



ARTICLE

Is the Hamilton regression filter really superior to Hodrick–Prescott detrending?

Reiner Franke¹, Jiri Kukacka^{2,3} , and Stephen Sacht^{4,5} 

¹Chair of Monetary Economics and International Finance, University of Kiel, Kiel, Germany

²Faculty of Social Sciences, Institute of Economic Studies, Charles University, Prague 1, Czechia

³Institute of Information Theory and Automation, Czech Academy of Sciences, Prague 8, Czechia

⁴Hamburg Institute of International Economics (HWWI), Hamburg, Germany

⁵Institute of Economics, University of Kiel, Kiel, Germany

Corresponding author: Jiri Kukacka; Email: jiri.kukacka@fsv.cuni.cz

Abstract

An article published in 2018 by J.D. Hamilton gained significant attention due to its provocative title, “Why you should never use the Hodrick–Prescott filter.” Additionally, an alternative method for detrending, the Hamilton regression filter (HRF), was introduced. His work was frequently interpreted as a proposal to substitute the Hodrick–Prescott (HP) filter with HRF, therefore utilizing and understanding it similarly as HP detrending. This research disputes this perspective, particularly in relation to quarterly business cycle data on aggregate output. Focusing on economic fluctuations in the United States, this study generates a large amount of artificial data that follow a known pattern and include both a trend and cyclical component. The objective is to assess the effectiveness of a certain detrending approach in accurately identifying the real decomposition of the data. In addition to the standard HP smoothing parameter of $\lambda = 1600$, the study also examines values of λ^* from earlier research that are seven to twelve times greater. Based on three unique statistical measures of the discrepancy between the estimated and real trends, it is evident that both versions of HP significantly surpass those of HRF. Additionally, HP with λ^* consistently outperforms HP-1600.

Keywords: Business cycles; smoothing parameter; trend concept; growth regimes

JEL classifications: C18; C32; E32

1. Introduction

Macroeconomic research focusing on the long-term growth of business cycles often involves breaking down time series data into a trend component, possibly exhibiting limited flexibility or breaks and a cyclical component that reflects recurrent fluctuations at a business cycle frequency. Many economists consider detrending a routine task, opting for widely accepted methods that require minimal effort. The Hodrick–Prescott (HP) filter meets these criteria, providing a transparent and widely applied solution in both academic research and practical business cycle analysis.

Since its introduction, the HP filter has been criticized. Singleton (1988) is cautious about the two-sided HP filter for distorting the properties of seasonally adjusted versus unadjusted time series, which might result in inconsistent parameter estimates. Not accounting for the stochastic properties of the data could also lead to spurious cycles when the HP filter is applied, as noted by Jaeger (1994). Diebold and Kilian (2001) introduce a metric based on evaluating short-run and long-run forecasts. They discuss the likelihood of the predictability of various model outputs being

undisguisable if the underlying time series are HP-filtered. In the context of dynamic stochastic general equilibrium modeling and estimation, Canova (2014) emphasizes the consequences of using model-based transformed versus statistically filtered data, where the latter is the outcome of applying the HP or other (linear) bandpass filtering techniques.

Hamilton's (2018) influential paper, titled "Why you should never use the Hodrick-Prescott filter," has recently received much more attention for criticizing the HP filter. Beyond its compelling title, the paper introduces a superior alternative with a comparable level of complexity: Hamilton's regression filter (HRF). This alternative addresses the limitations of the HP approach and has quickly gained popularity, being applied in various empirical studies (Van Zandweghe, 2017; López-Salido *et al.*, 2017; Danielsson *et al.*, 2018; Richter *et al.*, 2019; Blecker and Setterfield, 2019, p. 247; Araújo *et al.*, 2019; Cauvel, 2022; Reissl, 2020; Richter *et al.*, 2021).¹

Our study ties into the strand of literature that emerged after the launch of the HRF versus HP debate. In particular, our work is closely linked to Moura (2022), who provides direct commentary on Hamilton (2018). Moura noted the mechanical delay between the data and the estimated trend as one of the drawbacks of HRF, discussing the potential artificial serial correlation and predictability in HRF-generated cycles. In contrast, Hamilton (2018) argues that these observations are more pronounced when the HP filter is applied. Unlike our study, which focuses on quarterly real gross value added, as macroeconomists typically decompose GDP time series, Moura (2022) concentrates on detrending consumption and stock prices, which resemble random walks. In a direct follow-up to Phillips and Shi (2021), Mei *et al.* (2022) introduced the boosted HP method, a machine-learning type of filtering. This approach involves filtering the data twice, extracting trend elements that remain part of the cyclical component after the initial filtering. The authors assert the robustness of the boosted HP method for various values of the smoothing parameter λ . Therefore, they follow Ravn and Uhlig (2002) and consider a standard value of $\lambda = 1600$ for quarterly data.²

Following Franke *et al.* (2022, 2023), we also question the commonly used smoothing parameter $\lambda = 1600$ of the HP filter for quarterly data analysis. Our research does not rely on asymptotic econometric theory, as in many studies on HP, but rather builds on a more elementary approach so that its results can also be easily assessed by nonspecialists. It is based on generating artificial data with a specified trend and cycle to provide a controlled environment for the analysis. This simulation-based approach allows us to carefully evaluate the properties of the filters via Monte Carlo experiments and to draw clear conclusions regarding the potential superiority of one filter over another.³

Our findings suggest that HRF generally performs worse than HP detrending and that values of the HP smoothing parameter λ in the range of seven to twelve times higher than the conventional choice of 1600 are more appropriate for accurate trend estimation. While one might raise the question of whether the value of the smoothing parameter should be set to a value below unity (Hamilton, 2018) or is even somewhat negligible (Mei *et al.*, 2022), our paper adds additional insights stemming from high values of λ and shows how they may improve the performance of the HP filter. To ensure practical relevance, our data-generation process is tailored to the aggregate output over the US business cycle.

2. Methodological distinctions

Before delving into the potential superiority of one detrending procedure over the other, it is crucial to acknowledge their distinct methodological foundations. HP's perspective posits that time series can be decomposed into separately defined growth and cyclical components. According to this view, the growth component evolves smoothly over time, driven by a largely independent process involving demographic and technological factors. In this context, the HP filter aims to isolate the cyclical component, although it is less intricate than other procedures that share this perspective. For a survey of such methods, see Alexandrov *et al.* (2012).

