Discussion Paper

Deutsche Bundesbank
No 11/2020

On adjusting the one-sided Hodrick-Prescott filter

Elias Wolf
(Freie Universität Berlin)

Frieder Mokinski
(Deutsche Bundesbank)

Yves Schüler
(Deutsche Bundesbank)

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 978–3–95729–680–1 (Printversion)
ISBN 978–3–95729–681–8 (Internetversion)

Non-technical summary

Research question

The two-sided Hodrick-Prescott filter (HP-2s) is a popular tool for detrending macroeconomic time series, such as the gross domestic product (GDP). The one-sided Hodrick-Prescott filter (HP-1s) is a version of this filter, which is often used for predictive tasks, where there are no future values of a time series available yet. Furthermore, Basel III regulations recommend using HP-1s to construct a reference indicator for setting the countercylical capital buffer: the credit-to-GDP gap. HP-1s is usually applied under the assumption that the properties of HP-2s carry over one-to-one to HP-1s.

Contribution

In this study, we explore whether the properties of HP-1s resemble those of HP-2s and find pronounced differences. Therefore, we propose adjustments to HP-1s that align its properties closely to those of HP-2s.

Results

We find pronounced differences in the properties of HP-1s and HP-2s. However, two easy-to-implement adjustments to HP-1s strongly reduce these deviations: a lower value for the smoothing parameter and a multiplicative rescaling of the detrended component. For instance, considering HP-2s with the common choice of 1,600 as the value of the smoothing parameter, the adjusted HP-1s employs a value of 650 instead. Moreover, it rescales the detrended component by a factor of 1.15. Using simulated and empirical data, we illustrate the relevance of these adjustments. For instance, financial cycles may appear 1.7 times more volatile than business cycles, where in fact volatilities differ only marginally.

Nichttechnische Zusammenfassung

Fragestellung

Der zweiseitige Hodrick-Prescott-Filter (HP-2s) ist ein gängiges Verfahren, um makroökonomische Zeitreihen wie das Bruttoinlandsprodukt (BIP) um ihren Trend zu bereinigen. Der einseitige Hodrick-Prescott-Filter (HP-1s) ist eine Variante dieses Filters, die zum Beispiel für Prognoseaufgaben verwendet wird, da in diesem Fall noch keine zukünftigen Werte der Zeitreihe verfügbar sind. Der HP-1s wird unter Basel III vorgeschlagen, um den Indikator zur Ermittlung des antizyklischen Kapitalpuffers zu bestimmen: die Kredit/BIP-Lücke. Gewöhnlich wird der HP-1s unter der Annahme verwendet, dass seine Eigenschaften sich praktisch nicht von denen des HP-2s unterscheiden.

Beitrag

In dieser Studie untersuchen wir, inwieweit die Eigenschaften des HP-1s tatsächlich denen des HP-2s ähneln und finden ausgeprägte Unterschiede. Daher schlagen wir Anpassungen des HP-1s vor, die seine Trendbereinigung jener des HP-2s angleichen.

Ergebnisse

Die Verwendung des HP-1s erzeugt trendbereinigte Komponenten, deren Eigenschaften oft stark von denen des HP-2s abweichen. Durch zwei einfache Anpassungen des HP-1s können diese Unterschiede jedoch drastisch reduziert werden: die Verwendung eines kleineren Werts für den Glättungsparameter und eine Skalierung der trendbereinigten Komponente. Wenn man zum Beispiel einen HP-2s mit dem häufig genutzten Wert 1600 als Glättungsparameter betrachtet, benutzt der entsprechende adjustierte HP-1s statt des Wertes 1600 den Wert 650. Zudem multipliziert er die trendbereinigte Komponente mit einem Faktor von 1,15. Anhand von simulierten und empirischen Daten verdeutlichen wir die Relevanz dieser Anpassungen. Ohne Anpassungen des HP-1s könnten zum Beispiel Finanzzyklen 1,7-mal volatiler erscheinen als Konjunkturzyklen, wenn sich die wahren Volatilitäten nur marginal unterscheiden.

Deutsche Bundesbank Discussion Paper No 11/2020

On adjusting the one-sided Hodrick-Prescott filter*

Elias Wolf† Frieder Mokinski‡ Yves Schüler§

February 14, 2020

Abstract

We show that one should not use the one-sided Hodrick-Prescott filter (HP-1s) as the real-time version of the two-sided Hodrick-Prescott filter (HP-2s): First, in terms of the extracted cyclical component, HP-1s fails to remove low-frequency fluctuations to the same extent as HP-2s. Second, HP-1s dampens fluctuations at all frequencies – even those it is meant to extract. As a remedy, we propose two small adjustments to HP-1s, aligning its properties closely with HP-2s: (1) a lower value for the smoothing parameter and (2) a multiplicative rescaling of the extracted cyclical component. For example, for HP-2s with $\lambda = 1{,}600$ (value of smoothing parameter), the adjusted one-sided HP filter uses $\lambda^* = 650$ and rescales the extracted cyclical component by a factor of 1.1513. Using simulated and empirical data, we illustrate the relevance of the adjustments. For instance, financial cycles may appear 1.7 times more volatile than business cycles, where in fact volatilities differ only marginally.

Keywords: Real-time analysis, detrending, business cycles, financial cycles

JEL classification: C10, E32, E58, G01.

*This paper has benefited from valuable comments and suggestions by Robert De Jong, Alessandra Luati, Elmar Mertens, Christian Schumacher, and Lars Winkelmann. The views expressed in this paper are those of the authors and do not necessarily coincide with the views of the Deutsche Bundesbank or the Eurosystem.

†Contact address: Freie Universität Berlin, Chair of Econometrics, Boltzmannstraße 20, 14195 Berlin, Germany. E-mail: eliaswolf@zedat.fu-berlin.de.

‡Contact address: Deutsche Bundesbank, DG Financial Stability, Wilhelm-Epstein-Straße 14, 60431 Frankfurt am Main, Germany. E-mail: frieder.mokinski@bundesbank.de.

§Contact address: Deutsche Bundesbank, Research Centre, Wilhelm-Epstein-Straße 14, 60431 Frankfurt am Main, Germany. E-mail: yves.schueler@bundesbank.de.

1 Introduction

The one-sided Hodrick and Prescott (1981, 1997) filter (HP-1s) is used as the real-time version of the regular two-sided HP filter (HP-2s), a popular tool for detrending macroeconomic time series. As a "one-sided" or "real-time" filter, HP-1s uses only observations dated t and earlier to filter the time-series observation y_t. By contrast, HP-2s also uses information beyond period t to filter y_t. As a consequence, HP-2s revises its inference on all observations in the sample as new observations become available, whereas the inference of HP-1s on past observations does not change with new observations. HP-1s is attractive for predictive tasks, where it is reasonable to use only information that is available in real time. It is also appealing for policy making, where decisions are faced in real time and where there is a preference not to revise past estimates. Related to this, HP-1s has recently gained popularity through Basel III regulations that recommend the use of HP-1s to construct a so-called credit-to-GDP gap.[1]

It is common to use HP-1s under the implicit assumption that the properties of HP-2s carry over one-to-one, see, for example, Stock and Watson (1999). Other examples include Orphanides and van Norden (2002), Christiano and Fitzgerald (2003), Edge and Meisenzahl (2011), and Hamilton (2018), who compare properties of HP-2s and HP-1s using the same values for the smoothing parameter in both cases. The purpose of this paper is (i) to illustrate that the similar-properties assumption does not hold and (ii) propose adjustments for HP-1s so that it can be used as the real-time version of HP-2s.

