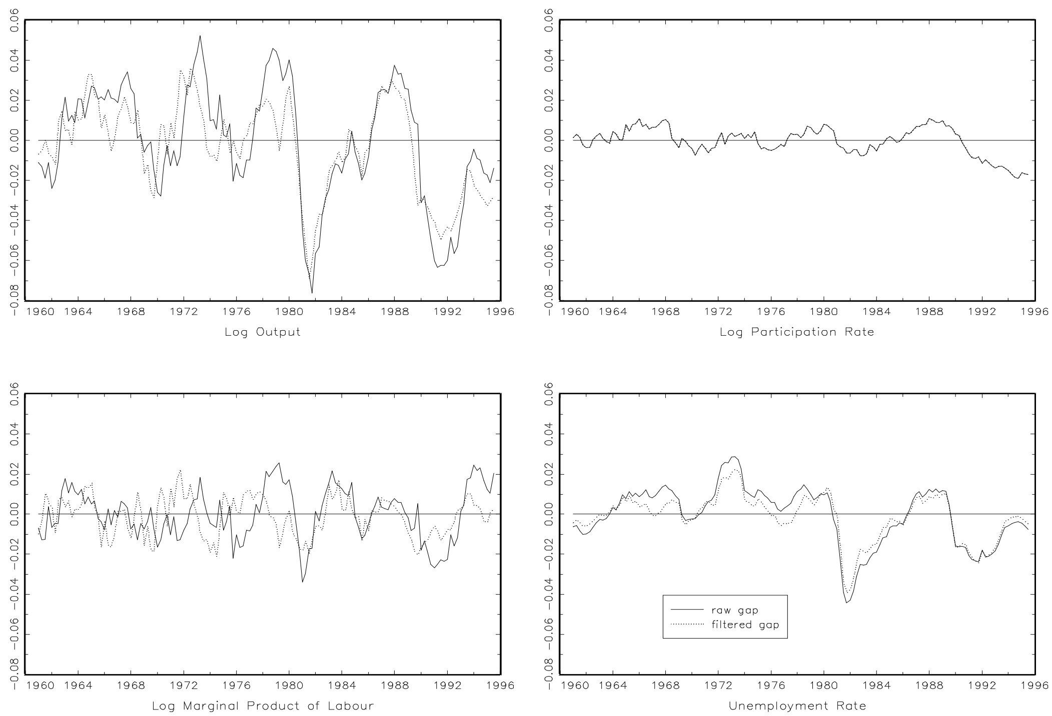

3.3 The end-of-sample problem

While it has been demonstrated that both the filter and the structural information play important roles in the EMVF's estimate of recent output gaps, so far, the end-of-sample problems of HP filters noted in Section 2.4 have not been addressed. However, the EMVF contains two novel features intended to modify its end-of-sample behaviour.

First, the EMVF contains an additional growth-rate restriction. If we temporarily ignore the structural information for expositional simplicity, the modified filter solves the problem

$$\{y^g{}_t\}_{t=0}^{T+1} = \operatorname{argmin}\left(\left(\sum_{t=1}^{T}(y_t - y_t^g)^2 + \lambda_g \cdot (\Delta^2 y_{t+1}^g)^2\right) + \sum_{t=T-j}^{T} \lambda_{ss} \cdot (\Delta\tau_t - \mu_{ss})\right) \tag{27}$$

where μ_{ss} is a constant equal to the steady-state growth rate of potential output and λ_{ss} is the weight put on the growth-rate restriction. The key feature is that λ_{ss} only penalizes deviations from the steady-state growth rate in the last j periods of the sample, effectively "stiffening" the filter. This restriction assumes that the growth rate of potential reverts towards a constant, whereas the theoretical justification for HP filters as optimal filters (noted in Section 2) assumes that this growth rate contains a stochastic trend and therefore will not show any such reversion. Whether such a restriction leads to more accurate estimates of the output gap depends on the accuracy with which the appropriate value of μ_{ss} can be determined.

The second novel feature of the EMVF's treatment of end-of-sample problems is the introduction of a recursive updating restriction. This simply adds an additional term to equation (27), giving

$$\{y^g{}_t\}_{t=0}^{T+1} = \tag{28}$$

$$\operatorname{argmin}\left(\left(\sum_{t=1}^{T}(y_t - y_t^g)^2 + \lambda_g \cdot (\Delta^2 y_{t+1}^g)^2 + \lambda_{pr} \cdot (y_t^g - {}_{pr}y_t{}^g)^2\right) + \sum_{t=T-j}^{T} \lambda_{ss} \cdot (\Delta\tau_t - \mu_{ss})\right)$$

where ${}_{pr}y_t{}^g$ is the tth element of $\{y^g{}_t\}_{t=0}^{T}$. This restricts the filter to choosing $\{y^g{}_t\}$ and maximizes the degree to which new observations modify estimates of y^g based on shorter spans of observations. Not surprisingly, perhaps, Butler (1996) shows that this gives a one-sided estimate of the output gap that behaves more like the subsequent two-sided estimate. While this makes estimates of the output gap behave in a more "orderly" fashion at the end of sample, the net effect on the accuracy of the estimated output gap is unclear.

One way to understand better the effects of these two changes at the end of samples is to compare the resulting one-sided filter with the one- and two-sided HP filters examined in Section 2.[19] As shown in Figure 10, these modifications cause the EMVF to put much less weight on the last few observations of the sample than the one-sided HP filter, and they bring

19. The remainder of this section expands the analysis of the EMVF filter properties presented in Butler (1996) to include the effects of the recursive updating restriction.

its weights more closely in line with those of the two-sided HP filter. If we look in the frequency domain, however, we see that this change causes the one-sided EMVF to pass more of the undesired low-frequency or “trend” components than either of the HP filters. In fact, Figure 11 shows that the squared gain of the filter is greater than 0.2 for all frequencies.

The end result is shown clearly in Figure 12, where both the EMVF output gap and each of its three components appear to be dominated by low-frequency movements not normally associated with business cycles. Compared with Figure 7, the end-of-sample modifications of the EMVF impair the filter’s ability to isolate fluctuations at business-cycle frequencies compared with its simple HP filter counterpart. One way to quantify this effect is to use the estimated spectrum to calculate the correlation of the “measured” EMVF gap with an “ideally filtered” gap that perfectly isolates business-cycle frequencies. The measured EMVF output gap has a correlation of 31.4 per cent with the “ideally filtered” gap, while its two filtered components—the unemployment gap and the labour productivity gap—have correlations of 24.9 per cent and 44.0 per cent, respectively.

The differences in weights between the one-sided HP filter and the one-sided EMVF also imply that since relatively more of its weight comes from observations with greater lags, the EMVF must have a greater phase shift than the HP filter at the end of sample. This in turn implies that the measured EMVF output gap will tend to lag the true output gap by more than the measured HP output gap. The extent of this difference depends on the frequency of the data series, as shown in Figure 13. For all but the lowest of the business-cycle and the sub-business-cycle frequencies, the difference between the two is small, with phase shifts roughly constant at a lag of about two quarters. For lower frequencies, however, the phase shift of HP falls to zero and then becomes negative, while that of the EMVF reaches five quarters by the lower bound of the business-cycle frequencies and increases rapidly thereafter. If we weight these different phase shifts by the relative importance of various frequencies in measured output gaps, we obtain a weighted average measure of the overall phase lag for the different measures. This gives an overall phase lag of roughly 0 for the one-sided HP filter, compared with a lag of 3.3 quarters for the EMVF output gap, 3.8 for the EMVF employment-rate gap, and 2.1 for the EMVF labour-productivity gap.

