**Central European Journal of Economic Modelling and Econometrics**

# Real-Time vs. Full-Sample Performance of One-Sided and Two-Sided HP Filters. An Application to 27 EU Member States' GDP Data

Kaloyan Ganev*

Submitted: 28.06.2019, Accepted: 19.06.2020

**Abstract**

The paper makes a comparison of the results of the application of two-sided and one-sided versions of the Hodrick-Prescott filter on GDP data concerning 27 EU Member States. Based on the results, the overall finding is that, contrary to its assumed advantages, the one-sided filter does not help overcome endpoint unbiasedness. Quite the opposite, it rather spreads and consolidates the endpoint bias that plagues the two-sided version over the entire filtered data. In addition, regression-based results on the influence of the second, third, and fourth moments of the GDP acceleration rates on the differences between one-sided and two-sided HP trends are presented.

**Keywords:** One-sided and two-sided Hodrick-Prescott filters, endpoint bias

**JEL Classification**: C32, E37

*Sofia University "St. Kliment Ohridski"; e-mail: k_ganev@feb.uni-sofia.bg; ORCID: 0000-0002-2966-6303

# 1 Introduction

Trend-cycle decompositions are among the most widely used techniques of economic time series analysis. Their main goal is to separate the permanent and the transitory components of time series. In macroeconomic research, one of the main areas of application is business cycle analysis. There, usually the goal is to identify and extract the long-term (growth) component of aggregate output, respectively to isolate its short-term (cyclical) component. Each of those can then be analysed either separately, in their own specific context, or jointly, when a more complete description of output dynamics is necessary.

Since its first appearance in a working paper in 1980, the Hodrick-Prescott (HP) filter (Hodrick and Prescott, 1997[1980]) has become one of the standard tools used for the purpose of trend-cycle decompositions of real GDP. In business cycle analysis, it has often been used “on-the-fly” (i.e. without too much thinking) to detrend data series. The latter has been facilitated by the fact that the filter is also pre-programmed in mainstream data modelling software packages (e.g. Stata, EViews, Matlab, or R).

Despite its attractiveness and ease of usage, the HP filter has received a significant amount of criticism. For example, it has been shown that the filter can possibly generate spurious cycles (Harvey and Jaeger, 1993; Cogley and Nason, 1995), i.e. it is not always clear whether the presence of cyclical behaviour in the detrended non-stationary time series corresponds to the true cyclical patterns of data. The latter depends on whether or not the spectrum of the data is maximized at the business cycle frequencies, or, which is usually the case, at lower ones. Another obvious issue associated with the application of the filter is that it is a two-sided one. This implies a serious issue: when the sample endpoint data are being filtered, the information needed to perform the calculation is only partially available. Therefore, a bias tends to arise specifically at those data locations; such a bias is not present when “inner” datapoints are filtered. If the endpoints coincide with the available most recent (“current”) data, the consequence is that there will be biased real-time trend and cycle estimates.

Most of the drawbacks of the HP filter arise largely due to the fact that the ratio of the trend and cycle variances is not estimated from the data but manually imposed. If this ratio is estimated, the filter construct leads to values that are very different from the manually imposed ones for the corresponding data frequencies. The mentioned problems have even led to the recommendation by some authors to strictly refrain from using this trend-cycle decomposition technique (see for example Hamilton, 2018).

# 2 Filter specifications

The purpose of the HP filter is to extract a smooth trend from a non-stationary time series. Smoothness is defined as the square of the second difference (i.e. square of the rate of acceleration) of the trend component $\tau_t$, i.e.:

$$(\Delta^2\tau_t)^2 = (\tau_t - 2\tau_{t-1} + \tau_{t-2})^2. \tag{1}$$

In the standard two-sided version of the HP filter, the following problem is solved to achieve the filter purpose:

$$\min_{\tau_t}\left[\sum_{t=0}^{T}(y_t - \tau_t)^2 + \lambda\sum_{t=2}^{T-1}(\tau_{t+1} - 2\tau_t + \tau_{t-1})^2\right], \tag{2}$$

where $T$ is the time series length (sample size). The parameter $\lambda$ is a penalizing factor controlling for the smoothness of the trend. In fact, this parameters allows to manually specify the value of the ratio of the variances of the trend and the cycle.
For quarterly data, the recommended value for $\lambda$ is 1600, while the "standard" values of this parameter for annual and monthly frequencies are respectively 100 and 14400. The general rule that is followed in determining the value that corresponds to a selected frequency is

$$\lambda = \frac{1600}{n^2}, \tag{3}$$

where $n$ is the number of quarters present in a unit period of the frequency. For example, there are 4 quarters in a year, therefore $\lambda = \frac{1600}{4^2} = 100$, etc.
Ravn and Uhlig (2002) proposed using a higher power in the denominator (4 recommended), therefore the rule becomes

$$\lambda = \frac{1600}{n^4}. \tag{4}$$

The specification of the one-sided version of the Hodrick-Prescott filter appeared first in Stock and Watson (1999). It consists of the following two equations:

$$y_t = \tau_t + \varepsilon_t \tag{5}$$

$$(1-L)^2\tau_t = \eta_t \tag{6}$$

where $\tau_t$ once again denotes the trend component of $y_t$, and $\varepsilon_t$ is identified with the cycle (a residual).
Equation (6) can be alternatively written as

$$\tau_t = 2\tau_{t-1} - \tau_{t-2} + \eta_t. \tag{7}$$

Noting that the trend is an unobservable state variable, and that equation (5) serves the purpose of an observation (signal) equation, the model can be put in state-space form:

$$\begin{aligned} y_t &= \mathbf{F}_t\boldsymbol{\theta}_t + \boldsymbol{\nu}_t \\ \boldsymbol{\theta}_t &= \mathbf{G}_t\boldsymbol{\theta}_{t-1} + \boldsymbol{\omega}_t \end{aligned} \tag{8}$$

where $\mathbf{F}_t = \begin{pmatrix}1 & 0\end{pmatrix}$, $\boldsymbol{\theta}_t = \begin{pmatrix}\tau_t \\ x_t\end{pmatrix}$, $\mathbf{G}_t = \begin{pmatrix}2 & -1 \\ 1 & 0\end{pmatrix}$, $\nu_t = \varepsilon_t$, $\boldsymbol{\omega}_t = \begin{pmatrix}\eta_t \\ 0\end{pmatrix}$, and $x_t = \tau_{t-1}$.
The latter is then subject to estimation using the Kalman filter, and the filtered series is identified with the underlying unobservable trend component. In this approach to the filter specification and estimation, if the Kalman smoother is also applied (on the filtered series), it would produce the two-sided HP trend.
In the system of equations (8), the random disturbance terms $\varepsilon_t$ and $\eta_t$ (i.e. the innovations in the observation and the state equations, respectively) are assumed to be Gaussian, with variances $\sigma_\varepsilon^2$ and $\sigma_\eta^2$. When estimation is performed, the ratio of those variances, $q$ is fixed to a predefined value:

$$q = \frac{\sigma_\eta^2}{\sigma_\varepsilon^2}. \tag{9}$$

In fact, $q$ is the reciprocal of $\lambda$ as defined above. Stock and Watson (1999) use $q = 0.000675$ which corresponds to the spectral gain of the HP filter for quarterly data; this implies a value of $\lambda$ approximately equal to 1481.5. From the reverse perspective, if $\lambda$ is set to the "standard" value of 1600 for quarterly frequencies, this would imply $q = 0.000625$. Setting $q$ (or $\lambda$) manually is obviously a mechanistic approach which does not "let the data speak". The alternative would be to estimate $q$ from the data but practice shows that estimates vastly deviate from the standard values that are used commonly in applications.
A distinct advantage of using the state-space formulation and Kalman-filter estimation is that forecasts of both the filtered and the smoothed states can be produced, which is generally impossible in the "standard" implementations of the Hodrick-Prescott filter in software packages.

# 3 Data used in the current application

In the empirical analysis that follows, Eurostat seasonally and calendar adjusted data on 27 (out of 28) EU Member States data on quarterly GDP are used. All data used are at 2010 constant prices. Data on the Slovakian economy are missing at the time of writing, possibly due to technical reasons. It would be possible to seasonally adjust the not-seasonally adjusted series for Slovakia which is provided. However, the specifics of the seasonal adjustment method applied by Eurostat are not available, therefore a risk exists that the output data could not be in compliance with those on the remaining Member States. That the seasonal component of a time series is removed beforehand is a standard requirement for the application of trend-cycle decompositions such as the Hodrick-Prescott filter.
The earliest available data are on the French and the UK economies, starting in the first quarter of 1975. The shortest series are observed for Malta and Croatia, starting in the first quarter of 2000. For all Member States, the data samples end in the

fourth quarter of 2019. In general, this means that the sample size differs across Member States. However, as estimation is applied individually to each one of them, no truncation to the least common sample is necessary for the application of the HP filters.
In order to be able to directly interpret the first differences of data as growth rates, respectively the second differences as rates of acceleration, all data are natural-log-transformed.

# 4 Estimation and results

## 4.1 Comparison of one-sided and two-sided HP filter results

In the current application, the first task is to compare the real-time and the full-sample performance of each of the versions (one-sided and two-sided) of the Hodrick-Prescott filter. The real-time versions of the one-sided and the two-sided HP filter are implemented by using expanding data windows in estimation. More precisely, in the first step, estimation is performed on the smallest possible data window required by the computational procedure in the beginning of the available sample. In every next step, an observation is added, and estimation is carried out until the full sample is reached. From each step's estimation results, only the endpoint estimate is recorded in a newly generated series, called from this point onwards "the real-time HP trend estimate".
Since the frequency of the data used in estimation is quarterly, the power rule of Ravn and Uhlig is irrelevant ($n = 1$ so any power of it would also yield 1). Therefore, the usual values of 1600 for $\lambda$ and 0.000625 for $q$ are adopted. The software used for estimation is R. Specifically, for the two-sided HP filter, the mFilter package is utilized, while the one-sided version is programmed using the facilities built in the dlm package.
The results concerning the real-time estimates are visualized in Figure 1 where for each Member State the two HP trends are plotted alongside the logs of actual GDP data. The two extracted HP trends are denoted respectively by `HP1s_RT` and `HP2s_RT`; the suffix `_RT` signifies "real-time". The solid gray lines in the subplots correspond to actual GDP figures. The abbreviations of the names of Member States follow the Eurostat nomenclature. The sample that is used for plotting the results is 2001 Q2 - 2019 Q2. It is based on the least common sample available for all Member States (2000 Q4 - 2019 Q4) but it excludes two points at the beginning and two points at the end so that potentially biased results are not plotted. Choosing a common sample for the plots also avoids distorting the visualization due to different length of scales in different subplots.
The numerical comparisons show that the real-time estimates from the application of both the one-sided and the two-sided version of the HP filter are identical, therefore

Figure 1: Results from the application of the one-sided and two-sided HP filter in real time

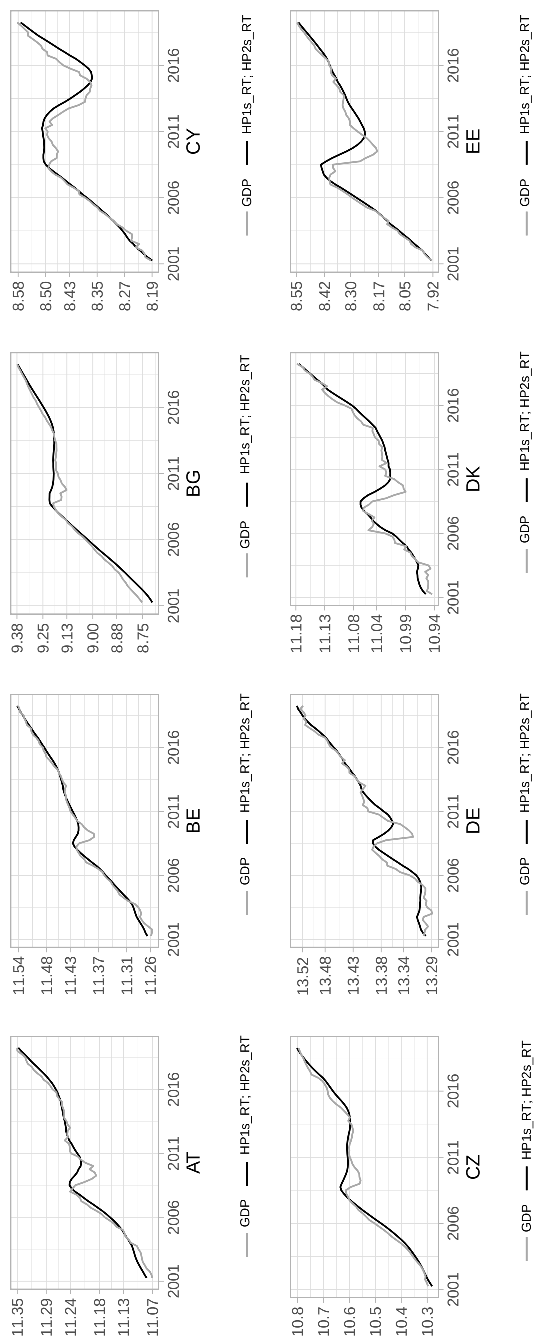

Figure 1 (cont.): Results from the application of the one-sided and two-sided HP filter in real time

