



Does Hamilton's OLS Regression Provide a "better alternative" to the Hodrick-Prescott Filter? A New Zealand Business Cycle Perspective

Viv B. Hall^{1,2} · Peter Thomson³

Received: 12 October 2020 / Accepted: 17 September 2021 / Published online: 2 November 2021
© The Author(s), under exclusive licence to Springer Nature Switzerland AG 2021

Abstract

Within a New Zealand business cycle context, we assess whether Hamilton's (H84) OLS regression methodology produces stylised business cycle facts which are materially different from the Hodrick–Prescott (HP) and Baxter–King (BK) measures, and whether using the H84 predictor for forecast-extension improves the HP filter's properties at the ends of series. Stylised business cycle facts were computed for a set of key New Zealand macroeconomic variables. In general, H84 produces greater volatilities and less credible trend movements during key economic periods than either HP or BK, and so for this purpose there is no material advantage in using H84 over HP or BK. At the ends of series, we evaluate the performance of the forecast-extended HP filter for three representative business cycle environments. The forecast-extension methods compared include the H84 predictor, the informed forecasts of three leading New Zealand economic agencies, two methods based on models of past data, and the HP filter with no extension. As expected, the better the forecast-extension the more accurate the HP filter at the ends of series and, as reported elsewhere in the literature, the HP filter with no extension performed poorly. However, in all cases considered the H84 predictor performed worst.

Keywords Hamilton regression filter · Stylised business cycle facts · New Zealand · Ends of series

JEL Classification E32 · E37 · C10 · G01

✉ Viv B. Hall
viv.hall@vuw.ac.nz

Peter Thomson
peter@statsresearch.co.nz

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

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

³ Statistics Research Associates, Wellington, New Zealand

1 Introduction

Hamilton (2018) makes a case for why you should never use the Hodrick-Prescott filter (Hodrick & Prescott, 1997), with his key arguments being the following.

- (a) The Hodrick-Prescott (HP) filter introduces spurious dynamic relations that have no basis in the underlying data-generating process.
- (b) Filtered values at the end of the sample are very different from those in the middle and are also characterised by spurious dynamics.
- (c) A statistical formalization of the problem typically produces values for the smoothing parameter vastly at odds with common practice.
- (d) There is a better alternative.

His better alternative is to use the regression of a variable at date t on the four most recent values as of date $t-h$ since this would achieve “... *all the objectives sought by users of the HP filter with none of its drawbacks*”.

Hamilton provides illustrative empirical results for a long-run quarterly U.S. total employment series (1947q1–2016q2), and for quarterly GDP and GDP component series of similar length. He uses his proposed regression method (H84) and his companion 8-lag difference method (H8). The regression can also be used to project forward eight more quarters of observations. Growth cycle volatility and contemporaneous cross-correlation results are provided (Hamilton 2018, Table 2), and benchmarked against random walk results, but perhaps surprisingly not against results from using the HP filter. Nor does he provide standard errors or non-contemporaneous cross-correlation results and there is also no explicit assessment of the economic circumstances in which end-point issues might be empirically concerning.

An increasing number of studies have provided theoretical and empirical evaluations of Hamilton’s proposed methodology. See, for example, Phillips and Shi (2021), Hodrick (2020), Jönsson (2020a, 2020b), Quast and Wolters (2020), Schüler (2018), and Drehmann and Yetman (2018) among others.

Phillips and Shi (2021) provide detailed responses to Hamilton’s (2018) critique of the HP filter, and analyse the performance of Hamilton’s regression approach relative to their boosted HP filter (bHP). Their findings show a clear preference for the bHP filter over the Hamilton regression, and they also conclude that the HP filter may continue to be used as a helpful empirical device for the estimation of trends and cycles.

Hodrick (2020) has used simulations approximating the natural logarithm of U.S. real GDP to examine the properties of the HP filter and Baxter-King filter (Baxter & King, 1999) (BK), relative to those from Hamilton’s H84 filter. He finds that the HP and BK filters are quite similar, that the H84 filter performs much better in environments with straightforward first difference stationary time series, and that the reverse is true for more complex time series. This leads him to conclude that the choice of methodology might depend on one’s priors about the nature of the series. But for the purpose of developing stylised business cycle

facts to assist in building and evaluating growth cycle models, he has expressed a preference to use the HP and BK methods.¹

Jönsson (2020a) has assessed the extent to which HP and Hamilton filters introduce spurious dynamics to a business cycle's cyclical component, and report that similar dynamics can be found in the cyclical component of HP-filtered and Hamilton-filtered series. The paper takes no stand on which of the two methods should be used when decomposing a series into trend and cycle components, and concludes that choosing between the two filters may turn out to be harder than at first thought. Jönsson (2020b) compares the HP and Hamilton filters with respect to real-time stability in US GDP gap estimation and finds that the Hamilton filter outperforms the HP filter when it comes to real-time revisions. The source of the inferior performance of the HP filter is that trend and cycle estimates close to the end of the sample are revised to a large extent as more data are added to the series, a finding also documented by other authors.

For U.S. log GDP and credit-to-GDP data, Schüler (2018) has compared the cyclical properties of Hamilton's regression filter with those from the HP filter. Overall, he finds that while Hamilton's filter is not subject to the same drawbacks as the HP filter, it too reflects ad hoc underlying assumptions. Specifically, he singles out the two-year regression filter for excluding two-year cycles and emphasising cycles which are longer than typical business cycle fluctuations, thereby being at odds with stylised business cycle facts, such as the one-year duration of a typical recession.

Drehmann and Yetman (2018) have assessed whether an HP trend or a Hamilton linear projection of the credit-to-GDP gap performed the better in providing an early warning indicator for crises. While acknowledging that it is an empirical question as to which of a range of measures performs best on this question, they also find that no other gap outperforms their baseline measure (a one-sided HP filter with smoothing parameter $\lambda = 400,000$). Further, they find that credit gaps based on linear projections in real time perform poorly.

Against the above background and for New Zealand data, we assess whether Hamilton's H84 filter has provided a "better alternative" to the HP filter. In particular, we evaluate comparative performance in two areas. In the body of the series we compare the stylised business cycle facts produced by the H84 filter relative to those from the HP filter and also the BK filter. The latter provides direct estimates of the business cycle based on the definitions of Burns and Mitchell (1946). At the ends of series, we compare the relative performance of the HP filter in the case where the series has been augmented with forecasts (HP forecast-extension) from the H84 regression and other forecast-extension methods. Our stylised facts evaluation is for a wider set of key macroeconomic model variables than the real GDP, output gap and credit-to-GDP gap variables investigated by others, and both evaluations use

¹ On the importance of developing important empirical regularities/stylised facts as benchmarks for assessing theoretical macroeconomic models and guiding macroeconomic policy, see Stock and Watson (1999).

data from a small, open economy rather than from the considerably larger economies examined by others.

Specifically, we address two questions.

- For post-1987q2 New Zealand, does Hamilton's H84 filter produce business cycle volatility and bivariate cross-correlation measures which are materially different from those obtained using the HP and BK filters?
- At the ends of series, does HP forecast-extension using the Hamilton H84 predictor perform better than other forecast-extension methods, including the informed forecasts of three leading New Zealand economic agencies, two methods based on models of past data, and the HP filter with no extension?

The latter question is investigated for three representative end-point environments: New Zealand's post-2009q1 business cycle expansion path, and two business cycle turning point periods encompassing the peak and trough associated with New Zealand's five quarter classical Global Financial Crisis (GFC) recession 2008q1–2009q1.² We are not aware of other evaluations which have focussed directly on end-of series issues at business cycle turning point periods.

Section 2 provides a brief description of our methodological framework. Empirical results are presented in Sects. 3 and 4, and Sect. 5 concludes.

2 Methodological Framework

Consider a non-seasonal quarterly economic time series x_t , possibly log transformed, and assume that x_t admits the additive decomposition

$$x_t = g_t + c_t \quad (1)$$

where g_t is an unobserved or hidden trend and c_t is the deviation from the trend. The decomposition and its conceptual components are identified by assuming that g_t is smooth and follows the secular general movement of the time series concerned, whereas c_t reflects shorter-term fluctuations and cyclical behaviour not accounted for by the trend. In essence, g_t is the base mean level around which shorter-term deviations, such as the business cycle, are estimated.

Typically, g_t is estimated by a linear trend filter of the form

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

with the trend deviations estimated by

² For investigative work on the role of an HP1600 filter for the U.S. GFC period, see Phillips and Jin (2021).

$$\hat{c}_t = x_t - \hat{g}_t = \sum_s \tilde{w}_t(s)x_{t-s}$$

where $\tilde{w}_t(s) = -w_t(s)$ ($s \neq 0$) and $\tilde{w}_t(0) = 1 - w_t(0)$. The filter weights $w_t(s)$ can be time-varying or time-invariant where $w_t(s) = w(s)$. Many filters used in business cycle analysis can be put into this general form including the Hodrick-Prescott filter (Hodrick & Prescott, 1997) where the $w_t(s)$ are time-varying, the Baxter-King filter (Baxter & King, 1999) which directly estimates the business cycle following the definitions proposed by Burns and Mitchell (1946), and simple moving average trend filters with time-invariant weights. The latter include the Henderson filters (Henderson, 1916) used by the seasonal adjustment procedure X-12-ARIMA (Findley et al., 1998) to estimate the so-called trend-cycle which includes the business cycle.

The empirical modelling framework (1) has a long history in official statistics and economic analysis going back to Macaulay (1931) if not earlier. It forms the basis for the trend-seasonal-irregular decomposition procedure X-12-ARIMA (a development of X-11 and X-11-ARIMA) which remains the dominant seasonal adjustment procedure used by official statistical agencies around the world, despite many attempts to supplant it by modern parametric dynamic models and methods. The use of the HP filter in economic time series analysis appears to be following a similar and parallel development, i.e. it is widely used in practice, but challenged by academic researchers who seek to replace it by a suitable dynamic model or equivalent.

