

Moura, Alban

Article

Why You Should Never Use the Hodrick-Prescott Filter. A Comment on Hamilton (The Review of Economics and Statistics, 2018)

Journal of Comments and Replications in Economics (JCRE)

Suggested Citation: Moura, Alban (2024) : Why You Should Never Use the Hodrick-Prescott Filter. A Comment on Hamilton (The Review of Economics and Statistics, 2018), Journal of Comments and Replications in Economics (JCRE), ISSN 2749-988X, ZBW - Leibniz Information Centre for Economics, Kiel, Hamburg, Vol. 3, pp. 1-17,
<https://doi.org/10.18718/81781.31>

This Version is available at:

<https://hdl.handle.net/10419/295067>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



<https://creativecommons.org/licenses/by/4.0/>

Why You Should Never Use the Hodrick-Prescott Filter

A Comment on Hamilton (*The Review of Economics and Statistics*, 2018)

Alban Moura*

Journal of Comments and Replications in Economics, Volume 3, 2024-1, DOI: [10.18718/81781.31](https://doi.org/10.18718/81781.31)

JEL: B41, C22, E32

Keywords: HP Filter, Hamilton Filter, Business Cycles, Detrending, Filtering, Comment

Data Availability: The data and Matlab code to reproduce the results of this comment can be downloaded at JCRE's data archive (DOI: [10.15456/j1.2024051.1449561565](https://doi.org/10.15456/j1.2024051.1449561565)).

Please Cite As: Moura, Alban (2024). Why You Should Never Use the Hodrick-Prescott Filter: A Comment on Hamilton (2018). *Journal of Comments and Replications in Economics*, Vol 3(2024-1). DOI: [10.18718/81781.31](https://doi.org/10.18718/81781.31)

Abstract

Hamilton (2018) argues that one should never use the Hodrick-Prescott (HP) filter to detrend economic time series and proposes an alternative approach. This comment reconsiders Hamilton's case against the HP filter, emphasizing two simple points. First, in the empirical example Hamilton considers, the HP and Hamilton filters yield cyclical estimates with very similar dynamic properties, questioning the notion that one decomposition outperforms the other. Second, there is a mechanical lag in the Hamilton trend, which might cast doubt on the economic plausibility of the trend-cycle decomposition. It follows that the Hamilton filter might not constitute a systematically better alternative to the HP filter.

*Banque centrale du Luxembourg, Département Économie et Recherche, 2 boulevard Royal, L-2983 Luxembourg, alban.moura@bcl.lu.

Acknowledgements: I thank James Hamilton for detailed and helpful feedback on earlier drafts. For useful comments, I also thank the editor (W. Robert Reed) and one anonymous referee, as well as Patrick Fève, Paolo Guarda, Robert Hodrick, Olivier Pierrard, Xiaoxia Shi, and colleagues at the Banque centrale du Luxembourg. The views presented in this paper are personal and should not be reported as those of the Banque centrale du Luxembourg or the Eurosystem. The author declares that he has no material financial interest related to the results reported in this study.

Received December 12, 2022; Revised November 17, 2023; Accepted January 09, 2024; Published February 20, 2024.
©Author(s) 2024. Licensed under the Creative Common License - Attribution 4.0 International (CC BY 4.0).

1 Introduction

In an important paper, Hamilton (2018) argues that one should never use the HP filter, proposed by Hodrick and Prescott (1981, 1997) to decompose a time series into separate trend and cyclical components. Hamilton makes his point in two steps. First, he highlights three drawbacks of the HP filter: (a) It introduces spurious dynamic relations that have no basis in the underlying data-generating process (DGP). (b) The estimates at the boundaries of the sample are not reliable. (c) Common choices for the smoothing parameter are not supported by the data. Second, he proposes an alternative regression-based strategy, since known as the Hamilton filter, that, he argues, extracts plausible cyclical components while eschewing the pitfalls of the HP filter. According to Hamilton, these elements close the case against using the HP filter.

But is the case really closed? Hamilton's dismissal of the HP filter runs counter to widespread empirical practice. Limitations of the HP filter are well known, but economists still use it for lack of a better alternative. Therefore, the main question is whether the Hamilton filter indeed improves on the HP filter. Using two simple points, this comment argues that the improvement is not so clear.

Hamilton (2018) motivates his criticism of the HP filter from a real-world example: the detrending of consumption and stock prices. Both series resemble random walks, but the cyclical components extracted by the HP filter feature complex dynamics. Hamilton considers that these patterns reflect the filter rather than true properties of the data. Surprisingly, Hamilton does not discuss the properties of the cycles extracted by his alternative approach. This comment fills this gap and finds that the Hamilton cycles exhibit persistence and comovements that closely mirror those found in the HP cycles. Thus, the very example Hamilton invokes to dismiss the HP filter actually fails to establish the superiority of his preferred strategy. If the dynamics found in the HP cycles are spurious, then the similar dynamics found in the Hamilton cycles must be equally misleading. On the contrary, if the dynamics found in the Hamilton cycles are authentic, then the similar dynamics found in the HP cycles imply that the HP filter provides at least a reasonable cyclical estimate.

