

OPTIMAL HP FILTERING FOR SOUTH AFRICA

LEON DU TOIT*

Abstract

Among the various methods used to identify the business cycle from aggregate data, the Hodrick-Prescott filter has become an industry standard – it ‘identifies’ the business cycle by removing low frequency information, thereby smoothing the data. Since its inception in 1980 the value of the smoothing constant for quarterly data has been set at a ‘default’ of 1600, following the suggestion of Hodrick and Prescott (1980). This paper argues that this ‘default value’ is inappropriate due to its *ad hoc* nature and problematic underlying assumptions. Instead this paper uses the method of optimal filtering, developed by Pedersen (1998, 2001, and 2002), to determine the optimal value of the smoothing constant for South Africa. The optimal smoothing constant is that value which least distorts the frequency information of the time series. The result depends on both the censoring rule for the duration of the business cycles, and the structure of the economy. The paper raises a number of important issues in the practical use of the HP filter, and provides an easily replicable method in the form of MATLAB code.

Keywords: Hodrick-Prescott filter; spectral analysis; ideal filtering; optimal filtering; distortionary filtering; business cycles; MATLAB

JEL Codes: C22; E32

*Masters student, Department of Economics, University of Stellenbosch. This paper benefited from comments by Stan du Plessis, Kevin Kotzé, and help from Francois Schonken in writing the code.

1. INTRODUCTION

This paper uses the method of optimal filtering to determine a new set of values for the smoothing constant in the Hodrick-Prescott filter. The proposed values for quarterly data are 1338 or 352, while the choice between which one to use is driven primarily by the duration of business cycles. This method has the merit of being transparent and replicable, and since the HP filter is widely used in business cycle and macro-econometric modelling in general, contributes to a more careful study of business cycles.

Although there is no definitive way of identifying the business cycle, the various methods can be grouped into three broad approaches. Firstly, the classical cycles methodology identifies the relative expansion and contraction of aggregate economic activity and is mainly concerned with determining turning points, developed in extraordinary detail by Burns and Mitchell (1946). Secondly, the method of deviation cycles which includes the class of statistical filters such as the HP filter (Hodrick and Prescott, 1997), generate a new stochastic variable, which is the stationary cyclical component in the original trended series. Lastly, model-based approaches, for example, SVAR analysis (Du Plessis, Smit and Sturzenegger, 2007) or Markov-switching models (Moolman, 2004), are used to identify the business cycle based on theoretical priors.

The HP filter is widely used in descriptive research and in econometric modelling, and since its inception in 1980, most applications have been guided, implicitly or explicitly, by Hodrick and Prescott's original calibration. This paper argues that their method yields sub-optimal HP filtering for two reasons; firstly, their specific method for selecting the value of the smoothing parameter is *ad hoc* and not transparent or replicable; and secondly, the probability model underlying their *ad hoc* choice is problematic. Section 2 will develop these claims and guide the argument into its main thesis: that the theory of optimal filtering stated and implemented by Pedersen (2002) is the most appropriate alternative to the method used by Hodrick and Prescott (1997). The paper is organised to explain exactly how the theory of optimal filtering is used to determine the optimal value of the smoothing parameter of the HP filter for South Africa.

Section 3 contextualises the use of the HP filter and provides the reader with the necessary concepts to understand how the method is developed and applied. It shows the relation between time- and frequency-domain representations of time-series, and the general characteristics of filters. Section 3 discusses the different approaches to selecting the smoothing parameter and motivates the choice to use the theory of optimal filtering.

Section 4 shows that the optimal value of the smoothing constant is that value which minimises the distortionary effect of the HP filter. The optimal HP filter depends on both the specified duration of the business cycle, and on the cyclical information of the input series – the structure of the economy. The results suggest optimal values of lambda well below the ‘default’ of 1600, and motivate the use of a censoring rule which specifies business cycles as frequencies which occur at six years or less, following the results of Pedersen (1998), Rand and Tarp (2001), Du Plessis (2006) and the South African Reserve Bank (2007). This optimal value is 352. Section 5 summarises the main results on optimal HP filtering, and suggestions to its implementation, and section 6 applies the results to the South African business cycle.

2. BUSINESS CYCLES AND FILTERING IN FREQUENCY DOMAIN

2.1 THE HP FILTER IN THE BUSINESS CYCLE LITERATURE

Although the two approaches of classical and deviation cycles have different specific conceptualisations of the business cycle, they agree on a set of general characteristics (Du Plessis, 2006: 1 – 8). The business cycle is a pathology with four main aspects, all mentioned by Burns and Mitchell (1946: 3); it is a succession of relatively prosperous and depressed periods with two phases (peak-to-peak or trough-to-trough); it is recurrent but not periodic; it varies in duration from 1 to 12 years; and the cycle refers to total economic activity (Agenor, 2004).

The classical cycles methodology was originally developed Burns and Mitchell (1946), later to be automated in a very simplified form by Bry and Boschan (1971) and recently used by Pedersen (1998) and Harding and Pagan (2001). The so-called BBQ algorithm preserves the classical cycles methodology's focus on duration and turning points. The essence of this method is to identify turning points and separate periods of relative expansion and contraction in aggregate economic activity. The analysis depends crucially on the implementation of censoring rules (usually about the duration of cycles) which are applied to yield a consistent set of turning points.

In contrast to classic cycles, the currently dominant approach is the study of deviation cycles (Du Plessis, 2006: 4); this approach conceptualises the business cycle as "serially correlated deviations of output from its trend". By applying a statistical filter (for example, Band-Pass, Beveridge-Nelson or HP filter) to the data, a new stochastic variable is created, identified as the cyclical component of GDP. These methods may differ in their conceptualisation of the business cycle, but an important conclusion drawn from this paper is that both of their results depend on a censoring rule that captures the researcher's view of the duration of cycles.

Secondly, the HP filter should be seen in the light of recent applications, in business cycle studies, and econometric modelling. In recent South African academic literature it is used in well over 30 published journal articles, confirming the popularity of the method. Examples in international literature are Du Rand (2006), Rand and Tarp (2001), Razzak & Dennis (1999), and Agenor, McDermott and Prasad (1999), while some recent examples in South African literature are Liu and Gupta (2007), Burger and Marinkov (2006), Fedderke and Schaling (2005). But there is little discussion of the issues surrounding its application. This paper will contribute to existing research practices by providing a systematic analysis of optimal HP filtering for South Africa, following the method developed by Pedersen (1998, 2001, and 2002).

2.2 SPECTRAL ANALYSIS AND THE EFFECTS OF FILTERING

A simple example will clarify the relation between time domain and frequency domain representations of time series. We know that the business cycle is a phenomenon which occurs alongside the long term growth of GDP, and we wish to isolate the medium- and short-term fluctuations in the observed series of total economic activity, to identify it. But we cannot observe the business cycle directly, and must use an analytical framework to identify the cycle from the aggregate data.

Now suppose one is interested in listening to the sopranos in a choir, but can only observe the whole choir, with all of the differently pitched voices singing simultaneously. In time domain, one would observe the music as sang by the choir. In frequency domain, one would decompose the music (the unfiltered time series) into voices which correspond to different frequencies, and rank them on a scale from low to high. Combining the information about the different frequencies into a new measure will present all the frequency information of the choir, just as a frequency domain representation presents the cyclical information of time series.

