BIS

BIS Working Papers
No 1033

The Holt-Winters filter and the one-sided HP filter: A close correspondence

by Rodrigo Alfaro and Mathias Drehmann

Monetary and Economic Department

July 2022

This publication is available on the BIS website (www.bis.org).

ISSN 1020-0959 (print)
ISSN 1682-7678 (online)

The Holt-Winters filter and the one-sided HP filter: A close correspondence*

Rodrigo Alfaro
Central Bank of Chile

Mathias Drehmann
Bank for International Settlements

Abstract

We show that the trend of the one-sided HP filter can be asymptotically approximated by the Holt-Winters (HW) filter. The later is an elegant, moving average representation and facilitates the computation of trends tremendously. We confirm the accuracy of this approximation empirically by comparing the one-sided HP filter with the HW filter for generating credit-to-GDP gaps. We find negligible differences, most of them concentrated at the beginning of the sample.

*We would like to thank Stijn Claessens and James Yetman for helpful comments. The views expressed herein are those of the authors and should not be attributed to the Central Bank of Chile or to the Bank for International Settlements.

Rodrigo Alfaro: ralfaro@bcentral.cl (corresponding author); Mathias Drehmann: mathias.drehmann@bis.org.

1 Introduction

The Hodrick and Prescott (1980, 1997; HP hereafter) filter is heavily used in economics and by policy makers. In many cases, the two-sided, full sample version is applied to determine trends. But the one-sided filter that is estimated recursively has also important applications.[1] Most notably, it is embedded in the Basel III framework for the countercyclical capital buffer (Basel Committee, 2010), where the credit gap forms the starting point for discussions of the setting of the buffer.[2] The credit gap, in turn, is calculated as the difference of the credit-to-GDP ratio from a long-term trend computed by a one-sided HP filter.

In this paper, we show that the trend of the one-sided HP filter can be asymptotically approximated by the Holt-Winters (HW) filter. This implies an elegant, moving average representation. And it also facilitates the computation of trends tremendously. Instead of running the standard HP filter recursively with an expanding sample, all what is needed is a simple lag operation. Ultimately, computation is so simple that it can be implemented by any spreadsheet user.

To derive this result, we combine several insights from the literature. It is well known that applying the (two-sided) HP filter is equivalent to a Kalman smoother of a local-linear-trend (LLT) model (Hodrick and Prescott, 1980, 1997).[3] Somewhat less known is the equivalence between the one-sided HP filter and the Kalman filter (ie the one-step ahead forecasts) of the LLT model (Hamilton, 2018). At the same time, the Holt-Winters (HW) filter is equivalent to the steady-state Kalman filter of the LLT model (Harvey, 1989). As such, the one-sided HP filter can be very closely approximated by the HW filter, that in turn, has a closed-form solution.

For practical purposes, we also highlight three further results. First, we derive a mapping between the smoothing parameter of the HP filter (typically called λ) and the key parameters of the HW filter. Second, we provide an exact formula for the weights of the HW filter, that complements the result for the asymptotic two-sided HP filter reported in Hodrick and Prescott (1980, 1997). Third, we discuss different

[1] The one-sided HP filter is also known as the real-time HP filter because it is derived by a standard HP filter that is applied recursively using an increasing sample-size. This mirrors how new observations would have arrived if the filter would have been applied in a live setting.

[2] The use of the one-sided HP filter with a large smoothing parameter lambda to uncover the build-up of credit booms was first proposed by Borio and Lowe (2002).

[3] Given that the HP filter is defined over the second difference of the trend component, it should be classified as a smoothed trend model which, in turn, is a particular case of the LLT model (StataCorp, 2017).

ways to set the initial values for the HW filter. And we find that there is quick convergence independent of the start point decision.

We illustrate our approach by generating credit-to-GDP gaps using the different filters. Unsurprisingly, we find that the HW filter closely approximates the one-sided HP filter. There are some differences at the beginning of the sample, in particular for countries which experienced large swings in the credit-to-GDP ratio at that point. In addition, after major swings in the credit-to-GDP ratio some small difference between both filters also open up, such as in Spain after the Great Financial Crisis (GFC).

In the remainder of the paper, we first show formally the link between the HW and the HP filter in Section 2. In Section 3, we apply the filters to generate credit-to-GDP gaps. And in the last section we conclude.

2 Filters and Structural Models

In this section, we cast the one-sided Hodrick-Prescott (HP) filter as the structural time-series models following Harvey (1989) to show that the one-sided filter can be approximated by the Holt-Winters (HW) filter.