FIGURE 10. MA Representation of the EMV and HP ($\lambda = 1,600$) filters

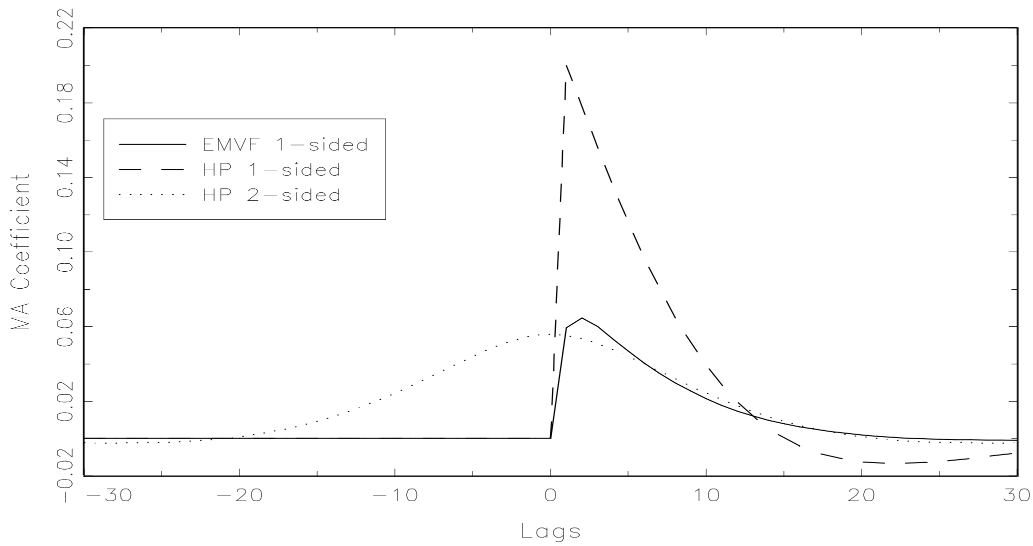

FIGURE 11. Squared gain of the EMV and HP ($\lambda = 1,600$) filters

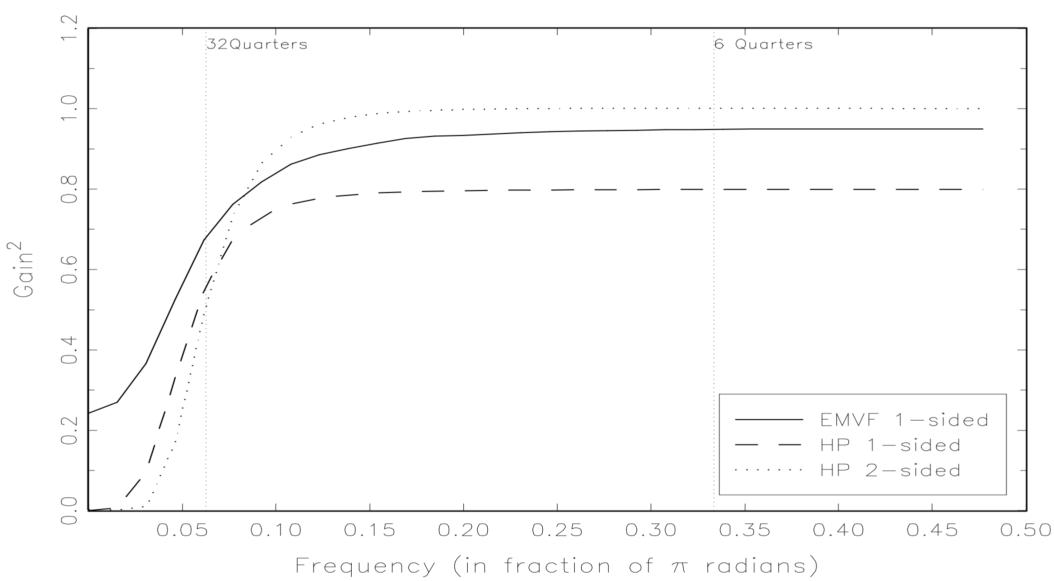

FIGURE 12. Spectrum of the EMVF gap and its components (Ar(3) fit for 1954Q4-1996Q4)

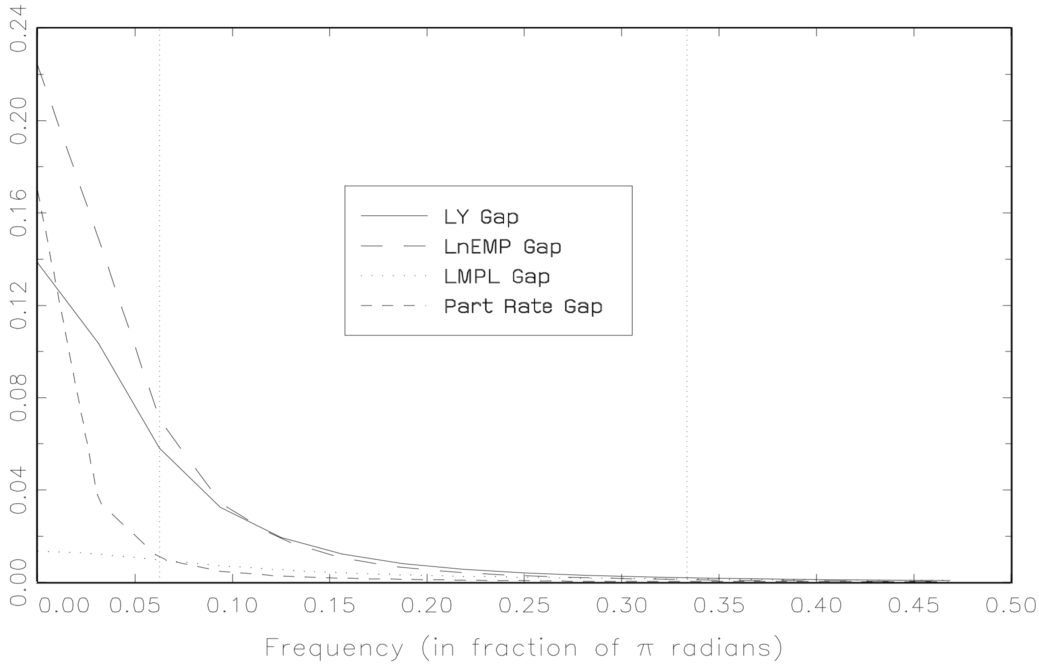

FIGURE 13. Phase shift of the EMV and HP ($\lambda = 1,600$) filters

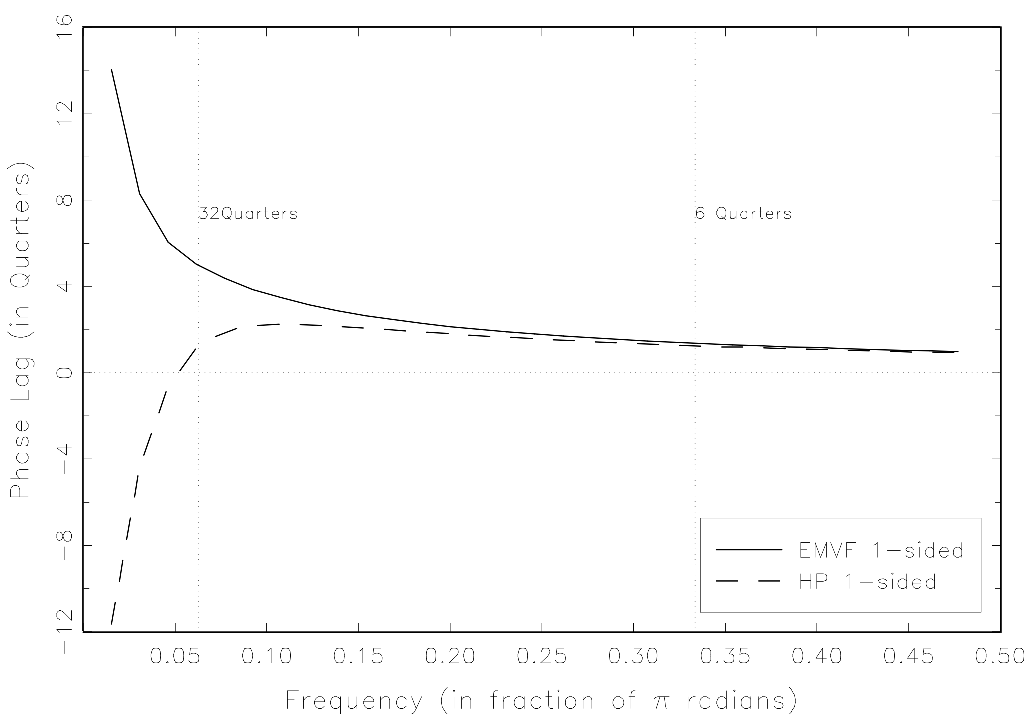

3.4 TOFU

In Section 2, we investigated how well a time-series method (in this case, the HP filter) measured output gaps, and we concluded that by itself it did not produce very reliable estimates for policymakers. So far in Section 3, we have looked at how adding sources of structural information to HP and related filters might improve the situation. One of the strengths of the MHPF approach is that it clearly states the problem that the resulting estimate of the output gap solves. However, some components of that optimization problem may be easier to accept than others. It seems reasonable to us for the estimated output gap to be as consistent as possible with one or more structural economic relationships. The justification for the smoothing portion of the filter is more difficult, as noted in Section 2, and it is unclear how helpful the particular assumptions of the HP filter are in identifying output gaps in the presence of structural information. Furthermore, as the complexity of the filter increases, the question of how to choose the parameters controlling the filter's behaviour becomes more difficult. While it is conceptually straightforward to estimate the filter parameters jointly with the structural relationships (as in Côté and Hostland 1994), this can be quite difficult to implement. Accordingly, it would be useful to restate the optimization problem in a way that allows for easier estimation and easier justification.[20] This leads to an alternative to the HP filter.

