

A boosted HP filter for New Zealand business cycle analysis?

Viv B. Hall^{1,2} and Peter Thomson^{3,1}

¹ *School of Economics and Finance, Victoria University of Wellington, New Zealand*

² *Centre for Applied Macroeconomic Analysis, ANU, Canberra, Australia*

³ *Statistics Research Associates Ltd, Wellington, New Zealand*

25 May 2022

Abstract

We investigate whether the boosted HP filter (bHP) proposed by Phillips and Shi (2021) might be preferred for New Zealand trend and growth cycle analysis, relative to using the standard HP filter (HP1600). We do this for a representative range of quarterly macroeconomic time series typically used in small theoretical and empirical macroeconomic models, and address the following questions.

Tradition dictates that business cycle periodicities lie between 6 and 32 quarters (e.g. Baxter and King, 1999) (BK). In the context of more recent business cycle durations, should periodicities up to 40 quarters or more now be considered?

Phillips and Shi (2021) propose two stopping rules for selecting a bHP trend. Does it matter which is applied? For growth cycle analysis, we propose other trend selection criteria based on the cut-off frequency and sharpness of the trend filter.

Are stylised business cycle facts from bHP filtering materially different to those produced from HP1600? In particular, does bHP filtering lead to New Zealand growth cycles which are noticeably different from those associated with HP1600 or BK filtering?

HP1600 is commonly used as an omnibus filter across all key macroeconomic variables. Does the greater flexibility of bHP filtering provide better alternatives?

We conclude that the 6 to 32 quarter business cycle periodicity is sufficient to reflect New Zealand growth cycles and determine stylised business cycle facts and, for our representative 13-variable macroeconomic data set, using a bHP filter (2HP1600) as an omnibus filter is preferable to using the HP1600 filter.

JEL Classification: E32, E37, C10, G01

Keywords: Boosting; Hodrick-Prescott filter; Business cycles; New Zealand

Acknowledgements:

Corresponding author (email): Viv B. Hall (viv.hall@vuw.ac.nz)

Declarations:

Funding – None

Conflicts of interest/Competing interests – None

Availability of data and material – Available on request from corresponding author

1 Introduction

Globally and in New Zealand, there has been a lengthy tradition of using the standard Hodrick-Prescott filter for business cycle analysis. See, for example, Hodrick and Prescott (1997) (HP)¹, Canova (1994, 1998), Pagan (1997), Harding and Pagan (2016), Hodrick (2020) and, for New Zealand, Kim, Buckle and Hall (1994, 1995), Hall, Kim and Buckle (1998), McCaw (2007), Hall and McDermott (2016), Hall, Thomson and McKelvie (2017), Lienert (2018), and Hall and Thomson (2021).

There have also been a significant range of contributions advocating alternative filters and providing critiques of the HP filter. These include Beveridge and Nelson (1981), Nelson and Kang (1981), Nelson and Plosser (1982), Harvey and Jaeger (1993), King and Rebelo (1993), Cogley and Nason (1995), Baxter and King (1999) (BK), Christiano and Fitzgerald (2003) (CF), Hamilton (2018), Phillips and Jin (2021), and Phillips and Shi (2021).

A recent forceful rejection of the HP filter has been Hamilton’s publication entitled “Why you should never use the Hodrick-Prescott filter”. Hall and Thomson (2021) have recently evaluated the use of Hamilton’s proposed OLS regression method (H84) in a New Zealand business cycle context, and concluded that there is no material advantage in using the H84 regression over the HP filter for the purpose of presenting stylised business cycle facts; nor does the H84 predictor improve on other forecast extension methods at the ends of series, including the HP filter with no extension.

Also for New Zealand, in the context of assessing the robustness of business cycle facts across alternative filters, Hall, Thomson and McKelvie (2017) have concluded that, on balance, stylised business cycle facts from standard HP1600 filtering could be preferred over those from BK, CF, and two loess (local regression) trend filtering methods.

So, while it has been widely acknowledged in the literature that under certain conditions the HP filter can have limitations such as introducing spurious dynamic relations and reflecting filtered values at the ends of series that are very different from those in the body of the series, the HP filter has been and remains widely used in practice for trend and cycle determination, including in macroeconomic models used to guide policy formation.²

However, Phillips and Jin (2021) have recently shown that standard settings for the HP filter may not be adequate for completely removing stochastic trends in macroeconomic time series, and may also tend to over penalise shorter time series. Against that background, Phillips and Shi (2021) have proposed the use of a boosted HP filter (bHP).

Phillips and Shi (2021) develop limit theory to show that the use of a bHP filter asymptotically removes trends involving unit root processes, deterministic polynomial trends and structural breaks, and provide empirical results illustrating applications of bHP to

¹As noted in Hodrick (2021, fn. 2), the original paper by Hodrick and Prescott (1980) was started in 1978. This work was released as Carnegie-Mellon University, Working Paper 451.

²For example, Phillips and Shi (2012, fn. 3) have established that, as of August 2020, Hodrick and Prescott’s (1997) publication had over 9500 citations in Google Scholar, and Baxter and King’s (1999) paper had nearly 4000 citations.

macroeconomic time series exhibiting a range of key business cycle characteristics. They also provide detailed responses to the critique of Hamilton (2018), present numerical and empirical results that show a clear preference for the bHP filter over an autoregressive approach for trend and cycle determination, and conclude that the HP filter and boosted enhancements may validly be used as a helpful empirical device for trend and cycle determination.

Specifically, bHP is proposed as an iterative HP algorithm, terminated by two data-driven stopping rules. These are an augmented Dickey-Fuller (ADF) unit root test, and a new version of the Bayesian information criterion (BIC) developed for use in this context. Phillips and Shi (2021) consider the ADF approach to be appropriate for applications such as business cycle analysis, where trend deviations are required to be stationary (or near stationary), and they illustrate ADF for a p-value of 0.05. Their new version of BIC takes into account sample fit and effective degrees of freedom after each iteration, and is consistent with current usage of BIC-type information criteria as stopping rules in econometric work. They show that both the ADF and BIC stopping rules lead to significantly reduced mean squared errors relative to those emanating from the HP filter.

For many years, standard HP applications for quarterly macroeconomic time series have generally involved choosing λ to be 1600, not only for GDP variables but also for all other major macroeconomic variables, and usually with no prior attention to either the particular variable, country or length of time series. In contrast, applications of bHP allow for choices to be made on which stopping criterion to adopt, whether to impose a maximum number of iterations, and for applications using the ADF test what p-value to specify. A potential advantage of the bHP algorithm is therefore that the volatility, persistence, and cross correlation properties of heterogeneous macroeconomic time series may not be solely dependent on a common value of λ and a single HP iteration, as traditionally has been the case when HP1600 is used.

Against this background, we evaluate circumstances in which a bHP filter could be preferred for New Zealand business cycle analysis, relative to using an HP filter with standard setting of $\lambda = 1600$. We do this for a representative range of quarterly New Zealand macroeconomic time series typically used in small theoretical and empirical macroeconomic models and address the following key questions.

Tradition has dictated that business cycle periodicities lie between 6 and 32 quarters (e.g. Baxter and King, 1999). In the context of more recent business cycle durations, should periodicities of up to 40 quarters or more now be allowed for?

Does it matter which stopping rule (bHP-ADF or bHP-BIC) is applied? Are there other criteria that might lead to better trend selection in the context of business cycle analysis?

Are stylised business cycle facts from bHP filtering materially different from those produced from HP1600? In particular, does bHP filtering lead to New Zealand growth cycles which are noticeably different from those associated with HP1600 or BK filtering?

HP1600 is commonly used as an omnibus filter across all key macroeconomic variables. Does the greater flexibility of bHP filtering provide better alternatives?

Section 2 makes it clear that we are principally concerned with those business cycles often

termed growth or deviations-from-trend cycles. Section 3 provides our methodological underpinnings. Empirical results are presented in Section 4, and Section 5 concludes.

2 What is the business cycle, and should its duration lie between 6 and 32 quarters or 6 and 40 quarters?

Beaudry *et al.* (2020) have recently raised the possibility that post-war U.S. business cycles may reflect periodicities of as much as 36 to 40 and possibly even 50 quarters. If so, this could mean that business cycle analysis should allow for cycle durations of between 6 and (say) 40 quarters instead of the traditionally accepted definition of periodicities sitting between 6 and 32 quarters utilised by Baxter and King (1997) and Stock and Watson (1999), on the basis of work by Burns and Mitchell (1946).

Background to this possibility, and consistent with Baxter and King (1999, p 1) and Stock and Watson (1999, fn. 4), is the Beaudry *et al.* (2020, fn. 14) observation that National Bureau of Economic Research (NBER) chronology lists 30 complete cycles for the U.S. since 1858, with the shortest full cycle having been 6 quarters, the longest being 39 quarters, and 90 percent being no longer than 32 quarters. However, Beaudry *et al.* (2020) go on to suggest that a cut off of 32 quarters may no longer be appropriate, as the two most recent NBER cycles have been 43 quarters and at least 40 quarters respectively.

More recently, Kulish and Pagan (2021) have concluded that there is little merit in the Beaudry *et al.* (2020) possibility, and find that while Beaudry *et al.* (2020)'s limit cycle model for U.S. post-war hours per capita does produce an oscillation of 9–10 years, the model both fails to match the data at other frequencies and is restricted to looking only for oscillations.³ Moreover, if such an oscillation exists, Kulish and Pagan conclude that it accounts for less than 1 percent of the variance of the series.

In the context of this paper, the possibility raised by Beaudry *et al.* (2020) gives rise to two questions: which form of business cycle should we be considering; and should its duration be constrained to lie between 6 and 32 quarters or between 6 and (say) 40 quarters?

On the issue of which business cycle should be considered, one first needs to be clear on the distinction between a *classical* business cycle and a *growth* cycle. A classical cycle is associated with the pioneering work of Burns and Mitchell (1946), and is a cycle in the level (or log level) of an aggregate economic activity variable such as real gdp or as reflected in the turning points and recessions published by the U.S. NBER. A growth cycle would reflect fluctuations in an aggregate activity variable relative to an appropriate *trend* in that series, is often referred to as a “deviations-from-trend” cycle, and its turning points can therefore be observed from a wide range of detrending methods such as those put forward by Hodrick and Prescott (1971) and many others.⁴

³For explanation of oscillations, cycles, and fluctuations, see Harding and Pagan (2016, ch. 2).

⁴For further explanation of classical and growth cycles, and the methodology associated with each, see Zarnowitz (1992, ch. 7), Stock and Watson (1999, Section 2) and Harding and Pagan (2016).

On the second issue, are empirical business cycle *turning points* derived from classical and growth cycles materially different, and is this also the case for the *average durations* of classical and growth cycles? For the U.S., Stock and Watson (1999, p 13) have concluded that the distinction between classical cycle declines in aggregate activity and growth cycle recessions leads to slight differences in official NBER peaks and local maxima in bandpass filtered data. There are, though, somewhat greater differences for the average durations of post-Second World War U.S. cycles, as their classical cycle lengths have averaged around six years while the average length of their growth cycles has been around three to four years, depending on the detrending method adopted. However, many business cycle researchers seem to have failed to differentiate between classical and growth cycles for their business cycle modelling and empirical findings. As Kulish and Pagan (2021, p 2) put it: “It is therefore possibly somewhat odd that the Burns and Mitchell conclusion that business cycles were of duration between 2 and 8 years has been translated by academics into the proposition that this means one should study *oscillations* over the range of 2-8 years in the series”. The Beaudry *et al.* (2020) analysis seems to come within the latter category, as they utilise recent classical-type NBER durations to support their consideration of a business cycle having periodicity of up to 40 or 50 quarters rather than up to 32 quarters.

If there are “slight differences” between U.S. classical and growth cycle turning points, and somewhat greater differences between U.S. classical and growth cycle average durations, is this also the case for New Zealand’s small open economy? Preliminary empirical evidence from Hall and McDermott (2016, Figure 1 and Table 1), indicates that differences between New Zealand’s post-Second World War classical and growth cycles for real gdp are considerable. This is the case for the number and dates of their turning points, and also for their average cycle durations. For example, from the Hall and McDermott (2016) classical cycle turning points reflecting the Bry and Boschan (1971) (BB) dating algorithm, and their growth cycle turning points reflecting HP1600 detrending and BB assisted dating, they concluded that there have been eight completed classical business cycles of average duration 7.5 years, but 15 completed growth cycles with average duration of only four years. This classical average duration of 30 quarters sits just within the Burns and Mitchell, and Baxter-King upper bounds of 6 to 32 quarters, but this growth cycle average duration of only 16 quarters is a long way short of providing support for a periodicity beyond 32 quarters.

Hence, given that there have been more than slight differences between average durations for New Zealand’s post-Second World War classical and growth cycles, and that a principal focus of this paper is on whether or not a boosted bHP filter is more appropriate for New Zealand business cycle analysis than the frequently used HP1600 filter, the business cycles referred to in this paper will be *growth cycles* and it is in this context that we address the important empirical question of whether one should allow for cycle durations of up to (say) 40 quarters rather than standard durations of 6 to 32 quarters.

The methodology underpinning our investigation is presented in Section 3, and the associated empirical results are reported in Section 4.

3 Methodology

Throughout this paper attention is restricted to non-seasonal (or seasonally adjusted) quarterly macro-economic time series, although the methods and analyses considered are readily adapted to other time frequencies (annual, monthly etc). As noted in Section 2, our primary focus is on the estimation and analysis of growth cycles measured by the deviations from a suitably chosen trend (the HP and BK trend filters are common choices).

Let y_t be such a quarterly time series, possibly log transformed, and assume that y_t admits the additive decomposition

$$y_t = g_t + d_t \tag{1}$$

where g_t is an unobserved trend and d_t is the corresponding trend deviation. The decomposition and its conceptual components are identified by assuming that g_t is smooth, yet follows the secular general movement of the time series y_t , whereas d_t has zero mean and reflects shorter-term fluctuations and cyclical behaviour not accounted for by the trend. Ideally g_t should capture the local mean level and direction of the time series over a suitable time scale (medium to long term) and the shorter-term trend deviation d_t should be stationary or near-stationary.

