DEUTSCHE BUNDESBANK
EUROSYSTEM

Discussion Paper

Deutsche Bundesbank
No 03/2018

On the cyclical properties of Hamilton's regression filter

Yves S. Schüler

Deutsche Bundesbank, Wilhelm-Epstein-Straße 14, 60431 Frankfurt am Main,
Postfach 10 06 02, 60006 Frankfurt am Main

Tel +49 69 9566-0

Please address all orders in writing to: Deutsche Bundesbank,
Press and Public Relations Division, at the above address or via fax +49 69 9566-3077

Internet http://www.bundesbank.de

ISBN 978–3–95729–426–5 (Printversion)
ISBN 978–3–95729–427–2 (Internetversion)

Non-technical summary

Research question

Many economic variables, such as the gross domestic product, grow over time. To analyse the cyclical properties of such variables, as for example the business cycle, they are often corrected for a trend. For this purpose, there are several filtering techniques available. Recently, the widely known Hodrick-and-Prescott filter (HP filter) has been criticised by Hamilton (2017), as it induces spurious cycles and suffers from an end-of-sample bias when applied to a typical economic time series. Furthermore, a smoothing parameter must be chosen, which is most often selected on an ad hoc basis. As an alternative, Hamilton proposes a filter that is based on a regression (Hamilton filter). In this study, the properties of the Hamilton filter are analysed and compared to those of the HP filter.

Contribution

I contribute by providing theoretical and empirical insights on how the Hamilton filter addresses or attenuates the indicated drawbacks to the HP filter. Thus, I illustrate the properties of the Hamilton filter so that users can sensibly apply the filter.

Results

The Hamilton filter is also based on ad hoc assumptions and also induces a certain cyclical structure in a typical economic time series. The Hamilton filter most strongly emphasises cycles that exceed the duration of regular business cycles, i.e., longer than eight years, and completely mutes certain shorter-term fluctuations. Due to this, the Hamilton filter falls short of reproducing the chronology of US business cycles. Nonetheless, the amplification of cycles longer than business cycles can be regarded as a desirable property for some applications. For instance, a credit-to-GDP gap derived via the Hamilton filter indicates that imbalances prior to the global financial crisis started earlier than shown by the Basel III credit-to-GDP gap, which is derived using the HP filter. In general, I find that at the end of the sample the Hamilton filter produces more robust cycle estimates than the HP filter, which can be important if policy measures draw upon these estimates.

Nichttechnische Zusammenfassung

Fragestellung

Viele ökonomische Variablen, wie zum Beispiel das Bruttoinlandsprodukt, wachsen im Zeitverlauf. Um die zyklischen Eigenschaften solcher Variablen, also zum Beispiel den Konjunkturzyklus, analysieren zu können, werden sie häufig um einen Trend bereinigt. Dafür stehen verschiedene Filterverfahren zur Verfügung. Der weit verbreitete Hodrick-und-Prescott-Filter (HP-Filter) wurde vor Kurzem von Hamilton (2017) kritisiert, da bei seiner Anwendung auf typische ökonomische Variablen Scheinzyklen und Verzerrungen am Beginn und am Ende der Stichprobe entstehen. Zudem muss ein Glättungsparameter gewählt werden, was meist ad hoc geschieht. Als Alternative schlägt Hamilton einen auf einer Regression basierenden Filter vor (Hamilton-Filter). In dieser Studie werden die Eigenschaften des Hamilton-Filters analysiert und mit denen des HP-Filters verglichen.

Beitrag

Ich arbeite sowohl theoretisch, als auch empirisch heraus, inwiefern der Hamilton-Filter die Nachteile des HP-Filters vermeiden oder abmildern kann. Somit gebe ich Einblicke in die Eigenschaften des Hamilton-Filters, wodurch Nutzer in die Lage versetzt werden sollen, den Filter sinnvoll zu verwenden.

Ergebnisse

Auch der Hamilton-Filter basiert auf Ad-hoc-Annahmen und erzeugt eine spezifische zyklische Struktur in typischen ökonomischen Zeitreihen. Der Hamilton-Filter hebt Zyklen besonders hervor, die länger als reguläre Konjunkturzyklen sind, d. h. länger als acht Jahre. Außerdem löscht der Filter gewisse kürzerfristige Schwankungen. Daher kann der Hamilton-Filter beispielsweise die US-amerikanischen Konjunkturzyklen nicht korrekt reproduzieren. Die Hervorhebung von Zyklen, die länger als Konjunkturzyklen sind, kann jedoch für manche Anwendungen wünschenswert sein. Zum Beispiel signalisiert eine Kredit/BIP-Lücke, die mit Hilfe des Hamilton-Filters bestimmt wurde, Ungleichgewichte vor der globalen Finanzkrise früher als die Kredit/BIP-Lücke nach Basel III, die auf dem HP-Filter basiert. Im Allgemeinen zeigt sich, dass der Hamilton-Filter am Stichprobenende robustere Zyklus-Schätzungen liefert als der HP-Filter, was von großer Bedeutung sein kann, wenn sich Politikmaßnahmen auf diese Schätzungen stützen.

Deutsche Bundesbank Discussion Paper No 03/2018

On the cyclical properties of Hamilton's regression filter*

Yves S. Schüler
Deutsche Bundesbank

Abstract

Hamilton (2017) criticises the Hodrick and Prescott (1981, 1997) filter (HP filter) because of three drawbacks (i. spurious cycles, ii. end-of-sample bias, iii. ad hoc assumptions regarding the smoothing parameter) and proposes a regression filter as an alternative. I demonstrate that Hamilton's regression filter shares some of these drawbacks. For instance, Hamilton's ad hoc formulation of a 2-year regression filter implies a cancellation of two-year cycles and an amplification of cycles longer than typical business cycles. This is at odds with stylised business cycle facts, such as the one-year duration of a typical recession, leading to inconsistencies, for example, with the NBER business cycle chronology. Nonetheless, I show that Hamilton's regression filter should be preferred to the HP filter for constructing a credit-to-GDP gap. The filter extracts the various medium-term frequencies more equally. Due to this property, a regression-filtered credit-to-GDP ratio indicates that imbalances prior to the global financial crisis started earlier than shown by the Basel III credit-to-GDP gap.

Keywords: Detrending, spurious cycles, business cycles, financial cycles, Basel III

JEL classification: C10, E32, E58, G01.

*Contact address: Yves S. Schüler, Research Centre, Deutsche Bundesbank, Wilhelm-Epstein-Strasse 14, 60431 Frankfurt am Main, Germany. E-Mail: yves.schueler@bundesbank.de. This paper has benefited from valuable comments and suggestions by Jochen Mankart, Emanuel Mönch, and Morten Ravn. Furthermore, I would like to thank Fabio Canova, Claudia Foroni, James Hamilton, and Mark Watson for their helpful discussions.

1 Introduction

A series of studies points out that the Hodrick and Prescott (1981, 1997) filter (HP filter) induces spurious cycles: in theory, when applied to difference stationary data, and empirically, when applied, for instance, to GDP, total credit, the credit-to-GDP ratio, house prices, equity prices and bond prices across G7 countries (see Harvey and Jaeger (1993); King and Rebelo (1993); Cogley and Nason (1995); A'Hearn and Woitek (2001); Pederson (2001); Schüler (2018)). In line with these studies, Hamilton (2017) argues that the typical economic time series is best approximated by a random walk, i.e., a difference stationary process, and thus one should not use the HP filter as a general approach for detrending economic time series. Hamilton (2017) also stresses that the HP filter suffers from an end-point bias and is commonly used with an ad hoc smoothing parameter that is at odds with a statistical formalisation of the problem. As an alternative, he proposes a regression filter.

However, when applied to a random walk, Hamilton's (2017) regression filter reduces to a difference filter. In the case of difference filters, we know that certain cycle frequencies are cancelled and others emphasised. However, there has been no discussion of how these difference filters, and thus Hamilton's regression filter, behave when applied to difference stationary data and how this compares to the characteristics of the HP filter. The purpose of this paper is to address this gap by discussing the cyclical properties of Hamilton's regression filter, emphasising potential distortions and their implications for applied economic analyses.