This first point is mostly rhetorical and only questions the bite of Hamilton's empirical criticism of the HP filter. The second point is more substantial: when the original series is persistent, there is a mechanical delay between the data and the estimated Hamilton trend. At a basic level, this is expected: Hamilton *defines* the trend as an 8-quarter-ahead forecast for quarterly series, so that the trend component reacts to data movements with an automatic two-year gap. While Hamilton does not discuss this timing, this comment shows that it might result in implausible trend-cycle decompositions. For instance, the estimated trend for stock prices rises during most stock-market contractions, before falling abruptly two years after the actual drop in prices, when valuations are already recovering. A similar issue arises when detrending real output and interpreting the trend as potential output: in this case, potential output rises mechanically during most recessions, before sharply falling during the recovery. Of course, given the additive trend-cycle decomposition, questioning the timing of trend estimates necessarily leads to doubts about cyclical estimates as well.

In light of these two points, this comment concludes that the Hamilton filter might not offer a systematically superior alternative to the HP filter. A more balanced assessment is that the two filters provide different views of the data, and that which view is more useful is likely to depend on the application. More broadly, the classic question of how to best extract a stationary component from a potentially non-stationary time series remains open and economists can only benefit from

viewing alternative approaches as complementary rather than substitute, as Canova (1998) argued more than twenty years ago.

1.1 Literature

There is an infinite number of ways to detrend a time series, which unsurprisingly led to a large literature comparing, evaluating, and proposing detrending methods. It is beyond the scope of this comment to review this literature. Here, the focus is more narrowly restricted to the HP and Hamilton filters.

Hamilton (2018) summarizes the literature about the HP filter. Important papers include Harvey and Jaeger (1993) and Cogley and Nason (1995), who showed that the HP filter can produce cyclical components with dynamic properties absent from the original DGP, and de Jong and Sakarya (2016), who reviewed in detail the econometric properties of the HP filter.

Some authors have also evaluated the theoretical properties of Hamilton's regression filter. For instance, Schuler (2021) studies its spectral properties and shows that it amplifies longer-term cycles and mutes shorter-term fluctuations. As a result, the Hamilton filter typically attributes more medium-term movements to the cycle compared to more standard definitions of the business-cycle phenomenon (see, e.g., Stock and Watson, 1999). Schuler also argues that the Hamilton filter may alter the dynamic relationship between several variables because it induces potentially different phase shifts in the cyclical components. This comment concurs with Schuler that the timing properties of the Hamilton decomposition may be problematic.

Jonsson (2020a) documents numerically that the HP and Hamilton filters produce cyclical components with similar dynamic properties in several univariate setups, including the random-walk case. Jonsson concludes that, if one views the properties of HP cycles as problematic, then the same properties found in Hamilton cycles can only be viewed as problematic. This comment complements Jonsson (2020a) at three levels. First, it emphasizes that Jonsson's critique applies to Hamilton's own empirical example, questioning his argumentation. Second, it considers a bivariate setup, shedding light on the cross-correlations and joint dynamics discussed by Hamilton (2018). Third, it highlights the timing properties of the Hamilton trend, which affect the economic plausibility of the decomposition.

Other authors compare the properties of the two filters when applied to actual or simulated data. For instance, Hodrick (2020) simulates various time-series models approximating the U.S. real gross domestic product (GDP) and evaluates the cyclical components recovered by the HP filter, the Hamilton filter, and a band-pass filter. Hodrick argues that the Hamilton filter outperforms the HP filter for simple DGPs and that the HP and band-pass filters perform better for complex models. Also focusing on real GDP, Hall and Thomson (2021) and Dritsaki and Dritsaki (2022) argue that the HP filter provides more plausible trend-cycle decompositions than the Hamilton filter for New Zealand and Greece. Of course, an issue in interpreting these results is the ambiguity related to the definition of the "true" cyclical component of the data. Another study by Jonsson (2020b) confirms the appealing real-time performance of the Hamilton filter, which behaves better than the HP filter in presence of data revision.

Finally, a third group of authors propose extensions of the HP and Hamilton filters. For instance, Phillips and Shi (2021) and Mei, Phillips, and Shi (2022) suggest that repeated applications of the HP filter result in improved asymptotic ability to recover a variety of trends. This iterative procedure, dubbed the boosted HP filter, is grounded in the machine-learning theory of boosting. Quast and Wolters (2022), on the other hand, suggest that smoothing the Hamilton trend by averaging across estimates obtained from different forecast horizons leads to cyclical components with better properties. The results reported in this comment indicate that the Quast-Wolters approach does not solve the timing issue associated with Hamilton trends. Lastly, Hamilton and Xi (2023) show how the Hamilton filter can be used to make non-stationary time series amenable to Principal Component Analysis. They also document the robustness of the Hamilton filter to the extreme COVID-19 outliers.

2 The HP and Hamilton Filters

For completeness, this section provides a brief characterization of the HP and Hamilton filters and reviews some of their properties. More details can be found in the original papers (Hodrick and Prescott, 1981, 1997; Hamilton, 2018).

Both the HP and the Hamilton filters decompose a time series x_t into the sum of two components: $x_t = g_t + v_t$, where g_t is the trend and v_t is the cycle. The difference between the two filters lies in the statistical restrictions used to identify the trend component.

