



Australian
National
University

Crawford School of Public Policy

CAMA

Centre for Applied Macroeconomic Analysis

Does Hamilton's OLS regression provide a "better alternative" to the Hodrick-Prescott filter? A New Zealand Business Cycle Perspective

CAMA Working Paper 71/2020
July 2020

Viv B Hall

School of Economics and Finance, Victoria University of Wellington
Centre for Applied Macroeconomic Analysis, ANU

Peter Thomson

School of Economics and Finance, Victoria University of Wellington
Statistics Research Associates, Wellington

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 HP1600 measures, and whether using the H84 predictor and other forecast extensions improves the HP filter's properties at the ends of series.

In general, H84 produces exaggerated volatilities and less credible trend movements during key economic periods so there is no material advantage in using H84 de-trending over HP1600. At the ends, the forecast-extended HP filter almost always performs better than the HP filter with no extension which performs slightly better than H84 forecast extension.

Keywords

Hamilton regression filter, stylised business cycle facts, New Zealand, end-point issues

JEL Classification

E32, E37, C10, G01

Address for correspondence:

(E) cama.admin@anu.edu.au

ISSN 2206-0332

[The Centre for Applied Macroeconomic Analysis](#) in the Crawford School of Public Policy has been established to build strong links between professional macroeconomists. It provides a forum for quality macroeconomic research and discussion of policy issues between academia, government and the private sector.

The Crawford School of Public Policy is the Australian National University's public policy school, serving and influencing Australia, Asia and the Pacific through advanced policy research, graduate and executive education, and policy impact.

Does Hamilton’s OLS regression provide a “better alternative” to the Hodrick-Prescott filter? A New Zealand Business Cycle Perspective

Viv B Hall^{a†} and Peter Thomson^{ab}

^a *School of Economics and Finance, Victoria University of Wellington,
and Centre for Applied Macroeconomic Analysis, ANU*

^b *Statistics Research Associates, Wellington
peter@statsresearch.co.nz*

July 2020

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 HP1600 measures, and whether using the H84 predictor and other forecast extensions improves the HP filter’s properties at the ends of series.

In general, H84 produces exaggerated volatilities and less credible trend movements during key economic periods so there is no material advantage in using H84 de-trending over HP1600. At the ends, the forecast-extended HP filter almost always performs better than the HP filter with no extension which performs slightly better than H84 forecast extension.

JEL Classification: E32, E37, C10, G01

Keywords: Hamilton regression filter; stylised business cycle facts; New Zealand; end-point issues

[†] Corresponding author: School of Economics and Finance, Victoria University of Wellington.
Email: viv.hall@vuw.ac.nz

1. Introduction

James Hamilton (2018, p 831) makes a case as to why you should never use the Hodrick-Prescott Filter, 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.” (p 831). In particular this alternative “ ... can isolate a stationary component from any $I(4)$ series, preserves the underlying dynamic relations, and consistently estimates well-defined population characteristics for a broad class of possible data-generating processes.” (pp 839-840).

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 diff). Underlying these results is an assumed benchmark business cycle of around two years which, together with the regression incorporating the four most recent values of the variable, leads to a reduction in the beginning-period H84 output values by 12 quarters and the corresponding H8 diff output values by eight quarters. The regression can also be used to project forward eight more quarters of observations. Growth cycle volatility and contemporaneous cross-correlation results are provided (2018, Table 2), and benchmarked against random walk results, but perhaps surprisingly not against results from using the HP1600 filter (HP); nor does he provide robust standard errors for these results or provide non-contemporaneous cross-correlation results. There is also no explicit assessment of the economic circumstances in which end-point issues might be empirically concerning.

To date, a limited number of studies have empirically evaluated Hamilton’s methodology. In a New Zealand context, Callaghan, Culling and Robinson (2018, Appendix D) have used both the HP1600 filter and Hamilton’s linear projection method to compute Labour Force Participation gaps, though on the basis that Schuler (2018) shows Hamilton’s methodology to be more suited to credit cycles than to business cycles, Callaghan *et al.* opted to use the HP filter for their main results.

For U.S. log GDP and credit-to-GDP data, Schuler (2018) has compared the cyclical properties of Hamilton’s regression filter with those from the HP filter. Overall, in the context of emphasising that the “correct” filter should depend on the investigator’s objective, he finds that

while Hamilton's filter is not subject to the same drawbacks as the HP filter, it too reflects some 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. More positively, though, Schuler concludes that the Hamilton regression filter should be preferred to the HP filter when constructing a credit-to-GDP gap.

Drehman 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 $\lambda = 400,000$ as smoothing parameter). Further, they find that credit gaps based on linear projections in real time perform poorly.

Phillips and Shi (2019) develop a boosted HP filter (bHP) and analyse its performance relative to Hamilton's regression approach. Their findings show a clear preference for the bHP filter over the Hamilton regression, and also support the conclusion that the HP filter may continue to be used as a helpful empirical device for the estimation of trends and cycles.

Against the above background, we evaluate Hamilton's proposed H84 and H8 diff methods for key New Zealand macroeconomic time series of considerably shorter duration and markedly greater volatility than the 70-year employment and GDP series illustrated by Hamilton. This is done in both a stylised business cycle facts context and with a significant focus on end-point implications. Accordingly, consistent with empirical stylised business cycle facts analysis and macroeconomic model-based approaches to business cycle work, the deviations from trend cycles assessed in this paper are *growth cycles* rather than *classical business cycles*ⁱ.

More specifically, the two broad questions we address are:

- For post-1987q2 New Zealand, does Hamilton's H84 regression methodology produce business cycle volatility and bivariate cross-correlation measures which are materially different from the HP1600 measures published in Hall *et al.* (2017) (HTM)?
- In an end-point context, does forecast extension improve the properties of the HP filter at the ends of series and, in this context, is the Hamilton H84 regression (an autoregressive predictor) a useful forecast extension method?

The latter question is investigated for two historical environments: New Zealand's post-2009q1 classical 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 – 2009q1ⁱⁱ.

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

2. Methodological framework

We consider a non-seasonal quarterly economic time series x_t , possibly log transformed. Following HTM (2017) we assume that x_t admits the additive decomposition

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

where g_t is an unobserved or hidden trend and d_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 d_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

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

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 and Prescott, 1997) where the $w_t(s)$ are time varying, the Baxter-King filter (Baxter and 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 X-12-ARIMA procedure 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 and Thomson, 2002, for example). Here it is known (see Geweke, 1978, 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. This approach takes the best of both modelling frameworks. Dynamic time series models can be used to provide out-of-sample forecasts that can be used to eliminate deficiency (d) and provide more stable estimates of the desired empirical trend at the ends of series which directly addresses deficiency (c). Such *forecast-extension* procedures have been used in business cycle analysis (see Christiano and Fitzgerald, 2003, for example), but have yet to be as routinely adopted as they have in seasonal adjustment 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*. This is our **target trend**. As a consequence, the trend deviations in the body of the series 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 are compared including the case of no extension (using the HP filter trend estimates at the ends).