Specifically, I contrast the cyclical properties of Hamilton's 2-year and 5-year regression filters to the characteristics of the HP filter with smoothing parameters 1,600 (HP (1600)) and 400,000 (HP (400000)) for the case of quarterly data. While the 2-year regression filter and the HP (1600) filter have been suggested for business cycle analysis, the 5-year regression filter and the HP (400000) filter have been recommended for the analysis of financial variables, such as cycles in total credit or the credit-to-GDP ratio as in Basel III (see Hamilton (2017) or the recommendation ESRB/2014/1).[1] Here, cycles that exceed the duration of regular business cycles, or medium-term cycles, are also of interest.

My findings suggest that Hamilton's regression filter is not subject to the exact same drawbacks as the HP filter. However, the regression filter still modifies the original cyclical structure of economic time series. The characteristics of a detrended component is strongly determined by ad hoc assumptions, such as the consideration of a 2-year or 5-year specification. In general, Hamilton's regression filter most strongly emphasises frequencies that are longer than typical business cycle frequencies, i.e., longer than eight years, and completely mutes shorter-term fluctuations, which has consequences for economic analyses. For instance, the 2-year regression filter amplifies the variance of medium-term and longer-term cycles by a maximum factor of around 64 and mutes 2-year cycles completely. However, neglecting 2-year cycles (and emphasising longer-term cycles) reflects a strong assumption, as a typical recession, i.e., a half cycle, lasts about one year, which would thus be cancelled; see, for instance, the NBER's Business Cycle Dating Committee, Watson (1994), Harding and Pagan (2002), or – for similar statistics across 17 advanced

[1] Furthermore, see Basel Committee on Banking Supervision (2010). "Guidance for National Authorities Operating the Countercyclical Capital Buffer", December.

economies – Jordà, Schularick, and Taylor (2017).[2] As a consequence, I find that the 2-year regression filter does not capture all expansionary and contractionary phases as classified by the NBER, such as the quick contraction and recovery in US GDP around 1980. Due to the amplification of medium-term frequencies, but also the dampening of short-term cycles, detrended US GDP is marked by a prolonged contraction over that period.

Nonetheless, the amplification of medium-term frequencies can be argued to be a desirable property for some applications. For instance, an advantage of the 5-year regression filter over the HP (400000) filter is that it extracts medium-term frequencies more equally. That is, 30 year cycles are amplified over 8 year cycles by a factor of only 3.9 in the former and by a factor of 7.9 in the latter case. This can be argued to be favourable for the analysis of medium-term frequencies, which is the goal of the Basel III credit-to-GDP gap. Due to this property, I find that the 5-year regression-filtered credit-to-GDP gap indicates that pronounced levels of imbalances prior to the global financial crisis started more than four years earlier than detected using the HP (400000) filter.[3] While the lead time reduces in a real time exercise, the 5-year regression-filtered series nevertheless signals pronounced levels of imbalances one year prior to the HP-filtered indicator. In general, differences between the two filters diminish in a real time setting, although this is mostly due to revisions of the HP-filtered real time series compared to the full sample ones. Thus, I find that the regression filter has more robust real time properties, suggesting it has a smaller end-of-sample bias.[4]

Above all, the "correct" filter depends on the researcher's objective, i.e., the feature of the data she would like to focus on. If the objective is to remain agnostic about the importance of the different cyclical characteristics, one should use first differences to detrend a typical economic time series. For difference stationary data, such transformation preserves all dynamics of a series, while the HP filter – but also the regression filter – extract specific frequencies of a time series, masking potentially relevant fluctuations.

The paper is structured as follows: In Section 2, I characterise the HP filter and Hamilton's regression filter and reflect on Hamilton's critique of the HP filter. Next, I compare the cyclical properties of the two filters when applied to a stationary series and to a difference stationary series. In Section 4, I discuss the filters' effects on US GDP and the US credit-to-GDP ratio. Section 5 concludes.

2 Characterising the filters

This section characterises the HP filter and Hamilton's regression filter. Furthermore, I review Hamilton's critique on the HP filter and reflect on how his regression filter improves on the issues raised.

[2] The NBER's Business Cycle Dating Committee reports that the average duration of a contraction is 11.1 months for the period 1945-2009.

[3] I define pronounced levels of imbalances as a value of each gap prior to the global financial crisis that is equal or exceeds its previous value during the savings and loan crisis.

[4] Weak real time properties of the HP filter have also been identified in the seminal study by Christiano and Fitzgerald (2003), in which they compare the HP filter to their optimal finite-sample approximations of the band-pass filter. In addition to their analysis, I also study the real time performance of a medium-term HP filter, i.e., the HP filter (400000), which allows medium-term frequencies to be analysed.

Hodrick and Prescott filter: The HP filter decomposes an observed, possibly non-stationary time series, y_t, into its cyclical and secular (or trend) component:

$$y_t = \tau_t + \psi_t, \tag{1}$$

where τ_t is the secular and ψ_t the cyclical component. To separate these two components, one minimises the variance of ψ_t subject to a penalty for variation in the second difference of τ_t,

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

where λ controls the smoothness of the extracted trend. The higher its value, the smoother is the trend. Using quarterly data, λ is most commonly set to 1,600 for analysing business cycles. Furthermore, a value of 400,000 is recommended in the Basel III regulations for constructing a credit-to-GDP gap. This value has also been used in several academic studies researching on financial imbalances (see, for instance, Behn, Detken, Peltonen, and Schudel (2013); Anundson, Gerdrup, Hansen, and Kragh-Sørensen (2016); Bauer and Granziera (2016)).

As shown by King and Rebelo (1993), the HP filter yields a stationary detrended component as long as the fourth differences of the original series are stationary.

Hamilton (2017) references three specific drawbacks to the HP filter. First, a series of studies indicates that the HP filter induces spurious cycles, or spurious dynamic relationships (see Harvey and Jaeger (1993); King and Rebelo (1993); Cogley and Nason (1995); A'Hearn and Woitek (2001); Pederson (2001); Schüler (2018)). Specifically, the HP filter induces spurious cycles when applied to difference stationary time series, which is a leading example of a typical economic time series. Hamilton (2017) argues that a typical economic time series is best described by a random walk, as suggested by simple economic theory or out-of-sample forecasting exercises.[5] Given this evidence, Hamilton (2017) concludes that we should not be using the HP filter as an all-purpose method for detrending economic time series.

Second, the HP filter is a two-sided (or symmetric) filter, i.e., it considers both future and past observations to determine the cyclical component at a given time period. Thus, an extracted cyclical component at time t has the artificial ability to predict its future values, provided it is not at the end of the sample. This implies an end-of-sample bias as filtered values in the middle of the sample and at the end are very different. It can lead to substantial biases in small-samples.[6] Applying the HP filter on an expanding sample may eliminate this problem. Nonetheless, Hamilton (2017) notes that due to the way in which cycle and trend are characterised, changes in the one-sided trend and its implied cycle are still forecastable to some degree.

Third, the common choice $\lambda = 1,600$ is ad hoc and at odds with a statistical formalization of the problem. That is, Hodrick and Prescott motivated their choice of 1,600

[5] Just recently, Schüler (2018) empirically supports this finding for a broad set of macroeconomic and financial variables in the context of G7 countries.

[6] For instance, De Jong and Sakarya (2016) note that there may still be non-stationarity in the cyclical component near the start or end of the sample.

based on prior beliefs about the magnitudes of changes in the cyclical component relative to the trend component.[7] There have been suggestions for finding an optimal value of λ, but these – relying on strong assumptions – produce values of the smoothing parameter that are at odds with the common choice. Specifically, an optimal value of λ may be rationalized when formulating the HP filter as an optimal filter in a structural time series framework (see Hodrick and Prescott (1981); Harvey and Jaeger (1993); King and Rebelo (1993)). Here, $\lambda = \sigma_\psi^2/\sigma_\nu^2$, where σ_ψ^2 is the variance of the cyclical component (see Equation (1)) and σ_ν^2 is the variance of the second difference of the trend component, i.e,

$$\tau_t = 2\tau_{t-1} - \tau_{t-2} + \nu_t. \tag{3}$$

This specific formulation of the HP filter requires that ψ_t and ν_t are white noise and $E[\psi_t \nu_t] = 0$, which are clearly unrealistic assumptions, as also noted by Hodrick and Prescott (1981, 1997). Estimating such an optimal λ, Hamilton (2017) finds that it should be close to 1 rather than 1,600, for a series of macroeconomic and financial variables. Next, I present Hamilton's alternative.