From equation (13), we see that we can estimate the output gap in terms of the observable variable Δy_t if we can identify $\tilde{A}(L)$. Presumably, if we know of an economic relationship that involves the output gap, we could use this to define an optimal estimate of $\tilde{A}(L)$, say $\hat{A}(L)$. For example, we might consider a Phillips curve of the form

$$\pi_t = \alpha_0 + \gamma \cdot (q_t - y_t) + B(L) \cdot \pi_{t-1} + C(L) \cdot z_t + e_t \, , \qquad (29)$$

where z_t is a vector of additional observable variables, e_t is an independent and identically distributed mean-zero error term, and $B(L), C(L)$ are one-sided polynomials in non-negative powers of L. We could substitute equation (13) into equation (29) to obtain

20. This section draws heavily on van Norden (1995).

$$\pi_t = \alpha_0 + \gamma \cdot (\tilde{A}(L) \cdot \Delta y_t) + B(L) \cdot \pi_{t-1} + C(L) \cdot z_t + \tilde{e}_t\,, \tag{30}$$

where $\tilde{e}_t = e_t + \gamma \cdot \varepsilon_t$. Equation (30) can now be estimated by conventional methods to obtain optimal estimates of $\tilde{A}(L)$, since it is now specified entirely in terms of observable variables. This would allow us to estimate $\hat{A}(L) \cdot \Delta y_t$ and thereby use equation (13) to estimate the output gap.[21]

We call this estimator of the output gap TOFU: a Trivial Optimal Filter that may be Useful. It is optimal in the sense that estimation by maximum likelihood is straightforward, so our estimates of $\tilde{A}(L)$ will be efficient. The estimator imposes quite general assumptions on the time-series properties of the series involved, so the restrictions should be reasonable. It incorporates a simple structural relationship in order to identify the output gap. Furthermore, if we wish to estimate the output gap at the end of sample, we can simply replace $\tilde{A}(L)$ with $\bar{A}(L)$ (i.e., use only lagged values of Δy_t). This estimator therefore potentially avoids some of the problems mentioned at the start of this section.

Of course, the TOFU estimate of the output gap is obviously related to the estimate obtained by “inverting” the structural relationships. The difference is simply that inversion calculates the implicit value of the output gap that would exactly fit the structural relationship. TOFU therefore lies halfway between such methods and the MHPF methods of Section 3.1. The HP methods are optimal filters only for quite special cases, whereas the TOFU methods provide the optimal linear filter estimate of the output

---

21. We should note two minor caveats. First, equation (30) identifies $\tilde{A}(L) \cdot \Delta y_t$ only up to the scaling factor γ. Strictly speaking, therefore, we only recover an index of the output gap. This should be reasonable for the purpose of, say, deciding whether interest rates should be higher or lower to achieve a given target, since the current value of the index can be readily compared with its historical values. Second, consistent estimation of this relationship requires an implicit assumption. Ordinary-least-squares (OLS) estimation requires $cov(e_t, \Delta y_{t-j}) = 0 \;\; \forall j$ for consistency. If this condition is not satisfied, then instruments for Δy will be required for estimation. Consistent instrumental variables (IV) estimation in turn will depend on the assumption that the chosen instruments are valid. This estimator may be extended in a number of ways. If the output gap were to enter the structural equation in a non-linear fashion, we could estimate the system via general method of moments (GMM) rather than least-squares techniques. If we had a series of structural equations involving the output gap, we could estimate them simultaneously, subject to cross-equation restrictions on the coefficients of Δy_t.

gap. TOFU offers less smoothing than MHPF methods but more than simple inversion of the structural equations.

Unfortunately, the information gained from structural relationships is contaminated by considerable “noise.” Inversion of structural equations is therefore rarely used as a guide to policy, since the resulting estimates of the output gap are usually considered to be too volatile to be of practical use. Whether estimated TOFU filters can reduce this noise enough to be a useful tool for policymakers remains to be seen. If so, they may offer a tractable alternative to the MHPF methods. If not, it suggests that MHPF estimates may be dominated by the arbitrary assumptions they impose on the dynamics of the output gap rather than on the information coming from structural relationships. This suggests that other sources of information on these dynamics should be investigated, which we turn to in the next section.

4 USING LONG-RUN RESTRICTIONS TO ESTIMATE THE OUTPUT GAP

In this section, we discuss approaches based on long-run restrictions imposed on a VAR. These approaches allow the identification of structural shocks and structural components on the basis of a limited number of economic restrictions imposed on an estimated VAR. The chosen restrictions can be those widely agreed upon in the literature. No arbitrary mechanical filter has to be imposed on the data. Compared with methods based on mechanical filters, the methods discussed in this section do not suffer from obvious end-of-sample problems and provide forecasted values of the output gap.

In Section 4.1, we discuss the method based on long-run restrictions imposed on output (LRRO) put forward by Blanchard and Quah (1989), Shapiro and Watson (1988), and King et al. (1991). This method is compared with two VAR-based alternatives: the multivariate Beveridge-Nelson method (MBN) and Cochrane's (1994) method (CO).[22] We argue that one important advantage of the LRRO approach over the MBN and CO approaches is that it allows the diffusion process of shocks to potential output to be estimated. Section 4.2 considers an application of the LRRO methodology to Canadian data. Many of the arguments presented in Sections 4.1 and 4.2 are drawn from Dupasquier, Guay, and St-Amant (1997).

In Section 4.3, we present a method involving restrictions on real output and inflation (LRROI) that yields an output gap corresponding to that part of the cyclical component of real output associated with the trend of inflation. This should be of interest to policymakers concerned about an output gap associated with movements in that trend as opposed to cyclical movements of output unrelated to that trend. This method is discussed in more detail in Lalonde, Page, and St-Amant (forthcoming).

22. See Cogley (1996) for another comparison of the MBN and CO methodologies.

4.1 The LRRO, CO, and MBN methodologies

Let Z_t be an n x 1 stationary vector including an n_1-vector of I(1) variables and an n_2-vector of I(0) variables such that $Z_t = (\Delta X_{1t}', X_{2t}')'$. By the Wold decomposition theorem, Z_t can be expressed as the following reduced form:

$$Z_t = \delta(t) + C(L)\varepsilon_t, \qquad (31)$$

where $\delta(t)$ is deterministic, $C(L) = \Sigma_{i=0}^{\infty} C_i L^i$ is a matrix of polynomial lags, $C_0 = I_n$ is the identity matrix, the vector ε_t is the one-step-ahead forecast errors in Z_t given information on lagged values of Z_t, $E(\varepsilon_t) = 0$, and $E(\varepsilon_t \varepsilon_t') = \Omega$ with Ω positive definite. We assume that the determinantal polynomial $|C(L)|$ has all its roots on or outside the unit circle, which rules out the non-fundamental representations emphasized by Lippi and Reichlin (1993).

Beveridge and Nelson (1981) show that equation (31) can be decomposed into a long-run component and a transitory component:

$$Z_t = \delta(t) + C(1)\varepsilon_t + C^*(L)\varepsilon_t, \qquad (32)$$

with $C(1) = \Sigma_{i=0}^{\infty} C_i$ and $C^*(L) = C(L) - C(1)$. We define $C_1(1)$ as the long-run multiplier of the vector X_{1t}. If the rank of $C_1(1)$ is less than n_1, there exists at least one linear combination of the elements in X_{1t} that is I(0). In other words, there exists at least one cointegration relationship between these variables.

The LRRO approach assumes that Z_t has the following structural representation:

$$Z_t = \delta(t) + \Gamma(L)\eta_t, \qquad (33)$$

where η_t is an n-vector of structural shocks, $E(\eta_t) = 0$, and $E(\eta_t \eta_t') = I_n$ (a simple normalization). We can retrieve the structural form equation (33) from the estimated reduced form by using the following relationships: $\Gamma_0 \Gamma_0' = \Omega$, $\varepsilon_t = \Gamma_0 \eta_t$, and $C(L) = \Gamma(L)\Gamma_0^{-1}$.

The long-run covariance matrix of the reduced form is equal to $C(1)\Omega C(1)'$. From equation (31) and equation (33) we can derive:

$$C(1)\Omega C(1)' = \Gamma(1)\Gamma(1)'. \tag{34}$$

This relation suggests that we can identify matrix Γ_0 with an appropriate number of restrictions on the long-run covariance matrix of the structural form. Blanchard and Quah (1989) and Shapiro and Watson (1988) use long-run restrictions to identify shocks with $C(1)$ having full rank. King et al. (1991) work in a context where the rank of $C(1)$ is less than n_1 and use cointegration restrictions.

Let us assume that the log of real output is the first variable in the vector Z_t. It is then equal to

$$\Delta y_t = \mu_y + \Gamma_1^p(L)\eta_t^p + \Gamma_1^c(L)\eta_t^c, \tag{35}$$

where η_t^p is the vector of permanent shocks affecting output, η_t^c is the vector of shocks having only transitory effects on output, and $\left\{ \Gamma_1^c(L), \Gamma_1^p(L) \right\}$ reflects the dynamic effects of these shocks. Potential output growth based on the LRRO method can then be defined as:

$$\Delta y_t^p = \mu_y + \Gamma_1^p(L)\eta_t^p. \tag{36}$$

Thus, "potential output" corresponds to the permanent component of output. The part of output due to purely transitory shocks is defined as the "output gap."