In Section 2.1 we briefly discuss the criticisms of Hamilton (2018) within the above framework and, in Section 2.2, provide further details on the forecast extension procedures we have adopted.

2.1 Hamilton critiqueⁱⁱⁱ

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 Section 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 d_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

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

by OLS to get forecasts and prediction errors given by

$$\hat{x}_t = \hat{\beta}_0 + \hat{\beta}_1 x_{t-h} + \hat{\beta}_2 x_{t-h-1} + \hat{\beta}_3 x_{t-h-2} + \hat{\beta}_4 x_{t-h-3}, \quad \hat{v}_t = x_t - \hat{x}_t \quad (4)$$

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 (forecast) determined from data up to 8 quarters (2 years) earlier, whereas the trend deviations \hat{d}_t measure the departure of x_t from its local level \hat{g}_t , 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.

In Section 3, we apply the Hamilton (2018) methodology to a subset of the New Zealand macroeconomic time series for which HTM (2017) have previously published HP filter results. The resulting differences are noted. In Section 4 we use the Hamilton robust predictor \hat{x}_t for forecast extension and compare it to other contenders.

2.2 Further details of forecast extension procedures adopted

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

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

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 is F the closer \hat{g}_t follows the data, and the smaller is S 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 end of the series are always time-varying, those in the body of the series are essentially 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 our target trend filter which we call the **central HP filter**. While the HP trend estimates at the ends of the series can be motivated by signal extraction and attributed to a suitable structural time series model (see Harvey and Jaeger, 1993), here we make no such attribution. They will be referred to as the HP end filters which, in this case, define the *HP filter with no forecast extension*.

In addition to forecast extension using the *Hamilton robust predictor* (4) we consider forecast extension using the simple random walk model

$$x_t = x_{t-1} + \delta_t + \varepsilon_t$$

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 is intended to provide a *benchmark forecast extension* and will be called the *naïve predictor*.

We also consider forecast extension using predictions published by leading New Zealand public sector economic forecasters (Reserve Bank of New Zealand (RBNZ) and New Zealand Treasury (Treasury)), and a prominent private sector economic forecasting entity (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.

In section 4.2 below, we present comparative results for six forecast extension cases: RBNZ, Treasury, NZIER, the Hamilton robust predictor, the naïve predictor based on a simple random walk model, and the HP filter with no extension (the standard HP1600 end filter).

3. Empirical results: a comparison of HP and H84 business cycle facts

3.1 Data

The HTM raw data series have been sourced from Statistics New Zealand (SNZ), the RBNZ and Treasury, as documented in McKelvie and Hall (2012, Appendix C). Additional to the series in that data set, the CPI tradables and non-tradables series are from the RBNZ. Series were seasonally adjusted as required and then log transformed, with the exception of those containing negative observations (e.g. net exports share, CPI tradables inflation rate, real 90-day Bank Bill rate) or those already expressed as a percentage (e.g. unemployment, gross government debt/GDP).

3.2 Results

HP1600 and Hamilton H84 volatility, persistence and cross correlation characteristics are presented in Tables 1 and 2 for a representative subset of the preferred HP1600 results published in HTM (2017)^{iv}. For ease of comparison, these results are presented for the same sample period as in HTM, i.e. 1987q2 to 2015q3.

Comparative trend and trend deviation paths are illustrated in Figures 1, 2 and 3 for real expenditure-based gdp (gdpe), real residential investment (invres) and real gross fixed capital formation (gfcf). The underpinning H84 regression output for these three variables can be found in Appendix 1.

[Figures 1, 2, & 3 about here]

Key findings here are that:

- H84 trends and cycles lead to considerably greater cycle volatilities than those computed from HP1600. In particular, for the variables gdpe, invres and gfcf, H84 volatilities are at least double those of the HP1600 volatilities (Table 1)^v. It is also clear from Figures 1 to 3 that the H84 trends do not represent at all well the corresponding SNZ outcomes for the periods associated with the 1991-92 and GFC recessions.
- For almost all variables, the most statistically significant H84 cross correlations are greater in (absolute) magnitude than those for HP1600 (Table 2). They also have the disadvantage of losing 12 observations for the beginning of sample periods, a not unimportant issue for New Zealand's relatively short quarterly macroeconomic time series.

[Tables 1 & 2 about here]

- Hence, representative stylised New Zealand business cycle facts emanating from use of H84 OLS regression methodology are materially different from those obtainable from HP1600 de-trending.

4. Empirical results: a comparison of HP forecast extension methods

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

- New Zealand’s post-2009q1 classical business cycle expansion path with no turning points (**NTP**); and
- two illustrative business cycle turning point environments associated with the peak (**TPP**) and trough (**TPT**) of New Zealand’s five quarter classical GFC recession 2007q4 – 2009q1.

For each historical period (NTP, TPP and TPT) we consider quarterly GDPP time series data up to a given time point in the given 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, but use more recent GDPP data (as of 2019q4) to augment the available data to provide “true” values of GDPP for the period after each series end point.

Historical 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. Examining results for the periods TPP and TPT is important. The various forecast methods are unlikely to vary greatly along an ongoing classical business cycle expansion path, but are likely to differ in the neighbourhood of turning points. In the latter case, significantly different results will directly address the issue of whether credible forecast extension methods can significantly improve the end point accuracy and volatility concerns associated with use of the HP filter at the ends of series.

Further details on the data series and performance measures adopted are given in Section 4.1. Results are presented in Section 4.2.

4.1 Data and performance measures adopted

For period NTP we use log transformed New Zealand quarterly GDPP data from 1987q2 to 2015q3 released by Statistics New Zealand in December 2015. This data was analysed in HTM and used by RBNZ, Treasury and NZIER for their forecasts. The Hamilton robust predictor and the naïve predictor base their forecasts on the last 8 quarters of this data.

For period TPP we use log transformed New Zealand quarterly GDPP data from 1987q2 to 2006q4 released by Statistics New Zealand in March 2007. This data precedes the 2007q4 GFC business cycle peak by 4 quarters. For period TPT we use log transformed New Zealand GDPP data from 1987q2 to 2008q1 released by Statistics New Zealand in June 2008. This data

precedes the 2009q1 GFC business cycle trough by 4 quarters. The TPT and TPT data sets were taken from the real-time data sets compiled by the RBNZ (see Sleeman, 2006). As before, the Hamilton robust predictor and the naïve predictor base their forecasts on the last 8 quarters of each data set.

The RBNZ, Treasury and NZIER forecasts 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 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 regression 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 which is a symmetric moving average filter of the form (2) whose time-invariant weights are given by

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

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

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

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. 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. This analysis provides 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 suitable 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 HP1600 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 HP1600 trend of the original log GDPP data available at that time, augmented by stable (fully revised) ex-post log GDPP data. The latter were obtained by applying the growth rates of GDPP data to 2019q4 (released by Statistics New Zealand in 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 HP1600 trend we wish to better estimate at the ends of series.