Once we obtain these different frequencies, a filter can be applied to separate the higher from the lower, thereby enabling us to listen more clearly to the sopranos. This is exactly the process employed when applying a filter to a time series in frequency domain, and optimal filtering theory just aims to minimise the noise captured when filtering all of the frequency information. This section develops these concepts formally.

2.2.1 THE SPECTRUM OF TIME SERIES

The aim of this section is to derive the theoretical *spectrum* and the *spectral density function* for a stochastic process, since it represents the information about the cyclical pattern of the series. We are interested in thinking about time series in this way, since the deviations cycles method conceptualises the business cycle as serially correlated deviations of output from its trend.

A stochastic process $\{Y_t\}_{t=-\infty}^{\infty}$ with realizations $\{y_t\}_{t=0}^T$ and moments,

$$\mu_t = E y_t = \sum_{t=0}^T y_t, \quad (1)$$

$$\gamma_t = E y_t y_{t-j} = \sum_{t=0}^T (y_t - \mu)(y_{t-j} - \mu), \quad (2)$$

is covariance (weakly) stationary if μ_t and γ_t are independent *for all t and j*. The dynamic behaviour or cyclical pattern of an economic time series is summarized by the autocovariance - or autocorrelation function, which contains all the information about the time dependence of individual observations in the series (Pedersen, 2002: 14). A *generating function* is a way of recording the information of some sequence. If $\{\gamma_j\}_{j=-\infty}^{\infty}$ is the sequence of autocovariances of $\{Y_j\}_{j=-\infty}^{\infty}$, and $\sum_{j=-\infty}^{\infty} |\gamma_j| < \infty$, then the *autocovariance generating function* is,

$$g_y(z) = \sum_{j=-\infty}^{\infty} \gamma_j z^j, \quad (3)$$

where the argument z for the autocovariance generating function is any value that lies on the complex unit circle. Now (3) is the *Fourier transform* of the *covariances* (the cyclical behaviour) of the stochastic process. From (3) we get the spectrum (5) of y_t using De Moivre's Theorem (4),

$$z = \cos(\omega) - i \sin(\omega) = e^{-i\omega}, \quad (4)$$

$$s_y'(\omega) = \frac{1}{2\pi} \sum_{j=-\infty}^{\infty} \gamma_j e^{-i\omega j}, \quad (5)$$

where $i = \sqrt{-1}$, ω is the frequency (or radian angle of z with the real axis), and where dividing (5) by 2π normalizes it to sum to 1. So (5) explicitly represents the cyclical information of the process y_t . Importantly, the *spectrum* contains all of the information about the cyclical behaviour of the time series. Spectral representation makes it easier to

grasp how the series is composed of different fluctuations. The spectrum has some important properties. Since $\gamma_j = \gamma_{-j}$,

$$s_y'(\omega) = \frac{1}{2\pi} (\gamma_0 + \sum_{j=1}^{\infty} \gamma_j [e^{i\omega} + e^{-i\omega}]) \quad (6)$$

$$= \frac{1}{2\pi} (\gamma_0 + 2 \sum_{j=1}^{\infty} \gamma_j \cos(\omega j)), \quad (7)$$


using Euler's formula ($\frac{e^{iz} + e^{-iz}}{2} = \cos z$), the spectrum is always real valued, is a sum of infinite cosines, and all the relevant information is concentrated in the interval $[0, \pi]$.

Sometimes, however, the *autocovariance generating function* is replaced by the *autocorrelation generating function*. This is equivalent to dividing the spectrum (5) by γ_0 so that we get the *spectral density function* of the time series:

$$s_y(\omega) = \frac{1}{2\pi} (1 + 2 \sum_{j=1}^{\infty} \rho_j \cos(\omega j)), \quad (8)$$

where $\rho_j = \frac{\gamma_j}{\gamma_0}$. The *spectral density* is the function that will be used to capture the information about the cyclical behaviour of real GDP when developing an optimal HP filter for South Africa. Spectral representation is, therefore, nothing more than a useful way of representing time series data that makes the cyclical composition of the series explicit (the different voices in the choir). Figure 1 below shows the *spectrum* of the log of real GDP for South Africa (quarterly, seasonally adjusted, 1960q1 to 2007q2), estimated using the Yule-Walker autoregressive technique (Berk, 1974).

Figure 1: Spectrum of log real GDP



The *spectral density* exhibits the typical shape for macroeconomic data noted by Granger (1966) with most of the information being bunched in the lower frequencies, referring to the long term fluctuations in the series. As is seen on the x-axis, the spectral density lies within $[0, \pi]$. The next section develops the ideas of ideal and optimal filtering, and presents the HP filter, its characteristics and distortionary effects in frequency domain.

2.2.2 BENCHMARK FILTERING: LINEAR TIME INVARIANT FILTERS

Given a process x_t , a linear time invariant filter (that is, covariance stationary) is an operator from the space of sequences into itself that generates a new process y_t of the form

$$y_t = H(L)x_t, \tag{9}$$

where,

$$H(L) = \sum_{j=-\infty}^{\infty} h_j L^j, \tag{10}$$

where h_j are *Fourier coefficients* and L^j is the lag operator, and the operation of the filter on the input series is summarised by the *frequency response function*, $H(L)$ (Fernandez-Villaverde, 2007: 19 – 22).

Then the spectrum of the filtered series y_t is given by,

$$s_y'(\omega) = |H(e^{-i\omega})|^2 s_x'(\omega), \quad (11)$$

where $H(e^{-i\omega})$ is the *Fourier transform* of the coefficients of the lag operator or the frequency response function. To gauge the precise nature of the effects of filtering we look at the polar decomposition of the frequency response function,

$$H(e^{-i\omega}) = G(\omega)e^{-iPh(\omega)}, \quad (12)$$

where $|H(e^{-i\omega})|$ is the *gain* (the relative importance of various cyclical components), and $Ph(\omega)$ is the *phase* (the displacement of frequencies). The operation of a linear filter on the input process, is completely described by the *gain* and *phase* of a filter. But when we look at how the spectrum is changed in (11) we see that the operation of the filter is captured by the *power transfer function* defined as,

$$H(\omega) \equiv |H(e^{-i\omega})|^2. \quad (13)$$

The power transfer function indicates how much the spectrum of the series is changed *at each particular frequency* when filtered. This summarises the effect of the filter on the input series. Now, by adding *ideal* and *optimal filters*, and distortionary effects of filters to the analysis in the next section, the *spectral density* and the *power transfer function* will completely describe the method determining the optimal λ .

2.2.3 IDEAL AND OPTIMAL FILTERS

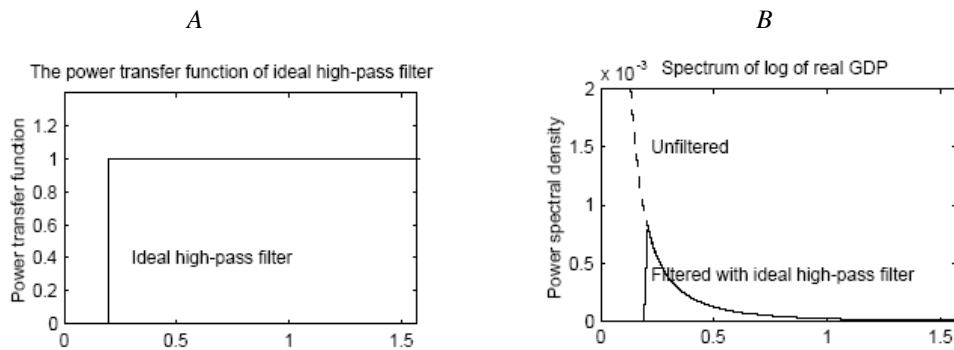
An *ideal filter*, such as a linear time invariant filter described in the previous section, is an operator $H(L)$ such that the new process Y only has a positive spectrum in some specified part of the domain.

