How Reliable is the Band Pass Filter?

Ashton de Silva*[a]
Farshid Vahid*

May 2001

*DEPARTMENT OF ECONOMETRICS
AND BUSINESS STATISTICS
MONASH UNIVERSITY AUSTRALIA
[a]: Email: Ashton.deSilva@BusEco.monash.edu.au

Title: How Reliable Is the Band-Pass Filter?

Abstract

The aim of this paper is to examine the ability of the Band Pass (BP) filter to extract "business cycles" from non-stationary data. A "business cycle" is defined as a completed cyclical movement of 6 to 32 quarters' duration. We simulate time series from data generating processes (DGPs) that have stochastic trends and 6-period to 32-period cycles, pass them through the BP filter and compare the filter output with the actual cycle. Using the four estimated models in Harvey and Jaeger (*Journal of Applied Econometrics*, 1993) as the basis of our Monte Carlo simulations, we find no significant correlation between the actual cycles and BP filter outputs for three out of the four DGPs. This finding entices further research into the relationship between the success of BP filter in isolating the cyclical component and specific features of the DGP.

Sensitivity analysis is conducted with respect to four characteristics of the DGP: the amplitude of the cycle, the duration of the cycle, the ratio of the variance of the trend innovation to the variance of cycle innovation, and the relative variance of measurement noise. The relative size of the trend error appears to impede the ability of the BP filter greatly. The relative size of the irregular dispersion significantly deteriorates the accuracy of the BP filter, however less severely than the trend counterpart. There is a positive relationship between the accuracy of the BP filter and the amplitude of the cycle. The relationship between duration of the cycle and BP filter's success in extracting it is non-monotone. Other things equal, the BP filter is most successful when the true cycle has a period of 10 quarters.

JEL Classification: C22, E32.

1 Introduction

Time series analysis has been traditionally involved with the decomposition of time series into trend, cycle, seasonal and irregular components. Time series decomposition methods have found their way into applied macroeconomics, because macroeconomists are interested in identifying the importance of innovations that have permanent effect, relative to those that have transitory effect, in the total variation of economic variables. Baxter and King (1999) propose an approximate band-pass filter (which we call the BK filter) that is designed to extract cyclical components of a band of frequencies from stationary time series. Our paper quantifies the effect of trend and irregular components on the performance of the BK filter.

The design of the BK filter is based on the theory of linear filters and the spectral representation of stationary time series. This filter is designed to pass periodic components of certain frequencies of interest, and suppress all components of other frequencies. The frequency band that Baxter and King use corresponds to the "business cycle frequencies" (those with periods of one and a half to eight years). The discussion of whether or not the output of a band-pass filter can be called the "business cycle component" of a macroeconomic time series, is outside the scope of this paper. Interested readers can find some idea about the controversy surrounding the use of time series decomposition methods in Canova (1998a,b), Burnside (1998) and Harding and Pagan (1999).

Band-pass filtering as a method of extracting cyclical components has several appealing features. First, it is easy to implement. It will not be surprising if widely used econometric time series packages, such as EVIEWS, add a BK filtering procedure to their menu in their future updates. Secondly, because band-pass filters are not "model-based", in the sense that they do not require that a statistical model be fitted to the data first, band-pass filtering can facilitate communication among researchers who use the same data. This means that two researchers who use the same data set and use BK filter with the same truncation lag to isolate components of frequency band $\left(\underline{\lambda}, \overline{\lambda}\right)$, will obtain exactly the same data. In contrast, for example, two researchers who use the same data set and extract Beveridge-Nelson cycles from fitted models to their time series, may obtain different results depending on their modelling strategy (e.g. univariate or multivariate, allow or not allow moving average terms, method of lag selection, etc.). For the same reason, as long as historical data are not revised, addition of new observations will not affect the historical filtered output, whereas it will affect model-based estimates of cyclical component for all time periods. Finally, band-pass filtering has a sounder theoretical basis than exponential smoothing and Hodrick-Prescott filters that are used for isolating the cyclical components of economic time series. Band-pass filters are designed to suppress all components with frequencies outside a user specified frequency band, whereas smoothing filters at best can only remove the trend (very low frequency components). It is true that the BK filter is an approximate band-pass filter (the ideal filter is a practically infeasible double infinite filter), but the approximation is also based on reasonable grounds (see Baxter and King 1999).