Using frequency domain analysis, we show that there are important differences between HP-1s and HP-2s in terms of extracted cyclical components. While HP-2s provides a good approximation to the ideal band pass filter – it eliminates most fluctuations below some frequency threshold (i.e. the trend) and passes most fluctuations above the threshold (i.e. the cycle) without distortions – HP-1s does worse: First, HP-1s fails to eliminate low-frequency fluctuations to the same extent, yielding a cyclical component that is more strongly contaminated with low-frequency fluctuations. Second, HP-1s fails to pass fluctuations above the threshold without distortions, dampening even those fluctuations it is meant to extract.[2]

In order to harmonize the properties of HP-1s with HP-2s, we propose the adjusted one-sided HP filter (HP-1s*). Specifically, we make two small adjustments to HP-1s, minimizing the squared distance of its power transfer function (PTF) with the PTF of HP-2s: (1) a lower value for the smoothing parameter and (2) a multiplicative re-scaling of the cyclical component. For instance, consider HP-2s with 1,600 as the value of the smoothing parameter. Instead of 1,600, HP-1s* uses 650 and rescales the extracted cyclical component by a factor of 1.1513. Table 1 gives an overview of the most commonly used smoothing parameters of HP-2s together with the corresponding parameters for HP-1s*. Furthermore, Table 4 in Appendix A gives a quick overview of adjustment parameters for

[1] The purpose of this indicator is to inform the calibration of the countercyclical capital buffer, a cyclical bank capital requirement. See Basel Committee on Banking Supervision (2010).

[2] There is yet another notable difference between HP-1s and HP-2s. In contrast to HP-2s, HP-1s – like any one-sided filter – shifts phases. This potentially affects the timing relationships of fluctuations at different frequencies. Given that any one-sided filter necessarily introduces phase shifts, we decided not to harmonize properties of HP-1s with HP-2s along this dimension. For a more detailed discussion, please see Section 5 and Appendix D.

a broad range of smoothing parameters of HP-2s.[3]

Table 1: Typical parameters of the adjusted one-sided Hodrick-Prescott filter

	Two-sided filter (HP-2s)	Adjusted one-sided filter (HP-1s*)	
	λ	λ^*	κ
Business cycles (yearly data)	6.25	2.45	1.7962
Business cycles (quarterly data)	1,600	650	1.1513
Financial cycles (quarterly data)	400,000	163,101	1.0360

Notes: λ denotes smoothing parameter of the two-sided HP filter; λ^* and κ denote, respectively, the corresponding smoothing parameter and scaling factor of the adjusted one-sided HP filter. The scaling factor is multiplied by the extracted cyclical component of the one-sided HP filter.

In an application to both simulated and empirical data, we show that our adjustments succeed at harmonizing the properties of HP-1s with HP-2s: The extracted cyclical components resemble one another more closely in terms of their persistence and variability. Furthermore, we provide evidence that adjusting HP-1s can be relevant in empirical analyses: Given discussions on volatile financial cycles (see, for instance, Claessens, Kose, and Terrones (2011, 2012); Aikman, Haldane, and Nelson (2015); Schüler, Hiebert, and Peltonen (2015, 2020)), we show that financial cycles may appear 1.7 times more volatile than business cycles, where in fact volatilities are only marginally different.

2 The Hodrick-Prescott filter

The HP filter decomposes the time series $y = (y_1, \ldots, y_T)'$ into a cyclical component $\psi = (\psi_1, \ldots, \psi_T)'$ and a trend component $\tau = (\tau_1, \ldots, \tau_T)'$:

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

where T denotes sample size.

The regular, two-sided filter: The *two-sided* HP filter (HP-2s) estimates the trend component by solving the following minimization problem:

$$\left\{\widehat{\tau}_{1|T,\lambda}, \ldots, \widehat{\tau}_{T|T,\lambda}\right\} = \arg\min_{\tau_1,\ldots,\tau_t} \left(\sum_{s=1}^{T} (y_s - \tau_s)^2 + \lambda \sum_{s=2}^{T-1} (\tau_{s+1} - 2\tau_s + \tau_{s-1})^2 \right), \tag{2}$$

where λ controls the smoothness of the trend estimate $\widehat{\tau}_{t|T,\lambda}$: The higher its value, the smoother the extracted trend component will be. $\lambda = 1{,}600$ is a common choice to extract business cycle fluctuations in quarterly data.

Notice that we have denoted the trend estimate for period t by $\widehat{\tau}_{t|T,\lambda}$ to illustrate that it depends on the full sample of data $(1, \ldots, T)$ and on the choice of the smoothing parameter λ, see Equation (2). Accordingly, $\widehat{\psi}_{t|T,\lambda}$ is the estimate of the cyclical component obtained as $\widehat{\psi}_{t|T,\lambda} = y_t - \widehat{\tau}_{t|T,\lambda}$.

[3]In addition, software implementing the adjusted HP filter can be downloaded from `https://sites.google.com/site/yvesschueler/research`

The one-sided filter: By contrast, the idea of the *one-sided* HP filter (HP-1s) is to decompose y_t into trend (τ_t) and cycle (ψ_t) based only on observations dated t and earlier (and not beyond t, as with HP-2s). To stress this idea, we denote the corresponding estimates by $\widehat{\tau}_{t|t,\lambda}$ and $\widehat{\psi}_{t|t,\lambda}$. The trend component is extracted by solving the following expression for all values of t:

$$\widehat{\tau}_{t|t,\lambda} = \arg\min_{\tau_t} \left(\min_{\tau_1,\ldots,\tau_{t-1}} \left(\sum_{s=1}^{t} (y_s - \tau_s)^2 + \lambda \sum_{s=2}^{t-1} (\tau_{s+1} - 2\tau_s + \tau_{s-1})^2 \right) \right). \tag{3}$$

This procedure is equivalent to applying HP-2s recursively on an expanding sample and keeping, from each recursion step, only the trend estimate for the latest period. Analogously to HP-2s, the cyclical component is obtained as $\widehat{\psi}_{t|t,\lambda} = y_t - \widehat{\tau}_{t|t,\lambda}$.

2.1 Hodrick-Prescott filter as a linear moving average

Both HP-1s and HP-2s are linear filters. This means that one can express the trend components and the cyclical components as weighted averages of the data. Specifically, the trend components are

$$\text{HP-1s:} \quad \widehat{\tau}_{t|t,\lambda} = \textstyle\sum_{s=1}^{t} w_{t|t,s,\lambda} \cdot y_s \quad = W_{t|t,\lambda}(L) \cdot y_t, \tag{4}$$

$$\text{HP-2s:} \quad \widehat{\tau}_{t|T,\lambda} = \textstyle\sum_{s=1}^{T} w_{t|T,s,\lambda} \cdot y_s \quad = W_{t|T,\lambda}(L) \cdot y_t, \tag{5}$$

where L is the lag operator with $L^k \tau_s = \tau_{s-k}$. $W_{t|t,\lambda}(L) = \sum_{s=1}^{t} w_{t|t,s,\lambda} L^{t-s}$ and $W_{t|T,\lambda}(L) = \sum_{s=1}^{T} w_{t|T,s,\lambda} L^{t-s}$ are linear filter polynomials. Analogously, the cyclical components are

$$\text{HP-1s:} \quad \widehat{\psi}_{t|t,\lambda} = y_t - \textstyle\sum_{s=1}^{t} w_{t|t,s,\lambda} \cdot y_s = (1 - W_{t|t,\lambda}(L)) \cdot y_t \quad = \overline{W}_{t|t,\lambda}(L) \cdot y_t, \tag{6}$$

$$\text{HP-2s:} \quad \widehat{\psi}_{t|T,\lambda} = y_t - \textstyle\sum_{s=1}^{T} w_{t|T,s,\lambda} \cdot y_s = (1 - W_{t|T,\lambda}(L)) \cdot y_t \quad = \overline{W}_{t|T,\lambda}(L) \cdot y_t. \tag{7}$$

The notation of the weights ($w_{t|t,s,\lambda}$, $w_{t|T,s,\lambda}$) illustrates that they depend on the observation t to be filtered, the sample length t (or T in the two-sided case), the position of the weighted observation in the sample s, and the value of λ.

Several papers derive analytical expressions for these filter weights in finite samples (see, for instance, Danthine and Girardin (1989); De Jong and Sakarya (2016); Cornea-Madeira (2017); Hamilton (2018)). We use these insights to derive the frequency domain properties of HP-1s and HP-2s analytically. Specifically, we use the filter polynomials that we summarize in Table 2 for a quick overview.