The power transfer function of an *ideal high-pass filter* (used because the HP filter is a high-pass filter) is such that,

$$H^*(\omega) \equiv \begin{cases} 0 & \text{if } |\omega| < \omega_{cut-off}, \\ 1 & \text{if } |\omega| \geq \omega_{cut-off}, \end{cases} \quad (14)$$

where $\omega_{cut-off}$ is specified to match the maximum average duration of the business cycle. Panel 1 (Pedersen, 2002: 83) shows the power transfer function of an ideal filter (A) and the filtered spectrum after such an ideal filter has been applied to the data (B).

PANEL 1: Ideal filters and the power transfer function

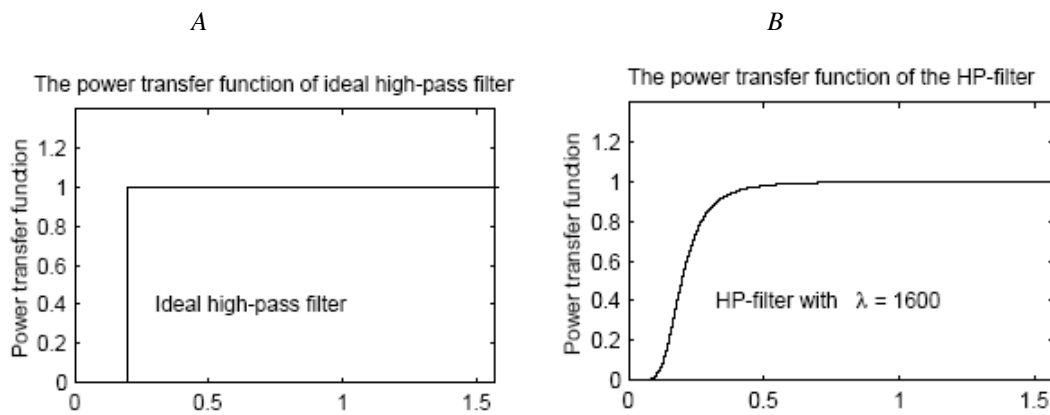


The distinction between *ideal* and *optimal filters* is that an *optimal filter* is an approximation of an *ideal filter*. This is necessary because with finite data it is not possible to build an *ideal filter*. The HP filter is, therefore, an example of such a finite sample approximation. The power transfer function of the HP filter (Pedersen, 2002) is

$$H_{HP}(\omega) = \left| \frac{4\lambda(1-\cos(\omega))^2}{1+4\lambda(1-\cos(\omega))^2} \right|^2. \quad (15)$$

Panel 2 (Pedersen, 2002: 85) shows the different shapes of the power transfer function of an ideal high-pass filter, $H^*(\omega)$, and the HP filter, $H_{HP}(\omega)$.

PANEL 2: Ideal and optimal power transfer functions



As the smoothing constant increases, so the power transfer function moves left, cutting off less of the low frequency information. Correspondingly, as lambda increases, the less is the difference between the trend and the cycle – by including more low frequency information as lambda increases, so more of the trend information is included in the cycle itself. Apart from the role of lambda, the power transfer function of the HP filter highlights important features of the HP filter itself (Pedersen, 2001: 1088). The cyclical filter has zero gain at zero frequency, and near unit gain at $\omega = \pi$, since $\cos(\omega) = -1$. It is a symmetric filter which induces no phase shift, and lastly, there is no cycle in the power transfer function itself. This makes the HP filter a close approximation to the *ideal high-pass filter*. Baxter and King (1999) highlight important finite sample consequences related to the power transfer function, which will be discussed in section 5.

By analysing the HP filter in frequency domain we have derived the building blocks for measuring the distortionary effect of the HP filter, whereby the optimal value of the smoothing parameter is determined. These building blocks are the *spectral density* of the log of GDP, and the *power transfer function* of the *ideal* and the *optimal* HP filter.

3. THE HP FILTER

3.1 HP FILTER DESIGN AND ISSUES

Hodrick and Prescott (1997: 2) propose the use of an “easily replicable technique that incorporates our prior knowledge about the economy” as an input to understanding the “features of the economy that an equilibrium theory should incorporate.” The maintained hypothesis for their study is “that the growth component of aggregate economic time series varies smoothly over time”. The study also explicitly states that “no one approach dominates all the others and that it is best to examine the data from a number of different perspectives”. So, given the maintained hypothesis, the conceptual framework is that a given, seasonally adjusted, time series y_t is the sum of a growth component g_t and a cyclical component c_t ,

$$y_t = g_t + c_t \quad \text{for } t = 1, \dots, T. \quad (16)$$


Since c_t are deviations from g_t the problem is summarised in the following formula for determining growth components;

$$\min_{\{g_t\}_{t=1}^T} \{ \sum_{t=1}^T (y_t - g_t)^2 + \lambda \sum_{t=1}^T [(g_t - g_{t-1}) - (g_{t-1} - g_{t-2})]^2 \}. \quad (17)$$

The unconstrained minimisation problem presented in (17), given the structure of the time series in (16) yields a stationary series that is identified as the cyclical component of GDP, where $(y_t - g_t)^2 = c_t^2$. Here λ is the smoothing parameter which penalises variability in the growth component of the series. If $\lambda = 0$ then $g_t = y_t$, and the solution is trivial – there is no cycle, only trend growth. And if $\lambda = \infty$ then there is a linear trend. The latter possibility is ruled out on the basis that the growth rate of output is not constant over the sample period, while the first case yields no business cycle, making it useless. Practical studies implement values between these extreme cases. The typical output is presented in figure 2 below, using the value of 1600 proposed by Hodrick and Prescott (1997: 4) for quarterly data, where the cyclical component is the stationary series of the

the log of seasonally adjusted real GDP for quarterly South African data from 1960q1 to 2007q2.

Figure 2: Log real GDP and the cyclical component



The value of the smoothing parameter plays a pivotal role in the results generated by the HP filter. It will be argued below that choosing the appropriate value of λ is equivalent to using the HP filter optimally. Before explaining this idea the next section considers Hodrick and Prescott's approach to find the smoothing parameter.

3.2 THE SMOOTHING PARAMETER

Hodrick and Prescott (1997: 4) proceed as follows; "the following probability model is useful for bringing to bear prior knowledge in the selection of the smoothing parameter λ . If the cyclical component and the second differences of the growth components in (17) were independently and identically distributed, normal variables with means zero and variances σ_g^2 and σ_c^2 (which they are not), the conditional expectation of the g_t , given the observations, would be the solution to the formula presented in (17) when $\lambda = \sigma_g^2/\sigma_c^2$."