On closer scrutiny, however, the above features can be looked at as the *dangers* of band-pass filtering. The ease of implementation may result in band-pass filters being used indiscriminately and without much thought, the same way that X11 and X12 filters are used for removing seasonal components. The claim that the approximate band-pass filters are not model-based is also not strictly true. As we shall see later, the BK filter performs best when the data generating process is of a particular type. When we apply BK filter to different time series, we are implicitly assuming (or hoping) that this fixed model type is appropriate for all time series. The existing problem in communicating our fitted time series models with sufficient detail so that others can reproduce our results, is not a compelling reason to advocate fitting a fixed model to all time series. Finally, the theoretical foundation of band-pass filters is the spectral representation theorem of *stationary* time series. When time series have dominant stochastic trends, as in most macroeconomic time series, the output of the filter is influenced by the trend, and may not be a good reflection of the cyclical component of the series.

This paper is an investigation to this latter point. We generate time series with stochastic trends and cycles. The cycles are designed to have a single dominant frequency within the "business cycle" frequency band. We then compare the cyclical component with the output of the BK filter. The comparison is made using several methods. The result are sorted in a way that relative success of the BK filter as a function of the parameters of the DGP can be understood.

The plan of the paper is as follows. In Section 2, we outline the design and the theoretical foundation of the BK filter. In Section 3, we provide the details of our simulation exercise, and present our results. Section 4 concludes.

2 The BK filter

The basis of frequency domain analysis is the Cramér Representation Theorem, which states that any zero mean stationary stochastic process y_t can be written as

$$y_t = \int_0^\pi u(\lambda)\cos(\lambda t)\,d\lambda + \int_0^\pi v(\lambda)\sin(\lambda t)\,d\lambda,$$

where the amplitudes of the periodic component of frequency λ, i.e. $u(\lambda)$ and $v(\lambda)$, have mean zero, have equal variances denoted by $2f(\lambda)$, are uncorrelated with each other and are uncorrelated with the amplitudes at other frequencies[1]. This implies that the variance of the process is

$$\gamma(0) = \mathrm{Var}(y_t) = 2\int_0^{\pi} f(\lambda)\,d\lambda$$

The function f, which shows the contribution of cycles of different frequencies to the total variation in the process is the *spectral density function* or the *power spectrum* of the series. When this time series is passed through a linear time invariant filter $\alpha(L) = \sum_{j=-r}^{s} \alpha_j L^j$, the output, $z_t = \alpha(L)\, y_t$, is a stochastic process whose spectral density at frequency λ is

$$f_z(\lambda) = \left|\alpha\left(e^{-i\lambda}\right)\right|^2 f_y(\lambda),$$

where the function $\alpha\left(e^{-i\lambda}\right) = \sum_{j=-r}^{s} \alpha_j e^{-i\lambda j}$ is called the *frequency response function* of the filter. This shows that a filter with frequency response function that is equal to 1 on a desired frequency band and zero elsewhere will isolate the periodic components of the time series with frequencies within the desired band.

Baxter and King (1998) use the above theory and design a symmetric band-pass filter for frequency band $\left(\underline{\lambda}, \overline{\lambda}\right)$. The theoretical band pass filter is infinite with weights,

$$\alpha_0 = \frac{\overline{\lambda}}{\pi} - \frac{\underline{\lambda}}{\pi} \text{ and } \alpha_j = \frac{\sin\left(j\overline{\lambda}\right)}{j\pi} - \frac{\sin\left(j\underline{\lambda}\right)}{j\pi} \text{ for } j = \pm 1, \pm 2, \pm 3, \ldots .$$

In order to make the filter operational, Baxter and King truncate the filter at a finite lag K in such a way that the "distance" between the spectral density functions of the output of the theoretical filter and that of the truncated filter is smallest, subject to a constraint. This constraint restricts that the spectral density of the output of the truncated filter at frequency zero to be the same as that of the ideal filter, i.e. zero. This constraint ensures that no frequency zero component passes through the truncated filter. The resulting $\mathrm{BK}\left(\underline{\lambda}, \overline{\lambda}, K\right)$ filter has weights,

$$\alpha_j^K = \begin{cases} \alpha_j + \theta & \text{for } j = 0, \pm 1, \pm 2, \ldots, \pm K \\ 0 & \text{otherwise,} \end{cases}$$

$$\text{where } \theta = \frac{-\sum_{j=-K}^{+K} \alpha_j}{2K+1}.$$

Truncation can distort the dynamic properties of the series. As the following graph illustrates, truncation *exacerbates* the relative importance of cyclical

