

Performance and risk analysis of the Hodrick-Prescott filter

Jonas Bergroth

Department of Mathematics
Stockholm, June 2011

Abstract

We evaluate the consistent estimator of the noise-to-signal ratio parameter, the so called DDR estimator, in the Hodrick-Prescott filter introduced in Dermoune *et al.* (2008, [2]) and suggest two ways to update it to make the probability distribution of the cyclical component as closer to Gaussian as possible. When comparing the results between the different estimators, both the trend and the cyclical component are analyzed to decide which of the three estimators generates the best result. It appears that in most cases the DDR estimator gives the best trend component. We then apply the filter to standard risk calculations in the sense that we compare risk figures such as Value-at-Risk and expected volatility obtained for the original time series and the filtered one. The observed variation in the ratio of these two risk figures may be useful to enhance the performance of the underlying optimal portfolio.

Keywords: Hodrick-Prescott filter, trend, noise-to-signal ratio, Value-at-Risk, expected volatility

Acknowledgements

I would like to thank my supervisor at KTH Mathematical Statistics, Boualem Djehiche, for great feedback and guidance as well as my supervisors at IPM Informed Portfolio Management; Anders Lindell, Mattias Jansson and Hanna Sahle, for feedback and ideas.

Stockholm, June 2011

Jonas Bergroth

Table of Contents

1	Introduction	1
1.1	Noise-to-signal estimators	1
1.2	Trend	1
1.3	Outline	2
2	Theoretical Background	3
2.1	The Hodrick-Prescott filter	3
2.2	The DDR estimator of the noise-to-signal ratio	5
2.2.1	Noise-to-signal ration when using D^3	6
2.2.2	Data fitted noise-to-signal ratio α	7
2.3	Risk measures	8
2.3.1	Expected volatility (ETE)	8
2.3.2	Value-at-Risk (VaR)	9
3	Time series under study	11
3.1	Macroeconomic time series	11
3.2	Price series	12
4	Performance of the filter	13
4.1	Applying the filter	13
4.2	Analysis of the trend component x	14
4.3	Analysis of the Cyclical component u	16
4.4	Analysis of the noise-to-signal ratio α	17
5	Computing α for different time series	19
5.1	Macroeconomic time series	19
5.2	Seasonally adjusted macroeconomic time series	22
5.3	Price series	25
5.4	Inflection time series	28
5.5	Overall choice of α	31
6	Risk calculations	33
6.1	The portfolio	33
6.2	Constructing the filtered data	33
6.3	Distribution of daily returns	33
6.4	Risk calculation	35
	Bibliography	39

Introduction

There are two fairly known methods to fit a smooth trend to a time series. By means of moving average and by fitting regression curves. Unfortunately, both have known drawbacks, in the regression curves the first possibility to fit a straight line is often a too big of an assumption. The next question is how many degrees of freedom that are needed now to get a smooth trend which would persist even in the future.

Because of this the preferred choice is to use moving averages. It is shown, e.g. in Macaulay (1931, [6]), that moving averages in some special cases can be seen as moving regression curves. The drawback with this method is that the time series needs to be long, since the beginning and end parts of the time series does not yield trend values.

A third way to generate a trend was introduced by Whittaker (1923, [13]), where he minimizes two sums of squares. The first sum is the deviation of the trend from the original time series and the second sum specifies the degree of smoothness of the trend. This method was later on introduced to finance with the Hodrick-Prescott filter, by Hodrick & Prescott (1997, [4]), where some fixed numbers were proposed for the parameter, the noise-to-signal ratio, that suites to financial data with different time intervals.

1.1 Noise-to-signal estimators

Different papers have argued that these fixed values, for the noise-to-signal ratio, are not optimal and new estimation methods have been introduced. One of these methods is the Generalized Cross- Validation method, see Weinert (2007, [11]), another one based on likelihood is introduced in Schlicht (2005, [10]).

1.2 Trend

One thing that can be good to ask yourself before reading this thesis is: "What is a trend". If you have an answer to that question it will probably not be the same as any friend of yours, since there is no right answer to that question, because what a trend is has never been decided. The article White *et al.* (2011, [12]) discuss the subject and their opinion of what a trend is

- it should have a direction - that is, it should generally be higher at one end of the series than the other, so that it will seem generally to increase or decrease throughout
- it should be somewhat smooth (although there may be an unsmooth component)

This is the idea of a trend that we will use in this thesis, with the addition that a time series can be broken in to smaller parts and thereby fulfills these criteria.

1.3 Outline

Chapter 2 recalls the mathematical framework that we will use in the thesis, such as the Hodrick-Prescott filter, the different noise-to-signal estimators and risk measures. In Chapter 3 the various time series that we will use are presented, with the way to construct some of them since they consists of different assets that have expiry dates. In Chapter 4 and Chapter 5 we analyze the trend, the cyclical component and the smoothing parameter α and in of the Hodrick-Prescott filter for these time series. Finally in Chapter 6 we use filtered time series to calculate risk values as specified in the abstract.

Theoretical Background

This chapter recalls the theoretical background described in Schlicht (2005, [10]).

2.1 The Hodrick-Prescott filter

The purpose with the Hodrick-Prescott filter is to decompose a time series of observations, $x = (x_1, \dots, x_T) \in \mathbf{R}^T$, into a non stationary trend, $y = (y_1, \dots, y_T) \in \mathbf{R}^T$, and a cyclical residual component, $u = (u_1, \dots, u_T) \in \mathbf{R}^T$:

$$x = y + u. \tag{2.1}$$

To decompose x into y and u a positive smoothing parameter α is needed, then the weighted sum of squares can be minimized,

$$\|x - y\|^2 + \alpha \|D^n y\|^2, \tag{2.2}$$

with respect to y and usually $n = 2$ as e.g. in Dermonue *et al.* (2008, [2]), even though Araujo (2003, [1]) found that the third order shift operator were more appropriate for some FX rate series. In this paper the second order shift operator will be used if nothing else is mentioned.

As pointed out, e.g. in Pedersen (2001, [8]), the first term in (2.2) measures a goodness-of-fit by minimizing the deviation between the trend y_t and the observation x_t , and the second term, when $n = 2$, is a measure of the degree-of-smoothness.

D^2 is, as mentioned earlier, the second order forward shift operator which here is applied on the trend y ,

$$D^2 y_t := (y_{t+2} - y_{t+1}) - (y_{t+1} - y_t), \quad t = 1, \dots, T - 2.$$

This measures the deviation between the value of the trend at $t + 1$, y_{t+1} , and the linear interpolation between y_t and y_{t+2} .