Hodrick and Prescott (1997: 4) mention that this model has a state space representation and can be solved with efficient Kalman filtering, but do not present this formally. Pedersen (2002: 85 - 86) summarises the analysis of Harvey and Jaeger (1993) who show that the procedure used by Hodrick and Prescott yields an optimal filter for a very specific class of unobserved component models. King and Rebelo (1993) argue that the relevant unobserved component model is an uninteresting class of economic models, raising critique against the procedure and the HP filtering method itself. The Hodrick and Prescott probability model is stated explicitly as follows (Pedersen, 2002: 85 – 86):

$$y_t = g_t + c_t, \quad c_t \sim NID(0, \sigma_c^2), \quad (18)$$

where c_t is the business cycle component and g_t is given by

$$g_t = g_{t-1} + \beta_{t-1}, \quad (19)$$

$$\beta_t = \beta_{t-1} + \varepsilon_t^g, \quad \varepsilon_t^g \sim NID(0, \sigma_{\varepsilon^g}^2). \quad (20)$$

Here (19) and (20) are the state equations, while (18) is the measurement equation. The innovations to trend growth are assumed to be orthogonal to the business cycle components, ($E[c_t \varepsilon_t^g] = 0$). As Pedersen notes (2002: 85) this is a crucial assumption, “since fluctuations in both the business cycle and growth components are the effect of stochastic growth of technology in real business cycle theory”. The orthogonality assumption is likely to be violated, making the underlying probability model problematic.

Furthermore, the stochastic growth component g_t has a random walk growth rate, which implies that the second difference of the trend is white noise,

$$[(g_t - g_{t-1}) - (g_{t-1} - g_{t-2})] = \Delta^2 g_t = \Delta \beta_{t-1} = \varepsilon_t^g. \quad (21)$$

So in this probability model, if one assumes that σ_g^2 and σ_c^2 or equivalently ε_t^g and σ_c^2 are *i.i.d.* it is equivalent to assuming that the business cycle is white noise. And this runs

contrary to the earlier conception of the business cycle. If it were the case then the formula in (17) could be solved with the assumption that $\lambda = \sigma_g^2/\sigma_c^2$.

The procedure described above is the method employed by Hodrick and Prescott (1997: 4). “Our prior view is that a 5 percent cyclical component is moderately large, as is a one eighth of 1 percent change in the growth rate in a quarter. This led us to select $\lambda = 1600$.” The original method of selecting the smoothing constant is, however, problematic in a number of ways. As was argued, the assumptions necessary for the underlying probability model to hold are likely to be violated, leaving the method for determining the smoothing constant suboptimal. Given this, the *ad hoc* procedure used by Hodrick and Prescott (1997: 4) to arrive at the value of 1600 does not use the information of the structure of the economy or the process of filtering transparently.

This is important because of the critical role played by smoothing constant in generating the serially correlated deviations of output from its trend, which is then identified as the business cycle. The subsequent output is also used when studying the co-movements of various series with the business cycle. The next section will investigate alternative ways of selecting the smoothing parameter in a data-based and transparent way, leading the paper into a detailed discussion of the method of optimal filtering.

4. APPROACHES TO SELECTING THE SMOOTHING PARAMETER

Two lines of thought are reflected in the literature which analyses the HP filter. One strand uses a time domain approach, while the other analyses the HP filter in frequency domain. Schlicht (2004), for example, follows a time-domain approach developing both a maximum-likelihood and method of moments estimator for the ratio of the variances so that $\lambda = \sigma_g^2/\sigma_c^2$. This approach, however, has the same drawback as the method proposed by Hodrick and Prescott (1997). Marcet and Ravn (2003) develop a method for choosing lambda in cross-country comparisons, but do not investigate the initial choice itself. Treating $\lambda = \sigma_g^2/\sigma_c^2$ relies on the same questionable probability model.

The HP filter was first analysed in frequency domain by Singleton (1988), and due to subsequent research (Ravn and Uhlig (2002), Harvey and Jaeger (1993), King and Rebelo (1993), Cogley and Nason (1995), and Guay and St-Amant (1997)) ideas surrounding optimal filtering have been developed within frequency domain analysis of time series, discussed in the seminal work by Koopmans (1974), Warner (1998) and Fernandez-Villaverde (2007). Various ideas formalised in the spectral analysis of time series and consequently the HP filter, are captured in the method for choosing the optimal value of λ developed by Pedersen (1998, 2001, 2002), and used recently by Rand and Tarp (2001) in their study of developing country business cycles. This method determines the optimal value of λ as that which minimises a measure of the distortionary effect of the HP filter.

Before discussing the metric of the distortionary effect of the HP filter (which is minimised to find the optimal λ) a few concepts will be clarified in the next section; general spectral representation and the analysis of time series; linear time invariant filters, ideal and optimal filters in frequency domain; and the properties of the HP filter and its distortionary effects. These concepts are the building blocks of the method used to determine the optimal value of the smoothing parameter.

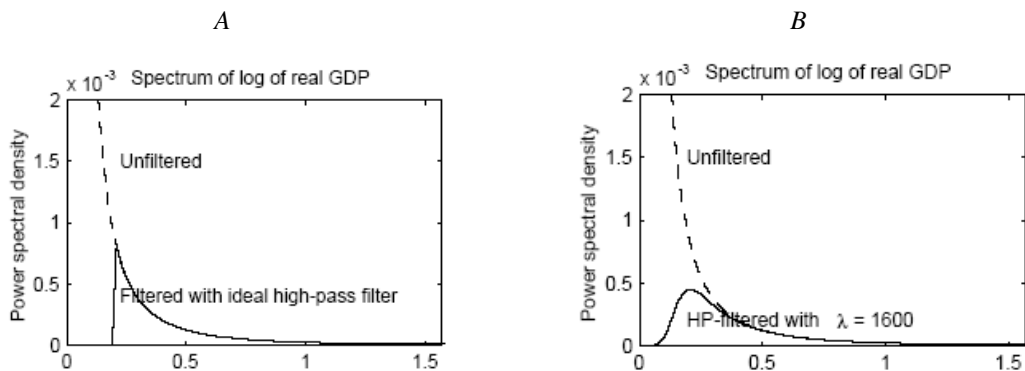
5. OPTIMAL HP FILTERING

5.1 THE DISTORTIONARY EFFECTS OF THE HP FILTER

Firstly, the application of the HP filter alters the computed second moments of the input series (Pedersen, 2002: 81). Among these are the standard deviation, the autocorrelation of individual series, and the correlations between different series and various lags and leads. Despite the claim by King and Rebelo (1993) and Guay and St-Amant (1998) that this is evidence of a failure of the HP filter, Pedersen (2002: 82) argues that this is a general characteristic of filters, since even an *ideal filter* changes the computed second moments.

Secondly, as noted by Granger (1996: 152 – 154), and restated in Harvey and Jaeger (1993), Cogley and Nason (1995) and Guay and St-Amant (1997), when filtering economic time series with the “typical spectral shape” (figure 1), the output series’ spectral shape has a rounded hump at business cycle frequencies which is not present in the original series. These authors interpret this as a spurious cycle. Panel 3 (Pedersen, 2002: 83) shows the different shapes in the filtered spectra when using an ideal (A) or a distortionary filter (B).