Hamilton's regression filter: Hamilton (2017) proposes an OLS regression of the observed non-stationary time series, y_t, at date $t+h$ on a constant and its four most recent values as of date t, i.e.,

$$y_{t+h} = \beta_0 + \beta_1 y_t + \beta_2 y_{t-1} + \beta_3 y_{t-2} + \beta_4 y_{t-3} + v_{t+h}. \tag{4}$$

The stationary, or cyclical, component is then obtained from the residuals,

$$\hat{v}_{t+h} = y_{t+h} - \hat{\beta}_0 - \hat{\beta}_1 y_t - \hat{\beta}_2 y_{t-1} - \hat{\beta}_3 y_{t-2} - \hat{\beta}_4 y_{t-3}. \tag{5}$$

In the case of quarterly data, Hamilton (2017) suggests employing $h = 8$ for analyses concerned with business cycles and $h = 20$ for studies interested in credit or financial cycles.[8]

This procedure rests on results by Den Haan (2000), who shows that forecast errors are stationary for wide range of non-stationary processes. Thus, the filter has the advantage that we do not have to know the true data-generating process before applying it. As it uses four lags of the observed time series, it yields stationary residuals as long as the fourth differences of the original time series are stationary.

How does the filter improve on the drawbacks of the HP filter? Regarding Hamilton's (2017)'s first critique – appropriateness of a filter for a typical economic time series – it is important to note that the regression filter reduces to a difference filter when applied to a random walk. In this case the OLS estimates of Equation (4) converge, in large samples, to $\beta_1 = 1$ and all other $\beta_j = 0$ (Hamilton, 2017, pp. 16-17). Thus, the forecast error is

[7] "Our prior view is that a five percent cyclical component is moderately large as is a one-eighth of one percent change in the growth rate in a quarter. This led us to select $\sqrt{\lambda} = 5/(1/8) = 40$ or $\lambda = 1600$ as a value of the smoothing parameter" Hodrick and Prescott (1981, p. 6)

[8] Actually, Hamilton (2017) recommends $h = 5$ years, as he refers to the analysis of debt cycles on the basis of the dataset developed by Jordà, Schularick, and Taylor (2016), which has a yearly sampling frequency. Thus, $h = 20$ is the analogue for a quarterly sampling frequency.

simply the difference of y_{t+h} and y_t, i.e.,

$$\tilde{v}_{t+h} = y_{t+h} - y_t. \tag{6}$$

However, it is known that difference filters also have the potential to induce spurious cycles, as is the case, for instance, with the Kuznets (1961) filter. Applied to a stationary quarterly series, this filter amplifies cycles of duration just above 5 years, while relatively muting others (see, for example, Pederson (2001)).

Regarding his second criticism – symmetry of the HP filter – the regression filter is an asymmetric filter, i.e., it does not use information of future realisations. For instance, when applied to a random walk in large samples, it only uses information at time $t + h$ and t to determine the value of the cyclical component at time $t + h$. However, Hamilton (2017) recommends to determine the β-coefficients of his regression filter using the entire sample. Thus, the filter exploits the entire sample to construct estimates of the cyclical component as well. In this vein, the filter is also subject to potential small-sample biases, such as finite-sample approximations of filter weights or structural breaks.

Related to the ad hoc formulation of the HP smoothing parameter, Hamilton's regression filter is also based on ad hoc assumptions that influence the properties of the extracted cyclical component by assuming h to be 8 or 20.

Thus, Hamilton's regression filter also has specific properties and is based on certain assumptions that determine the characteristics of the detrended component. In light of this, I discuss in detail the cyclical properties of Hamilton's regression filter in large samples in the next section, contrasting them to the ones of the HP filter.

3 Contrasting the cyclical properties

To discuss the cyclical properties of the HP filter and Hamilton's regression filter, I first introduce the concept of a power transfer function (PTF) that completely describes the way these filters modify the cyclical characteristics of a series given a large sample; both for the case of a stationary series and a difference stationary series.

I study the characteristics of Hamilton's regression filter by examining the cyclical properties of difference filters. I justify this approach by the fact that Hamilton's regression filter reduces to a difference filter when applied to a random walk process (see Equation (6)). And as the typical economic time series is best approximated by a random walk process (see Hamilton (2017)), this is the empirically relevant case for studying the cyclical properties of Hamilton's regression filter.[9]

3.1 Power transfer functions for stationary data and difference stationary data

Assume that ψ_t is a stationary stochastic process with an autocovariance-generating function defined as $g_\psi(z) \equiv \sum_{t=-\infty}^{\infty} \gamma_t z^t$, where z denotes a complex scalar and γ_t the autoco-

[9]I assess the validity of this approach in the empirical exercise.

variances. The spectral density of ψ_t is then defined as

$$S_{\psi}(\omega) \equiv \frac{1}{2\pi} g_{\psi}(e^{-\mathrm{i}\omega}), \tag{7}$$

where i is the imaginary unit and $\omega \in [-\pi, \pi]$ the cycle frequency in radians. Filtering ψ_t with a time-invariant filter that has absolutely summable weights, say $\xi_t = \sum_{j=-\infty}^{\infty} h_j \psi_{t-j}$, implies that the spectral density is altered via a *transfer function*. This transfer function can be denoted by $h(e^{-\mathrm{i}\omega})$ and the exact relation between the spectral density of ψ_t and the spectral density of the filtered series, say ξ_t, is

$$S_{\xi}(\omega) = H(\omega) \cdot S_{\psi}(\omega), \tag{8}$$

where $H(\omega) \equiv |h(e^{-\mathrm{i}\omega})|^2$ is the *power transfer function*. The PTF completely describes the change in the relative importance of the cyclical components in ψ_t.[10] When $H(\omega) > 1$, the amplitude of the cycle component ω of ψ_t is increased. When $H(\omega) < 1$, the amplitude of the respective cycle component is dampened. Such amplification and dampening of frequencies in the stationary component ψ_t may give rise to spurious or artificial cycles.

If the observed time series, say y_t, is difference stationary (for example, follows a random walk process), the filters amplify and dampen other cyclical characteristics of ψ_t. To illustrate, let

$$y_t = \alpha + y_{t-1} + \psi_t, \tag{9}$$

where α is a scalar intercept. Assuming ψ_t is white noise, we obtain a random walk with possible drift (when $\alpha \neq 0$). The PTF illustrating the effects of a filter on ψ_t can be derived by the following steps: first, y_t is differenced to render the series stationary and, second, smoothed with the filter, adjusted for the use of a difference operator. Or put differently, the application of a filter on y_t "uses up" one difference operator (see, for example, Cogley and Nason (1995) or Murray (2003)). Thus, the filter operates as a two-step linear filter. The PTF of interest, say $J(\omega)$, which describes the effects on the stationary component ψ_t, can be derived by

$$S_{\xi}(\omega) = H(\omega) S_y(\omega) = J(\omega) S_{\Delta y}(\omega) = J(\omega) S_{\psi}(\omega) \tag{10}$$

where $J(\omega) = H(\omega)/H^{\Delta}(\omega)$ and $H^{\Delta}(\omega)$ is the power transfer function of the first difference filter.[11]

Next, I discuss the PTFs of the HP and the difference filters in both settings, i.e., when applied to a stationary series and a difference stationary series.

3.2 The cyclical properties of the HP filter and difference filters

Figure 1 shows the PTFs relevant for stationary data. Specifically, it introduces the PTFs of the HP filters (black lines) with a smoothing parameter of 1,600 and 400,000 and of the 2-year (Δ_8) and the 5-year (Δ_{20}) difference filters (red lines). Figure 2 presents the

[10] The transfer function can be decomposed into gain and phase, where the square of the gain is the power transfer function, i.e., $h(e^{-\mathrm{i}\omega}) = |h(e^{-\mathrm{i}\omega})| e^{-\mathrm{i}\Theta(\omega)}$. $\Theta(\omega)$ refers to the phase. Note that changes in the importance of frequencies are fully described by the power transfer function.