The HP filter defines the trend component as a smooth variable that does not differ much from the observed series. Given a sample of data $\{x_t\}_{t=1}^T$, this objective is formalized by choosing g_t as the solution to the following program:

$$\min_{\{g_t\}_{t=-2}^T} \left\{ \sum_{t=1}^T (x_t - g_t)^2 + \lambda \sum_{t=1}^T [(g_t - g_{t-1}) - (g_{t-1} - g_{t-2})]^2 \right\}, \quad (1)$$

where $\lambda \geq 0$ is a smoothing parameter penalizing large changes in the slope of the trend g_t . The HP trend reduces to the original series when there is no smoothness penalty ($\lambda \rightarrow 0$) and it corresponds to a linear time trend when the penalty is extreme ($\lambda \rightarrow \infty$). At each period, the HP cycle verifies $v_t = x_t - g_t$.

Looking at the minimization program (1), it is clear that the value of the HP trend at any given date depends on the full set of available observations on x_t . This can be formalized in two ways. Given a finite sample of data, stacking all observations on x_t in a column vector $x = (x_T, x_{T-1}, \dots, x_1)'$ allows expressing the HP trend as a column vector $g = A^*(\lambda)x$, where $g = (g_T, g_{T-1}, \dots, g_1)'$ stacks the trend values and $A^*(\lambda)$ is a $(T+2, T)$ matrix whose entries are functions of the smoothing parameter λ . In population, the HP trend admits a symmetric two-sided representation, $g_t = h^*(L)x_t$, where L is the lag operator and where the filter weights $\{h_j^*\}_{j=-\infty}^{\infty}$ are determined only by the value of λ . Given the additive trend-cycle decomposition, there exist similar matrix and filter representations for the HP cycles.¹

¹One can force the trend and cyclical components at each period to load only on current and past observations of the data. This comment does not discuss this one-sided HP filter.

Hodrick and Prescott select the value of the smoothing parameter based on prior assumptions about the relative volatility of the trend and cyclical components for typical macroeconomic time series. This leads them to advocate the use of $\lambda = 1,600$ for quarterly data. Ravn and Uhlig (2002) show how to extend the logic to other frequencies.

Turning to the Hamilton filter, it defines the trend component as the value that we would expect for the original series at date t , based on its behavior up to date $t - h$. This is formalized using a simple linear regression of x_t on a constant, the realization h periods ago x_{t-h} , and $p - 1$ additional lags $x_{t-h-1}, \dots, x_{t-h-p+1}$. For quarterly time series, Hamilton (2018) suggests using $h = 8$ quarters and $p = 4$ lags, so that the regression has the following form:

$$x_t = b_0 + b_1x_{t-8} + b_2x_{t-9} + b_3x_{t-10} + b_4x_{t-11} + u_t. \quad (2)$$

The fitted values and residuals from this linear regression correspond to the estimated Hamilton trend and cycle: $g_t = \hat{x}_t$ and $v_t = \hat{u}_t$.

Contrasting the two filters highlights the drawbacks Hamilton and others find in the HP filter. First, the HP filter performs the trend-cycle decomposition using a matrix/filter that depends solely on the smoothing parameter λ , and not on the properties of the series under consideration. For instance, given a value for λ , the HP filter will detrend a white noise and a random walk with the same exact filter. Second, the HP filter is two-sided, so that the trend and cyclical estimates at any given period depend on past, present, and future values of the original series. In finite samples, this implies that the filtered values at the bounds of the sample are defined differently from those in the middle. Third, the value of the smoothing parameter λ is typically chosen without reference to the observed features of the data.

Hamilton designs his alternative approach so as to avoid these drawbacks. His regression filter estimates a population property of the DGP, the linear regression of the variable on a constant and its past values, avoiding the use of a fixed detrending operator as done by the HP filter. Simply put, the Hamilton cycle recovers a feature of the original DGP, while the HP cycle is a feature of the filtered DGP. The Hamilton filter's one-sided nature also ensures that the trend-cycle decomposition at date t solely relies on the information available up to that date.² Finally, the regression coefficients are estimated from the data, instead of being imposed as in the HP filter.

3 Cyclical Dynamics of Stock Prices and Consumption

Section III.A in Hamilton (2018) illustrates the issues arising when one applies the HP filter to detrend typical economic time series using an empirical example, based on stock prices and consumption. This section reexamines this example by submitting the Hamilton filter to the same evaluation as the HP filter.

Figures 1 and 2 below reproduce Hamilton's Figures 2 and 3 using an extended sample. Data definitions and sources are the same as in Hamilton (2018). Stock prices are measured as 100 times the natural log of the end-of-quarter value for the S&P 500 composite stock price index published

²This is a population statement. Given a finite sample of data, all observations contribute to estimating the regression coefficients, so that each period's trend-cycle decomposition relies on a small amount of future information.

