THE BUDGET BALANCE IN THE MEDIUM TERM AND THE HODRICK-PRESCOTT FILTER

Gerrit Van Den Dool

1. Introduction

The Pact for Stability and Growth obliges the EMU countries to strive for a budgetary position in the medium term that is 'close to balance or in surplus'. Such a position allows them to cope with cyclical downturns and to keep the budget deficit below the 3% ceiling without having to resort to special measures. The European Commission therefore has proposed that Member States should concentrate on the budget balance 'corrected for the influence of cyclical fluctuations in economic activity'. This proposal was agreed with by the Monetary Committee and the Ecofin in the opinion of 24 September: 'It is therefore clear that the assesment of the appropriateness of Member States' medium-term objectives and the examination of their fulfilment has to take explicit account of the cyclical position and its effect on the budget. The time frame for interpreting the medium term would be the length of the business cycle[1]'. Apart from a cyclical safety margin, leeway is needed to cope with 'other sources of variability and uncertainty in the budgets, the need to ensure a rapid decline in high debt ratios and the the need to cater for the costs associated to population ageing'.

Given the importance which the Pact attaches to the budgetary position in the medium term, paragraph 2 goes into this subject in greater detail. Structural budget balances are usually calculated by adjusting the

---

[1] Opinion on the content and format of stability and convergence programmes; MC/II/482/98.

actual balance for cyclical influences and not for any other temporary factors. This means that the structural balance may not be as structural as the word suggests. *For analytical purposes* one might therefore need to go further and eliminate other components (due to temporary or one-off measures) as well. In many calculations a major role is played by the Hodrick-Prescott filter. Therefore paragraph 3 discusses the end point problem, which is inherent to this filter. Paragraph 4 gives some examples of the peculiarities of the HP-filter. Finally, some tentative conclusions are drawn in paragraph 5.

2. The budgetary position in the medium term

Budgetary balances are calculated differently by different international institutions[2]. The IMF, for example, adjusts only that part of public revenues for cyclical conditions that represents the same ratio to GDP as in a cyclically-neutral base year in the past. However, as regards expenditure on unemployment, every change (expressed as a percentage of GDP) since the base year is eliminated, so that all changes in subsequent years are considered changes of a temporary nature. This is only one cause of differences vis-à-vis figures calculated by the European Commission and the OECD. The differences also depend on the way the output gap is calculated, as well as on the reliability of elasticities, such as those derived by the OECD for a large number of countries from the tax rate structure prevailing in the years 1978-1991, in combination with the income distribution in 1993. Elasticities are sometimes estimated by the Member States themselves by way of simulations with macro-economic models. In this context, the question arises to what extent the elasticities found also reflect anti-cyclical policy

---

[2] *European Commission*: Technical note: The Commission services' method for the cyclical adjustment of government budget balances, Broad Guidelines 1995. *OECD*: Giorno, Potential output, output gaps and structural budget balances, OECD Economic Studies, no. 24, 1995/I.

IMF: Heller, A review of the fiscal impulse measure, IMF Occasional paper 44, 1986.

reactions and whether the elasticities may be assumed to be constant across the entire economic cycle.

In practice, the budget balance is influenced not just by cyclical factors but also by other temporary factors and by continuous changes of an institutional nature. These institutional changes too may reinforce or compensate each other and are hard to forecast. After all, successive governments each put their own stamp on public spending, which not seldom includes undoing or adjusting measures taken by their predecessors. It also happens that a government alternates tax increases with tax cuts within the course of its term of office. That is why a 'snap shot' of the budget balance such as that recorded at year's end is not always a proper indicator of the deficit in the medium term, not even when adjusted for cyclical influences.

In a broader analysis of public finances the possibly temporary nature of many influences may therefore be a good reason to adjust the actual budget position not just for cyclical conditions, but for other fluctuations as well, irrespective of whether these ensue from temporary policy changes or from measures which are *meant* to be "permanent". Actually, it is often only afterwards that they can be described as permanent or turn out to be nothing of the sort at all.