[11] The power transfer function of the first difference filter is shown in Appendix A.1.

effects of these filters on difference stationary data.[12] In all graphs, the y-axis shows the squared gain induced by each filter across frequencies (PTF) as well as the values of the spectral density ($S(\omega)$) of a hypothetical time series (blue line), a white noise process with spectral density of one at all frequencies to illustrate the distortions induced by the filters. It is this case for which the PTFs of filters indicate exactly how the cyclical structure of the underlying series would be modified. The x-axis gives the cycle length in radians.[13] The shortest cycle length is at π (2 quarters, if using quarterly data) and the longest at 0 radians (∞). Furthermore, the scale of cycle duration is non-linear, i.e., the closer to 0 radians, the larger the increase in cycle length. To make it easier to read the graphs, I mark business cycle frequencies (1.5-8 years) in blue and medium-term cycle periods (8-30 years) in purple, assuming a quarterly sampling frequency. The latter area is argued to be important for financial variables, such as the credit-to-GDP ratio.[14]

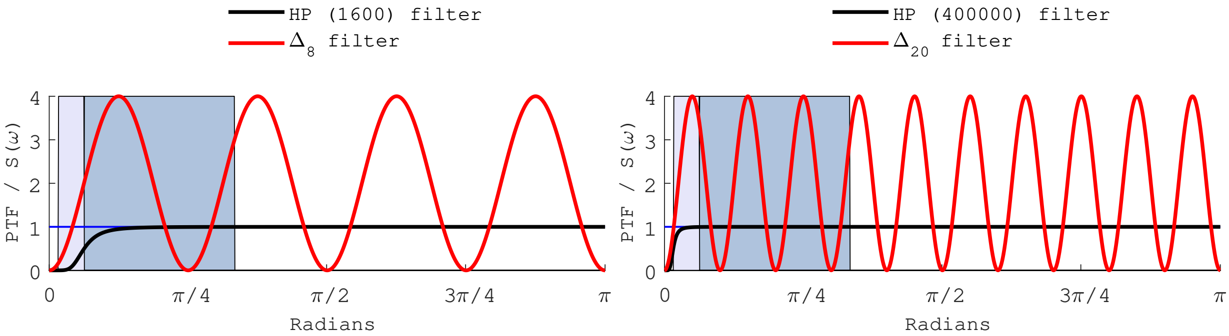

Figure 1: Power transfer functions of HP and difference filters: Effects on (trend) stationary data

Notes: The blue area indicates business cycle frequencies (1.5-8 years) and the purple area medium-term cycles (8-30 years), assuming quarterly data. PTF denotes the power transfer function and $S(\omega)$ is the spectral density. To illustrate the distortions of the filters, the blue line shows the spectral density of a white noise process with a value of one at all frequencies. Filtering this white noise process with the indicated procedures would lead to amplified and dampened frequencies as shown by the black and red lines.

Figure 1 illustrates that the HP filter is a high-pass filter, where shorter-term frequencies pass without distortion (PTF of one) up to a threshold that depends on the chosen smoothing parameter. The higher the smoothing parameter is, the longer the duration of cycles that pass the filter. That is, in the case of the HP (400000) filter, longer-term frequencies are apparent in the detrended component: up to cycles with a duration of around 30 years. In contrast, the difference filters amplify certain frequencies by a factor of 4, while completely cancelling others (PTF of zero). In the case of the Δ_8 filter, these amplified cycles are of a duration of around 4 years, 1.3 years, 0.8 year and 0.58 year. Furthermore, the removed cycles are around the duration of ∞ (the trend), 2 years, 1 year, 0.68 year, and 0.5 year. Thus, if one passes a white noise process through this filter,

[12]The exact formulas of the power transfer functions are given in Appendix A.1.

[13]The conversion to cycle duration in years using quarterly data is $\pi/(2\omega)$, where ω is the value of the x-axis, e.g., $\pi/(2\pi) = 0.5$ years, which is the shortest cycle duration that can be measured using quarterly data.

[14]See, for instance, Drehmann, Borio, and Tsatsaronis (2012); Borio (2014); Aikman, Haldane, and Nelson (2015); Hiebert, Klaus, Peltonen, Schüler, and Welz (2014); Schüler, Hiebert, and Peltonen (2015, 2017); Strohsal, Proaño, and Wolters (2015a,b); Rünstler and Vlekke (2016); Galati, Hindrayanto, Koopman, and Vlekke (2016); Verona (2016)).

it would have amplified and cancelled cycles of these durations. Similar results hold for the Δ_{20} filter, except that more cycle frequencies are amplified and cancelled.[15] Such amplification and dampening of cycles may lead to spurious dynamics, as it modifies the properties of the original stationary series.

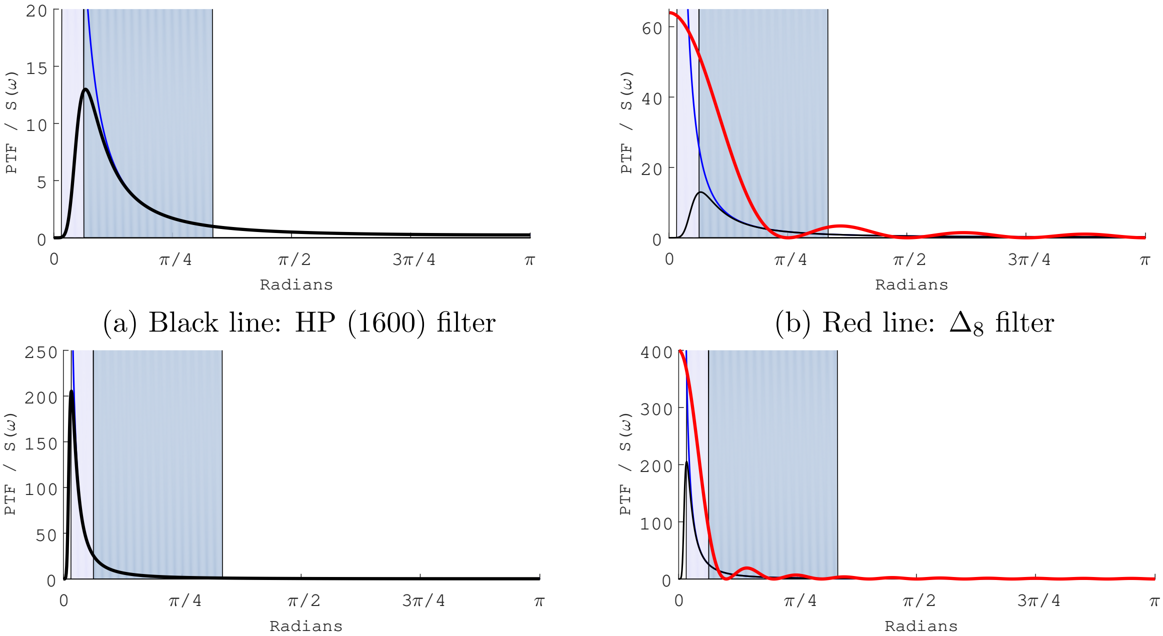

(a) Black line: HP (1600) filter

(b) Red line: Δ_8 filter

(c) Black line: HP (400000) filter

(d) Red line: Δ_{20} filter

Figure 2: Power transfer functions: The filters' effects on difference stationary data

Notes: The blue area indicates business cycle frequencies (1.5-8 years) and the purple area medium-term cycles (8-30 years), assuming quarterly data. PTF denotes the power transfer function and $S(\omega)$ is the spectral density. To illustrate the distortions of the filters, the blue line shows the spectral density of a random walk process whose first difference has a spectral density with a value of one at all frequencies. Filtering this random walk process with the indicated procedures would lead to amplified and dampened frequencies as shown by the black and red lines.

By contrast, if these filters are applied to a difference stationary series such as a random walk, the effects on the stationary component take the form as shown in Figure 2. I mark in blue the assumed time series process, i.e., a random walk whose first difference has a spectrum of one at all frequencies. For comparison purposes, I keep the PTFs of the corresponding HP filter in all graphs.