Table 2: Notation for the filter polynomials of the Hodrick-Prescott filter

Filter	Component	Filter polynomial
HP-1s	Trend ($\widehat{\tau}_{t\|t,\lambda}$)	$W_{t\|t,\lambda}(L)$
HP-1s	Cycle ($\widehat{\psi}_{t\|t,\lambda}$)	$\overline{W}_{t\|t,\lambda}(L)$
HP-2s	Trend ($\widehat{\tau}_{t\|T,\lambda}$)	$W_{t\|T,\lambda}(L)$
HP-2s	Cycle ($\widehat{\psi}_{t\|T,\lambda}$)	$\overline{W}_{t\|T,\lambda}(L)$

2.2 Filtering from a frequency-domain perspective

Below, we give a brief overview of the frequency domain methods used in this paper. Let y_t be a stationary stochastic process with autocovariances $\gamma_k = \text{Cov}(y_t, y_{t-k})$ and define the autocovariance-generating function of y_t as $g_y(z) = \sum_{k=-\infty}^{\infty} \gamma_k z^k$, where z denotes a complex scalar. Evaluating the autocovariance generating function at $z = e^{-i\omega}$ and dividing by 2π yields the spectral density of y_t:

$$S_y(\omega) = \frac{1}{2\pi} g_y(e^{-\mathrm{i}\omega}), \tag{8}$$

where $\mathrm{i} = \sqrt{-1}$. Integrating the spectral density over the interval $[-\pi, \pi]$ gives the variance of y_t, i.e.

$$\text{Var}(y_t) = \int_{-\pi}^{\pi} S_y(\omega) d\omega = 2\int_{0}^{\pi} S_y(\omega) d\omega,$$

where we can interpret the value of $\omega \in [0, \pi]$ as a cycle frequency measured in radians.[4] This suggests that we can decompose the variance of y_t into portions related to movements at different frequencies. For instance, integrating the spectral density over the interval $[0, \omega_1]$ with $\omega_1 < \pi$, i.e. $2\int_0^{\omega_1} S_y(\omega) d\omega$ would give the portion of variance related to movements at frequencies less than or equal to ω_1.

We use the concept of the *power transfer function (PTF)* to study how filtering changes the spectral density of y_t. Assuming that $\mathcal{W}(L)$ is a linear filter polynomial with $\mathcal{W}(L) = \sum_{j=-\infty}^{\infty} w_j L^j$ and absolutely summable polynomial coefficients, it is possible to show that the spectral densities of y_t and the filtered series $x_t = \mathcal{W}(L) y_t$ are related through

$$S_x(\omega) = PTF_{\mathcal{W}}(\omega) \cdot S_y(\omega), \tag{9}$$

where $PTF_{\mathcal{W}}(\omega) = |\mathcal{W}(e^{-\mathrm{i}\omega})|^2$ is the power transfer function of the linear filter polynomial $\mathcal{W}(L)$. $PTF_{\mathcal{W}}(\omega)$ is a non-negative and real-valued scalar function that measures how $\mathcal{W}(L)$ dampens ($PTF_{\mathcal{W}}(\omega) < 1$), passes ($PTF_{\mathcal{W}}(\omega) = 1$), or amplifies ($PTF_{\mathcal{W}}(\omega) > 1$) movements at specific frequencies ω in y_t. In the following, we use the concept of the PTF to study the extent to which different variants of the HP filter succeed at eliminating lower frequencies and preserving higher frequencies.

3 Why adjust the one-sided HP filter?

The top two panels of Figure 1 show the PTFs of HP-1s, $\overline{W}_{t|t,1600}(L)$, and HP-2s, $\overline{W}_{t|T,1600}(L)$, for $\lambda = 1{,}600$. $\lambda = 1{,}600$ is the value of the smoothing parameter regularly used to extract business cycle fluctuations from quarterly data.[5] The PTF of HP-2s is that of a high pass

[4] Given the periodicity of $e^{-i\omega}$, cycles with frequencies higher than two periods are indistinguishable, a phenomenon commonly called the aliasing effect (see, for example, Hamilton (1994)). Hence, the analysis of any PTF is limited to $[0, \pi]$.

[5] To obtain the PTF of the one-sided filter ($PTF_{\overline{W}_{t|t,\lambda^*}}(\omega)$), we compute the filter coefficients of HP-2s at the sample boundary, as given by Hamilton (2018), using a sample size of $T = 1{,}000$. We use this large sample size to avoid problems related to the small sample properties of the filter weights. Subsequently, we cast the resulting filter polynomial into the frequency domain using the Finite Fourier Transform (FFT). For the PTF of the two-sided HP filter ($PTF_{\overline{W}_{t|T,\lambda}}(\omega)$), we use the large sample results given in

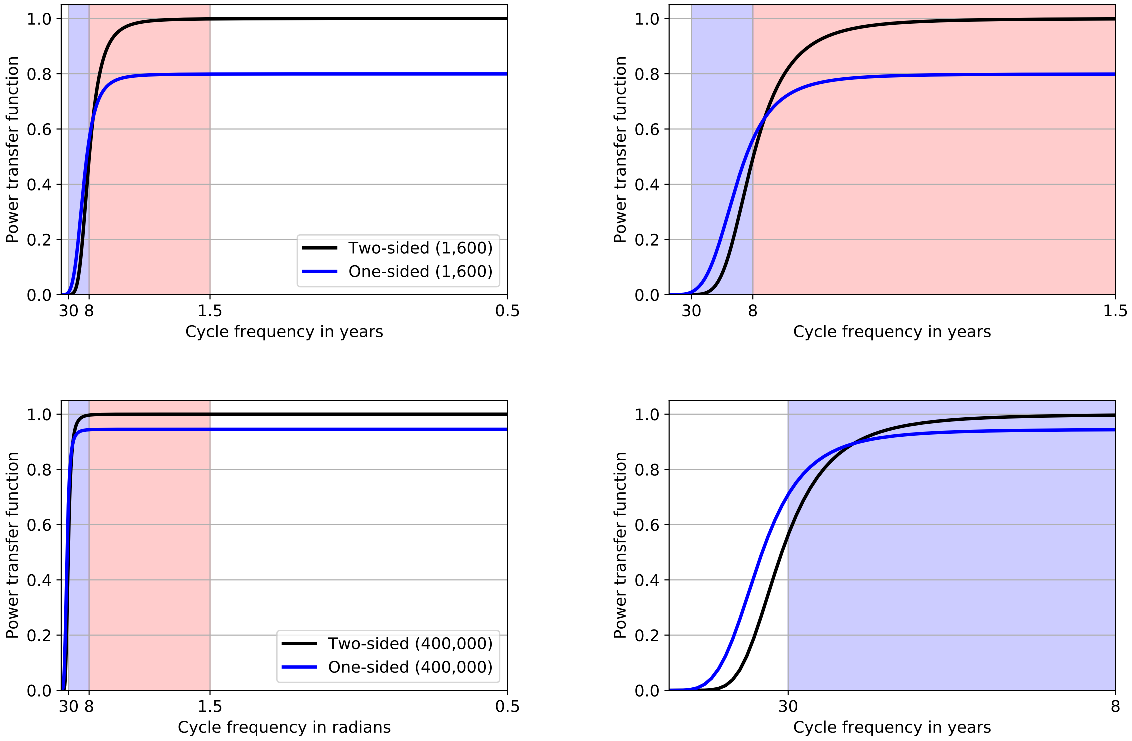

Figure 1: PTFs of HP-1s and HP-2s for $\lambda = 1{,}600$ (top) and $\lambda = 400{,}000$ (bottom)

Notes: Panels depict power transfer functions of HP-1s ("One-sided") and HP-2s ("Two-sided") for different values of the smoothing parameter λ. Top panels refer to $\lambda = 1{,}600$, bottom panels refer to $\lambda = 400{,}000$. Panels to the left and to the right differ in domains. Red shaded areas highlight business cycle frequencies (1.5 to 8 years), purple areas highlight financial cycle frequencies (8 to 30 years).

filter: Higher-frequency fluctuations in the range from 8 to 0.5 years pass the filter with only moderate or virtually no dampening. Lower frequency fluctuations, by contrast, are dampened almost fully: At a cycle length of approximately 17 years, for instance, its PTF reaches a value of 0.01, meaning that the filter dampens variations at this frequency by 99 percent.