The additive decomposition (1) has a long history dating back to Macaulay(1931), if not earlier, and continues to enjoy wide-spread acceptance in practice. However the concepts that underpin the identification of its components are not well-defined and remain an active topic of discussion and research (see, for example, the Royal Statistical Society discussion paper by Kenny and Durbin (1982) and the wide-ranging discussion that follows, as well as Harvey (1997)). This discussion has led to parametric structural time series models based on (1) that have been developed by Akaike (1980), Harvey (1989), Kitagawa and Gersch (1996) and Durbin and Koopman (2001) among many others.

Typically g_t is estimated by a *linear filter* of the form

$$\hat{g}_t = \sum_s w_t(s) y_{t-s}$$

where the filter weights $w_t(s)$ can be time-varying or time-invariant ($w_t(s) = w(s)$) and d_t is estimated by

$$\hat{d}_t = y_t - \hat{g}_t = \sum_s \tilde{w}_t(s) y_{t-s}$$

where $\tilde{w}_t(0) = 1 - w_t(0)$ and $\tilde{w}_t(s) = -w_t(s)$ ($s \neq 0$). Many filters used in business cycle analysis can be put into this general form including the HP filter (Hodrick and Prescott, 1997), the BK filter (Baxter and King, 1999) and simple moving average trend filters such as the Henderson filters (Henderson, 1916) used in the seasonal-trend-irregular decomposition procedure X-11 and its derivatives (see Findley *et al.* 1998). More recently Phillips and Shi (2021) have proposed an iterated form of the Hodrick-Prescott filter, called the boosted-HP filter (bHP), which builds directly on and is a generalisation of the HP filter. In these and almost all other cases, such trend filters reduce to symmetric time-invariant moving average filters in the body of the series, with trends at the ends of series estimated by asymmetric end filters, forecast extension or signal extraction using

fitted parametric models. For any chosen symmetric moving-average trend filter, forecast extension provides a general technique for extending the trend to the ends of series (see Hall and Thomson (2021) for a discussion).

Here we focus on the performance and properties of the symmetric time-invariant moving-average linear filters that underpin the bHP filter in the body of the series. These central moving-average filters define the nature of the historic trend in the body of the time series which is unaffected by the addition of new observations and any recent data revisions. Key properties of these central filters will be evaluated and compared with those of the HP filter and also those of the BK filter. The latter is a widely-used alternative to the HP filter that is directly based on a band-pass filter closely tailored to the Burns and Mitchell (1946) paradigm.

The central moving-average trend filters considered are of the form

$$\hat{g}_t = \sum_s w(s)y_{t-s} \quad (2)$$

where the time-invariant filter weights $w(s)$ are symmetric ($w(s) = w(-s)$) and sum to unity ($\sum_s w(s) = 1$). These seemingly innocuous assumptions confer a number of useful properties. Write

$$\hat{g}_t = g(L)y_t = \sum_s w(s)L^s y_t$$

where L is the backward shift operator with $Ly_t = y_{t-1}$. Then the moving-average trend filter $g(L)$ given by (2)

P1: passes constant and linear deterministic time trends without distortion so that

$$g(L)(\alpha + \beta t) = \alpha + \beta t$$

for all choices of α and β ;

P2: has trend deviation filter $1 - g(L)$ that renders $I(1)$ time series stationary with zero mean, and $I(2)$ time series stationary (not necessarily with zero mean).

These simple properties have the following ramifications. If y_t is corrected for a deterministic linear time trend estimated by ordinary least squares, then the OLS residuals

$$\hat{e}_t = y_t - \hat{\alpha} - \hat{\beta}t$$

are a time series that has zero mean and whose appearance will often appear stationary, or near stationary, at least to a rough approximation. In particular, using property P1 above,

$$\hat{g}_t = g(L)y_t = \hat{\alpha} + \hat{\beta}t + g(L)\hat{e}_t$$

so that the estimated trend \hat{g}_t is the sum of the fitted OLS time trend plus the trend of the approximately stationary OLS residuals. This observation was discussed in King and Rebelo (1993) and shows that filtering y_t is equivalent to filtering its OLS time trend

residuals. It also provides a justification for considering the frequency domain properties of the filter (2) which are normally only available for stationary (or near stationary) time series that admit a spectral representation (a sum, or superposition, of sinusoidal or Fourier components). A central concept in this regard is the *transfer function* of a linear filter.

From the theory of stationary time series, any zero mean stationary time series x_t has spectral representation

$$x_t = \int_0^\pi \cos t\omega \, dU(\omega) + \int_0^\pi \sin t\omega \, dV(\omega) = \int_{-\pi}^\pi e^{-it\omega} dZ(\omega) \quad (3)$$

where the increments $dU(\omega)$, $dV(\omega)$ (or complex $dZ(\omega)$) are mutually uncorrelated random amplitudes whose size (root mean square) measures the degree to which the component at *angular frequency* ω predominates. Note that $\omega = 2\pi f$ where f is *frequency* in cycles per unit time, or cycles per quarter in the case of the quarterly series considered here. If mainly low frequency components are sizeable then the observed time series would present a very smooth appearance, whereas sizeable components at mainly high frequencies would lead to time series that display more oscillatory behaviour. Note that (3) embodies two equivalent specifications with the one involving the complex exponential ($\exp(i\theta) = \cos \theta + i \sin \theta$) being the most mathematically tractable. This representation is given in standard texts such as Hannan (1970), Brockwell and Davis (1991) among many others.

From the spectral representation, the action of the linear filter (2) on x_t is now given by

$$g(L)x_t = \int_{-\pi}^\pi e^{-it\omega} G(\omega) dZ(\omega) \quad (4)$$

where

$$G(\omega) = g(e^{i\omega}) = \sum_s w(s) e^{is\omega} \quad (|\omega| \leq \pi) \quad (5)$$

is the *transfer function* of the filter. The filter modifies the original spectral representation of x_t by replacing the spectral amplitudes $dZ(\omega)$ by $G(\omega)dZ(\omega)$ so that frequency components where $G(\omega)$ is very small will be suppressed or *filtered out* and those where $G(\omega)$ are large will be amplified. If the transfer function $G(\omega)$ is written in polar form as

$$G(\omega) = |G(\omega)| e^{i\theta(\omega)}$$

then $|G(\omega)|$ is the *gain* and $\theta(\omega)$ the *phase* of the filter. The latter is a direct measure of any *phase shifts* or leads and lags induced by the filter $g(L)$. In the case of central moving-average trend filters considered here, the symmetric weights $w(s)$ ensure that $G(\omega)$ will always be real with

$$G(\omega) = w(0) + 2 \sum_{s>0} w(s) \cos(s\omega)$$

and $G(0) = 1$ since the weights sum to unity. However $G(\omega)$ can still be negative. Components at such frequencies will have phase $\theta(\omega) = \pm\pi$ and suffer phase shifts of half

the period of the frequency component concerned (leads or lags of π/ω). This undesirable behaviour is mitigated if there are few such cases and their gain $|G(\omega)|$ is close to zero. If $G(\omega)$ is non-negative for all ω then $G(\omega) = |G(\omega)|$ and the filter is said to be *non-negative definite*.

Moving average trend filters of the form (2) are examples of *low-pass filters* which are designed to pass frequency components below a given *cut-off frequency* and eliminate the rest. In this context the *ideal low-pass filter* with cut-off frequency ω_0 and transfer function

$$B_{\omega_0}(\omega) = \begin{cases} 1 & (|\omega| \leq \omega_0) \\ 0 & (\omega_0 < |\omega| \leq \pi) \end{cases} \quad (6)$$

serves as a benchmark to categorise and measure the effectiveness of any given trend filter $g(L)$. The intervals $|\omega| \leq \omega_0$ and $\omega_0 < |\omega| \leq \pi$ are called the *pass band* and *stop band* respectively. In particular, the ideal low-pass filter that best approximates the transfer function $G(\omega)$ given by (5) in terms of minimum mean-squared error can be shown to have a cut-off frequency ω_c that satisfies $G(\omega_c) = 0.5$ and, as a consequence, this value is said to define the cut-off frequency of the filter $g(L)$. Figure 1 shows the transfer functions of the HP filter with cut-off frequency $\omega_c = 2\pi/40$ (a period of 40 quarters or 10 years), the BK trend filter⁵ and the ideal low-pass filter, where the latter two filters have cut-off frequencies of $\omega_c = 2\pi/32$ (a period of 32 quarters or 8 years). It can be seen that the HP filter is a non-negative definite filter whereas the BK trend filter is not.




Figure 1: Transfer functions of the HP filter (red) with cut-off frequency $\omega_c = 2\pi/40$ (a period of 40 quarters or 10 years), the BK trend filter (green) and the ideal low-pass filter (black), where the latter two filters have cut-off frequencies of $\omega_c = 2\pi/32$ (a period of 32 quarters or 8 years).

⁵The BK filter estimates the trend using a finite 25 point moving average filter approximation to the ideal low-pass filter (6) with cut-off frequency $\omega_c = 2\pi/32$.

Now consider Property P2 above. Following Baxter and King (1999) we note that

$$1 - g(L) = \sum_{s>0} w(s)(1 - L^s)(1 - L^{-s}) = (1 - L)(1 - L^{-1})\Psi(L)$$

where $\Psi(L)$ is a symmetric moving average filter and $(1 - L)(1 - L^{-1})$ is the centred form of the second-difference filter Δ^2 ($\Delta = 1 - L$). In particular

$$\Psi(L) = \sum_{s>0} s^2 w(s) \Lambda_{2s-1}(L)$$

where $\Lambda_{2s-1}(L)$ is the symmetric *triangular moving average filter* given by

$$\Lambda_{2s-1}(L) = \frac{1}{s^2} \sum_{j=-s+1}^{s-1} (s - |j|) L^j$$

which arises when a time series is smoothed twice with the simple averaging filter $\frac{1}{s} \sum_{j=1}^s L^j$ (one forward pass and one backward pass). Thus the trend deviation is a weighted sum of triangular smoothing filters applied to the series

$$(1 - L)(1 - L^{-1})y_t = -L^{-1}\Delta^2 y_t$$

and $\Delta^2 y_t$ will transform non-stationary $I(1)$ time series to zero mean stationary time series, and $I(2)$ time series to stationary time series with a mean that will not necessarily be zero. If $\Delta^2 y_t$ is stationary with non-zero mean μ then the trend deviations $(1 - g(L))y_t$ will have mean $\mu \sum_{s>0} s^2 w(s)$ resulting in trend bias and violation of the the assumption that the trend deviations should have zero mean. In such cases, the trend deviations can be mean corrected with the mean deviation used to correct the bias in the level of the trend. Alternatively, this problem can be circumvented if the trend deviation filter is iterated (repeatedly applied) to give the filter $(1 - g(L))^n$ where the number of iterates n exceeds 1. This leads to higher-order differencing with the filter Δ^{2n} applied to y_t before smoothing. When $n = 2$ for example, Δ^4 will reduce $I(1)$, $I(2)$ and $I(3)$ time series to zero mean stationary time series and yield zero mean stationary trend deviations. In this way iteration (or boosting) provides a mechanism for eliminating any stochastic trend that might be present in the trend deviations $(1 - g(L))y_t$. This observation underpins the boosted HP filter (Phillips and Shi, 2021). However, since most economic time series are typically $I(1)$ or $I(2)$ and rarely higher, there would appear to be little need for more than two iterations in practice.

The following sections consider the properties of the HP filter, the bHP filter and a version of the bHP filter tailored for business (growth) cycle analysis.

3.1 HP filter

The HP filter is widely used as a general-purpose empirical trend filter for quarterly non-seasonal (or seasonally adjusted) macroeconomic time series that follow the additive

decomposition (1). Its original purpose was to decompose such time series into a growth component (trend) and cyclical component (trend deviation), but its general utility and applicability has seen it widely used in many different contexts. Here the trend g_t is estimated by \hat{g}_t which minimises the criterion

$$F + \lambda S = \sum_t (y_t - \hat{g}_t)^2 + \lambda \sum_t (\Delta^2 \hat{g}_t)^2 \quad (7)$$

where λ is a trade-off parameter balancing the fidelity F of \hat{g}_t to the data y_t with the smoothness S of \hat{g}_t . The smaller λ is the closer \hat{g}_t follows the data and the larger λ is the closer \hat{g}_t is to a simple linear trend. For quarterly data λ is normally chosen to be $\lambda = 1600$ (the standard Hodrick-Prescott filter), but other choices are possible, depending on the balance of smoothness and fidelity required.

For observed data $\mathbf{y} = (y_1, y_2, \dots, y_T)'$, minimising (7) yields the solution

$$\hat{\mathbf{g}} = H\mathbf{y}, \quad H = (I + \lambda D'D)^{-1} \quad (8)$$

where $\hat{\mathbf{g}} = (\hat{g}_1, \hat{g}_2, \dots, \hat{g}_T)$, I is the identity matrix and the $(T-2) \times T$ matrix D has typical element $D_{ij} = 1$ ($j = i, i+2$), $D_{ij} = -2$ ($j = i+1$) and $D_{ij} = 0$ otherwise. For large T , the rows of H give the weights of the HP filter with the central rows corresponding to the time-invariant weights in the body of the series and the remaining rows corresponding to the weights of the time-varying asymmetric HP end filters. Hall and Thomson (2021) show that the time-invariant weights in the body of the series are given by (2) with

$$w(s) = \alpha \sin(|s|\phi + \psi) \rho^{|s|} \quad (9)$$

where

$$\rho = 1/(\sqrt{1+\delta} + \sqrt{\delta}), \quad \alpha = 1/\sqrt{\lambda(\rho^2 + 1/\rho^2 - 2\cos 2\phi)}$$

and

$$\delta = \frac{1 + \sqrt{1 + 16\lambda}}{8\lambda}, \quad \phi = \tan^{-1} \frac{1 + \rho^2}{2\sqrt{\lambda}(1 - \rho^2)}, \quad \psi = \tan^{-1}(2\sqrt{\lambda} \tan^2 \phi)$$

with $0 < \phi, \psi < \pi/2$. These are simplified versions of the formulae given in McElroy (2008) and De Jong and Sakarya (2016). When λ is 1600 the value of ρ is 0.8941 so the weights $w(s)$ decay slowly to zero as $|s|$ increases. These weights define the *central HP filter*.