Separating an unknown cyclical component from an unknown trend, in line with common economic growth theories, differs somewhat from the objective of HRF. This approach poses the following question: “How different is the value at date $t+h$ from the value that we would have expected to see based on its behavior through date t ?” Forecasting, achieved through simple linear regressions, serves to address factors such as cyclical fluctuations, notably the occurrence of recessions and the timing of recovery. The “cycle” is then defined by prediction errors, ensuring their stationarity (Hamilton, 2018, p. 836).⁴

Strictly speaking, the HP filter and HRF represent distinct concepts lacking a common basis or superiority criterion. Despite this, the aforementioned authors employing HRF often simply replace the HP filter without detailed methodological reflection. Notably, the authors using HRF often find it adequate to restate the widely known critique of HP without providing a rationale for this decision that would have weighed the advantages and disadvantages of the two filters, while their interpretations of HRF outcomes largely mirror those of HP. This is particularly true for output series such as GDP, where deviations from estimated trends signify the economic cycle.

Critical issues raised against HP by Hamilton (2018) include (a) the introduction of spurious dynamic relations that are unrelated to the underlying data-generating process, which is the most often repeated criticism of HP, and (b) notable distortions in filtered values at sample endpoints, which are the most serious for real-time quantitative assessments of the current stage of the business cycle. However, these criticisms often stem from a specific view of stochasticity, disregarding alternative theoretical possibilities.

On the other hand, HRF also encounters criticism. Hodrick (2020, p. 11) emphasized that HRF “is designed for classical time series environments, and it may not work well when samples are small relative to the changes taking place in the underlying economic process or when the changes in the underlying trend do not allow for straightforward time series models.” Moreover, the regression forecasts may not capture all the information in the past history of the process. Finally, the following section provides a basic example illustrating that the HRF trend, when interpreted similarly to an HP trend, may introduce spurious dynamic relations without foundation in the underlying data-generating process. An additional discussion on finer methodological issues distinguishing HRF from the HP filter can be found in Franke et al., (2023, Section 2).

3. Elementary observations of HRF properties

3.1 A smooth sine wave scenario

Given a series $\{y_t\}_{t=1}^{t^{max}}$, HRF specifies the trend value y_t^* in period t as linear regression forecast made h periods earlier based on a small number of lags p up to y_{t-h} . The “cycle” $\{c_t\}_{t=h+p}^{t^{max}}$ is then defined as the prediction error. Regression coefficients homogeneous across all periods t are obtained from a preceding OLS estimation over the entire sample. Regarding h , Hamilton (2018) recommends a value corresponding to roughly a quarter of a typical business cycle. Specifically, for quarterly data, he fixes $h = 8$ and $p = 4$. The regression reads

$$y_t = \beta_0 + \sum_{k=1}^p \beta_k y_{t-h-k+1} + u_t, \quad t = h + p, \dots, t^{max}$$

and with the resulting estimates $\widehat{\beta}_0, \widehat{\beta}_1, \dots, \widehat{\beta}_p$, the HRF trend y_t^* and the cyclical component c_t in period t are given by

$$y_t^* = y_t^{*,HRF} = \widehat{\beta}_0 + \sum_{k=1}^p \widehat{\beta}_k y_{t-h-k+1}, \quad t = h + p, \dots, T \tag{1}$$

$$c_t = c_t^{HRF} = y_t - y_t^* = \widehat{u}_t$$

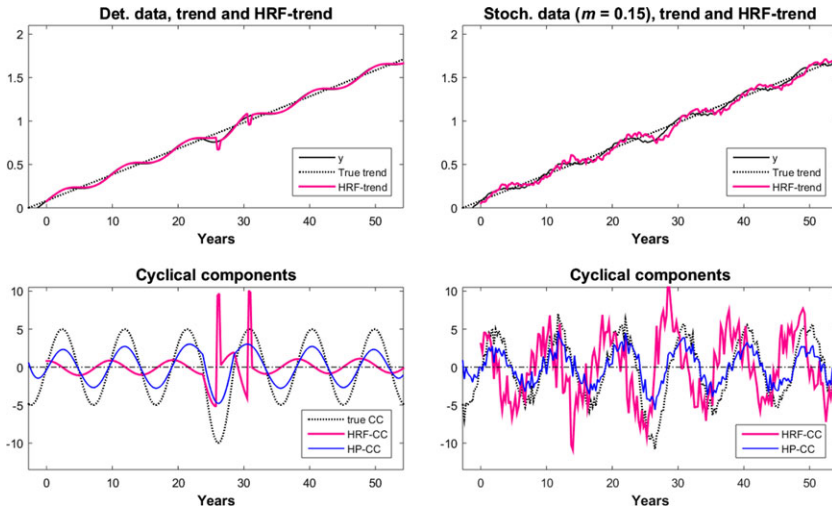


Figure 1. Application of HRF in a sine wave scenario.

To illustrate HRF limitations, consider the examples depicted in Figure 1. The upper-left panel plots a sine wave with a period of $T = 9.50$ years around a straight trend line y_t^* with a slope of 3% per year. The series to be examined is thus

$$y_t = y_t^* + c_t. \tag{2}$$

The bold red solid line in the upper-left panel represents the estimated HRF trend, which covers the data most of the time. Excluding the spikes in the middle of the sample, the lower-left panel illustrates the significantly lower amplitude of the cyclical component derived from HRF *vis-à-vis* the sine wave oscillations (the dotted line); the proportion is approximately 1:5. Importantly, a quarter-cycle phase shift between the two series further indicates that HRF might lead to a significant misjudgment in accurately timing the true underlying business cycles.⁵ For a similar point, for applied work, see Cauvel (2022).