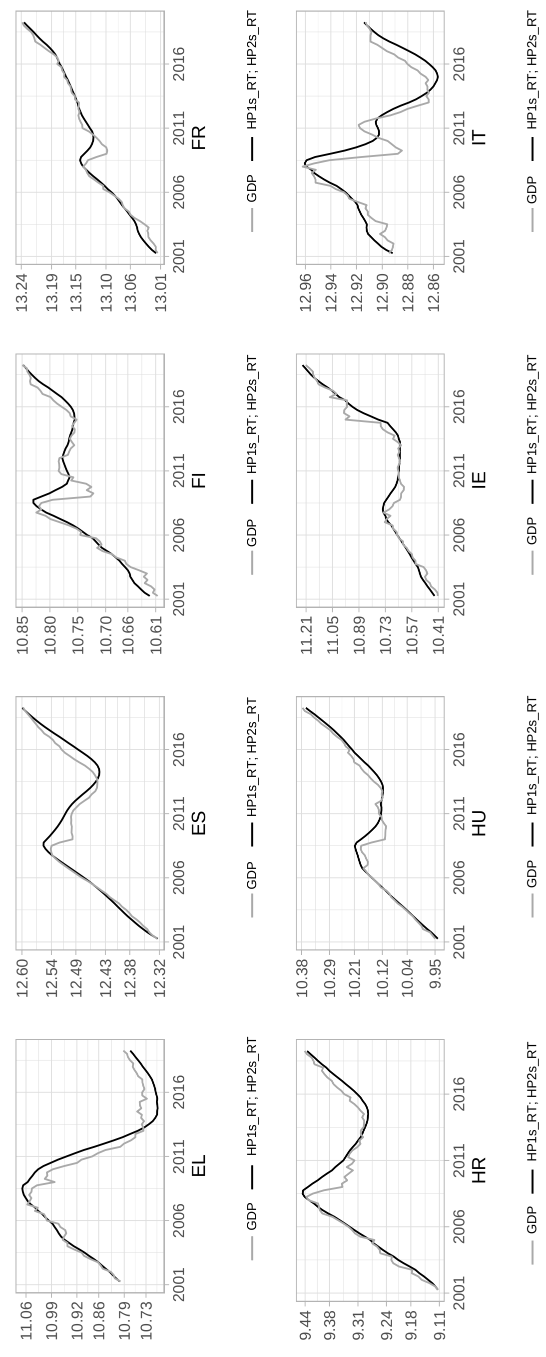

Figure 1 (cont.): Results from the application of the one-sided and two-sided HP filter in real time

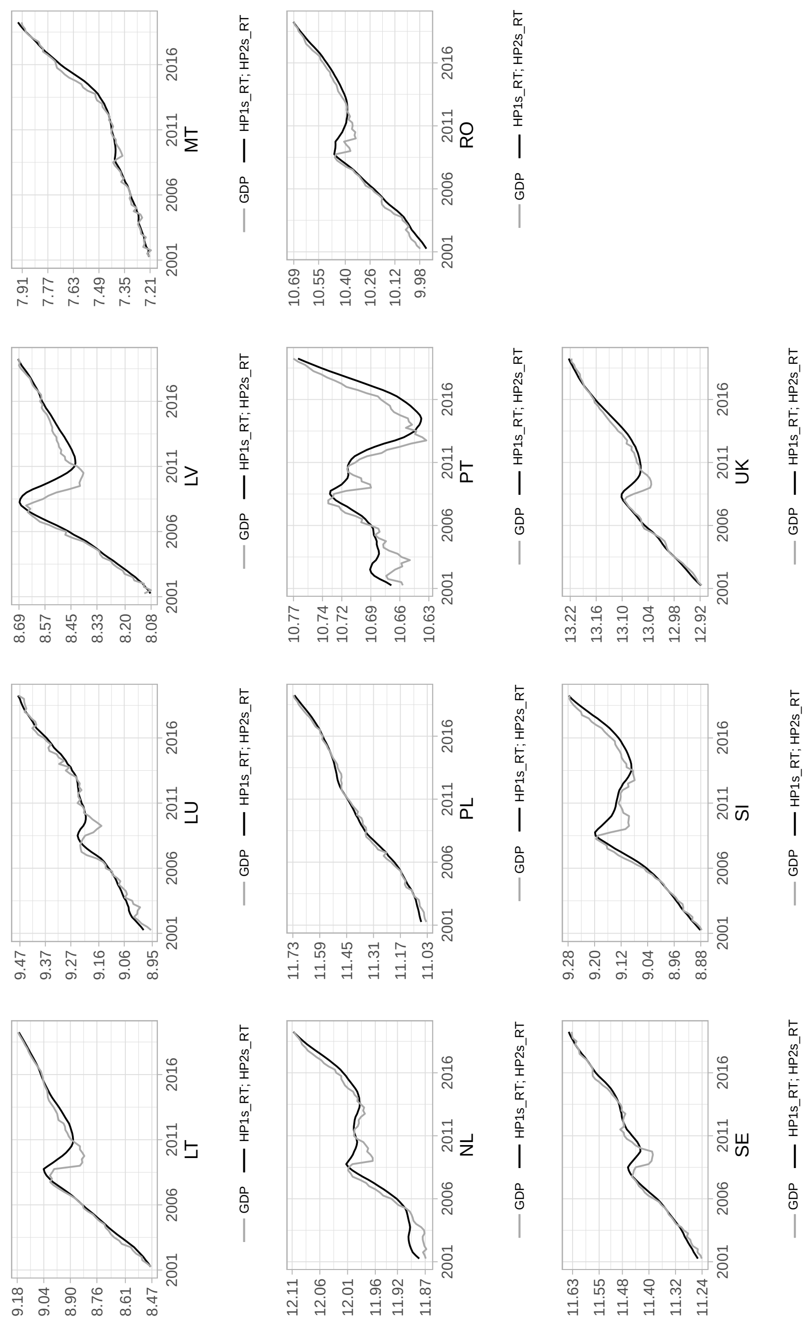

the corresponding two trend lines fully overlap. This is reflected by a single solid black line that stands for both trends in each Member State subplot. This feature of one-sided HP filters was explicitly commented for the first time (after a referee remark) in Mehra (2004), however without further qualification concerning the implications for biasedness; instead, the paper just maintained the idea that "...the standard HP filter could generate biased results..." (ibid., p. 7). What can be inferred from the latter is that the author had the implicit understanding that the one-sided filter is providing the desired unbiased estimates.

Returning to the identity between the results of the two versions of the HP filter when applied in real time, however, points to a completely different conclusion. If the standard two-sided filter leads to biased results at sample endpoints, then running it in real time implies biasedness for each filtered datapoint as in each iteration they correspond to endpoints. The identical results produced by the one-sided filter cannot then be unbiased but should be qualified as biased for each filtered datapoint.