This can be written in vector form as:

$$Py(t) = D^2 y_t, \quad t = 1, \dots, T - 2,$$

where, P is the following $(T - 2) \times T$ -matrix

$$P := \begin{pmatrix} 1 & -2 & 1 & 0 & \dots & 0 \\ 0 & & \ddots & & \ddots & \vdots \\ \vdots & \ddots & & \ddots & & 0 \\ 0 & \dots & 0 & 1 & -2 & 1 \end{pmatrix}.$$

But P is of rank $T - 2$, which implies that $v := Py$ does not determine a unique y , rather a set of solutions

$$y := \{P'(PP')^{-1}v + Z\gamma; \quad \gamma \in \mathbf{R}^2\},$$

Chapter 2. Theoretical Background

where Z is a $T \times 2$ -matrix and satisfies

$$PZ = 0, \quad Z'Z = I_2,$$

I_2 is the 2×2 identity matrix. The time series x in equation (2.1) may now be described in terms of (u, v) as

$$x = u + P'(PP')^{-1}v + Z\gamma, \quad (2.3)$$

for some $\gamma \in \mathbf{R}^2$.

The first-order condition to (2.2) gives,

$$(I_T + \alpha P'P)y = x, \quad (2.4)$$

and since $(I + \alpha P'P)$ is positive definite, the second order condition is satisfied as well and (2.4) can be uniquely solved as

$$y(\alpha, x) = (I_T + \alpha P'P)^{-1}x. \quad (2.5)$$

where I_T denotes the $T \times T$ identity matrix. Equation (2.5) defines how the trend y is decomposed from the time series y using the HP-filter, this equation is dependent on the smoothing parameter α and the disturbance operator P .

One way to estimate the smoothing parameter α , the one that Schlicht (2005, [10]) used, is to let the optimal solution $y(\alpha, x)$ in equation (2.5) be the best predictor of any trend y given the time series x , i.e.

$$y(\alpha, x) = E[y|x]. \quad (2.6)$$

Unfortunately, a conditional expected value use to be hard to calculate, but if we follow e.g. Schlicht (2005, [10]) or Dermoune *et al.*(2008, [2]), a widely used model, and assume that u and v are independent and normally distributed. This turns (x, y) into a normally distributed vector and makes estimations using (2.5) and (2.6) feasible.

$$\begin{pmatrix} u \\ v \end{pmatrix} \sim \mathcal{N}(0, \Sigma_{uv}), \quad (2.7)$$

with covariance matrix

$$\Sigma_{uv} := \begin{pmatrix} \sigma_u^2 I_T & 0 \\ 0 & \sigma_v^2 I_T \end{pmatrix}.$$

This makes the increments of the trend y following a Gaussian random walk, since, $v = Py$ gives $y_{t+2} - y_{t+1} = y_{t+1} - y_t + v_t$. From this we get that the time series x is a sum of the trend y which is generated by a Gaussian random walk and the normal distributed vector u . We can now describe (x, y) as a normal distributed vector using (2.3), as

$$\begin{pmatrix} x \\ y \end{pmatrix} \sim \mathcal{N}\left(\begin{pmatrix} Z \\ Z \end{pmatrix} \gamma, \Sigma_{xy}\right),$$

with covariance matrix

$$\Sigma_{xy} := \begin{pmatrix} \sigma_u^2 I_T + \sigma_v^2 Q & \sigma_v^2 Q \\ \sigma_v^2 Q & \sigma_v^2 Q \end{pmatrix},$$

where,

$$Q := P'(PP')^{-1}(PP')^{-1}P.$$

This yields an explicit expression for equation (2.6), namely

$$E[y|x] = Z\gamma + \sigma_u^2 Q[\sigma_u^2 I_T + \sigma_v^2 Q]^{-1}(x - Z\gamma),$$

if we now insert this into (2.5), the smoothing parameter α and the parameter γ satisfy

$$Z\gamma + \sigma_u^2 Q[\sigma_u^2 I_T + \sigma_v^2 Q]^{-1}(x - Z\gamma) = (I_T + \alpha P'P)x.$$

We then recall Theorem 1 in Schlicht (2005, [10]) saying that equation (2.6) holds if and only if $\gamma = Z'x$ and $\alpha = \sigma_u^2/\sigma_v^2$.

2.2 The DDR estimator of the noise-to-signal ratio

This section will present the explicit unbiased consistent DDR estimator of the noise-to-signal operator α , the one used in Dermoune *et al.* (2008, [2]), it will converge in probability as the length T of the time series tends to infinity.

To this purpose, we will use observations from the time series Px :

$$Px = v + Pu \sim \mathcal{N}(0, \sigma_v^2 I_{T-2} + \sigma_u^2 PP').$$

The elements $V(i, j)$ of the covariance matrix satisfy

$$V(i, j) = \sigma_v^2 \delta_i^j + \sigma_u^2 (PP')_{ij} = r_{|i-j|},$$

where,

$$r_k = \begin{cases} \sigma_v^2 + 6\sigma_u^2 & \text{if } k = 0 \\ -4\sigma_u^2 & \text{if } k = 1 \\ \sigma_u^2 & \text{if } k = 2 \\ 0 & \text{otherwise.} \end{cases} \quad (2.8)$$

We then use the well established fact that (see e.g. Proposition 2.1 in Giurcanu *et al.* (2002, [3]))

$$\hat{r}_k = \frac{1}{(T-2)-k} \sum_{j=1}^{T-2-k} Px(j)Px(j+k), \quad k = 0, 1, 2, \quad (2.9)$$

is an unbiased estimator of $r_k = E[Px(s)Px(s+k)]$ in the sense that $E[\hat{r}_k] = r_k$.

If we combine (2.8) and (2.9) we get some easily checked relations. We start with $k = 1$

$$E[\hat{r}_1] = -4\sigma_u^2,$$

this gives the following consistent unbiased estimator of σ_u^2

$$\hat{\sigma}_u^2 = -\frac{1}{4}\hat{r}_1 = -\frac{1}{4(T-3)} \sum_{j=1}^{T-3} Px(j)Px(j+1). \quad (2.10)$$

If we instead look at $k = 0$ we get

$$E[\hat{r}_0] = \sigma_v^2 + 6\sigma_u^2,$$

together with (2.10), we get a consistent unbiased estimator of σ_v^2

$$\hat{\sigma}_v^2 = \hat{r}_0 + \frac{3}{2}\hat{r}_1,$$

or

$$\hat{\sigma}_v^2 = \frac{1}{(T-2)} \sum_{j=1}^{T-2} Px(j)^2 + \frac{3}{2(T-3)} \sum_{j=1}^{T-3} Px(j)Px(j+1). \quad (2.11)$$

We now recall that $\alpha = \sigma_u^2/\sigma_v^2$, by using (2.10) and (2.11) we are now able to get a consistent estimator of the smoothing parameter α

$$\hat{\alpha} = -\frac{1}{4} \left(\frac{3}{2} + \frac{(T-3) \sum_{j=1}^{T-2} Px(j)^2}{(T-2) \sum_{j=1}^{T-3} Px(j)Px(j+1)} \right)^{-1}. \quad (2.12)$$

2.2.1 Noise-to-signal ration when using D^3

Sometimes the estimations above generate a negative α , it usually comes from the fact that the time series includes a lot of inflection points, see Figure 2.1




Figure 2.1: The Japanese Gross Domestic Product from 1980 to the third quarter of 2010, in first plot the whole time series is plotted and in the second some of the inflection points are highlighted in red.