Using the weights (9), the central HP filter can be shown to have the form

$$h(L) = \sum_{-\infty}^{\infty} w(s)L^s = \frac{1}{1 + \lambda(1-L)^2(1-L^{-1})^2} \quad (10)$$

with transfer function

$$H(\omega) = \frac{1}{1 + 4\lambda(1 - \cos \omega)^2} \quad (|\omega| \leq \pi) \quad (11)$$




Figure 2: Transfer functions of the HP filters with cut-off frequencies ω_c of $2\pi/48$ (red), $2\pi/40$ (black), $2\pi/32$ (green), $2\pi/16$ (blue) and $2\pi/8$ (cyan) corresponding to periods of 48, 40, 32, 16 and 8 quarters (12, 10, 8, 4 and 2 years).

where $H(\omega)$ is monotonically decreasing from unity at $\omega = 0$ to $1/(1 + 16\lambda)$ at $\omega = \pi$. Thus the HP filter is a non-negative definite filter with no phase shifts and no ripples. Moreover $H(\omega)$ has cut-off frequency ω_c that satisfies $H(\omega_c) = 0.5$ and

$$\lambda = 1 / (4(1 - \cos \omega_c)^2), \quad \omega_c = \cos^{-1} \left(1 - 1/(2\sqrt{\lambda}) \right) \quad (12)$$

provided $\lambda \geq 1/16$. When $\lambda = 1600$ the HP cut-off frequency $\omega_c = 2\pi/39.7$ (a period of 9.9 years) which is close to 40 quarters or 10 years ($\lambda = 1649$ gives a 10 year period) and the cut-off frequency $\omega_c = 2\pi/32$ (a period of 8 years or 32 quarters) corresponds to $\lambda = 677$. These relationships are also given by Kaiser and Maravall (1999), Gomez (2001), Harvey and Trimbur (2008) among others. In particular $H(\omega)$ has high contact with $H(0) = 1$ over the pass band since the first three derivatives of $H(\omega)$ at $\omega = 0$ are zero, and $H(\omega)$ is close to zero over the stop band when λ is large. These desirable properties, among others, are inherited from the Butterworth filter of which the HP filter is a special case (see Gomez, 2001, Harvey and Trimbur, 2003, for example).

Figure 2 shows the transfer functions of the central HP filter for the cut-off frequencies ω_c given by $2\pi/48$, $2\pi/40$, $2\pi/32$, $2\pi/16$ and $2\pi/8$ (λ values of 3416, 1649, 677, 43 and 3 respectively) corresponding to periods of 48, 40, 32, 16 and 8 quarters (12, 10, 8, 4 and 2 years). As noted earlier, the HP filter with $\lambda = 1649$ is essentially the same as the standard HP filter with $\lambda = 1600$. Evidently the degree to which HP filters are well-approximated by ideal low-pass filters is a function of the cut-off frequency ω_c with the quality of the approximation better for lower values of ω_c . One simple measure of the quality of the approximation of a filter to an ideal low-pass filter is the steepness (slope) of the filter's transfer function at ω_c . We call this measure the *sharpness* of the filter. For

the HP transfer function $H(\omega)$ this is given by

$$\beta_c = -1/(2 \tan(\omega_c/2)) \quad (13)$$

and the larger $|\beta_c|$, the better the approximation. Note that when ω_c is small $2\pi|\beta_c|$ is approximately $2\pi/\omega_c$ which is just the period of the cut-off frequency.

Since the HP filter is a central moving-average trend filter of the form (2) it has Properties P1 and P2 listed earlier. However more can be said. For any time series y_t , note that the HP trend deviation is given by

$$(1 - h(L))y_t = \lambda h(L)(1 - L)^2(1 - L^{-1})^2 y_t$$

where $(1 - L)^2(1 - L^{-1})^2$ is the centred form of the fourth-difference filter Δ^4 . Thus the HP trend deviation is an HP-smoothed version of the fourth-differences $\Delta^4 y_t$ scaled by λ . It will transform non-stationary $I(1)$, $I(2)$ and $I(3)$ time series to zero-mean stationarity and transform $I(4)$ time series to stationarity, but with a mean that is not necessarily zero. It also implies that the central HP trend filter $h(L)$ passes constant, linear, quadratic and cubic deterministic time trends without distortion.

As noted in Hall and Thomson (2021), the HP filter (8) and the central HP filter weights (9) with transfer function (11) can be computed using alternative methods. King and Rebelo (1993) show that the HP filter can be given a model-based interpretation with (1) comprising a stochastic trend g_t that satisfies

$$\Delta^2 g_t = \epsilon_t \quad (14)$$

where ϵ_t , d_t are mutually independent Gaussian white noise processes. Under these assumptions \hat{g}_t can be computed using the Kalman filter and smoother (see Harvey and Jaeger, 1993). Kaiser and Maravall (1999, 2012) use this model and Wiener-Kolmogorov filtering to forecast missing values at the ends of the series (forecast-extension) and then apply a computationally efficient form of the central HP filter (9) to the extended series. Gomez (1999) shows that these three procedures are equivalent.

Furthermore, Mise et al (2005) build on King and Rebelo (1993) to show that the HP filter is the optimal (minimum mean-squared error) estimator of g_t in the body of the series for two general classes of stochastic trend models. Let y_t be given by (1) with components that follow the models ($d = 1$ or $d = 2$) given by

$$\Delta^d g_t = C(L)\epsilon_t, \quad d_t = C(L)(1 - L)^{2-d}\eta_t \quad (15)$$

where $C(z) = \sum_{j=0}^{\infty} c_j z^j$ ($c_0 = 1, \sum_{j=0}^{\infty} c_j^2 < \infty$) is non-zero for $|z| \leq 1$ and ϵ_t , η_t are mutually independent Gaussian white noise processes. The reduced models for y_t are given by

$$\Delta^d y_t = C(L)(1 - 2\rho \cos \theta L + \rho^2 L^2)u_t$$

where ρ and θ are given by (9) and u_t is Gaussian white noise. *For both models ($d = 1$ or $d = 2$) the central HP filter generates the optimal estimator of g_t in the body of the series.* Suitable choices for $C(L)$ allow for a more general class of data generation processes than

just the case $C(L) = 1$ given in (14). In particular, y_t follows an $ARIMA(p, d, q)$ model when $C(L) = A(L)/(B(L)(1 - 2\rho \cos \theta L + \rho^2 L^2))$ where the invertible moving average operator $A(L)$ has order q and the stationary autoregressive operator $B(L)$ has order p . These models are examples of the $I(1)$ and $I(2)$ economic time series models commonly met in practice.

In summary, the central HP filter:

- is an infinite, symmetric, central moving average filter with coefficients $w(s)$ given by (9) that sum to unity and decay exponentially to zero;
- passes constant, linear, quadratic and cubic deterministic time trends without distortion;
- has a trend deviation filter that renders $I(1)$, $I(2)$, and $I(3)$ time series stationary with zero mean, and $I(4)$ time series stationary (not necessarily with zero mean);
- is a non-negative definite filter (no phase distortion) with a transfer function that is monotonically decreasing from 1 ($\omega = 0$) to $1/(1 + 16\lambda)$ ($\omega = \pi$);
- is a special case of the Butterworth filter and inherits many of its excellent properties;
- has a cut-off frequency ω_c given by (12) which decreases as λ increases;
- has sharpness β_c given by (13) whose magnitude (quality) increases as ω_c decreases (λ increases) with $|\beta_c|$ well-approximated by the period of ω_c when ω_c is small;
- can be calculated efficiently and conveniently by a variety of methods;
- generates the optimal trend estimator for a wide class of $I(1)$ and $I(2)$ data generation processes commonly met in practice.

These properties, among others, go some way to explaining why the HP filter has been so widely used in macro-economic applications and more generally. However, despite its wide-spread use in practice, the HP filter has attracted considerable criticism in the academic literature. At one extreme, Hamilton (2018) makes a case for why you should never use the HP filter, a case not universally accepted by others (see the discussions in Phillips and Shi (2021), Hall and Thomson (2021) for example). At the other extreme, Phillips and Jin (2021) review and further extend the theoretical properties of the HP filter and show why it is often successful in practice. In addition to quantifying the dependence of choice of λ on series length, they show that the standard HP filter with $\lambda = 1600$ will often fail to include all of any stochastic trend present in the data, leaving part of the stochastic trend in the HP trend deviation. The latter deficiency is directly addressed by the bHP filter discussed in Section 3.2.

A key issue with the HP filter is that it is an infinite central moving-average trend filter which is almost always applied to finite length time series. As with all central moving average filters (finite or infinite) special techniques are needed to provide suitable trend estimates at the ends of series. In such cases model-based methods provide a suitable

solution with forecast extension being more generally appropriate. The latter provides useful improvements over the standard HP filter defined by (7) and (8) (see Hall and Thomson, 2021, and the references therein) and is expected to provide similar advantages when applied to the bHP filter.

3.2 Boosted HP filter

Phillips and Jin (2021) show that a key deficiency of the standard HP filter is that it often fails to completely remove stochastic trend components from the HP trend deviation. They conclude that it is likely that remnants of stochastic trends are present, to some degree, in the HP trend deviations (estimated cycles) of much applied business cycle analysis. This observation provides a key rationale for the *boosted HP filter* (Phillips and Shi, 2021) which builds on and extends the HP trend filter by iteration. The latter involves repeated application of the HP trend deviation filter to the data until the resulting trend deviation is free of any stochastic trend components that may be present.

As noted earlier, our focus is on the performance and properties of the bHP filter in the body of the series. Here the bHP filter is given by

$$h_n(L) = 1 - (1 - h(L))^n \quad (16)$$

with transfer function

$$H_n(\omega) = 1 - (1 - H(\omega))^n \quad (|\omega| \leq \pi) \quad (17)$$

where n is the number of iterations, $h(L)$ is the HP filter (10) and $H(\omega)$ is the HP transfer function (11). When $n = 1$ the bHP trend is the same as the HP trend since $h_1(L) = h(L)$ and

$$h_n(L) = h_{n-1}(L) + h(L)(1 - h_{n-1}(L)) \quad (18)$$

for $n > 1$. This recursion shows that the bHP trend at iteration n is just the bHP trend at the previous iteration $n - 1$ plus the HP trend of the bHP trend deviation at iteration $n - 1$. This process repeats until the bHP trend deviation shows no evidence of any stochastic trend. This trend deviation is then subtracted from the data to give the final bHP trend.

Phillips and Shi (2021) set $\lambda = 1600$ and propose two stopping criteria for determining the number of iterations needed. Both criteria are applied to the bHP trend deviations. The first method applies the Augmented Dickey-Fuller test for unit roots with significance level set to 0.05 and selects the first iteration n when the alternative hypothesis of stationarity is accepted. The resulting filter is called the bHP-ADF filter. The second method is based on minimising a specially constructed Bayesian information criterion (BIC) with the resulting bHP filter called the bHP-BIC filter. An empirical evaluation of the properties of the bHP-ADF and bHP-BIC filters is given in Section 4.1 where these filters are applied to a selection of key New Zealand macro-economic time series.

However the bHP filters can also be regarded as a family of trend filters indexed by the two parameters λ ($\lambda > 0$) and n (n a positive integer) with values that can be freely

chosen by the economic analyst. The bHP trend filters are not especially designed for business cycle analysis. Rather, they are a general family of trend filters whose fidelity to the data and smoothness are controlled by suitable choice of n and λ . While Phillips and Shi (2021) advocate $\lambda = 1600$ with n chosen by either the ADF or BIC stopping rules, other criteria could also be used. An example of a criterion tailored for business cycle analysis is given in Section 3.3.

The rest of this section considers the general properties of the central bHP trend filters in the body of the series with these filters indexed by λ and n . Since the bHP filters are based on the HP filter, these properties build on those for the HP filter given in Section 3.1.

The central bHP filter is an infinite, symmetric, central moving average filter of the form (2) with weights given by (9) in the case $n = 1$ (the HP filter) and calculated recursively from (18) for $n > 1$. The latter involves convolution of the HP filter with the bHP trend deviation filter at the previous iteration. The bHP transfer function $H_n(\omega)$ given by (17) is non-negative definite (no phase distortion) and decreases monotonically from $\omega = 0$ to $\omega = \pi$ with

$$H_n(0) = 1, \quad H_n(\pi) = 1 - (1 - 1/(1 + 16\lambda))^n.$$

From (17) it can be seen that for ω near 0, $H_n(\omega)$ is closer to 1 (the ideal) than $H(\omega)$. Also $H_n(\pi) \geq H(\pi)$ for all n (it is approximately $nH(\pi)$ for large λ) and $H_n(\pi)$ increases as n increases to a limit of $H_\infty(\pi) = 1$. Compared to the HP filter, the bHP filter is closer to the ideal filter in the pass band of the HP trend filter, but further away in the stop band. This is unlikely to be of any consequence when λ is large or n is small. However, since the stop band of a trend filter is the pass band of the corresponding trend deviation filter, use of the bHP filter could potentially lead to less accurate estimates of the cycle when n is large and λ small.