The one-sided HP filter has the attractive property that it produces the same results both in real time and when applied to the full sample. However, having in mind the previous conclusion, the full-sample one-sided HP filter implies biasedness for each filtered datapoint, too.

Given the above considerations, it turns out that the standard two-sided HP filter remains the better, although still very far from perfect, alternative as it is unbiased at least for the non-endpoint observations of the sample (i.e. the ones that remain after the endpoint estimates implying biasedness are eliminated).

The question that follows is whether the one-sided filter is useful at all, as the above reasoning suggests that it should be totally discarded. Still, based on the logic followed so far, it could possibly be used to assess the total bias that accumulates over time if the HP filter is used for real-time assessment of a variable's deviations from its trend. To gain a first impression concerning the presence of such a bias, see Figure 2 in which the plots give an idea of the differences between the full-sample standard two-sided HP filter, and the one-sided one. As already mentioned, it would not matter whether the full-sample or the real-time estimates using the one-sided filter are plotted, as those are virtually the same. The visual inspection points to two observations. First, towards the sample ends the differences between the two trends tend to be smaller, even negligible; put differently, differences tend to be larger for the interior sample datapoints. Second, the differences between the two trends tend to be larger when the trends are less smooth. A lower degree of smoothness practically means larger shifts in the trend direction, and is itself implied by by changes in the growth rate of the filtered series.

Looking at the calculated output gaps using correspondingly the one-sided and the two-sided HP filters (Figure 3), some additional observations can be made. Assuming the one-sided version is the one bearing bias at all dates, then for example it could be said that it fails to capture well the height of the pre-Great-Recession boom in most of the Member States (except only maybe for Poland). Also, it tends to overestimate

Figure 2: Comparison of the full-sample two-sided HP filter with the one-sided one

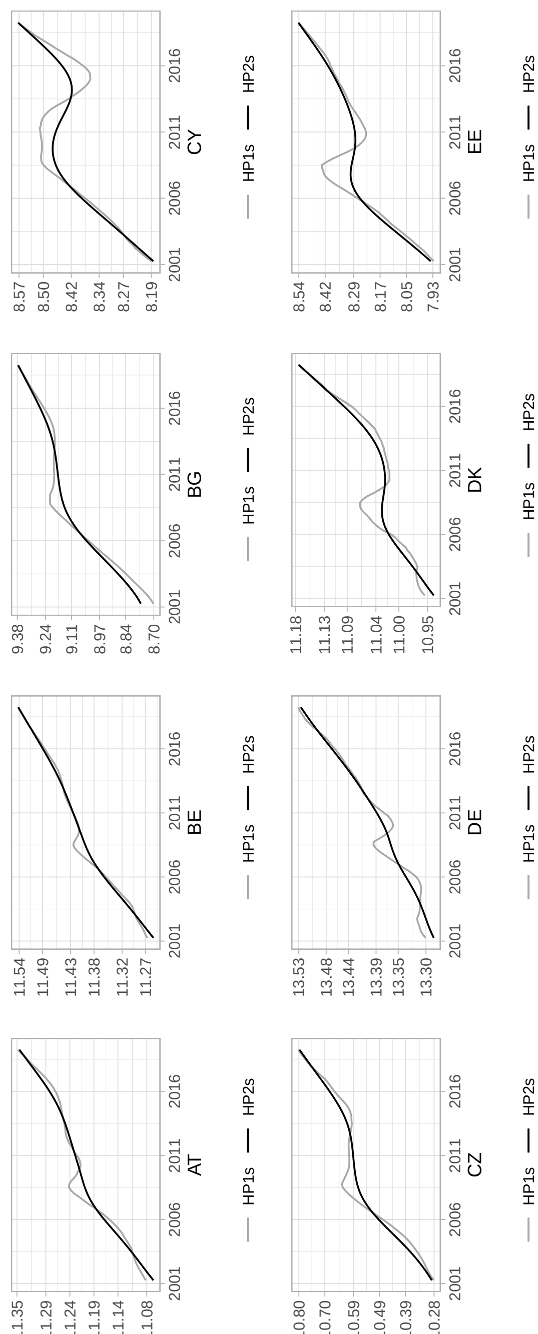

Figure 2 (cont.): Comparison of the full-sample two-sided HP filter with the one-sided one

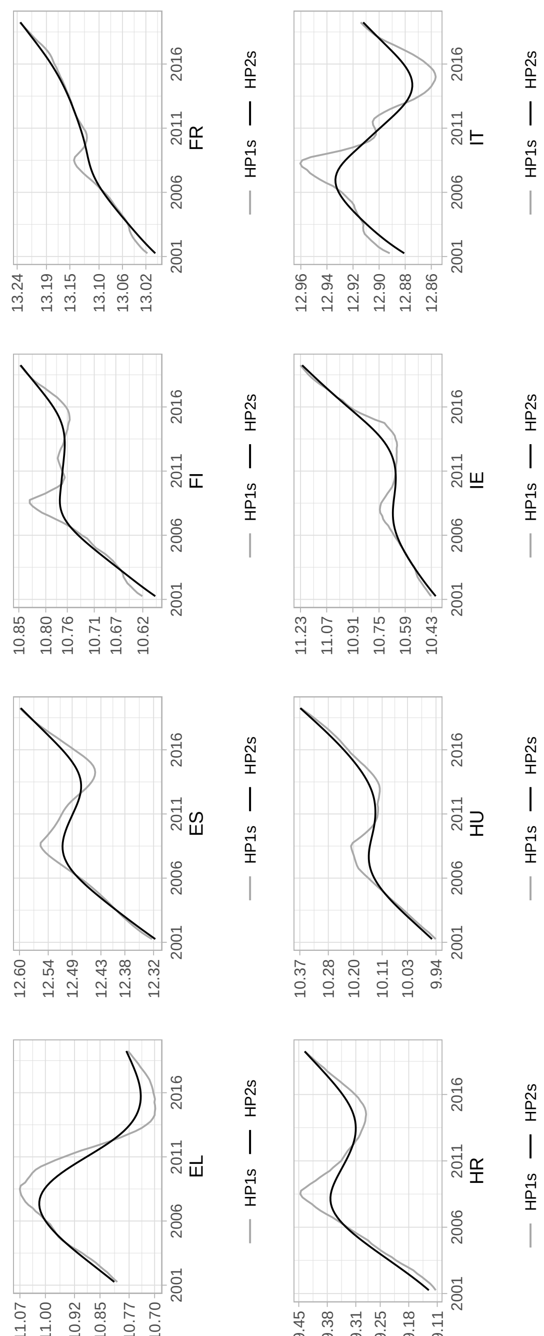

Figure 2 (cont.): Comparison of the full-sample two-sided HP filter with the one-sided one