In this occasions we use another estimator to calculate α , an estimator that corresponds to the third order shift operator instead. This is done in the same way as for the second order but with the difference that we have a new P matrix

$$P := \begin{pmatrix} 1 & -3 & 3 & -1 & 0 & \dots & 0 \\ 0 & & \ddots & \ddots & & \ddots & \vdots \\ \vdots & \ddots & & \ddots & \ddots & & 0 \\ 0 & \dots & 0 & -1 & -3 & 3 & 1 \end{pmatrix}.$$

As before we start with observations from the time series Px :

$$Px = v + Pu \sim \mathcal{N}(0, \sigma_v^2 I_{T-2}) + \sigma_u^2 PP',$$

with following elements, $V(i, j)$, in the covariance matrix

$$V(i, j) = \sigma_v^2 \delta_i^j + \sigma_u^2 (PP')_{ij} = r_{|i-j|},$$

Since we use a new P matrix, r_k is new as well

$$r_k = \begin{cases} \sigma_v^2 + 20\sigma_u^2 & \text{if } k = 0 \\ -15\sigma_u^2 & \text{if } k = 1 \\ 6\sigma_u^2 & \text{if } k = 2 \\ -\sigma_u^2 & \text{if } \frac{1}{2} k = 3 \\ 0 & \text{otherwise.} \end{cases}$$

We then use Proposition 2.1 in Giurcanu *et al.* (2002, [3]) and the fact that

$$\begin{aligned} E[\hat{r}_1] &= -16\sigma_u^2 \\ E[\hat{r}_0] &= \sigma_v^2 + 20\sigma_u^2, \end{aligned}$$

to get

$$\begin{aligned} \hat{\sigma}_u^2 &= -\frac{1}{15}\hat{r}_1 = -\frac{1}{15(T-4)} \sum_{j=1}^{T-4} Px(j)Px(j+1) \\ \hat{\sigma}_v^2 &= \hat{r}_0 - 20\sigma_u^2 = \frac{1}{(T-3)} \sum_{j=1}^{T-3} Px(j)^2 + \frac{4}{3(T-4)} \cdot \sum_{j=1}^{T-4} Px(j)Px(j+1), \end{aligned}$$

All of this gives the following α

$$\hat{\alpha}_{D^3} = -\frac{1}{15} \left(\frac{4}{3} + \frac{(T-4) \sum_{j=1}^{T-3} Px(j)^2}{(T-3) \sum_{j=1}^{T-4} Px(j)Px(j+1)} \right)^{-1}. \quad (2.13)$$

2.2.2 Data fitted noise-to-signal ratio α

We have also constructed two new ways to estimate a optimal α , where we use the calculated α from above and then try to find a local minimum point to some function in the region around this α . The first estimator fits the distribution of the cyclical component u with the normal distribution, according to (2.7). The other one minimizes the autocorrelation within the cyclical component u .

Normal fitted α

To estimate this α_{norm} we start by estimating α with (2.12), or in special cases (2.13), and then use the MATLAB function *fminsearch* to find a local minimum to a certain function $f(\alpha, u)$ with respect to α , starting at the estimated α .

The way this works is that we first calculate u by combining (2.1) and (2.4). We then calculate the standard deviation, σ_u , of u and finally we calculate the following sum of squares

$$\sum_{i=-300}^{300} \left(L \left(\frac{i}{100} \sigma_u \right) - \Phi \left(\frac{i}{100} \right) \right)^2,$$

where,

$$L \left(\frac{i}{100} \sigma_u \right) = \frac{1}{N} n_{\frac{i}{100}, \sigma_u}$$

$\Phi(x)$: Standard normal cumulative distribution function

and

$$\begin{aligned} n_{\frac{i}{100}, \sigma_u} &= \# \text{ of } u \text{ that is lower then } \frac{i}{100} \sigma_u \\ N &= \text{The length of } u \end{aligned}$$

These calculations are repeated until *fminsearch* finds the α that generates the lowest sum of squares and returns it as α_{norm} . This α_{norm} is the α that generates the u -vector that best approximates a normally distributed vector.

Autocorrelation fitted α

The α_{acorr} is mostly calculated in the same way as the normal fitted α but with another sum of squares

$$\sum_{l=1}^{100} r_u(l)^2,$$

where,

$r_u(l)$: Is the autocorrelation of u with lag l .

As above, this is repeated until the optimal α_{acorr} is found. What this α_{acorr} has generated is a u -vector that has the least dependence within it, in other words noise.

2.3 Risk measures

We will use two different types of risk measures in this thesis, expected volatility and Value-at-Risk. To calculate both of these risk measures we first need to calculate the weighted covariance matrix with the following elements

$$C_t(i, j) = 0, \quad t = 0$$

$$C_t(i, j) = 0.96C_{t-1} + 0.04 \log(1 + r_{i,t}) \log(1 + r_{j,t}), \quad t = 1, 2, \dots, T,$$

where,

$r_{i,t}$: Return of asset i at time t

T : Date of the day when the risk is calculated.

This will generate the $n \times n$ matrix, where n is the amount of assets within the portfolio, that will be used to calculate the two risk measures that we will use.

2.3.1 Expected volatility (ETE)

First the market value, MV , of each position is calculated, using

$$MV_{i,t} = P_{i,t} X_{i,t}, \tag{2.14}$$

where,

$P_{i,t}$: Current positions in asset i at time t

$X_{i,t}$: Price of asset i at time t .

Then the percentage of notional, pn , is calculated from

$$pn_{i,t} = \frac{MV_{i,t}}{n_t}, \tag{2.15}$$

where,

n_t : Total value of the whole portfolio is closed down at time t

Now the expected volatility can be calculated, using

$$ETE_t = \sqrt{252}(pn_t^T \times C_t \times pn_t). \tag{2.16}$$

2.3.2 Value-at-Risk (VaR)

Value-at-Risk is mostly calculated as expected volatility, we start by calculating each positions marked value with

$$MV_{i,t} = P_{i,t}X_{i,t}, \quad (2.17)$$

where,

$P_{i,t}$: Current positions in asset i at time t

$X_{i,t}$: Price of asset i at time t .

With these we calculate the Value-at-Risk, as

$$VaR_t = 1.645(MV_t^T \times C_t \times MV_t). \quad (2.18)$$

Time series under study

This chapter covers the various types of time series that we will apply the filter to and also what characterize each type of time series. The types of time series that we will use can be divided into two groups, these are; macroeconomic time series and "ordinary" price series. The main difference between the two types is that the macroeconomic time series has a much more obvious seasonal dependence.

3.1 Macroeconomic time series

The most eye catching with these type of time series is that they tend to have a quite obvious seasonal component. This becomes even more obvious since they are usually presented monthly or quarterly. Because of this distinct seasonal component, they also tend to be available in a seasonally adjusted type of time series. One example with a time series containing the seasonal component and one that has been seasonally adjusted can be seen in Figure 3.1




Figure 3.1: The Japanese Gross Domestic Product from 1980 to the third quarter of 2010, both seasonally adjusted and not.

We will apply the filter to a lot of different macroeconomic time series like; Gross Domestic Product, Current Account, Industrial Production and Trade Balance.