In general, as n increases $H_n(\omega)$ approaches $H_\infty(\omega) = 1$ and the bHP trend approaches the original series (zero trend deviations). This behaviour is inconsistent with the normal definition of a trend so stopping rules, penalty functions or other criteria, are necessary to determine suitably parsimonious values for n . The transfer function $H_n(\omega)$ has cut-off frequency $\omega_c^{(n)}$ given by

$$\omega_c^{(n)} = \cos^{-1} \left(1 - 1/(2\sqrt{\lambda(2^{1/n} - 1)}) \right) \quad (19)$$

where $H_n(\omega_c^{(n)}) = 0.5$ and $\omega_c^{(1)} = \omega_c$ given by (12). The sharpness of $H_n(\omega)$ (the slope of the transfer function at the cut-off frequency $\omega_c^{(n)}$) is given by

$$\beta_c^{(n)} = -n(1 - 2^{-1/n})/\tan(\omega_c^{(n)}/2) \quad (20)$$

with $\beta_c^{(1)} = \beta_c$ given by (13). The values of $\omega_c^{(n)}$ monotonically increase with n and the corresponding values of $|\beta_c^{(n)}|$ monotonically decrease provided $n < \log 2/\log(1 + 1/(16\lambda))$ (17745 for $\lambda = 1600$). Table 1 gives numerical values of these quantities for the first 10 iterations in the case where $\lambda = 1600$. Here the period of the cut-off frequency decreases from approximately 10 years ($n = 1$) to 8 years ($n = 2$), 7 years ($n = 3$), 6 years ($n = 5$)




Figure 3: Transfer functions of bHP filters with $\lambda = 1600$ for $n = 1$, the HP filter (black), $n = 2$ (red), $n = 11$ (green), $n = 26$ (blue) and $n = 381$ (cyan) iterations with cut-off frequencies $\omega_c^{(n)}$ corresponding to periods of 40, 32, 20, 16 and 8 quarters (10, 8, 5, 4 and 2 years).

Table 1: Values of the period $2\pi/\omega_c^{(n)}$ (quarters) of the cut-off frequency and the absolute sharpness ($|\beta_c^{(n)}|$) of the bHP filter with $\lambda = 1600$ for $n = 1, \dots, 10$ iterations.

n	1	2	3	4	5	6	7	8	9	10
$2\pi/\omega_c^{(n)}$	39.7	31.8	28.3	26.1	24.6	23.4	22.5	21.7	21.1	20.5
$ \beta_c^{(n)} $	6.3	5.9	5.6	5.3	5.0	4.9	4.7	4.6	4.4	4.3

or 6) and less than 6 years ($n > 5$) with absolute sharpness deteriorating by more than 20% from its value for $n = 1$ (the HP filter) when $n \geq 5$.

Figure 3 shows the transfer functions of the central bHP filters for $\lambda = 1600$ and n iterations where $n = 1$ (the HP filter), $n = 2, 11, 26$ and 381 where these values of n correspond to cut-off frequencies $\omega_c^{(n)}$ with periods 40, 32, 20, 16 and 8 quarters (10, 8, 5, 4 and 2 years). A measure of the quality of these low-pass filters as approximations to ideal low-pass filters is given by the sharpness $\beta_c^{(n)}$ (20). In this case these values (the slopes of the transfer function at $\omega_c^{(n)}$) are -6.3, -5.9, -4.2, -3.4, and -1.7 respectively. The first two, the HP filter with $n = 1$ and the bHP filter with $n = 2$, are essentially the same with the quality of the others worsening as n increases and as expected from (20). Thus, for $\lambda = 1600$, the HP trend filter and the bHP trend filter with $n = 2$ are of much the same quality, but one (the HP filter) passes frequency components with periods greater than 40 quarters (10 years) whereas the other passes frequency components with periods greater than 32 quarters (8 years).

Like the HP filter, the bHP filter passes polynomial time trends and renders integrated

time series stationary. From (16), the bHP trend deviation for any time series y_t is

$$(1 - h_n(L))y_t = (1 - h(L))^n y_t = (\lambda h(L))^n ((1 - L)(1 - L^{-1}))^{2n} y_t$$

where $((1 - L)(1 - L^{-1}))^{2n}$ is the centred form of the iterated difference filter Δ^{4n} and so the central bHP trend filter $h_n(L)$ passes deterministic polynomial time trends of degree less than $4n$ without distortion. For $n = 2$, for example, the bHP trend deviation filter involves differencing the original series 8 times with the degree of differencing ramping up rapidly for larger values of n . As a consequence, the bHP trend deviation filter will transform non-stationary $I(p)$ series ($p < 4n$) to zero-mean stationarity and transform $I(4n)$ time series to stationarity, but with a mean that is not necessarily zero. Whereas the bHP trend deviation for $n = 1$ (the HP filter) may not always remove all the stochastic trend present in economic data, the case $n = 2$ almost always will.

In summary, the central bHP filter:

- is an infinite, symmetric, central moving average filter with coefficients $w(s)$ given by (9) and (18) where these weights sum to unity and decay exponentially to zero;
- passes deterministic polynomial time trends of degree $p < 4n$ without distortion;
- has trend deviation filter that renders $I(p)$ ($p < 4n$) time series stationary with zero mean, and $I(4n)$ time series stationary (not necessarily with zero mean);
- is a non-negative definite filter (no phase distortion) with a transfer function that is monotonically decreasing from 1 ($\omega = 0$) to $1 - (1 - 1/(1 + 16\lambda))^n$ ($\omega = \pi$);
- has a cut-off frequency $\omega_c^{(n)}$ given by (19) which increases as n increases;
- has sharpness $\beta_c^{(n)}$ given by (20) whose magnitude (quality) decreases as n increases;
- has much the same sharpness (quality) as the HP filter when $n = 2$ and $\lambda = 1600$, with the HP filter having cut-off frequency $2\pi/40$ (10 year period) and the bHP filter ($n = 2, \lambda = 1600$) having cut-off frequency $2\pi/32$ (8 year period);
- has trend deviation filter that will almost always remove any stochastic trend present in economic data when $n = 2$ or greater, but may not when $n = 1$ (HP filter).

Like the HP filter on which it is based, the central bHP filter is an infinite central moving-average trend filter that will need to be adapted for finite length time series. The current implementation of the bHP filter (Phillips and Shi, 2021) recursively applies the standard HP filter to the data. However there is no guarantee that this will improve on the well-documented poor performance of the HP filter at the ends of series. This potential deficiency of the bHP filter is not investigated here, but left for future research. Forecast extension is a standard technique applied in such situations (see Hall and Thomson, 2021, and the references therein) and is expected to provide similar advantages when applied to the bHP filter.

3.3 Boosted HP filter with constant cut-off frequency

In general, the bHP filter provides suitable trend estimation procedures for removing any stochastic and deterministic trends present in data with these procedures placed on a sound theoretical footing in Phillips and Jin (2021) and Phillips and Shi (2021). However, in a business cycle context, the resulting trend filter may end up with a higher cut-off frequency than expected from theoretical economic considerations and applied econometric practice. If n is large (say $n > 5$) the bHP filter will lead to estimated cycles that retain frequency components with periods less than 6 years, but suppress (filter out) components with periods 6 years or greater. This is at variance with the discussion given in Section 2 and the Burns and Mitchell (1946) paradigm.

To address this concern we now consider a sub-class of the bHP filters (bHPc) indexed by the number of iterations n where the cut-off frequency is held fixed at some chosen value ω_c . In line with Section 2, we have in mind the case where $\omega_c = 2\pi/32$ (a cut-off frequency of 32 quarters or 8 years), but other cut-off frequencies could be adopted such as the cut-off frequency of the HP filter (40 quarters or 10 years) among others. To maintain a constant cut-off frequency the value of λ used in the bHP filter will now need to depend on n . This value is given by

$$\lambda_n = \frac{\lambda_1}{2^{1/n} - 1}, \quad \lambda_1 = \frac{1}{4(1 - \cos \omega_c)^2} \quad (21)$$

where n is the number of iterations ($n \geq 1$) and ω_c is the chosen cut-off frequency. Note that λ_n increases with n and, for large n , is asymptotically equivalent to $n\lambda_1/\log 2$. The sharpness of the bHPc filter is given by

$$\beta_c^{(n)} = -n(1 - 2^{-1/n})/\tan(\omega_c/2) \quad (22)$$

so that the absolute sharpness increases with n to a limiting value of $(\log 2)/\tan(\omega_c/2)$. The ratio of the absolute sharpness to its limit is independent of the choice of ω_c and converges reasonably quickly to 1.

Table 2 gives the first 10 values of λ_n and the absolute sharpness $|\beta_c^{(n)}|$ for the bHPc filter with cut-off frequency $\omega_c = 2\pi/32$ (32 quarter period) and also $\omega_c = 2\pi/40$ (40 quarter period). The limiting absolute sharpness values are shown in each case. Note that the absolute sharpness of the bHPc filter is approximately 85% of its limiting value when $n = 2$, 90% of its limiting value when $n = 3$ and greater than 90% when $n > 3$. This suggests that the associated transfer functions will be very similar for $n \geq 2$.

From (17) and (21) the transfer function of the bHPc filter is given by

$$H_n(\omega) = 1 - \left(1 - \frac{1}{1 + 4\lambda_n(1 - \cos \omega)^2}\right)^n \quad (|\omega| \leq \pi) \quad (23)$$

which converges to the limiting transfer function

$$H_\infty(\omega) = 1 - 2^{-((1 - \cos \omega_c)/(1 - \cos \omega))^2} \quad (|\omega| \leq \pi) \quad (24)$$

Table 2: Values of λ_n and the absolute sharpness $|\beta_c^{(n)}|$ for the bHP filter with constant cut-off frequency $\omega_c = 2\pi/32$ (32 quarter period) and $\omega_c = 2\pi/40$ (40 quarter period) for $n = 1, \dots, 10$ iterations. The limiting values ($n = \infty$) for the absolute sharpness are also shown in each case as are the absolute sharpness values for the sharpened HP filter sHP with $\lambda = 677$ (sHP677) and $\lambda = 1600$ (sHP1600).

		$\omega_c = 2\pi/32$											
n		1	2	3	4	5	6	7	8	9	10	∞	sHP677
λ_n		677	1635	2605	3579	4554	5529	6505	7481	8458	9434		
$ \beta_c^{(n)} $		5.1	5.9	6.3	6.5	6.6	6.6	6.7	6.7	6.8	6.8	7.0	7.6

		$\omega_c = 2\pi/40$											
n		1	2	3	4	5	6	7	8	9	10	∞	sHP1600
λ_n		1649	3982	6345	8717	11092	13468	15845	18223	20601	22980		
$ \beta_c^{(n)} $		6.4	7.4	7.9	8.1	8.2	8.3	8.4	8.4	8.5	8.5	8.8	9.5

as the number of iterations n increases. Figure 4 shows these transfer functions when the number of iterations is $n = 1, 2, 3$ and $n = \infty$ for cut-off frequencies $\omega_c = 2\pi/32$ and $\omega_c = 2\pi/40$ corresponding to periods of 32 and 40 quarters (8 and 10 years). As expected, the bHPc filters with $n \geq 2$ provide a good approximation to the limiting filter (24).

In practice it will be very difficult to differentiate the outputs of the bHPc filters when $n \geq 2$. This suggests that the only cases of practical importance for growth cycle analysis are $n = 1$ and $n = 2$ with the latter only marginally different from (24). Of these two filters the sharpest (the twice iterated case $n = 2$) is the obvious choice.

Other ways of using iteration to sharpen a trend filter while maintaining a fixed cut-off frequency are available. As noted in Phillips and Shi (2021), the bHP filter with $n = 2$ is an example of “twicing”, a data smoothing procedure advocated by Tukey (1977) which also has its counterparts in the design of digital filters where sequential application of filters (iteration) can lead to better sharpness properties. Kaiser and Hamming (1977) (see also Section 6.6 of Hamming, 1989) explores “twicing” and develops a general technique for sharpening a given moving average trend filter $g(L)$ of the form (2) using linear combinations of its iterates $g(L)^n$. For non-negative integers r and s , the sharpened filter takes the functional form $P_{rs}(g(L))$ where

$$P_{rs}(x) = x^{s+1} \sum_{k=0}^r \binom{s+k}{k} (1-x)^k$$

is a polynomial in x of degree $r + s + 1$, $P_{rs}(0) = 0$, $P_{rs}(1) = 1$ and $P_{rs}(x)$ has zero derivatives to order s at $x = 0$ and order r at $x = 1$ (implying $P_{rs}(x)$ has high contact to $P_{rs}(0) = 0$ near $x = 0$ and $P_{rs}(1) = 1$ near $x = 1$). In particular, the transfer function of $P_{rs}(g(L))$ is $P_{rs}(G(\omega))$ where $G(\omega)$ is the transfer function of the moving average trend filter $g(L)$.

For the important symmetric case $r = s$

$$P_{ss}(0.5) = 0.5, \quad P'_{ss}(0.5) = (2s + 1) \binom{2s}{s} (0.5)^{2s}$$




Figure 4: Transfer functions of bHPc filters with constant cut-off frequency $\omega_c = 2\pi/32$ (left plot) and $\omega_c = 2\pi/40$ (right plot) corresponding to periods of 32 and 40 quarters respectively. The number of iterations chosen is $n = 1$ (black), $n = 2$ (red), $n = 3$ (green), $n = \infty$ (cyan) and only the lower frequencies are shown since the values of the transfer functions at higher frequencies are negligible. The simple symmetric sharpening filter sHP is also shown (blue).

where $P'_{ss}(0.5)$, the derivative of $P_{ss}(x)$ at $x = 0.5$, is monotonically increasing in s with $P'_{00}(0.5) = 1$ and $P'_{ss}(0.5)$ well-approximated by $\sqrt{1 + 4s/\pi}$. It follows that *the transfer function $P_{ss}(G(\omega))$ has the same cut-off frequency as $G(\omega)$* , the transfer function of the moving average trend filter $P_{ss}(g(L))$ is based on, with sharpness given by

$$\tilde{\beta}_c^{(s)} = (2s + 1) \binom{2s}{s} (0.5)^{2s} G'(\omega_c)$$

where ω_c is the cut-off frequency of $G(\omega)$ and $G'(\omega_c)$ is its sharpness. Moreover $P_{ss}(G(\omega))$ approaches the ideal filter with cut-off frequency ω_c as s increases. Two important special cases of these symmetric sharpening filters are $P_{00}(g(L)) = g(L)$ (no sharpening) and $P_{11}(g(L)) = 3g(L)^2 - 2g(L)^3$ which uses iterates of order 2 and 3. Kaiser and Hamming (1977) refer to the latter case as *simple symmetric sharpening* which produces a 50% improvement in the sharpness of $G(\omega)$ since $P'_{11}(0.5) = 1.5$.