Attention to this fact has been drawn before by the Österreichische Nationalbank[3] and the European Monetary Institute[4]. They proposed to calculate trends in public finances as moving weighted averages over a medium term and to make use of the Hodrick-Prescott filter, which indeed is well suited to such calculations. The only preconditon is that additional data are available, to extrapolate a series and to make possible the calculation of centred HP averages for the last year of the series as well. In order to make additional forecasts for budgetary positions, one can make use of the stability and convergence programmes which the EU Member States are required to draw up regularly, or of other official

---

[3] Peter Brandner, Structural budget deficits and sustainability of fiscal positions in the European Union, February 10, 1998 (meeting of Fiscal Experts Group, May 1998).

[4] L. Cano and A.Z. Kanutin, Estimation of structural deficits in EU Countries, EMI, MESD, Background studies for the Public Finance Report, March 1996.

sources. By their very nature forecasts are attended by uncertainty, but there is no escape from it, if one desires to assess a country's current budgetary position in the perspective of past as well as future developments.

This method of calculating a trend in the budgetary position does not in any way preclude the determination of cyclical components in that position. The ÖNB and the EMI dit not elaborate on it, but cyclical components may be derived in the usual ways, i.e. from output gaps through multiplication by elasticities for public revenues and expenditures or by a (net) elasticity for the balance. So for analytical purposes one may decompose the actual budget balance as follows:

Structural balance -	Derived from the actual balance with an HP filter
+Cyclical component -	Derived from the output gap with elasticities
+Other components -	Calculated as a residual item
Actual balance	

By contrast with the cyclically adjusted balance, as calculated by the international institutions (i.e. the actual position minus cyclical components), the structural balance in the above scheme does not contain any residual components. This makes it much less sensitive to distortions ensuing from all sorts of temporary factors, which in the alternative apporoach figure under "other components". This advantage goes hand in hand, however, with the disadvantage that changes which are to be said to be structural by politicians show up only gradually in the structural balance, because the latter is a moving average. However, a change of this sort does show up directly in the residual item, which can therefore fulfil a major role as early warning indicator, in the case of both temporary as well as structural changes.

A problem inherent to the analysis of the residual item is that its sign and size depend on the way the cyclical component of the budgetary

position is calculated. However, this does not mean that it cannot fulfil a signalling function; after all, it may be strongly positive or negative, no matter how it is calculated. Unfortunately, however, it is often impossible to tell what underlies the differences between the figures calculated by national bodies and those of international institutions. Identification of the causes is hindered by the complexity of the calculations or by insufficiently detailed explanations of the methodology applied and the sources used.

All this explains the need to fall back some times to a simple alternative. This need is met, for instance, by the simple but effective way of calculating the output gap with the aid of an HP filter. In the same way one could attempt to estimate elasticities by means of a simple regression analysis. Further refinement might be considered if its additional value can be demonstrated convincingly and in a transparent manner. As an example of the above mentioned approach some preliminary results are presented for the Netherlands in annex 1. At first sight there is no clear-cut correlation for the Netherlands between cyclical conditions, on the one hand, and the actual deficit or public revenues and expenditure, on the other. Broken down by category, however, a correlation seems to show up for direct taxation and expenditure on unemployment.

But a purely statistical approach such as this one has also its limitations. Charts like those in annex 1 do not show, for instance, that in practice the elasticity of Dutch spending has been limited for a number of years because ceilings have been imposed on real spending. That is a factor international institutions so far have not made any allowance for. From the charts it cannot be deduced either that the elasticity of revenues could be dampened over the next few years if the government, as it has announced, were to compensate for windfalls or setbacks on the revenue side of the budget with tax cuts or increases. This underscores once again the relative value of detailed refinements if no allowance is made for this type of specifically national circumstances.

3. The merits of the Hodrick Prescott filter

In the simple alternative approach described above, a major role is fulfilled by the HP filter, as is also the case in the more refined techniques used to calculate the output gap (take, for example, the

calculation of trends for the NAWRU and for technical progress made by the OECD to determine production capacity).

However, the merits of the HP filter depend largely on the way the end point problem is dealt with. The European Commission solves this problem through extrapolation with estimations and forecasts with the aid of ARIMA models. The need to do this might seem a disadvantage of the HP filter, but an important advantage of it is its simplicity, making it possible to assess accurately, with the aid of simple simulations, the sensitivity to forecast errors and the choice of parameters. This will be demonstrated below.