4.2 Results

We consider each of the three historical periods NTP, TPP and TPT in turn. In each case the log GDPP forecasts are first evaluated before comparing the accuracy of the corresponding forecast-extended HP filters at the ends of the series.

4.2.1 Historical period NTP

Figure 4 shows the log GDPP forecasts produced by the forecast extension methods over the 8 quarter forecast window from 2015q4 to 2017q3. Also shown are the “true” log GDPP values derived from the actual log GDPP series available at that time (2015q3) augmented by the ex-post log GDPP growth rates released in March 2020. The latter time series (“true” log GDPP) shows a near-linear expansion path over the assessment and forecast windows. All forecasts are below the “true” log GDPP values with the exception of the naïve predictor which provides the best forecast. The mean (bias), mean absolute error (MAE) and root mean squared error (RMSE) for all forecast extension methods are given in Table 3a. Since 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. Of the forecast methods, the naïve method performs best in terms of RMSE followed by the forecasts from the RBNZ, NZIER and Treasury. The Hamilton robust predictor H84 is worst. The accuracy of these forecasts has a direct bearing on the quality of their associated forecast-extended HP1600 trends.

[Figure 4 about here]

[Table 3a, 3b, 3c about here]

Figure 5 shows the log GDPP forecast-extended HP1600 trends and their trend deviations from log GDPP over the 8 quarter assessment window from 2013q4 to 2015q3. Also shown are the HP1600 trend with no forecast extension and the target trend (the HP1600 trend of the “true” log GDPP). Although the various trends are much the same at the beginning of the assessment window (as expected), they show greater divergence at the end. The forecast-extended HP1600 trend using the naïve predictor provides the most accurate estimate of the target trend, reflecting the relatively benign near-linear expansion path of log GDPP over period NTP. Table 5a gives the mean (bias), mean absolute error (MAE) and root mean squared error (RMSE) for the differences between the log GDPP forecast-extended trends (including the HP1600 trend with no extension) and the target log GDPP trend. Again, these measures involve differences in logarithms and can be regarded as proportionate errors or percentage errors of the untransformed trends. In terms of RMSE, the forecast-extended HP1600 trends show a worsening performance in the same RMSE rank order as their forecasts in Table 4a. The worst trends are the HP1600 trend with no extension and the Hamilton H84 forecast-extended HP1600 trend which are comparable (the former being slightly better). The respective trend deviations also reflect these findings and rankings.

[Figure 5 about here]

[Table 4a, 4b, 4c about here]

4.2.2 Historical period TPP

Figure 6 shows the log GDPP forecasts produced by the forecast extension methods over the 8 quarter forecast window from 2007q1 to 2008q4. The “true” log GDPP values are also shown where these now comprise the data available at that time (2006q4) augmented by the ex-post

log GDPP growth rates released in March 2020. The “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 mean (bias), mean absolute error (MAE) and root mean squared error (RMSE) for all forecast extension methods are given in Table 3b where the forecast errors can be regarded as percentage errors of the untransformed data. The Treasury forecast, naïve predictor and the NZIER forecast provide the best forecasts and the Hamilton robust predictor H84 the worst.

[Figure 6 about here]

Figure 7 shows the log GDPP forecast-extended HP1600 trends and their trend deviations from log GDPP over the 8 quarter assessment window from 2005q1 to 2006q4. The HP1600 trend with no forecast extension and the target trend are also shown. The latter is the HP1600 trend of the “true” log GDPP (data to 2006q4 augmented by ex-post log GDPP growth rates released in March 2020). In this case it does not run through the middle of the log GDPP data in the assessment window (also evident in Figure 1) since it is already turning to accommodate the contraction phase just ahead. As for period NTP, the various near-linear trends are much the same at the beginning of the assessment window, but show divergence at the end. Table 4b gives the mean (bias), mean absolute error (MAE) and root mean squared error (RMSE) for the differences between the log GDPP forecast-extended trends (including the HP1600 trend with no extension) and the target log GDPP trend where these differences can be regarded as percentage errors of the untransformed trends. The forecast-extended HP1600 using the naïve predictor is closest to the target trend, with Treasury and NZIER forecast-extended HP1600 trends not far behind. The worst trends are the HP1600 trend with no extension and the Hamilton H84 forecast-extended HP1600 trend with the former being the better. Again, the ranking of the forecast-extended trends broadly matches their forecast ability and forecast rankings.

[Figure 7 about here]

4.2.3 Historical period TPT

Figure 8 shows the log GDPP forecasts produced by the forecast extension methods over the 8 quarter forecast window from 2008q2 to 2010q1. The “true” log GDPP values are also shown where now these comprise the data available at that time (2008q1) augmented by the ex-post log GDPP growth rates released in March 2020. Here the “true” log GDPP shows near-linear expansion over the assessment window followed by 5 quarters of near-linear contraction to 2009q1 when it turns (a trough) and 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 closest and did predict earlier turning

points (in expectation of a shorter recession). The mean (bias), mean absolute error (MAE) and root mean squared error (RMSE) for all forecast extension methods are given in Table 3c where the forecast errors can be regarded as percentage errors of the untransformed data. The RBNZ and NZIER forecasts were best in terms of RMSE and the Hamilton robust predictor H84 the worst.

[Figure 8 about here]

Figure 9 shows the log GDP forecast-extended HP1600 trends and their trend deviations from log GDP over the 8 quarter assessment window from 2006q2 to 2008q1. The HP1600 trend with no forecast extension and the target trend are also shown. As for period TPP, the target trend does not run through the middle of the log GDP data in the assessment window (see Figure 1 for example). 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 (a trough). Again, the various near-linear trends are much the same at the beginning of the assessment window, but show divergence at the end. The RBNZ and NZIER forecast-extended HP1600 trends are much the same and are the closest, but not close to, the target trend. The worst trends are the HP1600 trend with no extension and the Hamilton H84 forecast-extended HP1600 trend (the former being the better) which are both markedly different from the target trend. These observations are supported by Table 4c which gives the mean (bias), mean absolute error (MAE) and root mean squared error (RMSE) for the differences between the log GDP forecast-extended trends (including the HP1600 trend with no extension) and the target log GDP trend where these differences can be regarded as percentage errors of the untransformed trends. As before, the ranking of the forecast-extended trends broadly matches their forecast ability and forecast rankings.

[Figure 9 about here]

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 the key findings.

- Forecast extension can markedly improve the accuracy of the HP filter at the ends of series and, as a consequence, minimise the volatility of trend estimation at the ends.
- As a general rule, the more accurate the forecast extension, the more accurate and less volatile the forecast-extended HP1600 trend at the end.
- Using the forecast-extended HP filter is almost always better than using the HP filter with no extension.
- In turning point environments (TPP, TPT) the forecast-extended HP filter using informed forecasts (RBNZ, Treasury and NZIER) typically performs better than using forecasts based only on past data (the naive predictor and the Hamilton robust predictor), particularly for forecasts up to 1 year ahead.

- In benign environments (no turning points) forecast extension using the naïve predictor is more than competitive with other forecast extension methods; it also provides a useful benchmark method in more challenging environments.
- In accord with usage reported in the literature, the HP filter with no extension does not always perform well at the ends of series; it is comparable to, but better than, the use of forecast extension with the Hamilton robust predictor.