Applying simple symmetric sharpening to the HP filter $h(L)$ with cut-off frequency ω_c yields

$$P_{11}(h(L)) = 3h(L)^2 - 2h(L)^3 = 3h_2(L) - 2h_3(L) \quad (25)$$

where $h_n(L)$ is the bHP trend filter (17). This filter (sHP) is a linear combination of the bHP filters with $n = 2$ and $n = 3$ and has the same cut-off frequency ω_c as the HP filter it is based on.⁶ A plot of the transfer function of the sHP filter is given in Figure 4 for

⁶All sharpened HP filters $P_{rs}(h(L))$ can be written as linear combinations of bHP filters, but will only have the same cut-off frequency as the HP filter $h(L)$ when $r = s$.

the two cases where the HP filter being sharpened has cut-off frequency ω_c corresponding to periods of 32 or 40 quarters. The sHP trend filter is essentially the same as the bHPc filter with $n = 3$ in the pass band, but is significantly better than all bHPc filters in the stop band (its sharpness values are given in Table 2). Whether this improvement in the stop band translates to practical improvements in the estimate of the cycle is considered in Section 4.

In summary, the central bHPc trend filter:

- is an infinite, symmetric, central moving average filter whose weights sum to unity and decay exponentially to zero;
- passes deterministic polynomial time trends of degree $p < 4n$ without distortion;
- has trend deviation filter that renders $I(p)$ ($p < 4n$) time series stationary with zero mean, and $I(4n)$ time series stationary (not necessarily with zero mean);
- is a sub-class of the bHP filters indexed by number of iterations n and constant cut-off frequency ω_c ;
- has values of λ given by (21) that increase with n and are asymptotically proportional to n ;
- has sharpness $\beta_c^{(n)}$ given by (22) whose magnitude (quality) increases as n increases with limiting absolute sharpness $(\log 2)/\tan(\omega_c)$;
- is a non-negative definite filter with transfer function (23) which converges to the limiting form (24) as n increases;
- includes two primary cases of practical importance, $n = 1$ (HP filter) and $n = 2$ (bHP filter with $n = 2$), where the latter procedure is also referred to as “twicing”;
- can be further improved using simple symmetric sharpening (a linear combination of the bHP filters with $n = 2$ and $n = 3$) to give the sHP filter (25) whose sharpness exceeds that of all bHPc trend filters.

The stopping criteria proposed in Phillips and Shi (2021) for selecting the number of iterations in the bHP filter have less relevance for the bHPc filter where the cut-off frequency is constant and, in practice, only a few low values of n will be considered. However the use of the Augmented Dickey Fuller test may assist with the choice of cut-off frequency which is likely to remove more of any stochastic trend component the larger the cut-off frequency becomes. Such issues will be discussed further in Section 4.1.

Throughout Section 3 our discussion has focussed on the properties of the various filters (HP, bHP and bHPc) in the body of the series and does not address the important case of the ends of series. However, as noted earlier, forecast extension is expected to provide a suitable strategy in such cases.⁷ In Section 4 we present a variety of empirical evaluations of selected HP, bHP and bHPc trend filters applied to representative New Zealand macroeconomic time series.

⁷For some end of series results in the context of HP1600 and Hamilton (2018) filters, see Hall and Thomson (2021).

4 Results

Utilising the methodology presented in Section 3, we now report results from our empirical evaluations of the bHP and bHPc trend filters, contrasting and comparing these with results from the more traditional HP and BK filters. The key questions set out in Section 1 are a primary focus. An evaluation of bHP trends and stopping rules is given in Section 4.1, and stylised business cycle facts for HP, bHP and bHPc trend deviations (estimated cycles) are compared in Section 4.2. Section 4.3 considers what material differences, if any, can be discerned between bHP, HP and BK growth cycles.

Table 3: New Zealand quarterly macroeconomic data considered: code, start and end dates.

Code	Start	End	Series
gdpe	1987q2	2019q4	log GDP (expenditure)
gdpp	1987q2	2019q4	log GDP (production)
cons	1987q2	2019q4	log consumption (private)
invr	1987q2	2019q4	log investment (residential)
gfcf	1987q2	2019q4	log gross fixed capital formation
gcon	1987q2	2019q4	log government consumption expenditure
nxsh	1987q2	2019q4	net exports share (%)
mtot	1987q2	2019q4	log imports goods & services
empl	1987q2	2019q4	log employment
unem	1987q2	2019q4	unemployment rate (%)
cpix	1987q3	2019q4	CPI inflation (annual % change)
cpin	1989q1	2019q4	CPI nontradables (annual % change)
nine	1987q3	2019q4	real 90-day Bank Bill rate (%)

Our quarterly, seasonally adjusted data set of key New Zealand macroeconomic variables reflects variables typically included in theoretical or empirical macroeconomic models of small open economies. They have been sourced from Statistics New Zealand (SNZ), the RBNZ and Treasury, and are as documented in McKelvie and Hall (2012, Appendix C) with the exception of the CPI non-tradables series which comes from the RBNZ. Series were log-transformed with the exception of those containing negative observations (e.g. net exports share, CPI inflation rate, real 90-day Bank Bill rate) or those already expressed as a percentage (e.g. unemployment). Table 3 lists the series considered, their start and end dates, and the code used to identify them in the sections that follow.

Throughout the rest of this paper we adopt the following notations for the particular HP, bHP or bHPc filters used. References to HP1600, HP677, 2HP1600 or 5HP1600 filters, for example, refer to HP or bHP filters where the number of iterations selected is the numerical prefix (assumed to be 1 if missing) and the numerical suffix gives the value of λ . Similarly, references to sHP677 or sHP1600 filters refer to the sharpened HP filter given by (25) with the numerical suffix denoting the value of λ chosen. Furthermore, all computations and graphical analysis were carried out in the R statistical environment (R Development Core Team, 2004) and, in particular, we acknowledge use of the R function *BoostedHP(.)* made available as part of Phillips and Shi (2021).

4.1 An empirical evaluation of bHP trends and stopping rules.

Phillips and Shi (2021) propose two stopping rules for selecting the iterations of a bHP trend filter both of which are applied to the bHP trend deviations. The bHP-ADF filter uses the Augmented Dickey-Fuller test for unit roots with significance level set to 0.05 and selects the first iteration when the alternative hypothesis of stationarity is accepted. This stopping rule is considered appropriate for situations where stationary trend deviations are deemed a prerequisite, such as for business cycle analysis. The bHP-BIC filter is based on minimising a specially constructed Bayesian information criterion (BIC) which takes into account sample fit and effective degrees of freedom after each iteration. Both filters provide smooth trends that always fit the data at least as closely, if not more closely, than the standard HP1600 filter on which they are based.

In this section an empirical evaluation of the properties of the bHP-ADF and bHP-BIC filters based on the HP1600 filter is undertaken in a New Zealand context with the filters applied to our seasonally-adjusted data set of 13 New Zealand macroeconomic time series. The nature of the trends fitted is considered as is the stability of the stopping rules and the number of iterations selected.

Figure 5 shows a selection of four of the New Zealand macroeconomic time series considered (log GDPE, log residential investment, unemployment and real 90-day bank bill rate) with bHP-ADF, bHP-BIC and 2HP1600 trends superimposed. The latter (bHP with $n = 2$ and $\lambda = 1600$) is chosen as a benchmark trend due to its excellent theoretical properties (see Sections 3.2, 3.3) and its cut-off frequency $\omega_c = 2\pi/32$ corresponding to a period of 8 years (32 quarters). As expected, the bHP-BIC trend always fits the data best (has highest fidelity), with the bHP-ADF filter being the HP1600 filter (log GDPE, log investment (residential), unemployment) or the 2HP1600 (90-day bank bill rate).

Table 4: Number of iterations selected by the bHP-ADF and bHP-BIC trend filters for each of the New Zealand macroeconomic series considered.

	gdpe	gdpp	cons	invr	gcf	gcon	nxsh	mtot	empl	unem	cpix	cpin	nine
ADF	1	1	3	1	1	1	1	1	1	1	1	1	2
BIC	7	6	6	5	7	3	5	6	5	9	13	8	5

The number of iterations selected by the bHP-ADF and bHP-BIC trend filters for all 13 macroeconomic time series are given in Table 4. While the bHP-ADF trend filter almost always selects the HP1600 filter, the exceptions being log consumption ($n = 3$) and real 90-day bank bill rate ($n = 2$), the selections for the bHP-BIC trend filter range from 3 to 13 with a median of 6 and a mode of 5. From Table 1 the latter correspond to cut-off frequencies with periods around 6 years (24 quarters) compared to the HP1600 (10 years) and the 2HP1600 (8 years). In all cases the bHP-BIC filter selects more iterations than the bHP-ADF filter, in most cases considerably more, and so the bHP-BIC trend deviations will have reduced autocorrelation structure compared to that of the corresponding bHP-ADF trend deviations.

These observations are consistent with the advice given in Phillips and Shi (2021) that the bHP-ADF filter is the more appropriate stopping rule for business cycle analysis.

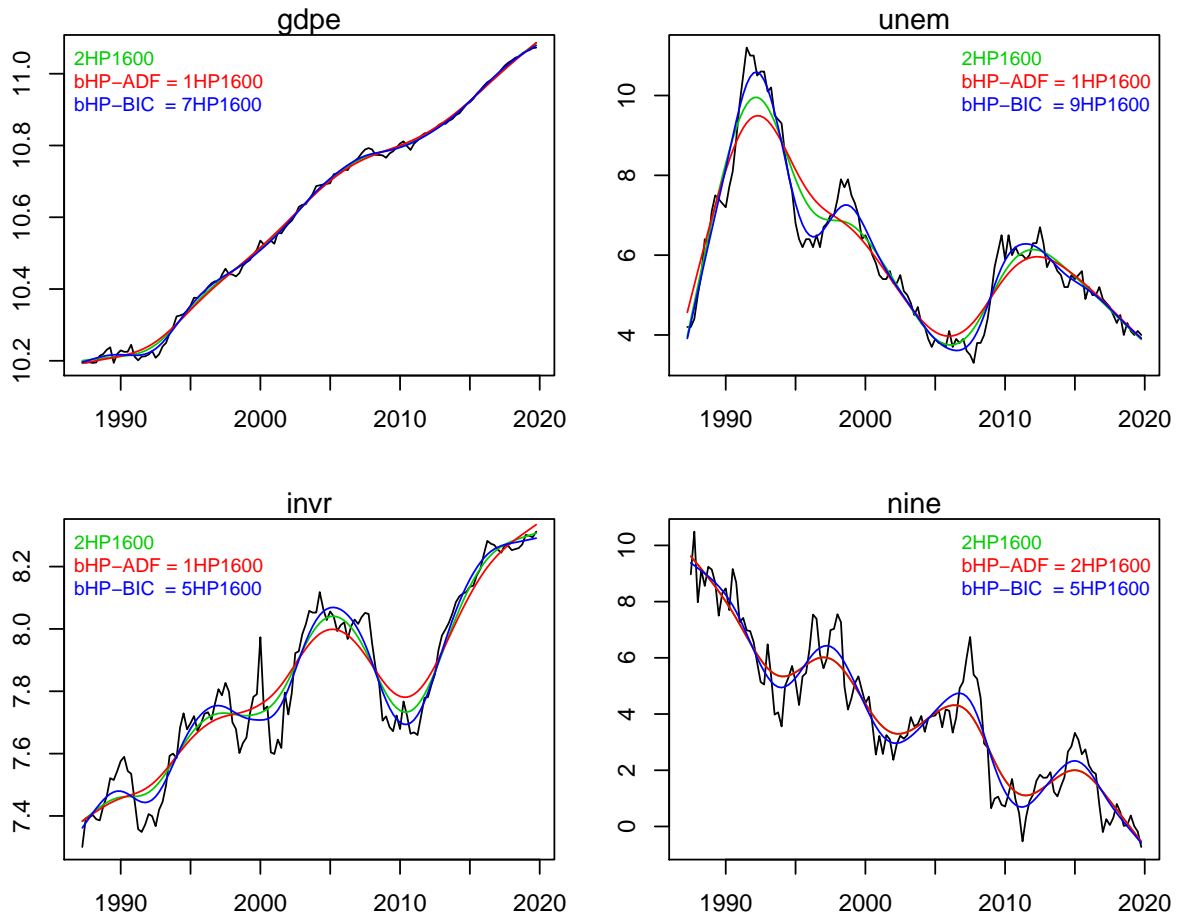


Figure 5: bHP trends for New Zealand log GDPE, log residential investment, unemployment (%) and real 90-day bank bill rate (%). Three trends are shown: 2HP1600 (green), bHP-ADF (red), bHP-BIC (blue) with the bHP models selected by the ADF and BIC stopping rules shown in each case.

This rule is directly based on the Augmented Dickey-Fuller test for unit roots and, as a consequence, should ensure that the resulting bHP-ADF trend deviations are free of any residual stochastic trend. Table 5 shows the p-values of the ADF test applied to the HP1600, 2HP1600 and 3HP1600 bHP trend deviations with p-values for the HP677 and the sharpened HP filter sHP677 (both with cut-off period 8 years or 32 quarters) also shown for comparison.

From Table 5 the bHP-ADF stopping rule (an ADF test with significance level 0.05) selects $n = 1$ (the HP1600 filter) for the majority of the 13 series considered. In most cases the result is reasonably unequivocal with the exceptions being CPI nontradables (cpin), where the unit root null hypothesis is marginally rejected, and real 90-day bank bill rate (nine), where it is marginally retained. The HP1600 trend deviation for log private consumption (cons) shows clear evidence of a unit root. By contrast, the p-values

Table 5: p-values of the ADF test applied to the HP1600, 2HP1600 and 3HP1600 bHP trend deviations for each of the 13 New Zealand macroeconomic quarterly series considered. The p-values for the HP677 and the sharpened HP filter sHP677, both with a cut-off period of 8 years (32 quarters), are also shown.

	gdpe	gdpp	cons	invr	gfcf	gcon	nxsh	mtot	empl	unem	cpix	cpin	nine
HP1600	0.040	0.022	0.141	0.010	0.012	0.010	0.010	0.010	0.039	0.014	0.010	0.047	0.053
2HP1600	0.010	0.010	0.054	0.010	0.010	0.010	0.010	0.010	0.010	0.010	0.010	0.015	0.013
3HP1600	0.010	0.010	0.023	0.010	0.010	0.010	0.010	0.010	0.010	0.010	0.010	0.010	0.010
HP677	0.012	0.010	0.080	0.010	0.010	0.010	0.010	0.010	0.010	0.010	0.010	0.017	0.017
sHP677	0.010	0.010	0.039	0.010	0.010	0.010	0.010	0.010	0.010	0.010	0.010	0.015	0.010

of the ADF test on the 2HP1600 trend deviations all convincingly reject the unit root null-hypothesis in favour of stationarity with the possible exception of cons where the p-value is only just greater than 0.05. As expected, there is no evidence of any residual stochastic trend in the 3HP1600 trend deviations (recall that the 3HP1600 filter has a 7 year cut-off period).