To demonstrate another potentially spurious outcome of HRF between years 20 and 30, we introduce a recession that is twice as deep as that in the other cycles. While HRF predictions y_t^* reflect this change, the corresponding downward revisions significantly overshoot the true contraction. Consequently, the cyclical component exhibits the first positive spike. Two subsequent spikes, negative and positive, occur when the sine wave oscillations return to their original amplitude. Thus, HRF suggests a sudden and extraordinary recovery just in the deepest recession and a sizeable contraction in the following recovery, around $t = 30$.⁶

The lower-left panel of Figure 1 also contrasts the HRF cyclical component with that from the ordinary HP filter (the thinner blue solid line), which succeeds almost perfectly in tracing out the regular pattern of the sine wave oscillations; the correlation is as high as 0.975. Conversely, the HP filter tends to incorporate too much of the cycle into the trend, resulting in an estimated cyclical component with only three-fifths of the true cycle’s amplitude.

3.2 The role of stochastic noise

The example of a deterministic sine wave shows that HRF seems unsuitable for filtering smoothly evolving time series. It is useful to add stochastic noise to this example to assess the potential improvements the filter may offer.

Table 1. Statistics of the estimated cyclical components c_t^{HP} and c_t^{HRF} from (3), varying with the noise level m

m	HP		HRF					
	std.	C(0)	std.	C(0)	C(5)	C(6)	C(7)	C(8)
0.00	0.523	0.98	0.440	0.05	0.36	0.42	0.47	0.44
0.05	0.525	0.97	0.924	0.07	0.67	0.74	0.80	0.82
0.15	0.544	0.96	1.073	0.24	0.74	0.79	0.82	0.83
0.25	0.578	0.95	1.165	0.42	0.76	0.77	0.79	0.77
0.40	0.641	0.94	1.234	0.59	0.73	0.71	0.71	0.70
0.50	0.683	0.94	1.248	0.65	0.68	0.66	0.65	0.62
0.60	0.721	0.93	1.248	0.70	0.63	0.60	0.59	0.56
0.70	0.754	0.93	1.242	0.73	0.57	0.54	0.53	0.50

Note: "std." is the standard deviation of c_t^{HP} and c_t^{HRF} , respectively, in relation to that of the true cycle c_t . "C(k)" for $k = 0, 5, 6, 7, 8$ quarters stands for the cross correlation coefficient $\text{Corr}(c_t, c_{t-k}^j), j = \text{HP, HRF}$. Boldface figures indicate maximal correlation across all lags.

The deterministic oscillations in (2) with an amplitude $\pm\bar{\alpha} = 0.05$ are disturbed by random draws from a normal distribution with a standard deviation related to $\bar{\alpha}$ by a factor m . For $-2 \leq t \leq 55$, the cyclical component c_t is given by

$$c_t = \alpha_t \sin(\omega t) + \varepsilon_t, \quad \alpha_t = \begin{cases} 2\bar{\alpha} & 23.75 \leq t \leq 28.50 \\ \bar{\alpha} = 0.05 & \text{else} \end{cases} \tag{3}$$

$$\omega = 2\pi/9.50, \quad \varepsilon_t \sim N(0, \sigma^2), \quad \sigma = m\bar{\alpha}$$

Simulations in the two panels on the right-hand side of Figure 1 are based on $m = 0.15$. The slightly perturbed cycle c_t is the black dotted line in the lower-right panel. What is particularly noticeable is the substantial, overproportional noise in the cyclical component c_t^{HRF} indicated by HRF-CC, the bold red line. Moreover, its cycle variability not only increases compared to that of the deterministic case but is now of a similar magnitude to that of the true cycle. In this regard, HRF is better than the HP filter (blue), whose cyclical component exhibits only minimal alterations compared to the deterministic case. However, considerable doubts arise when attempting to identify a general pattern in these fluctuations. It appears that the contractions and expansions of c_t^{HRF} end considerably before those of the true cycle. In essence, the phase shift highlighted in the lower-left panel seems to persist in the stochastic environment. Contemporaneous and lagged cross-correlations support this visual impression: a disappointingly low $\text{Corr}(c_t, c_t^{HRF}) = 0.237$, whereas $\text{Corr}(c_t, c_{t-k}^{HRF}) = 0.828$ for lag $k = h = 8$ (compare to a large $\text{Corr}(c_t, c_t^{HP}) = 0.964$ for the HP filter).

Table 1 examines the impact of varying noise levels m . The random sequences $\tilde{\varepsilon}_t$, drawn from the unit normal so that $\varepsilon_t = \sigma \tilde{\varepsilon}_t = m\bar{\alpha} \tilde{\varepsilon}_t$, are uniform across all m . Differences in statistics (i.e., standard deviation and correlation with the true cycle) are thus solely attributed to changes in the noise level. For the HP filter, the statistics remain consistent across noise levels. Although the contemporaneous correlation slightly decreases with higher m , even for elevated levels, the cyclical pattern maintains close alignment with the true pattern. Conversely, the impacts on HRF are more pronounced. Initially, c_t^{HRF} has insufficient variability at $m = 0$, aligning more closely with the true cycle as m increases. At higher levels, however, some overshooting is observed. However, in the deterministic case, both contemporaneous and lagged correlations are unsatisfactorily low, and a modest increase in randomness notably improves these statistics. Thus, at high noise levels such as $m = 0.50$, the contemporaneous correlation becomes a stronger indicator of the connection between the estimated and true cyclical components. Nevertheless, the correlation coefficients with c_t do not surpass those of the HP filter.

These simple experiments provide an initial heuristic to assess the relative merits of HRF and the HP filter. They caution against using HRF when dealing with data showing excessive regularity. Regarding stochastic perturbations, the findings demonstrate that HP is more reliable than HRF in recognizing the true cyclical pattern, even under high noise levels (m). However, relying on HP results in a systematic underestimation of the cycle's overall variability, an aspect seemingly better captured by HRF.

4. Generation of the artificial data

According to the stylized examples in the previous section, HRF may be misleading in practice. We now examine the accuracy of HRF via Monte Carlo simulation. We aim to assess to what extent the insights from the previous section relate to real-world data and noise levels closer to reality. Additionally, we explore smoothing parameter values greater than 1600 for HP, reducing the trend's flexibility and potentially yielding a more variable estimated cyclical component.

To provide guidance on the suitability of HRF, we generate a large number of draws from an artificial stochastic process with a known trend and a cyclical component that reflects the basic features of the US business cycle. Measuring the distance between the true and estimated trend series allows an examination of a filter's precision to recover the true trend.