While not wishing to rehearse all the arguments for and against the use of this empirical modelling framework, primary reasons for the popularity of X-12-ARIMA and the HP filter are that they:

- provide a *conceptually simple* framework that is readily understood by analysts and users alike;
- enjoy *broad consensus* with the methods widely used in practice and found useful by the international official statistics and applied economics communities;
- adopt a *common standard method* used across broad classes and collections of time series both nationally and internationally with little data-dependent fitting involved.

As a consequence, these *omnibus methods* have, for the most part, provided useful descriptions of individual economic time series and informative national and international economic comparisons across series, over a long period of time (post WW2).

On the other hand, all trend estimation methods involving moving average filters:

- (a) necessarily *modify the dynamic structure* of the original time series, sometimes significantly;
- (b) only provide *stable historical trend estimates in the body* of the series;
- (c) provide *more volatile trend estimates at the ends* of series than those in the body since they are subject to revision as more data values are added to the series;
- (d) *cannot match the forecasting performance of individual dynamic time series models* identified and fitted to each series.

Of these deficiencies, arguably the greatest is (c). In the case of seasonal adjustment, considerable effort has been invested into minimising the revisions of trend estimates at the ends of series (see Gray & Thomson, 2002, for example). Here it is known (see Geweke, 1978, Burman, 1980, and Pierce, 1980) that the optimal strategy is to augment the time series with optimal forecasts and then apply the moving average trend filter to the augmented time series. Such *forecast-extension* methods take the best of both modelling frameworks. Dynamic time series models can be used to provide out-of-sample forecasts that largely eliminate deficiency (d) and provide more stable estimates of the desired empirical trend at the ends of series which directly addresses deficiency (c).

Kaiser and Maravall (1999, 2012) would appear to be the first to propose and implement forecast-extension procedures for the HP trend filter within a business cycle framework. As they acknowledge, their procedures are derived from the Statistics Canada seasonal adjustment program X-11-ARIMA (Dagum, 1980) which implemented forecast extension procedures for the X-11 trend-seasonal-irregular decomposition procedure that underpins X-12-ARIMA used by most official statistical agencies worldwide. Forecast-extension has also been used more widely in business cycle analysis: Christiano and Fitzgerald (2003) use forecast-extension with the Baxter-King filter, Garratt et al. (2008) have used forecast extension for estimating the output gap, and a number of economic agencies such as the European Commission have used such methods (see Kaiser & Maravall, 2012, for example, and the references therein). Nevertheless, forecast-extension methods for the HP trend filter have yet to be as routinely and universally adopted as they have within the official statistics community and trend-seasonal-irregular decomposition procedures such as X-12-ARIMA.

In this paper we follow the consensus and *assume that the HP filter provides a useful, and economically meaningful, empirical trend in the body of the series*. As a consequence, the HP trend deviations are deemed to provide reliable estimates of the business cycle in the body of the series. At the ends of the series we consider a variety of methods of forecast extension for trend estimation at key dates in the evolution of our New Zealand time series. The quality of the various forecast extensions is compared, including the case of no extension (using the HP filter trend estimates at the ends).

In Sect. 2.1 we briefly review some technical aspects of the HP filter, in Sect. 2.2 we discuss the criticisms of Hamilton (2018) within the above framework and, in Sect. 2.3, we provide further details on the forecast extension procedures we have adopted.

2.1 Comments on the HP Filter

The HP filter is an empirical trend filter whose trend \hat{g}_t minimises the criterion

$$F + \lambda S = \sum_t (x_t - \hat{g}_t)^2 + \lambda \sum_t (\Delta^2 \hat{g}_t)^2 \quad (3)$$

where Δ is the first difference operator $\Delta x_t = x_t - x_{t-1}$ and λ is a trade-off parameter balancing the fidelity F of \hat{g}_t to the data x_t with the smoothness S of \hat{g}_t . The smaller F is the closer \hat{g}_t follows the data, and the smaller S is the closer $\Delta^2 \hat{g}_t$ is to zero and the closer \hat{g}_t is to a simple linear trend. For most quarterly applications the standard choice of λ is 1600 although this value can be tuned, if necessary, to better reflect the balance of smoothness and fidelity desired.

As noted earlier, the HP filter is a linear trend filter of the form (2) with time-varying weights. De Jong and Sakarya (2016) show that, while the weights at the ends of the series (defining the HP end filters) are always time-varying, those in the body of the series are time-invariant provided the time series is long enough (around 50 quarters or greater for quarterly data and λ equal to 1600). These time-invariant weights define the *central HP filter*, which is a symmetric, non-negative definite, moving average filter of the form (2) whose weights are given by

$$w(s) = \frac{1}{\alpha} \sin(|s|\theta + \varphi) \rho^{|s|} \tag{4}$$

where

$$\rho = \frac{1}{\sqrt{1 + \delta} + \sqrt{\delta}}, \quad \alpha = \sqrt{\lambda \left(\rho^2 + \frac{1}{\rho^2} - 2 \cos 2\theta \right)}$$

and

$$\delta = \frac{1 + \sqrt{1 + 16\lambda}}{8\lambda}, \quad \theta = \tan^{-1} \left(\frac{1}{2\sqrt{\lambda}} \frac{1 + \rho^2}{1 - \rho^2} \right), \quad \varphi = \tan^{-1} \left(2\sqrt{\lambda} (\tan \theta)^2 \right).$$

These are simplified versions of the formulae given in McElroy (2008) and De Jong and Sakarya (2016). In our case $\lambda = 1600$ and ρ takes the value 0.8941 so the weights $w(s)$ decay slowly to zero as $|s|$ increases. It is the *central HP filter* that we adopt as our target trend filter.

Given observations $x = (x_1, x_2, \dots, x_T)'$, minimising (3) yields the solution

$$\hat{g} = Hx, H = (I + \lambda D'D)^{-1} \tag{5}$$

where $\hat{g} = (\hat{g}_1, \hat{g}_2, \dots, \hat{g}_T)'$ 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 time-varying weights of the HP filter with the central rows corresponding to the time-invariant weights (4) and the remaining rows corresponding to the asymmetric HP end filters.

Alternative methods for computing \hat{g} are available. 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

$$\Delta^2 g_t = \varepsilon_t \tag{6}$$

and ε_t, c_t being mutually independent Gaussian white noise processes. Under these assumptions \hat{g} can be computed using the Kalman filter and smoother (see Harvey & 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 (4) to the extended series. Gomez (1999) shows that these three procedures are equivalent.

Although we adopt the central HP filter as our target filter, irrespective of the data generating process that x_t or its components g_t, c_t might follow, it is noted 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. Here we follow Mise et al (2005) who build on King and Rebelo (1993). Let x_t be given by (1) with components that follow the models

$$\Delta^2 g_t = C(L)\varepsilon_t, \quad c_t = C(L)\eta_t \quad (7)$$

or

$$\Delta g_t = C(L)\varepsilon_t, \quad c_t = C(L)(1 - L)\eta_t \quad (8)$$

where L is the backward shift operator, $C(z) = \sum_{j=0}^{\infty} \alpha_j z^j$ ($\alpha_0 = 1, \sum_{j=0}^{\infty} \alpha_j^2 < \infty$) is non-zero for $|z| \leq 1$ and ε_t, η_t are mutually independent Gaussian white noise processes. *In both cases the central HP filter generates the optimal estimator of g_t in the body of the series.* Moreover, the reduced model for x_t is given by

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

where $d = 2$ for model (7), $d = 1$ for model (8), ρ and θ are given by (4), and u_t is Gaussian white noise. Judicious choice of $C(L)$ allows for a more general class of data generation processes than just the case $C(L) = 1$ given in (6). In particular, x_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 and, in turn, imply non-trivial models for the cycle c_t .

2.2 Hamilton Critique

Hamilton (2018) argues against the routine use of the HP filter in business cycle analysis and suggests an alternative procedure. Some of his reasons, such as (c) in Sect. 1, relate specifically to the HP filter itself. However most apply to the generic model framework (1). In this sense Hamilton (2018) can be seen as a more general argument against the use of structural time series models such as (1) for business cycle analysis. This is a more serious challenge, especially given the wide-spread and long-standing use of the empirical framework (1) and the more recent development of parametric structural time series models based on (1) that are exemplified in the literature by Akaike (1980), Harvey (1989), Kitagawa and Gersch (1996) and

Durbin and Koopman (2001) among many others. Like the HP filter, the latter have their genesis in the much earlier work of Whittaker (1923) and Henderson (1924).

What is the alternative procedure proposed in Hamilton (2018)? In essence Hamilton eschews the model framework (1) with its unobserved trend g_t and trend deviation c_t . Instead, he argues that business cycle information can be gleaned directly from the time series x_t using suitably chosen OLS prediction errors. In particular, he fits the (auto) regression forecasting model

$$x_t = \beta_0 + \beta_1 x_{t-h} + \beta_2 x_{t-h-1} + \beta_3 x_{t-h-2} + \beta_4 x_{t-h-3} + v_t \quad (9)$$

by OLS to get forecasts and prediction errors given by

$$\hat{x}_t = \hat{\beta}_0 + \hat{\beta}_1 x_{t-h} + \hat{\beta}_2 x_{t-h-1} + \hat{\beta}_3 x_{t-h-2} + \hat{\beta}_4 x_{t-h-3}, \hat{v}_t = x_t - \hat{x}_t \quad (10)$$

respectively. Here the $\hat{\beta}_j$ are the OLS regression coefficients determined from all the data and the forecast horizon h is recommended to be $h = 8$ quarters (2 years) ahead for quarterly time series as is the case considered here. Hamilton (2018) shows that \hat{x}_t is a robust (largely model independent) predictor that yields consistent forecasts of x_t for a wide variety of nonstationary processes. In practice it would appear that this predictor is close to

$$\hat{x}_t = \tilde{\beta}_0 + x_{t-h}$$

where, as before, $\tilde{\beta}_0$ is determined by OLS regression and so any business cycle analysis would be undertaken on the mean-corrected lag- h differences $x_t - x_{t-h}$.