3.2 Price series

Price series are traded continuously and are therefore available in time spans that macroeconomic time series aren't. These kinds of series are often found in everything from tick, the time points when a trade is made, up to whatever time length you want.

Financial price series can be found in all sort of financial assets that are traded in a market. The ones that we will mostly use are the daily prices of futures and forwards, because the portfolios that we later on will calculate risk within are trading in these sorts of financial assets. However, there is one drawback with these type of securities and that is that they have an expiry date. So what we do is that we overlap each contract by one day. We then calculate the daily return for all days within each contract, this gives us returns for all days. We then put together a time series by starting at one level, say 100, and then take it times the first return, or more exactly the return plus one, and keep doing this for all returns to create a continuous time series. For a graphical description of this see Figure 3.2




Figure 3.2: Future contracts on the Swedish OMX30 index from 2008-07-09 to 2008-09-26, first plot are daily closing prices for four different contracts, in the middle plot are the daily returns from these contracts and in the last plot a time series put together using the daily returns.

The reason why we are able to start the new series at whatever point we want is that the filter will generate the same kind of time series even if you divide the original time series with any real number. The proof of this is presented in Chapter 4.4.

We will use a wide range of different price series, like; Stock index futures, future contracts on 10 years treasures and FX forwards.

Performance of the filter

In this chapter we will go through the returns from the filter and what happens when we change the input time series. We start by just applying the filter to some time series and look at the different parameters that the filter returns. We then look at what happens with α when we scale the input time series.

4.1 Applying the filter

To get a feeling how the filter works we apply it to prices of futures contracts on the German DAX index, constructed as in Chapter 3.2 with the price at 2010-06-16 set to 100, see Figure 4.1




Figure 4.1: Futures contract prices on the German DAX index from 2010-06-16 to 2010-11-03, the blue line is the original time series and the red line is the trend estimated by the Hodrick-Prescott filter.

As we can see, the trend component from the Hodrick-Prescott filter, the red line, follows the original time series, the blue line. But when the original time series makes some quick changes, the trend component has a damped response. This comes from the fact that the second term in (2.2) tries to counteract rapid changes.

4.2 Analysis of the trend component x

One thing that might be good to know about the trend component, x in (2.1), is how each data point changes when more data points are added to the original time series. To test this the filter is applied to same time series as above, the prices of future contracts on the German DAX index. Then we analyze the data of three time points, the time points are 2010-08-06, 2010-08-27 and 2010-09-17, highlighted in Figure 4.2




Figure 4.2: Futures contract prices on the German DAX index from 2010-06-16 to 2010-11-03, the blue line is the original time series, the red line is the trend estimated by the Hodrick-Prescott filter and the three green circles are the time points that are analyzed.

To see how each data point changes when new date points are added we start with only data until the selected time point and apply the filter to them, the last datapoint is saved and then the next data point is added on the time series. Then we apply the filter again and save the value at the same time point as we saved before, then we do this for all the following data points that we have in the time series. The three saved time series can be seen in Figure 4.3




Figure 4.3: Price changes of future contracts on the German DAX index when more time points are added. The first plot is the price of the contract at 2010-08-06, the middle is at 2010-08-27 and the last is at 2010-09-17

As we can see in Figure 4.3 there is a big change in the price when the first data point is added, this is caused by the second part of (2.2) tries to minimize the change in the second derivative over the middle of three data points, this means that it needs at least one extra data point to affect a singular data point. After the first correction, the time series starts to stabilize.

From the fact that there is a big change after adding one extra data point, we got the idea of making a new kind of time series where we just have end points. With end points we mean that we only save the last data point after each filtration, this is done at the same time series as earlier, the prices of the German DAX index, this new time series can be seen in Figure 4.4




Figure 4.4: Futures contract prices on the German DAX index from 2010-06-16 to 2010-11-03, the blue line is the original time series and the red line is the last data point in each estimation by the Hodrick-Prescott filter.

From Figure 4.4 can we see that the last filtrated data point is not that good if we want a consistent trend, this means that the Hodrick-Prescott filter does not generate a good trend estimation of the last time point.

4.3 Analysis of the Cyclical component u

The filter returns more then just the trend component. Another thing is the cyclical component u , which is the difference between the original time series and the trend returned by the filter. What we know about this cyclical component, from (2.7), is that it is supposed to be normally distributed. This is one of the things we look for in the cyclical component, it can be done by plotting a histogram of the component. Another thing we know is that if there does not exist a seasonal part in the original time series then there should be no autocorrelations within the cyclical component, in other words any cyclical component should be just noise.

We have, as in previous section, applied the filter to the prices of futures contracts on the German DAX index, we used data from 2008-03-31 to 2010-11-03 and the price at 2010-06-16 has been set to 100, see the results in Figure 4.5

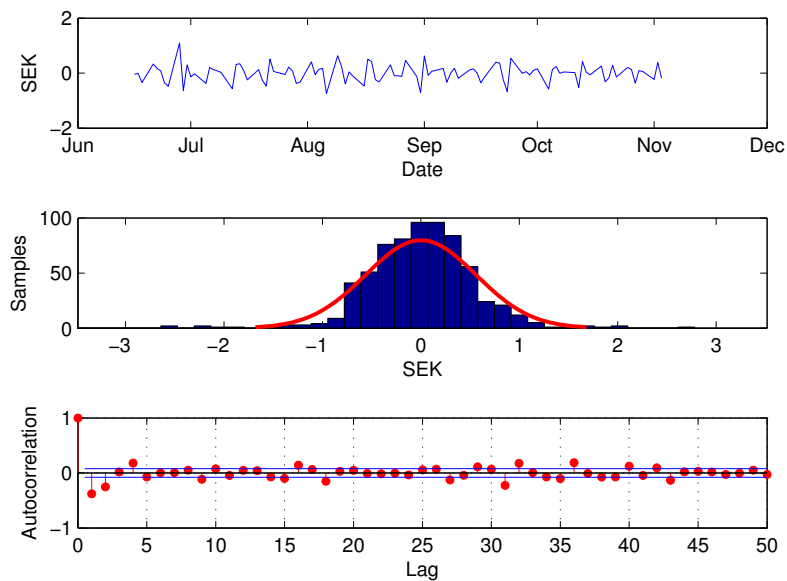


Figure 4.5: Futures contract prices on the German DAX index from 2008-03-31 to 2010-11-03 with 2010-06-16 set at 100, in the first plot the corresponding cyclical component to the trend in Figure 4.1 has been plotted. In the middle the distribution of the cyclical component over the whole interval has been plotted together with a fitted normal distribution and in the last plot the corresponded autocorrelation is plotted.

As we can see in the middle graph the distribution of the cyclical component have a similar appearance as the fitted normal distribution. The last plot shows that there is almost no autocorrelation in the cyclical component, as we would have guessed, since futures contracts tend to have no seasonality, because if it did it would create an arbitrage that the market would trade away.

4.4 Analysis of the noise-to-signal ratio α

The noise-to-signal ratio α is the parameter that decides how much weight the filter should put on the smoothing part according to the minimizing of the quadratic error. The only restriction we have for α is that it is supposed to be a positive real number, see the third section of Chapter 2, this is easily checked. But one thing that might be good to investigate is what happens with α when the input time series is scaled.