PANEL 3: Effects of different filters

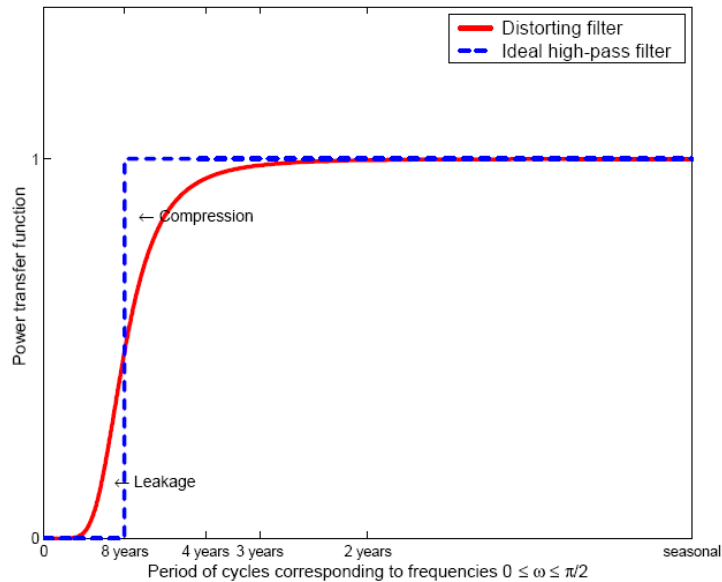


Pedersen (2002: 83), however, argues that the rounded shape is the result of leakage and compression – the difference between the *power transfer function* of an ideal filter (14) and that of a distortionary filter, such as the HP filter (15). The conclusion is that the humped shape is not evidence of a ‘spurious cycle’, since the *power transfer function* of the HP filter does not have a cyclical component, seen in (15).

Pedersen (2002: 61) states that the “goal of the theory of optimal filter design” is to “construct filters which minimise distortionary effects.” This is necessary because when computing business cycle stylised facts, and when using HP filtered data in a model, the representation of the business cycle cannot be said to be exactly “true”. But the application of optimal filter design can improve the confidence in these business cycle computations. “In general, a filter is distorting by passing frequencies which it was supposed to attenuate (*leakage*), and by compressing frequencies which should pass the filter (*compression*)” (Pedersen 2002: 61). This is illustrated in figure 3, and is the

consequence of the difference between the *power transfer function* of the *ideal filter* and the HP filter.

Figure 3: Leakage and Compression



Importantly, the combination of leakage and compression is not, in itself, a measure of the distortionary effect of a filter. It depends on the *power transfer function* of the filter and the *spectrum* of the input series. That is, the leakage and compression, at each frequency ω is weighted by the relative power of the input series at the same frequency (Pedersen, 2002: 61). And because the power transfer function of the HP filter is dependent of the value of the smoothing parameter, as seen in (15), one can find the value of λ which minimises the distortionary effect of the HP filter – this is the same as using the HP filter optimally. The method is transparent, and uses the data to compute the optimal value of λ , while circumventing the problems associated with the original *ad hoc* method and similar methods relying on the same problematic probability model, as discussed in section two. The next section shows the derivation of the optimal HP filter.

5.2 MEASURING DISTORTIONARY EFFECTS

This section discusses the method of optimal filtering developed by Pedersen (2001: 1091 – 1092) and uses the concepts developed in the previous sections of this paper to derive a

metric to determine the optimal value of the smoothing parameter. Let the ideal high-pass filter have a power transfer function $H^*(\omega)$, where $\omega \in W = (\omega_1 < \omega_2 < \dots < \omega_n)$ with $\omega_1 = 0$ and $\omega_n = \pi$, (25). Also, let the power transfer function of the distortionary filter (in this case the HP filter) be $H_{HP}(\omega)$, as stated in (21). The spectral representation of the true cyclical component for $\omega \in W$ is given in (22) and the cyclical component of the distortionary cyclical component is given by (23),

$$H^*(\omega)2s_y'(\omega), \quad (22)$$

$$H_{HP}(\omega)2s_y'(\omega). \quad (23)$$

The distortionary effect is based on the leakage and compression, for the given input series, that arise due to the difference between the transfer function of the ideal and distortionary filters (figure 7). The distortionary effect of the filter (Q) is the sum of the absolute value of the difference between the cyclical component of the ideal and the distortionary filter, weighted by spectrum of the input process and the size of the grid $\Delta\omega = \omega_i - \omega_{i-1}$.

$$Q = \sum_{\omega \in W} |H^*(\omega) - H_{HP}(\omega)| \cdot 2s_y'(\omega)\Delta\omega, \quad (24)$$

A new set of weights $v(\omega)$ is constructed so that they are normalised to sum to 1, where the weights are defined as the ratio of the spectrum of the input process to the variance of the series. Applying the new weights to (24) and minimising with respect to λ , in $H_{HP}(\omega)$ yields the metric for determining the optimal value of the smoothing constant in the HP filter,

$$\text{argmin}_{\lambda} Q = \sum_{\omega \in W} |H^*(\omega) - H_{HP}(\omega)| \cdot v(\omega), \quad (25)$$

where,

$$v(\omega) = \frac{2s_y'(\omega)\Delta\omega}{\sum_{\omega \in W} 2s_y'(\omega)\Delta\omega}. \quad (26)$$

The formula in (25) is fully described by the power transfer functions of the ideal and the distortionary filters, and the spectrum of the input series. Firstly, the ideal filter needs to be specified with a cut-off frequency corresponding to the maximum length of the business cycle. This is analogous to the censoring rules applied in the identification of classical cycles (Burns & Mitchell, 1946), where the researcher must make an explicit choice based on the information provided by the data, and possibly a theoretical prior. Secondly, (25) depends on the value of the smoothing constant via the power transfer function of the HP filter. Minimising Q with respect to λ means minimising leakage and compression, yielding a more reliable business cycle component. It is in this sense that using the HP filter as explained here is optimal.

Lastly, Q is weighted by the spectrum of the input process (log GDP) which contains all the relevant information about the cyclical behaviour of the series. In this way, the data adds directly to the determination of the optimal smoothing parameter. These points along with the program itself are also systematically described in the appendices (PlotQ.m and HP.m), explaining practical implementation with MATLAB files.

5.3 RESULTS

This section presents the results of taking the method to the data. It implements the programs presented in the appendices. The data is the log of seasonally adjusted quarterly real GDP for South Africa for the period 1960q1 to 2007q2. When working with real, finite data, the spectrum of the series must be estimated – we use *autoregressive spectral estimation* (Parzen, 1961 and Shibata, 1981). Berk (1974) shows that this is an unbiased, consistent and asymptotically normal estimator. Estimation proceeds by approximating the true data generating process by an $AR(p)$ process, and using the coefficients to estimate the power spectral density (8).

The second step in generating the results involves choosing an appropriate cut-off frequency for the ideal and the distortionary filter. Pedersen (1998) uses the method developed by Bry and Boschan (1971) to determine the average length of business cycles for a group of 11 OECD countries, and finds that for most countries, it is shorter than six years, *on average*. Rand and Tarp (2001) use the results of Pedersen (1998) to study

business cycles for the developing world, and the results for Emerging Market Economies generated by Du Plessis (2006), also using the BBQ algorithm, suggest that the duration of business cycles is less than six years. This is also supported by the turning points published by the South African Reserve Bank (SARB, 2007: S-135), using the method described by Venter (2005). According to the SARB's method the average duration is 4.79 years.