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 Tables 3a, 3b and 3c. Using the RMSE values in Tables 4a, 4b and 4c for the HP1600 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.

Table 5 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. For the turning point periods TPP and TPT in particular, the three economic forecasters (RBNZ, Treasury and NZIER) typically show practically useful reductions in trend volatility over the benchmark naïve forecast extension.

[Table 5 about here]

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 compelling. They are also in accord with the more extensive practical and theoretical evidence in the seasonal adjustment literature that underpin procedures such as X-12-ARIMA. 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

Using log transformed New Zealand quarterly gdp data post 1987q2, we have assessed whether Hamilton's H84 regression methodology produces stylised business cycle facts that are materially different from the HP1600 measures published in Hall *et al.* (2017). In an end-point context, we have also assessed whether using the Hamilton robust predictor (H84), along with other forecast extension procedures, improves the HP filter's properties at the ends of series.

In general, the H84 regression methodology produces exaggerated volatilities and fails to produce credible trend movements during key economic periods. On comparative business cycle properties we find that H84 trends and cycles lead to greater business cycle volatilities than HP1600 and, for almost all variables, the most statistically significant H84 cross correlations are greater in (absolute) magnitude than those for HP1600. H84 also has the disadvantage of losing 12 observations from the beginning of sample periods, a not unimportant issue for New Zealand's relatively short quarterly macroeconomic time series. Accordingly, our preferred growth cycle stylised facts remain those obtained from HP1600 methodology.

Three historical sample periods associated with New Zealand's five-quarter classical GFC recession (2008q1 – 2009q1) were chosen to be representative of end point environments met in practice. These included a case of no turning points (NTP), a turning point peak (TPP) and a turning point trough (TPT) in the two year forecast window following the end of the series. Of these, TPT presented the greatest forecasting challenge. A two year assessment window at the end of each series was used to evaluate the performance of the various forecast extended HP1600 trends by comparison to the 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 HP1600 trend at the end. The only exception to this rule was forecast extension using the Hamilton H84 predictor which produced comparable results to, but was worse than, the HP filter with no extension. In GFC turning point environments (TPP, TPT), the forecast-extended HP filter using informed forecasts (RBNZ, Treasury and NZIER) generally performed better than using forecasts based only on past data. In benign environments (NTP), forecast extension using the naïve predictor was more than competitive with other forecast extension methods and also provided a useful benchmark method in more challenging environments. These findings are in accord with the more extensive practical and theoretical evidence in the seasonal adjustment literature that underpin procedures such as X-12-ARIMA.

On the basis of the evidence presented there is no material advantage in using H84 de-trending over HP1600 in a New Zealand business cycle context. Furthermore, forecast extension can be used to minimise trend volatility at the ends of series with the credibility of competing forecast extension methods benchmarked against both naïve forecast extensions and (ex-post) against the trends from SNZ actual outcomes. In practice, the gains from forecast extension are likely to be considerable and largely eliminate many of the deficiencies associated with the HP filter, especially at the ends.

Acknowledgements

We acknowledge valuable comments on an earlier version from Bob Buckle.

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. doi: 10.1162/003465399558454
- Burns, A.F. and Mitchell, W.C. (1946) *Measuring business cycles*. New York, NY: NBER.
- Callaghan, Michael, Culling, Jamie, & Finn Robinson (2018). Ageing is a drag: projecting labour force participation in New Zealand. *Reserve Bank of New Zealand Analytical Note AN2018/10*, November.
- Christiano, L.J. and T.J. Fitzgerald (2003). The band pass filter. *International Economic Review*, 44(2), 435-465.
- De Jong, R.M. and Sakarya, N. (2016). The Econometrics of the Hodrick-Prescott Filter. *Review of Economics and Statistics*, 98(2), 310-317.
- Drehmann, M. and 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. 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.
- Geweke, J. (1978) The revision of seasonally adjusted time series. In *Proceedings of the Business and Economics Section, American Statistical Association*, 320-325.
- Gray, A.G. and Thomson, P.J. (2002) On a family of moving-average trend filters for the ends of series. *Journal of Forecasting*, 21, 125-149.
- Hall, Viv. B., & C. John McDermott (2016). Recessions and recoveries in New Zealand's post-Second World War business cycles. *New Zealand Economic Papers*, 50(3), December, 261-280.

- Hall, Viv. B., Thomson, Peter, & Stuart McKelvie (2017). On the robustness of stylised business cycle facts for contemporary New Zealand. *New Zealand Economic Papers*, 51(3), December, 193-216.
- Hamilton, James D. (2018). Why you should never use the Hodrick-Prescott filter. *Review of Economics and Statistics*, 100(5), December, 831-843.
- Harvey, A.C. (1989) *Forecasting, structural time series models and the Kalman filter*. Cambridge University Press.
- Harvey, A.C., and 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. and Prescott, E.C. (1997). Postwar US business cycles: An empirical investigation. *Journal of Money, Credit, and Banking*, 29, 1-16.
- Joutz, F. & H. O. Stekler (2000). An evaluation of the predictions of the federal reserve. *International Journal of Forecasting*, 16, 17-38.
- Kitagawa, G. and Gersch, W. (1996). *Smoothness priors analysis of time series*. Springer-Verlag, New York.
- Labbe, F. and 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, New York.
- 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.

Wellington: Victoria University of Wellington. Retrieved from <http://researcharchive.vuw.ac.nz/handle/10063/2364>.

Phillips, Peter C. B. and Sainan Jin (2015). Business Cycles, Trend elimination, and the HP Filter. *Cowles Foundation Discussion Paper CFDP 2005*, June.

Phillips, Peter C. B. and Zhentao Shi (2019). Boosting: Why you Can Use the HP Filter. [arXiv:1905.00175v2](https://arxiv.org/abs/1905.00175v2) [econ.EM]

Pierce, D.A. (1980) Data revisions with moving average seasonal adjustment procedures. *Journal of Econometrics*, 14, 95-114.

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.

Whittaker, E.T. (1923) On a new method of graduation. *Proceedings of the Edinburgh Mathematical Society*, 41, 63-75.