4.1 Specification of stochastic processes

We calibrate the data-generating process to an empirical business cycle variable. We work with the quarterly real gross value added (GVA) of the US nonfinancial corporate business, measured in logarithms, instead of GDP.⁷ GVA is preferred over GDP because it aligns closely with the output variable in many small-scale macro models, which typically represents the output of the firm sector alone.

For the remainder of the paper, the underlying time unit will shift to a year from a quarter. The simulation's time horizon remains constant at 60 years, a standard duration for many post-WWII empirical investigations.

4.1.1 Three growth regimes

In specifying the output trend, two versions will be distinguished, considering a pivotal distinction in economic theory: the concept of a deterministic *versus* a stochastic trend. The former is typically depicted as an increasing line with a fixed slope, while a random walk with a constant drift represents the latter. However, for the postwar US economy, maintaining the assumption of a constant growth trend over an extended period is unrealistic; indeed, growth rates were consistently declining. A simplified narrative identifies three growth regimes. Drawing on a recent discussion by Hall (2020), these regimes include an era of modernization from 1950 to 1975, an era of liberalism from 1980 to 2000, and subsequently, an era of knowledge-based growth. Associated with the concept of productivity slowdown, growth rates progressively decrease from one regime to the next.

Correspondingly, we postulate three regimes, each spanning 20 years, with growth rates of 5%, 3.5%, and 2%, respectively. The abrupt transitions between regimes are gently smoothed using a moving average around the break dates to mitigate any undue disadvantage to the HP filter. Specifically, we employ the two-sided moving average TS-MA($x_t, 4$) with equal weights, considering a series extended by 4 quarters on each side of a value x_t . Applying this to the step function representing the three growth rates, we obtain a continuous relationship of (actual or expected) trend growth rates g_t^e in periods $t = 0, 0.25, 0.50, \dots, 60.00$, albeit with kinks during regime changes:

$$g_t^e = \text{TS-MA}(\tilde{g}_t^e, 4) \quad (4)$$

$$\tilde{g}_t^e = \begin{cases} 0.050 & \text{if } t < 20 \\ 0.035 & \text{if } 20 \leq t < 40 \\ 0.020 & \text{if } t \geq 40 \end{cases} \tag{5}$$

Typically, $g_t^e = \tilde{g}_t^e$. We employ the acronyms DT and ST to denote the deterministic and stochastic trend concepts, respectively. Initialized with a $y_0^* = 0$ and variance σ_ε^2 for the random walk, the two trend series are expressed as follows:

$$y_t^* = y_{t-0.25}^* + 0.25 g_t^e \tag{DT}$$

$$y_t^* = y_{t-0.25}^* + 0.25 g_t^e + \varepsilon_t \quad \varepsilon_t \sim N(0, \sigma_\varepsilon^2), \text{ iid} \tag{ST}$$

4.1.2 Specification of the cyclical component

In outlining our methodology for assessing the HRF and the HP filter using artificial data from a stochastic process, our approach aligns with recent works by Hodrick (2020, Section 10) and Schüler (2021, Section 4). However, we contend that our experiments are more suitable for studying business cycle data. This implies that we directly impose a desired periodicity on the cyclical component, which contrasts with standard time series analysis. In such analyses, there is often little discussion on whether the resulting cyclical component adequately captures such a fundamental characteristic of the business cycle. Our findings, therefore, complement the investigations conducted by Hodrick (2020) and Schüler (2021).⁸

To construct the cyclical component c_t , we adopt a sine wave with suitable random disturbances to ensure the desired spectrogram shape. Two noise sources are considered: one randomly varies the period of a sine wave from one full cycle run to the next, and the other adds a stochastic moving average process to these waves.

We begin by describing the first concept of a regular oscillatory motion, denoted s_t . It features a constant amplitude α and an average period T . The random periods of a single cycle run are drawn from an interval $[T - \Delta T, T + \Delta T]$ with equal probabilities. Denoting the uniform distribution as U , ϕ is a parameter that may shift the waves in time, and $t \in \mathbb{R}$ is a point in time. The relationship $s_t = s_t(T, \Delta T, \phi, \alpha)$ is defined as follows:

$$\begin{aligned} s_t(T, \Delta T, \phi, \alpha) &= \alpha \sin [\omega_t (t - \tau_t) + \phi] \quad \text{where:} \\ \omega_t &= 2\pi/T(k) \quad \text{for } t_{k-1} \leq t < t_k \\ T(j) &\sim U(T - \Delta T, T + \Delta T) \quad j = 1, 2, \dots \\ t_k &= \sum_{j=1}^k T(j) \quad t_0 = 0 \\ \tau_t &= t_{k-1} \quad \text{for } t_{k-1} \leq t < t_k \end{aligned} \tag{6}$$

In a continuous-time scenario with a transient $\phi = 0$, a cycle initiates at $t = t_{k-1}$ when the argument of the sine function in the first row is $\omega_t(t - t_{k-1}) = 0$. Note that $t_k - t_{k-1} = T(k)$. Thus, as t approaches t_k from the left, the argument tends toward $\omega_t(t_k - t_{k-1}) = [2\pi/T(k)] T(k) = 2\pi$, and the function converges to $\sin(2\pi) = 0$, from whereon the next cycle starts. With a nonzero shift factor ϕ , we have $\sin(\phi)$ at these connection points. In the context of specifying the cyclical component c_t below, the function s_t only needs to be evaluated in quarterly intervals at $t = 0, 0.25, 0.50, \dots, 60$.

Function (6) prompts the question of selecting a “typical” business cycle period, conventionally considered no longer than 8 years in discussions around Hodrick–Prescott detrending or band-pass filters. Contrary to this perspective, recent research by Beaudry et al. (2020) and Barrales-Ruiz and von Arnim (2021) challenges this notion. Their analysis of spectrograms for cycle-related

empirical variables without long-run growth (e.g., working hours *per capita* or job finding rates) reveals a distinct peak occurring between 38 and 40 quarters. This is also the order of magnitude that guides our investigation.

Motions s_t constitute the core of our cyclical component, but the concern arises that a single motion might exhibit excessive regularity, potentially disadvantaging HRF relative to HP. Thus, we combine two sine waves, each with a unique (expected) period T : one with $T = 9.50$ years and the second with $T = 7.00$ years but half the amplitude of the former. This choice is motivated by a second minor peak for a shorter period in several of the aforementioned spectrograms.