The following examples are set up for a fictitious economy whose GDP shows an eight-year business cycle (a sinusoid) around a trend which is, for the time being, kept constant at 100. The question is how reliable the calculation of the GDP trend with an HP filter is for the last year of a 25-year series (say 1974-1998), if GDP is extrapolated correctly for four extra years (1999-2002) (Figure 1). The fluctuations vary between plus and minus 3%.

Let us first take a look at the distribution of weights with which the trend is calculated for each of the years from 1974 up to and including 1998/2002.

Fig. 1

Hypothetical Business Cycle (trend – GDP = 100)

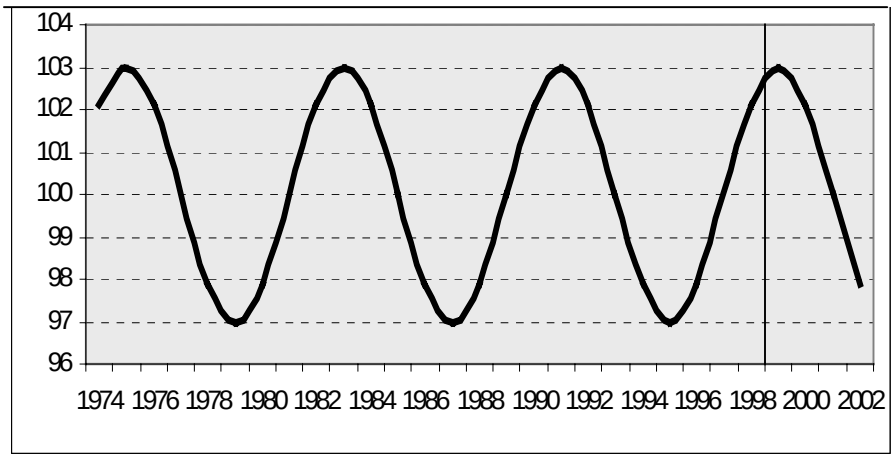

Fig. 2

HP-Weights Before and After Extrapolation (λ = 100)