The PTF of HP-1s differs in two respects. First, to the left of the two curves' intersection, the PTF of HP-1s runs above that of HP-2s. This implies that HP-1s fails to dampen lower frequencies to the same extent as HP-2s. In fact, differences are relatively pronounced: At a cycle length of 17 years, for instance, the value of the PTF of HP-1s is 0.08, which is eight times the value of the PTF of HP-2s. As a consequence, cyclical components extracted using HP-1s feature low-frequency fluctuations to a much larger extent than cyclical components extracted using HP-2s. Put differently, they tend to be more persistent.[6] Clearly, this is an undesirable property of HP-1s, as it implies that its output is more strongly contaminated with the fluctuations that one aims to remove.

Second, to the right of the two curves' intersection, the PTF of HP-1s starts to run

King and Rebelo (1993).

[6]We use the word "tend" because persistence, for example, as measured by first-order autocorrelation, reflects only a proxy for cycle length. First-order autocorrelation crucially depends on the actual DGP of the series, such as the presence of a unit root.

horizontally at a level of approximately 0.8. The fact that this value is smaller than one implies that HP-1s dampens higher frequencies. As a consequence, cyclical components extracted using HP-1s feature higher-frequency fluctuations to a smaller extent than they are present in the original data. By contrast, these fluctuations pass HP-2s without dampening, as its PTF approaches a value of one as the frequency rises. This is yet another drawback of HP-1s compared with HP-2s: HP-1s dampens precisely the fluctuations that one aims to extract.

Jointly, these two differences imply that the variability of filtered series is likely to differ between HP-1s and HP-2s. Yet the direction is not clear and depends on the spectral density of the series we intend to filter. We elaborate on this in the empirical part of the paper (see Section 5).

Finally, the bottom two panels of Figure 1 show the PTFs of HP-1s and HP-2s for $\lambda =$ 400,000, which is the parameter value recommended in Basel III regulations to construct the credit-to-GDP gap. The respective PTFs resemble those for $\lambda = 1{,}600$ although there are quantitative differences: Both the lack of dampening of lower frequencies and the excessive dampening of high frequencies that occur when using HP-1s appear less pronounced for $\lambda = 400{,}000$ than for $\lambda = 1{,}600$. This suggests that differences between HP-1s and HP-2s diminish as λ grows larger.

In this section, we have shown that the PTFs of HP-1s and HP-2s differ and that differences are more pronounced for small values of the smoothing parameter λ. So why are these differences important? The reason is that HP-1s is regularly used as the real-time version of HP-2s, assuming that the two filters have the same properties. However, this is not true, as we show above. So what can we do about it? In the following, we propose two adjustments to HP-1s that harmonize its properties with HP-2s: First, in order to eliminate the (relatively constant) dampening of higher frequency fluctuations, we rescale the cyclical component. Second, to harmonize PTFs in the range of lower frequencies, we choose a lower value of λ for HP-1s than for HP-2s.

4 The adjusted one-sided HP filter

The filter polynomial for the cyclical component of the *adjusted one-sided HP filter* (HP-1s*) is:

$$\widetilde{W}_{t|t,\lambda}(L) = \kappa \cdot \overline{W}_{t|t,\lambda^*}(L), \tag{10}$$

where we use $\widetilde{W}_{t|t,\lambda}(L)$ to denote the filter polynomial of HP-1s* that is harmonized with the two-sided HP filter (HP-2s) with smoothing parameter λ, $\kappa = k(\lambda)$ is the scaling factor, and $\overline{W}_{t|t,\lambda^*}(L)$ is the standard filter polynomial of the one-sided HP filter (HP-1s) with the adjusted smoothing parameter $\lambda^* = l(\lambda)$. The power transfer function of HP-1s* is given by

$$PTF_{\widetilde{W}_{t|t,\lambda}}(\omega) = \kappa^2 \cdot PTF_{\overline{W}_{t|t,\lambda^*}}(\omega).$$

To clarify notation, consider HP-1s* for $\lambda = 1,600$. As we show below, in this case we have

$$\widetilde{W}_{t|t,1600}(L) = 1.1513 \cdot \overline{W}_{t|t,650}(L),$$

i.e. HP-1s* uses the weight polynomial of the unadjusted one-sided HP filter (HP-1s) with smoothing parameter $\lambda^* = 650$ and scales the weights by the factor $\kappa = 1.1513$.[7]

In general, we obtain $\kappa = k(\lambda)$ and $\lambda^* = l(\lambda)$ by solving the following minimization problem

$$\min_{\kappa,\lambda^*} \left(\int_0^{\pi} \left(PTF_{\overline{W}_{t|T,\lambda}}(\omega) - \kappa^2 \cdot PTF_{\overline{W}_{t|t,\lambda^*}}(\omega) \right)^2 d\omega \right). \tag{11}$$

We thus minimize the squared distance between the PTF of HP-2s with smoothing parameter λ and the PTF of HP-1s*. This strategy is equivalent to the approach used by Baxter and King (1999) in finding an approximate band pass filter for economic time series.

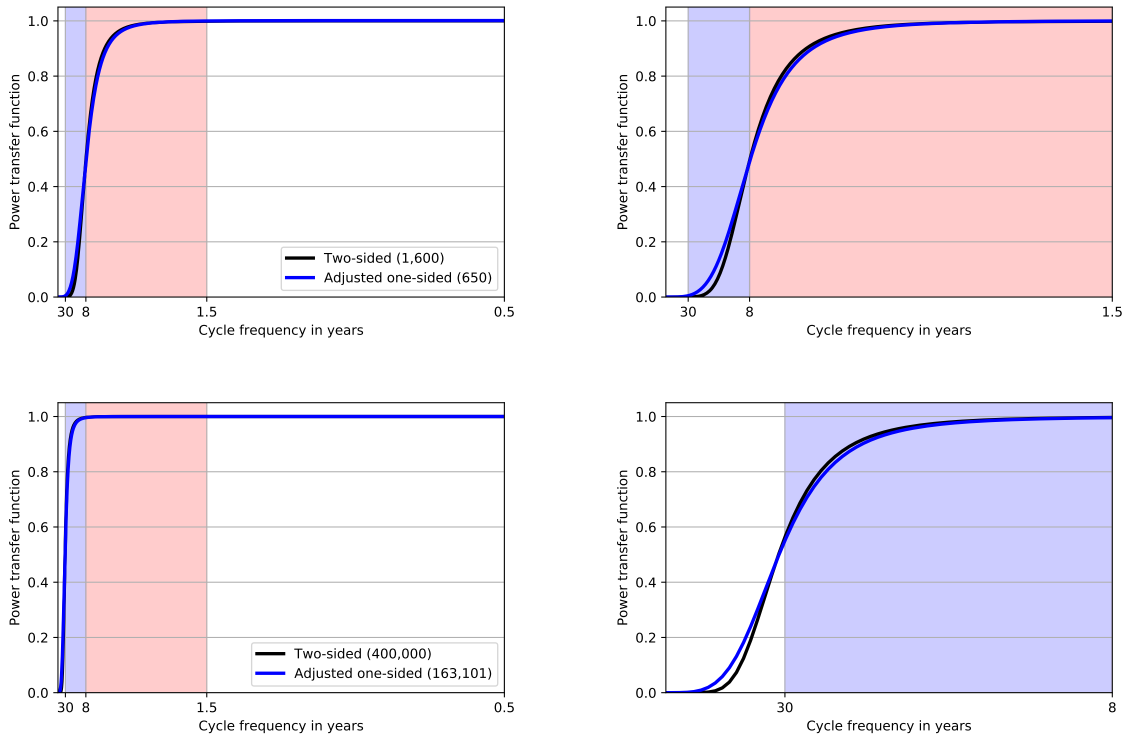

Figure 2: PTFs of HP-1s* and HP-2s for $\lambda = 1{,}600$ (top) and $\lambda = 400{,}000$ (bottom)

Notes: Panels depict power transfer functions of HP-1s* ("Adjusted one-sided") and HP-2s ("Two-sided") for different values of the smoothing parameter λ. Top panels refer to $\lambda = 1{,}600$, bottom panels refer to $\lambda = 400{,}000$. Panels to the left and to the right differ in domains. Red shaded areas highlight business cycle frequencies (1.5 to 8 years), purple areas highlight financial cycle frequencies (8 to 30 years).