The MBN decomposition defines potential output as the level of real output that is reached after all transitory dynamics have worked themselves out. With reference to equation (32), where real output is the first element of Z_t, we can write the following decomposition:

$$\Delta y_t = \mu_y + C_1(1)\varepsilon_t + C_1{}^*(L)\varepsilon_t. \tag{37}$$

Potential output can be defined as the first two terms on the right-hand side of equation (37):

$$\Delta y_t^p = \mu_y + C_1(1)\varepsilon_t . \tag{38}$$

It is thus simply a random walk with drift.

Note that the MBN approach gives an output gap that is sensitive to the choice of variables included in the VAR. In general, the more information that is brought into the VAR, the more important the transitory component. This is not the case with the LRRO approach. Adding additional information may or may not add to the importance of the cyclical component.

Cochrane (1994) uses a two-variable VAR including GNP and consumption to identify the permanent and transitory components of GNP. The bivariate representation is augmented with lags of the ratio of consumption to GNP. The permanent income theory implies that consumption is a random walk (for a constant real interest rate). In addition, if we assume that GNP and consumption are cointegrated, then fluctuations in GNP with consumption unchanged must be perceived as transitory. It is on that basis that Cochrane decomposes real GNP into permanent and transitory components. To extract potential output, the errors of the VAR are orthogonalized so that consumption does not respond contemporaneously to GNP shocks.

Cochrane shows that, if GNP and consumption are cointegrated and consumption is a random walk, identification based on the LRRO method and conventional orthogonalization (i.e., a Choleski decomposition) are essentially equivalent. Moreover, if consumption is a pure random walk, Cochrane's decomposition corresponds exactly to the Beveridge-Nelson decomposition based on output and consumption.

To extend our comparison of the LRRO approach with the CO and MBN approaches, let us first write the structural form equation (33) in terms of the log of real GDP (y_t) and the log of real consumption (c_t) decomposed between permanent and transitory shocks (assuming that y_t

and c_t are cointegrated):

$$\Delta y_t = \mu_y + \Gamma_y^p(1)\eta_t^p + \Gamma_y^{p*}(L)\eta_t^p + \Gamma_y^c(L)\eta_t^c \tag{39}$$

$$\Delta c_t = \mu_c + \Gamma_c^p(1)\eta_t^p + \Gamma_c^{p*}(L)\eta_t^p + \Gamma_c^c(L)\eta_t^c , \tag{40}$$

where $\Gamma^p(1)$ is the long-run multiplier of permanent shocks and $\Gamma_y^{p*}(L) = \Gamma^p(L) - \Gamma^p(1)$ is their transitory component. The MBN method considers only the first component of the permanent shocks plus the drift term, i.e., $\mu + \Gamma_y^p(1)\eta_t^p$. The LRRO approach is different in that it also includes the dynamics of permanent shocks to real output ($\Gamma_y^{p*}(L)$) in potential output.

With the CO approach, potential output is constrained to be a random walk to the extent that consumption is a random walk. Indeed, the validity of the permanent-income hypothesis would imply that the last two terms of equation (39) are equal to zero and that $\Gamma_y^p(1) = \Gamma_c^p(1)$. It is not clear what the CO decomposition corresponds to if consumption is not a random walk.[23]

As pointed out by Lippi and Reichlin (1994), modelling the trend in real output as a random walk is inconsistent with most economists' interpretation of productivity growth. Indeed, it is generally believed that technology shocks are absorbed gradually by the economy. Adjustment costs for capital and labour, learning and diffusion processes, habit formation, and time to build are factors that imply richer dynamics than a random walk for these shocks. Again, a crucial advantage of the LRRO approach is that it lets the data determine the shape of the diffusion process of permanent shocks.[24]

---

23. Stochastic growth models—such as in King, Plosser, and Rebelo (1988) or King et al. (1991)—imply that the ratio of the log of GNP to the log of consumption is stationary but that consumption is not a random walk because the real interest rate is not constant. In these models, the transitory component of permanent shocks to consumption is not equal to zero. The LRRO decomposition is compatible with the predictions of these models.

24. Kuttner (1994) proposes a method based on the univariate unobserved stochastic-trend decomposition of Watson (1986) augmented with a Phillips-curve equation. As with the Beveridge-Nelson decomposition, Kuttner's approach constrains potential output to follow a random-walk process.

One implication of defining potential output as a random walk with drift is that when the contemporary effect of a positive permanent shock is smaller (greater) than its long-run effect, the output gap, defined as observed output minus potential, is negative (positive). For example, a positive technological shock whose short-term impact is smaller than its long-term impact will cause a transitory negative output gap. Many researchers and policymakers will find that this feature reduces the attractiveness of the MBN and CO approaches (in the latter case, under the assumption that consumption is a random walk). It will often appear preferable to include the diffusion process associated with permanent shocks in potential output, since the economy is likely to remain on its production possibility frontier as adjustments unfold. There should be no reason for the trend of inflation to change during that adjustment process.

4.2 An application of the LRRO approach to Canadian data

For our applications of the LRRO methodology to Canadian data, we assume that the growth rate of real output (Δy) follows a stationary stochastic process responding to two types of structural shocks: permanent (ε_P) and transitory (ε_T). Also included in the estimated VARs are the first differences of: inflation ($\Delta\pi$), the unemployment rate (Δu), and the real interest rate (Δr). We assume that these series are I(0) and that there is no cointegration involved.[25]

We verified that adding money or the exchange rate to the estimated VARs would have little impact on the results. In selecting the variables, the aim is to include in the VAR the information necessary to identify the structural components of interest. Of course, there is a cost in adding information in terms of lost degrees of freedom and less-precise estimates.

It is important to note that the assumptions made on the level of integration of the series could be changed and that they are not part of the

25. Unit root tests support these assumptions. Results are available on request.

LRRO methodology per se (except for real output, which has to be I(1)).[26] Changing these assumptions could significantly affect the results.

Our objective here is to illustrate the methodology. A practitioner interested in using the approach to estimate the output gap might well choose to incorporate additional information to obtain a better estimate of the output gap.

The structural shocks and the variables used in the VAR can be expressed in the following vector form:

$$\eta_t = \begin{bmatrix} \varepsilon_P \\ \varepsilon_{T1} \\ \varepsilon_{T2} \\ \varepsilon_{T3} \end{bmatrix} \text{ and } Z_t = \begin{bmatrix} \Delta y \\ \Delta \pi \\ \Delta u \\ \Delta r \end{bmatrix} . \tag{41}$$

We use quarterly data on real GDP. Our measure of inflation is the total CPI. The real interest rate is proxied by the overnight rate (see Armour et al. 1996 for a discussion of that series) minus inflation (quarterly growth rate). Our sample extends from the first quarter of 1970 to the fourth quarter of 1996 in order to focus on the flexible exchange rate period. The estimation of the output gap based on the LRRO methodology is fairly robust to the choice of the sample period, however.

The autoregressive reduced form of the model is first estimated as

$$Z_t = \sum_{i=1}^{q} \Pi_i Z_{t-i} + e_t , \tag{42}$$

where q is the number of lags and e_t is a vector of estimated residuals with $E(e_t e_t) = \Sigma$.

It is crucial that the estimated VARs include a sufficient number of lags. Indeed, Monte Carlo simulations by DeSerres and Guay (1995) show that using a lag structure that is too parsimonious can significantly bias the estimation of the structural components in a structural VAR. We decided to

26. Indeed, DeSerres, Guay, and St-Amant (1995) assume that money growth is I(0).

use eight lags. However, we verified that using six or ten lags had little impact on the results.

The LRRO approach involves the identification of structural shocks (ε_t) from reduced-form shocks (e_t) and their variance. For this, we need to provide enough identifying restrictions to evaluate the 16 elements in Γ_0. Given that Σ is symmetric, we need to impose six additional restrictions. The matrix of long-run effects of reduced-form shocks, $C(1)$, is related to the equivalent matrix of structural shocks, $\Gamma(1)$, as follows:

$$\Gamma(1) = C(1)\Gamma_0, \tag{43}$$

where the matrix $C(1)$ is calculated from the estimated VAR. To identify the system we simply impose the condition that $\Gamma(1)$ is triangular, i.e., that three shocks have no long-run effect on real output and two have no long-run effect on inflation. There are then three transitory components of real output that do not need to be identified separately.