We apply the filter to the prices of futures contracts in the German DAX index, from 2008-03-31 to 2010-11-03. With the difference that the time series is scaled with five different factors, [-100 000, -100, 1, 100, 100 000].

The results is that in all five cases the α that the filter use is 0.4048. This means that we can scale the time series without changing the outcome of the filter. It is this fact that we use when we construct price series as in Chapter 3.2.

Computing α for different time series

We will in this chapter test what kind of α to use for which time series, four sort of different time series will be tested; Macroeconomic time series, Seasonally adjusted Macroeconomic time series, ordinary Price series and also time series with lots of Inflection points. We will test the three types of estimators that were described in Chapter 2.2. The way we will analyze the different α estimators is to look at the trend component x and then analyze the cyclical u since we know what to expect from this component.

5.1 Macroeconomic time series

The first type of time series is the Macroeconomic time series, as described in Chapter 3.1 this type of time series has a quite obvious seasonal component. The time series that we use is the Japanese Current Account with monthly data from July 2002 until October 2010.




Figure 5.1: The Japanese Current Account from July 2002 until October 2010. The black line is the original time series, the green line is the trend component using the DDR noise-to-signal ratio, the blue line is the trend component using the normal fitted noise-to-signal ratio and the red line is the trend component using the autocorrelation fitted noise-to-signal ratio.

As can be seen in Figure 5.1 there is almost no difference between the normal fitted time series, the green line, and the correlation fitted time series, the blue line, but there is quite a big difference to the time series that used the DDR estimator, the green line. The trend that the DDR estimator generates is the best, since both the normal and autocorrelation fitted time series still has some cyclical components.

DDR	15.348
Normal	1.9125
Autocorrelation	1.4886

Table 5.1: The α s used in the calculations of the time series in Figure 5.1.

Table 5.1 confirms what we saw in Figure 5.1, that there is a small differences between the two fitted time series since both have similar α s. Something more that we will check is the behavioral of the different cyclical components, this will be done as described in Chapter 4.3.




Figure 5.2: Analysis of the cyclical component to the Japanese Current Account using the DDR noise-to-signal ratio, in the first plot the time series are plotted. In the middle the distribution of the cyclical component is plotted together with a fitted normal distribution and in the last plot the autocorrelation of the cyclical component is plotted.




Figure 5.3: Analysis of the cyclical component to the Japanese Current Account using the normal fitted noise-to-signal ratio, in the first plot the time series are plotted. In the middle the distribution of the cyclical component is plotted together with a fitted normal distribution and in the last plot the autocorrelation of the cyclical component is plotted.




Figure 5.4: Analysis of the cyclical component to the Japanese Current Account using the autocorrelation fitted noise-to-signal ratio, in the first plot the time series are plotted. In the middle the distribution of the cyclical component is plotted together with a fitted normal distribution and in the last plot the autocorrelation of the cyclical component is plotted.

If we start with the first plots in Figures 5.2, 5.3 and 5.4, we see that all of them have a similar shape but that the cyclical component corresponding to the DDR estimator has almost twice the amplitude. This comes from the fact that the noise-to-signal ratio is bigger which will give a smoother time series.

In the two last plots in each Figure we can easily see the difference between the plots in Figure 5.2 and the plots in the two other Figures, 5.3 and 5.4, these differences comes from the fact that we optimize the results that corresponds to what can be seen in both of the plots. But even here it is hard to see the differences between the two optimizations because both of them returns similar α .

It is not easy to determine which of the estimators that is superior. If we just want a good trend the DDR is the best estimator, but we still have the restriction that the cyclical component shall be normal distributed and in that case both the other two estimators are superior.

5.2 Seasonally adjusted macroeconomic time series

As mentioned earlier there are two types of macroeconomic time series, the original in the section above and the one that comes from the original but where the seasonally dependent component has been removed. We will use the same time series as in Section 5.1, the Japanese Current Account with monthly data from July 2002 until October 2010 but with the difference that this time it is the seasonally adjusted time series.




Figure 5.5: The seasonally adjusted Japanese Current Account from July 2002 until October 2010. The black line is the original time series, the green line is the trend component using the DDR noise-to-signal ratio, the blue line is the trend component using the normal fitted noise-to-signal ratio and the red line is the trend component using the autocorrelation fitted noise-to-signal ratio.

5.2. Seasonally adjusted macroeconomic time series

From Figure 5.5 we can see that the seasonally adjusted time series, the black line, is still noisy, it jumps back and forth all the time. There is a larger difference between the three calculated time series this time, where we can see that the autocorrelation fitted time series, the red line, will have the largest α and that the normal fitted time series, the blue line, will have the smallest α . But the time series that have been calculated with the DDR estimation still has the best shape, not too much nor to little information from the original time series.

DDR	10.483
Normal	1.8672
Autocorrelation	202.13

Table 5.2: The α s used in the calculations of the time series in Figure 5.5.

As we can see in Table 5.2 the difference between the three α s is quite large, both the normal fitted and DDR estimated α s are close to the ones estimated in section 5.1 while the autocorrelation fitted shows major changes. As before, we will analyze all three of the cyclical components as well.

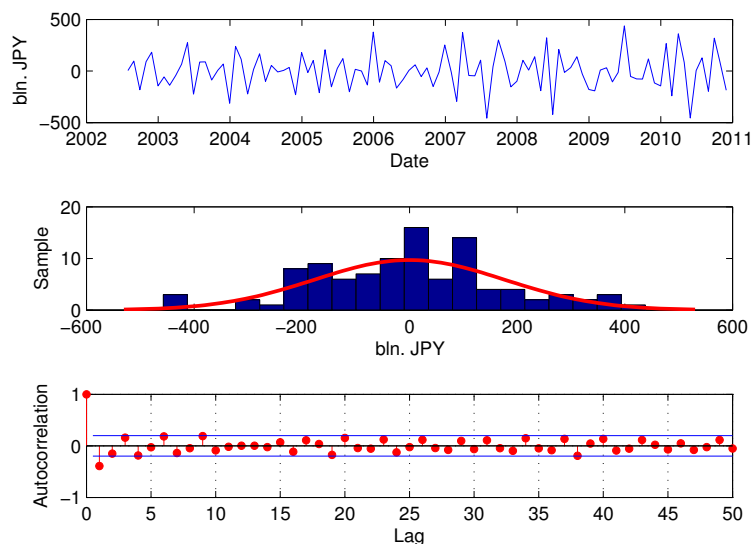


Figure 5.6: Analysis of the cyclical component to the seasonally adjusted Japanese Current Account using the DDR noise-to-signal ratio, in the first plot the time series are plotted. In the middle the distribution of the cyclical component is plotted together with a fitted normal distribution and in the last plot the autocorrelation of the cyclical component is plotted.




Figure 5.7: Analysis of the cyclical component to the seasonally adjusted Japanese Current Account using the normal fitted noise-to-signal ratio, in the first plot the time series are plotted. In the middle the distribution of the cyclical component is plotted together with a fitted normal distribution and in the last plot the autocorrelation of the cyclical component is plotted.