Appendix 1. Hamilton H84 regression output for gdpe, invres, and gfcf

gdpe: Residuals:

Min	1Q	Median	3Q	Max
-0.095224	-0.017731	0.002971	0.022662	0.066344

Coefficients:

	Estimate	Std. Error	t value	Pr(> t)
(Intercept)	0.22938	0.15747	1.457	0.148412
X(t-8)	1.07624	0.28828	3.733	0.000317 ***
X(t-9)	0.10605	0.40493	0.262	0.793943
X(t-10)	0.08704	0.40542	0.215	0.830459
X(t-11)	-0.28641	0.28819	-0.994	0.322757

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Residual standard error: 0.03405 on 98 degrees of freedom

Multiple R-squared: 0.978, Adjusted R-squared: 0.9771

F-statistic: 1087 on 4 and 98 DF, p-value: < 2.2e-16

invres: Residuals:

Min	1Q	Median	3Q	Max
-0.31385	-0.11033	0.00961	0.10264	0.34645

Coefficients:

	Estimate	Std. Error	t value	Pr(> t)
(Intercept)	2.96852	0.63759	4.656	1.01e-05 ***
X(t-8)	0.62070	0.29313	2.117	0.0368 *
X(t-9)	-0.03145	0.39252	-0.080	0.9363
X(t-10)	0.11650	0.39166	0.297	0.7668
X(t-11)	-0.08291	0.28899	-0.287	0.7748

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Residual standard error: 0.1674 on 98 degrees of freedom

Multiple R-squared: 0.3796, Adjusted R-squared: 0.3543

F-statistic: 14.99 on 4 and 98 DF, p-value: 1.361e-09

gfcf: Residuals:

Min	1Q	Median	3Q	Max
-0.34634	-0.07787	0.03010	0.09908	0.20757

Coefficients:

	Estimate	Std. Error	t value	Pr(> t)
(Intercept)	0.6244	0.3597	1.736	0.08572.
X(t-8)	1.0040	0.3001	3.346	0.00116 **
X(t-9)	-0.1416	0.4251	-0.333	0.73973
X(t-10)	0.1908	0.4252	0.449	0.65468
X(t-11)	-0.1147	0.3007	-0.381	0.70368

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Residual standard error: 0.1356 on 98 degrees of freedom

Multiple R-squared: 0.8496, Adjusted R-squared: 0.8435

F-statistic: 138.4 on 4 and 98 DF, p-value: < 2.2e-16

ⁱ Turning points in deviations from trend growth cycles can be seen as approximating turning points in *economic growth rates*, whereas turning points in classical business cycles reflect turning points in the *level of aggregate economic activity* such as real gdp. For measures of New Zealand's classical business cycle recessions and recoveries, and the extent to which these differ from those derived from growth cycles, see Hall and McDermott (2016, Figure 1 and Table 2). For example, in New Zealand's post-Second World War period, 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.

ⁱⁱ For investigative work on the role of an HP1600 filter for the U.S. GFC period, see Phillips and Jin (2015).

ⁱⁱⁱ Phillips and Shi (2019, section 5) provide summary arguments in support of their bHP filter and detailed responses to Hamilton's critique of the HP filter.

^{iv} Volatility, persistence and cross correlation results for Hamilton's H8 diff procedure are not presented in Tables 1 and 2, as it is clear from the H84 and H8 diff time paths that move very closely together in Figures 1 to 3 that their statistical magnitudes will be very similar and have the same overall characteristics.

^v For the three business cycle sample periods investigated in Section 4, the volatilities produced from H84 detrending of log gdp are 2.6 times greater than those obtained from HP1600 detrending. Similarly, Schuler (2018, p11, Table 1) reports that, for the U.S over the period 1953q4 to 2017q2, the volatility of Hamilton two-year regression-filtered log GDP is more than double that of HP1600-filtered series (3.25 relative to 1.50). This 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).




Figure 1. Logarithms of gdpe, trends, and trend deviations. (a) Top panel shows log gdpe (black), HP1600 trend (red), H84 trend green and H8 diff trend (blue). (b) Bottom panel shows gdpe trend deviations for HP1600 (red), H84 (green) and H8 diff (blue).

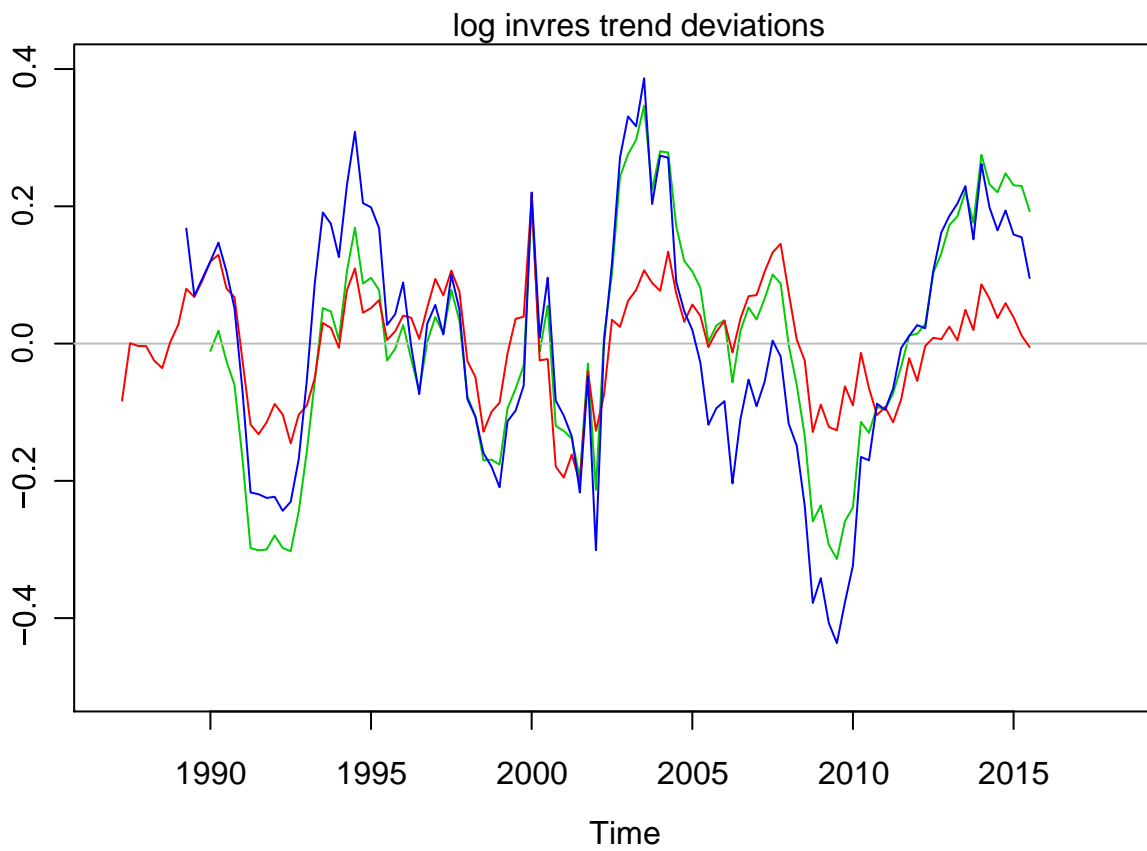
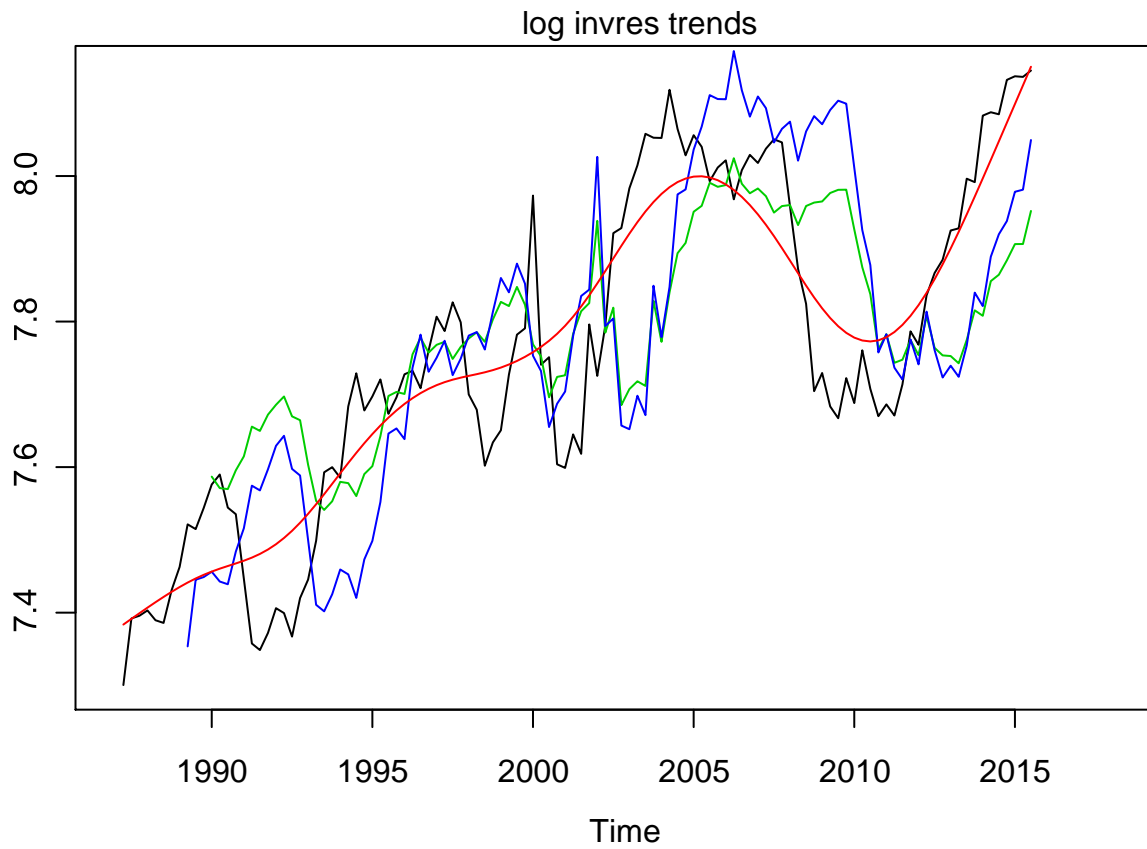