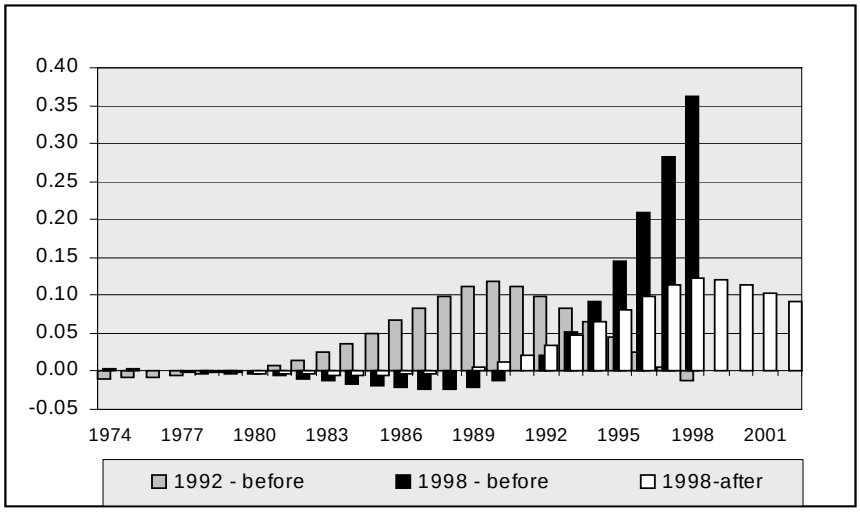

Fig. 3

HP-Weights 1998 at Different λ's

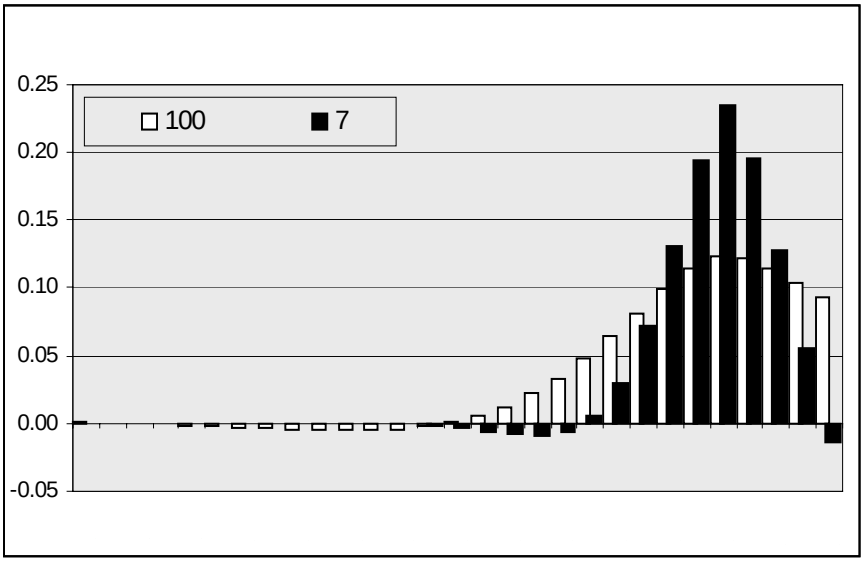

Figure 2 shows that the distribution of weights for the trend in 1990 is rather symmetrical, but becomes increasingly skewed towards the end point 1998.

In 1990, only around 12% of the GDP trend is determined by GDP in 1990 itself, while the influence from earlier and later years is about equal. In 1998, however, two-thirds of the trend is determined by GDP over the last two years. Consequently, depending on the phase of the business cycle, the calculated trend can turn out to be much lower or higher than the actual trend. However, Figure 3 shows that after extrapolation the trend for 1998 is based on a much less skewed distribution.

So the difference between the calculated and the actual trend is determined by the skewness of the distribution of weights and by the phase of the business cycle. But a role is also played by the width of the distribution, which depends on the parameter λ. After all, a symmetrical distribution can also be so narrow, by comparison with the length of the business cycle, that individual years are assigned an unduly large weight. Compare, in Figure 3, the distribution of weights for 1998 at a λ of 100, and a λ of 7 (which was used by the ΦNB in its above mentioned paper). Finally, the distribution also depends on the length of the period for which trends are calculated. For 1998, however, that factor is virtually insignificant, as the trend based on a figures for 1974-2002 hardly differs from the trend based on the years 1979-2002 or 1984-2002 (Figure 4).

Simulations with various λ's

For a GDP developing as supposed in Figure 1, the corresponding trend, calculated with the aid of the HP filter, is shown in Figure 5 for a λ of 100 and a λ of 7. Like the actual trend, the trend for a λ of 100 is fairly constant, except towards the 'end' point 1974. At a distance of about four years from the end points, the deviation vis-à-vis the actual trend is no larger than around ¼%.However, the trend calculated with the aid of ?=7 fluctuates markedly, the difference against the actual trend sometimes coming out at nearly 1% even in the middle of the 29-year period. A λ of 7 is thus unsuited to the calculation of trends in the present case.