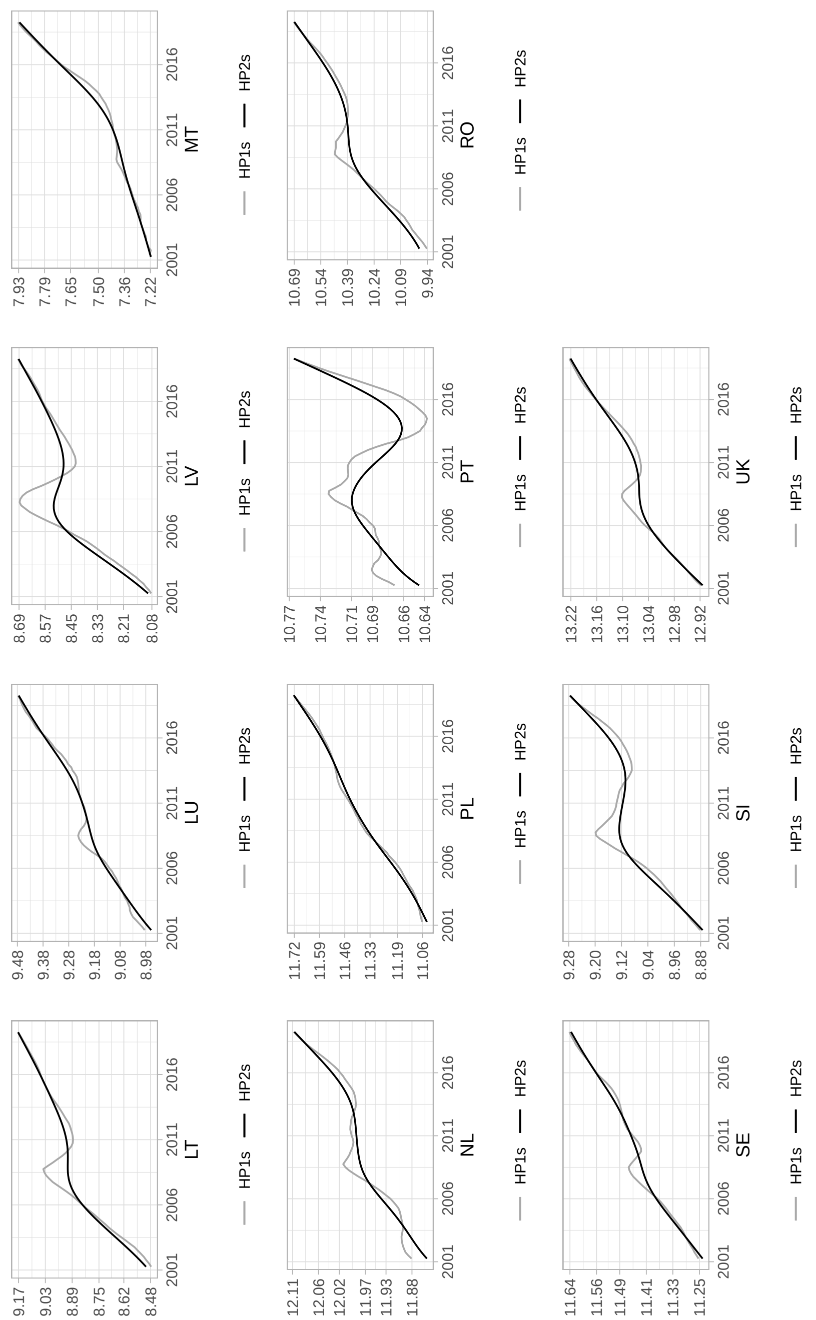

Figure 3: Comparison of output gaps obtained from the application of the one-sided and the two-sided HP filters

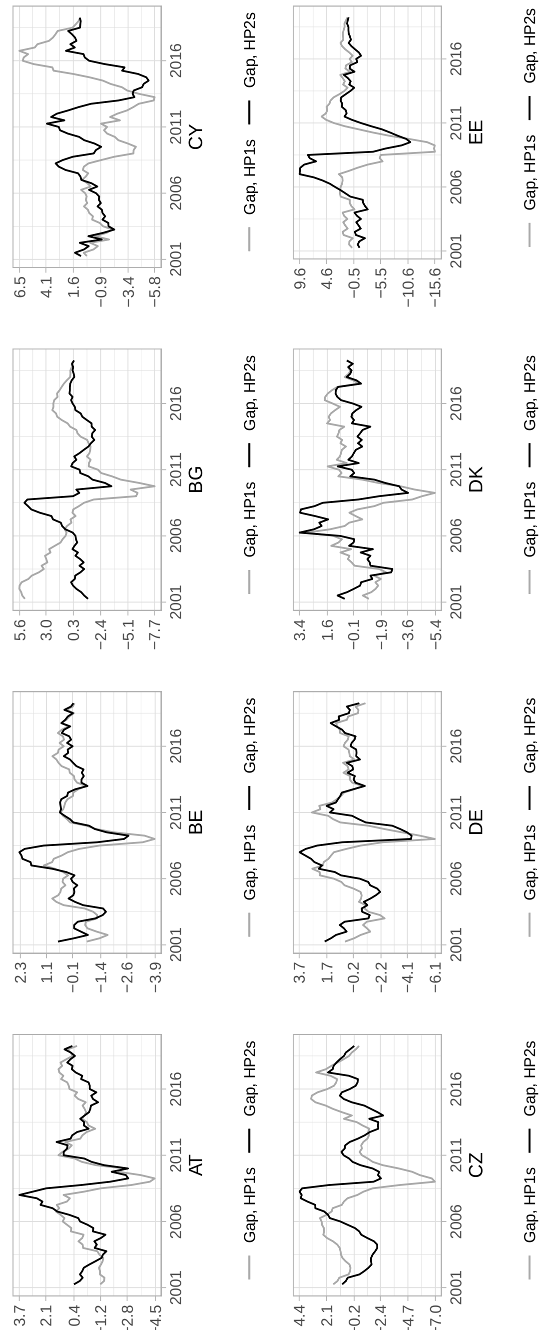

Figure 3 (cont.): Comparison of output gaps obtained from the application of the one-sided and the two-sided HP filters