In part, the near-uniform rejection of the unit root null hypothesis for the 2HP1600 trend deviations is a result of the filter’s 8 year cut-off period which places more low frequency variation in the trend by comparison to the HP1600 filter with 10 year cut-off period. Sharpness also plays a role since sharper trend filters will better extract low frequency and trend components. While the HP677 trend filter has an 8 year cut-off period, it is not as sharp as the HP1600 filter (10 year cut-off period) or the 2HP1600 filter (8 year cut-off period) which both have much the same sharpness, but are not as sharp as the sHP677 (8 year cut-off period) which is the sharpest. The p-values in Table 5 reflect these observations. In particular, the 2HP1600 and sHP677 trend filters have 8 year cut-off periods, are as sharp or sharper than the HP1600 with 10 year cut-off period and, unlike the HP1600, have trend deviations that are largely free of any residual stochastic trend.

The previous comments and discussion have a direct bearing on which, if any, of the bHP filters would be suitable as an omnibus filter. The HP1600 filter is commonly used in practice as an omnibus filter applied to all quarterly macroeconomic variables considered. While this strategy has met with some criticism, Kaiser and Maravall (1999, 2012) show that application of the HP1600 filter generally has little effect on the estimation of lag-zero cross-correlations between series. Since the HP1600 filter is a non-negative definite filter (no phase distortion), similar results are to be expected for cross-correlations at lags other than zero. However, as noted above and in Phillips and Shi (2021), HP1600 trend deviations are not always free of residual stochastic trend and so the HP1600 filter with 10 year cut-off period will not always be appropriate for use as an omnibus filter. Based on the properties of the bHP filters given in Section 3.2 and as borne out by the empirical results in Table 5, the 2HP1600 filter with 8 year cut-off period is a much better candidate for use as an omnibus filter.

Finally, we consider the sensitivity of the bHP-ADF and bHP-BIC stopping rules to series length and data augmentation. For long stable time series one might expect that

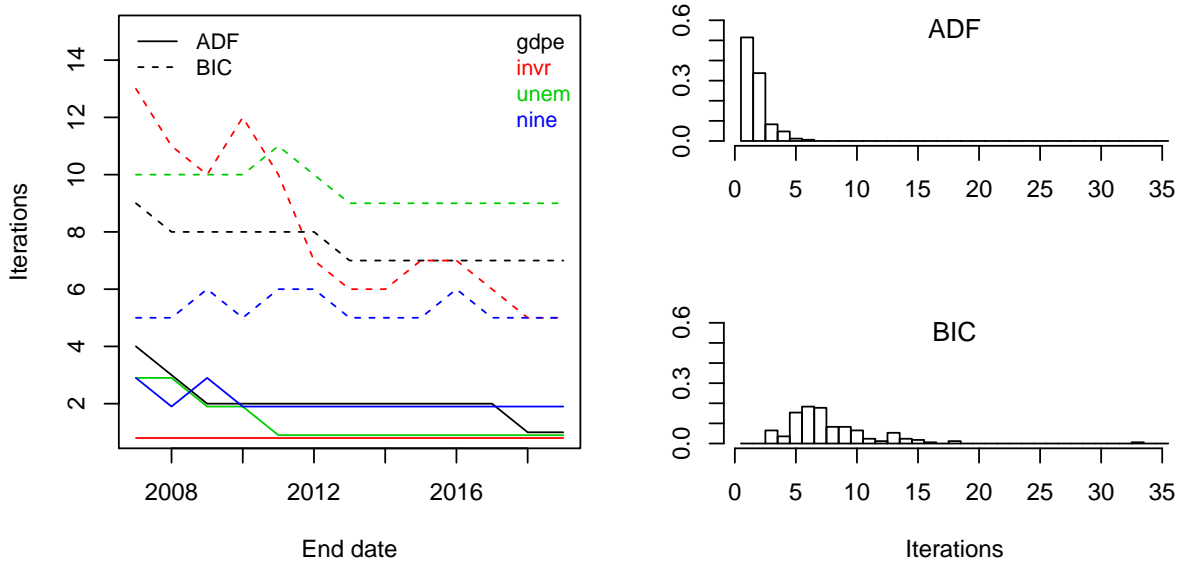


Figure 6: The left plot shows the number of iterations selected by the bHP-ADF and bHP-BIC trend filters for quarterly New Zealand log GDPE, log residential investment, unemployment (%) and real 90-day bank bill rate (%). In each case 13 increasing length series are considered with start date 1987q2 and annual end dates from 2007q4 to 2019q4. The right plots show the histograms of the iterations selected by the bHP-ADF (top) and bHP-BIC (bottom) trend filters for all 13 New Zealand macroeconomic series considered and all 13 series lengths.

bHP trend selection (ADF or BIC) would vary little with series length. If that were not the case then earlier historical trend estimates could be revised, possibly significantly, as a consequence of augmenting the time series with recent data. Such revisions are undesirable, especially in any official statistics setting where historic trends should remain unchanged if at all possible. To investigate this issue we consider applying the bHP trend filters to our New Zealand macroeconomic quarterly data where the starting date is 1987q2 and the end dates are successively chosen to be 2007q4, 2008q4, . . . 2019q4. The base period of 1987q2 to 2007q4 is approximately 20 years of quarterly data which would seem a minimum expectation, in practice, for stable stopping times. The 13 annual data augmentations considered give rise to time series whose lengths range from 20 to 33 years. Applying the bHP-ADF and bHP-BIC trend filters to these series should provide a guide, not only for the series lengths required for stopping rule stability, but also their volatility once stability has been achieved.

Figure 6 plots the number of iterations selected by the bHP-ADF and bHP-BIC trend filters for the four New Zealand macroeconomic time series shown in Figure 5 (log GDPE, log residential investment, unemployment and real 90-day bank bill rate). In each case 13 increasing length series are considered with start date 1987q2 and annual end dates from 2007q4 to 2019q4. For these series the bHP-BIC iterations generally show greater variability than the bHP-ADF iterations with both showing greater stability from around 2010 (ADF) and 2012(BIC) or approximately 25 years (100 quarters) of data. Some

iteration series show reasonably dramatic change as series length is augmented with others very little. For example, bHP-BIC log residential investment shows considerable variation over the period 2007 to 2019, yet bHP-ADF for the same series shows none since it selects HP1600 throughout. Note also that the bHP-ADF trend for log GDPE is essentially the 2HP1600 filter from 2009 to 2017, but the HP1600 filter from 2018. Such variability is typical of the other series in our New Zealand macroeconomic data.

Table 6: Summary statistics for the number of iterations selected by the bHP-ADF and bHP-BIC trend filters for each of the 13 New Zealand macroeconomic quarterly series considered. For each series, 13 increasing length series are considered with start date 1987q2 and annual end dates from 2007q4 to 2019q4.

		gdpe	gdpp	cons	invr	gfcf	gcon	nxsh	mtot	empl	unem	cpix	cpin	nine	All
ADF	Min	1	1	2	1	1	1	1	1	1	1	1	1	2	1
	LQ	2	2	2	1	1	1	1	1	2	1	1	2	2	1
	Med	2	2	2	1	1	1	1	1	2	1	1	3	2	1
	UQ	2	2	3	1	1	1	1	1	2	2	1	4	2	2
	Max	4	4	4	1	2	1	2	2	4	3	5	6	3	6
BIC	Min	7	6	6	5	7	3	4	6	4	9	13	8	5	3
	LQ	7	7	6	6	7	3	5	6	5	9	13	9	5	5
	Med	7	7	6	7	7	3	5	6	5	9	14	10	5	7
	UQ	8	8	6	10	8	3	5	11	6	10	15	12	6	9
	Max	9	9	7	13	9	4	7	33	6	11	18	13	6	33

Table 6 provides summary statistics for all 13 New Zealand macroeconomic quarterly series considered and all 13 series lengths with start date 1987q2 and annual end dates 2007q4 to 2019q4. It also provides overall summary statistics across all 13 series. Here the bHP-ADF trend filter selects iterations with a median and lower quartile of 1 (HP1600 with 10 year cut-off) and an upper quartile of 2 (2HP1600 with 8 year cut-off) whereas the bHP-BIC trend filter selects iterations with a median of 7 (7HP1600 with 5.6 year cut-off), lower quartile of 5 (5HP1600 with 6.2 year cut-off) and upper quartile of 9 (9HP1600 with 5.3 year cut-off). Figure 6 also shows the histograms of all the iterations selected by the bHP-ADF and bHP-BIC trend filters. Evidently, bHP-ADF selects far fewer iterations than bHP-BIC in general with the former being considerably less variable (predominantly $n = 1$ or $n = 2$).

In summary and as noted in Phillips and Shi (2021), the bHP-ADF trend filter is more appropriate for business cycle analysis although in some cases its variation with series augmentation may prove to be a limitation in practice. The bHP-BIC trend filter would not normally be a candidate for business cycle analysis, but may well be appropriate in other contexts. Of the bHP-ADF trend filters selected by our New Zealand macroeconomic data, the two predominant choices are the HP1600 and 2HP1600 with cut-off periods of 10 years (40 quarters) and 8 years (32 quarters) respectively. Of the two filters, only the 2HP1600 trend filter almost always delivers trend deviations free of any residual stochastic trend. As a consequence, there is a stronger case for using the 2HP1600 trend filter as an omnibus filter rather than the HP1600 trend filter.

Table 7: Stylised business cycle facts for New Zealand macroeconomic data over the period 1987q2–2019q4: volatility, persistence and most significant cross correlations with log GDPE. Numbers in parentheses are robustly estimated standard errors; volatility and cross-correlation measures assume constant volatility trend deviations.

Series	Volatility				Persistence			
	HP1600	2HP1600	sHP677	bHP-BIC	HP1600	2HP1600	sHP677	bHP-BIC
gdpe	1.30 (0.17)	1.12 (0.13)	1.14 (0.13)	0.87 (0.09)	0.72 (0.06)	0.62 (0.07)	0.62 (0.07)	0.46 (0.08)
gdpp	1.18 (0.21)	0.90 (0.14)	0.95 (0.17)	0.67 (0.09)	0.85 (0.04)	0.77 (0.05)	0.80 (0.05)	0.65 (0.07)
cons	1.48 (0.17)	1.16 (0.12)	1.08 (0.10)	0.83 (0.08)	0.79 (0.05)	0.68 (0.06)	0.64 (0.07)	0.45 (0.08)
invr	7.97 (0.88)	6.36 (0.74)	6.40 (0.72)	5.51 (0.58)	0.79 (0.05)	0.71 (0.06)	0.69 (0.06)	0.62 (0.07)
gfcf	5.46 (0.72)	4.57 (0.51)	4.64 (0.51)	3.49 (0.35)	0.76 (0.06)	0.68 (0.06)	0.67 (0.07)	0.52 (0.07)
gcon	1.29 (0.15)	1.15 (0.13)	1.18 (0.13)	1.12 (0.12)	0.47 (0.08)	0.36 (0.08)	0.35 (0.09)	0.32 (0.08)
nxsh	1.23 (0.15)	1.12 (0.13)	1.12 (0.14)	1.00 (0.12)	0.70 (0.06)	0.65 (0.07)	0.63 (0.07)	0.57 (0.07)
mtot	4.16 (0.58)	3.75 (0.48)	3.73 (0.47)	3.26 (0.43)	0.77 (0.05)	0.73 (0.06)	0.73 (0.06)	0.67 (0.06)
empl	1.30 (0.17)	0.91 (0.10)	0.94 (0.12)	0.69 (0.06)	0.90 (0.04)	0.80 (0.05)	0.82 (0.05)	0.68 (0.06)
unem	0.55 (0.13)	0.45 (0.09)	0.48 (0.10)	0.30 (0.04)	0.87 (0.04)	0.81 (0.05)	0.84 (0.05)	0.62 (0.07)
cpix	0.86 (0.10)	0.81 (0.09)	0.84 (0.10)	0.65 (0.06)	0.76 (0.06)	0.73 (0.06)	0.77 (0.06)	0.62 (0.07)
cpin	0.81 (0.16)	0.73 (0.15)	0.75 (0.15)	0.61 (0.11)	0.76 (0.06)	0.71 (0.07)	0.73 (0.06)	0.61 (0.07)
nine	0.87 (0.12)	0.77 (0.09)	0.77 (0.10)	0.68 (0.06)	0.70 (0.06)	0.61 (0.07)	0.66 (0.07)	0.49 (0.08)

Series x_t	Cross-correlation with log GDPE				Lag
	HP1600	2HP1600	sHP677	bHP-BIC	
gdpp	0.86 (0.18)	0.82 (0.18)	0.84 (0.18)	0.76 (0.18)	x_t
cons	0.73 (0.16)	0.68 (0.16)	0.64 (0.14)	0.61 (0.16)	x_t
invr	0.68 (0.13)	0.65 (0.13)	0.63 (0.13)	0.61 (0.14)	x_t
gfcf	0.74 (0.15)	0.67 (0.14)	0.66 (0.14)	0.57 (0.14)	x_t
gcon	0.42 (0.13)	0.33 (0.13)	0.24 (0.11)	0.29 (0.12)	x_{t+5}
nxsh	-0.47 (0.14)	-0.41 (0.14)	-0.34 (0.13)	-0.35 (0.15)	x_{t+2}
mtot	0.54 (0.14)	0.49 (0.14)	0.45 (0.14)	0.41 (0.15)	x_{t+2}
empl	0.55 (0.13)	0.41 (0.13)	0.50 (0.14)	0.23 (0.12)	x_{t+2}
unem	-0.59 (0.15)	-0.48 (0.14)	-0.51 (0.14)	-0.27 (0.11)	x_{t+1}
cpix	0.46 (0.13)	0.48 (0.13)	0.47 (0.15)	0.46 (0.14)	x_{t+4}
cpin	0.60 (0.16)	0.57 (0.14)	0.52 (0.14)	0.48 (0.14)	x_{t+4}
nine	0.41 (0.12)	0.33 (0.12)	0.31 (0.13)	0.28 (0.12)	x_{t+2}

4.2 Are bHP stylised business cycle facts materially different from those produced by a standard HP filter?

We consider stylised business cycle facts represented by volatility, persistence and most significant cross correlation measures for our 13-variable data set. Our particular interest is in the results for the HP1600 trend filter with 10 year cut-off period and the HP677, 2HP1600 and sHP677 trend filters which all have 8 year cut-off periods. With the exception of HP677⁸, these results are presented in Table 7 where the outcomes for bHP-BIC are also given and standard errors are calculated using the same methodology as that given in Hall *et al.* (2017). Differences between all outcomes are assessed.