Fig. 4

HP-Weights 1998 at Different Lengths of Series

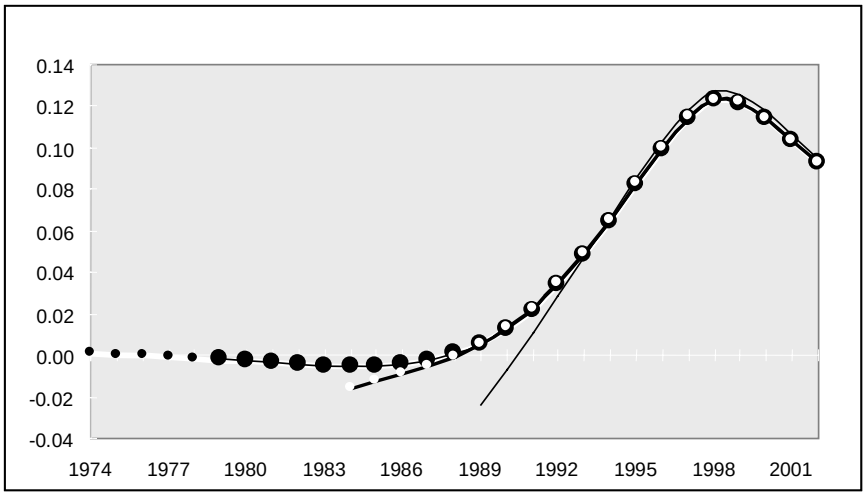

Fig. 5

Trend GDP at λ = 100 and λ = 7

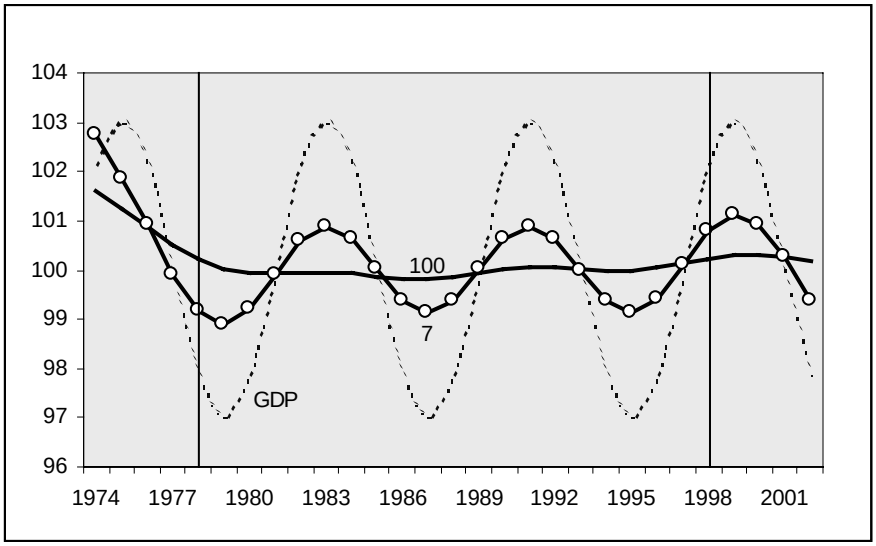

Shifts in phase

The fact that the trend errors, i.e. the differences between the actual and the calculated trend, are larger near one end point (1974) than towards the other (2002), has to do with the assumed development of GDP, whose beginning and end also differ in Figure 1. When the beginning in 1974 is placed in a different cyclical phase, i.e. if the GDP curve in Figure 1 is shifted one or more years, the trend errors and consequently the output gap errors also turn out differently, not just at the end points, but also in the intermediate years. Figure 6 shows that the error in the calculation of the output gap in 1998 is inversely correlated to the size of the output gap in that year.

It appears that the trend calculated with the HP filter for 1998 can deviate as much as 0.3% from the actual trend. Consequently, at an elasticity of about 0.6 cyclical components in the budget balance may also be distorted by nearly 0.2%.

Fig. 6

Difference Between Actual and HP-Output Gap in 1998

at Different Gaps/Phases of the Business Cycle in 1998

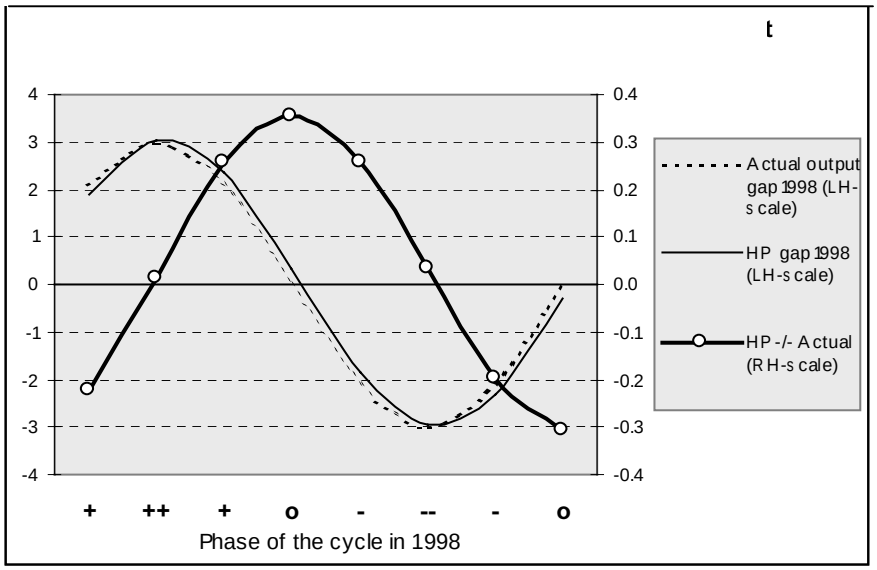

Fig. 7

Maximum Estimation Error of the 1998-Output Gap at Different Lenghts of the Business Cycle