Hamilton's arguments for using these prediction errors for business cycle analysis instead of a more conventional analysis based on trend deviations seem less convincing. The prediction errors \hat{v}_t measure the departure of x_t from its expected level, a forecast determined from data up to 8 quarters (2 years) earlier. By contrast, the trend deviations \hat{c}_t measure the departure of x_t from its local level \hat{g}_t where the latter is either a trend determined largely by consensus or, in the case of parametric structural time series models, the expected value of the hidden trend given all the available data. The two estimates of level are quite different, both conceptually and in terms of their time series properties. Whereas \hat{g}_t is expected to be smooth, this is not necessarily the case for \hat{x}_t which will be much more variable (typically of the same order as the original series x_t). The use of a common $h=8$ quarter forecast horizon is also arguable since some economic variables or jurisdictions may need different horizons to produce useful results. These issues are addressed by Quast and Wolters (2020) who consider a modified Hamilton filter which averages the prediction errors from the Hamilton regressions (9) for $h=4$ to $h=12$. This produces a much smoother Hamilton trend which they argue gives better measures of real-time output gaps for the U.S., U.K. and Germany than those produced by the HP filter or the bandpass filter of Christiano and Fitzgerald (2003).

In Sect. 3, we apply the Hamilton (2018) methodology to establish stylised business cycle facts for our set of key New Zealand macroeconomic variables. In Sect. 4 we use the Hamilton robust predictor \hat{x}_t for forecast extension and compare it to other contenders.

2.3 HP Forecast Extension Procedures Adopted

We consider the performance of forecast extension of the HP filter using predictions published by two leading New Zealand public sector economic forecasters, the Reserve Bank of New Zealand (RBNZ) and the New Zealand Treasury (Treasury), and a prominent private sector economic forecasting entity, the New Zealand Institute of Economic Research (NZIER). The predictions of these three institutions have been chosen because their quarterly predictions are both publicly available and have been published for the three historical business cycle-related sample periods we investigate. It is to be expected that these forecasts are the result of extensive modelling and other, possibly nonlinear, procedures which are informed by information about future events that is not encapsulated in just the past data of the series concerned. These institutional predictions will be termed *informed forecasts*.

In Sect. 4.2 the performances of the informed forecasts are benchmarked against a number of simple forecast-extension procedures based only on past observations available at the time. A number of candidate benchmark procedures could be chosen. Here we have chosen four simple procedures; the HP filter with no-extension, forecasting using the Hamilton predictor, ARIMA forecast extension, and forecasting using a simple adaptive random walk model.

2.3.1 HP with No Extension

Kaiser and Maravall (1999, 2012) show that the HP filter defined by (3) and (5) can be calculated by applying the central HP filter to the original series augmented by forecasts from model (7) in the simplest case where $C(L) = 1$ and g_t follows (6). Mise et al. (2005) show that, for models (7) and (8) where the HP filter is optimal in the body of the series, forecast-extension using these models always performs better than the HP end filters, except for the case when $C(L) = 1$. Thus, the HP filter is consistent with forecast-extension using forecasts only from this special model where g_t follows (6). Since this simple special model is unlikely to hold in practice, other forecast-extension methods are likely to be more efficient and better minimise revisions.

2.3.2 Hamilton Robust Predictor

Here we adopt the Hamilton method as a robust auto-regressive predictor that yields consistent forecasts of x_t for a wide variety of nonstationary processes. Use of the Hamilton methodology for robust forecast-extension does not seem to have been considered elsewhere in the literature.

2.3.3 ARIMA Forecast Extension

Following in the footsteps of Dagum (1980), Kaiser and Maravall (1999, 2012) and many others, we augment x_t with forecasts generated by a best-fitting ARIMA model. The results of Mise et al. (2005) and the discussion following models (7) and (8) in Sect. 2.1, suggest that this approach should reduce revision mean-squared

errors in practice, particularly at the ends of series. However, the extent of any reduction will depend in no small part on how well the fitted model captures the dynamics of the data.

2.3.4 Naïve Predictor

Here we consider forecast extension using the simple random walk model

$$x_t = x_{t-1} + \delta_t + \varepsilon_t \quad (11)$$

where δ_t measures a smooth, slowly evolving, drift and ε_t is stationary noise. In effect, this model assumes that the trend g_t in (1) is locally linear. Here we have chosen to estimate δ_t as the median of the first differences $x_t - x_{t-1}$ over the most recent 8 quarters (2 years), but other simple robust location estimators could also be chosen and applied over alternative local time windows. In the case of log data, note that this estimator is just the median quarterly growth rate of the untransformed time series over the last two years. This simple robust predictor provides a suitable benchmark forecast extension which will be called the *naive predictor*. It is an $I(1)$ model that is conceptually simple, easy to implement and understand, adaptive and robust. It is similar in philosophy to the rolling $IMA(1,1)$ model used in Stock and Watson (2007), but more adaptive.

3 Empirical Results: H84, HP and BK Stylised Business Cycle Facts

We compute stylised business cycle facts for a set of key New Zealand macroeconomic variables typically included in theoretical or empirical macroeconomic models of small open economies. Quarterly, seasonally adjusted data have been sourced from Statistics New Zealand (SNZ), the RBNZ and Treasury. They 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).

In addition to computing and comparing H84 and HP stylised business cycle facts, we also compare these with the stylised business cycle facts obtained using the Baxter-King (BK) filter which has a different rationalisation to either the H84 or HP filters. The BK filter is directly based on a band-pass filter (see Baxter & King, 1999; Christiano & Fitzgerald, 2003, for example) and is closer in spirit to the Burns and Mitchell (1946) paradigm than either the H84 or HP filters.

In this and subsequent sections, any reference to the HP filter assumes that the trade-off parameter chosen is $\lambda = 1600$, any reference to the Hamilton filter or predictor (without qualification) refers to the Hamilton H84 regression filter described in Sect. 2.2 with $h=8$, and any reference to the BK filter assumes that it is the 25 point moving average band-pass filter recommended in Baxter and King (1999) which is designed to pass frequency components with periods between 6 and 32

quarters. All computations and graphical analysis were carried out in the R statistical environment (R Development Core Team, 2004).

3.1 Results

Stylised business cycle facts for the cycles (trend deviations) estimated by the Hamilton H84, HP and BK filters are compared for 13 New Zealand macroeconomic series over the period 1987q2 to 2019q4, the latter quarter being the business cycle peak immediately prior to New Zealand's Covid-19 recession. This analysis complements and extends that given in Hall et al. (2017) who compute stylised business cycle facts for a more extensive set of series over the shorter period 1987q2 to 2015q3 using the HP, BK, Christiano and Fitzgerald (2003) and *loess* (local regression) trend filtering methods, but not the Hamilton H84 filter. We use the same methodology as Hall et al. (2017) where further details are given on the underpinning theory, including standard error calculations. We are not aware of any other study assessing use of the Hamilton H84 method for New Zealand business cycle analysis.

Earlier studies of stylised business cycle facts for New Zealand macroeconomic data include Kim et al. (1994), McCaw (2007) as well as the more recent Hall et al. (2017). In common with experience reported elsewhere in the world, the HP filter remains a commonly used technique for business cycle and economic analysis in New Zealand. It is typically the method against which other methods are compared and contrasted. This is the case for our study where, as noted earlier, we follow the consensus and assume that the HP filter provides a useful, and economically meaningful, empirical trend *in the body of the series* where its trend deviations provide reliable estimates of the business cycle.

Hamilton H84, HP and BK volatility, persistence and cross correlation statistics are presented in Tables 1 and 2.³ For each of our 13 data series the H84 volatility always exceeds the HP or BK volatility where these differences are, for the most part, statistically significant. On the other hand, while the HP volatility typically exceeds BK volatility, these differences are always not statistically significant. In terms of *persistence* (measured by the lag 1 autocorrelation), BK cycles are always more persistent than either the HP or H84 cycles, and the H84 cycle is almost always more persistent than the HP cycle. However, while the HP and BK persistence measures are almost always not significantly different from the H84 persistence measure, significant differences are more typical between the HP and BK persistence measures. The *cross correlations* of the H84, HP and BK cycles for each of our data series with the log GDPE cycle show broad agreement (Table 2). All are of the same sign, much the same magnitude (their differences are not significantly different), and they have significant non-contemporaneous correlation at the same lag.

³ Volatility, persistence and cross correlation results for Hamilton's H8 procedure are not presented in Tables 1 and 2, since the differences from H84 are small and their respective trends and cycles are very similar.