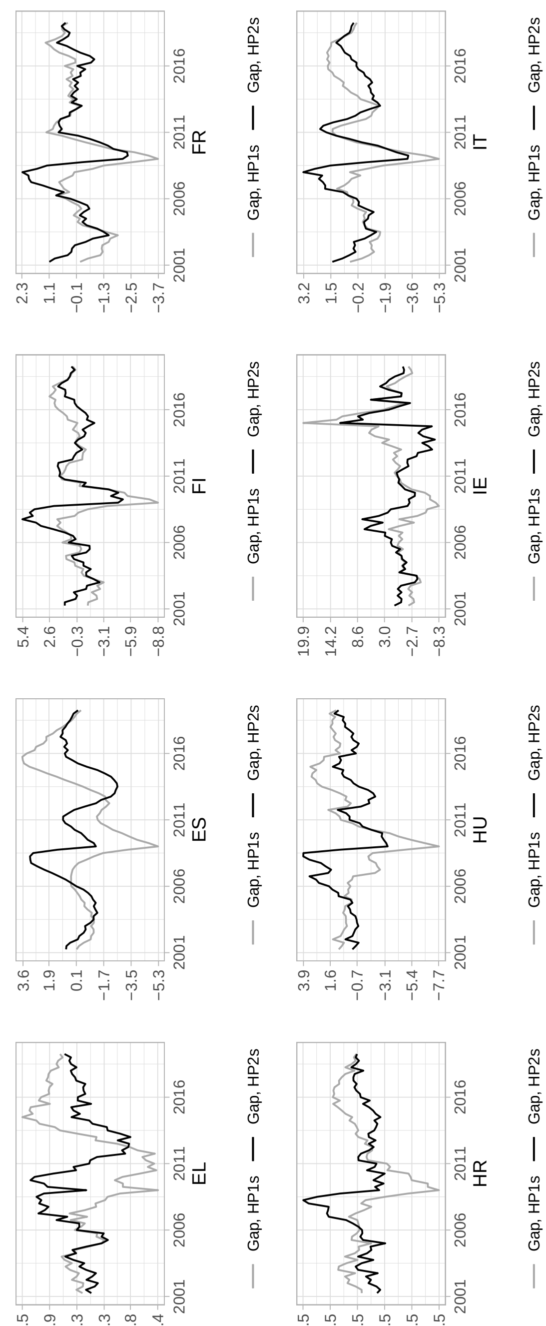

Figure 3 (cont.): Comparison of output gaps obtained from the application of the one-sided and the two-sided HP filters

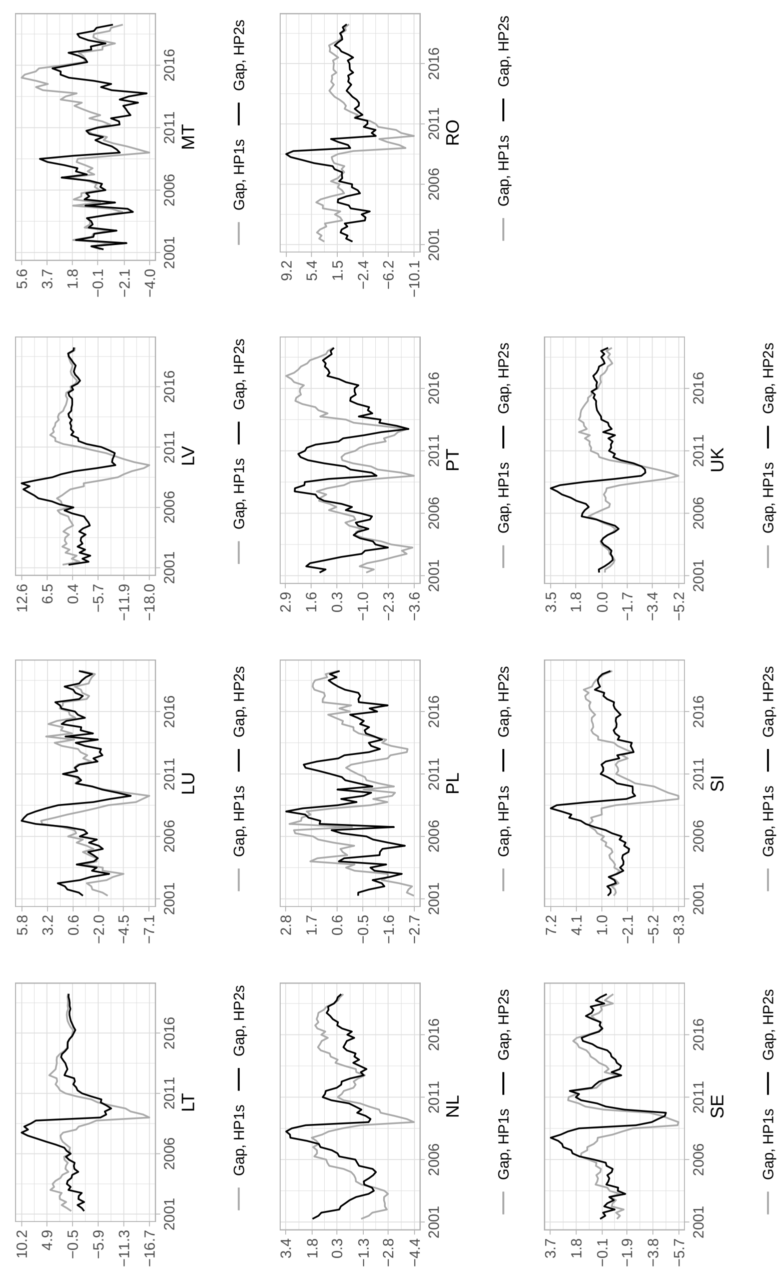

the negative output gaps during the following economic downturn. Illogical behaviour could also be traced for other periods, too. For example, the one-sided filter estimates a very strong boom period for Bulgaria in the beginning of the 2000s, and similar odd situations around 2015-2016 in Cyprus, Greece, Spain, and Portugal, among others.

## 4.2 The difference between one-sided and two-sided HP filter results as a function of empirical moments

A potential unevenness of the growth rate of the analysed series should be reflected in the shape of the statistical distribution of the acceleration rate, i.e. whether extreme values are more frequent, tails are fatter, etc.

While working directly with growth rates is in principle possible, the properties of the data might imply the impossibility to objectively assess the characteristics of their distributions. Specifically, in the current dataset, the growth rate series for some of the Member States do not comply with the requirement for stationarity (stationarity was tested using the KPSS test). This is largely due to the presence of structural breaks (regime shifts) for the series concerning some Member States resulting from the Great Recession of 2007-2009. The presence of non-stationarity renders the calculation of the empirical distribution moments one of questionable value.

As a first step in the assessment of the overall bias introduced by the one-sided filter, we hypothesize that the magnitude of the differences between the one-sided and the full-sample two-sided trends is determined by the properties of the distribution of the GDP acceleration rates (the growth acceleration series are stationary for all Member States in the data sample). Here, we assess those properties on the basis of their second, third, and fourth empirical moments, and more specifically we calculate the empirical distribution variance, skewness, and kurtosis of each series.