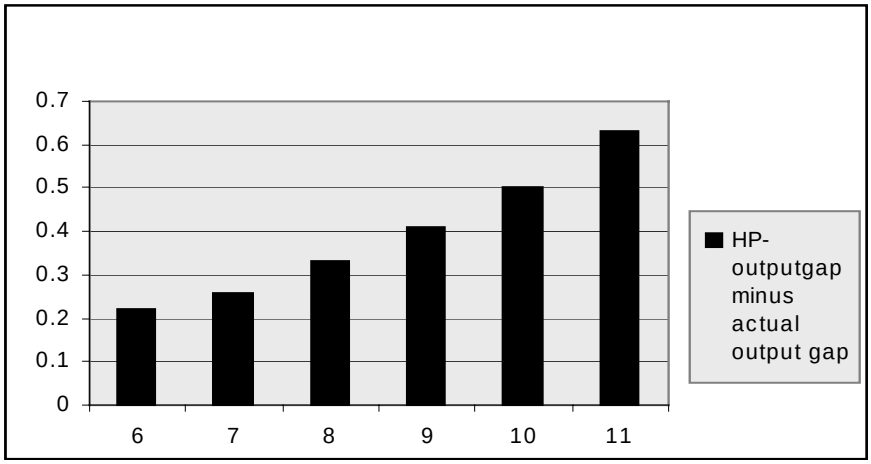

Variations of the length of the business cycle

Similar deviations as above could be found for 1998 by varying the length of the cycle in stead of the phase of the business cycle in 1974. For the sake of completeness, the maximum 1998 trend errors for this option are shown in Figure 7. The larger the length of the business cycle, the larger the possible trend errors. At a cycle of eleven years the trend error can even rise to 0.6%.

Trend rate of growth

Furthermore it is of vital importance for the calculation of the trend that allowance is made for any trend rate of economic growth. If the HP filter is applied to the levels of real GDP growing at an annual rate of, for instance, 3%, substantial differences vis-à-vis the actual trend may arise, not just at the end points, but also in 1998, as illustrated by Figure 8. However, by calculating the trend for the logarithm of GDP instead of its level, this type of trend errors can easily be prevented.

Fig. 8

HP Trend, Based on GDP Level Versus log GDP, Minus Actual Trend (1998)

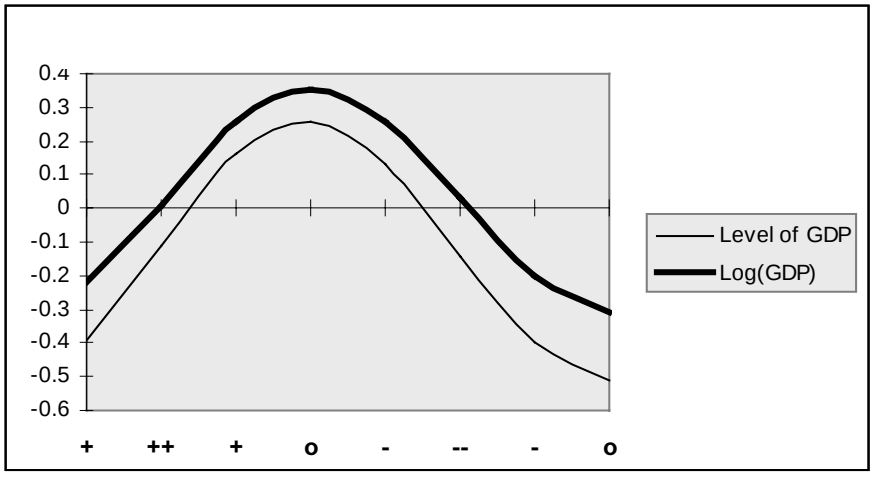

Fig. 9

Difference Between HP-Trend and Actual Trend in 1998 at Different Growth Forecasts (y*) for 1999-2002