Figure 2 shows the PTFs of HP-1s* and HP-2s. Though some visible differences remain between the two PTFs, these are clearly smaller than for HP-1s (see Figure 1). Notably, in the immediate surroundings of the intersection of the two PTF curves, HP-1s* continues to show a slight lack of dampening at lower frequencies (left of intersection)

[7] Table 1 in Section 1 provides adjustment parameters for the most commonly used values of the smoothing parameter λ of HP-2s. For a more comprehensive collection of values, see Appendix A. In addition, software implementing the adjusted HP filter can be downloaded from `https://sites.google.com/site/yvesschueler/research`

and a slight excess of dampening at higher frequencies (right of intersection). Despite these differences, in terms of the proximity of PTFs, HP-1s* is a more closely harmonized real-time version of HP-2s than HP-1s.

Why do we adjust HP-1s in the manner of Equation (10)? Suppose we adjust only the smoothing parameter λ^* but apply no scaling factor, i.e. set $\kappa = 1$ in Equation (10). This adjustment does not produce a good fit, because a change in the smoothing parameter alone always has opposing effects on the degree to which PTFs are harmonized in the range of higher frequencies vs. the range of lower frequencies. This can be seen from Figure 1: Choosing a higher value of λ reduces the undesirable dampening of higher frequencies but, at the same time, increases the undesirable lack of dampening at lower frequencies. By contrast, in the specification of Equation (10), the scaling factor κ eliminates the dampening of higher frequencies, where the PTF is almost horizontal, whereas the adjusted smoothing factor λ^* harmonizes PTFs in the low-frequency band.

5 Applying the adjusted one-sided HP filter

Below, we apply HP-1s* to both simulated and actual data and show that it extracts cyclical components that closely resemble those from HP-2s in terms of persistence and volatility. By contrast, we find no improvement over HP-1s in terms of contemporaneous correlation with the cyclical component extracted using HP-2s. Interestingly, this correlation can be raised by adjusting the smoothing parameter in the opposite direction, i.e. we find the highest correlations for smoothing parameter values greater than those of HP-2s. We suspect that this finding relates to the fact that HP-1s induces phase shifts, and that these phase shifts change with the value of the smoothing parameter λ.

Specifically, we use HP-1s*, HP-1s, and HP-2s to detrend four series with distinct properties:

1. 3,000 randomly sampled observations from a **white noise** process with $y_t = \varepsilon_t$ and $\varepsilon_t \sim N(0, 1)$.

   In a white noise process, all frequencies contribute equally to the overall variance, i.e. the spectral density is uniform.

2. 3,000 randomly sampled observations from a **random walk** process with $y_t = y_{t-1} + \varepsilon_t$ and $\varepsilon_t \sim N(0, 1)$.

   This data-generating process is a good representation of the time series behaviour of many macroeconomic variables (see, for instance, Hamilton (2018)). In a random walk process, the lower the frequency, the higher its contribution to the overall variance, i.e. the spectral density falls monotonically for $\omega \in [0, \pi]$.

3. The natural logarithm of **quarterly US real GDP** for the period 1952Q2 to 2018Q3, retrieved from FRED, and the natural logarithm of **yearly US real GDP** for the years 1880 to 2016, retrieved from the Jordà-Schularick-Taylor Macrohistory Database (see Jordà, Schularick, and Taylor (2017)).

   The distinct feature of real GDP is that business cycle frequencies (1.5 to eight years) contribute heavily to its overall variance.

4. The **quarterly US credit-to-GDP ratio** for the period 1952Q2 to 2018Q3, retrieved from the BIS website, and the **yearly US credit-to-GDP ratio** for the years 1880 to 2016, obtained from the Jordà-Schularick-Taylor Macrohistory Database.[8]

   The distinct feature of the credit-to-GDP ratio is that medium-term frequencies (eight to 30 years) contribute heavily to its overall variance.[9]

We set the value of the smoothing parameter λ for HP-2s to

- either $\lambda = 6.25$, $\lambda = 1{,}600$, or $\lambda = 400{,}000$ for the two simulated series in order to demonstrate how filter properties change with the value of the smoothing parameter;
- $\lambda = 1{,}600$ and $\lambda = 6.25$ respectively for quarterly and yearly real GDP, which are the values typically used to extract business cycle and higher frequencies in quarterly and yearly data;
- $\lambda = 400{,}000$ and $\lambda = 1{,}562.25$ respectively for quarterly and yearly US credit-to-GDP ratio, which are the values typically used to extract financial cycle and higher frequencies in quarterly and yearly data.[10]

Table 3 reports summary statistics for the detrended series for each of the three filters, i.e. HP-2s, HP-1s, and HP-1s*.[11] The block of columns depicting standard deviations shows that HP-1s* more closely resembles HP-2s in terms of the variability of the detrended series. Not in a single case do we find that the discrepancy in standard deviations is greater in absolute terms for HP-1s* than for HP-1s. That said, for some series there are still notable differences in standard deviations between HP-1s* and HP-2s. This is because PTFs are not fully harmonized, and the degree of harmonization varies over frequencies. The consequences of this are more severe for series that have important variation at frequencies that are not well harmonized.

Table 3 also reports autocorrelations, i.e. the correlation of each detrended component and its own first lag. This statistic offers one way of summarizing the extent to which the different filters extract a similar mix of frequencies. Based on the fact that HP-1s* has a PTF that is harmonized with that of HP-2s, we would expect only minor differences between the two. As expected, differences are small with somewhat higher autocorrelations for HP-1s*, which is a consequence of the fact that HP-1s* fails to dampen lower frequency fluctuations as strongly as HP-2s. Therefore, the extracted cyclical component of HP-1s* has a higher share of persistent fluctuations. This is even more pronounced for HP-1s, which produces a considerably more persistent cyclical component than both HP-1s* and HP-2s. The table also suggests that differences in autocorrelations are bigger for smaller values of λ (see, for instance, the yearly US real GDP series).

[8] Note that definitions differ somewhat for the quarterly and yearly credit-to-GDP ratio series: Whereas credit represents all loans to the non-financial private sector in the yearly series, the quarterly series uses a more comprehensive definition that also includes debt securities.

[9] See, for example, Galati, Hindrayanto, Koopman, and Vlekke (2016); Schüler (2018, 2019).

[10] The smoothing parameters 6.25 and 1,562.5 reflect the yearly counterpart of 1,600 and 400,000 in the context of a quarterly sampling frequency (see Ravn and Uhlig (2002)).

[11] Appendix B shows the filtered series.

Table 3: Extracted cyclical components: Standard deviations, autocorrelations, and correlations with the two-sided filter

	Standard deviation			Autocorrelation			Correlation with HP-2s		
λ	HP-2s	HP-1s*	HP-1s	HP-2s	HP-1s*	HP-1s	HP-2s	HP-1s*	HP-1s
DGP 1: $y_t = \varepsilon_t$									
6.25	0.83	0.84	0.55	-0.30	-0.29	-0.23	1.00	0.79	0.84
1,600	0.96	0.96	0.86	-0.05	-0.05	-0.04	1.00	0.94	0.95
400,000	0.99	0.99	0.96	0.01	0.01	0.01	1.00	0.98	0.98
DGP 2: $y_t = y_{t-1} + \varepsilon_t$									
6.25	0.62	0.65	0.47	0.11	0.16	0.29	1.00	0.50	0.58
1,600	1.29	1.34	1.36	0.72	0.74	0.80	1.00	0.47	0.54
400,000	2.72	2.79	3.01	0.93	0.94	0.95	1.00	0.49	0.58
Quarterly US real GDP									
1,600	1.50	1.53	1.53	0.85	0.86	0.89	1.00	0.44	0.54
Yearly US real GDP									
6.25	3.56	3.73	2.80	0.34	0.39	0.50	1.00	0.31	0.44
Quarterly US credit-to-GDP ratio									
400,000	5.90	5.79	6.12	0.99	0.99	0.99	1.00	0.54	0.65
Yearly US credit-to-GDP ratio									
1,562.50	4.37	4.59	4.75	0.87	0.88	0.91	1.00	0.36	0.46

Notes: The table compares properties of cyclical components extracted using the two-sided HP filter (columns 'HP-2s'), the adjusted one-sided HP filter (columns 'HP-1s*'), and the unadjusted one-sided HP filter (columns 'HP-1s'). Column λ gives the value of the smoothing parameter used with the two-sided HP filter and the unadjusted one-sided HP filter. The adjusted one-sided HP filter uses the corresponding smoothing and scaling parameter given in Table 1 and Table 4.