Figure 2. Logarithms of invres, trends, and trend deviations. (a) Top panel shows log invres (black), HP1600 trend (red), H84 trend green and H8 diff trend (blue). (b) Bottom panel shows invres trend deviations for HP1600 (red), H84 (green) and H8 diff (blue).





Figure 3. Logarithms of gfcf, trends, and trend deviations. (a) Top panel shows log gfcf (black), HP1600 trend (red), H84 trend green and H8 diff trend (blue). (b) Bottom panel shows gfcf trend deviations for HP1600 (red), H84 (green) and H8 diff (blue).




Figure 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), and naïve predictor (grey).




Figure 5. Forecast-extended HP1600 trends (top) and their trend deviations (bottom) for log GDPP data to 2015q3 (top dotted) over the 8 quarter assessment window from 2013q4 to 2015q3. Also shown are the target HP1600 trend (top black) based on ex-post log GDPP 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), naïve predictor (grey) and no extension (magenta). The latter corresponds to using the HP filter with no extension.




Figure 6. 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), and naïve predictor (grey).

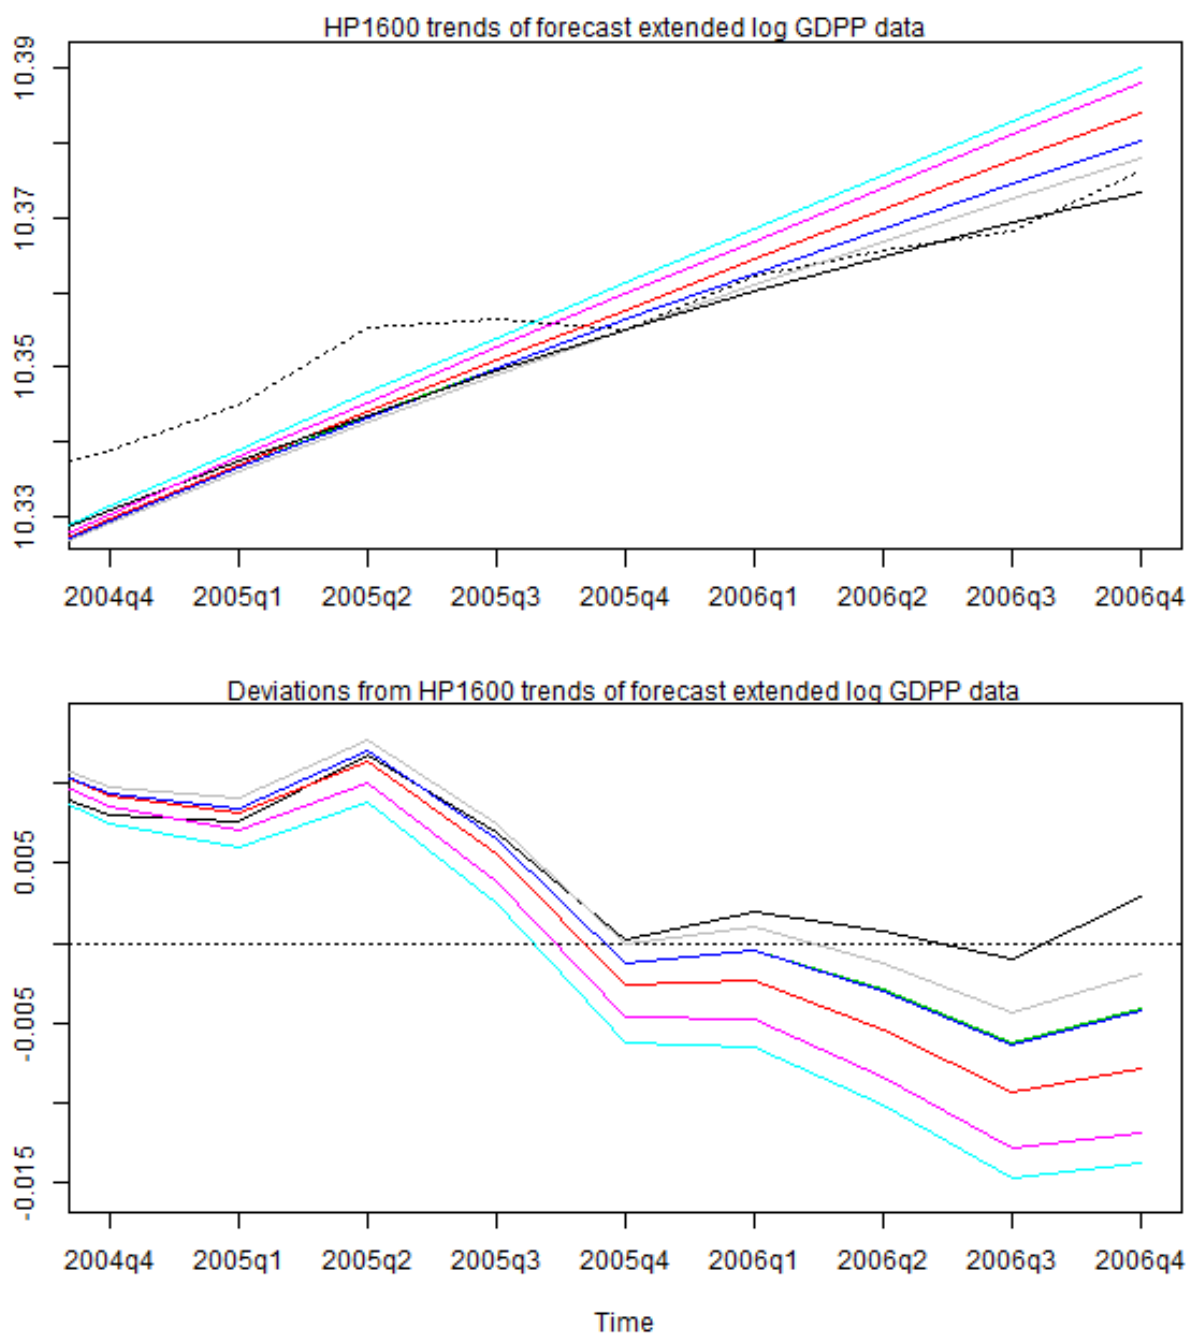


Figure 7. Forecast-extended HP1600 trends (top) and their trend deviations (bottom) for log GDP data to 2006q4 (top dotted) over the 8 quarter assessment window from 2005q1 to 2006q4. Also shown are the target HP1600 trend (top black) based on ex-post log GDP data to 2019q4 and its corresponding target trend deviation (bottom black). The forecast extensions used are RBNZ (red), Treasury (green), NZIER (blue), Hamilton H84 predictor (cyan), naïve predictor (grey) and no extension (magenta). The latter corresponds to using the HP filter with no extension.

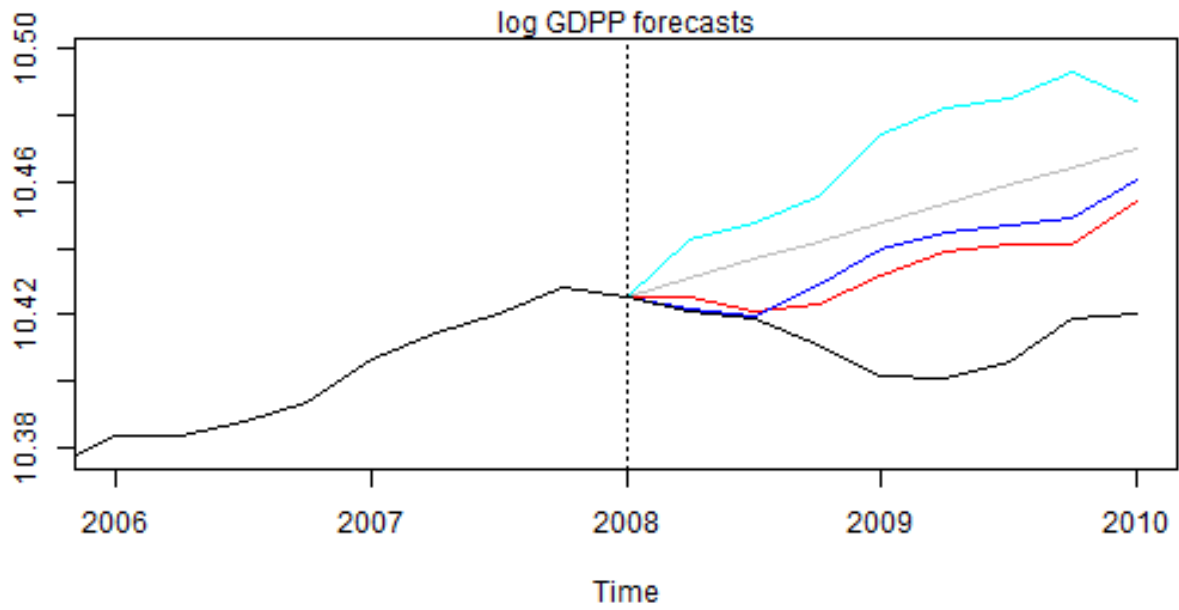


Figure 8. Log GDPP (black) and log GDPP forecasts based on data to 2008q1 and an 8 quarter forecast window: RBNZ (red), NZIER (blue), Hamilton H84 predictor (cyan), and naïve predictor (grey).