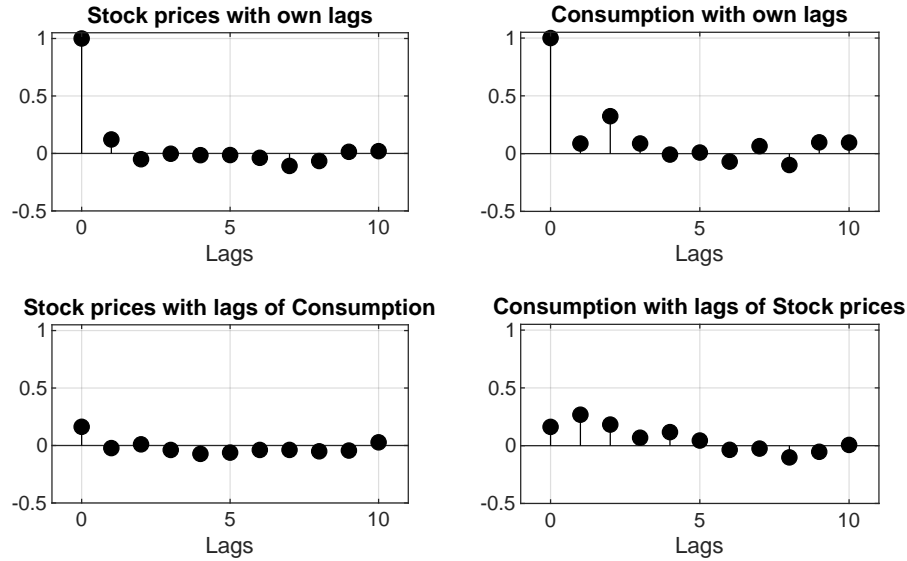


Figure 1: Autocorrelations and cross-correlations for the first differences of stock prices and real consumption.

Notes: Upper left: Autocorrelations of the first difference of end-of-quarter value for log S&P 500 composite stock price index. Upper right: Autocorrelations of the first difference of log real consumption. Lower panels: Cross-correlations.

by Robert Shiller, available online from <http://www.econ.yale.edu/~shiller/data.htm>. Consumption is measured as 100 times the natural log of real personal consumption expenditures from the U.S. National Income and Product Accounts. The data are quarterly and run from 1950Q1 to 2019Q4.

Figure 1 reports the autocorrelation structure for the first differences of log stock prices and real consumption, as well as their cross-correlations. The top panels show that growth in either series is essentially unpredictable, while the bottom panels indicate that after first differencing neither series has strong predictive power for the other. These features are in line with the idea that both variables resemble random walks.

Figure 2 reports the same statistics for the HP cycles extracted from the two series when the smoothing parameter takes the standard value $\lambda = 1,600$. Hamilton emphasizes the presence of a rich auto-regressive structure in the cyclical components of stock prices and real consumption. The cycles are strongly persistent, so that they are predictable from their past values. The cross-correlations also indicate that the two cycles forecast each other. This discrepancy between the original properties of the data and those of the HP cycles embodies Hamilton's claim that the HP filter distorts the series: "The rich dynamics in [the cyclical components] are purely an artifact of the filter itself and tell us nothing about the underlying data-generating process. Filtering takes us

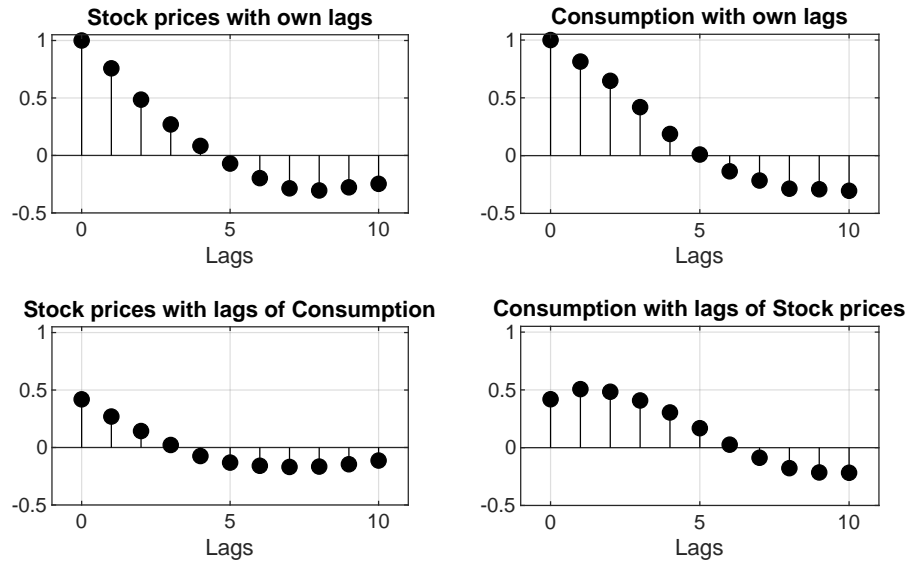


Figure 2: Autocorrelations and cross-correlations for HP-filtered stock prices and real consumption.

Notes: Upper left: Autocorrelations of HP-filtered end-of-quarter value for log S&P 500 composite stock price index. Upper right: Autocorrelations of HP-filtered log real consumption. Lower panels: Cross-correlations. Smoothing parameter: $\lambda = 1,600$.

from the very clean understanding of the true properties of these series [...] to the artificial set of relations [found in the cycles, which] summarize the filter, not the data.”

According to Hamilton, two characteristics of the HP filter combine to generate these spurious dynamics. First, because the HP filter is two-sided, the estimate at each date loads on past, present, and future shocks. It follows that the cyclical component “is both highly predictable (as a result of the dependence on [lagged shocks]) and will in turn predict the future (as a result of dependence on future [shocks]).” Second, the coefficients relating the cyclical estimate to the underlying shocks “are determined solely by the value of λ ,” so that the HP filter effectively imposes dynamics on the data. As noted above, Hamilton overcomes these deficiencies by designing his detrending method as an estimated backward-looking regression. Because the coefficients b_0, \dots, b_4 in (2) are estimated from the data, the filter adapts to the underlying DGP. Because the regression uses only past information, the estimated trend and cyclical components will not depend on future shocks.