Table 1 Stylised business cycle facts, 1987q2—2019q4: comparative volatilities and persistence

Variable x	Volatility			Persistence		
	H84	HP	BK	H84	HP	BK
Log GDP (production)	3.05 (0.52)	1.18 (0.21)	1.09 (0.35)	0.85 (0.05)	0.85 (0.04)	0.94 (0.03)
Log GDP (expenditure)	3.05 (0.55)	1.30 (0.17)	1.18 (0.32)	0.86 (0.04)	0.72 (0.06)	0.91 (0.04)
Log Consumption (private)	3.12 (0.66)	1.48 (0.17)	1.31 (0.31)	0.86 (0.04)	0.79 (0.05)	0.94 (0.03)
Log Gross Fixed Capital Formation	11.72 (2.42)	5.46 (0.72)	5.55 (1.23)	0.88 (0.04)	0.76 (0.06)	0.93 (0.03)
Log Investment (residential)	16.13 (2.85)	7.97 (0.88)	8.16 (1.42)	0.89 (0.04)	0.79 (0.05)	0.92 (0.04)
Log Govt. Consumption expend.	3.21 (0.60)	1.29 (0.15)	1.22 (0.29)	0.85 (0.05)	0.47 (0.08)	0.94 (0.04)
Net Exports Share (%)*	2.22 (0.39)	1.23 (0.15)	1.09 (0.25)	0.87 (0.04)	0.70 (0.06)	0.90 (0.04)
Log Imports goods & services*	7.04 (1.66)	4.16 (0.58)	4.00 (0.89)	0.87 (0.04)	0.77 (0.05)	0.90 (0.04)
Log Employment	2.81 (0.38)	1.30 (0.17)	1.10 (0.28)	0.84 (0.05)	0.90 (0.04)	0.96 (0.03)
Unemployment (%)	1.22 (0.25)	0.55 (0.13)	0.58 (0.20)	0.88 (0.04)	0.87 (0.04)	0.95 (0.03)
CPI Inflation (annual % change)**	1.16 (0.17)	0.86 (0.10)	0.85 (0.14)	0.87 (0.05)	0.76 (0.06)	0.87 (0.05)
CPI Nontradables (ann. % change)***	0.99 (0.12)	0.81 (0.16)	0.67 (0.19)	0.84 (0.05)	0.76 (0.06)	0.88 (0.05)
Real 90-day Bank Bill (%)**	1.60 (0.24)	0.87 (0.12)	0.84 (0.17)	0.85 (0.05)	0.70 (0.06)	0.88 (0.04)

H84 statistics are for standard Hamilton OLS regression estimates; HP statistics are for Hodrick-Prescott estimates with standard assumption for quarterly data of $\lambda=1600$; BK statistics are for Baxter-King estimates using the standard assumption of cycles between 6 and 32 quarters

Volatility is % standard deviation; persistence is represented by first order serial correlation

Numbers in parentheses for volatility and short-term persistence are robustly estimated standard errors

* SNZ National Accounts series adjusted for frigate purchases recorded in 1997q2 and 1999q4

** Sample period 1987q3–2019q4

*** Sample period 1989q1–2019q4

For government consumption expenditure, CPI inflation and CPI non-tradables, all three cycles judge the cross-correlation with log GDPE as not significantly different from zero.

Comparative trend and trend deviation (cycle) paths are illustrated in Figs. 1, 2 and 3 for real expenditure-based gdp (gdpe), real residential investment (invres) and real gross fixed capital formation (gfcf). It is clear from the top panels in these Figures that, while the HP and BK trends are in close agreement and provide a good description of the local level of the series, the H84 trend does not. Rather, the H84 trend is close to a mean-corrected phase-shifted version of the series (a lag of 8 quarters) reflecting the discussion following (10). In particular, the phase-shifted H84 trends fail to capture the local levels of the series during periods of substantive level change such as those associated with the 1991–92 and GFC recessions.

Table 2 Stylised business cycle facts, 1987q2–2019q4: comparative cross correlations with log real GDP

Variable x	Contemporaneous Cross Correlation			Most significant Non-contemporaneous Cross Correlation		
	H84	HP	BK	H84	HP	BK
Log Consumption (private)	0.84 (0.19)	0.73 (0.16)	0.79 (0.21)	–	–	–
Log Gross Fixed Capital Formation	0.81 (0.17)	0.74 (0.15)	0.83 (0.19)	–	–	–
Log Investment (residential)	0.70 (0.15)	0.68 (0.13)	0.77 (0.14)	–	–	–
Log Govt. Consumption expend.	0.14 (0.13)	0.07 (0.11)	0.12 (0.13)	0.52 (0.14)	0.42 (0.13)	0.46 (0.14)
Net Exports Share (%) [*]	–0.33 (0.14)	–0.20 (0.09)	–0.32 (0.12)	–0.52 (0.16)	–0.47 (0.14)	–0.44 (0.14)
Log Imports goods & services [*]	0.60 (0.15)	0.43 (0.11)	0.55 (0.13)	0.70 (0.16)	0.54 (0.14)	0.58 (0.16)
Log Employment	0.64 (0.13)	0.52 (0.14)	0.68 (0.18)	0.63 (0.14)	0.55 (0.13)	0.75 (0.18)
Unemployment (%)	–0.60 (0.17)	–0.58 (0.16)	–0.71 (0.20)	–0.62 (0.17)	–0.59 (0.15)	–0.71 (0.19)
CPI Inflation (annual % change) ^{**}	0.10 (0.14)	0.13 (0.14)	0.19 (0.16)	0.44 (0.12)	0.46 (0.13)	0.56 (0.15)
CPI Nontradables (ann. % change) ^{***}	0.13 (0.12)	0.28 (0.16)	0.38 (0.21)	0.36 (0.13)	0.60 (0.16)	0.68 (0.21)
Real 90-day Bank Bill (%) ^{**}	0.26 (0.13)	0.32 (0.14)	0.38 (0.15)	0.33 (0.13)	0.41 (0.12)	0.52 (0.15)

H84 statistics are for standard Hamilton OLS regression estimates; HP statistics are for Hodrick-Prescott estimates with standard assumption for quarterly data of $\lambda = 1600$; BK statistics are for Baxter-King estimates using the standard assumption of cycles between 6 and 32 quarters

Numbers in parentheses are robustly estimated standard errors

"-" Denotes most significant cross correlation is Contemporaneous

^{*} SNZ National Accounts series adjusted for frigate purchases recorded in 1997q2 and 1999q4

^{**} Sample period 1987q3–2019q4

^{***} Sample period 1989q1–2019q4




Fig. 1 Logarithms of gdpe, trends, and trend deviations. **a** Top panel shows log gdpe (black), HP trend (red), H84 trend (green) and BK trend (blue). **b** Bottom panel shows log gdpe trend deviations for HP (red), H84 (green) and BK (blue)

Figures 1, 2 and 3 also show that the HP and BK trend deviations (the estimated cycles) are in good agreement, but both differ markedly from the Hamilton H84 estimated cycle. All three cycles show stationary behaviour about a zero mean, but the H84 cycle is much more volatile (larger standard deviation) than either the HP or BK cycles.

To better quantify the differences between any two competing cycles we consider the root mean squared difference (RMSD) between them. Here RMSD is used as a *dissimilarity measure*. It is readily shown that the mean squared difference (MSD) of two time series X_t and Y_t (the square of RMSD) follows the simple identity




Fig. 2 Logarithms of invres, trends, and trend deviations. **a** Top panel shows log invres (black), HP trend (red), H84 trend (green) and BK trend (blue). **b** Bottom panel shows log invres trend deviations for HP1600 (red), H84 (green) and BK (blue)

$$MSD = \frac{1}{T} \sum_{t=1}^T (X_t - Y_t)^2 = (\bar{X} - \bar{Y})^2 + (s_X - s_Y)^2 + 2s_X s_Y (1 - r_{XY}) \quad (12)$$

where $\bar{X}, \bar{Y}, s_X, s_Y$ and r_{XY} are the sample means, sample standard deviations and sample correlation of series X_t and Y_t with the last 3 measures having divisor T (series length) rather than $T-1$. Clearly X_t and Y_t are identical when RMSD or MSD is zero which can only occur when $\bar{X} = \bar{Y}, s_X = s_Y$ and $r_{XY} = 1$. The larger RMSD, the greater the dissimilarity of X_t and Y_t with each of the three non-negative terms in (12) measuring the respective contributions to MSD of the difference in the means,




Fig. 3 Logarithms of gfcf, trends, and trend deviations. **a** Top panel shows log gfcf (black), HP1600 trend (red), H84 trend (green) and BK trend (blue). **b** Bottom panel shows log gfcf trend deviations for HP1600 (red), H84 (green) and BK (blue)

Table 3 Dissimilarity measure RMSD between HP, H84 and BK estimated cycles for log GDPE, together with their means, standard deviations (SD) and correlations

	Mean	SD	Correlations			RMSD		
			HP	H84	BK	HP	H84	BK
HP	-0.001	0.014	1	0.649	0.911	0	0.027	0.006
H84	-0.000	0.033	0.649	1	0.611	0.027	0	0.028
BK	-0.002	0.013	0.911	0.611	1	0.006	0.028	0

the difference in the standard deviations, and the difference of the correlation from one (a measure of the difference in dynamics).

Table 3 shows an example of the values of the dissimilarity measure MSD between the HP, H84 and BK estimated cycles for log GDPE. The respective means, standard deviations and correlations are also given. For our 13 data series, the dissimilarity measure RMSD between the HP and H84 cycles or the BK and H84 cycles is always considerably larger (over twice, and typically over four times) than the RMSD between the HP and BK cycles. While the means of the three cycles are always very similar and close to zero as expected, the standard deviation for the H84 cycle is typically over twice that of either the HP or BK cycles which are very similar. In terms of the contributions of the components in (12) to the MSD between the HP and H84 cycles, the difference in means is negligible with the difference in standard deviations and the difference in dynamics (difference of the correlation from 1) each typically contributing around 50 per cent. For the MSD between the HP and BK cycles, the difference in the dynamics typically contributes over 90 per cent to MSD with the differing standard deviations under 10 per cent.

Others have also shown that H84 cycle volatilities for U.S. real GDP are typically over twice those of HP cycles (e.g. Hodrick, 2020; Schüler, 2018). We find that this is further the case for New Zealand real GDP when H84 volatilities are set against BK volatilities.⁴

3.2 Key Findings

Our key findings on stylised facts are as follows.

- H84 *trends* are subject to phase-shifts and generally fail to provide a good description of the local level of our series, unlike the HP and BK trends which are in close agreement.
- Business *cycles* estimated using the HP or BK filters are very similar by comparison to the markedly different cycles estimated using the H84 regression methodology.
- H84 cycles have considerably greater volatilities than their HP or BK counterparts.
- HP cycles are generally less persistent than the corresponding H84 cycles which, in turn, are generally less persistent than the corresponding BK cycles, although these differences are, for the most part, not statistically significant.
- Cross correlations of H84, HP and BK cycles with real GDPE show broad agreement. Any differences are not statistically significant.