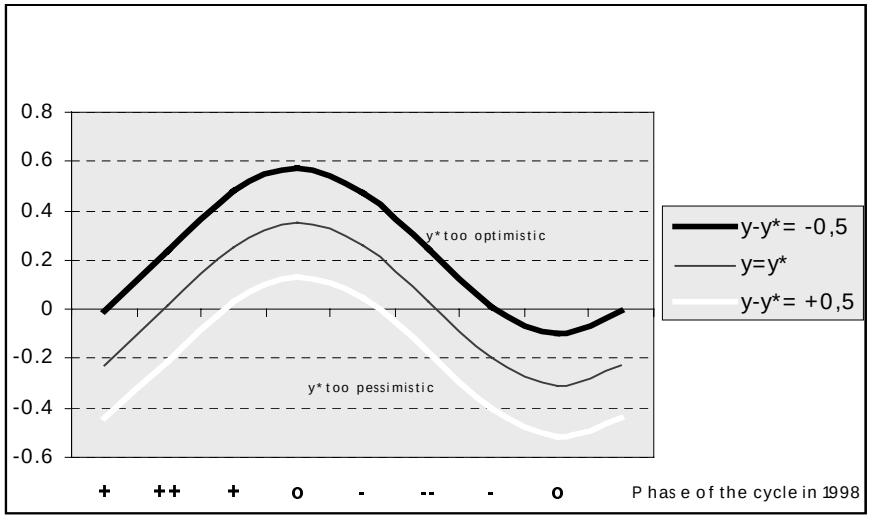

Forecasting errors

Finally, a major question is how forecasting errors for the years 1999-2002 can feed through to the calculation of the trend for 1998. This, too, clearly depends on the phase of the business cycle assumed, as well as on the size of λ, the amplitude, the length of the business cycle and so on. Figure 9 shows that extrapolation with reliable forecastst limits the trend error to about 0.3% of the trend (and one tenth of the amplitude of the business cycle). This error doubles, however, if the forecast of GDP growth is systematically overestimated or underestimated by 0.5%.

4. Some peculiarities of the HP filter in practice

As was made clear in the foregoing the trend calculated for 1998 depends largely on the forecast for the years 1999-2002 (and beyond). A very optimistic forecast, for instance, means that the trend in 1998 is lifted. This leads to a less positive or more negative output gap in 1998 which in turn, for a given budget balance, leads to a more favourable structural budget balance. Conversely, a pessimistic growth forecast may lead to a less favourable structural balance in 1998. The dependence of the structural balance in 1998 on forecasts of average GDP growth for the period 1999-2002 is illustrated for the Netherlands in Figure 10[5]. From the graph it can be concluded that *the* structural deficit in 1998 does not exist, not yet at least. It is only after a number of years (depending on λ and actual GDP growth in current and future years) that a definite assessment of the *level* of the structural deficit in 1998 (1997, ...) becomes possible. It is important to note that the same holds for calculations of the structural balance based on the production function approach, in which the HP filter also plays an important role.

The provisional character of the structural deficit might seem a serious problem in assessing structural developments. However, *changes*

---

[5] Based on data available in October 1998 and on the assumption that the government will partially compensate setbacks or windfalls on the revenu side if gdp growth differs from 2¼% (gdp growth after 2002 has been kept constant at 2¾%).

of the structural balance in individual years turn out to be much less sensitive to different growth assumptions, as is illustrated in Figure 10.

The HP filter implies also that even at a sufficiently large value of λ a temporarily high growth rate of GDP may lift the HP trend of GDP during some years while in fact no real, structural changes in the economy has taken place. After all, production capacity and productivity may have continued to grow at their usual rates. An example of interpretation problems which might arise can be found in the Commission's assessment of the Stability Program of France[6].

For the period 2000-2002 the French program distinguishes a cautious scenario with an annual GDP growth of 2.5% and a favourable scenario with a growth rate of 3%. Not surprisingly, in the latter scenario the budget balance improves more, even after half of the extra improvement has been used for a structural reduction of taxes (as has been announced by the French governement). What is surprising, however, is that the *improvement* of the budget balance turns out to be completely structural, in spite of the structural *reduction* of taxes and in spite of the absence of any explicitly known *structural* changes in the French economy. To understand this seeming anomaly one has to realize that in the favourable scenario the trend too is lifted. This causes the output gaps in both scenarios to differ only slightly. Consequently, the cyclical component of the government balance in 2002 is approximately the same in both scenarios, leaving the structural component to explain any differences. In the favourable scenario the government balance improves by 0,4% more (about half of the improvement of 0,75% which would be possible without tax reduction[7]).

Peculiarities like this one underline once more the need to examen figures on structural deficits carefully. It is important to keep in mind that a trend which seems to be higher from a *statistical* point of view does not imply an equally higher trend in the real world.