Turning to our second noise concept, introducing additional random perturbations in these oscillations, a three-lag moving average process, MA(3), proves sufficient for our purpose. It is unnecessary to include an autoregressive part, as sine waves already fulfill their role. Therefore, with respect to the amplitude α , the MA coefficients $\theta_1, \theta_2, \theta_3$, and the variance of the innovations σ_η^2 , our specification for the cyclical component c_t is as follows, at times $t = 0, 0.25, 0.50, \dots, 60$:

$$c_t = s_t(9.50, 0.00, 0, \alpha) + s_t(7.00, \Delta T, 0.70 \cdot 7.00, \alpha/2) + \epsilon_t + \sum_{j=1}^3 \theta_j \eta_{t-0.25j} \quad \eta_t \sim N(0, \sigma_\eta^2) \quad \Delta T = 1.00 \tag{7}$$

A shift $\phi = 0.70 \cdot 7.00$ in the second sine wave in (7) introduces more irregularity in the composite motion than for example, $\phi = 0.50 \cdot 7.00$. Note that the parameters $\theta_1, \theta_2, \theta_3, \sigma_\eta^2$ will generally differ when (7) is combined with (DT) and (ST). Overall, the series serves as a representative example of what detrending procedures typically produce in empirical analyses of GDP and similar macroeconomic variables. Therefore, we posit that the series

$$y_t = y_t^* + c_t \quad t = 0, 0.25, 0.50, \dots, 60 \tag{8}$$

resulting from the combination of (7) with the trends (DT) and (ST) provided a realistic test bed that HRF and HP should be able to cope with.

4.2 Calibration of the numerical coefficients

Before delving into the details, it is crucial to outline the procedure used to determine the parameters of the data-generation processes, $\beta := (\alpha, \theta_1, \theta_2, \theta_3, \sigma_\eta, \sigma_\epsilon)$. The design of this procedure and its outcomes are adopted from Franke *et al.* (2023). It has already been pointed out that independent of the specific values assigned to β , the criterion of a realistic periodicity of the simulated series y_t is directly fulfilled by (7) by construction.

Turning to the other features shared with the empirical GVA, we focus on the series of its first differences. Specifically, we consider five moments: their standard deviation and the first four autocorrelations. For each scenario (DT) and (ST), our goal is to determine a parameter set β that, on average, brings these simulated moments as close as possible to their empirical counterparts. Closeness, in this context, is measured by a loss function $L(\cdot)$ that computes the mean quadratic distance between simulated and empirical moments across 10 simulation runs of (7) and (8). Averaging enhances the precision of parameter estimates compared to the ideal case where expected values of moments could be determined analytically; for comprehensive details and the improved precision achieved by 10 repetitions, see, e.g., Duffie and Singleton (1993, p. 945).

The objective function’s value depends on the parameter vector β and a random seed, denoted by natural numbers b , initializing the stochastic simulations; thus, $L = L(\beta, b)$ in general. In total, $b = 1, 2, \dots, B = 100$ samples prove sufficient for the following two-stage procedure. In the first stage, for each b , coefficients $\hat{\beta}^b$ are computed, which minimize the loss function over all admissible β . Interestingly, in each case, all the empirical moments are (almost) perfectly matched, so that $L(\hat{\beta}^b, b) \approx 0$ for all b . For the stochastic trend scenario, it should be added that this holds when the standard deviation σ_ϵ of the random walk is fixed at values less than or equal to 0.0050.

Table 2. Numerical coefficients β^o (rounded) obtained from optimization (9)

	α	θ_1	θ_2	θ_3	$100 \cdot \sigma_\eta$	$100 \cdot \sigma_\varepsilon$
DT :	0.0433	1.3885	0.7083	0.6475	0.7175	—
ST :	0.0421	1.2354	0.7515	0.6310	0.7117	0.500

Any increase in σ_ε beyond this threshold would progressively deteriorate the match. This finding led us to choose an upper limit of $\sigma_\varepsilon = 0.0050$ for the calibration. Within the framework of (ST) combined with (7), we consider this value as providing a maximal role to the randomness in the trend, one that still remains compatible with our current requirements for the first differences of GVA.

Having obtained the optimal parameters $\{\hat{\beta}^b\}_{b=1}^B$ in the first stage, it is crucial to consider that a single vector $\hat{\beta}^b$ is tailored to a particular random seed b . Using $\hat{\beta}^b$ for simulations initialized with a different seed $c \neq b$ may generally lead to a suboptimal match, with $L(\hat{\beta}^b, c) > L(\hat{\beta}^b, b) \approx 0$. To account for random variations, we opt for the random seed b^o that yields the lowest mean loss across all realizations. The final parameter vector is therefore specified as:

$$\beta^o = \hat{\beta}^{b^o} \quad \text{where} \quad b^o := \operatorname{argmin}_b \left\{ \frac{1}{B} \sum_{c=1}^B L(\hat{\beta}^b, c) \right\} \tag{9}$$

The outcome of this optimization procedure is reported in Table 2. The minimum losses $(1/B) \sum_c L(\beta^o, c)$ in (9) show no significant difference between the deterministic and stochastic trend scenarios. Thus, both scenarios are suitable dynamic processes for generating artificial data in the subsequent section to test HRF and HP detrending in the context of GVA cyclical growth.

5. Comparing HRF and HP

In this section, the stochastic processes with scenarios (DT) and (ST) are simulated numerous times to generate data for applying the HRF and HP filter. As the trend and cyclical components of these time series are known, they offer a robust foundation for assessing, on average, the efficacy of the two filters in recovering the true trend-cycle decomposition.

5.1 Qualitative visual analysis

Expectations lean towards HRF being more effective in handling a random walk trend rather than a deterministic trend line. We, therefore, begin with a visual analysis of a typical example derived from (7) and (8) in the context of a stochastic trend (ST). In Figure 2, the upper-left panel displays the data y_t itself, depicted by the thin solid (black) line closely aligned with the thinner (blue) line representing its stochastic trend y_t^* . While these lines are challenging to differentiate due to their overall increase, the other three panels provide alternative perspectives. What can be better seen, however, is the main characteristic of the estimated trend $y_t^{*,\text{HRF}}$, the bold red line: it exhibits distinct swings, with the upper and lower turning points occurring slightly delayed compared to those in the data, evident at approximately $t = 5, t = 33, t = 37,$ and $t = 56$.