These graphs can be used to illustrate the phenomenon of spurious cycles in HP-filtered data. First, the PTFs of the HP filters peak and so a detrended series would be mainly characterised by the frequencies around this peak. These are cycles of a duration around 7.5 years for the HP (1600) filter and about 30 years for the HP (400000) filter. Hence, the duration of expansions and contractions of cycles in the detrended component are biased towards the peak frequencies. Second, the larger the smoothing parameter, the stronger the distortions are, as shorter-term frequencies are muted relatively more strongly (see Schüler (2018)). That is, the PTF of the HP (1600) filter peaks at a value of around 13, but at a value of about 204 for the HP (400000) filter, implying that in

[15] Here, cycles of a duration of around 10 years, 3.3 years, 2 years, 1.4 years, 1.1 years, 0.9 year, 0.8 year, 0.68 year, 0.6 year, and 0.52 year are amplified, while those of duration ∞ (the trend), 5 years, 2.5 years, 1.7 years, 1.2 years, 1 years, 0.8 year, 0.7 year, 0.6 year, 0.56 year, and 0.5 year are cancelled.

the latter case shorter-term frequencies become relatively less important. For instance, in the case of the HP (1600) filter, cycles of duration 7.5 years are amplified over yearly cycles by a factor of around 26. Using the HP (400000) filter entails 30-year cycles being amplified over yearly cycles by a factor of around 408.

For the Δ_8 and Δ_{20} filters, spurious cycles take a different form, as neither PTF peaks. Here, cycles shorter than 2 years (Δ_8 filter) and shorter than 5 years (Δ_{20} filter) are almost completely cancelled. Furthermore, cycles slightly longer than 2 years and 5 years are amplified beyond their presence in the original series, as the red lines are located above the blue line. Thus, in some sense, the higher-order difference filters act as low-pass filters in this setting, except that amplification increases with longer cycle durations.

A close comparison of the HP and difference filters suggests important differences in the amplification and dampening of certain cycle frequencies. Overall, both difference filters have a maximum factor of amplification (around 64 and 400) that is larger than the maximum factor of amplification of the corresponding HP-filter specification (about 13 and 204). Furthermore, while the Δ_8 filter most strongly amplifies cycles that are longer than business cycle frequencies, these are cancelled by the HP (1600) filter. As a result, fluctuations in the business cycle range are dampened relatively more strongly (and even erased) by the Δ_8 filter. By contrast, the medium-term filters do not differ to such an extent. Only the relative amplification of medium-term versus business cycle frequencies stands out. For instance, the Δ_{20} filter amplifies 30-year cycles over 8-year cycles by a factor of about 3.8, compared to a factor of around 7.9 for the HP (400000) filter.

Overall, both filters lead to distortions when applied to a typical economic time series. Clearly, in the case of a difference stationary series, the only way to remain agnostic about the relative significance of frequencies is to use a first difference filter. Figure 3 illustrates in green the PTF of the first difference filter when applied to a stationary process (left) and the PTF of the first difference filter when applied to a difference stationary process (right). While in the left panel we see the usual result that the first difference filter amplifies cycles shorter than 1.5 years when applied to a stationary series, the right panel indicates that, when applied to a difference stationary process, none of the cycle frequencies of the stationary component are altered (PTF of one), while removing the unit root.

4 Detrending US GDP and the US credit-to-GDP ratio

In this section, I compare HP-filtered series to regression-filtered series using actual time series in order to highlight the filters' implications for applied research. Specifically, I apply the filters to US log GDP and the US credit-to-GDP ratio.[16] Compared to the large-sample properties discussed in the previous section, additional small sample issues

[16] I obtain real US GDP spanning from 1947Q1 to 2017Q2 from FRED. I download the US credit-to-GDP ratio from the BIS. It spans from 1952Q1 to 2016Q4. Credit reflects loans to the private non-financial sector.

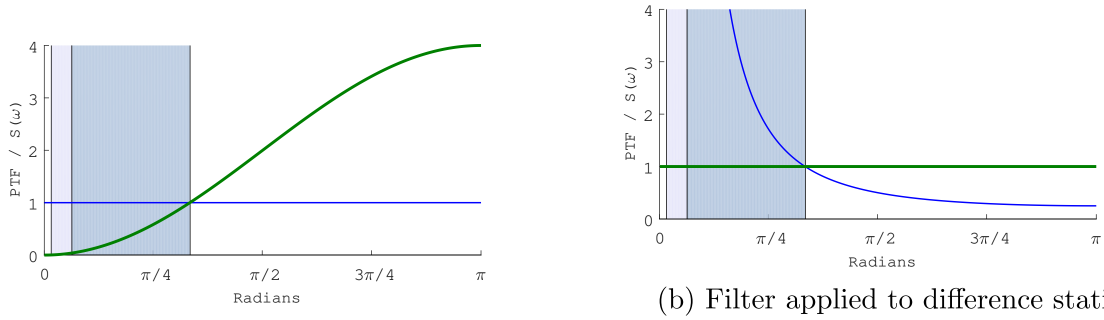

(a) Filter applied to stationary series

(b) Filter applied to difference stationary series

Figure 3: Green line: Power transfer functions of the first difference filter

Notes: The blue area indicates business cycle frequencies (1.5-8 years) and the purple area medium-term cycles (8-30 years), assuming quarterly data. PTF denotes the power transfer function and $S(\omega)$ is the spectral density. To illustrate the distortions of the filters, the blue line shows the spectral density of a white noise process with a value of one at all frequencies in the left panel. In the right panel, the blue line shows the spectral density of a random walk process whose first difference has a spectral density with a value of one at all frequencies. Filtering these processes with the indicated procedure would lead to amplified and dampened frequencies as shown by the green lines.

of the filters become relevant in this empirical exercise.[17]

Similarly to the previous section, I compare the HP (1600) filter to the version of the regression filter that regresses each variable at date $t + 8$ on the four most recent values as of date t (2-year regression filter). Furthermore, I contrast the HP (400000) filter with a regression of each variable at date $t + 20$ on the four most recent values as of date t (5-year regression filter). Given that policy indicators such as the Basel III credit-to-GDP gap are real time indicators, I report for all filters the detrended series using full sample information as well as information from an expanding sample, mimicking a real time situation.

In Figure 4, I present the detrended series obtained through the HP (1600) filter (black) and the 2-year regression filter (red). In the upper panel, I show the results of US log GDP and in the lower panel, the detrended credit-to-GDP ratio. Furthermore, the left panel uses full sample information, while the right panel presents results derived on an expanding window. Figure 5 presents the same charts but contrasts the detrended components obtained through the HP (400000) filter and the 5-year regression filter. In addition to these charts, I report descriptive statistics in Table 1, such as standard deviations, first-order autocorrelations, but also the correlations of each filter's full sample and real time estimate. Furthermore, Table 2 shows the contemporaneous correlations of the HP-filtered and regression-filtered components using the full sample and the real time estimates.

Using full sample information: As suggested by the power transfer functions of the previous section, the detrended series of the 2-year regression filter have greater volatility

[17]To support the assumption that the theoretical insights of the difference filters in large samples can be transferred to the regression filters for this empirical exercise, I show in Appendix A.2 that the regression filters (also in real time) indeed produce a detrended series that is highly correlated with the respective difference filter. In addition, the exercise indicates that there are some small-sample issues, as, for instance, the correlation of the Δ_8 difference filter and the 2-year regression filter decreases from 0.92 to 0.81 for log GDP when considering the real time estimate instead of the full sample one.