FIGURE 14. Response of real GDP to a permanent output shock

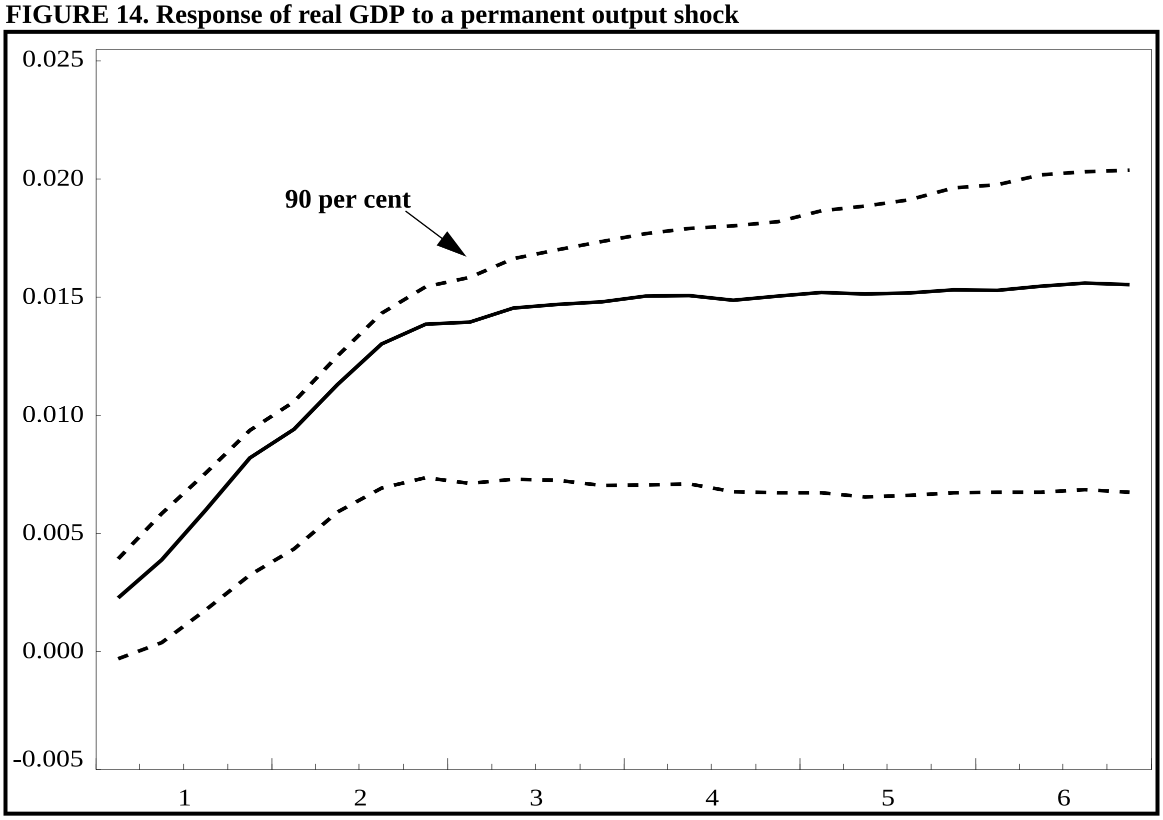

Figure 14 presents the impulse response of real output to a one-standard-deviation permanent output shock. (Blanchard and Quah suggest interpreting this as an aggregate supply shock.) The horizontal axis represents the number of years. Confidence intervals were generated using Monte Carlo simulations in RATS with 1,000 replications.

The most important finding illustrated in Figure 14 is that permanent shocks are characterized by statistically significant dynamics; in other words, potential output has richer dynamics than a simple random walk.[27] As mentioned above, this result could reflect such factors as adjustment costs on capital and labour, learning, habit formation, and time to build. One implication of the rejection of the random-walk assumption is that methods that do not take into account the diffusion process of permanent shocks could miss an important part of potential output. Indeed, Dupasquier, Guay, and St-Amant (1997) show that the correlation is relatively small between the output gaps calculated on the basis of the LRRO, CO, and MBN approaches applied to U.S. data. Part of the reason is the different treatments of the diffusion process of permanent shocks.

Figure 15 shows the output gaps calculated on the basis of the LRRO methodology as applied in this paper, together with 90 and 67 per cent confidence intervals.

Another significant point shown in Figure 15 is that there is a high degree of uncertainty surrounding the estimation of the output gap. Dupasquier, Guay, and St-Amant (1997) reach the same conclusion using U.S. data. Staiger, Stock, and Watson (1996), using a different methodology, also arrive at a similar conclusion concerning the estimation of the NAIRU for the United States. Some of that uncertainty is attributable to the large number of lags that have to be included in the estimated VARs. As mentioned above, DeSerres and Guay (1995) show that many lags have to be used to provide an unbiased decomposition into permanent and transitory components with structural VARs. To a large extent, the purpose of these

27. Dupasquier, Guay, and St-Amant (1997), Blanchard and Quah (1989), and Gali (1992), among others, report similar results for the U.S. economy.

lags is to approximate the moving-average part of the underlying DGP. Preliminary results obtained at the Bank of Canada suggest that the estimation of VARMAs instead of VARs could reduce parameter uncertainty by allowing the use of more parsimonious models.

FIGURE 15. LRRO-based Canadian output gap

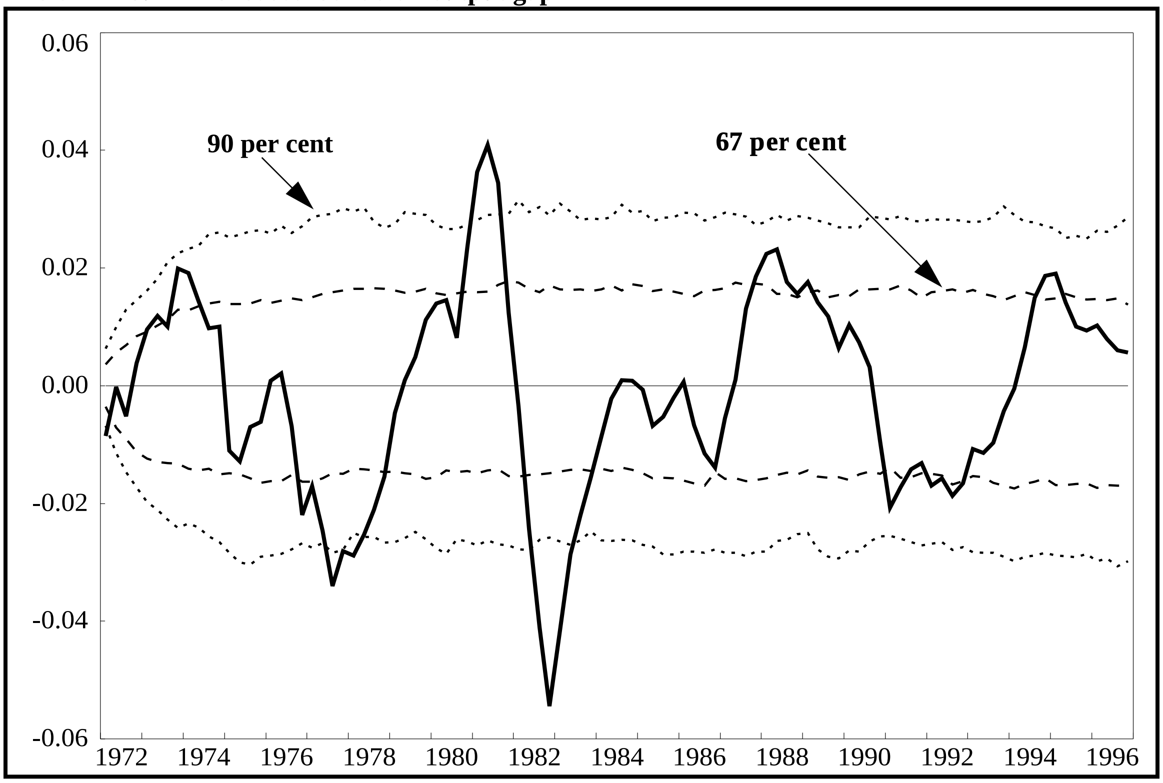

Still, there are episodes of significant output gaps at either the 90 per cent or the 67 per cent levels. These output gaps appear reasonable, in that positive output gaps are associated with episodes of accelerating inflation, while negative output gaps correspond to episodes of decelerating inflation.

It would be of interest to see whether the structural shocks used to calculate the output gap account for an important part of the variance of inflation at different horizons. That information is presented in Table 2.

TABLE 2: Variance decomposition of Canada's inflation rate (LRRO method) (relative contribution of the different types of shocks, per cent)

Horizon (quarters)	Permanent output shock	Transitory output shock
1	11 (0-43)[a]	89 (57-100)
4	8 (2-36)	92 (64-98)
8	6 (4-30)	94 (69-97)
16	13 (5-49)	87 (51-95)
32	23 (5-66)	77 (34-95)
Long-term	37 (3-84)	63 (16-97)

a. 90 per cent confidence interval

We can see that transitory shocks affecting real output account for a large part of the variance of Canadian inflation. This result suggests that the component of output that we identify includes much of the information that one might want to include in the output gap. Note, however, that there appears to be some additional information related to permanent shocks, especially at longer horizons. Lalonde, Page, and St-Amant (forthcoming) discuss another approach that does not require that the output gap to be part of the cyclical component of output.