---

[1] This is an oversimplified, intuitive and mathematically imprecise explanation of the Spectral Representation Theorem. For a mathematically correct exposition, refer to a graduate time series textbook, such as Brockwell and Davis (1991).

components at some frequencies and *compresses* the importance of some others within the desired frequency range. It also shows that there is some *leakage* of periodic components with frequencies outside the desired range. The restriction on frequency zero ensures that we will not have any zero frequency components leak through the filter.

Frequency Response Function of the Truncated Filter

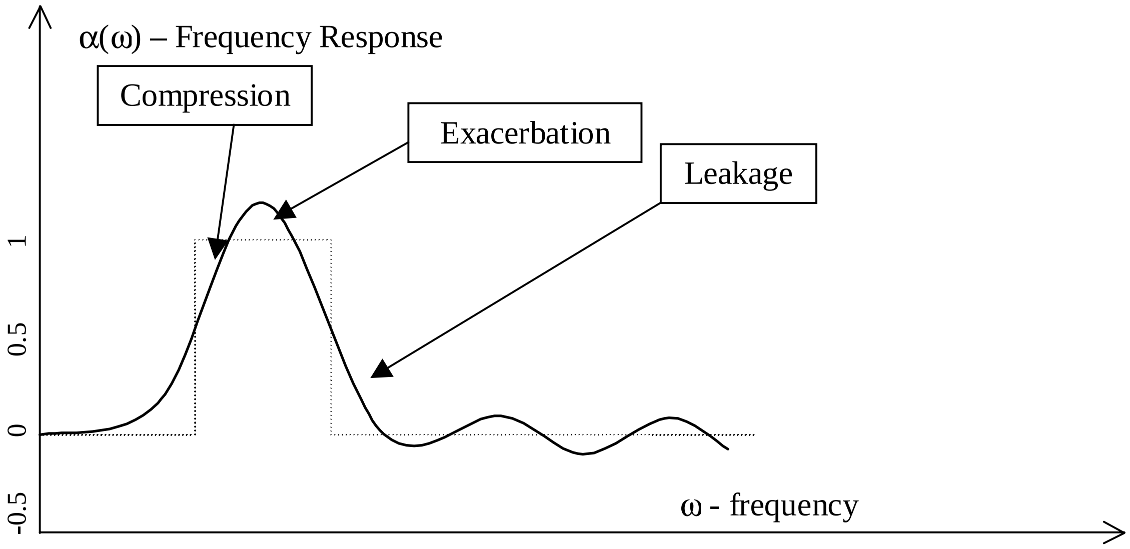

Figure 1

Using the US data and experimenting with different truncation lags, Baxter and King establish that the statistical properties of the filtered series stabilizes by the truncation lag 12, in the sense that there is not a significant change in covariance properties of the filter output as truncation lag is increased beyond 12. Athanasopoulos (2000) comes to the same conclusion using the Australian data. Anticipating that most applied researchers are going to follow Baxter and King and use BK filter with truncation lag 12, we base our investigation on this particular filter.

3 Investigation

The strategy of the investigation is two-fold. The first objective is to examine the consistency of the BK filter with respect to a set of empirical specifications. The second exercise examines the sensitivity of the BK filter with respect to specific parameters of the structural time series model.