Because of the cyclical nature of y_t , it is evident that this affects the resulting cyclical component $c_t^{\text{HRF}} = y_t - y_t^{*,\text{HRF}}$. As depicted in the upper-right panel of Figure 2, on these dates c_t^{HRF} exhibits troughs and peaks before they actually occur in the true cyclical component $c_t = y_t - y_t^*$.⁹ This phase shift is reminiscent of the lower-right panel in Figure 1, which was obtained in a more stylized setting by applying HRF to a strictly linear trend line and mildly disturbed regular sine waves around it. This deficiency is likely to be more general in nature, especially for growth series with a similar, relatively low noise level in their components.

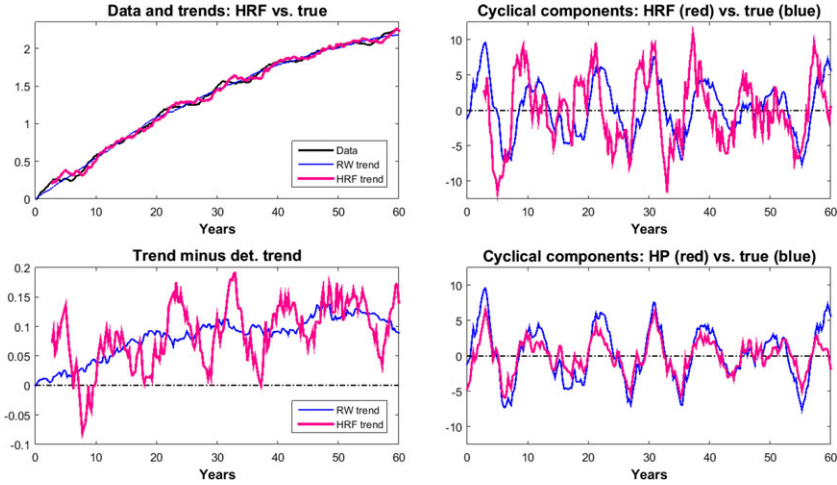


Figure 2. HRF and HP were applied to a sample from (ST).

The contrast between the HRF trend and the true random walk trend becomes more evident in the lower-left panel of Figure 2, which plots their difference from the deterministic trend (DT). While the relatively steady (blue) line illustrates the random walk’s gradual deviation from its deterministic counterpart, the HRF estimate captures this tendency with heightened volatility, mirroring the data series fluctuations that contribute to the observed phase shifts.

The random walk trend itself may appear relatively smooth in relation to the HRF trend, but nevertheless, its slope ranges between -3 and $+10$ percent per year, where the latter order of magnitude is not a rare exception. As another crude statistic, the standard deviation of the random walk trend increases to 2.35%, compared to 1.22% for the deterministic trend. These figures illustrate that there is indeed a substantial difference between the two frameworks of (DT) and (ST).

In summary, applying HRF to the current series does not yield the typical detrending outcome. The lower-right panel in Figure 2 illustrates that the conventional HP filter outperforms HRF in recovering the true cyclical pattern. Quantitatively, the contemporaneous correlation coefficient between c_t^{HP} and c_t is 0.903, which is significantly greater than the 0.413 for c_t^{HRF} . However, HP exhibited a downward bias in overall variability, with a standard deviation ratio of 0.640 compared to HRF’s overestimation, resulting in a ratio of 1.289.

5.2 Monte Carlo simulation study

We can now evaluate HRF and the two HP filters more systematically through the quantitative experiments outlined at the beginning of this section. Simultaneously, we take the opportunity to broaden the perspective on the Hodrick–Prescott approach, motivated by occasional suggestions in the literature that the HP filter may excessively incorporate the cycle into the trend (Gordon, 2003, p. 218; Franke *et al.*, 2006, p. 460, fn 10). Therefore, in addition to the conventional smoothing parameter $\lambda = 1600$, we explore the performance of an alternative, higher value. To this end, we make use of the recommendations of Franke *et al.* (2023). Accordingly, in addition to HRF, we tested two versions of HP: HPC, denoted as the “conventional” $\lambda = 1600$, and HPe, with “e” for “enhanced.” The alternative parameters, however, differ for (DT) and (CT). Table 3 reports the alternative parameters, which differ for (DT) and (ST). In both cases, they are substantially greater than $\lambda = 1600$, thus reducing the flexibility of the estimated trend and potentially increasing the overall variability of the cyclical component.¹⁰

Table 3. Smoothing parameters λ for HPC (conventional) and HPe (enhanced)

	HPC	HPe
DT :	1 · 1600	11.71 · 1600 = 18736
ST :	1 · 1600	7.51 · 1600 = 12016

To assess and contrast HRF, HPC, and HPe, a statistical measure must be utilized to quantify the disparity between the estimated and true trends or between the estimated and true cyclical components. This can be done along three dimensions. The first statistic is the distance between the series of the estimated trend $y_t^{*,est}$ and the true trend y_t^* , which is most naturally specified by the root mean square deviation (RMSD). To mitigate the effects of HP’s end-of-sample bias, the first and last four quarters are discarded in the corresponding summations. For HRF, the first $h + p = 8 + 4 = 12$ unavailable values of the sample period are omitted from these sums. Given our samples of 241 quarters from $t = 0$ to $t = 60$ and two series x_t, z_t , their RMSD is defined as:

$$RMSD(x, z) := \sqrt{\frac{1}{241 - 8} \sum_{q=1+4}^{241-4} (x_{t(q)} - z_t)^2}, \quad t(q) := (q - 1)/4 \tag{10}$$

and correspondingly modified for HRF. To make the distance between $y_t^{*,est}$ and y_t^* independent of the size of the variations in the cyclical component $c_t = y_t - y_t^*$, we scale $RMSD(y^{*,est}, y^*)$ by the latter’s standard deviation, which equals $RMSD(y, y^*)$. For a trend estimation $y_t^{*,est}$, we therefore define our distance measure as:

$$d = d(y^{*,est}, y^*; y) := \frac{RMSD(y^{*,est}, y^*)}{RMSD(y, y^*)} \tag{11}$$

Hence, d measures the deviations of the estimated trend from the true trend in the percent of the cyclical component’s variability, making values of d considerably below unity desirable. This normalization provides a better sense of the overall quality of the estimation. It may be noted that the distance between $y_t^{*,est}$ and y_t^* is equivalent to the distance between the corresponding cyclical components $c_t^{est} = y_t - y_t^{*,est}$ and $c_t = y_t - y_t^*$.