Surprisingly, Hamilton (2018) does not report the autocorrelation function for the cycles extracted from stock prices and real consumption by his alternative approach. Yet, evaluating both filters on the same dataset would be a fair comparison. It would also clarify how moving from the two-sided, calibrated HP filter to the one-sided, estimated Hamilton filter affects the cyclical dynamics extracted from the data. Figure 3 fills this gap. Following Hamilton’s recommendation

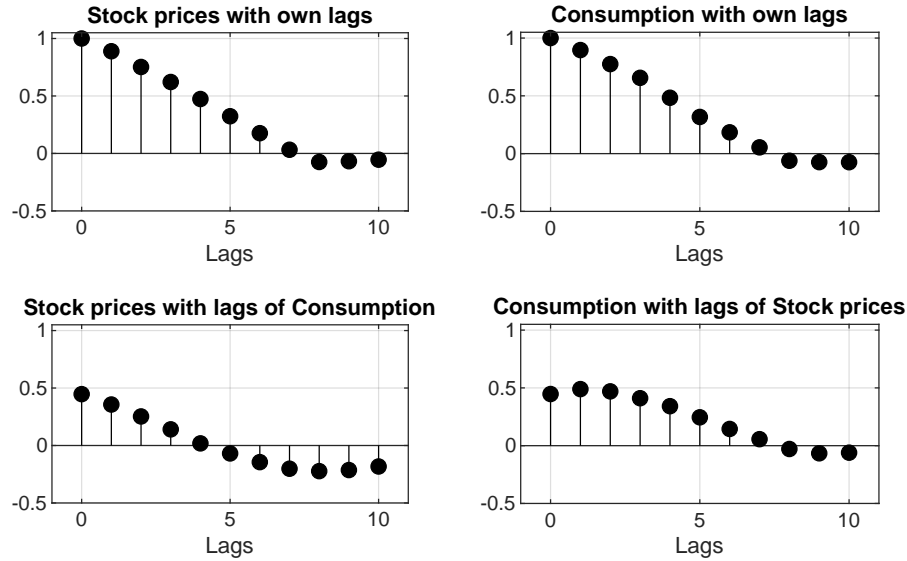


Figure 3: Autocorrelations and cross-correlations for Hamilton-filtered stock prices and real consumption.

Notes: Upper left: Autocorrelations of Hamilton-filtered end-of-quarter value for log S&P 500 composite stock price index. Upper right: Autocorrelations of Hamilton-filtered log real consumption. Lower panels: Cross-correlations. Regression parameters: $p = 4$ and $h = 8$.

for quarterly series, the filter uses $p = 4$ and $h = 8$, so that the cyclical components are obtained by regressing each series at date t on the four most recent observations available at date $t - 8$.

A striking finding is that the Hamilton cycles display virtually the same dynamic behavior as the HP cycles: the cyclical components are very persistent (the autocorrelations decay slowly toward zero); they have strong forecasting power for each other (the cross-correlations are high at several lags); and there are complex dynamics in cross-correlations that are very similar to those found in HP cycles. Observing the magnitude of the correlations, we see that the Hamilton cycles display even *more persistence* and *more cross-variable predictability* than the HP cycles. This is confirmed by the business-cycle statistics reported in Table 1: the first-order autocorrelations of Hamilton-filtered series are 0.89 for stock prices and 0.90 for real consumption, larger than the corresponding values computed from HP-filtered series (0.76 and 0.81). Of course, the HP and Hamilton cycles extracted from stock prices and real consumption are different, but the important point is that they share very similar dynamics.³

These are surprising results, which weaken Hamilton's case against the HP filter. Given Hamil-

³For instance, Table 1 shows that the Hamilton cycles are about twice as volatile as the HP cycles, even though they have similar persistence properties.

ton's claims that the dynamics present in the HP cycles are "spurious [...] relations that have no basis in the underlying data-generating process" (abstract, p. 831) and that HP-filtering "takes us from the very clean understanding of the true properties of these series" to an "artificial set of relations [which] summarize the filter, not the data," then it is difficult not to draw the same conclusions about the similar dynamics found in the Hamilton cycles. Conversely, if one adopts Hamilton's definition of cycles, then comparing Figures 2 and 3 suggests that the HP filter characterizes the cyclical properties of the data relatively well.

The example also clarifies that Hamilton's argument linking the persistence and predictability found in the HP cycles to the calibrated, two-sided nature of the HP filter is partly misleading, since his estimated, one-sided filter produces cyclical estimates with similar persistence and predictability.⁴

From a broader perspective, Hamilton sees the HP decomposition as spurious because it applies the same filter to all series and because it does not adapt the smoothing parameter to data properties. But why would that be a problem? Band-pass filters, which isolate time-series components with specific frequency properties, also present coefficients that depend only on the frequency band of interest, and not on the time series they are applied to (see, e.g., Baxter and King, 1999). In fact, Baxter and King show that, when the smoothing parameter takes the usual value $\lambda = 1,600$, the HP filter closely resembles a high-pass filter targeting frequencies higher than or equal to $\pi/16$. Under this perspective, the fixed-weight design of the HP filter appears neither problematic nor spurious.