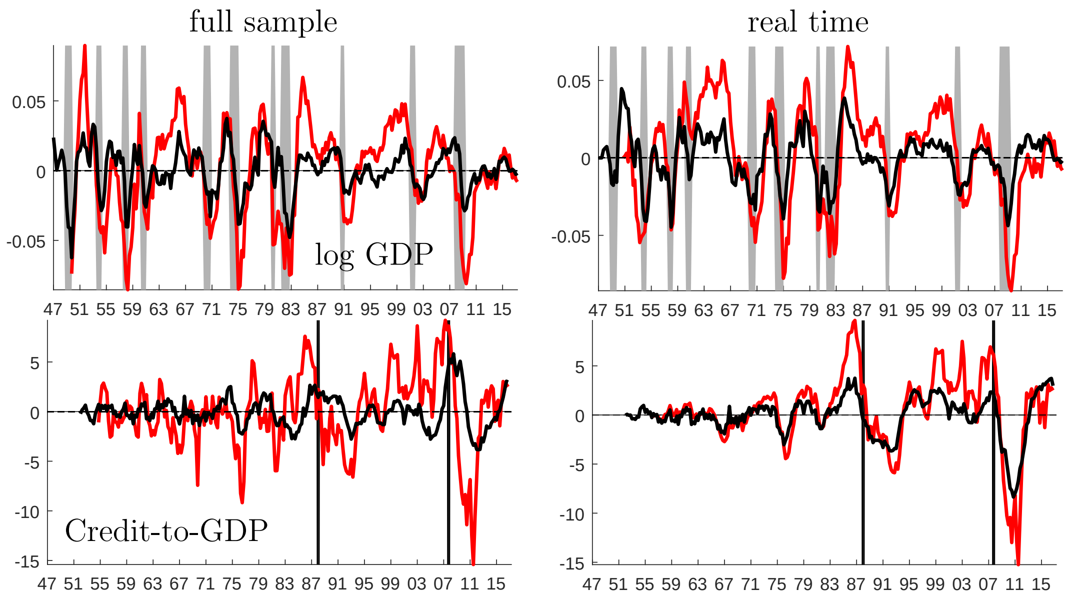

Figure 4: HP (1600) filter (black) versus 2-year regression filter (red)

Notes: Grey area depicts NBER dates of recessions. Black vertical lines indicate the onset of systemic banking crises as defined by Laeven and Valencia (2012).

Table 1: Descriptive statistics

	Standard deviations				First-order autocorrelations				Correlation	
	log GDP (×100)		Credit-to-GDP		log GDP		Credit-to-GDP		*fs* and *rt*	
Filter	*fs*	*rt*	*fs*	*rt*	*fs*	*rt*	*fs*	*rt*	log GDP	credit-to-GDP
HP (1600)	1.50	1.62	1.69	2.19	**0.85**	**0.89**	**0.92**	**0.96**	0.54	0.17
2-year regression	3.25	3.19	3.96	3.83	**0.91**	**0.93**	**0.87**	**0.94**	0.90	0.91
HP (400000)	2.91	3.02	6.04	6.24	**0.96**	**0.96**	**0.99**	**0.99**	0.68	0.70
5-year regression	4.65	4.50	9.05	8.78	**0.95**	**0.96**	**0.96**	**0.98**	0.77	0.90

Notes: *fs* denotes full sample and rt stands for real time. Bold numbers indicate significance at least at the 10% level. Statistics are derived using HAC standard errors. The sample period is 1953Q4 to 2017Q2 for GDP and 1958Q4 to 2016Q4 for the credit-to-GDP ratio, reflecting the largest common sample for each indicator across detrending methods.

than the HP (1600)-filtered ones, for instance, in the case of full sample GDP (see upper left panel of Figure 4). Here, the historical volatility of 2-year regression-filtered GDP is more than double the figure of the HP (1600)-filtered series (3.25 vs. 1.50, see Table 1). Due to the fact that the 2-year regression filter most strongly emphasises frequencies that are longer than typical business cycle frequencies (longer than 8 years) and cancels fluctuations around a 2-year duration, it does not capture all phases as classified by the NBER. For instance, it is a stylised business cycle fact that a typical recession lasts around 1 year (see, for instance, the NBER's Business Cycle Dating Committee, Watson (1994), Harding and Pagan (2002), or – for similar statistics across 17 advanced economies – Jordà et al. (2017)), representing a half cycle that is eliminated by applying this filter.[18] This can be seen, for instance, when considering the full-sample regression-filtered GDP series. It does not indicate the quick contraction and recovery around the 1980 recession, but

[18]The NBER's Business Cycle Dating Committee reports that the average duration of a contraction is 11.1 months for the period 1945-2009.

instead suggests a prolonged contraction. The amplification of cycles that are longer than typical business cycles seems to be more consequential for the full-sample credit-to-GDP ratio. While shorter duration cycles are apparent for the 2-year regression-filtered credit-to-GDP ratio, the classification of phases is primarily driven by medium-term cycles. Such observations are in line with recent empirical evidence. Studies find that the credit-to-GDP ratio has important variation at medium-term frequencies, which is in contrast with GDP (e.g., Galati et al. (2016); Schüler (2018)). These are amplified by the 2-year regression filter and cancelled by the HP (1600) filter. Still, first-order autocorrelations of the detrended indicators, i.e., their persistence, differ only marginally and range between 0.85 and 0.96 (see Table 1).

Table 2: Contemporaneous correlations of HP-filtered and regression-filtered series

Variable Filter	2-year regression *fs*	 *rt*	5-year regression *fs*	 *rt*
log GDP				
HP (1600)	0.72	0.79	–	–
HP (400000)	–	–	0.75	0.84
Credit-to-GDP				
HP (1600)	0.24	0.86	–	–
HP (400000)	–	–	0.69	0.92

Notes: fs denotes full sample and rt stands for real time. The *fs* columns refer to the correlations of both detrended components, i.e., using HP and regression filters, that exploit the entire sample. The *rt* columns refer to the correlations of both detrended components, i.e., using HP and regression filters, that are constructed on an expanding sample. The sample period is 1953Q4 to 2017Q2 for GDP and 1958Q4 to 2016Q4 for the credit-to-GDP ratio, reflecting the largest common sample for each indicator across detrending methods.

Using the medium-term filters, Figure 5 and Table 1 indicate that the historical volatility of the detrended components is, again, larger using the regression filter. However, differences are smaller than in the case of the 2-year regression filter and the HP (1600) filter. For instance, volatilities only differ by a factor of around 1.5 in the case of the full sample credit-to-GDP ratio. The smaller difference in volatilities is in line with the PTFs presented in the previous section. Furthermore, and also in line with the PTFs, the 5-year regression filter leaves more shorter-term fluctuations in the detrended components. For instance, the recovery of the full-sample credit-to-GDP gap around 1995 or after the global financial crisis is faster using the 5-year regression filter. As such, the 5-year regression filter can be argued to be preferable when extracting medium-term cycles, as for instance suggested in Basel III. The 5-year regression filter extracts a range of cycles, rather than focussing on a few frequencies. Due to this property, the regression filter indicates imbalances prior to the global financial crisis much earlier than the HP filter. This is illustrated by comparing the red and black lines in the lower left-hand panel of Figure 5, where the red line reaches a magnitude equal to or higher than at the savings and loan crisis already in 2001Q4, while the black line does so only in 2006Q1. Still, both filters capture the onset of both systemic banking crises (black vertical lines). In the case

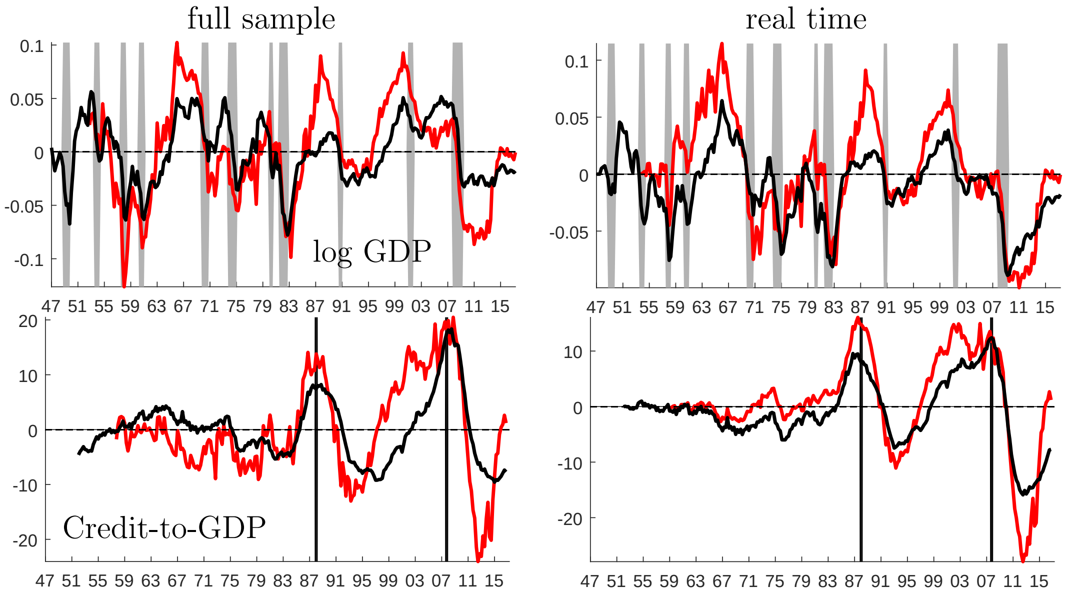

Figure 5: HP (400000) filter (black) versus 5-year regression filter (red)