Figure 9. Forecast-extended HP1600 trends (top) and their trend deviations (bottom) for log GDPP data to 2008q1 (top dotted) over the 8 quarter assessment window from 2006q2 to 2008q1. Also shown are the target HP1600 trend (top black) based on ex-post log GDPP data to 2019q4 and its corresponding target trend deviation (bottom black). The forecast extensions used are RBNZ (red), NZIER (blue), Hamilton H84 predictor (cyan), naïve predictor (grey) and no extension (magenta). The latter corresponds to using the HP filter with no extension.

Table 1. Stylised Business Cycle Facts, 1987q2 - 2015q3: Comparative Volatilities, Relative volatilities, and Persistence

Variable x	Volatility				Relative Volatility		Persistence			
	HP1600		H84		HP1600	H84	HP1600		H84	
GDP (production)	1.33	(0.23)	3.44	(0.55)			0.86	(0.04)	0.85	(0.05)
GDP (expenditure)	1.38	(0.19)	3.18	(0.62)	1.00	1.00	0.71	(0.06)	0.87	(0.05)
Consumption (private)	1.49	(0.20)	3.15	(0.73)	1.08	0.99	0.78	(0.06)	0.86	(0.05)
Gross Fixed Capital Formation	5.89	(0.75)	13.26	(2.36)	4.27	4.17	0.76	(0.06)	0.88	(0.04)
Investment (residential)	8.26	(0.91)	16.18	(3.22)	5.99	5.09	0.77	(0.06)	0.89	(0.04)
Govt. Consumption expend.	1.36	(0.16)	3.37	(0.66)	0.99	1.06	0.50	(0.08)	0.85	(0.05)
Gross Govt. Debt/GDP**	2.67	(0.61)	6.02	(0.85)	1.93	1.89	0.82	(0.06)	0.84	(0.05)
Net Exports Share*	1.30	(0.16)	2.38	(0.40)	0.94	0.75	0.70	(0.07)	0.88	(0.04)
Imports goods & services*	4.50	(0.61)	7.90	(1.68)	3.26	2.48	0.77	(0.06)	0.87	(0.05)
Employment	1.31	(0.19)	2.83	(0.38)	0.95	0.89	0.90	(0.04)	0.84	(0.05)
Unemployment	0.63	(0.13)	1.27	(0.26)	0.46	0.40	0.88	(0.04)	0.88	(0.04)
CPI Nontradables (ann. % change)***	0.90	(0.17)	1.05	(0.12)	0.65	0.33	0.76	(0.06)	0.83	(0.06)
CPI Tradables (annual % change)***	1.41	(0.22)	1.80	(0.29)	1.02	0.57	0.71	(0.07)	0.83	(0.06)
Real 90-day Bank Bill****	0.88	(0.13)	1.49	(0.29)	0.64	0.47	0.68	(0.07)	0.84	(0.05)

Notes:

Volatility is % standard deviation; relative volatility is relative to GDP(expenditure) volatility; 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 (as for NZ Treasury series) for frigate purchases recorded in 1997q2 and 1999q4.