2.1 The one-sided Hodrick-Prescott filter

To understand the one-sided HP filter, we need to consider the original, two-sided HP filter. This filter assumes that the original series (x_t) can be divided into two additive components: the trend and the cycle. The trend can be obtained by solving the following problem:

$$\min_{\{g_t\}_{t=1}^{T}} \frac{1}{2}\left[\sum_{t=1}^{T}(y_t - g_t)^2 + \lambda \sum_{t=2}^{T-1}(g_{t+1} - 2g_t + g_{t-1})^2\right] \tag{1}$$

where λ is the smoothing parameter.

The first term in the loss function penalizes the variance of the cyclical component, while the second puts a penalty on the lack of smoothness in the trend. Hence, the solution to the problem is a trade-off between the smoothness of the trend and how well it fits the original series.

Taking first order conditions of (1) we have for the central cases (ie $2 < t < T-1$): $-(y_t - g_t) + \lambda(g_t - g_{t-1} + g_{t-2}) - 2\lambda(g_{t+1} - 2g_t + g_{t-1}) + \lambda(g_{t+2} - 2g_{t+1} + g_t) = 0.$

Using $f(L)$ as a function of lag operators, we can arrange the previous expression as: $y_t = g_t[1 + \lambda f(L)]$, where $f(L) = L^{-2} - 4L^{-1} + 6 - 4L + L^2$.[4] For the first and last observations, we have different expressions: $-(y_1 - g_1) + \lambda(g_3 - 2g_2 + g_1) = 0$, and $-(y_T - g_T) + \lambda(g_T - 2g_{T-1} + g_{T-2}) = 0$. Similar expressions arise for the second and for $T - 1$ observations. This is known as the end-point issue.

More generally, if we define $\mathbf{y} = (y_T, \ldots, y_1)'$ and $\mathbf{g} = (g_T, \ldots, g_1)'$ then a general solution can be written in matrix form as $\mathbf{g} = \mathbf{Hy}$, where $\mathbf{H} = (\mathbf{I_T} + \lambda\mathbf{F})^{-1}$, where $\mathbf{F}$ is the following $T \times T$ matrix (Kim, 2004):

$$\mathbf{F} = \begin{bmatrix} 1 & -2 & 1 & 0 & 0 & 0 & 0 & & & & & & \\ -2 & 5 & -4 & 1 & 0 & 0 & 0 & & & & & & \\ 1 & -4 & 6 & -4 & 1 & 0 & 0 & & & & & & \\ 0 & 1 & -4 & 6 & -4 & 1 & 0 & & & & & & \\ & & & & & & & & & & & & \\ & & & & & & & 1 & -4 & 6 & -4 & 1 & 0 \\ & & & & & & & 0 & 1 & -4 & 6 & -4 & 1 \\ & & & & & & & 0 & 0 & 1 & -4 & 5 & 2 \\ & & & & & & & 0 & 0 & 0 & 1 & -2 & 1 \end{bmatrix}$$

Given that the one-sided HP filter takes account of information available at the current time, only the first row of the matrix $\mathbf{H}$ is relevant. We will return to this in order to get a numerical approximation for the weights of the one-sided HP filter.

2.2 Local linear trend model

The LLT model assumes that the variable of interest (x_t) can be decomposed as follows: $x_t = \alpha_t + e_t$, where α_t is a stochastic trend and e_t is the cycle component with mean zero and variance σ_e^2. The dynamics of the trend are defined by a system of state variables:

$$\begin{pmatrix} \alpha_t \\ \beta_t \end{pmatrix} = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} \alpha_{t-1} \\ \beta_{t-1} \end{pmatrix} + \begin{pmatrix} u_t \\ v_t \end{pmatrix}.$$

[4] Hodrick and Prescott (1980,1997) consider this expression as the main argument of the filter, discussing its asymptotic properties. In particular, weights of the asymptotic two-sided filter can be expressed in closed-form: $w_i = 0.8941^i[0.056168 \cdot \cos(0.11168 \cdot i) + 0.055833 \cdot \sin(0.11168 \cdot i)]$.

where u_t and v_t are distributed independently of each other and over time with mean zero and variances σ_u^2 and σ_v^2, respectively.

Three features of the LLT model are relevant for our discussion:

- Estimates of α_t obtained by applying the Kalman smoother with parameters $\sigma_u^2 = 0$ and $\sigma_v^2/\sigma_e^2 = 1/\lambda$ are equivalent to the one computed by the standard HP filter (Hodrick and Prescott, 1980, 1997; Harvey and Trimbur, 2008).
- Estimates of α_t obtained by applying the Kalman filter (one-step ahead forecast) with parameters $\sigma_u^2 = 0$ and $\sigma_v^2/\sigma_e^2 = 1/\lambda$ are equivalent the one-sided HP filter (Hamilton, 2018). This implies that the one-sided HP filter is optimal when the variable of interest is well approximated by the LLT model.
- Estimates of α_t obtained by applying the Kalman filter (one-step ahead forecast) can be approximated by using the HW filter because the HW filter is equivalent to the steady-state Kalman filter of the LLT model (Harvey, 1989).