The differences between the two types of trends can be positive or negative, therefore, in order to get an average measure we use the mean of squares of per-observation differences, $MSD$:

$$MSD = \frac{1}{T}\sum_{t=1}^{T}(HP1s_t - HP2s_t)^2. \tag{10}$$

The acceleration rate of the log of GDP, $y_t$, is identified in the present context with the change in the year-on-year growth rate of the variable. Denoting the latter by $\alpha$, we have $\alpha = \Delta_4^2 y_t = \Delta(y_t - y_{t-4})$ for quarterly series. Denote the variance of this growth rate by $\sigma_\alpha^2$, its skewness by $S_\alpha$, and its kurtosis by $K_\alpha$. (The calculated empirical moments are presented in Table 1.) Then, in order to study the hypothesised relationship, the following regression equation could be specified:

$$MSD_i = \beta_0 + \beta_1 \cdot \sigma_{\alpha,i}^2 + \beta_2 \cdot S_{\alpha,i} + \beta_3 \cdot K_{\alpha,i} + \varepsilon_i, \quad i = 1, 2, \ldots, 27, \tag{11}$$

where $i$ indexes the Member States, $\beta_0, \beta_1, \beta_2$ and $\beta_3$ are parameters, and $\varepsilon_i$ is a random disturbance term.

Table 1: Moments of acceleration rates

| | Member State | Variance*$10^5$ | Skewness | Kurtosis |
|---|---|---|---|---|
| 1 | AT | 5.72 | -0.24 | 6.03 |
| 2 | BE | 4.69 | 0.38 | 5.51 |
| 3 | BG | 9.39 | -0.95 | 7.95 |
| 4 | CY | 10.82 | -0.02 | 2.47 |
| 5 | CZ | 10.65 | -0.24 | 4.04 |
| 6 | DE | 11.82 | 0.74 | 6.12 |
| 7 | DK | 7.92 | 0.12 | 4.87 |
| 8 | EE | 47.49 | 0.40 | 6.28 |
| 9 | EL | 13.76 | -0.40 | 3.39 |
| 10 | ES | 4.70 | -0.45 | 4.37 |
| 11 | FI | 18.43 | 0.27 | 5.81 |
| 12 | FR | 3.49 | 0.27 | 4.96 |
| 13 | HR | 11.51 | -1.15 | 6.76 |
| 14 | HU | 11.05 | -0.35 | 5.21 |
| 15 | IE | 57.65 | -0.70 | 4.81 |
| 16 | IT | 7.73 | 0.29 | 5.63 |
| 17 | LT | 47.42 | -0.25 | 6.32 |
| 18 | LU | 21.38 | 0.05 | 3.78 |
| 19 | LV | 41.72 | -0.38 | 4.20 |
| 20 | MT | 11.42 | -0.23 | 2.98 |
| 21 | NL | 5.76 | -0.66 | 4.97 |
| 22 | PL | 5.28 | -0.47 | 3.29 |
| 23 | PT | 7.34 | 0.23 | 3.06 |
| 24 | RO | 24.35 | -1.42 | 6.28 |
| 25 | SE | 12.93 | 0.70 | 5.13 |
| 26 | SI | 18.35 | -0.65 | 6.37 |
| 27 | UK | 6.23 | -0.06 | 6.58 |

Due to the fact that both in the left-hand side and in the right-hand side there are squared magnitudes, and having in mind that scales can differ substantially for the two variables, it is possible that heteroskedasticity appears in the regression residuals. Therefore, it would be a better alternative to work in terms of square roots of the two variables, respectively the square root of $MSD$, $RMSD$, and the standard deviation, $\sigma_\alpha$. The regression equation then becomes:

$$RMSD_i = \beta_0 + \beta_1 \cdot \sigma_{\alpha,i} + \beta_2 \cdot S_{\alpha,i} + \beta_3 \cdot K_{\alpha,i} + \varepsilon_i, \quad i = 1, 2, \ldots, 27. \tag{12}$$

The left-hand-side variable, $RMSD$, can however take theoretically only non-negative values. In the present case, all its values are positive. Therefore, the distribution of the disturbance term cannot be Gaussian. A more adequate specification of the regression equation would be then the following:

$$RMSD_i = \exp(\beta_0 + \beta_1 \cdot \sigma_{\alpha,i} + \beta_2 \cdot S_{\alpha,i} + \beta_3 \cdot K_{\alpha,i} + \varepsilon_i), \quad i = 1, 2, \ldots, 27. \tag{13}$$

Taking natural logs of both sides of this equation, the following is obtained:

$$\ln(RMSD_i) = \beta_0 + \beta_1 \cdot \sigma_{\alpha,i} + \beta_2 \cdot S_{\alpha,i} + \beta_3 \cdot K_{\alpha,i} + \varepsilon_i, \quad i = 1, 2, \ldots, 27. \tag{14}$$

There are generally two options to estimate this equation: either through OLS or through GLM. The former is justified if the regression residuals are homoskedastic and their distribution is symmetric (Manning and Mullahy, 2001). Therefore, first OLS estimation is applied, and then tests for heteroskedasticity and normality are applied to the resulting regression residuals in order to decide whether to go for GLM estimation or not.

Table 2: OLS estimation output (standard errors in parentheses)

| | *Dependent variable:* |
|---|---|
| | $\ln(RMSD)$ |
| $\sigma_\alpha$ | $25.279^{***}$<br>(3.383) |
| $S_\alpha$ | $-0.426^{***}$<br>(0.101) |
| $K_\alpha$ | $-0.078^{*}$<br>(0.040) |
| Constant | $-4.500^{***}$<br>(0.228) |
| Observations | 27 |
| $R^2$ | 0.782 |
| Adjusted $R^2$ | 0.754 |
| Residual Std. Error | 0.268 (df = 23) |
| $F$ Statistic | $27.563^{***}$ (df = 3; 23) ($p = 0.00000$) |

Note: $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

The results from the OLS estimation procedure are displayed in Table 2. At the 5% level of significance, all estimated coefficients are statistically significant except the coefficient before the kurtosis, $K_\alpha$. The signs of the significant slope coefficients show that the chosen measure of difference between the two types of filters, $RMSD$, is positively associated with the standard deviation of the growth acceleration rates, and negatively with their skewness. The first of those positive slope coefficients shows that the higher the variance of the acceleration rates, the larger the difference between the two HP trends. Its value looks very high but this can be explained by the fact that the scales of the regressand and the regressor are quite different. In particular, the acceleration rates are expressed in fractions of one, therefore their standard deviation will be on the same scale. At the same time, $RMSD$ is a difference of natural logs, therefore it would produce percentage deviations. By the logic of equation (14), however, taking a log of those percentages leads to much larger (negative) values. The negative slope coefficient before the skewness statistic shows that larger differences between the two HP trends are associated with cases where the distribution of growth acceleration rates is negatively skewed. Taken together, the two significant slope coefficients point to the conclusion that the difference between the two trends tends to be larger for economies where less frequent but more severe downturns are experienced. The insignificant coefficient before the kurtosis indicates that the fourth moment plays no defining role in the present setting.