Finally, Table 3 reports correlations with the extracted cyclical component of HP-2s. This statistic offers a perspective on the extent to which two filters extract the same signal contemporaneously. This correlation depends not only on a filter's PTF but also on its phase shift, i.e. the extent to which the filter dislocates fluctuations at different frequencies in the time domain. This phase shift is an unavoidable property of any one-sided filter (see Priestley (2001)), and it has not been targeted by our adjustment.[12] As the one-sided filter's phase shift is more pronounced for smaller values of λ (see Figure 5 in Appendix D), we would expect the HP-1s* to perform somewhat worse than HP-1s in terms of this correlation. Indeed, Table 3 shows exactly that.

Table 3 offers an interesting insight for the debate on the relative variability of financial

[12]Figure 5 in Appendix D shows the phase shift as a function of cycle frequency and illustrates how the phase shift of the one-sided HP filter is closer to zero over a wide range of frequencies for higher values of λ. In a robustness exercise in Table 5 in Appendix C, we show that the correlation can be raised relative to the unadjusted one-sided filter by adjusting the smoothing parameter in the opposite direction, i.e. by choosing a higher value than for the two-sided filter. For each of the series considered in this section, it lists the value of λ that maximizes the correlation with the two-sided filter's cyclical component along with the respective correlations. For purposes of comparison, the table also features the same correlation for the unadjusted one-sided filter. Potential improvements in correlations can be sizable. For instance, in the case of the quarterly credit-to-GDP ratio, the correlation increases from 0.65 for the unadjusted filter to 0.86 for the one-sided filter with optimized smoothing parameter value. That said, this improvement in correlation comes at the cost of a much higher contamination of the extracted cyclical component with exactly those low-frequency fluctuations that the filter is meant to eliminate.

versus business cycles (see, among others, Claessens et al. (2012); Borio (2014); Aikman et al. (2015)). Suppose we measure these variabilities by the standard deviations of the cyclical component extracted from the yearly US credit-to-GDP ratio (financial cycle) and yearly US real GDP (business cycle) as shown in Table 3. This allows us to base our insights on a time series that extends from 1880 to 2016 and thus covers a number of financial crisis episodes. In terms of relative variability, both HP-1s* and HP-2s suggest that financial cycles are about 1.2 times as variable as business cycles. By contrast, HP-1s would suggest a far higher ratio of 1.7. Clearly, such differences are material and they may matter, for instance, in the calibration of macro-econometric models and in the debate on financial market regulation, whereby more volatile financial cycles may call for a stricter regulation.

6 Conclusion

Should the cyclical component obtained from the standard one-sided HP filter be used as the real-time version of the two-sided HP filter's cyclical component? This paper argues that it should not. The reason is that important properties of the standard one-sided filter are quite different from the two-sided filter. Namely, the standard one-sided filter (1) fails to remove low-frequency fluctuations to the same extent as the two-sided filter and (2) has the undesirable feature of dampening exactly those fluctuations that one wishes to extract. As a remedy, this paper proposes two easy-to-implement adjustments to the one-sided filter: (1) a lower smoothing parameter and (2) a multiplicative rescaling of the cyclical component. Jointly, these two adjustments address the above-mentioned problems of the standard one-sided HP filter. This is confirmed in applications of the adjusted one-sided HP filter to both simulated and empirical data.

References

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

Basel Committee on Banking Supervision (2010). Guidance for national authorities operating the countercyclical capital buffer. *Bank of International Settlements*.

Baxter, M. and R. King (1999). Measuring business cycles: Approximate band-pass filters for economic time series. *The Review of Economics and Statistics 81*, 575–593.

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.

Claessens, S., M. Kose, and M. Terrones (2011). Financial cycles: What? How? When? *IMF Working Paper WP/11/76*.

Claessens, S., M. Kose, and M. Terrones (2012). How do business and financial cycles interact? *Journal of International Economics 87*, 178–190.

Cornea-Madeira, A. (2017). The explicit formula for the Hodrick-Prescott filter in a finite sample. *The Review of Economics and Statistics 99*, 314–318.

Danthine, J.-P. and M. Girardin (1989). Business cycles in Switzerland: A comparative study. *European Economic Review 33*, 31–50.

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

Edge, R. and R. Meisenzahl (2011). The unreliability of credit-to-GDP ratio gaps in real time: Implications for countercyclical capital buffers. *International Journal of Central Banking 7*, 261–298.

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. (1994). *Time Series Analysis*. Princeton: Princeton University Press.

Hamilton, J. (2018). Why you should never use the Hodrick-Prescott filter. *The Review of Economics and Statistics: Just Accepted MS 100*, 831–843.

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. Taylor (2017). Macrofinancial history and the new business cycle facts. In M. Eichenbaum and J. A. Parker (Eds.), *NBER Macroeconomics Annual 2016*. 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.

Orphanides, A. and S. van Norden (2002). The unreliability of output gap estimates in real time. *Review of Economics and Statistics 84*, 569–583.

Priestley, M. (2001). *Spectral Analysis and Time Series*. London: Academic Press.

Ravn, M. and H. Uhlig (2002). On adjusting the Hodrick-Prescott filter for the frequency of observations. *Review of Economics and Statistics 84*, 371–380.

Schüler, Y. (2018). Detrending and financial cycle facts across G7 countries: Mind a spurious medium term! *ECB Working Paper 2138*.

Schüler, Y. (2019). How should we filter economic time series? *SSRN*.

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 (2020). Financial cycles: Characterisation and real-time measurement. *Journal of International Money and Finance, forthcoming*.

Stock, J. and M. Watson (1999). Forecasting inflation. *Journal of Monetary Economics 44*, 293–335.

A Adjustment parameters for the one-sided Hodrick-Prescott filter

Table 4: Given a value of λ, i.e. the HP-2s smoothing parameter, HP-1s* uses the value of λ^* and κ