2.3 Holt-Winters filter

Under the HW filter the trend estimate (g_t) for a given series x_t is $g_t = a_{t-1} + b_{t-1}$, where: $a_t = \theta_1 x_t + (1-\theta_1)(a_{t-1} + b_{t-1})$, and $b_t = \theta_2(a_t - a_{t-1}) + (1-\theta_2)b_{t-1}$. These state variables are associated with the intercept and slope of the LLT model (Durbin and Koopman, 2008). Replacing a_t in the slope equation, we get $b_t = \theta_1\theta_2(x_t - a_{t-1}) + (1-\theta_1\theta_2)b_{t-1}$. Hence, for $\theta_0 = \theta_1\theta_2$ we have $b_t = \theta_0(x_t - a_{t-1}) + (1-\theta_0)b_{t-1}$. Summarizing in matrix form:

$$\begin{pmatrix} a_t \\ b_t \end{pmatrix} = \begin{pmatrix} 1-\theta_1 & 1-\theta_1 \\ -\theta_0 & 1-\theta_0 \end{pmatrix} \begin{pmatrix} a_{t-1} \\ b_{t-1} \end{pmatrix} + \begin{pmatrix} \theta_1 \\ \theta_0 \end{pmatrix} \cdot x_t.$$

Thus, for initial values a_1 and b_1, the trend estimate can be obtained recursively.

Harvey (2001) shows that he parameters θ_1 and θ_0 are consistent with steady-state Kalman filter solution of the LLT model if $\sigma_u^2/\sigma_e^2 = (\theta_1^2 + \theta_1\theta_0 - 2\theta_0)/(1-\theta_1)$ and $\sigma_v^2/\sigma_e^2 = \theta_0^2/(1-\theta_1)$. The HP filter imposes $\sigma_u = 0$. Hence, $\theta_0 = \theta_1^2/(2-\theta_1)$ and only one parameter is needed in the HW filter to accommodate the one-sided HP filter. Indeed, we have the following relationship between the parameters of both filters:

$$\lambda = \frac{(2-\theta_1)^2(1-\theta_1)}{\theta_1^4} \tag{2}$$

2.4 Practical considerations of the HW filter

To apply the HW filter, the coefficient θ_1 and the starting values for a_1 and b_1 have to be known.

Setting θ_1 is simple given the mapping to the smoothing parameter λ of the HP filter (equation 2). In typical business cycle application, λ is set to 1600 for quarterly data. In contrast, the credit-to-GDP gap aims to uncover more medium term fluctuations with λ equalling 400,000. This implies that the credit cycle is assumed to be four times longer than the business cycle [5] Hence:

- $\theta_1 = 0.20055$ for business cycles (ie $\lambda = 1,600$)
- $\theta_1 = 0.0547$ for credit cycles (ie $\lambda = 400,000$)

Given $\theta_0 = \theta_1^2/(2-\theta_1)$ and $\theta_2 = \theta_0/\theta_1$, we have $\theta_0 = 0.00153811$ and $\theta_2 = 0.02811906$ when looking at credit cycles.

There are many ways to obtained estimates for starting values for a_1 and b_1, even though convergence is relatively quick. In the applications below, we concentrate our discussion on: (i) using the initial data, and (ii) fitting a linear trend, estimated with the first years of data (5 or 10). And we find that fitting a linear trend with 10 years of data works best. That said, convergence between the different approaches is relatively quick.

2.5 Weights of the HW filter

The recursion system of the HW filter can be solved backward implying that the trend estimate is a weighted average of past observations. Thus we can expressed the weights of this filter using the matrix of the state system as follows:

$$w_k = \begin{pmatrix} 1 & 1 \end{pmatrix} A^k \begin{pmatrix} \theta_1 \\ \theta_0 \end{pmatrix}, \text{ where } A = \begin{pmatrix} 1-\theta_1 & 1-\theta_1 \\ -\theta_0 & 1-\theta_0 \end{pmatrix}$$

Matrix A can be expressed using an alternative matrix (C) with similar eigenvalues (see Appendix A), therefore we can apply the following decomposition: $A = PCP^{-1}$.

[5] Ravn and Uhlig (2002) analyse how λ has to be adjusted if the frequency of the data changes (eg annual verus quarterly). They show for the two-sided filter that it is optimal to multiply λ with the forth power of the observation frequency ratio. Applying this insight implies that $400,000(\lambda_{credit-cycle}) \approx 4^4 * 1,600(\lambda_{business-cycle})$

Based on that we have $w_k = (1,1)PC^kP^{-1}(\theta_1, \theta_0)'$ which it has a closed-form expression because C is defined using trigonometric functions.