In this paper the focus remains on methods that impose a stationary output gap. Although Table 2 suggests that the output gap based on the LRRO approach accounts for a large fraction of fluctuations in the trend of inflation, a part of that gap may very well be unrelated to the trend. For example, it might include very-high-frequency cycles that have little to do with that trend. This suggests another method: imposing restrictions on both real output *and* inflation in order to produce an output gap that is constrained to be associated with movements in the trend of inflation.

4.3 The LRROI approach

Using the four-variable system presented in Section 4.2, it is possible to identify shocks with no long-run effects on real output but with an effect on the trend of inflation. Real output was decomposed in the following way:

$$\Delta y_t = \mu_y + \Gamma_y^p(1)\eta_t^p + \Gamma_y^{p*}(L)\eta_t^p + \Gamma_y^{cp}(L)\eta_t^{cp} + \Gamma_y^{cc}(L)\eta_t^{cc}, \tag{44}$$

with certain shocks having no long-run effect on real output but an effect on the trend of inflation (η_t^{cp}), certain shocks having no long-run effect on real output *or* inflation (η_t^{cc}),[28] and certain shocks having long-run effects on real output but an effect on inflation that is left unconstrained (η_t^p). The component $\Gamma_y^{cp}(L)\eta_t^{cp}$ could be used as a measure of the output gap. This is what we call the LRROI method. In Section 4.2, we were not interested in distinguishing $\Gamma_y^{cp}(L)\eta_t^{cp}$ from $\Gamma_y^{cc}(L)\eta_t^{cc}$, and so we simply added these components to form one component that was the LRRO output gap.

The LRROI method gives a measure of the output gap that is more constrained than the LRRO method, in that it combines restrictions on real output *and* inflation. With the LRRO approach, it is only necessary to assume that real output is I(1), but a necessary feature of the LRROI approach is the additional assumption that inflation is better characterized as being I(1) over the sample under consideration. Of course, that assumption is not uncontroversial. However, we think it is reasonable. Assuming that inflation is I(0), in contrast, is in some sense equivalent to assuming that inflation has to return to a constant mean, whatever the actions of the monetary authorities. We do not think that this is realistic. The mean of inflation can vary with factors such as the preferences of the monetary authorities, the political environment confronting these authorities, and the state of knowledge about the costs and benefits of targeting a certain inflation rate. Inflation has to be modelled as a process whose mean can change over time. One way to do so is to assume that it is I(1).

---

28. Since there are two shocks that have no long-run impact on the trend of output or inflation (the last two of equation 44), we simply consider their sum.

The LRROI method proposes a measure of the output gap that should be attractive for policymakers interested in that part of the cyclical component of real output that is associated with movements in the trend of inflation as opposed to short-run fluctuations of that series. On the other hand, the measure of the output gap provided by the LRRO method, which includes cyclical movements, may not be of interest for such policymakers. In some sense, the LRROI method could thus provide policymakers with a less "noisy" indicator of changes in the trend of inflation. Short-run fluctuations of inflation could be caused by factors such as transitory fluctuations in the exchange rate or changes in indirect taxes. Policymakers may not be interested in changing their policy stance in response to such movements. Of course, there might be instances when policymakers will be interested in reacting to fluctuations of output associated with transitory changes of inflation if these are expected to last for a long enough period. The LRROI gap would not include the effect of such shocks but the LRRO gap would.

Figure 16 presents the LRROI output gap calculated on the basis of the Canadian data used in Section 4.3. Figure 17 compares that output gap with the gap resulting from the LRRO approach and the gap from the one-sided HP filter. Table 3 presents the variance decomposition of inflation corresponding to the LRROI decomposition.

Figure 16 shows that there is as much uncertainty surrounding the estimation of the LRROI output gap as there is with the LRRO gap. We also see that the LRROI gaps are generally smaller than those produced by the LRRO method. This indicates that some cycles included in the LRRO gap are not related to the trend of inflation.

The variance decomposition of inflation presented in Table 3 suggests that the LRROI approach includes most of the information related to the cyclical component of output that is relevant for monitoring or forecasting medium- and long-term inflation. However, the LRRO gap appears to include non-negligible information about short-term movements in inflation. Table 4 shows the variance decomposition of output associated with the various structural shocks.

To conclude this section, we would like to emphasize that the LRRO and LRROI methods are two variants of a more general VAR-based approach. Other variants could be considered. For example, one might want to take into account possible changes in the inflation process due to changes in the monetary regime (see, for example, Fillion and Léonard 1997). Evans (1992) proposes a time-varying method that could allow for such non-linearities. One might also want to include different sets of restrictions, including cointegration and Bayesian types of restrictions, to ensure that the estimated gap is compatible with a particular macroeconomic model.

FIGURE 16. LRROI-based Canadian output gap

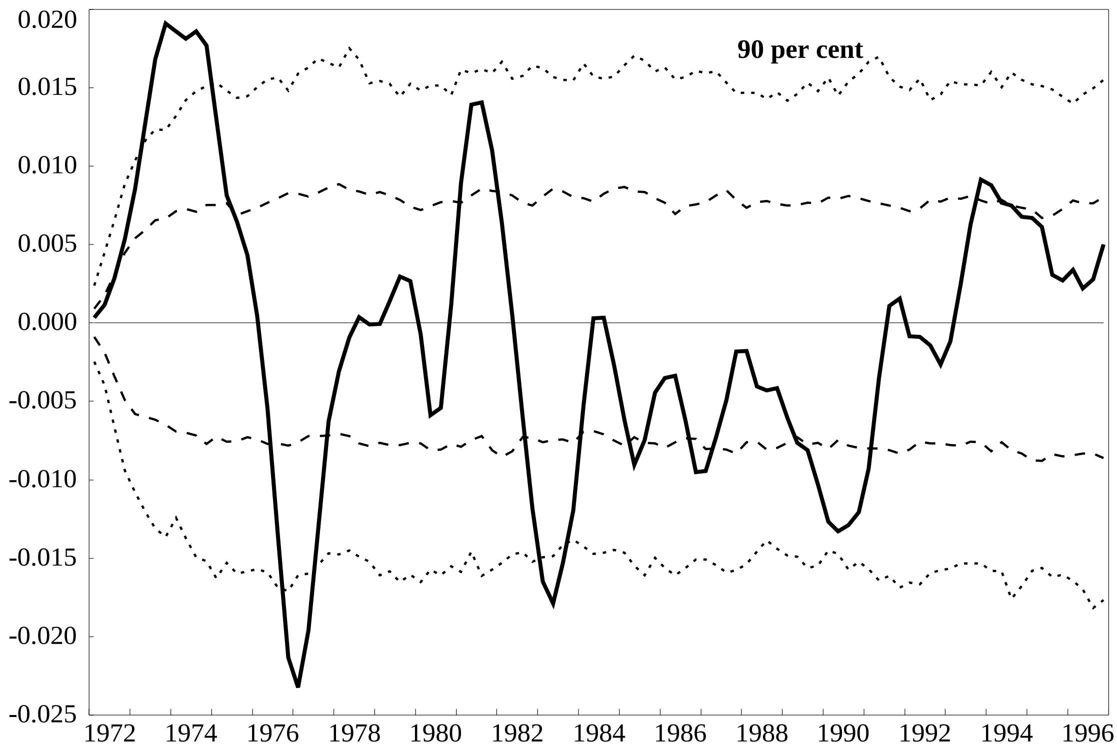

FIGURE 17. LRRO, LRROI, and one-sided HP-based Canadian output gaps

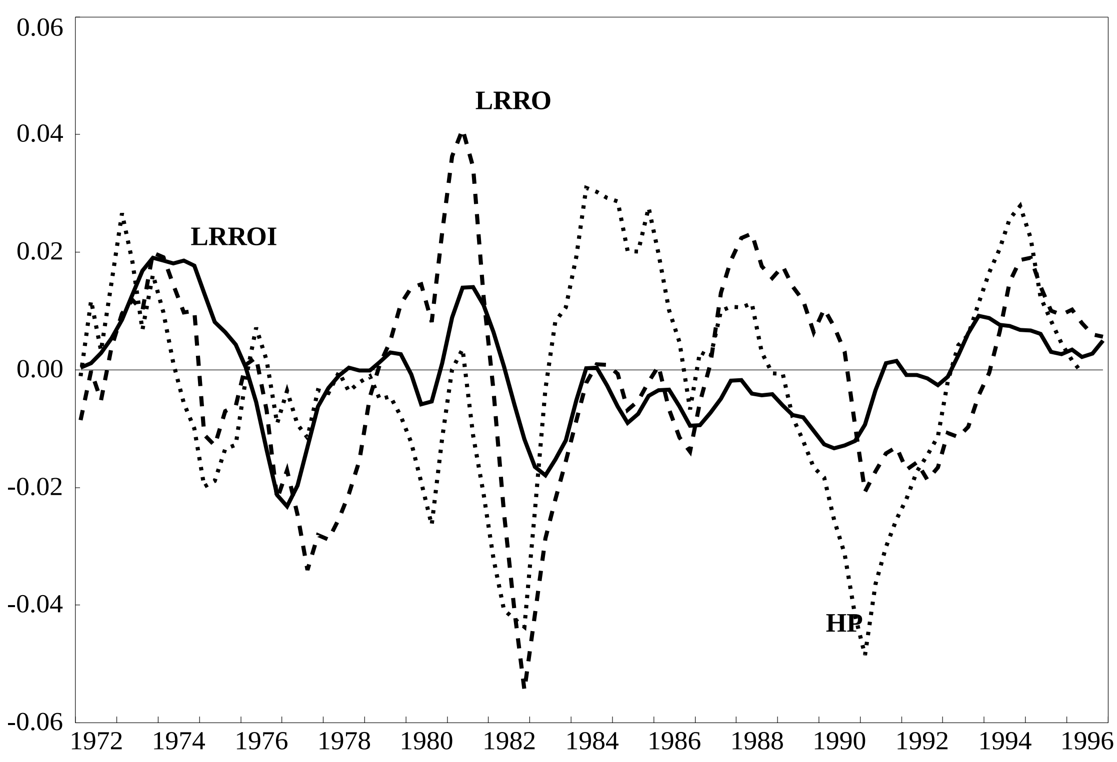

TABLE 3: Variance decomposition of Canada's inflation rate (relative contribution of the different types of shocks, in per cent)