Furthermore, Hamilton's definition of spuriousness is quite different from that usually adopted in the literature. Following Harvey and Jaeger (1993) and Cogley and Nason (1995), it is well known that applying the HP filter to integrated processes can generate cyclical dynamics when none are in fact present in the data. When economists argue that the HP filter yields spurious results, they generally have this classic Yule-Slutsky distortion in mind. As the empirical example illustrates, and as shown formally in Hamilton (2018) and below, the Hamilton filter also finds business-cycle dynamics when detrending integrated processes. Hence, this metric too seems to reject the notion that the HP filter is more spurious than the Hamilton filter.

4 Trend Estimates for Stock Prices and Consumption

Hamilton criticizes the two-sided nature of the HP filter, which allows the trend-cycle decomposition at each point in time to use past, current, and future information. As noted above, Hamilton addresses this issue by designing his alternative filter as a backward-looking regression. A convenient way to evaluate the strengths and weaknesses of the two approaches is to look at the filters' implications for the trend components. This is done in Figure 4, which compares the historical path of log stock prices with the estimated HP and Hamilton trends. Reporting the same figure for consumption would be less interesting because the data and the trends are more difficult to disentangle visually, but the conclusion would be similar.

The HP trend is smooth and lies well within the path of the original data, in line with Hodrick and Prescott's notion of a trend. In contrast, the Hamilton trend is quite volatile. As pointed out

⁴See Jonsson (2020a) for a similar view, based on Monte-Carlo simulations.

Table 1: Business-cycle statistics.

Panel A - Volatility and persistence

	Standard deviation	Autocorrelation
<i>Stock prices</i>		
First difference	7.20	0.12
HP filter	9.99	0.76
Hamilton filter	20.95	0.89
<i>Real consumption</i>		
First difference	0.81	0.09
HP filter	1.23	0.81
Hamilton filter	2.73	0.90

Panel B - Contemporaneous correlation

	Stock prices			Real consumption		
	FD	HP	Hamilton	FD	HP	Hamilton
<i>Stock prices</i>						
First difference (FD)	1.00					
HP filter	0.34	1.00				
Hamilton filter	0.31	0.71	1.00			
<i>Real consumption</i>						
First difference	0.16	0.25	0.29	1.00		
HP filter	−0.14	0.42	0.33	0.29	1.00	
Hamilton filter	−0.03	0.36	0.45	0.40	0.66	1.00

Notes: Stock prices: 100 times log S&P 500 composite stock price index. Consumption: 100 times log real consumption. Sample: 1950Q1 to 2019Q4. HP cycles computed with $\lambda = 1,600$; Hamilton cycles computed with $p = 4$ and $h = 8$.

by Hamilton (p. 835), “[m]aking use of unknowable future values [...] is in fact a fundamental reason that HP-filtered series exhibit the visual properties that they do, precisely because they impose patterns that [...] could not be recognized in real time.” For instance, during the 2000s, the stock market experienced two boom-bust episodes, related to the dot-com bubble and to the 2008 financial crisis. Knowledge of future developments allows the HP filter to attribute these temporary episodes to the cyclical component of stock prices. Instead, the Hamilton trend responds to both episodes.

Hamilton (p. 835) finds the smooth HP trend problematic: “Some researchers might be attracted by the simple picture of the ‘long-run’ component of stock prices [provided by the HP trend]. But that picture is just something that their imagination has imposed on the data.” However, it is not clear why we should willingly throw away sample information useful for the trend-cycle decomposition. With hindsight, we know that the boom-bust episodes were temporary, so why would we not use this knowledge to disentangle the trend from the cycle? Indeed, most time-series methods infer unobserved variables through similar two-sided estimates, from band-pass filters (Baxter and King, 1999; Stock and Watson, 1999; Christiano and Fitzgerald, 2003) to unobserved-components models (Harvey, 1985; Harvey and Jaeger, 1993; Harvey and Trimbur, 2003) and DSGE models (Edge, Kiley, and Laforte, 2008; Christiano, Trabandt, and Walentin, 2010; Justiniano, Primiceri,




Figure 4: HP, Hamilton, and Quast-Wolters trends for stock prices.

Notes: Thin black line: 100 times log S&P 500 composite stock price index. Thick blue line: HP trend ($\lambda = 1,600$). Dashed red line: Hamilton trend ($p = 4, h = 8$). Dotted black line: Quast-Wolters trend. Shaded areas: NBER recession dates.

and Tambalotti, 2013). There is nothing peculiar about the HP filter in this dimension.

Furthermore, it is not even clear that the Hamilton filter improves on the HP filter when preserving the information structure is especially important. Indeed, Figure 4 also reveals that the Hamilton trend reacts to movements in the original series with a mechanical two-year delay. This is especially apparent in the later part of the sample, when the Hamilton trend systematically lags the 1995-2000 rise in stock prices, the burst of the dot-com bubble in 2001-2002, the 2003-2007 rebound, and the 2008-2009 financial crisis by a constant window of 8 quarters. Of course, this behavior stems directly from Hamilton's choice to build the trend-cycle decomposition from a 2-year-ahead forecast.⁵

Few economists would view the Hamilton trend in Figure 4 as an attractive estimate of the trend in stock prices. It follows from the additivity of the decomposition that the Hamilton cycle may not be the most attractive estimate of the cyclical component. The mechanical lag in the trend amplifies