λ	λ^*	κ	λ	λ^*	κ	λ	λ^*	κ	λ	λ^*	κ
1.00	0.35	2.7174	55.00	22.15	1.3921	1000	406	1.1718	55000	22411	1.0598
1.25	0.45	2.5372	57.50	23.16	1.3869	1250	508	1.1617	57500	23431	1.0591
1.50	0.55	2.4111	60.00	24.17	1.3820	**1600**	**650**	**1.1513**	60000	24450	1.0584
1.75	0.65	2.3165	62.50	25.19	1.3774	1750	711	1.1477	62500	25469	1.0578
2.00	0.75	2.2420	65.00	26.20	1.3730	2000	813	1.1425	65000	26488	1.0572
2.25	0.85	2.1814	67.50	27.21	1.3688	2250	915	1.1381	67500	27507	1.0567
2.50	0.95	2.1308	70.00	28.23	1.3648	2500	1017	1.1342	70000	28526	1.0562
2.75	1.04	2.0877	72.50	29.24	1.3610	2750	1118	1.1309	72500	29546	1.0557
3.00	1.14	2.0503	75.00	30.25	1.3574	3000	1220	1.1279	75000	30565	1.0552
3.25	1.24	2.0176	77.50	31.26	1.3540	3250	1322	1.1252	77500	31584	1.0547
3.50	1.34	1.9885	80.00	32.28	1.3506	3500	1424	1.1227	80000	32603	1.0543
3.75	1.44	1.9625	82.50	33.29	1.3475	3750	1525	1.1205	82500	33623	1.0538
4.00	1.55	1.9390	85.00	34.30	1.3444	4000	1627	1.1184	85000	34642	1.0534
4.25	1.65	1.9177	87.50	35.32	1.3415	4250	1729	1.1166	87500	35661	1.0530
4.50	1.75	1.8981	90.00	36.33	1.3387	4500	1831	1.1148	90000	36680	1.0527
4.75	1.85	1.8802	92.50	37.35	1.3359	4750	1933	1.1132	92500	37700	1.0523
5.00	1.95	1.8636	95.00	38.36	1.3333	5000	2034	1.1116	95000	38719	1.0519
5.25	2.05	1.8483	97.50	39.37	1.3308	5250	2136	1.1102	97500	39738	1.0516
5.50	2.15	1.8339	100	40	1.3283	5500	2238	1.1089	100000	40758	1.0512
5.75	2.25	1.8206	125	51	1.3077	5750	2340	1.1076	125000	50951	1.0484
6.00	2.35	1.8080	150	61	1.2919	6000	2442	1.1064	150000	61145	1.0462
6.25	**2.45**	**1.7962**	175	71	1.2792	6250	2543	1.1053	175000	71340	1.0444
6.50	2.55	1.7851	200	81	1.2688	6500	2645	1.1042	200000	81534	1.0429
6.75	2.65	1.7746	225	91	1.2599	6750	2747	1.1032	225000	91730	1.0416
7.00	2.75	1.7647	250	101	1.2523	7000	2849	1.1022	250000	101925	1.0405
7.25	2.85	1.7552	275	111	1.2456	7250	2951	1.1012	275000	112120	1.0396
7.50	2.95	1.7463	300	122	1.2397	7500	3053	1.1003	300000	122316	1.0387
7.75	3.05	1.7377	325	132	1.2343	7750	3154	1.0995	325000	132512	1.0379
8.00	3.15	1.7296	350	142	1.2295	8000	3256	1.0986	350000	142708	1.0372
8.25	3.25	1.7218	375	152	1.2251	8250	3358	1.0978	375000	152904	1.0366
8.50	3.35	1.7143	400	162	1.2211	8500	3460	1.0971	**400000**	**163101**	**1.0360**
8.75	3.45	1.7072	425	172	1.2174	8750	3562	1.0964	425000	173297	1.0354
9.00	3.55	1.7004	450	183	1.2140	9000	3664	1.0956	450000	183494	1.0349
9.25	3.65	1.6938	475	193	1.2108	9250	3765	1.0950	475000	193691	1.0344
9.50	3.76	1.6875	500	203	1.2078	9500	3867	1.0943	500000	203887	1.0340
9.75	3.86	1.6814	525	213	1.2051	9750	3969	1.0937	525000	214084	1.0336
10.00	3.96	1.6755	550	223	1.2024	10000	4071	1.0930	550000	224281	1.0332
12.50	4.96	1.6265	575	233	1.2000	12500	5089	1.0878	575000	234478	1.0328
15.00	5.97	1.5897	600	243	1.1976	15000	6108	1.0837	600000	244675	1.0324
17.50	6.98	1.5607	625	254	1.1954	17500	7127	1.0804	625000	254873	1.0321
20.00	7.99	1.5370	650	264	1.1933	20000	8145	1.0776	650000	265070	1.0318
22.50	9.00	1.5172	675	274	1.1913	22500	9164	1.0753	675000	275267	1.0315
25.00	10.01	1.5001	700	284	1.1894	25000	10183	1.0733	700000	285465	1.0312
27.50	11.02	1.4853	725	294	1.1876	27500	11202	1.0715	725000	295662	1.0309
30.00	12.03	1.4723	750	304	1.1859	30000	12221	1.0699	750000	305860	1.0307
32.50	13.04	1.4606	775	315	1.1842	32500	13240	1.0685	775000	316057	1.0304
35.00	14.05	1.4502	800	325	1.1826	35000	14259	1.0672	800000	326255	1.0302
37.50	15.06	1.4407	825	335	1.1811	37500	15278	1.0660	825000	336453	1.0299
40.00	16.08	1.4320	850	345	1.1796	40000	16297	1.0649	850000	346650	1.0297
42.50	17.09	1.4241	875	355	1.1782	42500	17316	1.0639	875000	356848	1.0295
45.00	18.10	1.4167	900	365	1.1768	45000	18335	1.0629	900000	367046	1.0293
47.50	19.11	1.4099	925	376	1.1755	47500	19354	1.0621	925000	377244	1.0291
50.00	20.12	1.4036	950	386	1.1743	50000	20373	1.0612	950000	387442	1.0289
52.50	21.14	1.3977	975	396	1.1730	52500	21392	1.0605	975000	397640	1.0287
									1000000	407838	1.0285

Notes: λ denotes the smoothing parameter of the two-sided HP filter. λ^* is the corresponding adjusted smoothing parameter, used as an input to the one-sided HP filter. κ is the scaling factor by which the extracted cyclical component of the one-sided HP filter is multiplied. For instance, consider HP-2s with λ =1,600 (column three, in bold). Instead of 1,600, the adjusted HP-1s employs a smoothing parameter of value 650 (λ^*). In parallel, it multiplicatively rescales the extracted cyclical component by a factor of 1.1513. We also offer software implementing the adjusted HP-1s for given HP-2s smoothing parameter (see `https://sites.google.com/site/yvesschueler/research`)

B Hodrick-Prescott-filtered data

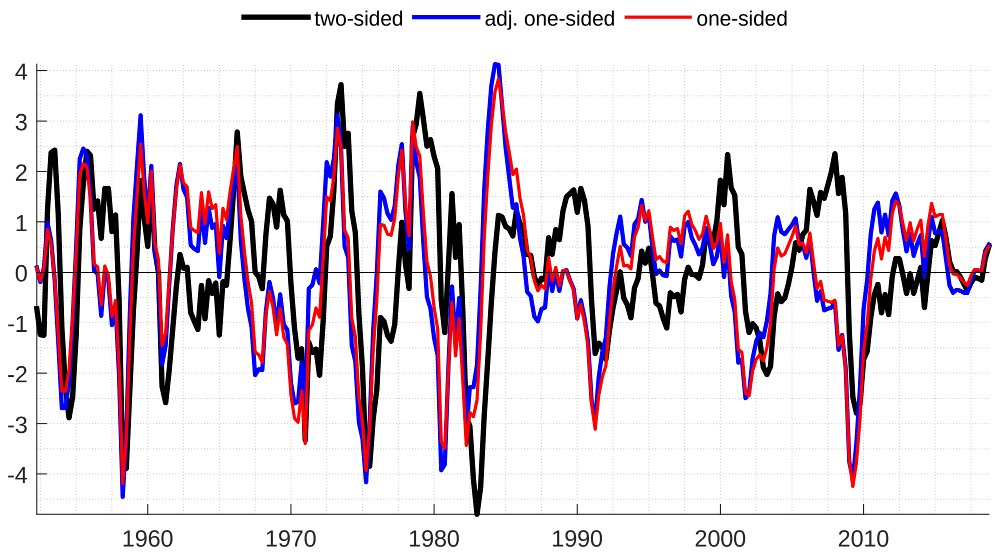

(a) Real GDP

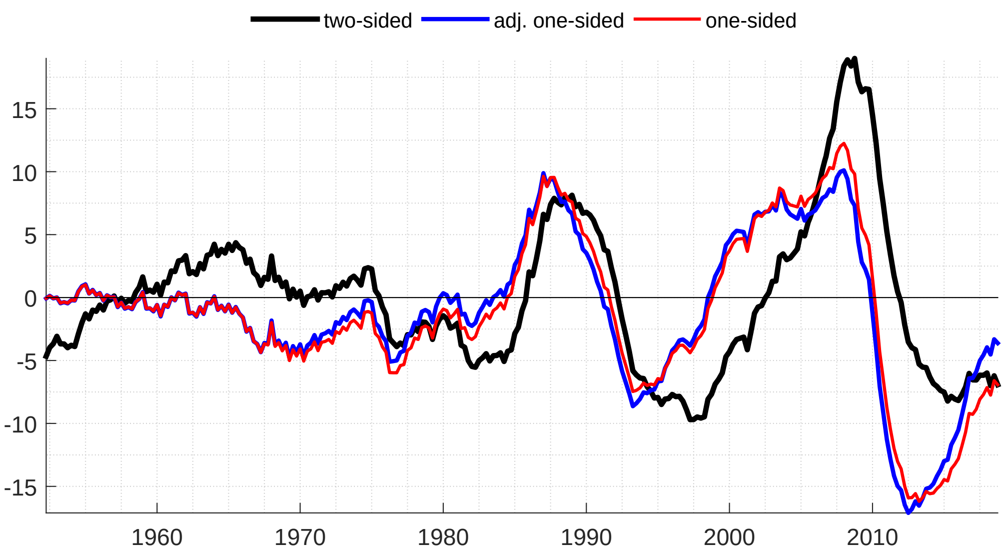

(b) Credit-to-GDP ratio

Figure 3: Quarterly HP-filtered data (1952Q2-2018Q3)

Notes: The graphs show HP-filtered data. "Two-sided" refers to data filtered with HP-2s, "adj. one-sided" to data filtered with HP-1s*, and "one-sided" to data filtered with HP-1s. The respective smoothing parameters can be inferred from Table 3.