Notes: Grey area depicts NBER dates of recessions. Black vertical lines indicate the onset of systemic banking crises as defined by Laeven and Valencia (2012).

of GDP, none of the detrended series closely matches NBER recession dates, which is, of course, due to the focus of the medium-term filters on non-business cycles frequencies.

Real time exercise: Differences between the HP-filtered and regression-filtered series become less prominent in the real time exercise. This is suggested by the contemporaneous correlations of the full sample and real time detrended components presented in Table 2. For instance, the correlation of HP (1600)-filtered and 2-year regression-filtered GDP increases from 0.72 (full sample) to 0.79 (real time). The change for the credit-to-GDP ratio is even stronger. In this case, the correlation increases from 0.24 (full sample) to 0.86 (real time). Similar results emerge for the medium-term filters.

Differences diminish when turning from the full sample to the real time estimate mostly due to revisions in the HP-filtered series. This is supported by the correlations between the full sample and real time estimates presented in Table 1. For instance, the two HP (1600)-filtered series are correlated by only 0.54 in case of log GDP, but by 0.90 for the 2-year regression filter. Similarly, for the credit-to-GDP ratio and the medium-term filters, the correlation is only 0.70 for the HP (400000) filter, but 0.90 for the 5-year regression filter. Thus, the regression filter's real time performance can be argued to be more robust than the one of the HP. It produces more similar estimates to the full sample case. This feature can be related to asymmetric formulation of the regression filter and the symmetric formulation of the HP filter.

Altogether, this supports the view that the 5-year regression filter should be preferred for studying medium-term fluctuations, as is the goal of the Basel III credit-to-GDP gap. For instance, also in the real time exercise of the credit-to-GDP ratio, the 5-year regression filter detects imbalances of a magnitude similar or higher than those during the savings and loan crisis as early as 2002Q1, while the HP (400000) filter does so only in 2003Q1. In case of log GDP, both the 2-year regression filter and the HP (1600) filter miss the double dip recession around 1980.

5 Conclusion

In this note, I compare the cyclical properties of the regression filter suggested by Hamilton (2017) to the HP filter. Overall, I find that Hamilton's (2017) regression filter is not subject to the exact same drawbacks as the Hodrick and Prescott (1981, 1997) filter (i. spurious cycles, ii. end-of-sample bias, iii. ad hoc assumptions regarding the smoothing parameter). However, I show that Hamilton's regression filter also modifies the original cyclical structure of a series and is based on ad hoc assumptions that determine the characteristics of the detrended component. While the HP filter suffers from the problem that, for instance, the duration of extracted cycles is to a large extent determined by the smoothing parameter chosen on an ex ante basis, the regression filter completely erases certain fluctuations and emphasises cycles that exceed regular business cycle frequencies. In a real time exercise, the differences between the HP filter and the regression filter are less severe. This, however, is due to revisions in the HP filter when switching to the one-sided version for the real time analysis. As such, the regression filter can be argued to be more robust in a real time setting, suggesting that it has a smaller end-of-sample bias.

Above all, the "correct" filter depends on the researcher's objective, i.e., the feature of the data she would like to focus on. If the objective is to remain agnostic about the importance of the different cyclical characteristics, one should use first differences to detrend a typical economic time series. For difference stationary data, such transformation preserves all dynamics of a series, while the HP filter – but also the regression filter – extract specific frequencies of a time series, masking potentially relevant fluctuations.

References

A'Hearn, B. and U. Woitek (2001). More international evidence on the historical properties of business cycles. *Journal of Monetary Economics 47*, 321–346.

Aikman, D., A. Haldane, and B. Nelson (2015). Curbing the credit cycle. *The Economic Journal 125*, 1072–1109.

Anundson, A., K. Gerdrup, F. Hansen, and K. Kragh-Sørensen (2016). Bubbles and crises: The role of house prices and credit. *Journal of Applied Econometrics 31*, 1291–1311.

Bauer, G. and E. Granziera (2016). Monetary policy, private debt and financial stability risks. *Bank of Canada Staff Working Paper 2016-59*.

Behn, M., C. Detken, T. Peltonen, and W. Schudel (2013). Setting countercyclical capital buffers based on early warning models: Would it work? *ECB Working Paper 1604*.

Borio, C. (2014). The financial cycle and macroeconomics: What have we learnt? *Journal of Banking and Finance 45*, 182–198.

Christiano, L. and T. Fitzgerald (2003). The band pass filter. *International Economic Review 44*, 435–465.

Cogley, T. and J. Nason (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.

De Jong, M. and N. Sakarya (2016). The econometrics of the Hodrick-Prescott filter. *Review of Economics and Statistics 98*, 310–317.

Den Haan, W. (2000). The comovement between output and prices. *Journal of Monetary Economics 46*, 3–30.

Drehmann, M., C. Borio, and K. Tsatsaronis (2012). Characterising the financial cycle: Don't lose sight of the medium term! *BIS Working Paper 380*.

Galati, G., I. Hindrayanto, S. Koopman, and M. Vlekke (2016). Measuring financial cycles with a model-based filter: Empirical evidence for the United States and the euro area. *Economics Letters 145*, 83–87.

Hamilton, J. (2017). Why you should never use the Hodrick-Prescott filter. *The Review of Economics and Statistics: Just Accepted MS*.

Harding, D. and A. Pagan (2002). Dissecting the cycle: A methodological investigation. *Journal of Monetary Economics 49*, 365–381.

Harvey, A. and A. Jaeger (1993). Detrending, stylized facts and the busines cycle. *Journal of Applied Econometrics 8*, 231–247.

Hiebert, P., B. Klaus, T. Peltonen, Y. Schüler, and P. Welz (2014). Capturing the financial cycle in the euro area. *Financial Stability Review: Special Feature B*.

Hodrick, R. and E. Prescott (1981). Postwar U.S. business cycles: An empirical investigation. *Discussion Paper 451, Northwestern University*.

Hodrick, R. and E. Prescott (1997). Postwar U.S. business cycles: An empirical investigation. *Journal of Money, Credit and Banking 29*, 1–16.

Jordà, O., M. Schularick, and A. M. Taylor (2016). The great mortgaging: housing finance, crises and business cycles. *Economic Policy 31*, 107–152.

Jordà, O., M. Schularick, and A. M. Taylor (2017). Macrofinancial History and the New Business Cycle Facts. In M. Eichenbaum and J. Parker (Eds.), *NBER Macroeconomics Annual 2016, Volume 31*, Chapter 3, pp. 213–263. Chicago: University of Chicago Press.

King, R. and S. Rebelo (1993). Low frequency filtering and real business cycles. *Journal of Economic Dynamics and Control 17*, 207–231.

Kuznets, S. (1961). *Capital and the American Economy: Its Formation and Financing.* Princeton, N.J.: Princeton University Press.

Laeven, L. and F. Valencia (2012). Systemic banking crises database: An update. *IMF Working Paper WP/12/163*.

Murray, C. (2003). Cyclical properties of Baxter-King filtered time series. *The Review of Economics and Statistics 85*, 472–476.

Pederson, T. (2001). The Hodrick-Prescott filter, the Slutzky effect, and the distortionary effects of filters. *Journal of Economic Dynamics and Contril 25*, 1081–1101.