---

[6] Stability Program of France, An assessment, 11/02/99, II/112/99-EN; see table 4.

[7] 0,75 = 0,5 (the elasticity for France) x 1,5 (the cumulated GDP growth difference).

5. Conclusion

The HP filter has proven to be a very useful tool to determine trends in variables like GDP, the budget balance, the NAWRU, labor productivity etc. However, one has to be aware of the different kinds of uncertainty connected to the HP filter. In the first place parameters like λ and the sample period can be arbitrarily chosen. A second source of uncertainty lies in the use of forecasts, to solve the endpoint problem. Furthermore, a statistical trend calculated by a HP filter does not necessarily reflect underlying developments in the real world.

On the other hand, in the previous paragraphs a closer look at the HP filter made clear, that these uncertainties should not be exagerated. It was shown, by means of various graphs and simulations, that many of the uncertainties mentioned can easily be quantified and need not to be as large as one might suppose. In any case, the simplicity of the HP filter, permitting a quick and transparent analysis of different alternatives, is one of its main advantages one should weigh against the extra efforts of other methods, which may be more sophisticated but not necessarily more reliable.

ANNEX
Estimating elasticities for the Netherlands; some preliminary results of a straightforward approach [8] [9]

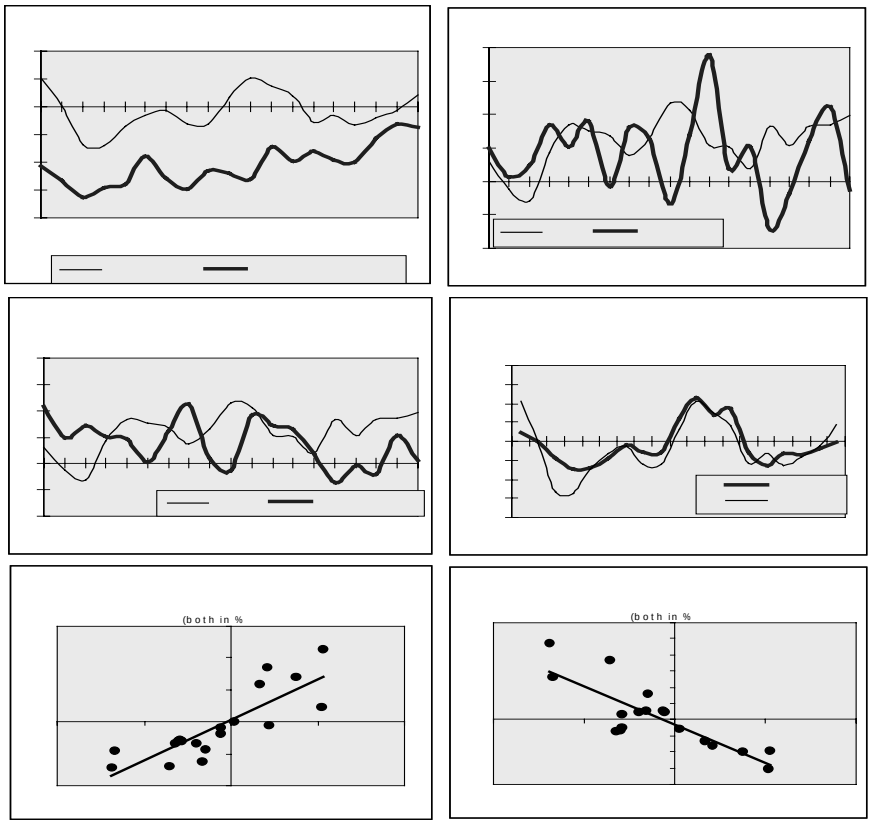

[8] Estimation of elasticities (slopes of regression lines) based on data and forecasts available in October 1998.

[9] It should be stressed that results of this approach always need a further examination. The elasticity of about **0,65** for direct taxes, for instance, which might be estimated from Chart 5, suggests a very high sensitivity of direct taxes to the business cycle (the share of direct taxes in GDP is only **13%**, so a one percent increase in GDP would have to correspond with a rise of **5%** in direct taxes). At first sight one might find this high sensitivity not very plausible. So the results of a simple regression analysis should also be checked by other means.