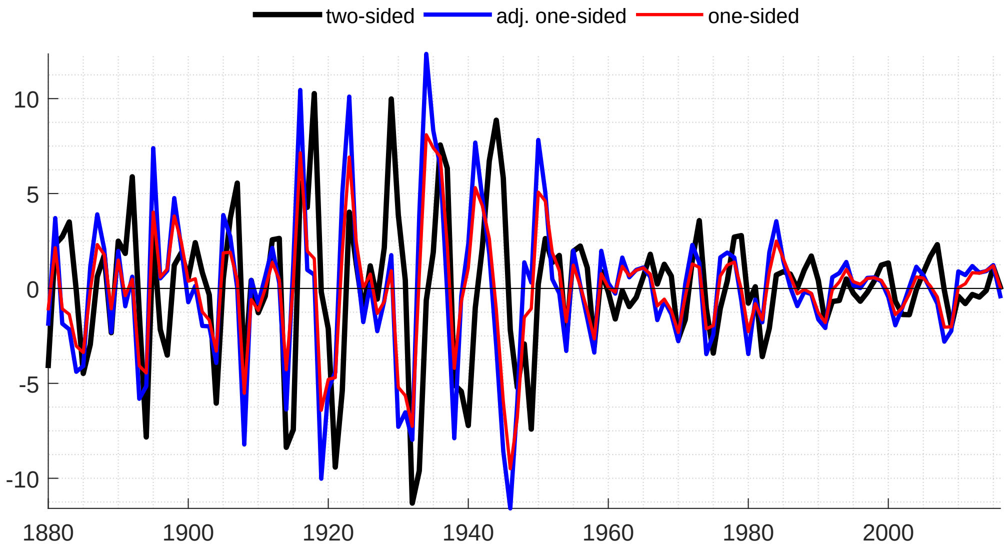

(a) Real GDP

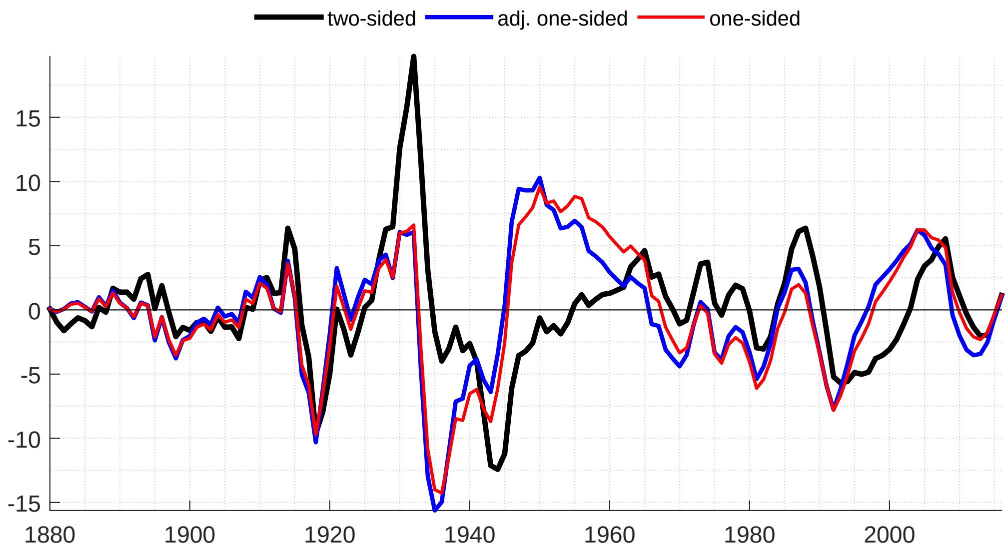

(b) Credit-to-GDP ratio

Figure 4: Yearly HP-filtered data (1880-2016)

Notes: The graphs show HP-filtered data. "Two-sided" refers to data filtered with HP-2s, "adj. one-sided" to data filtered with HP-1s*, and "one-sided" to data filtered with HP-1s. The respective smoothing parameters can be inferred from Table 3.

C Robustness exercise

Table 5: Maximum correlations of one-sided HP filters with two-sided HP filter

λ	HP-1s	HP-1s**	λ^{**}
	DGP 1: $y_t = \varepsilon_t$		
6.25	0.84	0.93	751.92
1,600	0.95	0.98	345,784
400,000	0.98	0.99	30,375,000
	DGP 2: $y_t = y_{t-1} + \varepsilon_t$		
6.25	0.58	0.67	58
1,600	0.54	0.63	16,732
400,000	0.58	0.65	2,463,711
	Quarterly US real GDP		
1,600	0.54	0.64	17,147
	Yearly US real GDP		
6.25	0.44	0.62	154.37
	Quarterly US credit-to-GDP ratio		
400,000	0.65	0.86	$39 \cdot 10^9$
	Yearly US credit-to-GDP ratio		
1,562.50	0.46	0.64	118,900.74

Notes: The table shows the correlations of the extracted cyclical components obtained using one-sided HP filters with the cyclical components obtained using the two-sided HP filter. HP-1s depicts the correlation when both filters, i.e. the one-sided filter and two-sided filter, use the same smoothing parameter λ. HP-1s** illustrates the maximum correlation that can be obtained by adjusting the smoothing parameter (λ^{**}) of the one-sided HP filter.

D Phase shift of the one-sided Hodrick-Prescott filter

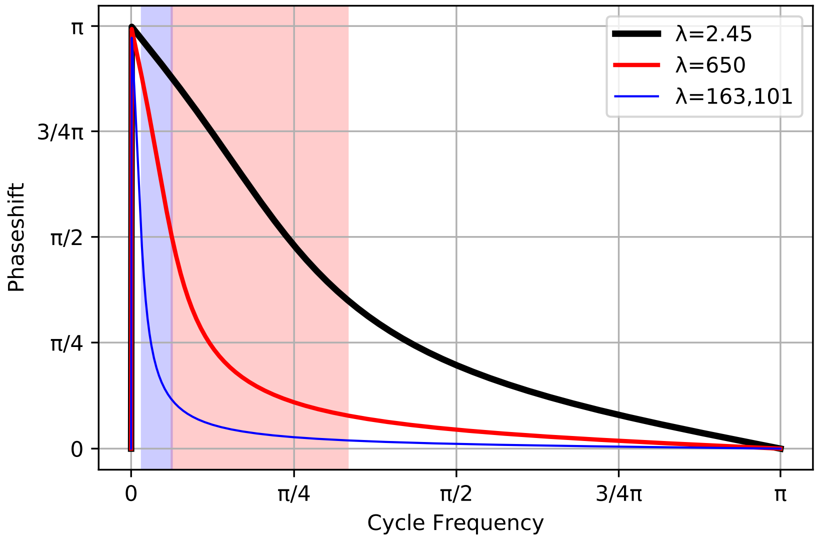

Figure 5: Phase diagram of HP-1s for different values of the smoothing parameter λ

Notes: The graph shows the phase shift that is induced by the one-sided HP filter for different values of the smoothing parameter λ at different frequencies. A positive phase (y-axis in radians) denotes a backward shift, i.e. introducing a lead. Based on a quarterly sampling frequency, the red and blue shaded areas indicate the frequency bands for business cycles (1.5 to 8 years) and financial cycles (8 to 30 years) respectively (x-axis in radians). The values close to zero of the y-axis, for instance, in the higher frequency bands, indicate that the respective frequency components are only mildly affected by the filter. However, the increasingly positive values in the lower frequency bands indicate that longer cycles are shifted backward in phase by a certain amount. For example, given a smoothing parameter of $\lambda = 650$, cycles with a duration of eight years are shifted backward in phase by two years.