⁵Figure 4 also shows the Quast-Wolters trend, represented by a dotted line. The estimate, obtained by averaging Hamilton regressions for 4- to 12-quarter-ahead horizons, presents the same lag in the trend. This is not surprising because averaging 4- to 12-quarter-ahead forecast errors results in the same average 8-quarter delay as found in the original Hamilton trend. A similar lag is present when detrending GDP instead of stock prices; see Figure 2, Panel A, in Quast and Wolters (2022).

the magnitude of the estimated cycle when the data present sharp movements. A more plausible trend component would dampen these movements and, consequently, the volatility of the cyclical component. More importantly, the lag matters when the trend has a direct economic interpretation. Consider, for instance, applying the Hamilton filter to decompose real GDP into potential output (trend) and the output gap (cycle), as Hall and Thomson (2021) and Dritsaki and Dritsaki (2022). Then, the two-year delay in the Hamilton trend implies that estimated potential output will rise during most downturns and start falling when the recovery is already ongoing, in a very mechanical fashion. The economic plausibility of such behavior is clearly dubious.

To illustrate this timing issue, Figure 5 looks at the behavior of stock prices before, during, and after the 2008-2009 recession. The top panel reports the log-level of stock prices between 2005 and 2013, together with the estimated HP and Hamilton trends. The bottom panel reports the corresponding cyclical components. As noted by Hamilton, the HP trend is essentially flat over the period, so that both the pre-recession boom and the post-recession bust in stock prices are viewed as cyclical events. The Hamilton trend, on the other hand, tracks stock prices with a two-year delay: it rises between 2005 and mid-2009, falls abruptly exactly two years after the 2008 stock-market drop, before rising again from 2011 on. This behavior largely shapes the estimated Hamilton cycle, which presents a much larger drop in 2009 compared to the HP cycle, with a negative deviation with respect to trend of about 40% for the HP filter versus 70% for the Hamilton filter. Another noticeable effect occurs in 2011, as the mechanical fall in the Hamilton trend two years after the crisis generates a surprising spike in the Hamilton cycle, that is present neither in the original data nor in the HP cycle.

This discussion suggests that, while the HP filter might rely excessively on *future* information, the Hamilton filter might under-utilize *recent and current* information, resulting in an economically implausible decomposition. Indeed, how plausible is it that agents in the economy, observing the sharp 2008 fall in stock prices, would maintain the view that the trend component is increasing up to mid-2009? And why would they suddenly realize in 2010, two years after the drop and while the economy is already recovering, that the contraction was in fact permanent? It is hard to imagine a credible information structure that would rationalize this inference, which nevertheless underlies Hamilton's definition of the trend-cycle decomposition. While the consequences of this timing issue might vary with the application, it seems important to highlight it to potential users of the Hamilton filter.

5 Simple Properties of the Hamilton Filter

This section demonstrates how the points made above find their origin in the design of the Hamilton filter. Hamilton (2018) presents most of the results discussed here, but he does not emphasize their interpretation. Following Hamilton, the discussion centers on the random-walk case, but the generalization is immediate.

Let x_t and y_t follow two random walks: $x_t = x_{t-1} + \epsilon_t$, $y_t = y_{t-1} + \eta_t$, with ϵ_t and η_t two white-noise processes with variances σ_ϵ^2 and σ_η^2 and covariance $\rho\sigma_\epsilon\sigma_\eta$. For instance, x_t might represent the log of stock prices and y_t the log of real consumption: the two variables have low forecasting




Figure 5: Stock prices around the 2008-2009 recession.

Notes: Black line: 100 times log S&P 500 composite stock price index. Blue lines: HP trend and cycle ($\lambda = 1, 600$). Dashed red lines: Hamilton trend and cycle ($p = 4, h = 8$). Shaded areas: NBER recession dates.

power for each other, and a common shock might induce contemporaneous comovement.⁶

Section IV.B in Hamilton (2018) shows that, in population, the Hamilton filter decomposes x_t and y_t as follows: the trend components are given by

$$g_t^x = x_{t-h}, \quad g_t^y = y_{t-h}, \quad (3)$$

⁶This bivariate random-walk representation provides a good approximation of the data. Letting x_t denote 100 times the log of stock prices and y_t 100 times the log of real consumption, estimating a simple first-order vector autoregression yields the following parameter values:

$$\begin{bmatrix} x_t \\ y_t \end{bmatrix} = \begin{bmatrix} -0.76 \\ 2.47 \end{bmatrix} + \begin{bmatrix} 0.98 & 0.03 \\ 0.00 & 0.99 \end{bmatrix} \begin{bmatrix} x_{t-1} \\ y_{t-1} \end{bmatrix} + \begin{bmatrix} \hat{\epsilon}_t \\ \hat{\eta}_t \end{bmatrix}, \quad \text{Var} \begin{bmatrix} \hat{\epsilon}_t \\ \hat{\eta}_t \end{bmatrix} = \begin{bmatrix} 51.89 & 0.99 \\ 0.99 & 0.62 \end{bmatrix}.$$

The implied correlation between the innovations is $\hat{\rho} = 0.17$.

while the cyclical components verify

$$v_t^x = x_t - x_{t-h} = \sum_{j=0}^{h-1} \epsilon_{t-j}, \quad v_t^y = y_t - y_{t-h} = \sum_{j=0}^{h-1} \eta_{t-j}. \quad (4)$$