Hence, primarily on the basis that H84 produces materially greater volatilities and less credible trend movements associated with H84's inherent phase-shift behaviour, particularly during key economic periods such as the 1991–92 and GFC

⁴ This finding is also the case for the longer-term Hall and McDermott (2016) New Zealand production-based GDP series, updated to 1947q2 to 2019q3, where the volatility from the H84 series is over twice that of the HP1600-filtered series (4.1 relative to 1.8, with commensurate standard errors 0.50 and 0.29).

recessions, we have a clear preference for measures of stylised business cycle facts produced by the HP or BK filters rather than those from Hamilton's H84 procedure.

For macroeconomic data from a small open economy such as New Zealand, we conclude that the HP filter remains our method of choice for business cycle analysis in the body of the series. With this established, Sect. 4 uses HP forecast extension to address the widely-acknowledged limitations of the HP filter at the ends of series.

4 Empirical Results: HP Forecast Extension Methods Compared

The objectives of this section are two-fold. The first objective is to assess the performance of the Hamilton filter as a robust predictor for HP forecast-extension and measure its performance against the other HP forecast-extension procedures given in Sect. 2.3. The second objective is to assess the performance of forecast extension using the informed forecasts provided by the RBNZ, Treasury and NZIER whose predictions might be expected to be influenced by additional information over and above the past data of the series concerned. This additional information should confer performance advantages over forecast-extension using just past data, especially where business cycle turning points are involved.

Further details on our data series and performance measures are given in Sect. 4.1. Results are presented in Sect. 4.2, and key findings are summarised in Sect. 4.3.

4.1 Data and Performance Measures Adopted

We consider log-transformed New Zealand quarterly real production-based gdp data (GDPP) post-1987q2 and focus on three illustrative historical periods and their associated data.

- **Period NTP** considers data to 2015q3 which falls in New Zealand's post-2009q1 classical business cycle expansion path which had no turning points until 2019q4.
- **Period TPP** concerns data to 2006q4 which is 4 quarters before the 2007q4 turning point peak of New Zealand's five quarter classical GFC recession 2007q4–2009q1.
- **Period TPT** concerns data to 2008q1 which is 4 quarters before the 2009q1 turning point trough of New Zealand's five quarter classical GFC recession 2007q4–2009q1.⁵

⁵ The TPP and TPT turning points, taken from Hall and McDermott (2016, Table 1) preferred classical real gdp turning points for New Zealand's post-1947q2 economy have been computed from the Bry and Boschan (1971) (BB) dating algorithm for quarterly data. Robustness of these turning points relative to those computed from Harding and Pagan's (2002) BBQ algorithm, along with other evaluations of robustness, is reported in Hall and McDermott (2016, Sect. 2.2). For 17 of 18 BB turning points, the BB and BBQ algorithms produced identical results, including for the peak at 2007q4 and the trough at 2009q1.

For each of the periods NTP, TPP and TPT, we consider quarterly GDPP time series data up to the given time point in the period and then evaluate the performance of the forecast-extended HP filters at the ends of these three series. In each case we use the data and forecasts available at that time (real-time forecasts). To assess the performance of the forecast extension procedures, more recent GDPP data (as of 2019q4) is used to augment the available data to provide ‘true’ values of GDPP for the period after each series end point. In particular, this data is used to provide the target historical estimates of the HP filter at the ends of series.

Periods TPP and TPT focus on the performance of the various forecast-extended HP filters at the ends of series in the important case of turning points, whereas period NTP has no turning points and presents fewer challenges. Results for the various forecast methods are unlikely to vary greatly for any NTP period chosen along an ongoing close-to-linear classical business cycle expansion path but are likely to differ in the neighbourhood of turning points. In the latter case, the results directly address the issue of which, if any, of the HP forecast extension methods significantly reduce revisions at the ends of series.

For period NTP the data is from 1987q2 to 2015q3 (SNZ release, December 2015) and is the data used by RBNZ, Treasury and NZIER for their forecasts from 2015q4. For period TPP, the data is from 1987q2 to 2006q4 (SNZ release, March 2007) and precedes the 2007q4 GFC business cycle peak by 4 quarters. For period TPT we use the data from 1987q2 to 2008q1 (SNZ release, June 2008) where this data precedes the 2009q1 GFC business cycle trough by 4 quarters. The TPP and TPT GDPP data sets were taken from the real-time data sets compiled by the RBNZ (see Sleeman, 2006). The specially compiled data set of real-time forecasts by the RBNZ, Treasury and NZIER have been sourced from the relevant publicly available Monetary Policy Statements (MPS), Treasury Budget and Half-year Economic and Fiscal Updates (BEFU/HYEFU), and NZIER Quarterly Predictions (QP) releases.

We have chosen to augment, or extend, each data set by forecasts over a *forecast window* of 8 quarters (two years) following the data’s last available quarterly observation. The HP filter is then applied to the extended data set to provide trend estimates over the times of the original data (the *data window*) as well as the forecast window. The trend estimates over the data window are the desired output of the *forecast-extended HP filter*. This means, for example, that the most recent trend value in the data window will be calculated by the HP end filter 8 quarters from the end of the forecast augmented data. Any gains in precision will depend on the quality of the forecasts and how closely this HP end filter agrees with the central HP filter.

The choice of an 8-quarter (two-year) forecast window needs more justification. In part this choice reflects expediency. The RBNZ publish quarterly forecasts up to 3 years ahead, NZIER up to 4 years ahead and Treasury up to 5 years ahead. However, Lees (2016), and Labbé and Pepper (2009) chose one-year and two-year ahead horizons for their comparisons of RBNZ and external forecaster performance which is consistent with our two-year forecast window. The Hamilton filter also provides a natural forecast over a two-year horizon. As already noted, with forecast extension there will always be a trade-off between forecast accuracy and trend volatility at the ends of series. Poor forecasts may well lead to greater trend volatility at the ends of series than the HP filter with no extension.

A further argument in favour of the 8-quarter forecast window relates to the difference between the HP end filters and the central HP filter that applies in the body of the series. The former are finite-window asymmetric approximations to the central HP filter with weights given by (4). If an HP end filter is a reasonable approximation to the central HP filter and we have accurate forecasts, then we would expect more accurate, less volatile, trend estimates at the ends of series. The following table shows the square root of the sum of squared differences (RSSD) between the weights of the central HP filter and the HP end filter located q quarters from the end of the series.

q	Quarters from end of series										
	0	4	8	12	16	20	24	28	32	36	40
RSSD	0.293	0.156	0.066	0.027	0.026	0.026	0.020	0.013	0.007	0.003	0.001

Note that the RSSD of the differences in filter weights is also a proxy measure of the root mean square difference between the outputs of the two filters. The RSSD is greatest, as expected, for the HP end filter at the end of the series where the RSSD is 0.293. It then falls off rapidly to 53%, 22% and 9% of the maximum RSSD, for HP end filters located at 4 quarters (1 year), 8 quarters (2 years) and 12 quarters (3 years) from the end of the series. The remaining RSSD are likely to lead to negligible trend differences in practice. Similar conclusions are reached by Kaiser and Maravall (1999, 2012) whose results are based on a more comprehensive analysis. These considerations provide further support for the use of the 8-quarter forecast window.

To assess the quality of the various HP forecast extension methods, including the case of no extension, we need to define a *target trend* and a suitable time interval (*assessment window*) over which measures of the size of deviations from the target trend are calculated. These measures include the mean deviation or bias, the root mean square error (RMSE) and the mean absolute error (MAE) of the respective deviations. The assessment window is focussed on the ends of the series since the differences between the forecasted-extended HP filter and the HP filter are negligible in the body of the series. The analysis given in the previous paragraph for the 8-quarter forecast window also applies to the assessment window which is now chosen to be the last 8 quarters of the data window. Note that for historical periods TPP and TPT, this places the GFC business cycle turning point in the middle of the forecast window.

For each historical period (NTP, TPP and TPT) we define the target trend to be the HP trend of the original log GDPP data available at that time, augmented by stable (fully revised) ex-post log GDPP data. A similar strategy has also been adopted in Kaiser and Maravall (1999, 2012) and Mise et al. (2005). The augmented log GDPP data were obtained by applying the growth rates of GDPP data to 2019q4 (SNZ release, March 2020) to the last GDPP value of the original data (the last observation in the data window). The target trend defines the stable historical HP trend we wish to better estimate at the ends of series.

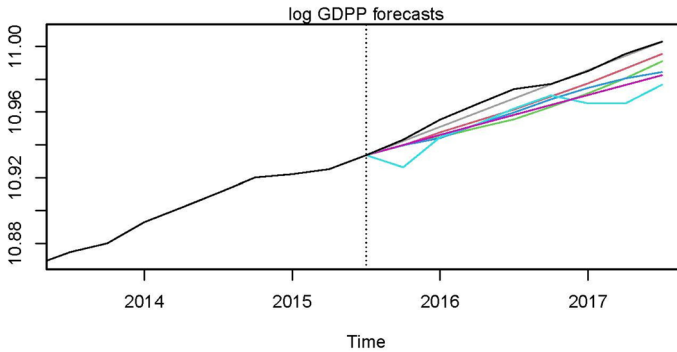


Fig. 4 Log GDPP (black) and log GDPP forecasts based on data to 2015q3 and an 8 quarter forecast window: RBNZ (red), Treasury (green), NZIER (blue), Hamilton H84 predictor (cyan), ARIMA predictor (magenta), and naïve predictor (grey)




Fig. 5 Log GDPP (black) and log GDPP forecasts based on data to 2006q4 and an 8 quarter forecast window: RBNZ (red), Treasury (green), NZIER (blue), Hamilton H84 predictor (cyan), ARIMA predictor (magenta) and naïve predictor (grey)

4.2 Results

Here we evaluate the accuracy of the various HP forecast-extension procedures for each of the NTP, TPP and TPT periods considered. The Hamilton robust predictor and ARIMA forecast extension fit models to all the data available for the period concerned, whereas the naïve predictor is more adaptive, using only the last 8 quarters of each data set. In the case of ARIMA forecast extension, a range of ARIMA models were fitted with model choice guided by information criteria such as the Akaike Information Criterion (AIC) and Bayesian Information Criterion (BIC). For each period (NTP, TPP, TPT) the first differences of the log GDPP data were

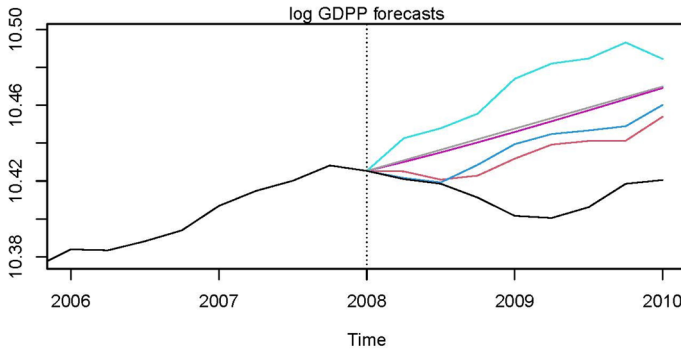


Fig. 6 Log GDPP (black) and log GDPP forecasts based on data to 2008q1 and an 8 quarter forecast window: RBNZ (red), NZIER (blue), Hamilton H84 predictor (cyan), ARIMA predictor (magenta), and naïve predictor (grey)

well-modelled by an $ARMA(1,1)$ process with mean. This was the model chosen for ARIMA forecast extension.