From Table 7 the values for bHP-BIC volatility, persistence and cross correlation are

⁸Results for HP677 are not presented in Table 7, as no results are statistically different from results for 2HP1600. Moreover, all cross correlation and persistence magnitudes are close to identical, and only the volatility magnitudes for log residential investment and log employment show minor differences (6.50 versus 6.36, and 0.97 versus 0.91 respectively).

almost always lower in magnitude than those for HP1600, 2HP1600 and sHP677. The two exceptions are the sHP677 cross correlations for *gcon* and *nxsh*, although differences are not statistically significant. All cross correlation differences between bHP-BIC estimates, and 2HP1600 and HP1600 estimates are not statistically significant, but bHP-BIC volatility and persistence estimates are statistically different from those for HP1600 for around half of our 13 variables, including for major real sector variables such as *gdpe*, *gdpp*, *cons*, *gfcf*, and *empl*. The considerably lower values for almost all of bHP-BIC's volatilities, persistence and cross correlations are consistent with this filter removing more low frequency and (stochastic) trend components than the other filters.

What are the relative merits of HP1600, 2HP1600 and sHP677 filters for stylised business cycle fact measures?⁹ On the basis of sharpness there is little to pick between the HP1600 and 2HP1600 filters, with both sharper than the HP677, but not as sharp as the sHP677. However the HP1600 filter has a 10 year cut-off period which places more low frequency components, including any residual stochastic trend, in the estimated cycle than the HP677, 2HP1600 or sHP677 filters which all have an 8 year cut-off period. Hence, as expected, the HP1600 measures of volatility, persistence and cross correlation are almost always greater in magnitude than those for 2HP1600 and sHP677, although no difference is statistically significant. Furthermore, all measures for the 2HP1600 and sHP677 filters are essentially the same with differences that are not statistically significant.

On balance, therefore, our preference is for the 2HP1600 (or very similar sHP677) measures which have our preferred 8 year cut-off period, rather than HP1600 measures with their less desirable 10 year cut-off period.

4.3 Are bHP New Zealand growth cycles noticeably different from the corresponding HP1600 or BK growth cycles?

Growth cycles can be produced using a wide range of detrending methods. Accordingly, as well as presenting results using the methodology adopted in Hall and McDermott (2016) establishing growth cycle turning points from BB dating of HP1600 trend deviations, we assess the average durations from growth cycles based on two *ocular* sets of assumptions. We term one ocular set *generous* and the other *conservative*.¹⁰

Section 2 cited preliminary evidence from a previous study that the differences between New Zealand's post-Second World War classical and growth cycles have been considerable, with the former having had an average cycle of 7.5 years and the latter only four years. This evidence was derived from HP1600 filtering with its associated 10 year cut-off period.

⁹For a comparative analysis of New Zealand stylised business cycle facts using HP1600, BK, CF and *loess* (local regression) trend filtering methods, see Hall *et al.* (2017); and for work finding a clear preference for measures produced by the HP1600 and BK filters rather than those from Hamilton's H84 procedure, see Hall and Thomson (2021).

¹⁰For an explanation and illustration of an ocular approach to determining growth cycle turning points for U.S. quarterly per capita hours, see Kulish and Pagan (2021, section 2.1.1). This approach, using the term *eyeballing*, has previously been illustrated for New Zealand in Kim, Buckle and Hall (1995) and Hall, Kim and Buckle (1998).




Figure 7: Log GDP trend deviations for HP1600 (black), BK (red), 2HP1600 (green) and sHP677 (blue) trend filters.

Given the evidence presented in favour of using a 2HP1600 filter with 8 year cut-off period, we now consider the extent to which HP1600, 2HP1600, sHP677 and BK trend deviations could lead to materially different growth cycles and average cycle durations.¹¹ Figure 7 shows these trend deviations for log GDP. Since they all display very similar movements and potential turning points, one might expect any differences in growth cycles and durations to be modest.

Using the Hall and McDermott (2016) methodology for our shorter, considerably updated sample period 1987q2 to 2019q4, we first confirm that HP1600-based growth cycles have remained significantly different from New Zealand’s classical cycles. Results are shown in Table 8 where completed growth cycle durations average 2.88 years compared to the classical cycle average of 8 years. In Table 9 BB-assisted growth cycle turning points are presented for 2HP1600 and BK to complement those presented in Table 8 for HP1600. Not surprisingly, the average growth cycle durations for 2HP1600 and BK are similar to HP1600 (2.58 years for 2HP1600 and 2.36 years for BK) with the 2HP1600 average

¹¹Many other transparent rule-based approaches would also be possible, including for example those advanced by Canova (1994). Use of some Canova rules have been illustrated for New Zealand and Pacific Rim cycles in Kim, Buckle and Hall (1995) and Hall, Kim and Buckle (1998). For this study, results from Canova’s (1994) Rule 1 did not shed any further light over and above the results presented in Tables 8, 9 and 10, so are not presented.

growth cycle durations adopting a compromise between the HP1600 and BK.¹²

Table 8: New Zealand real GDP business cycles 1987q2–2019q4. Classical cycle turning points reflect Bry and Boschan (1971) dating of updated Hall and McDermott (2016) series. Growth cycle turning points reflect HP1600 detrending and Bry-Boschan assisted dating.

Classical cycles					HP1600 growth cycles				
Peak	Trough	Duration (quarters)			Peak	Trough	Duration (quarters)		
		Expansion	Contraction	Peak to Peak			Expansion	Contraction	Peak to Peak
1987q4	1988q4		4		1987q4	1988q4		4	
					1989q2	1990q2	2	4	6
1990q4	1991q2	8	2	12	1990q4	1992q3	2	7	6
1997q2	1998q1	24	3	26	1996q1	1998q3	14	10	21
					2000q1	2001q1	6	4	16
					2002q4	2003q2	7	2	11
					2004q1	2004q4	3	3	5
2007q4	2009q2	39	6	42	2007q3	2009q2	11	7	14
					2010q2	2010q4	4	2	11
					2012q1	2013q4	5	7	7
					2016q3		11		18
2019q4		42		48					
Average (quarters)		28.25	3.75	32	Average (quarters)		6.5	5	11.5
Average (years)		7.06	0.94	8	Average (years)		1.63	1.25	2.88

Table 9: 2HP1600 and Baxter-King (BK) real GDP growth cycles 1987q2–2019q4: Bry-Boschan assisted turning points and cycle durations. Growth cycle turning points reflect 2HP1600 and BK detrending with Bry-Boschan assisted dating. BK turning points reflect 12 observations lost at beginning and end of series.

2HP1600 growth cycles					Baxter-King growth cycles				
Peak	Trough	Duration (quarters)			Peak	Trough	Duration (quarters)		
		Expansion	Contraction	Peak to Peak			Expansion	Contraction	Peak to Peak
1987q4	1988q4		4						
1989q2	1990q2	2	4	6					
1990q4	1992q3	2	7	6		1992q3			
					1995q3	1996q2	12	3	
1996q4	1998q3	17	7	24	1996q4	1998q4	2	8	5
2000q1	2001q1	6	4	13	2000q1	2001q2	5	5	13
2002q4	2003q2	7	2	11	2002q2	2003q1	4	3	9
2004q1	2005q4	3	7	5	2004q2	2005q1	5	3	8
2007q4	2009q2	8	6	15	2007q4	2009q2	11	6	14
2010q2	2010q4	4	2	10	2010q2	2011q1	4	3	10
2012q1	2013q4	5	7	7	2012q1	2013q4	4	7	7
2014q4	2015q2	4	2	11					
2016q3	2017q1	5	2	7					
2018q4		7		9					
Average (quarters)		5.83	4.5	10.33	Average (quarters)		5.88	4.75	9.43
Average (years)		1.46	1.13	2.58	Average (years)		1.47	1.19	2.36

Could adopting a less mechanistic ocular or eyeball approach to dating growth cycle turning points produce average cycle durations which are somewhat longer, and therefore potentially of greater usefulness for medium-term fiscal and monetary policy purposes? In the following we confine the illustrative ocular turning points and average durations

¹²Turning point and cycle duration results for sHP677 are not presented in Tables 9 and 10, as outcomes are the same as those for 2HP1600.

Table 10: 2HP1600 and Baxter-King (BK) real GDP growth cycles 1987q2–2019q4: ocular turning points and cycle durations.

2HP1600 conservative ocular growth cycles					2HP1600 generous ocular growth cycles				
		Duration (quarters)					Duration (quarters)		
Peak	Trough	Expansion	Contraction	Peak to Peak	Peak	Trough	Expansion	Contraction	Peak to Peak
1990q4	1992q3		7		1990q4	1992q3		7	
1997q2	1998q3	19	5	26	1997q2	1998q3	19	5	26
2000q1	2001q1	6	4	11	2000q1	2001q1	6	4	11
					2004q1	2005q4	12	7	16
2007q4	2009q2	27	6	31	2007q4	2009q2	8	6	15
					2010q2	2010q4	4	2	10
					2012q1	2013q4	5	7	7
2016q3		29		35	2016q3		11		18
Average (quarters)		20.25	5.50	25.75	Average (quarters)		9.29	5.43	14.71
Average (years)		5.06	1.38	6.44	Average (years)		2.32	1.36	3.68

BK conservative ocular growth cycles					BK generous ocular growth cycles				
		Duration (quarters)					Duration (quarters)		
Peak	Trough	Expansion	Contraction	Peak to Peak	Peak	Trough	Expansion	Contraction	Peak to Peak
1990q3	1992q3		8		1990q3	1992q3		8	
1996q4	1998q4	17	8	25	1996q4	1998q4	17	8	25
2000q1	2001q2	5	5	13	2000q1	2001q2	5	5	13
					2004q2	2005q1	12	3	17
2007q4	2009q2	26	6	31	2007q4	2009q2	11	6	14
					2010q2	2011q1	4	3	10
					2012q1	2013q4	4	7	7
2016q3		29		35	2016q3		11		18
Average (quarters)		19.25	6.75	26.00	Average (quarters)		9.14	5.71	14.86
Average (years)		4.81	1.69	6.50	Average (years)		2.29	1.43	3.71

presented in Table 10 to those from 2HP1600 and BK, as 2HP1600 is marginally preferred to HP1600, and both 2HP1600 and BK have 8 year cut-off periods. As noted earlier two ocular methods are considered, one *conservative* and the other *generous*.

Our *conservative* ocular turning points reflect only those which are most obvious to the eye, and are therefore relatively few in number. But even on this basis, growth cycle durations average only 6.44 years for 2HP1600 and 6.50 for BK, both well below 9–10 years. Determining a *generous* set of ocular turning points is more problematical, given the well-known volatility of New Zealand’s real GDP movements over certain periods. However, for the turning points offered in Table 10, it is not surprising that average 2HP1600 and BK durations are similar at 3.68 and 3.71 years, even lower than 9–10 years, and are also modestly above the BB-assisted HP1600, 2HP1600 and BK average durations in Tables 8 and 9.

In summary, our illustrative range of average growth cycle durations for New Zealand’s real production-based GDP, computed from HP1600, 2HP1600 and BK trend deviations, are much the same and provide confirmatory evidence that there is no need to consider cycles whose frequency components have periods less than 10 years. Restricting the cycle to have frequency components with periods less than 8 years is more than adequate and has the virtue of better removing any residual stochastic trend from the estimated cycle. This finding is consistent with that of Canova (1998), who used HP1600 and several other filters to find that for U.S. real GNP over the period 1955q3 to 1986q3, trend deviation

cycles had average durations in the range 4 to 6 years.

5 Conclusions

We investigate whether the boosted HP filter (bHP) proposed by Phillips and Shi (2021) might be preferred for New Zealand trend and growth cycle analysis, relative to using the HP filter with standard setting $\lambda = 1600$ (HP1600). The theoretical properties of the bHP filter in the body of a quarterly time series are explored and a variety of empirical evaluations undertaken.

In the body of the series, the bHP filter is equivalent to a symmetric time-invariant moving average filter whose time and frequency domain properties are given and transfer function determined. In particular, the bHP trend filter is a non-negative definite filter (no phase distortion) whose transfer function predominantly passes frequencies below its cut-off frequency and suppresses frequencies above it. An appropriate measure of the filter's quality is its sharpness which is defined here to be the slope of the transfer function at the cut-off frequency. Both the cut-off frequency and the sharpness of the trend filter are primary characteristics of the filter which describe its action and effectiveness with sharper filters better extracting low frequency or trend components. The HP1600 trend filter (bHP with number of iterations $n = 1$) has a cut-off frequency with 10 year period and much the same sharpness as the 2HP1600 filter which has a cut-off frequency with 8 year period. Both are suitable candidates for use in growth cycle analysis unlike the bHP filters for $n > 2$. The latter, while appropriate in other contexts, have cut-off periods which decrease as n increases.