The other two measures, as mentioned earlier, include the correlation between the estimated and true cyclical component, denoted as $Corr(c^{est}, c)$, and the standard deviation of c_t^{est} in relation to that of c_t , $\frac{std(c^{est})}{std(c)}$. The treatment of the beginning and end of the sample period remains consistent with the specification of the distance d .

In the main experiment, we generate 1000 sample runs of (7), (8) along with the DT or ST scenario. For each of these runs, the three performance statistics resulting from HRF, HPC, and HPe are computed, with Table 4 presenting the outcome. All the statistics exhibit a fairly symmetric distribution; therefore, the median values are sufficient to be reported. To provide context for the differences between the three detrending procedures, we also include the standard deviation (“std”) of the distributions. The results of our Monte Carlo experiment can, therefore, be summarized as follows:

1. Regarding the two measures of distance and correlation, the Hamilton regression filter (HRF) generally performs worse than conventional Hodrick–Prescott detrending (HPC).
2. On the other hand, in terms of the estimated variability of the cyclical component, HPC has a disadvantage compared to HRF. However, the tendency of HRF to overestimate the true variability needs to be emphasized.

Table 4. Performance statistics for HRF, HPc, HPe from 1000 sample runs

	d			$\text{Corr}(c^{est}, c)$			$\text{std}(c^{est})/\text{std}(c)$		
	HRF	HPc	HPe	HRF	HPc	HPe	HRF	HPc	HPe
Optimal:	0.000	0.000	0.000	1.000	1.000	1.000	1.000	1.000	1.000
<i>DT:</i>									
median :	1.157	0.435	0.199	0.504	0.947	0.983	1.284	0.635	0.944
std. :	0.030	0.015	0.029	0.039	0.011	0.007	0.030	0.018	0.012
<i>ST:</i>									
median :	1.208	0.469	0.304	0.492	0.913	0.953	1.336	0.657	0.940
std. :	0.067	0.029	0.036	0.052	0.016	0.012	0.063	0.032	0.047

3. HPc performs reasonably well in the first two measures, but its consistent underestimation of the cyclical component's variability also raises concerns.
4. The enhanced HP filter (HPe) with appropriately increased smoothing parameter values exhibits notable improvements. In particular, it outperforms HPc and, to a greater extent, HRF in the distance and correlation measures.
5. HPe, however, still exhibits a slight downward bias in the standard deviation of the cyclical component. Despite this, its small magnitude allows for potential bias correction in practical applications, as detailed in Franke *et al.* (2022, Section 5.3).

In summary, we find that the observed differences between the filters are significant, assuming the acceptance of (i) the study's focus on growing output series such as GDP or GVA, (ii) the experimental framework involving the decomposition into a predefined trend and cyclical component, and (iii) the defined task for detrending procedures to identify these series accurately. As highlighted throughout the text, the HP approach's central ambition aligns with the task at hand, and while HRF may adopt a somewhat different methodological perspective, users of HRF often share practical and interpretative views of the HP approach.

6. Conclusion

Our evaluation of the Hamilton regression filter (HRF) and the Hodrick–Prescott (HP) filter draws attention to crucial lessons for applied researchers. While HRF may yield results comparable to those of HP in practical scenarios, it carries the risk of serious misperceptions during certain stages of the economic cycle. The conventional HP filter, criticized for potential spurious dynamics, still proves effective at recognizing cyclical patterns, albeit with a tendency to underestimate cyclical variability. However, our study also introduces an enhanced version, HPe, with a higher smoothing parameter, outperforming HRF and the conventional HP filter. Therefore, applied researchers may approach HRF cautiously, reevaluate the conventional HP filter, and consider HPe a more credible detrending method tailored to the characteristics of their economic data.

Competing interests. The authors declare none.

Financial support. This work was supported by Charles University Research Centre program No. 24/SSH/020 and the Cooperatio Program at Charles University, research area Economics.

Use of artificial intelligence (AI) tools. During the preparation of this work, the authors used the online grammar checker [grammarly.com](https://www.grammarly.com) for standard grammar and spelling support, the paraphrasing tool [QuillBot](https://www.quillbot.com) and the generative AI system

ChatGPT-3.5 in order to improve the readability and language of the article. After using this tool/service, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication. All the authors have read and approved the final manuscript.

Notes

- 1 It is worth noting that there are modified versions of HRF, such as the one proposed by Quast and Wolters (2022). In their work, the authors emphasize a much smoother trend and a more meaningful economic interpretation of the filtered time series compared to the original HRF approach.
- 2 Ravn and Uhlig (2002) show from a time and frequency perspective, respectively, that λ should be set to the fourth power of a change in the frequency of observations. Setting $\lambda_{quarterly} = 1600$, for example, on a monthly frequency, we arrive at $\lambda_{monthly} = 1600 \cdot 3^4 = 129600$.
- 3 Incidentally, doing this for a great number of stochastic realizations of the data-generating process also allows insights into the small-sample variability of HP detrending.
- 4 Specifically, Hamilton (2018, p. 836) argues the “primary reason that we would be wrong in predicting the value of most macro and financial variables at a horizon of $h=8$ quarters ahead is cyclical factors such as whether a recession occurs over the next two years and the timing of recovery from any downturn.”
- 5 In the following, y_t can mean both the entire series $\{y_t\}_{t=1}^{t_{max}}$ or its value at a particular date t according to the context. In (11), we even prefer y to avoid confusion with the time indices.
- 6 This phenomenon persists even when selecting a period $T = 8$ consistent with HRF’s forecast horizon of $h = 8$ quarters.
- 7 The series was obtained from <https://fred.stlouisfed.org/series/B455RX1Q027SBEA>, covering the period 1960:1 to 2018:1.
- 8 For an additional analysis and justification of our approach to the cyclical component, refer to Franke et al. (2023, Section 4.1).
- 9 The cyclical components on the right-hand side of Figure 2 are presented as a percentage of the nonlogarithmic trend values. In the lower-left panel, the cyclical components are expressed as a percentage of the trend output Y_t^* . Thus, the y-axis is scaled as $100 c_t$, where $c_t = y_t - y_t^* = \ln Y_t - \ln Y_t^* = \ln(Y_t/Y_t^*) \approx (Y_t - Y_t^*)/Y_t^*$.
- 10 Hamilton (2018, pp. 835) adopts a different approach to derive more appropriate values for λ . Based on the assumption that not only the second differences in the trend but also the values of the cyclical component are normally distributed, he sets up a maximum likelihood problem to obtain optimal values for λ . For a range of macroeconomic time series, he reports values for λ below ten, with most of them even below unity. This contrasts starkly with our proposal for HPe.