The notable exception is, however, the current South African expansion, dating from August 1999 to the present. This exception, along with the finite sample issues of the HP filter (Baxter & King, 1999) warrants an investigation into the current expansion. This will be done in the next section. Table 1 below reports the optimal values of the smoothing parameter for South Africa's quarterly data, for various frequencies, based in minimising Q .

Table 1: Optimal Values of Lambda

Frequency	Duration	Lambda
$\pi/10$	5 years	177
$\pi/12$	6 years	352
$\pi/14$	7 years	524
$\pi/16$	8 years	1338

For a business cycle duration of six years the optimal value of lambda differs markedly from the 1600 proposed by Hodrick and Prescott (1997). It seems, however that if one specifies the business cycle as being 8 years and less, that the optimal lambda differs only slightly from 1600. Both these values are optimal but correspond to a different censoring rule. It is worth while investigating how these values compare in practice – this will be done in section 6. These results are generated by applying the theory of optimal filters, providing consistent way of evaluating different options.

The panel of figures below show the Q -statistics for durations of six and eight years.

Figure 4: Q for 6 Years





Figure 5: Q for 8 Years

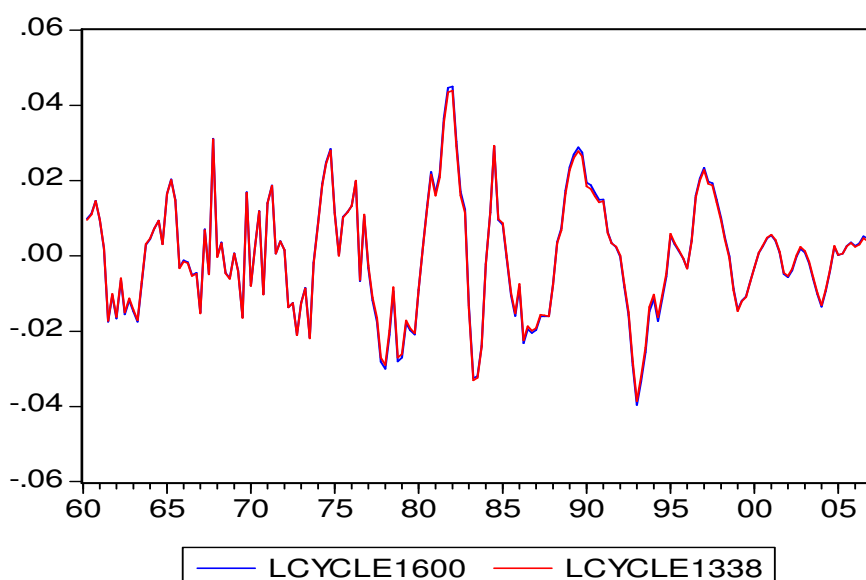


[more interpretation] The results confirm the influence of the censoring rule in the analysis of business cycles, and stresses that it plays just as big a role in deviation cycle analysis as in classic cycle analysis. It also shows that the results for a duration of six years suggest a value for lambda that is quantitatively different than the ‘default’ of 1600. For a duration of eight years, however, the difference is smaller. Now that the values for optimal HP filtering for South Africa have been determined, the next section applies these values to gauge how the results generated by applying optimal HP filtering compare to those generated by other methodologies.

6. APPLICATION TO THE SOUTH AFRICAN BUSINESS CYCLE

To gauge how the set of optimal values of the smoothing constant, derived in the previous section, translate into optimal filtering in practice, a few time domain applications are presented. It was shown above that the different optimal values of lambda arise due to different censoring rules regarding duration of business cycles and for durations less than 8 years the optimal values of lambda are well below 1600. Figure 6 shows the cyclical series generated by the optimal - and the default HP filter when the duration is 8 years, compared to 1600.

Figure 6: Cyclical series for different lambdas



The two series are almost indistinguishable, with standard deviations of 0.015471 and 0.015099 respectively. This is not to say, however, that the value 1338 is not optimal. It was derived based on the theory of optimal filtering, whilst the 'default' value of 1600 was specified in an *ad hoc* manner. Figure 7 shows that there is, however, an important difference in the amplitude of the cyclical series generated by an optimal lambda for business cycle durations of 6 years and less, when compared to cyclical series generated by the 8 year optimal value. The difference is important regarding the current expansion.

Figure 7: Cyclical series for optimal lambdas

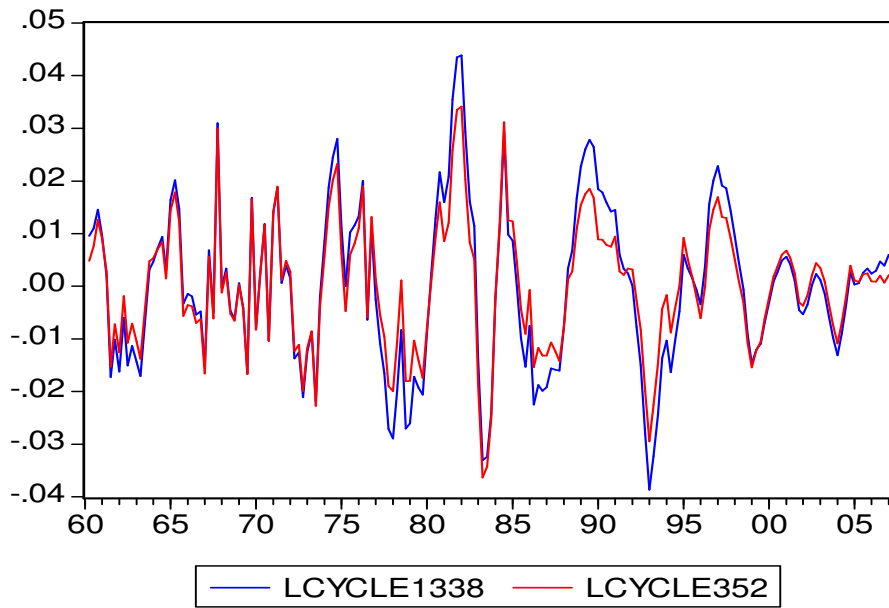


Table 2 shows a comparison of the business cycle turning points determined by various methods. Comparing to accuracy of various optimal HP filters, in relation to the SARB and BBQ turning points (SARB, 2007: S-135 and Du Plessis, 2006: 29), all fare more or less equally well.

Table 2: South African Business Cycle Turning Points

SARB	HP (352)		HP (1338)		HP (1600)		BBQ		
	Trough	Peak	Trough	Peak	Trough	Peak	Trough	Peak	
	1960q2		1960q4*		1960q4*		1960q4*		
1961q3	1965q2	1961q3*	1965q2	1961q3*	1965q2	1961q3*	1965q2		
1966q1	1967q1	1967q1	1967q4	1967q1	1967q4	1967q1	1967q4		1971q1*
1968q1	1971q1	1969q3	1971q2	1969q3	1971q2	1969q3	1971q2	1972q1	1974q3*
1972q3	1974q3	1973q3	1974q4	1973q3	1974q4	1973q3	1974q4	1975q1	1976q3
1978q1	1981q3	1977q4	1981q4	1978q1**	1981q4	1978q1**	1981q4	1977q3	1981q4
1983q2	1984q3	1983q2*	1984q3*	1983q2*	1984q3*	1983q2*	1984q3*	1983q1	1984q2
1988q2	1989q1	1987q4	1989q3	1986q2"	1989q3	1986q2"	1989q3	1988q1	1989q3
1993q2	1996q4	1993q1	1997q4	1993q1	1997q1"	1993q1	1997q1"	1992q4	1996q4*
1999q3		1991q1		1991q1		1991q1		1998q3	

* = coincidence with turning point identified by SARB

" = different than 352

The cyclical series generated by the censoring rule for 6 years is different from the 8 year and the default cyclical series, in amplitude and in the autocorrelations at various lags shown in figure 8 below.