Horizon (quarters)	Permanent output shock	Transitory output shocks affecting the trend of inflation	Transitory output shocks not affecting the trend of inflation
1	11 (0-43)[a]	54 (6-88)	35 (3-81)
4	8 (2-36)	63 (18-85)	29 (7-68)
8	6 (4-30)	70 (33-80)	24 (11-53)
16	13 (5-49)	72 (34-80)	15 (8-35)
32	23 (5-66)	68 (27-83)	9 (4-23)
Long-term	37 (3-84)	61 (16-95)	1 (0-4)

a. 90 per cent confidence interval

TABLE 4: Variance decomposition of Canada's real GDP (relative contribution of the different types of shocks, in per cent)

Horizon (quarters)	Permanent output shock	Transitory output shocks affecting the trend of inflation	Transitory output shocks not affecting the trend of inflation
1	12 (1-72)[a]	6 (0-25)	83 (23-94
4	44 (5-73)	12 (1-44)	45 (17-82)
8	78 (21-86)	5 (1-40)	17 (8-58)
16	92 (51-94)	2 (1-20)	6 (4-33)
32	97 (75-98)	1 (0-10)	2 (2-16)
Long-term	100 (97-100)	0 (0-1)	0 (0-2)

a. 90 per cent confidence interval

5 SUMMARY

Considerable research has been carried out on methods for measuring output gaps, and much remains to be done. In particular, future research might examine the economic reasonableness of the results obtained on the basis of different methodologies, something that we have not undertaken here. On the basis of our research, we think that there appear to be three main lessons.

i. Univariate time-series methods like the HP filter are not a reliable way to measure the output gap.

Univariate methods rely on an arbitrary decomposition of a series into a trend and a cyclical component. Changing the decomposition method, however, can significantly affect the measured output gap, and economic theory usually has little or nothing to say about which method should be favoured. In addition, the causal relationship between potential and the gap can limit the information obtained about the current gap, and some popular economic models imply that univariate filters will never be able to provide much information.

In the particular case of the HP filter, even though this filter is thought to be close to an ideal high-pass filter, it does not accurately measure the components from business-cycle frequencies when those series have the typical Granger shape. Although the HP filter can also be justified as an optimal filter for particular cases, these cases do not appear to be realistic approximations of output, and the filter is generally not a reliable way to estimate the “cyclical” component. HP filters also behave very differently at the end of the sample, which is the period policymakers care most about, and little is known about the trade-off between phase shift and smoothness at the end of sample.

ii. Existing hybrid methods that combine univariate dynamic methods and structural relationships are not a panacea.

In practical terms, existing hybrid methods have three problems: they have proved hard to estimate; they may not be robust to alternative reasonable calibrations; and it is difficult to calculate their appropriate confidence intervals. The ability of the EMVF, the filter used to estimate the output gap in the Bank of Canada's staff projections of the Canadian economy, to isolate business-cycle frequencies is worse than that of the HP filter, and its estimates lag the true output gap by just under a year. These problems seem to be the result of features introduced to improve the filter's estimate of current and recent output gaps.

More generally, the "hybrid" approach is driven not only by a desire to include structural relationships, but also by a pragmatic desire for "smooth" estimates of the output gap. If the structural relationships are very informative, however, then the "smoothing" assumptions may be unnecessary. Otherwise, it is hard to argue that these assumptions are innocuous, for the reasons mentioned above in the context of univariate methods. For that reason, and because they can incorporate the same sources of structural information, TOFU (Trivial Optimal Filter that may be Useful) methods may provide a good benchmark for hybrid methods. If TOFU estimates of the output gap are useful, then the strict (or ad hoc) filters used in existing hybrid methods are not required. If TOFU methods are not useful, then structural information alone is not sufficient to identify the output gap. Relying instead on ad hoc dynamic assumptions for identification raises questions about the reliability and significance of the resulting estimated output gaps.

iii. Methods that combine estimated dynamics with structural information offer an interesting alternative that merits further investigation.

This report has explored such methods in the form of VARs with long-run restrictions suggested by economic theory. Their advantages include an absence of arbitrary dynamic assumptions, straightforward estimation, and an ability to estimate both current and expected future output gaps. On the other hand, it is not always clear which variables have to be included in the VAR, and work to date suggests that the estimated output gaps have wide confidence intervals, comparable to those of other methods. More work is needed to evaluate the extension of these methods to VARMA models. It might also be of interest to explore VARs with time-varying parameters (as proposed by Evans 1992) in order to deal with possible non-linearities in structural or dynamic relationships.

The VAR-based methods considered in this report impose relatively little economic structure on the data and allow the dynamic properties of the estimated output gaps to be data-determined. This is an advantage, because economic theory provides little guidance on what the dynamics should be. Of course, it must be assumed that the estimated VAR includes the information relevant for the identification of the output gap.

On the other hand, there could be instances when practitioners want to impose more economic structure on their estimation of the output gap by incorporating more economic relationships. To some extent, this could be accommodated within the VAR (or VARMA or VECM) framework. Many types of economic or statistical restrictions have been discussed in the literature, including Bayesian priors and cointegration relationships. Such restrictions could be used to ensure that the estimated output gap is broadly compatible with some structural model of interest.

However, the TOFU approach might provide a better framework for those interested in imposing detailed economic relationships on the data (such as a specific Phillips curve or a NAIRU equation). One advantage of using detailed economic structure is that the derived measures of potential output can then be embedded in a model that is consistent with that structure. As noted briefly in van Norden (1995), extending the TOFU approach to multiple detailed structural relationships appears to be straightforward, although this question has yet to be explored in detail. It should be kept in mind, however, that the usefulness of TOFU output gaps depends on the validity of the economic structure imposed on the data and that economics provides very few non-controversial structural relationships.

Our findings suggest that the VAR-based and TOFU approaches deserve further research and that univariate and arbitrary smoothing methods should be avoided wherever practicable. Variants of the VAR-based and TOFU approaches also need to be investigated and their usefulness in terms of monitoring and projecting inflation evaluated.

REFERENCES

Armour, J., W. Engert and B. S. C. Fung. 1996. "Overnight Rate Innovations as a Measure of Monetary Policy Shocks in Vector Autoregressions." Working Paper 96-4. Bank of Canada, Ottawa.

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, Cambridge, MA.

Beveridge, S. and C. R. Nelson. 1981. "A New Approach to Decomposition of Economic Time Series into Permanent and Transitory Components with Particular Attention to Measurement of the 'Business Cycle'." *Journal of Monetary Economics* 7(2): 151–74.

Blanchard, O. J. and S. Fischer. 1989. *Lectures on Macroeconomics, 1989.* Cambridge, MA: MIT Press.

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

Boschen, J. and L. Mills. 1990. "Monetary Policy with a New View of Potential GNP." *Federal Reserve Bank of Philadelphia Business Review* (July-August): 3–10.

Burns, A. F. and W. C. Mitchell. 1946. *Measuring Business Cycles.* New York: National Bureau of Economic Research.

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.

Cochrane, J. H. 1994. "Permanent and Transitory Components of GNP and Stock Prices." *Quarterly Journal of Economics* 109(1): 241–65.

Cogley, T. 1996. “Evaluating Non-Structural Measures of the Business Cycle.” Draft Paper. Federal Reserve Bank of San Francisco.

Cogley, T. and J. Nason. 1995. “Effects of the Hodrick-Prescott Filter on Trend and Difference Stationary Time Series: Implications for Business Cycle Research.” *Journal of Economic Dynamics and Control* 19(1–2): 253–78.

Côté, D. and D. Hostland. 1994. “Measuring Potential Output and the NAIRU as Unobserved Variables in a Systems Framework.” In *Economic Behaviour and Policy Choice Under Price Stability,* proceedings of a conference held at the Bank of Canada, October 1993, 357–418. Ottawa: Bank of Canada.