References

- Alexandrov, T., S. Bianconcini, E. B. Dagum, P. Maass and T. S. McElroy. (2012) A review of some modern approaches to the problem of trend extraction. *Econometric Reviews* 31(6), 593–624.
- Araújo, R. A., M. J. Dávila-Fernández and H. N. Moreira. (2019) Some new insights on the empirics of Goodwin’s growth-cycle model. *Structural Change and Economic Dynamics* 51, 42–54.
- Barrales-Ruiz, J. and R. von Arnim. (2021) Endogenous fluctuations in demand and distribution: An empirical investigation. *Structural Change and Economic Dynamics* 58, 204–220.
- Beaudry, P., D. Galizia and F. Portier. (2020) Putting the cycle back into business cycle analysis. *American Economic Review* 110(1), 1–47.
- Blecker, R. A. and M. Setterfield. (2019) *Heterodox Macroeconomics: Models of Demand, Distribution and Growth*. Cheltenham, UK: Edward Elgar.
- Canova, F. (2014) Bridging DSGE models and the raw data. *Journal of Monetary Economics* 67, 1–15.
- Cauvel, M. (2022) The Neo-Goodwinian model reconsidered. *European Journal of Economics and Economic Policies: Intervention, Advance Access*.
- Danielsson, J., M. Valenzuela and I. Zer. (2018) Learning from history: Volatility and financial crises. *Review of Financial Studies* 31(7), 2774–2805.
- Diebold, F. X. and L. Kilian. (2001) Measuring predictability: theory and macroeconomic applications. *Journal of Applied Econometrics* 16(6), 657–669.
- Duffie, D. and K. J. Singleton. (1993) Simulated moments estimation of Markov models of asset prices. *Econometrica* 61(4), 929–952.
- Franke, R., P. Flaschel and C. R. Proaño. (2006) Wage-price dynamics and income distribution in a semi-structural Keynes-Goodwin model. *Structural Change and Economic Dynamics* 17(4), 452–465.
- Franke, R., J. Kukacka and S. Sacht. (2022). Reconsidering Hodrick-Prescott Detrending and Its Smoothing Parameter: Extended Version, Working paper, Universities of Kiel and Prague, Available at SSRN: 10.2139/ssrn.4147280.
- Franke, R., J. Kukacka and S. Sacht. (2023). Is the Hamilton Regression Filter Really Superior to Hodrick-Prescott Detrending? Extended Version, Working paper, Universities of Kiel and Prague, Available at SSRN: 10.2139/ssrn.4210446.

- Gordon, R. J. (2003) Exploding productivity growth: Context, causes, and implications. *Brookings Papers on Economic Activity* 2003(2), 207–279.
- Hall, P. A. (2020) The electoral politics of growth regimes. *Perspectives on Politics* 18(1), 185–199.
- Hamilton, J. D. (2018) Why you should never use the Hodrick-Prescott filter. *The Review of Economics and Statistics* 100(5), 831–843.
- Hodrick, R. J. (2020) An exploration of trend-cycle decomposition methodologies in simulated data. *NBER Working Paper* No. 26750.
- Jaeger, A. (1994) Mechanical detrending by Hodrick-Prescott filtering: A note. *Empirical Economics* 19(3), 493–500.
- López-Salido, D., J. C. Stein and E. Zakrajšek. (2017) Credit market sentiment and the business cycle. *Quarterly Journal of Economics* 132(3), 1373–1426.
- Mei, Z., P. C. B. Phillips and Z. Shi. (2022). The boosted HP filter is more general than you might think. Technical report, arXiv: 2209.09810.
- Moura, A. (2022). Why you should never use the Hodrick-Prescott filter: Comment, MPRA Paper 114922. Germany: University Library of Munich.
- Phillips, P. C. B. and Z. Shi. (2021) Boosting: Why we can use the HP filter. *International Economic Review* 62(2), 521–570.
- Quast, J. and M. H. Wolters. (2022) Reliable real-time output gap estimates based on a modified Hamilton filter. *Journal of Business & Economic Statistics* 40(1), 152–168.
- Ravn, M. O. and H. Uhlig. (2002) On adjusting the Hodrick-Prescott filter for the frequency of observations. *The Review of Economics and Statistics* 84(2), 371–376.
- Reissl, S. (2020) Minsky from the bottom up—Formalising the two-price model of investment in a simple agent-based framework. *Journal of Economic Behavior and Organization* 177, 109–142.
- Richter, B., M. Schularick and I. Shim. (2019) The costs of macroprudential policy. *Journal of International Economics* 118, 263–282.
- Richter, B., M. Schularick and P. Wachtel. (2021) When to lean against the wind. *Journal of Money, Credit and Banking* 53(1), 5–39.
- Schüler, Y. (2021). On the cyclical properties of Hamilton’s regression filter, Available at SSRN: 10.2139/ssrn.3559776.
- Singleton, K. J. (1988) Econometric issues in the analysis of equilibrium business cycle models. *Journal of Monetary Economics* 21(2-3), 361–386.
- Van Zandweghe, W. (2017) The changing cyclicity of labor force participation. *Federal Reserve Bank of Kansas City, Economic Review, Third Quarter* 2017, 5–34.