Figure 7: Autocorrelations of cyclical series




Table 3 compares the standard deviation of the trend components of log GDP for the various optimal HP filters to the standard deviation of the original unfiltered log GDP series. The value of lambda = 352 yields the standard deviation which is closest to the original, which is understandable since an increased value of lambda implies a smoother output series. The value of 352 most closely matches the long-run features of the data.

Table 3: Standard Deviation: input - and filtered series

log(gdp)	$\lambda = 352$	$\lambda = 1338$	$\lambda = 1600$
0.372464	0.372098	0.371964	0.371945


Regardless of the value of lambda, the HP filter has another important shortcoming, arising from small sample applications. Recall that the operation of a filter is described by its power transfer function, or the square of its *gain* (how much the spectrum of the input series is changed at each frequency when filtered), and that *phase* (the displacement

of the series with respect to time) is a sub-component of gain itself. Fernandez-Villaverde (2007) notes that the corner sections of the sample are difficult to deal with due to the limited one-sided application of the lag-operator in the filter.

Related to this general filtering problem, Baxter and King (1999: 589) note that the gain function of the HP filter differs markedly from the gain function of the ideal filter at the end-points of the sample. This is largest for the first and last 12 observations – that is, for three years of observations, at the ends of the sample the estimate of the cyclical component of GDP will be less reliable than the middle of the sample. The bias will be in the opposite direction of the current cyclical movement because the information will be bunched to the trend. Since the current expansion is of great importance to econometric modelling, it is well worth investigating the problem.

Panel 4 shows that from 1999 to 2007 the data shows trend growth (A) to be the dominant effect – it is increasing at an increasing rate. Due to the end-point sample problem, the cyclical component will be biased downward, as is clear for both values of lambda (B). This may account for some of the lack of upward movement in the cyclical series during the current expansion, but it is also possible that the GDP data is simply too smooth (or has been smoothed) to identify the cyclical fluctuations during this period – too much information is bunched in the trend¹. This issue warrants separate investigation.

Panel 4: Trend growth and the cyclical component



1. This was pointed out by Stan du Plessis and Kevin Kotzé

Interestingly, the value of 1338 seems to be better at capturing the expected upward cyclical movement in the current expansion than the value of 352. But this is expected, because, as stated before, the larger the value of lambda, the more of the low frequency information will be included in the cyclical series.

Concerning the duration of cycles – the censoring rule – the SARB turning points (SARB, 2007) suggest an average duration of 4.79 years since 1960, if the current expansion is included. The average suggested by the BBQ algorithm (Du Plessis, 2006) adjusting for the current expansion 4.54 years. Considering this evidence, and the turning points identified by the HP filter in table 2, this paper suggests that optimal HP filtering for South Africa is best implemented with a value of 352 for lambda. This is consistent with Pedersen (1998, 2002) and Rand and Tarp (2001) who apply optimal HP filtering to developing countries with similar business cycle durations. So the choice between the two values is up to the researcher. But these issues might be driven by peculiarities in the data.

Applying the optimal HP filter, therefore, requires that we think carefully about the process of filtering itself. This involves thinking about the role of filtering in the conceptual field of business cycles, and about the empirical properties of the specific economy in question. The possibility exists that results are driven to a large extent by finite sample problems, and data quirks, but regardless of those possibilities, the application of the theory of optimal filtering minimises the distortion of the frequency space of the particular time series.

Now that a transparent method has been used to determine the theoretically optimal values of the smoothing constant for the HP filter for South Africa, further research can focus on the finite sample - and data issues.

CONCLUSION

The paper started by considering the HP filter's place in business cycle research and econometric modelling in general, establishing that it is an important method for identifying the business cycle or de-trending data. The investigation into optimal HP filtering for South Africa is motivated by the inadequacy of the default value of the smoothing parameter, and the method used to choose it, as presented in Hodrick and Prescott (1997). There is, therefore, a need to revise the value of the smoothing parameter, and shed some light on practical applications and consequences of the filter.

This paper implemented the method developed by Pedersen (2002) to determine the optimal value of lambda for South Africa. This is by no means a strictly deterministic process, since the value is not only conditional on the data, but also on a censoring rule which specifies the duration of business cycles. This finding is interesting since it is consistent with the claim by Burns and Mitchell (1946) that the crucial step in business cycle identification is duration. This paper, therefore, establishes an important methodological link between the deviation cycle and classical cycles method. The author argued that the most appropriate censoring rule specifies business cycle frequencies as those which occur at frequencies less than 6 years. This choice was based on evidence presented on the duration of business cycles by Pedersen (1998), Rand and Tarp (2001), Du Plessis (2006) and SARB (2007). The proposed optimal value of lambda is 352.

Apart from issues related to the smoothing parameter and the input data, the study of the HP filter reveals that estimated of the cyclical component of GDP is less reliable at the end-points of the sample. Since the current business cycle expansion is the crucial part of the data that the filter needs to take to any model, special caution, and transparency about assumptions is necessary regarding this sample period. The overall message of the paper is that optimal HP filtering requires thinking explicitly about the structure of the economy, and the conceptualisation of the business cycle.

REFERENCES

Agenor, P. R. 2004. *The Economics of Adjustment and Growth, 2nd Edition*. Massachusetts: Harvard University Press.

Agenor, P. McDermott, J. C. & Prasad, E. S. 1999. 'Macroeconomic Fluctuations in Developing Countries: Some Stylized Facts', *International Monetary Fund Working Paper 99/35*.

Baxter, M. & King, R. G. 1999. 'Measuring business cycles: approximate band-pass filters for economic time series', *The Review of Economics and Statistics*. 81(4). 575 – 593.

Berk, K. N. 1974. 'Consistent Autoregressive Spectral Estimates', *The Annals of Statistics*. 2(3). 489 – 502.

Bry, G. & Boschan, C. 1971. *Cyclical analysis of time series: selected procedures and computer programmes*. New York: National Bureau of Economic Research.

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

Burger, P. & Marinkov, M. 2006. 'The South African Phillips Curve: How Applicable is the Gordon Model?', *South African Journal of Economics*. 74(2). 172 – 189.

Cogley, T. & Nason, J. M. 1993. 'Effects of the Hodrick-Prescott filter on trend and difference stationary time series: Implications for business cycle research', *Journal of Economic Dynamics and Control*. 19. 253 – 278.

Du Plessis, S. A. 2006. 'Business cycles in emerging market economies: a new view of the stylised facts', *Stellenbosch Working Paper 07/07*, Department of Economics.

Du Plessis, S. A. Smit, B. & Struzzenegger, F. 2007. 'Identifying aggregate supply and demand shocks in South Africa', *Stellenbosch Working Paper 11/07*, Department of Economics.

Du Rand, G. 2006. 'The Sensitivity of fit of DSGE models to different methods of detrending', *University of Columbia*.