————. 1996. “An Econometric Examination of the Trend Unemployment Rate in Canada.” Working Paper 96-7. Bank of Canada, Ottawa.

DeSerres, A. and A. Guay. 1995. “The Selection of the Truncation Lag in Structural VARs (or VECMs) with Long-Run Restrictions.” Working Paper 95-9. Bank of Canada, Ottawa.

DeSerres, A., A. Guay and P. St-Amant. 1995. “Estimating and Projecting Potential Output Using Structural VAR Methodology: The Case of the Mexican Economy.” Working Paper 95-2. Bank of Canada, Ottawa.

Dupasquier C., A. Guay and P. St-Amant. 1997. “A Comparison of Alternative Methodologies for Estimating Potential Output and the Output Gap.” Working Paper 97-5. Bank of Canada, Ottawa.

Dupasquier. C., R. Lalonde and P. St-Amant. 1997. "Les zones monétaires optimales : une application au Canada et aux États-Unis." Forthcoming.

Engle, R. F. and C. W. J. Granger. 1987. “Co-integration and Error Correction: Representation, Estimation, and Testing.” *Econometrica* 55(2): 251–76.

Evans, G. and L. Reichlin. 1994. “Information, Forecasts, and Measurement of the Business Cycle.” *Journal of Monetary Economics* 33(2): 233–54.

Evans, M. D. D. 1992. "Changing Nature of the Output-Inflation Trade-off." Working Paper 92-17. Leonard Stern School of Business, NYU.

Gali, J. 1992. “How Well Does the IS-LM Model Fit Postwar U.S. Data?” *Quarterly Journal of Economics* 107(2): pp. 709–38.

Giorno, C., P. Richardson, D. Roseveare and Paul van den Noord. 1995. “Estimating Potential Output, Output Gaps and Structural Budget Balances.” Working Paper 152. OCED Economic Department, Paris.

Gouriéroux, C. and A. Monfort. 1995. *Séries Temporelles et Modèles Dynamiques*. 2nd ed. Paris: Economica.

Guay, A. and P. St-Amant. 1996. *Do Mechanical Filters Provide a Good Approximation of Business Cycles?* Technical Report No. 78. Ottawa: Bank of Canada.

Hamilton, J. D. 1994. *Time Series Analysis.* Princeton: Princeton University Press.

Harvey, A. C. and A. Jaeger. 1993. “Detrending, Stylized Facts and the Business Cycle.” *Journal of Applied Econometrics* 8(3): 231–47.

Henderson, R. 1924. “A New Method of Graduation.” *Actuarial Society of America Transactions* 25: 29–40.

Hendry, S. 1995. "Long-Run Demand for M1." Working Paper 95-11. Bank of Canada, Ottawa.

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

Jeong, J. and G. S. Maddala. 1993. “A Perspective on Application of Bootstrap Methods in Econometrics.” In *Handbook of Statistics 11: Econometrics,* edited by G. S. Maddala, C. R. Rao and H. D. Vinod, 573–610. New York: North-Holland.

King, R. G., C. I. Plosser and S. T. Rebelo. 1988. “Production, growth and business cycles: II. New directions.” *Journal of Monetary Economics* 21(2–3): 309–41.

King, R. G., G. I. Plosser, J. H. Stock, and M. W. Watson. 1991. “Stochastic Trends and Economic Fluctuations.” *American Economic Review* 81(4): 819–40.

King, R. G. and S. Rebelo. 1993. “Low Frequency Filtering and Real Business-Cycles.” *Journal of Economic Dynamics and Control* 17(1–2): 207–31.

Kuttner, K. N. 1994. “Estimating Potential Output as a Latent Variable.” *Journal of Business and Economic Statistics* 12(3): 361–68.

Lalonde, R. J. Page and P. St-Amant. "Trois méthodes d’estimation de l’écart de production reposant sur des restrictions de long terme." Forthcoming.

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 R. Tetlow. 1993. “Is the Canadian Phillips curve non-linear?” Working Paper 93-7. Bank of Canada, Ottawa.

Lippi, M. and L. Reichlin. 1993. “The Dynamic Effects of Aggregate Demand and Supply Disturbances: Comment.” *American Economic Review* 83(3): 644–52.

———. 1994. “Diffusion of Technical Change and the Decomposition of Output into Trend and Cycle.” *Review of Economic Studies* 61(1): 19–30.

Nelson, C. R. and C. I. Plosser. 1982. “Trends and Random Walks in Macroeconomic Time Series: Some Evidence and Implications.” *Journal of Monetary Economics* 10(2): 139–62.

Perron, P. 1989. “The Great Crash, the Oil Price Shock, and the Unit Root Hypothesis. *Econometrica* 57(6): 1361–1401.

Priestley, M. B. 1981. *Spectral Analysis and Time Series.* Toronto: Academic Press.

Quah, D. 1992. “The Relative Importance of Permanent and Transitory Components: Identification and Some Theoretical Bounds.” *Econometrica* 60(1): 107–18.

Reinsel, G. C. and S. H. Ahn. 1988. Asymptotic Distribution of the Likelihood Ratio Test for Cointegration in the Nonstationary VAR Model. Technical Report. Madison, WI: University of Wisconsin.

Runkle, D. E. 1987. “Vector Autoregressions and Reality.” *Journal of Business and Economic Statistics* 5(4): 437–42, 454.

Sargent, T. J. 1987. *Macroeconomic Theory.* Boston: Academic Press.

Schwert, G. W. 1987. "Effects of model specification on tests for unit roots in macroeconomic data." *Journal of Monetary Economics* 20(1): 73–103.

Shapiro, M. D. and M. W. Watson. 1988. “Sources of Business Cycle Fluctuations.” Working Paper 2589. National Bureau of Economic Research, Cambridge, MA.

Singleton, K. 1988. “Econometric Issues in the Analysis of Equilibrium Business Cycle Models.” *Journal of Monetary Economics* 21(2–3): 361–86.

Staiger, D., J. H. Stock, and M. W. Watson. 1996. “How Precise are Estimates of the Natural Rate of Unemployment?” Working Paper 5477. National Bureau of Economic Research, Cambridge, MA.

van Norden, S. 1995. “Why Is It So Hard to Measure the Current Output Gap?” Bank of Canada. Mimeo.

Watson, M. W. 1986. “Univariate Detrending Methods with Stochastic Trends.” *Journal of Monetary Economics* 18(1): 49–75.

———. 1993. “Vector autoregressions and cointegration.” In *Handbook of Econometrics*, vol. 4, edited by R. F. Engle and D. L. McFadden, 2843–2915. New York: North-Holland.

Weber, C. E. 1995. “Cyclical Output, Cyclical Unemployment, and Okun’s Coefficient: a New Approach.” *Journal of Applied Econometrics* 10(4): 433–45.

Whittaker, E. T. 1923. “On a New Method of Graduation.” Proceedings of the Edinburgh Mathematical Society 41: 63–75.

Bank of Canada Technical Reports

1997

79	Measurement of the Output Gap: A Discussion of Recent Research at the Bank of Canada	P. St-Amant, S. van Norden

1996

78	Do Mechanical Filters Provide a Good Approximation of Business Cycles?	A. Guay and P. St-Amant
77	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	L. Butler
76	Excess Volatility and Speculative Bubbles in the Canadian Dollar: Real or Imagined?	J. Murray, S. van Norden and R. Vigfusson
75	The Dynamic Model: QPM, The Bank of Canada's New Quarterly Projection Model, Part 3	D. Coletti, B. Hunt, D. Rose and R. Tetlow
74	The Electronic Purse: An Overview of Recent Developments and Policy Issues	G. Stuber

1995

73	A Robust Method for Simulating Forward-Looking Models, The Bank of Canada's New Quarterly Projection Model, Part 2	J. Armstrong, R. Black, D. Laxton and D. Rose

1994

72	The Steady-State Model: SSQPM, The Bank of Canada's New Quarterly Projection Model, Part 1	R. Black, D. Laxton, D. Rose and R. Tetlow
71	Wealth, Disposable Income and Consumption: Some Evidence for Canada	R. T. Macklem
70	The Implications of the FTA and NAFTA for Canada and Mexico	W. R. White
69	From Monetary Policy Instruments to Administered Interest Rates: The Transmission Mechanism in Canada	K. Clinton and D. Howard
68	The Microstructure of Financial Derivatives Markets: Exchange-Traded versus Over-the-Counter	B. González-Hermosillo
67	The Role of House Prices in Regional Inflation Disparities	D. Maclean
66	Les sources des fluctuations des taux de change en Europe et leurs implications pour l'union monétaire	A. DeSerres and R. Lalonde

Copies of the above titles and a complete list of Bank of Canada technical reports are available from

PUBLICATIONS DISTRIBUTION, Bank of Canada, Ottawa, Ontario, Canada K1A 0G9
Tel: (613) 782-8248; fax: (613) 782-8874