For example for $\theta_1 = 0.0547$ ($\lambda \approx 400,000$) we have the following results:

$$A = \begin{bmatrix} 0.945300 & 0.945300 \\ -0.001538 & 0.994862 \end{bmatrix}, \text{ and } P = \begin{bmatrix} 0.999187 & 0.000000 \\ 0.028096 & -0.028898 \end{bmatrix}$$

$$C = \begin{bmatrix} 0.971881 & -0.027339 \\ 0.027339 & 0.971881 \end{bmatrix} = 0.972265 \begin{bmatrix} \cos(0.028123) & -\sin(0.028123) \\ \sin(0.028123) & \cos(0.028123) \end{bmatrix}$$

As matrix A has complex eigenvalues, the factorization using C is convenient in the sense that is populated with the real and imaginary parts of eigenvalues (diagonal and off-diagonal). Thus, we have the following expression for the weights when $\theta_1 = 0.0547$:[6]

$$w_k = 0.972265^k[0.05623811 \cdot \cos(0.028123 \cdot k) - 0.001582 \cdot \sin(0.028123 \cdot k)] \quad (3)$$

Figure 1 shows weights of one-sided HP filter with 100 and 200 observations and considering $\lambda = 400,000$. Weights of the HW filter (w_k) are also included with $\theta_1 = 0.0547$. We note that the weights of the HW filter are well approximated by the one-sided HP filter discussed before when $T = 200$.

3 Empirical application

In this section we apply the HW filter to derive credit-to-GDP gaps and compare the results with the gaps based on the one-sided HP filter. We use quarterly data for the credit-to-GDP ratio from the BIS.[7] The data covers 44 countries starting at the earliest in the last quarter of 1947 up to the third quarter of 2021. In total, we have more than 8600 observations.

We calculate in total 4 different gaps. First, the gaps based on the one-sided HP

[6] It is interesting to note that the coefficient in the sine term is almost zero. This is different to the result reported in Hodrick and Prescott (1980, 1997) for the two-sided HP filter. There, the coefficient in the sine term is similar to the one in the cosine term.

[7] More specifically, we use the ratio of total credit to the non-financial private sector to GDP. Data can be found on the BIS website: www.bis.org/statistics/totcredit.htm.

Figure 1: Weights of one-sided HP filters and HW filter

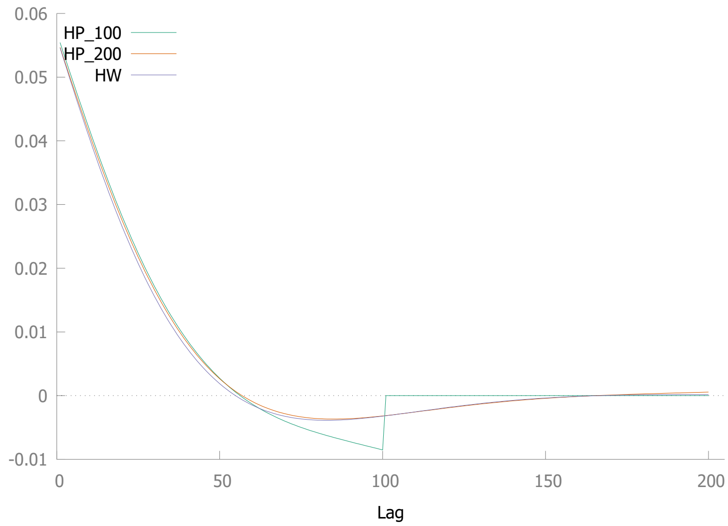

HP100: one-sided HP filter with $T = 100$ and $\lambda = 400,000$. HP200: one-sided HP filter with $T = 200$ and $\lambda = 400,000$. HW: HW filter with $\theta_1 = 0.0547$.

filter with $\lambda = 400,000$. Second, three different gaps based on the HW filter. In all of these cases $\theta_1 = 0.0547$. But the gaps differ in the starting values. For HW_{init} we use initial values and stet $a_{1,init} = credit_{t0}/GDP_{t0}$ where $t0$ is the country specific starting point for the credit-to-GDP ratio. The slope parameter is set for all countries to $b_{1,init} = 0$. For the other two, a_1 and b_1 are set to the constant and slope coefficient respectively from country specific regressions of the credit-to-GDP ratio on a time trend. Regressions use either the first five years or first ten years of data (HW_{5y} or HW_{10y}).