Fedderke, J. W. & Schaling, E. 2005. 'Modelling inflation in South Africa: A Multivariate cointegration analysis', *South African Journal of Economics*. 73(1). 79 – 92.

Fernandez-Villaverde, J. 2007. *Spectral Analysis*. Lecture Notes, Duke University. Available at <www.econ.upenn.edu>.

Guay, A. & St-Amant, P. 1997. 'Do the Hodrick-Prescott and Baxter-King Filters Provide a Good Approximation of Business Cycles?' *Center for Research on Economic Fluctuations and Employment Working Paper No. 53*.

Granger, C. W. J. 1966. 'The typical spectral shape of an economic variable', *Econometrica*. 34(1). 150 – 161.

Griffiths, J. D. 2005. *Quantum Mechanics*. New Jersey: Pearson Prentice Hall.

Harding, D. & Pagan, A. 2001. 'Extracting, analysing and using cyclical information', *Melbourne, Melbourne Institute of Applied Economics and Social Research*.

Hodrick, R. J. & Prescott, E. C. 1997 (1980). 'Postwar U.S. Business Cycles: An Empirical Investigation', *Discussion Paper no. 451*, Carnegie Mellon University. Reprinted in an updated version as 'Postwar U.S. Business Cycles: An Empirical Investigation', *Journal of Money, Credit and Banking*, 1997, 29(1). 1 – 16.

Jaeger, A. & Harvey, A. C. 1993. 'Detrending, Stylized Facts and the Business Cycle', *Journal of Applied Econometrics*. 8(3). 231 – 247.

King, R. G. & Rebelo, S. T. 1993. 'Low frequency filtering and real business cycles', *Journal of Economic Dynamics and Control*. 17. 207 – 231.

Koopmans, L. H. 1974. *The Spectral Analysis of Time Series*. Academic Press: New York.

Liu, D. & Gupta, R. 2007. 'A Small Scale DSGE Model for Forecasting the South African Economy', *South African Journal of Economics*. 75(2). 179 – 193.

Marcet, A. & Ravn, M. O. 2003. 'The HP-Filter in Cross-Country Comparisons', *CEPR Working Paper 4244*.

Moolman, E. 2004. 'A Markov Switching Regime Model of the South African Business Cycle', *Economic Modelling*. 21. 631 – 646.

Parzen, E. 1961. 'Mathematical Considerations in the Estimation of Spectra', *Technometrics*. 3(2). 167 – 190.

Pedersen, T. M. 2002. 'Spectral Analysis, Business Cycles, and Filtering of Economic Time Series; With MATLAB Applications', *Institute of Economics, University of Copenhagen & Ministry of Economic and Business Affairs. Phd Thesis, Unpublished Manuscript*. Correspondence: Torben Mark Pedersen, tmp@oem.dk.

Pedersen, T. M. 2001. 'The Hodrick-Prescott filter, the Slutsky effect, and the distortionary effect of filters', *Journal of Economic Dynamics and Control*. 25. 1081-1101.

Pedersen, T. M. 1998. 'How Long are Business Cycles? Reconsidering Fluctuations and Growth', *Institute of Economics Discussion Paper no. 98/24, University of Copenhagen*.

South African Reserve Bank. 2007. *Quarterly Bulletin* 3.

Rand, J. & Tarp, F. 2001. 'Business Cycles in Developing Countries: Are They Different?' *Centre for Research in Economic Development and International Trade (CREDIT) Research Paper 01/21*.

Razzak, W. A. & Dennis, R. 1999. 'The output gap using the Hodrick-Prescott filter with a non-constant smoothing parameter: an application to New Zealand', *Reserve Bank of New Zealand discussion paper G95/8*.

Schlicht, E. 2004. 'Estimating the Smoothing Parameter in the So-Called Hodrick-Prescott Filter', *Discussion Paper No. 1054, University of Munich*.

Shibata, R. 1981. 'An Optimal Autoregressive Spectral Estimate', *The Annals of Statistics*. 9(2). 300 – 306.

Singleton, K. J. 1988. 'Econometric issues in the analysis of equilibrium business cycle models', *Journal of Monetary Economics*. 21. 361 – 368.

Venter, J. C. 2005. 'Reference turning points in the South African business cycle: Recent developments', *South African Reserve Bank*.

Warner, R. M. 1998. *Spectral Analysis of Time-Series Data*. The Guilford Press: New York.

APPENCIDES: MATLAB CODE

a) PlotQ.m

```
function PlotQ(x, lamdaMax);

l = 1:lamdaMax;
for m = 1:length(l)
    Q(m) = HP(x, l(m));
end

min = Q(1);
lamdamin = 1;

for n = 1:length(Q)
    if (min > Q(n))
        min = Q(n);
        lamdamin = n;
    end
end

lamdamin %the function reports the lamda which minimises Q

plot (l, Q, 'DisplayName', 'Q vs lamda', 'XDataSource', 'lamda',
'YDataSource', 'Q');

%this function PlotQ(...,...) implements the procedure described in
%Pedersen (2001: 1090 - 1092). The variable Q is defined in HP.m and
allows
%for the use of different cut-off frequencies (business cycle duration)
and
%can use quarterly, annual or monthly data.

%to implement, save both HP.m and PlotQ.m in the work file in MATLAB,
load
%the gdp data, specify the cut-off frequency and run the PlotQ.m in the
%command window by typing the function PlotQ(...,...).

%The output shows the Q-stat for various values of lamda and calculates
the
%value that minimises Q. This is the optimal lamda in the sense
described
%by Du Toit (2007) and Pedersen (2001).

%Du Toit, L. C. 2007. 'Optimal Hodrick-Prescott Filtering for South
%Africa'. Unpublished manuscript. Stellenbosch: University of
Stellenbosch,
%Department of Economics. <14038919@sun.ac.za>.

%Pedersen, T. M. 2001. 'The Hodrick-Prescott filter, the Slutsky
effect,
%and the distortionary effect of filters',
%Journal of Economic Dynamics and Control. 25. 1081-1101.
```

b) HP.m

```
function Q = HP(y, lamda);

% lamda = smoothing constant;
% y = log real gdp;
% Q is the weighted absolute value of the difference between the ideal
and
% the distortionary filter
% Hhp is the power transfer function of the HP filter
% H is the power transfer function of the ideal filter

w1 = pi / 10; %cutoff frequency for ideal filter, specified according
to the view of the duration of business cycles

[S, w] = pyulear(y, 2); %this estimates the spectral density

deltaw = w(2) - w(1);
NormSum = 0;
Q = 0;

for n = 1:length(w)
    Hhp(n) = abs((4 * lamda * (1 - cos(w(n)))^2) / (4 * lamda * (1 -
cos(w(n)))^2 + 1))^2;
    if (abs(w(n)) < w1)
        H(n) = 0;
    else
        H(n) = 1;
    end
    NormSum = NormSum + 2 * S(n) * deltaw;
end

for n = 1:length(w)
    v(n) = 2 * S(n) * deltaw / NormSum;
end

for n = 1:length(w)
    Q = Q + abs(H(n) - Hhp(n)) * v(n);
end

%figure;
%plot (w, Hhp, 'DisplayName', 'S vs w', 'XDataSource', 'w',
'YDataSource', 'S');
%hold on;
%figure;
%plot (w, S, 'DisplayName', 'S vs w', 'XDataSource', 'w',
'YDataSource', 'S');

%figure(gcf);
```