The series not adjusted in this way show somewhat greater volatilities and less persistence, e.g. for HP1600 filtered series, imports of goods & services and net exports share volatilities are 4.74% and 1.34%, and persistences are 0.70 and 0.31.

** NZ Treasury series: Gross Govt. Debt/GDP sample period 1987q2 - 2013q3

*** Sample period 1989q1 - 2015q3

**** Sample period 1987q3 - 2015q3

Table 2. Stylised Business Cycle Facts, 1987q2 - 2015q3: Comparative Cross Correlations with real GDPE

Variable x	Contemporaneous Cross Correlation		Most significant Non-contemporaneous Cross Correlation			
	HP1600	H84	HP1600		H84	
Consumption (private)	0.73 (0.16)	0.85 (0.21)	-		-	
Gross Fixed Capital Formation	0.77 (0.15)	0.83 (0.18)	-		-	
Investment (residential)	0.72 (0.12)	0.71 (0.15)	-		-	
Govt. Consumption expend.	0.10 (0.11)	0.19 (0.14)	0.44 (0.13)	x_{t+5}	0.57 (0.15)	x_{t+5}
Gross Govt. Debt/GDP**	-0.51 (0.16)	-0.70 (0.16)	-0.53 (0.16)	x_{t+3}	-0.73 (0.17)	x_{t+3}
Net Exports Share*	-0.24 (0.10)	-0.34 (0.14)	-0.50 (0.14)	x_{t+2}	-0.52 (0.16)	x_{t+2}
Imports goods & services*	0.44 (0.11)	0.61 (0.15)	0.56 (0.14)	x_{t+2}	0.70 (0.16)	x_{t+2}
Employment	0.51 (0.14)	0.62 (0.14)	0.54 (0.14)	x_{t+2}	0.61 (0.15)	x_{t+2}
Unemployment	-0.60 (0.15)	-0.63 (0.17)	-0.61 (0.15)	x_{t+1}	-0.65 (0.18)	x_{t+1}
CPI Nontradables (ann. % change)***	0.30 (0.17)	0.17 (0.11)	0.61 (0.16)	x_{t+4}	0.40 (0.13)	x_{t+3}
CPI Tradables (annual % change)***	-0.18 (0.13)	-0.15 (0.09)	-		-0.14 (0.12)	x_{t-3}
Real 90-day Bank Bill****	0.38 (0.14)	0.28 (0.14)	0.49 (0.13)	x_{t+2}	0.41 (0.13)	x_{t+3}

Notes

Numbers in parentheses are robustly estimated standard errors

"-" Denotes most significant cross correlation is Contemporaneous

* SNZ National Accounts series, adjusted (as for NZ Treasury series) for frigate purchases recorded in 1997q2 and 1999q4.

The series not adjusted in this way show somewhat weaker (or relatively similar) cross correlations with real GDPE, e.g. for HP1600 filtered series, the statistically significant correlations for imports of goods and services and for net exports share are 0.54 (x_{t+2}) and -0.49(x_{t+2}).

** New Zealand Treasury Series: Gross Govt. Debt/GDP sample period 1987q2 - 2013q3

*** Sample period 1989q1 - 2015q3

**** Sample period 1987q3 - 2015q3

Table 3a. Percentage error measures for log GDPP forecasts, based on data to 2015q3 and an 8 quarter forecast window

Forecaster	Mean	MAE	RMSE	RMSE Rank
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	5
Naive	-0.21	0.23	0.32	1

Table 3b. Percentage error measures for log GDPP forecasts, based on data to 2006q4 and an 8 quarter forecast window

Forecaster	Mean	MAE	RMSE	RMSE Rank
RBNZ	1.22	1.49	2.15	4
Treasury	0.34	0.82	1.12	1
NZIER	0.39	1.17	1.46	3
H84	2.44	2.44	2.73	5
Naive	-0.11	1.08	1.23	2

Table 3c. Percentage error measures for log GDPP forecasts, based on data to 2008q1 and an 8 quarter forecast window

Forecaster	Mean	MAE	RMSE	RMSE Rank
RBNZ	2.23	2.23	2.61	1
NZIER	2.64	2.64	3.13	2
H84	5.83	5.83	6.23	4
Naive	3.82	3.82	4.12	3

Notes: Mean=Mean average error; MAE=Mean absolute error; RMSE=Root mean squared error

Table 4a. Percentage error measures for the differences between forecast-extended HP1600 log GDP trends and the target HP1600 trend over the 8 quarter assessment window from 2013q4 to 2015q3

Forecaster	Mean	MAE	RMSE	RMSE Rank
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	6
Naive	-0.10	0.10	0.11	1
HP1600 no extension	-0.44	0.44	0.48	5

Table 4b. Percentage error measures for the differences between forecast-extended HP1600 log GDP trends and the target HP1600 trend over the 8 quarter assessment window from 2005q1 to 2006q4

Forecaster	Mean	MAE	RMSE	RMSE Rank
RBNZ	0.42	0.43	0.56	4
Treasury	0.24	0.26	0.35	2=
NZIER	0.24	0.27	0.36	2=
H84	0.82	0.82	0.96	6
Naive	0.11	0.18	0.23	1
HP1600 no extension	0.66	0.66	0.81	5

Table 4c. Percentage error measures for the differences between forecast-extended HP1600 log GDP trends and the target HP1600 trend over the 8 quarter assessment window from 2006q2 to 2008q1

Forecaster	Mean	MAE	RMSE	RMSE Rank
RBNZ	0.78	0.78	0.84	1
NZIER	0.87	0.87	0.95	2
H84	1.70	1.70	1.87	5
Naive	1.19	1.19	1.30	3
HP1600 no extension	1.60	1.60	1.75	4

Notes: Mean=Mean average error; MAE=Mean absolute error; RMSE=Root mean squared error.

Table 5. Percentage reduction in RMSE of forecast-extended HP1600 filter at the ends of series compared to the HP1600 filter with no extension.

Forecast extension	NTP	TPP	TPT
RBNZ	43	30	52
Treasury	17	57	-
NZIER	25	56	46
H84	-4	-18	-7
Naive	78	71	26