Figure 5.8: Analysis of the cyclical component to the seasonally adjusted Japanese Current Account using the autocorrelation fitted noise-to-signal ratio, in the first plot the time series are plotted. In the middle the distribution of the cyclical component is plotted together with a fitted normal distribution and in the last plot the autocorrelation of the cyclical component is plotted.

Starting from the first plots in Figures 5.6, 5.7 and 5.8, we see that they all have a similar shape, something that looks like noise. In the second plot the DDR estimated α generates a cyclical component that does not seem to be normally distributed, but the other two cyclical components have distributions that are more or less normally distributed. The last plot in the three figures shows the autocorrelation in the cyclical component and since the time series should be seasonally adjusted there should not be any signs of autocorrelation within the component. This is also the case for all three plots, although the last plot, which minimizes the autocorrelation, has a smaller autocorrelation with lag 1 than the others.

In this case the DDR estimation generates the best looking trend component, but at the same time the cyclical component is the worst of all three. The only thing that is clear is that the autocorrelation fitted α generates a trend that lacks a lot of information in each time point. It is hard to decide which of the estimators to use, but since the DDR estimator generates the best looking trend component, it would be our preferable choice.

5.3 Price series

A common type of time series is the daily price series, it use to be closing prices of any type of asset. The price series that we will use is the constructed, as in Chapter 3.2, Swedish stock market index OMX 30 with daily closing prices, from 2010-02-26 until 2010-11-03.




Figure 5.9: The Swedish OMX30 index, constructed as in Chapter 3.2 with 2010-02-25 set to 100, from 2010-02-26 until 2010-11-03. The black line is the original time series, the green line is the trend component using the DDR noise-to-signal ratio, the blue line is the trend component using the normal fitted noise-to-signal ratio and the red line is the trend component using the autocorrelation fitted noise-to-signal ratio.

As seen in Figure 5.9 there seems to be just two lines but the fact is that behind the red line lies both the green and blue lines. One thing that the plot shows is that all three trends follow the original time series at almost all time points but skips some rapid changes.

DDR	0.6483
Normal	0.7139
Autocorrelation	0.8070

Table 5.3: The α s used in the calculations of the time series in Figure 5.9.

From Table 5.3 we understand why only the red line is visible in Figure 5.9. This is because all of the three α s are almost identical. But we still want to see if there exists any differences within the cyclical components, so as in previous sections we will make the same analyzing plots.




Figure 5.10: Analysis of the cyclical component to the constructed Swedish OMX30 index using the DDR noise-to-signal ratio, in the first plot the time series are plotted. In the middle the distribution of the cyclical component is plotted together with a fitted normal distribution and in the last plot the autocorrelation of the cyclical component is plotted.




Figure 5.11: Analysis of the cyclical component to the constructed Swedish OMX30 index using the normal fitted noise-to-signal ratio, in the first plot the time series are plotted. In the middle the distribution of the cyclical component is plotted together with a fitted normal distribution and in the last plot the autocorrelation of the cyclical component is plotted.




Figure 5.12: Analysis of the cyclical component to the constructed Swedish OMX30 index using the autocorrelation fitted noise-to-signal ratio, in the first plot the time series are plotted. In the middle the distribution of the cyclical component is plotted together with a fitted normal distribution and in the last plot the autocorrelation of the cyclical component is plotted.

As we would have guessed, after seeing Table 5.3, there is almost no differences between the three cyclical components as can be seen in Figures 5.10, 5.11 and 5.12. According to this the choice of estimator does not matter, but to use either the normal or autocorrelated fitted estimator the DDR estimations must first be calculated, this makes the choice of the DDR estimator quite obvious.

5.4 Inflection time series

The last type of time series that we are going to test is the one that has a lot of inflection points, we chose this one to show one cases where the DDR estimator, using the second order shift operator D^2 , does not work. The time series we will use is the Japanese Gross Domestic Product with quarterly data from the second quarter of 1980 until the second quarter of 2010.




Figure 5.13: The Japanese Gross Domestic Product from the second quarter of 1980 until the second quarter of 2010. The black line is the original time series, the green line is the trend component using the DDR noise-to-signal ratio corresponding to D^3 and the red line is the trend component using the DDR noise -to-signal ratio corresponding to D^2 .

As mentioned earlier the D^2 estimator does not handle time series with a lot of inflection points very well, as can be seen in Figure 5.13, the D^2 estimator generates something that looks like a sine curve, but with another period than the original time series. When this happens we use the third order shift operator, corresponding to D^3 . From Table 5.4 we see that the D^2 estimator generates a negative α , which we know is not allowed, that is why the D^2 estimated trend series looks as it does.

D^2	-1.4722
D^3	0.3896

Table 5.4: The α :s used in the calculations of the time series in Figure 5.13.

As seen in Figure 5.13 the D^3 estimator copes quite well with the inflection points. That is why we use that one, instead of the original D^2 estimator, when we calculate the two other α :s.




Figure 5.14: The Japanese Gross Domestic Product from the second quarter of 1980 until the second quarter of 2010. The black line is the original time series, the green line is the trend component using the DDR noise-to-signal ratio corresponding to D^3 , the blue line is the trend component using the normal fitted noise-to-signal ratio and the red line is the trend component using the autocorrelation fitted noise-to-signal ratio.

In Figure 5.14 the results from the Japanese Gross Domestic Product is shown. The thing that we first notice is that none of the trend series are really good, both the DDR estimator and the normal fitted estimator retains to much of the original time series, although the DDR estimator works better, and the autocorrelated fitted estimator generates a trend series that has too little information from the original time series.

DDR	0.3806
Normal	0.0858
Autocorrelation	168477

Table 5.5: The α :s used in the calculations of the time series in Figure 5.14.

In Table 5.5 we see that the normal fitted α is almost zero which means that the corresponding trend is almost the original time series, the autocorrelation fitted α is enormous and generates a trend that is too generalized.




Figure 5.15: Analysis of the cyclical component to the Japanese Gross Domestic Product using the DDR noise-to-signal ratio, in the first plot the time series are plotted. In the middle the distribution of the cyclical component is plotted together with a fitted normal distribution and in the last plot the autocorrelation of the cyclical component is plotted.