The consistency aspect of the investigation subscribes to an earlier paper by Harvey and Jaeger (1993). Harvey and Jaeger applied the Hodrick-Prestcott filter to four economic series. The results of which are presented in Table 1. The model postulated by Harvey and Jaeger is the structural time series model of the following form:

$$y_t = \mu_t + \psi_t + \epsilon_t, \quad \epsilon_t \sim NID(0, \sigma^2_\epsilon), \text{ for t = 1,....T}$$

In the above model y_t represents the observed series whilst μ_t, ψ_t and ε_t represent the trend, cycle and irregular component respectively. The prescribed trend component is of a local linear nature and defined by:

$$\begin{array}{ll} \mu_t = \mu_{t-1} + \beta_{t-1} + \eta_t & \eta_t \sim NID(0, \sigma^2_\eta) \\ \beta_t = \beta_{t-1} + \zeta_t & \zeta_t \sim NID(0, \sigma^2_\zeta) \end{array}$$

The stochastic cycle is defined by:

$$\begin{array}{ll} \psi_t = \rho \cos \lambda_c \psi_{t-1} + \rho \sin \lambda_c \psi^*_{t-1} + \varkappa_t, & \varkappa_t \sim NID(0, \sigma^2_\varkappa) \\ \psi^*_t = -\rho \sin \lambda_c \psi_{t-1} + \rho \cos \lambda_c \psi^*_{t-1} + \varkappa^*_t, & \varkappa^*_t \sim NID(0, \sigma^2_\varkappa) \end{array}$$

where

$0 \leq \rho \leq 1$ represents the dampening parameter

λ_c measures frequency cycle in radians

All disturbances are assumed to be independent of each other.

3.1 Empirical Performance

Table 1 Empirical Specifications of Structural Time Series Model[2]

Series	**Restrictions**	σ^2_η	σ^2_ζ	σ^2_ε	$\sigma^2_\varkappa$	ρ	$2\pi/\lambda_c$
DGP-M	$\sigma^2_\eta = 0$	-	47	0	25	0.73	6
DGP-P	$\sigma^2_\eta = 0$	-	19	3	79	0.94	13
DGP-G	None	0	8	0	625	0.92	22.2
DGP-A	$\sigma^2_\eta = 0$	-	21	438	36	0.97	13

*All variances are multiplied by 10^7

* $2\pi/\lambda_c$ Indicates the duration of the cycle

Table 1 presents the parameter values prescribed by Harvey and Jaeger. The value of the disturbance term η_t is either estimated or restricted to be zero in all in four cases. The restriction is motivated by two reasons. Firstly a better fit is obtained when the stochastic cyclical component is

---

[2] Extracted from Harvey and Jaeger (1993) Table 2 pg236
DGP-M: US Money Base (1959:1-89:4)
DGP-P: US Prices (1954:1-89:4)
DGP-G: US GNP (1954:1-89:4)
DGP-A: Australian GDP (1964:1-88:4)

enforced and secondly as ρ & $\lambda_c \longrightarrow 0$ the local linear trend model emerges as a limiting case of the smooth trend and cycle model.

To evaluate the consistency of the BK filter each DGP (as prescribed by table 1) is simulated 2000 times. In theory after the BK filter is passed through the generated model only the observed cyclical element should remain. Comparisons between the observed cyclical component and the BK extracted cycle is then conducted by means of the correlation coefficient. Naturally a correlation coefficient of 1 is desirable.

Table 2 Summary Statistics of the Correlation Coefficient

	DGP-M	DGP-P	DGP-G	DGP-A
AVERAGE	0.130864945	0.594831435	0.885023075	0.650452516
CoV	78.60439925	28.63231376	6.034835373	19.89312125
95% C.I.	[-0.07,0.329]	[0.185,0.851]	[0.757,0.958]	[0.357,0.856]
MAX	0.43312251	0.92382249	0.97842727	0.91226012
MIN	-0.23805945	-0.17323252	0.59858806	0.066261946
RANGE	0.67118196	1.09705501	0.37983921	0.845998174

Table two presents a summary of the distribution of the correlation coefficients with respect to the differing Data Generating Processes. Clearly the result raises more questions than it provides conclusions. The BK Filter is ineffective when applied to the DGP model of type M. The average correlation coefficient is 0.13 indicating little (if any) relationship between the extracted and observed cycle. Furthermore the value of the correlation coefficient corresponding to the 2.5th percentile is negative and continues to be negative till it reaches the 11th percentile. These results imply the BK filter performs poorly when applied to the parameter structure of type DGP-M, potentially extracting a cycle that is approximately the linear inverse of the true phenomena.