Before comparing the HW and the HP filter, it is useful to compare the convergence of the trends for different starting values for the HW filter. Figure 2 shows the average difference of the trend based on HW_{10y} from HW_{5y} or HW_{init}, relative to the HW_{10y}-trend. We also show the 5th /95th percentiles as shaded areas. Given that we use 10 years of data to compute HW_{10y}, the deviations are only shown from quarter 40 after the first observation of the credit-to-GDP ratio in each country.[8] In addition, the dotted line indicates the 20 year horizon, indicating the horizon when one full cycle has been completed (recall the very large value of λ assumes that the credit cycle is four times longer than the business cycle).

[8]Before the 10 year horizon, differences are at the maximum 30 percent for Korea.

Figure 2 shows that there is close convergence, in particular after one full cycle has been observed. Differences are larges at the beginning. Looking more closely at the underlying data, this is the case for countries, many of which are EMEs, that did not experience stable growth in the credit-to-GDP ratio but rather some large fluctuations at the early part of the sample (eg., Mexico, see Figure 4).

Turning to the main interest of the paper, a look at the country level indicates that differences between the HW and the HP filter are very small, which is unsurprising given that the HW filter approximates the one-sided HP filter (Figures 3 and 4). In particular in advanced economies, differences between the HW_{10y} and the HP filter are negligible. However, they do not only emerge at the beginning of the sample but also later after strong turning-points in the credit-to-GDP ratio such as after the GFC in Spain and to a lesser extent in the UK. For EMEs, differences are somewhat larger but again mainly at the beginning of the sample where several economies experienced relatively big swings in the credit-to-GDP ratio.

Figure 5 summarises this insights more succinctly.[9] Mirroring Figure 2, we show the difference of the trend based on HW_{10y} from the HP trend (relative to the HP-trend) as well as the 5th/95th percentiles as shaded areas. We can see that there are very small differences between both filters. While they are largest at the beginning of the sample before one full credit cycle has been observed, differences remain throughout the sample. As discussed for Spain, this reflects small differences after major swings in the credit-to-GDP ratio.

[9] In the main text, we concentrate on the trend based on HW_{10y}. A comparison with Figure 6 in Appendix B shows that convergence is strongest in the 10-year case.

Figure 2: Convergence of the HW filter for different initial conditions

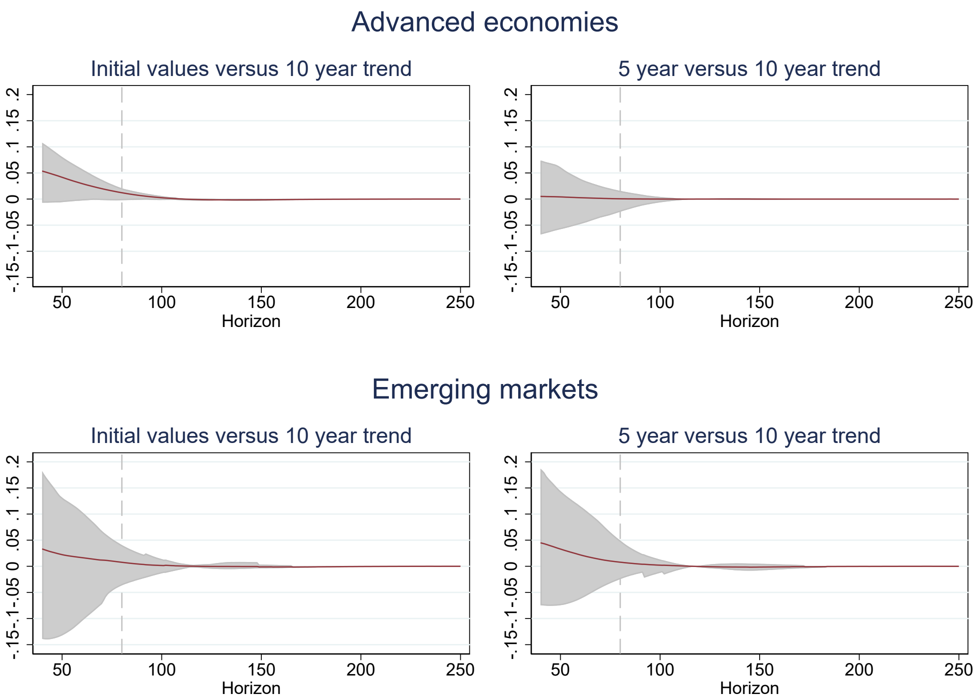

Solid line: Mean difference of the HW trend based on HW_{10y} versus HW_{init} or HW_{5y}. All relative to the HW_{10} trend. Shaded areas: 5th and 95th percent confidence bands. Horizon: quarters after the first observation of the credit-to-GDP ratio. Dotted line: quarter 80, indicating the horizon when one full cycle has been completed