In Sect. 4.2.1 we evaluate the accuracy of the log GDPP forecasts used prior to applying the HP filter to the extended series and, in Sect. 4.2.2, we evaluate the accuracy of the corresponding forecast-extended HP filters at the ends of series.

4.2.1 Evaluation of Forecast Extensions

Figures 4, 5 and 6 show the log GDPP forecasts over their 8-quarter forecast windows for each of the periods NTP, TPP, and TPT respectively. Also shown are the ‘true’ log GDPP values derived from the actual log GDPP series available at the time, augmented by the ex-post log GDPP growth rates released by SNZ in March 2020. The mean forecast error, MAE and RMSE for all forecast extension methods are given in Table 4.⁶ When the forecast errors are differences in logarithms, they can be regarded as proportionate errors or, when multiplied by 100, percentage errors of the untransformed data. The accuracy of these forecasts has a direct bearing on the quality of their associated forecast-extended HP trends. Specific comments on each historical period follow.

Period NTP: Here “true” log GDPP shows a near-linear expansion path over the assessment and forecast windows. All forecasts are below “true” log GDPP with the exception of the naïve predictor which provides the best forecast. Of the forecast methods, the naïve method performs best in terms of RMSE followed by the informed forecasts with the RBNZ forecasts better than comparable forecasts from NZIER and Treasury. ARIMA forecast extension was much closer to Treasury in performance than the Hamilton robust predictor H84 which was worst.

⁶ Results have also been computed for the modified Hamilton predictor recently evaluated by Quast and Wolters (2020). These results are not presented here, as for all three of our representative business cycle periods the root mean squared errors (RMSE) from using the Quast and Wolters modification are greater than those produced by the H84 filter.

Table 4 Percentage error measures for log GDPP forecasts

Forecaster	Mean	MAE	RMSE	RMSE Rank
<i>Period NTP: data to 2015q3 and an 8 quarter forecast window</i>				
RBNZ	− 0.82	0.82	0.85	2
Treasury	− 1.27	1.27	1.33	4
NZIER	− 1.18	1.18	1.25	3
H84	− 1.69	1.69	1.85	6
ARIMA	− 1.35	1.35	1.44	5
Naive	− 0.21	0.23	0.32	1
<i>Period TPP: data to 2006q4 and an 8 quarter forecast window</i>				
RBNZ	1.22	1.49	2.15	5
Treasury	0.34	0.82	1.12	1
NZIER	0.39	1.17	1.46	3
H84	2.44	2.44	2.73	6
ARIMA	0.35	1.17	1.48	4
Naive	− 0.11	1.08	1.23	2
<i>Period TPT: data to 2008q1 and an 8 quarter forecast window</i>				
RBNZ	2.23	2.23	2.61	1
NZIER	2.64	2.64	3.13	2
H84	5.83	5.83	6.23	5
ARIMA	3.68	3.68	3.99	3
Naive	3.82	3.82	4.12	4

Mean = Mean average error; MAE = Mean absolute error; RMSE = Root mean squared error; H84 percentage error measures are for standard Hamilton OLS regression estimates; ARIMA percentage error measures are for an $ARMA(1,1)$ model plus mean fitted to all available differences; Naive percentage error measures are for an $I(1)$ model with a smooth, slowly evolving drift and stationary noise, the drift calculated from the most recent 8 quarters of data

Period TPP: In this case “true” log GDPP shows near-linear expansion over the 4 quarters (one year ahead) to 2007q4 when it turns (a peak) and enters the contraction phase of the GFC recession (2008q1–2009q1). None of the forecasts have adequately managed to forecast the turning point with most below “true” log GDPP until 2007q4 and all well above it by the end of the forecast window. The Treasury forecast, naive predictor, NZIER forecast and ARIMA forecast extension provide the best forecasts with the Hamilton robust predictor H84 the worst.

Period TPT: Here “true” log GDPP shows near-linear expansion over the assessment window followed by 5 quarters of near-linear contraction to 2009q1 when its trough enters another expansion phase. In essence there are two turning points (2007q4 and 2009q1) rather than just the one for period TPP. Of our three sample periods, this provides the most challenging forecast environment. None of the forecasts have adequately managed to forecast log GDPP over the forecast window (all are well above “true” log GDPP) although RBNZ and NZIER forecasts are

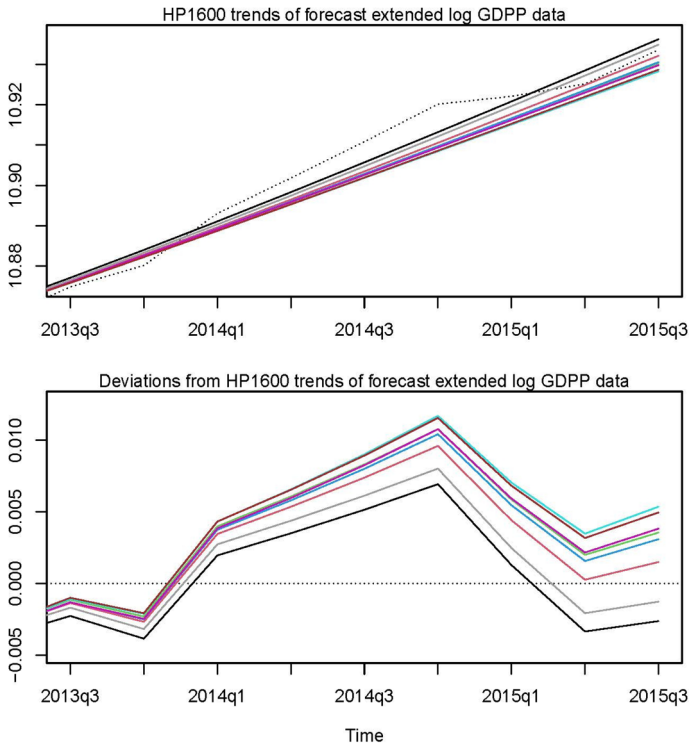


Fig. 7 Forecast-extended HP trends (top) and their trend deviations (bottom) for log GDP data to 2015q3 (top dotted) over the 8 quarter assessment window from 2013q4 to 2015q3. Also shown are the target HP trend (top black) based on ex-post log GDP data to 2019q4 and its corresponding target trend deviation (bottom black). The forecast extensions used are RBNZ (red), Treasury (green), NZIER (blue), Hamilton H84 predictor (cyan), ARIMA predictor (magenta), naïve predictor (grey) and no extension (brown). The latter corresponds to using the HP filter with no extension

closest and did predict earlier turning points (in expectation of a shorter recession).⁷ As a consequence, the informed forecasts were best in terms of RMSE, with RBNZ better than NZIER. Of the forecasts based on past data alone, the naïve predictor and ARIMA forecast extension performed comparably and were significantly better than the Hamilton robust predictor H84.

4.2.2 Evaluation of Forecast-Extended HP Filters at the Ends of Series

Figures 7, 8 and 9 show the log GDP forecast-extended HP trends and their trend deviations over the 8 quarter assessment windows for each of the periods NTP, TPP, and TPT respectively. Also shown are the HP trend with no forecast extension and

⁷ No results are presented for Treasury in the bottom panel of Table 4. Treasury's BEFU and HYEUFU forecasts are published at approximately six-monthly intervals, and so there are no directly comparable quarterly real time forecast observations available from Treasury for the TPT period.