The BK filter's ability to extract the business cycle from the DGP-P is also poor, however is an improvement when compared to the previously discussed DGP-M. The average correlation statistic has increased to a value of approximately 0.6 indicating a reasonably strong positive linear correlation between the extracted and observed cycle. There is a large dispersion however associated with the correlation statistic indicating the BK filter's ability to extract the observed cycle is hazardous at best.

The remaining two DGPs provide better conditions for the BK filter. Clearly when the BK filter is applied to DGP-G the extracted cycle is very close to the observed cycled as indicated by the high mean correlation coefficient of (0.88) and relatively tight dispersion (coefficient of variation 6%).

Although the DGPs are different along too many dimensions to allow a definitive conclusion from Table 2, we can see some patterns. The BK filter seems top be least successful in the DGPs with cycles at or close to the upper or lower bound of the frequencies of the BK filter (DGP-M and DGP-P). the greater variance of stochastic innovation in the cyclical component is probably responsible for the success of the BK filter in extracting the cyclical component of DGP-G. Observation noise σ_{ε} seems to matter relatively less than other factors as indicated by the perfomance when applied to DGP-A To ascertain these claims we perform a comprehensive sensitivity analysis in the next section.

DGP-G will be employed as a point of reference for the individual parameter effects. Essentially the analysis examines what conditions will cause the BK filter's performance to improve or more importantly deteriorate, (where the performance is measured by the correlation statistic). The parameters of interest are $\sigma_{\epsilon}^2, \sigma_{\zeta}^2, \rho$ & λ_c.

3.2 Sensitivity Analysis

The first parameter of interest is the dampening parameter ρ, The sensitivity analysis will involve ranging ρ from 0.05 to approximately 0.95. From the illustration below as ρ approaches 1 the amplitude of the cycle becomes more pronounced. Intuitively this suggests when the BK filter is applied to a DGP with ρ value close to 1 it will extract a cycle that would closely approximate the observed cycle. Applying this theory to the four empirical models investigated indicates DGP-A provides the most favourable conditions for the BK filter. From the results presented in Table 2 however it is hard to gauge the sensitivity of the BK filter with respect to ρ as all the ρ values are relatively close except for DGP-M. Interesting the parameter specification of DGP-M provides the least favourable conditions for BK filter.

Figure 3 provides the sensitivity results when the BK filter is applied to DGPs with varying dampening parameter values. The graph indicates a positive relationship between the performance variable r and the dampening parameter ρ. The success of the BK filter as ρ approaches 1 is further illustrated by the smaller confidence intervals(Lower Limit = 2.5^{th}Percentile & Upper Limit = 97.5^{th} Pecentile). Clearly however the value of ρ cannot be the only influential parameter on the performance of the BK filter (i.e. The most favourable DGP-G does not have the highest value of ρ.)

The second parameter of interest is λ_c. λ_c specifies the frequency of the cycle. The cycle band ranges from 6 to 32 quarters that is $0.2 \leq \lambda_c \leq 1$. Figure 1 provides some intuitive appeal when considering what effect λ_c.might have. The diagram illustrates as the cycle approaches either boundary of the band the filter is susceptible to compression. The compression effect is supported by the inferior performance of the BK filter when applied to DGP-M (where λ_c corresponds to the lower boundary value). Figure 4 displays the sensitivity results with respect to λ_c. The results are symmetric for λ_c but non-symmetric for $2\pi/\lambda_c$ (the actual frequency of the cycle,

Figure 5).