Figure 3: The credit-to-GDP ratio, trends and gaps for selected advanced economies

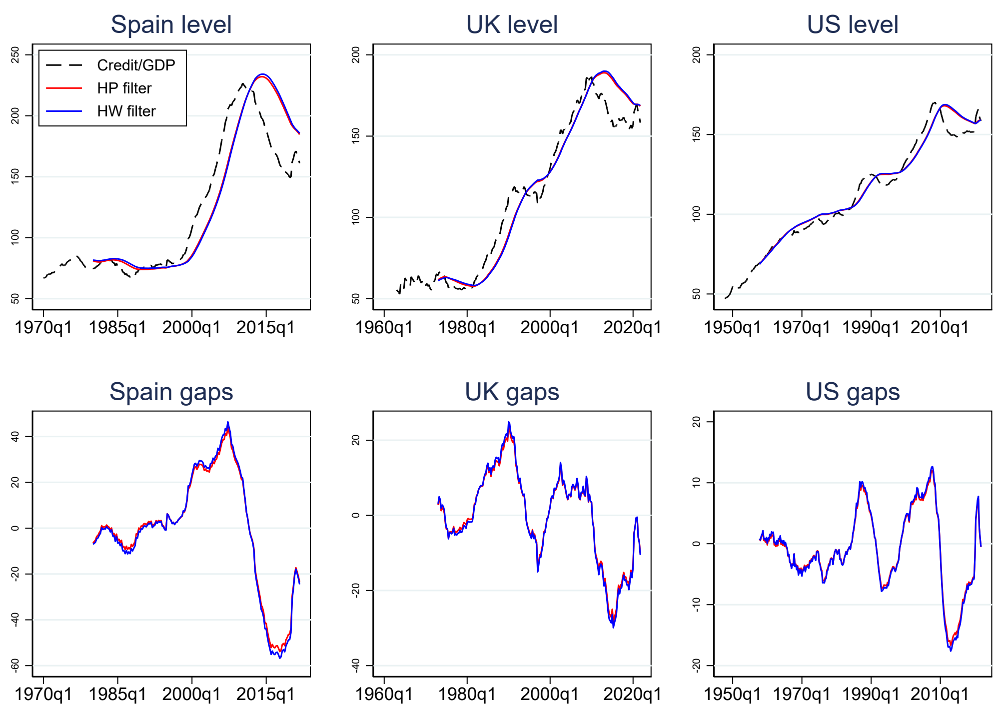

HP filter based on one-sided filter with $\lambda = 400,000$. HW filter based on HW_{10y} with $\theta_1 = 0.0547$.

Figure 4: The credit-to-GDP ratio, trends and gaps for selected emerging markets

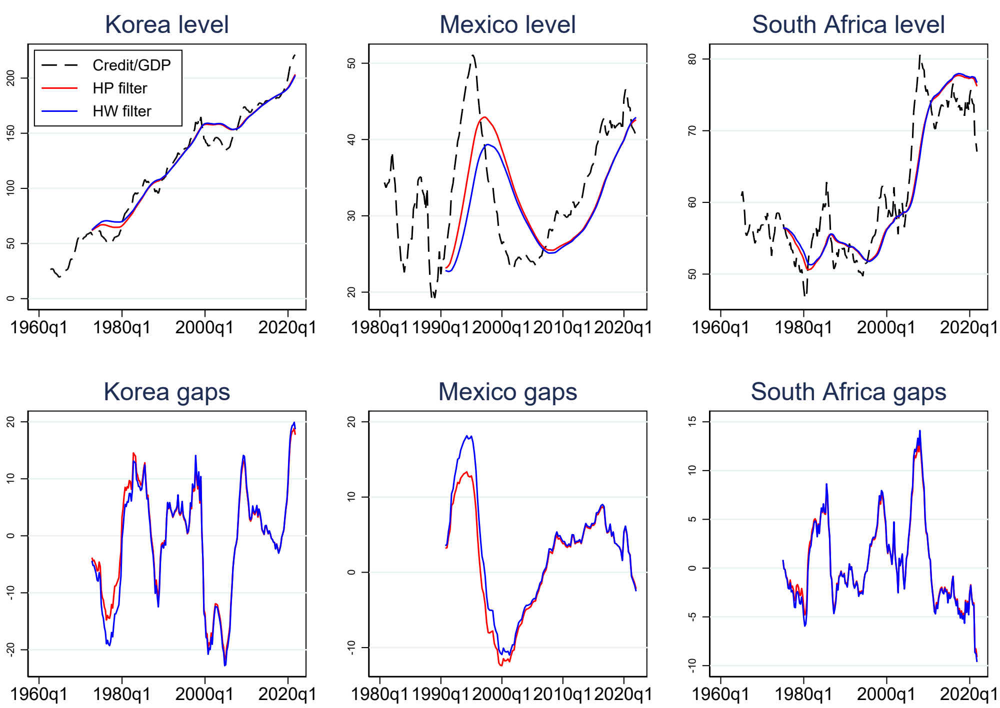

HP filter based on one-sided filter with $\lambda = 400,000$. HW filter based on HW_{10y} with $\theta_1 = 0.0547$.