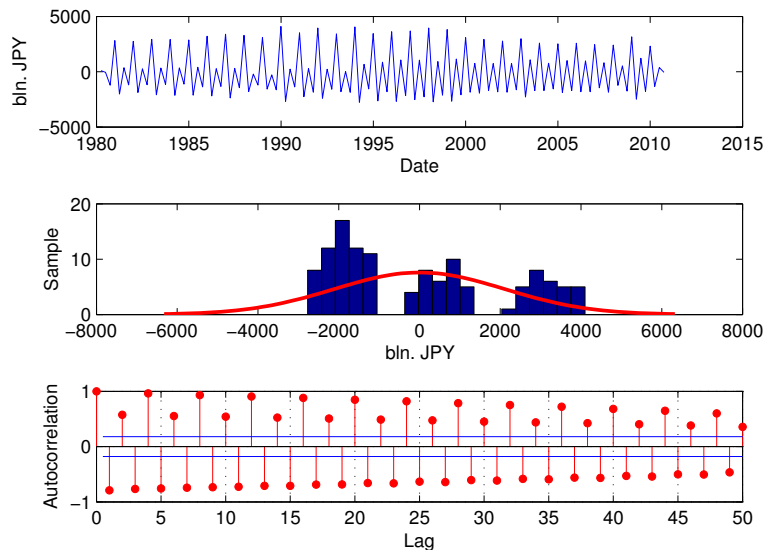


Figure 5.16: Analysis of the cyclical component to the Japanese Gross Domestic Product using the normal fitted noise-to-signal ratio, in the first plot the time series are plotted. In the middle the distribution of the cyclical component is plotted together with a fitted normal distribution and in the last plot the autocorrelation of the cyclical component is plotted.

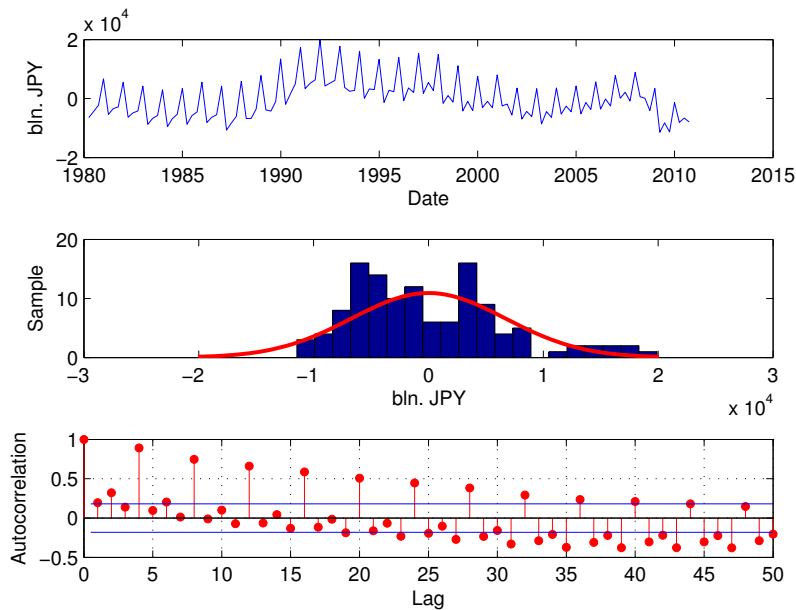


Figure 5.17: Analysis of the cyclical component to the Japanese Gross Domestic Product using the autocorrelation fitted noise-to-signal ratio, in the first plot the time series are plotted. In the middle the distribution of the cyclical component is plotted together with a fitted normal distribution and in the last plot the autocorrelation of the cyclical component is plotted.

All of the first plots in Figures 5.15, 5.16 and 5.17 shows a similar pattern. In the last figure, the one corresponding to the autocorrelation fitted estimation, there is also a trend in the component, not just fluctuations around zero, this is something we wanted to have in the trend series not the cyclical component

Since all of the cyclical components shows such a visible pattern, the distribution of the component can not be expected to be normally distributed as can be seen in the middle plot. It should also be a clear pattern in the autocorrelation, something that do not fit with the estimator that tries to minimize the autocorrelation.

If we have to decide what α to use in this case, we would go for the DDR estimator. Although it is not perfect and we think that the optimal α should be a bit larger, to get a perfect trend series.

5.5 Overall choice of α

As we have seen in the previous sections our choice of α estimator would be the DDR estimator, since it generates the trend component that does not lose to much information or is not smooth enough compared to the other trends. Even though it does not follow all the constraints on the cyclical component perfect.

As mentioned in the introduction, there exist other methods to estimate α . For example the standard way, the rule-of-thumb basis, where there exists some fixed values for different time intervals, this method was introduced together with the filter in Hodrick & Prescott (1997, [4]). Worth mentioning is also the maximum-likelihood estimation-based method introduced in Schlicht (2005, [10]) and the Generalized Cross-Validation method as in Weinert (2007, [11]). But, common for all these estimators is that they generate a too large α , that makes the trend too smooth, in the sense that the cyclical component is far from being Gaussian.

Risk calculations

This chapter covers all parts of how to calculate risk within a portfolio using filtered data, from the beginning when the time series is filtered through the risk calculations until the risk values are provided.

6.1 The portfolio

We calculate risk within a portfolio consisting of thirty assets, first thirteen future contracts on stock market indices; Australian S&P/ASX 200, Canadian S&P/TSE 60, French CAC 40, German DAX, Hong Kong based Hang Seng, Italian FTSE MIB, Japanese Topix, Dutch AEX, Spanish IBEX 35, Swedish OMX 30, Swiss SMI, British FTSE 100 and American S&P 500. It also consists of future contracts on 10 years treasuries in eight different countries; Australia, Canada, Germany, Japan, Sweden, Switzerland, UK and USA. Finally it has forward contracts in nine different currencies; Australian Dollar, Canadian Dollar, Hong Kong Dollar, Japanese Yen, Swedish Kronor, Swiss Francs, British Pound, American Dollar and New Zealand Dollar.

6.2 Constructing the filtered data

Since the portfolio consists of different types of contracts that each have expire dates, the portfolio has to change contracts from time to time. The way we cope with this is to use returns from each contract, by doing this it is possible to combine two different contracts to a new time series, see Chapter 3.2.

We start with time series of daily returns for all assets, we then construct a ordinary time series using

$$\begin{aligned}X(0) &= 1 \\X(t) &= X(t-1)r(t).\end{aligned}$$

Where $X(t)$ is the price series and $r(t)$ is the return series, look in Chapter 4.4 to see why it is possible to set the starting value to 1. After this we apply the Hodrick-Prescott filter to the time series and out of the returned trend we calculate the daily return.

6.3 Distribution of daily returns

What the filter does is that it evens out some of the rapid changes in the original time series, this makes the dispersion of the daily returns much lower. Two examples of this can be seen in Figure 6.1




Figure 6.1: Distributions of daily returns, the upper left plot is the original future contract on the German stock market index DAX and below it is the corresponding filtered series. The upper right plot is the original future contract on the British 10 years treasury and below it the corresponding filtered series.

Another way to see that the Hodrick-Prescott filter makes a dispersion of the distribution is to look at the standard deviation of the daily returns.

	$\sigma(x)$	$\sigma(HP(x))$	$\frac{\sigma(HP(x))}{\sigma(x)}$
S&P / ASX 200	0.0151	0.0086	0.5723
S&P / TSE 60	0.0192	0.0099	0.5135
CAC 40	0.0193	0.0102	0.5298
DAX	0.0187	0.0102	0.5462
Hang Seng	0.0218	0.0123	0.5660
FTSE MIB	0.0200	0.0115	0.5738
Topix	0.0211	0.0106	0.5003
AEX	0.0206	0.0112	0.5470
IBEX 35	0.0201	0.0112	0.5557
OMX 30	0.0203	0.0106	0.5237
SMI	0.0154	0.0084	0.5471
FTSE 100	0.0173	0.0092	0.5332
S&P 500	0.0197	0.0099	0.5019

Table 6.1: Standard deviation for the daily returns on future contracts on the stock indices with corresponding filtered time series and the ratio between the standard deviations.