Table 3: OLS residual tests

| Test | Test statistic | $p$-value |
|---|---|---|
| Breusch-Godfrey (up to order 4) | 3.059 | 0.548 |
| Breusch-Pagan (Koenker) | 4.965 | 0.174 |
| Jarque-Bera | 1.032 | 0.597 |
| Kolmogorov-Smirnov | 0.333 | 0.004 |

In order to establish whether the OLS estimation method is justified, the residuals from the OLS regression are tested for autocorrelation, heteroskedasticity, and normality, respectively using the Breusch-Godfrey, studentized Breusch-Pagan (Koenker), and Jarque-Bera tests. The results are presented in Table 3. They clearly show that the regression features neither autocorrelation, nor heteroskedasticity. Based on the Jarque-Bera statistic, it seems that the normality of residuals hypothesis cannot be rejected, too. However, the sample used here contains just 27 observations and its distribution could be problematic. Thadewald and Büning (2007) for example show that the Jarque-Bera test performs poorly when the distribution of the data is characterized with short tails. In the current setting, the kurtosis of the residuals equals approximately 2.06, i.e. indeed this is a case of a platycurtic distribution. Therefore, some normality cross-checking is due. Using e.g. the Kolmogorov-Smirnov test, and looking at the normal QQ-plot of the residuals (see Figure 4), normality, therefore symmetry of the distribution cannot be so light-heartedly accepted.

Thus, in order to cross-validate the results, the relationship stated in equation (12) is estimated as a GLM. A Gamma distribution with a log link is assumed. The estimation results are displayed in Table 4. In terms of point estimates, the differences to the OLS regression output are small. The only remarkable difference is that the slope coefficient before the kurtosis regressor becomes statistically significant ($p$-value just below 5%). The negative sign implies that the bigger the kurtosis of the empirical distribution, the smaller the difference between the two HP trends. More specifically, the fatter the tails of the distribution of acceleration rates are, the smaller the overall

Figure 4: Normal Q-Q plot of the OLS residuals

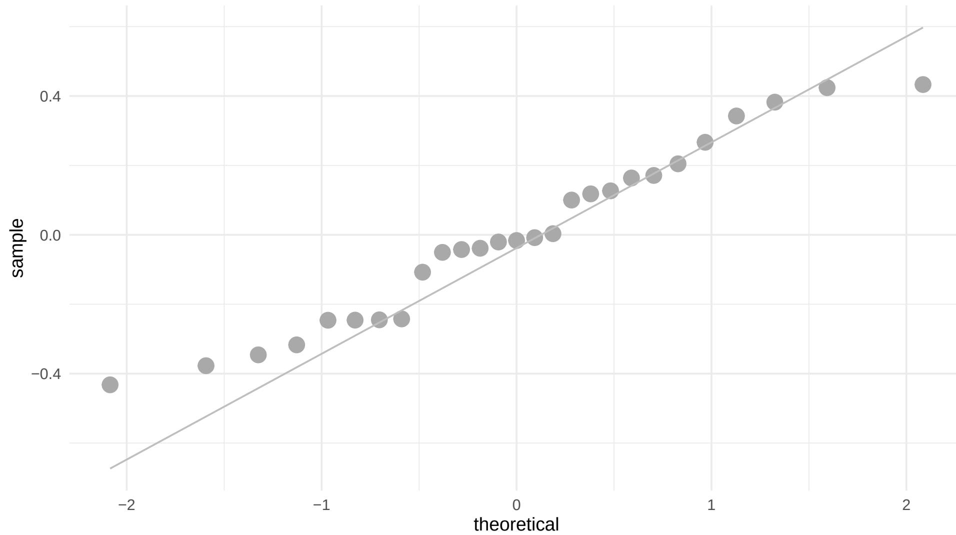


Table 4: GLM estimation output (standard errors in parentheses)

| | *Dependent variable:* |
|---|---|
| | *RMSD* |
| $\sigma_\alpha$ | 25.604*** |
| | (3.388) |
| $S_\alpha$ | −0.443*** |
| | (0.102) |
| $K_\alpha$ | −0.085** |
| | (0.040) |
| Constant | −4.445*** |
| | (0.229) |
| Observations | 27 |
| Log Likelihood | 103.826 |
| Akaike Inf. Crit. | −197.650 |
| Residual Deviance | 1.650 (df = 23) |
| Null Deviance | 7.512 (df = 26) |

Note: $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

bias introduced by the one-sided HP filter. In other words, the more outliers there are in the series of acceleration rates, the closer the results of the two filters will be.

# 5 Conclusions

The one-sided HP filter has been proposed to solve the endpoint issues present in the standard two-sided version. However, instead of solving them, it spreads and consolidates the endpoint bias over the entire filtered data sample. Given also that there is no evidence in literature that the one-sided filter solves the spurious cycles problem, too, it turns out that it is even worse suited for the task of trend-cycle decompositions. The results from the regression estimations that were carried out and presented in the article demonstrate that the higher the volatility inherent in the acceleration rate of the filtered series, the lower the reliability of the one-sided estimates for the entire filtered sample. Concerning the estimated linkage with the skewness parameter, it turns out that left-skewness of the distribution of acceleration rates is conducive to a higher bias produced by the one-sided HP filter. With respect to the slope parameter on the kurtosis parameter, higher values tend to lower the difference between the results of the two filters.

# References

[1] Cogley T., Nason J. M., (1995), Effects of the Hodrick-Prescott Filter on Trend and Difference Stationary Time Series. Implications for Business Cycle Research, *Journal of Economic Dynamics and Control* 19, 253–278.

[2] Hamilton J., (2018), Why You Should Never Use the Hodrick-Prescott Filter, *Review of Economics and Statistics* 100(5), 831–843.

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

[4] Hodrick R. J., Prescott E. C., (1997[1980]), Postwar U.S. Business Cycles: An Empirical Investigation, *Journal of Money, Credit and Banking* 29(1), 1–16.

[5] Manning W. G., Mullahy J., (2001), Estimating Log Models: To Transform or Not to Transform? *Journal of Health Economics* 20, 461–494.

[6] Mehra Y. P., (2004), The Output Gap, Expected Future Inflation and Inflation Dynamics: Another Look, *Topics in Macroeconomics* 4(1),

[7] Ravn M. O., Uhlig H., (2002), On Adjusting the Hodrick-Prescott Filter for the Frequency of Observations, *The Review of Economics and Statistics* 84(2), 371–376.

[8] Stock J. H., Watson M. W., (1999), Forecasting Inflation, *Journal of Monetary Economics* 44, 293–335.

[9] Thadewald T., Büning H., (2007), Jarque-Bera Test and its Competitors for Testing Normality – A Power Comparison, *Journal of Applied Statistics* 34(1), 87–105.