Figure 5: HP versus HW filter

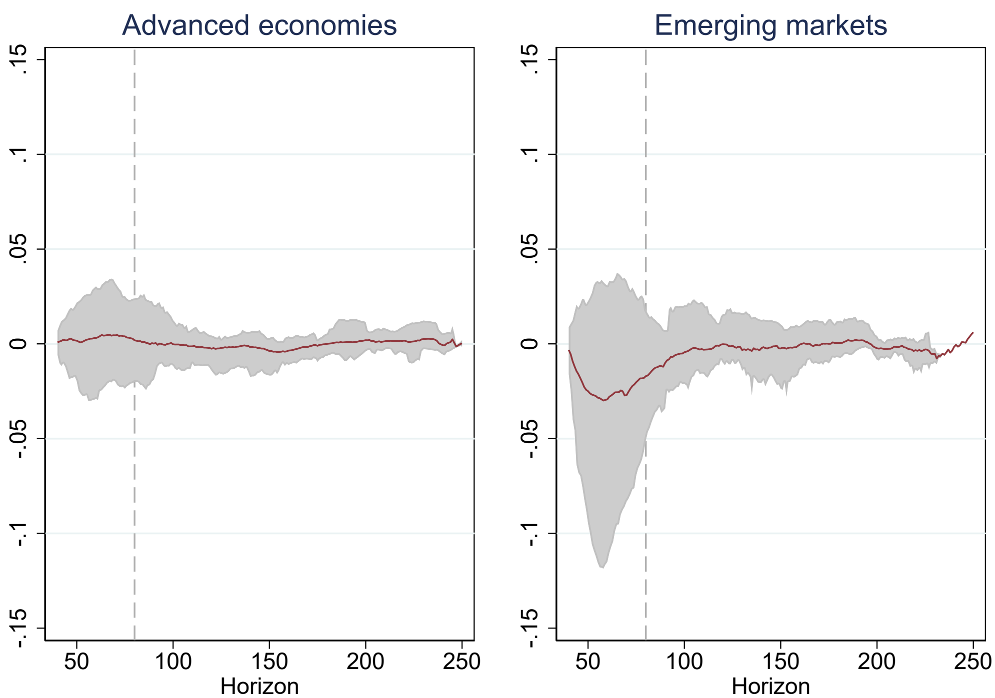

Solid line: Mean difference of the HW trend based on HP filter versus HW_{10} relative to the HP filtered trend. Shaded areas: 5th and 95th percent confidence bands. Horizon: quarters after the first observation of the credit-to-GDP ratio. Dotted line: quarter 80, indicating the horizon when one full cycle has been completed

4 Concluding remarks

In this paper we show that the Holt-Winters filter can closely approximate the one-sided Hodrick-Prescott filter. The reason is that both filters are based on Kalman filter solutions of the local linear trend model. Thus, the approximation relies on the existence of the steady-state solution of the Kalman filter, which is guaranteed under positive smoothing parameters. We also provide analytical expression for the weights of the Holt-Winters filter comparing those with ones obtained from the one-sided Hodrick-Prescott filter for a large sample-size. Hence, there is an asymptotic relationship between both filters.

As an empirical application, we examine the performance of both filters to derive credit-to-GDP gaps for 44 countries. We conclude that differences between gaps obtained from both filters are negligible, when the Holt-Winters filter is initiated properly. Some differences are observed at the beginning of the sample and after countries experienced large swings in the credit-to-GDP ratio. But even then difference are in the order at most 1-2 percent.

We conclude that the Holt-Winters filter does not provide serious disadvantages compared to the one-sided Hodrick-Prescott filter if there are at least 10 to 20 years of data. Given its computational simplicity and mathematical elegance, it is therefore a valid alternative.

A Appendix A: Alternative Matrix

Consider the following matrix:

$$C = \rho \begin{bmatrix} \cos(\varphi) & -\sin(\varphi) \\ \sin(\varphi) & \cos(\varphi) \end{bmatrix}.$$

Noting that $\sin^2(\varphi) + \cos^2(\varphi) = 1$ (Spiegel and Liu, 1999) then eigenvalues of matrix C are the complex numbers: $\rho[\cos(\varphi) + i\sin(\varphi)]$ and $\rho[\cos(\varphi) - i\sin(\varphi)]$.