From Table 6.1 we can see that all of the standard deviations of the future contracts on stock market indices are drastically decreasing when applying the Hodrick-Prescott filter. On average the filtered series has a standard deviation of 46% less than the original time series.

The same behaviour appears in the two other asset classes as well, this can be seen in Tables 6.2 and 6.3

	$\sigma(x)$	$\sigma(HP(x))$	$\frac{\sigma(HP(x))}{\sigma(x)}$
Australia	0.0058	0.0029	0.5069
Canada	0.0042	0.0023	0.5442
Germany	0.0040	0.0024	0.5979
Japan	0.0025	0.0014	0.5445
Sweden	0.0042	0.0026	0.6174
Switzerland	0.0034	0.0021	0.6283
UK	0.0049	0.0029	0.5957
USA	0.0052	0.0029	0.5543

Table 6.2: Standard deviation for the daily returns on future contracts on 10 years treasuries for different countries with corresponding filtered time series and the ratio between the standard deviations.

From Table 6.2 we get that the standard deviation within the future contracts on 10 years treasuries reduces on average with 43% after applying the filter.

	$\sigma(x)$	$\sigma(HP(x))$	$\frac{\sigma(HP(x))}{\sigma(x)}$
Australian Dollar	0.0103	0.0056	0.5488
Canadian Dollar	0.0079	0.0047	0.5911
Hong Kong Dollar	0.0080	0.0046	0.5711
Japanese Yen	0.0110	0.0064	0.5711
Swedish Kronor	0.0064	0.0037	0.5761
Swiss Francs	0.0050	0.0029	0.5816
British Pounds	0.0071	0.0042	0.5967
American Dollar	0.0081	0.0047	0.5760
New Zealand Dollar	0.0093	0.0055	0.5928

Table 6.3: Standard deviation for the daily returns on forward contracts on currencies with corresponding filtered time series and the ratio between the standard deviations.

In the forward contracts on currencies the standard deviation reduces on average with 42% after applying the filter, see Table 6.3, this gives an overall reduction for the standard deviation on average with 44% in the filtered time series.

6.4 Risk calculation

After the time series have been filtered the weighted covariance matrix is calculated. From this matrix both the expected volatility and Values-at-Risk can be calculated, how all calculations are made can be found in Chapter 2.3, from 2010-06-17 until 2010-11-03.

	Target	Original		HP Trend		$\frac{ETE(HP(x))}{ETE(x)}$
	ETE	ETE	VaR	ETE	VaR	
2010-06-17	13.33%	11.54%	\$2,881,753	6.78%	\$1,694,070	0.5879
2010-06-18	13.33%	11.30%	\$2,821,565	6.63%	\$1,656,964	0.5872
2010-06-21	13.33%	11.22%	\$2,801,652	6.18%	\$1,542,886	0.5507
2010-07-12	6.67%	5.85%	\$1,200,001	3.41%	\$698,753	0.5823
2010-07-13	6.67%	6.05%	\$1,240,389	3.42%	\$701,801	0.5658
2010-07-14	6.67%	5.88%	\$1,205,846	3.52%	\$721,938	0.5987
2010-10-04	12.00%	11.18%	\$2,304,727	6.59%	\$1,359,095	0.5897
2010-10-05	12.00%	11.05%	\$2,278,487	6.57%	\$1,354,127	0.5943
2010-10-06	12.00%	10.91%	\$2,249,421	6.73%	\$1,387,286	0.6167

Table 6.4: Risk values in the portfolio, both calculated with to original daily returns and with the filtered returns.

There is a noticeable difference between the risk values for the original data and the filtered data, as shown in Table 6.4. This difference can be seen in Figure 6.2 as well.




Figure 6.2: Chart over the ETE values in the portfolio, the black line is the target ETE. The blue line is the ETE level that the portfolio has with the original daily returns and the red line is the ETE levels with the filtered returns.

As we can see in Figure 6.2 there is a similar distance between the original ETE and the filtered ETE but it is not exactly the same over time. The distribution of the ratio between the two ETE levels can be seen in Figure 6.3.




Figure 6.3: Distribution of the ratio between the ETE level calculated with the filtered daily returns and the ETE level calculated with the original daily returns.

Table 6.3 confirms that there is a big spread between the ratios, reaching from 0.5020 up to 0.6772. This means that the trend risk does not follow the real risk perfect.

This variation in the ratio of the two risk figures may be useful to enhance the performance of the underlying optimal portfolio.

Bibliography

- [1] Araujo, F., Areaso, M. B. M. and Neto, J. A. R. (2003). r-filters: a Hodrick-Prescott Filter generalization, *Working Paper Series (69)*, Banco Central Do Brasil.
- [2] Dermoune, A., Djehiche, B. and Rahmania, N. (2008). A consistent estimator of the smoothing parameter in the Hodrick-Prescott filter, *J. Japan Statist. Soc.*, **38**(2), 225-241
- [3] Giurcanu, M. and Spokoiny, V. (2002). Confidence estimation of the covariance function of stationary and locally stationary processes, *Statistics & Decisions*, **22**(4), 283-300
- [4] Hodrick, R. and Prescott, E. C. (1997). Postwar U.S. business cycles: An empirical investigation, *Journal of Money, Credit and Banking*, **29**(1), 1-16
- [5] Leser, C. E. V. (1961). A simple method of trend construction, *J. Royal Stat. Society, Series B*, **23**(1), 91-107.
- [6] Macaulay, F. R. (1931), The Smoothing of Time Series, *New York: National Bureau of Economic Research*
- [7] Paige, R. L. and Trindade, A. A. (2010). The Hodrick-Prescott Filter: A special case of penalized spline smoothing, *Electronic J. of Statistics*, **4**, 856-874
- [8] Pedersen, T. M. (2001). The Hodrick-Prescott filter, the Slutsky effect, and the distortionary effect of filters, *J. Econom. Dynamics & Control.*, **25**, 1081-1101
- [9] Reeves, J. J., Blyth, C. M. and Small, J. P. (2000). The Hodrick-Prescott Filter, a generalization and a new procedure for extracting an empirical cycle from a series, *Studies in Nonlinear Dynamics and Econometrics*, **4**(1), 1-16
- [10] Schlicht, E. (2005). Estimating the smoothing parameter in the so-called Hodrick-Prescott filter, *J. Japan Statist. Soc.*, **35**(1), 99-119
- [11] Weinert, H. L. (2007). Efficient computation for Whittaker-Henderson smoothing, *Computational Statistics & Data Analysis*, **52**, 995-974
- [12] White, H. and Granger, C. W. J. (2011). Consideration of Trend in Time Series, *J. of Time Series Econometrics*, **3**(1), 8-46
- [13] Whittaker, E. T. (1923). On a new method of graduation, *Proc. Edinburgh Math. Soc.*, **41**, 63-75.