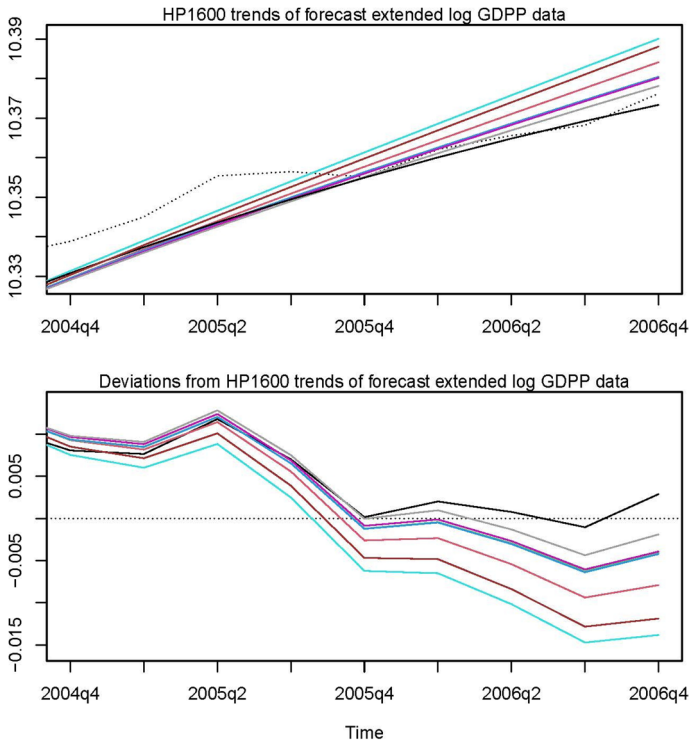


Fig. 8 Forecast-extended HP trends (top) and their trend deviations (bottom) for log GDP data to 2006q4 (top dotted) over the 8 quarter assessment window from 2005q1 to 2006q4. Also shown are the target HP trend (top black) based on ex-post log GDP data to 2019q4 and its corresponding target trend deviation (bottom black). The forecast extensions used are RBNZ (red), Treasury (green), NZIER (blue), Hamilton H84 predictor (cyan), ARIMA predictor (magenta), naive predictor (grey) and no extension (brown). The latter corresponds to using the HP filter with no extension

the target trend (the HP trend of “true” log GDP). Table 5 gives the mean, MAE and RMSE for the differences between the log GDP forecast-extended trends (including the HP trend with no extension) and the target log GDP trend. As before, when these measures involve differences in logarithms, they can be regarded as proportionate errors of the untransformed trends (percentage errors when multiplied by 100). For all periods the various trends are much the same at the beginning of the assessment window (as expected), but show greater divergence at the end. Moreover, in terms of RMSE, the forecast-extended HP trends show a worsening performance that broadly matches that of their associated forecast rankings given in Table 4. Specific comments on each historical period follow.

Period NTP: The forecast-extended HP trend using the naive predictor provides the most accurate estimate of the target trend, reflecting the relatively benign near-linear expansion path of log GDP over period NTP. This is followed by the informed forecast-extended HP trends (RBNZ, NZIER, Treasury) and the HP trend based on ARIMA forecast extension which is close to the Treasury forecast-extended




Fig. 9 Forecast-extended HP trends (top) and their trend deviations (bottom) for log GDP data to 2008q1 (top dotted) over the 8 quarter assessment window from 2006q2 to 2008q1. Also shown are the target HP trend (top black) based on ex-post log GDP data to 2019q4 and its corresponding target trend deviation (bottom black). The forecast extensions used are RBNZ (red), Treasury (green), NZIER (blue), Hamilton H84 predictor (cyan), ARIMA predictor (magenta), naïve predictor (grey) and no extension (brown). The latter corresponds to using the HP filter with no extension

HP trend. The worst trends are the HP trend with no extension and the Hamilton H84 forecast-extended HP trend which are comparable (the former being slightly better). The respective trend deviations also reflect these findings.

Period TPP: In this case the target trend does not run through the middle of the log GDP data in the assessment window since it is already turning to accommodate the contraction phase just ahead. The forecast-extended HP trend using the naïve predictor is closest to the target trend, with the informed forecast-extended HP trends (Treasury and NZIER) and the ARIMA forecast-extended HP trend all very close and not far behind. The worst trends are the HP trend with no extension and the Hamilton H84 forecast-extended HP trend with the former being the better.

Period TPT: As in the case of period TPP, the target trend does not run through the middle of the log GDP data in the assessment window. Its path is influenced by the two turning points (one in the assessment window and one in the forecast window) and so takes an intermediate course, tracking below the first turning point 2007q4 (a peak) and above the second turning point 2009q1

Table 5 Percentage error measures for the differences between forecast-extended HP log GDP trends and the target HP trend

Forecaster	Mean	MAE	RMSE	RMSE Rank
<i>Period NTP: measures for the 8 quarter assessment window from 2013q4 to 2015q3</i>				
RBNZ	− 0.25	0.25	0.27	2
Treasury	− 0.37	0.37	0.40	4=
NZIER	− 0.33	0.33	0.36	3
H84	− 0.45	0.45	0.50	7
ARIMA	− 0.37	0.37	0.40	4=
Naive	− 0.10	0.10	0.11	1
HP no extension	− 0.44	0.44	0.48	6
<i>Period TPP: measures for the 8 quarter assessment window from 2005q1 to 2006q4</i>				
RBNZ	0.42	0.43	0.56	5
Treasury	0.24	0.26	0.35	3
NZIER	0.24	0.27	0.36	4
H84	0.82	0.82	0.96	7
ARIMA	0.21	0.25	0.34	2
Naive	0.11	0.18	0.23	1
HP no extension	0.66	0.66	0.81	6
<i>Period TPT: measures for the 8 quarter assessment window from 2006q2 to 2008q1</i>				
RBNZ	0.78	0.78	0.84	1
NZIER	0.87	0.87	0.95	2
H84	1.70	1.70	1.87	6
ARIMA	1.15	1.15	1.26	3
Naive	1.19	1.19	1.30	4
HP no extension	1.60	1.60	1.75	5

Mean=Mean average error; MAE=Mean absolute error; RMSE=Root mean squared error; % reduction in RMSE is forecast-extended HP filter at the ends of series compared to the HP1600 filter with no extension; H84, ARIMA, and Naive measures are as defined in Table 4; HP no extension statistics are for Hodrick–Prescott estimates with standard assumption for quarterly data of $\lambda=1600$

(a trough). The informed forecast-extended HP trends (RBNZ and NZIER) are similar and performed best, being closest, but not close, to the target trend. Next were the HP trends based on the naive predictor and ARIMA forecast extension, both based on past data alone and both very similar. The worst trends are the HP trend with no extension and the Hamilton H84 forecast-extended HP trend (the former being the better) which are both markedly different from the target trend.

4.3 Key Findings

If the three sample periods and turning point environments chosen (NTP, TPP, TPT) are representative of those met in practice, then the following are key findings.

Table 6 Percentage reduction in RMSE of forecast-extended HP filters at the ends of series compared to the HP filter with no extension

Forecast extension	NTP	TPP	TPT	Mean
RBNZ	43	30	52	42
Treasury	17	57	–	37
NZIER	25	56	46	42
H84	–4	–18	–7	–10
ARIMA	16	58	28	34
Naive	78	71	26	59
Mean	29	42	29	34

- As expected, forecast extension can markedly improve the accuracy of the HP filter at the ends of series and, as a consequence, lessen the volatility of HP trend estimation at the ends of series.
- As a general rule, the more accurate the forecast extension, the more accurate and less volatile the forecast-extended HP trend at the end.
- Using the forecast-extended HP filter is almost always better than using the HP filter with no extension.
- For the most part, and especially for the most challenging forecast environment TPT, the best of the forecast-extended HP filters using informed forecasts (RBNZ, Treasury and NZIER) performs comparably to, or better than, the forecast-extended HP filters using forecasts based only on past data (Hamilton robust predictor, naive predictor and ARIMA forecast extension).
- In more benign environments (NTP and TPP) forecast extension using the naive predictor is more than competitive with other forecast extension methods; it also provides a useful benchmark in more challenging environments, particularly for forecast-extended HP filters using informed forecasts.
- In accord with usage reported elsewhere in the literature, the HP filter with no extension does not perform well at the ends of series, but for all representative business cycle environments considered the H84 robust predictor performed worse.

The three periods (NTP, TPP and TPT) each presented forecasting challenges of varying degrees of difficulty, with NTP the least challenging and TPT the most challenging. This is reflected in the size of the RMSE values in Table 4. Using the RMSE values in Table 5 for the HP filter with no extension as a measure of forecast difficulty, period TPP is almost twice as difficult, and period TPT almost 4 times as difficult, as the no turning point period NTP. Nevertheless, in all periods forecast extension typically led to practically significant trend improvements at the ends of series.

A measure of these trend improvements is given in Table 6 which shows the percentage reduction in RMSE using forecast extension over the assessment window, by comparison to using the HP filter with no extension. These reductions in RMSE translate directly to reductions in trend volatility at the ends of the series. Apart from the HP forecast extension using the naive predictor, the forecast-extended HP

filters using informed forecasts (RBNZ, Treasury, NZIER) were better on average (RMSE reductions of around 40 per cent) than the other methods based only on past data. For the most challenging turning point period TPT, the forecast-extended HP filters using informed forecasts dominated all the other methods. Overall, HP forecast extension using the naive predictor was best on average and dominated all other methods in the more benign environments (NTP, TPP). Since it and ARIMA forecast extension are based on $I(1)$ models with drift, it is likely that its success is due to the fact that it is the only truly adaptive forecast-extension method with drift determined locally (the median growth rate over the last 8 quarters).

We also note that the results for the two turning point periods TPP and TPT are consistent with the findings of Joutz and Stekler (2000) who found that, for four U.S. recessions during the period 1965 to 1989, Greenbook forecasts produced by the Federal Reserve staff generally failed to call an NBER business cycle peak in advance and tended to predict a cycle trough too early.

These findings are in accord with the more extensive practical and theoretical evidence in the forecast-extension literature. For macroeconomic business cycle analysis, (optimal) forecast extension should be used routinely to minimise trend volatility at the ends of series. The gains in practice are likely to be considerable and largely eliminate many of the deficiencies associated with the HP filter, especially at the ends.

5 Conclusions

In a New Zealand business cycle context, we assess whether Hamilton's H84 OLS regression method has provided a "better alternative" to the Hodrick-Prescott (HP) filter. In particular, we evaluate comparative performance in two areas: stylised business cycle facts produced by the H84 filter relative to those obtained using the HP and Baxter-King (BK) filters; and relative performance of the forecast-extended HP filter at the ends of series using the H84 predictor relative to other benchmark forecast-extension methods.

Firstly, for a set of key quarterly New Zealand macroeconomic variables typically included in a small theoretical or empirical macroeconomic model, H84 trends and cycles lead to considerably greater cycle volatilities than those computed from either the HP or BK filters which are comparable. Others have shown that H84 cycle volatilities for U.S. real GDP are typically over twice those of HP cycle volatilities (e.g. Hodrick, 2020; Schüler, 2018). Our findings confirm this effect which is even more pronounced when H84 cycle volatilities are compared to BK cycle volatilities and holds for the majority of our key macroeconomic series. Cycle persistence is generally less for HP than for H84 which, in turn, is generally less than BK, but the differences are, for the most part, not statistically significant. Cross correlations of H84, HP and BK cycles with real GDPE show broad agreement, with any differences being not statistically significant.