Rünstler, G. and M. Vlekke (2016). Business, housing and credit cycles. *ECB Working Paper 1915*.

Schüler, Y. (2018). Detrending and financial cycle facts across G7 countries: Mind a spurious medium term! *ECB Working Paper Series (forthcoming)*.

Schüler, Y., P. Hiebert, and T. Peltonen (2015). Characterising the financial cycle: A multivariate and time-varying approach. *ECB Working Paper 1846*.

Schüler, Y., P. Hiebert, and T. Peltonen (2017). Coherent financial cycles for G-7 countries: Why extending credit can be an asset. *ESRB Working Paper Series No. 43*.

Strohsal, T., C. Proaño, and J. Wolters (2015a). Characterizing the financial cycle: Evidence from a frequency domain analysis. *Deutsche Bundesbank Discussion Paper 22/2015*.

Strohsal, T., C. Proaño, and J. Wolters (2015b). How do financial cycles interact? Evidence for the US and the UK. *SFB 649 Discussion Paper 2015-024*.

Verona, F. (2016). Time-frequency characterization of the U.S. financial cycle. *Economics Letters 144*, 75–79.

Watson, M. (1994). Business-cycle durations and postwar stabilization of the U.S. economy. *American Economic Review 84*, 24–46.

A Appendix

A.1 Power transfer functions of filters

This section lays out the power transfer functions depicted in Section 3. Note that the power transfer functions $H(\omega)$ refer to those relevant for stationary data and $J(\omega)$ to those relevant for difference stationary data. The $J(\omega)$ are derived by $H(\omega)/H^{\Delta}(\omega)$, where $H^{\Delta}(\omega)$ is the power transfer function of the first difference filter.

The PTFs capturing the effects of the HP filters with $\lambda = 1,600$ and $\lambda = 400,000$ on ψ_t are (see King and Rebelo (1993))

$$H^{\mathrm{HP}(\lambda)}(\omega) = \left[\frac{4(1-\cos(\omega))^2}{4(1-\cos(\omega))^2 + 1/\lambda}\right]^2. \tag{11}$$

The PTFs of the difference filters are

$$H^{\Delta}(\omega) = |(1 - e^{-i\omega})|^2 = 2(1 - \cos(\omega)), \tag{12}$$

$$H^{\Delta_8}(\omega) = 2(1 - \cos(8\omega)), \text{ and} \tag{13}$$

$$H^{\Delta_{20}}(\omega) = 2(1 - \cos(20\omega)). \tag{14}$$

The $J(\omega)$ for the HP filter is

$$J^{\mathrm{HP}(\lambda)}(\omega) = H^{\mathrm{HP}(\lambda)}(\omega)[H^{\Delta}(\omega)]^{-1} = \frac{8(1 - \cos(\omega))^3}{(4(1 - \cos(\omega))^2)^2 + 1/\lambda}, \tag{15}$$

and $J(\omega)$s for the difference filters are

$$J^{\Delta}(\omega) = H^{\Delta}(\omega)[H^{\Delta}(\omega)]^{-1} = 1, \tag{16}$$

$$J^{\Delta_8}(\omega) = H^{\Delta_8}(\omega)[H^{\Delta}(\omega)]^{-1} = \frac{1 - \cos(8\omega)}{1 - \cos(\omega)}, \text{ and} \tag{17}$$

$$J^{\Delta_{20}}(\omega) = H^{\Delta_{20}}(\omega)[H^{\Delta}(\omega)]^{-1} = \frac{1 - \cos(20\omega)}{1 - \cos(\omega)}. \tag{18}$$

A.2 Comparison of difference filters and Hamilton's regression filters

Under the assumption of a random walk, the regression filter suggested by Hamilton (2017) reduces to a difference filter in large samples. In my analysis of the theoretical properties, I exploit this fact to shed light on the spectral properties of the regression filter. This Appendix explores whether the close link between the difference filters and regression filters holds for the empirical exercise, that is for actual economic data and given only small samples. For this purpose, Figure 6 shows three different detrended components in each panel: the one obtained by the respective difference filter in purple, the regression filter in solid red, and the real time regression filter in dotted red. For ease of comparison, all series are normalised. Furthermore, Table 3 provides information about the contemporaneous correlation of the detrended components shown in the graphs.

Overall, the results suggest that the differences between the series are marginal as, for instance, the correlations of the difference-filtered series with the others is always above 0.75. It is 0.92 for the 2-year regression filter and the 8-quarter difference filter in the case of GDP and 0.91 for the 5-year regression filter and the 20-quarter difference filter for credit-to-GDP. Still, the exercise points towards small sample issues of the regression filter, as, for example, considering the correlations of log GDP and the 2-year regression filter. In this case, the correlation of the 2-year regression-filtered series with the Δ_8-filtered series decreases from 0.92 in the full sample to 0.81 in the real time exercise.

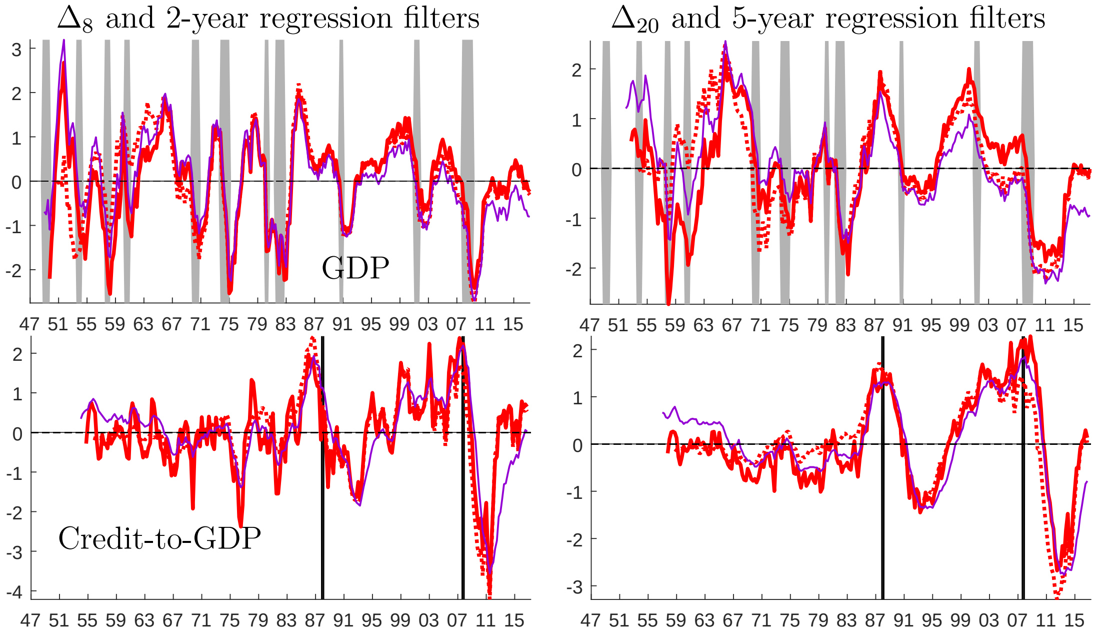

Figure 6: Comparing difference filter (purple), regression filter (solid red), and real time regression filter (dotted red)

Notes: Grey area depicts NBER dates of recessions. Black vertical lines indicate the onset of systemic banking crises as defined by Laeven and Valencia (2012). All series are standardised.

Table 3: Contemporaneous correlation between difference-filtered and regression-filtered series

Variable	Δ_8		Δ_{20}	
Filter	*fs*	*rt*	*fs*	*rt*
log GDP				
2-year regression	0.92	0.81	–	–
5-year regression	–	–	0.81	0.84
Credit-to-GDP				
2-year regression	0.77	0.77	–	–
5-year regression	–	–	0.91	0.88

Notes: fs denotes full sample and rt stands for real time. The *fs* columns refer to the correlations of the regression-filtered components that exploit the entire sample and the difference-filtered series. The *rt* columns refer to the correlations of the regression-filtered components constructed on an expanding sample and the difference-filtered series. Δ_8 and Δ_{20} refers to the 2-year and 5-year difference filters respectively.