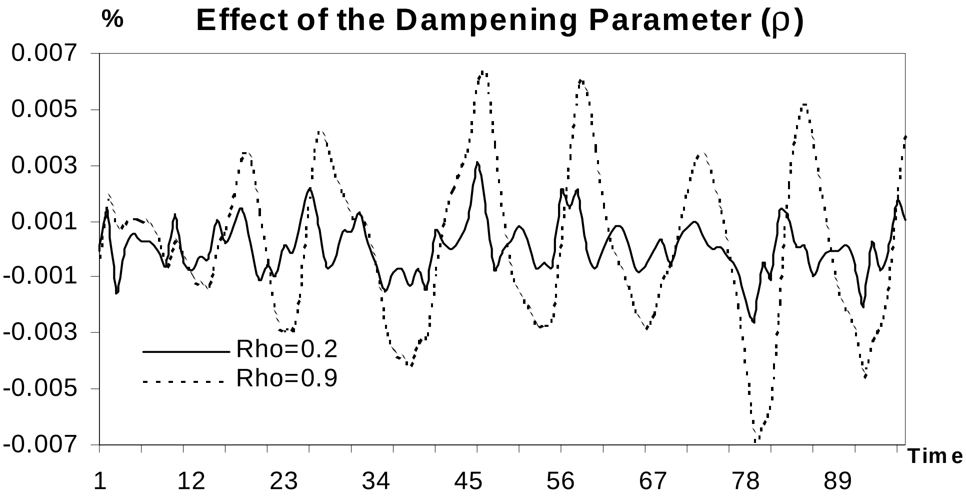

Figure 2

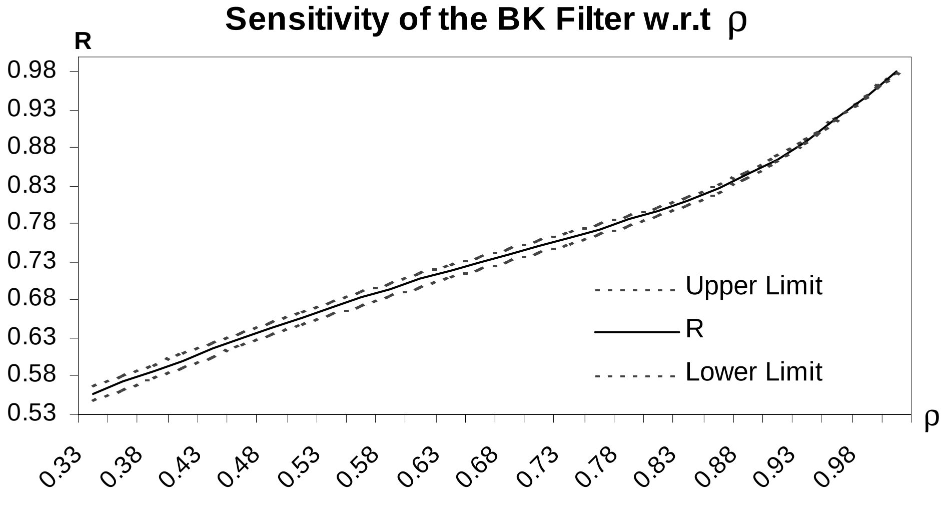

Figure 3

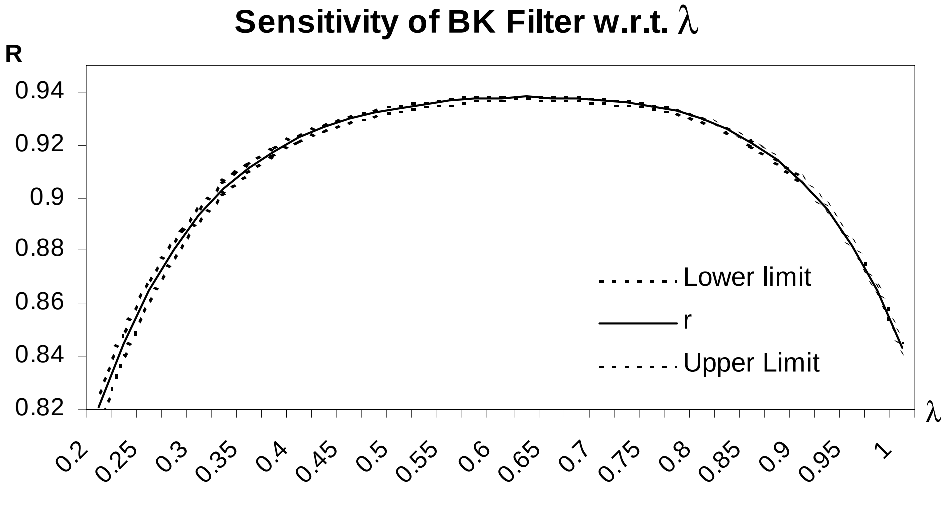

Figure 4

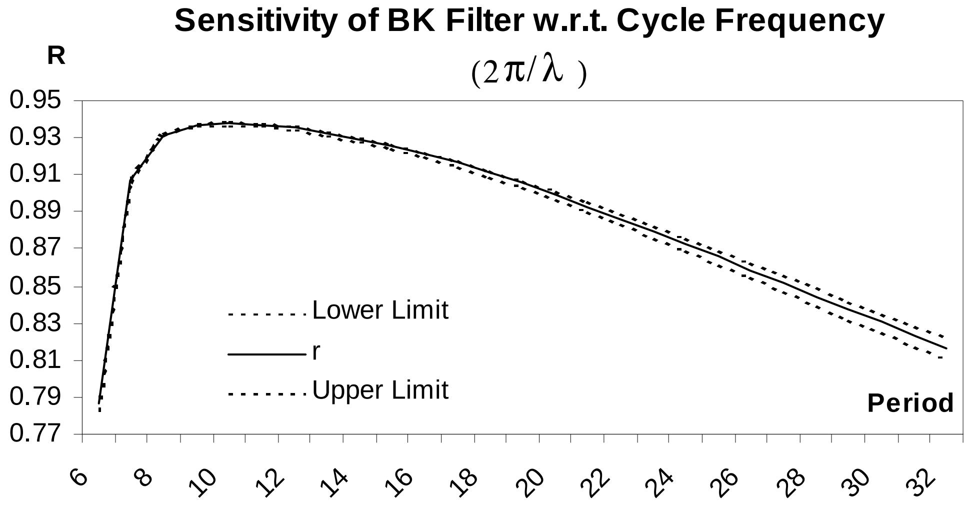

Figure 5

The previous two sensitivity exercises are informative and provide a good explanation for the empirical results displayed in Table 2, however are incomplete. The sensitivity analysis so far has ignored the disturbance effects of the trend, irregular and cyclical components components. The final two sensitivity exercises evaluate the BK filter's accuracy with respect to these parameters. Firstly the effect of the ratio $\sigma_\epsilon/\sigma_\varkappa$ is investigated followed by $\sigma_\zeta/\sigma_\varkappa$.