Accordingly, for a small open economy like New Zealand, and primarily on the basis that H84 produces materially greater volatilities and less credible trend movements associated with H84's inherent phase-shift behaviour, particularly during key

economic periods such as the 1991–92 and GFC recessions, we have a clear preference for measures of stylised business cycle facts produced by the HP or BK filters rather than those from Hamilton's H84 procedure. A similar conclusion has been reached by Hodrick (2020).

Secondly, at the ends of the series we evaluate the performance of the forecast-extended HP filter for real GDP and three representative business cycle environments. The forecast-extension methods compared include the H84 predictor, two methods based on models of past data, and the HP filter with no extension. They also included a specially compiled data set of real-time forecasts published by two leading New Zealand public sector institutions (RBNZ and the New Zealand Treasury) and a prominent private sector agency (NZIER). The three representative business cycle environments included a relatively undemanding close-to-linear expansion path with the other two being more demanding and involving business cycle turning points. For this paper, the latter are New Zealand's GFC-related business cycle peak at 2007q4 and its corresponding 2009q1 business cycle trough. A two-year assessment window at the end of each series was used to evaluate the performance of the various forecast-extended HP trends relative to a true HP trend based on ex-post data to 2019q4.

For all three end point environments considered, the HP filter with forecast extension was almost always markedly better than using the HP filter with no extension and led to practically significant trend improvements at the ends of series. As a general rule, the more accurate the forecast extension, the more accurate and less volatile the forecast-extended HP trend at the end. These results are in accord with findings reported elsewhere in the literature. The notable exception to these findings was forecast extension using the Hamilton H84 predictor which produced results that are inferior, in terms of root-mean-squared-error (RMSE), to those produced by the other forecast-extension methods and the HP filter with no extension. This outcome should not be regarded as surprising when the forecast extension comes from prominent economic forecasting agencies. Of greater significance is that H84 forecast-extension failed to out-perform two other univariate benchmark methods based solely on past data, and produced RMSEs which were inferior to those using the HP filter with no extension whose poor performance has been well-documented elsewhere in the literature.

Hence, in a New Zealand business cycle context, the evidence presented here suggests 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.

Acknowledgements We acknowledge valuable comments on an earlier version from four anonymous reviewers and from Bob Buckle.

Funding None.

Availability of Data and Material Available on request from corresponding author.

Code availability Available on request from corresponding author.

Declarations

Conflict of interest None.

References

- Akaike, H. (1980). Seasonal adjustment by a Bayesian modelling. *Journal of Time Series Analysis*, 1, 1–13.
- Baxter, M., & King, R. G. (1999). Measuring Business Cycles: Approximate band-pass filters for economic time series. *Review of Economics and Statistics*, 81(4), 575–593. <https://doi.org/10.1162/003465399558454>
- Bry, G., & Boschan, C. (1971). *Cyclical analysis of time series: Selected procedures and a computer program*. Columbia University Press.
- Burman, J. P. (1980). Seasonal adjustment by signal extraction. *Journal of the Royal Statistical Society A*, 143, 321–337.
- Burns, A. F., & Mitchell, W. C. (1946). *Measuring business cycles*. New York, NY: NBER.
- Christiano, L. J., & Fitzgerald, T. J. (2003). The band pass filter. *International Economic Review*, 44(2), 435–465.
- Dagum, E. B. (1980). *The X11 ARIMA Seasonal Adjustment method*. Statistics Canada Catalogue 12–564E.
- De Jong, R. M., & Sakarya, N. (2016). The Econometrics of the Hodrick–Prescott filter. *Review of Economics and Statistics*, 98(2), 310–317.
- Drehmann, M., & Yetman, J. (2018). Why you should use the Hodrick–Prescott filter—at least to generate credit gaps. BIS Working Papers No 744, September.
- Durbin, J., & 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., & Chen, B. (1998). New capabilities and methods of the X-12-ARIMA seasonal adjustment program. *Journal of Business and Economic Statistics*, 16, 127–177.
- Garratt, A., Lee, K., Mise, E., & Shields, K. (2008). Real-time representations of the output gap. *Review of Economics and Statistics*, 90(4), 792–804.
- Geweke, J. (1978). The revision of seasonally adjusted time series. In *Proceedings of the Business and Economics Section, American Statistical Association*, 320–325.
- Gomez, V. (1999). Three equivalent methods for filtering nonstationary time series. *Journal of Business and Economic Statistics*, 17, 109–116.
- Gray, A. G., & Thomson, P. J. (2002). On a family of moving-average trend filters for the ends of series. *Journal of Forecasting*, 21, 125–149.
- Hall, V. B., & McDermott, C. J. (2016). Recessions and recoveries in New Zealand’s post-Second World War business cycles. *New Zealand Economic Papers*, 50(3), 261–280.
- Hall, V. B., Thomson, P., & McKelvie, S. (2017). On the robustness of stylised business cycle facts for contemporary New Zealand. *New Zealand Economic Papers*, 51(3), 193–216.
- Hamilton, J. D. (2018). Why you should never use the Hodrick–Prescott filter. *Review of Economics and Statistics*, 100(5), 831–843.
- Harding, D., & Pagan, A. R. (2002). Dissecting the cycle: A methodological investigation. *Journal of Monetary Economics*, 49(2), 365–381.
- Harvey, A. C. (1989). *Forecasting, structural time series models and the Kalman filter*. Cambridge University Press.
- Harvey, A. C., & Jaeger, A. (1993). Detrending, stylised facts and the business cycle. *Journal of Applied Econometrics*, 8, 231–247.
- Henderson, R. (1916). Note on graduation by adjusted average. *Transactions of the Actuarial Society of America*, 17, 43–48.
- Henderson, R. (1924). A new method of graduation. *Transactions of the Actuarial Society of America*, 25, 29–40.
- Hodrick, R. J., & Prescott, E. C. (1997). Postwar US business cycles: An empirical investigation. *Journal of Money, Credit, and Banking*, 29, 1–16.
- Hodrick, R. J. (2020). An exploration of trend-cycle decomposition methodologies in simulated data. *NBER Working Paper 26750*, February.
- Jönsson, K. (2020a). Cyclical dynamics and trend/cycle definitions: Comparing the HP and Hamilton filters. *Journal of Business Cycle Research*, 16, 151–162.

- Jönsson, K. (2020b). Real-time US GDP gap properties using Hamilton's regression-based filter. *Empirical Economics*, 59, 307–314.
- Joutz, F., & Stekler, H. O. (2000). An evaluation of the predictions of the Federal Reserve. *International Journal of Forecasting*, 16, 17–38.
- Kaiser, R., & Maravall, A. (2012). *Measuring business cycles in economic time series*. Lecture notes in statistics 154. New York, NY: Springer Science & Business Media.
- Kaiser, R., & Maravall, A. (1999). Estimation of the business cycle: A modified Hodrick-Prescott filter. *Spanish Economic Review*, 1(2), 175–206.
- Kim, K., Buckle, R. A., & Hall, V. B. (1994). Key features of New Zealand business cycles. *Economic Record*, 70(208), 56–72.
- King, R., & Rebelo, S. T. (1993). Low frequency filtering and real business cycles. *Journal of Economic Dynamics and Control*, 17, 207–231.
- Kitagawa, G., & Gersch, W. (1996). *Smoothness priors analysis of time series*. Springer-Verlag.
- Labbé, F., & Pepper, H. (2009). Assessing recent external forecasts. *Reserve Bank of New Zealand Bulletin*, 72(4), 19–25.
- Lees, K. (2016). Assessing forecast performance. *Reserve Bank of New Zealand Bulletin*, 79(10), 1–19.
- Macaulay, F. (1931). *The smoothing of time series*. National Bureau of Economic Research.
- 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., & Hall, V. B. (2012). Stylised facts for New Zealand business cycles: A Post-1987 perspective. School of Economics and Finance Working Paper 12/2012. Wellington: Victoria University of Wellington. Retrieved from <http://researcharchive.vuw.ac.nz/handle/10063/2364>.
- Mise, E., Kim, T.-H., & Newbold, P. (2005). On suboptimality of the Hodrick-Prescott filter at time series endpoints. *Journal of Macroeconomics*, 27(1), 53–67.
- Phillips, P. C. B., & Jin, S. (2021). Business Cycles, Trend elimination, and the HP Filter. *International Economic Review*, 62(2), 469–520.
- Phillips, P. C. B., & Shi, Z. (2021). Boosting: Why you can use the HP filter. *International Economic Review*, 62(2), 521–570.
- Pierce, D. A. (1980). Data revisions with moving average seasonal adjustment procedures. *Journal of Econometrics*, 14, 95–114.
- Quast, J., & Wolters, M. (2020). Reliable real-time output gap estimates based on a modified Hamilton filter. *Journal of Business & Economic Statistics*. <https://doi.org/10.1080/07350015.2020.1784747>
- R Development Core Team (2004). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna. <http://www.R-project.org>
- Schüler, Yves S. (2018). On the cyclical properties of Hamilton's regression filter. *Deutsche Bundesbank Discussion Paper 03/2018*.
- Sleeman, C. (2006). Analysis of revisions to quarterly GDP—a real-time database. *Reserve Bank of New Zealand Bulletin*, 69(1), 31–44.
- Stock, J. H., & Watson, M. W. (1999). Business cycle fluctuations in US macroeconomic time series. *Handbook of Macroeconomics*, 1, 3–64.
- Stock, J. H., & Watson, M. W. (2007). Why has U.S. inflation become harder to forecast? *Journal of Money, Credit and Banking*, 39(s1), 3–33.
- Whittaker, E. T. (1923). On a new method of graduation. *Proceedings of the Edinburgh Mathematical Society*, 41, 63–75.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.