Further, the addition formulas of trigonometric functions sine and cosine are the following (Spiegel and Liu, 1999): $\sin(\varphi + \vartheta) = \sin(\varphi)\cos(\vartheta) + \cos(\varphi)\sin(\vartheta)$ and $\cos(\varphi + \vartheta) = \cos(\varphi)\cos(\vartheta) - \sin(\varphi)\sin(\vartheta)$. For example when $\vartheta = \varphi$ we have the double-angle formulas are $\sin(2\varphi) = 2\sin(\varphi)\cos(\varphi)$ and $\cos(2\varphi) = \cos^2(\varphi) - \sin^2(\varphi)$. Based on that we have the following property for any positive integer n:

$$C^n = \rho^n \begin{bmatrix} \cos(n\varphi) & -\sin(n\varphi) \\ \sin(n\varphi) & \cos(n\varphi) \end{bmatrix}.$$

B Appendix B: Additional graphs

Figure 6: Further comparisons of HP versus HW filter

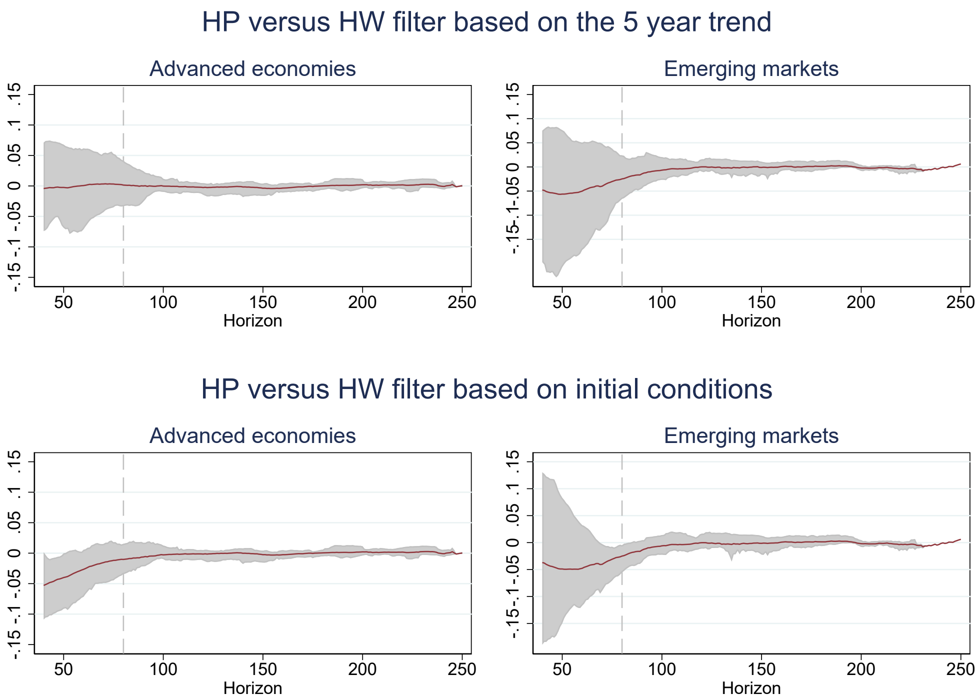

Solid line: Mean difference of the HW trend based on HP filter versus HW_{init} or HW_{5y}. All relative to the HP filtered trend. Shaded areas: 5th and 95th percent confidence bands. Horizon: quarters after the first observation of the credit-to-GDP ratio. Dotted line: quarter 80, indicating the horizon when one full cycle has been completed

References

Basel Committee on Banking Supervision (2010) *Guidance for National Authorities Operating the Countercyclical Capital Buffer.*

Borio, C. and P. Lowe (2002) "Assessing the Risk of Banking Crises" *BIS Quarterly Review.* 43-54.

Durbin, J., and S. J. Koopman (2008) *Time Series Analysis by State Space Methods* Oxford University Press.

Harvey, A. C. (1989) *Forecasting, structural time series models and the Kalman filter* Cambridge University Press.

Hodrick, R. J., and E. C. Prescott (1980) "Postwar U.S. Business Cycles: An Empirical Investigation," Carnegie Mellon University Discussion Paper no. 451.

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

Kim, H. (2004) "Hodrick-Prescott Filter" Mimeo, Auburn University.

Ravn, M. O. and H. (2002). "On Adjusting the Hodrick-Prescott Filter for the Frequency of Observations." *Review of Economics and Statistics* 84(2), 371-76

Spiegel, M., and J. Liu (1999) *Mathematical Handbook of Formulas and Tables,* Second Edition Schaum's Outline Series.

StataCorp (2017) *Reference Manual* version 15.

Previous volumes in this series

All volumes are available on our website www.bis.org.