A sub-class of bHP filters (bHPc) is proposed where the parameter λ now depends on n and a constant cut-off frequency of choice. Such filters are appropriate for growth cycle analysis and have sharpness that increases with n . However differences between bHPc filters for $n > 2$ are unlikely to be practically significant and so the primary cases of practical interest are the HP filter ($n = 1$) and the bHP filter with $n = 2$. Linear combinations of bHP filters can also be used to construct sharper trend filters with given cut-off frequency (see Kaiser and Hamming, 1977). A special case is simple symmetric sharpening of the HP filter (sHP) which yields a trend filter that is a linear combination of bHP filters ($n = 2$ and $n = 3$) whose sharpness exceeds all bHPc trend filters. In particular, the sHP677 and 2HP1600 trend filters have the same cut-off frequency (8 year period) with the sHP677 filter being almost 30% sharper than the 2HP1600 filter. In practice, however, it is likely that any small advantages conferred by the sharper filter will be outweighed by the computational simplicity of the 2HP1600 filter which is just two passes of the standard HP filter.

Our empirical evaluations are based on a quarterly, seasonally adjusted data set of 13 key New Zealand macroeconomic variables typically included in macroeconomic models of small open economies. These evaluations are primarily focused on the key questions given in Section 1 and have led to the following findings.

Two stopping rules for the bHP filter are proposed in Phillips and Shi (2021) with one

based on the ADF unit root test (bHP-ADF) and the other on a Bayesian information criterion (bHP-BIC). The bHP-ADF trend filter is more appropriate for growth cycle analysis although in some cases its variation with series augmentation is a possible limitation in practice. The bHP-BIC trend filter would not normally be a candidate for growth cycle analysis, but may well be appropriate in other contexts. Of the bHP-ADF trend filters selected by our New Zealand macroeconomic data, the two predominant choices are the HP1600 and 2HP1600 with cut-off periods of 10 years and 8 years respectively. Of these, only the 2HP1600 trend filter almost always delivers trend deviations free of any residual stochastic trend and, as a consequence, there is a stronger case for using the 2HP1600 trend filter as an omnibus filter rather than the HP1600 trend filter.

Are stylised business cycle facts from bHP filtering materially different to those produced from HP1600? Here the focus is mainly on the bHP filters most likely to be used for growth cycle analysis which are the HP1600 filter with 10 year cut-off period and the 2HP1600, sHP677 filters with 8 year cut-off period. For completeness, the bHP-BIC filter is also considered. As expected, measures of volatility, persistence and cross-correlation of the associated trend deviations are almost always lowest in magnitude for bHP-BIC compared to those for HP1600, 2HP1600 and sHP677. For the most part these are material differences which are statistically significant. On the other hand and again as expected, these measures are almost always highest for the HP1600 filter whose 10 year cut-off period places more low frequency components, including any residual stochastic trend, in the estimated cycle. However, the HP1600, 2HP1600 and sHP677 measures are not statistically significant, with the 2HP1600 and sHP677 measures being very similar. Our preference, therefore, is for measures based on the 2HP1600 or sHP677 filters which have 8 year cut-off period and are as sharp, if not sharper, than the HP1600 filter.

Are New Zealand bHP growth cycles noticeably different from the corresponding HP1600 or BK cycles? Here we note that business cycle periodicities are traditionally assumed to lie between 6 and 32 quarters (8 years). In the context of more recent business cycle durations, and as emphasised recently by Beaudry *et al.* (2020), should periodicities up to 40 quarters or more now be considered? New Zealand GDPP growth cycles for HP1600, 2HP1600, sHP677 and BK trend deviations are considered, and a range of growth cycle turning points derived using one rule-based and two less mechanistic methods. The trend deviation movements associated with the four filters look remarkably similar, yielding average cycle durations that vary from 2.4 to 6.5 years, all of which are significantly below 40 quarters (10 years). This provides confirmatory evidence that restricting the cycle to frequency components with periods less than 8 years, rather than 10 years, is costless in practice for growth cycle analysis. Moreover it has the virtue of better removing any residual stochastic trend from the estimated cycle.

Our main conclusions are that a 6 to 32 quarter business cycle periodicity is sufficient to reflect New Zealand growth cycles and determine stylised business cycle facts and, for our representative 13-variable New Zealand macroeconomic data set, using a bHP filter (2HP1600) as an omnibus filter for growth cycle analysis is preferable to using the standard HP filter (HP1600).

References

- Akaike, H. (1980) Seasonal adjustment by a Bayesian modelling. *Journal of Time Series Analysis* 1, 1-13.
- Baxter, M. and King, R.G. (1999) Measuring business cycles: approximate band-pass filters for economic time series. *Review of Economics and Statistics* 81(4), 575-593.
- Beaudry, P., Galizia, D. and Portier, F. (2020) Putting the Cycle Back into Business Cycle Analysis. *American Economic Review* 110(1), 1-47.
- Beverage, S. and Nelson, C.R. (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, 151-174.
- Brockwell, P.J. and Davis, R.A. (1991) *Time Series and Methods (2nd ed)*. Springer-Verlag, New York.
- Bry, G. and Boschan, C. (1971) *Cyclical analysis of time series: selected procedures and a computer program*. New York, NY: Columbia University Press.
- Burns, A.F. and Mitchell, W.C. (1946) *Measuring Business Cycles*. New York, NY: NBER.
- Canova, F. (1994) Detrending and turning points. *European Economic Review* 38, 614-623.
- Canova, F. (1998) Detrending and business cycle facts. *Journal of Monetary Economics* 41, 475-512.
- Christiano, L.J. and Fitzgerald, T.J. (2003) The band pass filter. *International Economic Review* 44(2), 435-465.
- Cleveland, W.S. and Devlin, S.J. (1988) Locally weighted regression: an approach to regression analysis by local fitting. *Journal of the American Statistical Association* 83, 596-610.
- Cleveland, W.S., Grosse, E. and Shyu, W.M. (1992) Local regression models. In Chapter 8 of *Statistical Models in S* (eds J.M. Chambers and T.J. Hastie), Wadsworth & Brooks/Cole.
- Cogley, T. and Nason, J.M. (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), 253-278.
- De Jong, R.M. and Sakarya, N. (2016) The econometrics of the Hodrick-Prescott filter. *Review of Economics and Statistics* 98(2), 310-317
- Durbin, J. and Koopman, S.J. (2001) *Time Series Analysis by State Space Methods*. Oxford University Press.

- Findley, D.F., Monsell, B.C., Bell, W.R., Otto, M.C. and Chen, B. (1998) New capabilities and methods of the X-12-ARIMA seasonal adjustment program. *Journal of Business and Economic Statistics* 16, 127-177.
- Gómez, V. (1999) Three equivalent methods for filtering nonstationary time series. *Journal of Business and Economic Statistics* 17, 109-116.
- Gómez, V. (2001) The use of Butterworth filters for trend and cycle estimation in economic time series. *Journal of Business and Economic Statistics* 19, 365-373.
- Gray, A.G. and Thomson, P.J. (1990) Invited commentary on 'STL: a seasonal-trend decomposition procedure based on loess' by R.B. Cleveland *et al.*, *Journal of Official Statistics* 6, 1, 47-54.
- Hall, V.B., Kim, K. and Buckle, R.A. (1998) Pacific Rim business cycle analysis: synchronisation and volatility. *New Zealand Economic Papers* 32(2), 129-159.
- Hall, V.B., and McDermott, C.J. (2016) Recessions and recoveries in New Zealand's post-Second World War business cycles. *New Zealand Economic Papers* 50(3), December, 261-280.
- Hall, V.B., Thomson, P.J., and McKelvie, S. (2017) On the robustness of stylised business cycle facts for contemporary New Zealand. *New Zealand Economic Papers* 51(3), December, 193-216.
- Hall, V.B. and Thomson, P.J. (2021) Does Hamilton's OLS Regression Provide a "better alternative" to the Hodrick-Prescott filter? A New Zealand Business Cycle Perspective. *Journal of Business Cycle Research*. <https://doi.org/10.1007/s41549-021-00059-1>.
- Hamilton, J.D. (2018) Why you should never use the Hodrick-Prescott filter. *Review of Economics and Statistics* 100(5), December, 831-843.
- Hamming, R.W. (1989) *Digital Filters (3rd ed)*. Prentice-Hall, Englewood Cliffs, NJ.
- Hannan, E.J. (1970) *Multiple Time Series Analysis*. Wiley, New York.
- Harding, D. and Pagan, A.R. (2002) Dissecting the cycle: a methodological investigation. *Journal of Monetary Economics* 49(2), 365-381.
- Harding, D. and Pagan, A.R. (2016) *The Econometric Analysis of Recurrent Events in Macroeconomics and Finance*. Princeton University Press.
- Harvey, A.C. (1989) *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge University Press.
- Harvey, A. C. (1997) Trends, cycles and autoregressions. *The Economic Journal* 107, 192-201.
- Harvey, A.C. (2010) The local quadratic trend model. *Journal of Forecasting* 29, 94-108.

- Harvey, A. C., and Jaeger, A. (1993) Detrending, stylised facts and the business cycle. *Journal of Applied Econometrics* 8, 231-247.
- Harvey, A. C., and Trimbur, T. (2003) General model-based filters for extracting cycles and trends in economic time series. *The Review of Economics and Statistics* 85, 244-255.
- Harvey, A.C., and Trimbur, T.M. (2008) Trend estimation and the Hodrick-Prescott filter. *Journal of the Japan Statistical Society* 38, 41-49.
- Henderson, R. (1916) Note on graduation by adjusted average. *Transactions of the Actuarial Society of America* 17, 43-48.
- Hodrick, R. J. (2020) An exploration of trend-cycle decomposition methodologies in simulated data. *NBER Working Paper* 26750, February.
- Hodrick, R.J. and Prescott, E.C. (1997) Postwar US business cycles: an empirical investigation. *Journal of Money, Credit, and Banking* 29 1-16.
- Kaiser, J.F. and R.W. Hamming (1977) Sharpening the response of a symmetric nonrecursive filter by multiple use of the same filter. *IEEE Transactions on Acoustics Speech and Signal Processing* ASSP-25, 415-422.
- Kaiser, R. and A. Maravall (1999) Estimation of the business cycle: a modified Hodrick-Prescott filter. *Spanish Economic Review* 1(2), 175-206.
- Kaiser, R. and A. Maravall (2012) *Measuring Business Cycles in Economic Time Series*. Lecture notes in statistics 154. New York, NY: Springer Science & Business Media.
- Kenny, P.B. and Durbin, J. (1982) Local trend estimation and seasonal adjustment of economic and social time series. *Journal of the Royal Statistical Society: Series A (General)* 145, 1-28.
- Kim, K., Buckle, R.A. and Hall, V.B. (1994) Key features of New Zealand business cycles. *Economic Record* 70(208), 56-72.
- Kim, K., Buckle, R.A. and Hall, V.B. (1995) Dating New Zealand business cycles. *New Zealand Economic Papers* 29, 143-171.
- King, R. and Rebelo, S.T. (1993) Low frequency filtering and real business cycles. *Journal of Economic Dynamics and Control* 17, 207-231.
- Kitagawa, G. and Gersch, W. (1996) *Smoothness Priors Analysis of Time Series*. Springer-Verlag, New York.
- Kulish, M. and Pagan, A. (2021) Turning point and oscillatory cycles: concepts, measurement, and use. *Journal of Economic Surveys* 1-30.

- Lienert, A. (2018) An introduction to Matai - The Treasury's new macro-economic forecasting model. Paper presented to New Zealand Association of Economists' Annual Conference, June 2018, <https://www.nzae.org.nz/events/nzae-conference-2018/2018-conference-papers/>
- Macaulay, F. (1931) *The Smoothing of Time Series*. National Bureau of Economic Research, New York.
- McCaw, S. (2007) Stylised facts about New Zealand business cycles. Discussion Paper No. DP2007/04. Wellington: Reserve Bank of New Zealand.
- McElroy, T. (2008) Exact formulas for the Hodrick-Prescott filter. *Econometrics Journal* 11, 209-217.
- McKelvie, S.D. and Hall, V.B. (2012) Stylised facts for New Zealand business cycles: a post-1987 perspective. *School of Economics and Finance Working Paper* 12/2012, Victoria University of Wellington, <http://researcharchive.vuw.ac.nz/handle/10063/2364>
- Mise, E., Kim, T-H and Newbold, P. (2005) On suboptimality of the Hodrick-Prescott filter at time series endpoints. *Journal of Macroeconomics* 27(1), 53-67.
- Nelson, C.R. and Kang, H. (1981) Spurious periodicity in inappropriately detrended time series. *Econometrica* 741-751.
- Nelson, C.R. and Plosser, C.R. (1982) Trends and random walks in macroeconomic time series: some evidence and implications. *Journal of Monetary Economics* 10(2), 139-162.
- Newey, W.K. and West, K.D. (1987) A simple positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica* 55, 703-708.
- Pagan, A.R. (1997) Towards an understanding of some business cycle characteristics. *Australian Economic Review* 30, 1-15.
- Phillips, P.C.B. and Jin, S. (2021) Business cycles, trend elimination, and the HP filter. *International Economic Review* 62(2), 469-520.
- Phillips, P.C.B., and Shi, Z. (2021) Boosting: why you can use the HP filter. *International Economic Review* 62(2), 521-570.
- R Development Core Team (2004) *R: A language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna. <http://www.R-project.org>
- Stock, J.H., and Watson, M.W. (1999) Business cycle fluctuations in US macroeconomic time series. *Handbook of Macroeconomics* 1, 3-64.
- Tukey, J.W. (1977) *Exploratory Data Analysis*. Addison-Wesley, Reading, PA.
- Whittaker, E.T. (1923) On a new method of graduation. *Proceedings of the Edinburgh Mathematical Society* 41, 63-75.

Zarnowitz, V. (1992) *Business cycles: Theory, History, Indicators, and Forecasting. Studies in Business Cycles* 27. University of Chicago Press for the NBER, Chicago, IL.