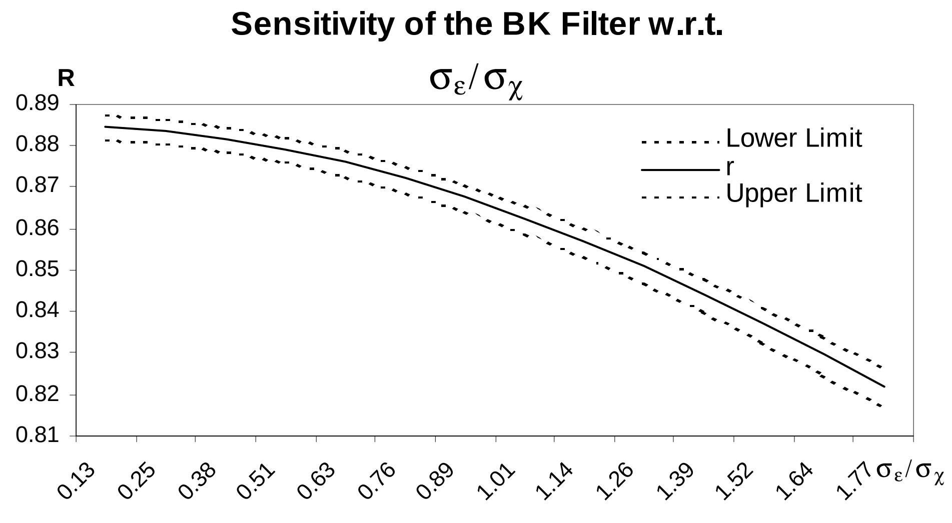

Figure 6

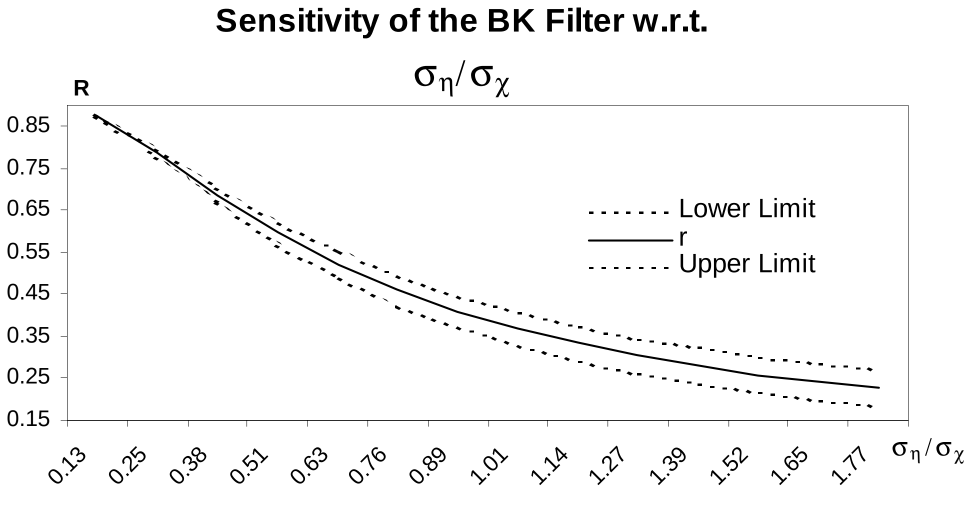

Figure 7

As expected both results suggest that as the disturbances of the irregular and trend components become more dominant the performance of the BK filter deteriorates The severity however is noteworthy, clearly the BK filter's performance is highly sensitive to the relative stochastic trend term $(\sigma_\eta/\sigma_\varkappa)$ penalising the performance heavily. Although the irregular component has the same directional impact on the BK filter performance it is much less influential than the trend disturbance term.

With respect to the structural time series model the BK filter appears not to be robust displaying highly sensitive behaviour to the parameter values of the structural times series model.

3.3 Conclusion

The aim of this paper was to analyse the effectiveness of the BK filter to extract the business cycle from a non-stationary time series model. Clearly the ability of the BK filter is sensitive to parameter values of the model. Generally as the stochastic terms of the model increase relative to the cyclical stochastic term the ability of the BK filter to extract the observed cycle deteriorates. The effect of the dampening parameter is positive, that is as $\rho \longrightarrow 1$ the ability of the BK filter to extract the observed cycle improves. Finally λ_c also influences the BK filter however in a symmetric fashion. From the evidence provided above it is obvious that the BK filter is not robust.

The above results help explain the outcome of the empirical evaluation conducted earlier. It is important to recognise however that this paper is limited with its respect to the rigidness of model specification. Further investigation needs to be conducted before conclusive evidence can be drawn on the effectiveness of the BK filter. Extending the analysis to I(1) models as well as the I(2) as investigated in this report would be a beneficial exercise. Another important question is to find out if the BK filter can be applied to two series with common stochastic cycles.

References

[1] Athanasopoulos, G., 2000. Extracting the Business Cycle: Band Pass Filters for Australian economic Time series

[2] Baxter, M. and King, R,. 1999. Measuring Business Cycles: Approximate Band-Pass Filters for Economic Time Series. The Review of Economic and Statistics Vol 81(4) 575-593

[3] Bergstrom, V., 1984. Measuring and interpreting Business Cycles. Oxford University Press

[4] Beveridge, S and Nelson, C., 1981. A New Approach to Decomposition of Economic Time Series into Permanent and Transitory Components with Particular Attention to Measurement of the 'Business cycle'. Journal of Monetary Economics. Vol 7

[5] Burnside, C., 1998 Detrending and business cycle facts: A comment. Journal of Monetary Economics Vol 41 pg 513-552

[6] Brockwell, P.J.. Davis, R, A., ..1991. *Time series : theory and methods* New York : Springer-Verlag

[7] Canova, F., 1998. Detrending and business Cycle Facts. Journal of Monetary Economics Vol 41 pg 533-540

[8] Canova, F., 1998. Detrending and business Cycle Facts. Journal of Monetary Economics Vol 41 pg475-512

[9] Harvey, A., 1993. Time Series Models. Sydney : Harvester Wheatsheaf.

[10] Harvey A, Jaeger A., 1993. Detrending, Stylised facts and the Business Cycle. Journal of Applied Econometrics Vol 8 231-247

[11] Harding, D and Pagan., A Dissecting the Cycle. Melbourne Institute Working Paper No. 13/99