Equation (3) indicates that, in the random-walk case, the Hamilton trend is exactly the value of the variable h periods ago. This property explains the mechanical two-year delay found in the trend estimates discussed in Section 4. With richer autoregressive processes, the Hamilton trend would be a more complex function of past observations, but the crucial point that the trend estimate at date t does not use information available from periods $t, t-1, \dots, t-h+1$ would remain true.⁷

Under the assumption that ϵ_t and η_t are white-noise processes, equation (4) implies the following second moments for the cyclical components:

$$\text{var}(v_t^x) = h\sigma_\epsilon^2, \quad \text{var}(v_t^y) = h\sigma_\eta^2, \quad (5)$$

$$\text{corr}(v_t^x, v_{t-j}^x) = \text{corr}(v_t^y, v_{t-j}^y) = \frac{h-j}{h} \text{ if } j = 0, 1, \dots, h, \quad = 0 \text{ if } j \geq h+1, \quad (6)$$

$$\text{corr}(v_t^x, v_{t-j}^y) = \text{corr}(v_t^y, v_{t-j}^x) = \frac{(h-j)\rho}{h} \text{ if } j = 0, 1, \dots, h, \quad = 0 \text{ if } j \geq h+1. \quad (7)$$

These expressions imply that: (i) The Hamilton filter extracts persistent cycles out of random walks. (ii) It extracts interrelated cycles out of correlated random walks. (iii) The variances, the persistence, and the joint dynamics of the cycles are entirely determined by the forecast horizon h . All three properties follow from Hamilton's definition of the cycle as a h -step-ahead forecast error: as shown by (4), two realizations of v_t^x and v_t^y separated by j periods share $h-j$ common innovations when $j \leq h$, necessarily leading to serial correlation and comovement. These properties also explain the dynamics found in Hamilton-filtered stock prices and consumption.

This analysis clarifies that, while the Hamilton filter seems to attribute shocks whose effects persist less than h periods to the cycle and shocks whose effects persist more than h periods to the trend, this distinction is not entirely accurate. Instead, what the Hamilton filter does is to attribute initially *all* unexpected shocks to the cyclical component, before assigning the effects that still persist after h periods to the trend. This sharp cutoff at lag h gives rise to the mechanical delay in the Hamilton trend.

6 Conclusion

Hamilton (2018) argues that economists should stop using the HP filter. This comment suggests that Hamilton's case is not entirely convincing: the Hamilton cycles present the same kind of filter-induced dynamics as Hamilton criticizes in the HP cycles, and the Hamilton trend lags the data by construction, which may threaten the economic plausibility of the decomposition. These two points clarify that the Hamilton filter provides a very particular trend-cycle decomposition, that might not outperform the HP filter on a systematic basis.

⁷The only case in which the Hamilton trend does not feature a mechanical delay relative to the data is when the value of x_t does not depend on the lagged observations $x_{t-h}, x_{t-h-1}, \dots, x_{t-h-p+1}$. Such a configuration is unlikely to arise in macroeconomic applications.

References

- Baxter, M., & King, R. G. (1999).** “Measuring Business Cycles: Approximate Band-Pass Filters For Economic Time Series.” *The Review of Economics and Statistics*, 81(4): 575–593. DOI: [10.1162/003465399558454](https://doi.org/10.1162/003465399558454).
- Canova, F. (1998).** “Detrending and Business Cycle Facts.” *Journal of Monetary Economics*, 41(3): 475–512. DOI: [10.1016/S0304-3932\(98\)00006-3](https://doi.org/10.1016/S0304-3932(98)00006-3).
- Christiano, L. J., & Fitzgerald, T. J. (2003).** “The Band Pass Filter.” *International Economic Review*, 44(2): 435–465. DOI: [10.1111/1468-2354.t01-1-00076](https://doi.org/10.1111/1468-2354.t01-1-00076).
- Christiano, L. J., Trabandt, M., & Walentin, K. (2010).** “DSGE Models for Monetary Policy Analysis” In: Friedman, B. M. & Woodford, M. (Eds.), *Handbook of Monetary Economics*, vol. 3, chap. 7, pp. 285–367. Elsevier. DOI: [10.1016/B978-0-444-53238-1.00007-7](https://doi.org/10.1016/B978-0-444-53238-1.00007-7).
- Cogley, T., & Nason, J. M. (1995).** “Effects of the Hodrick-Prescott Filter on Trend and Difference Stationary Time Series: Implications for Business Cycle Research.” *Journal of Economic Dynamics and Control*, 19(1-2): 253-278. DOI: [10.1016/0165-1889\(93\)00781-X](https://doi.org/10.1016/0165-1889(93)00781-X).
- de Jong, R.M., & Sakarya, N. (2016).** “The Econometrics of the Hodrick-Prescott Filter.” *The Review of Economics and Statistics*, 81(4): 575–593. DOI: [10.1162/REST_a_00523](https://doi.org/10.1162/REST_a_00523).
- Dritsaki, M., & Dritsaki, C. (2022).** “Comparison of HP Filter and the Hamilton’s Regression.” *Mathematics*, 10(8). DOI: [10.3390/math10081237](https://doi.org/10.3390/math10081237).
- Edge, R. M., Kiley, M. T., & Laforte, J.-P. (2008).** “Natural Rate Measures in an Estimated DSGE Model of the U.S. Economy.” *Journal of Economic Dynamics and Control*, 32(8): 2512-2535. DOI: [10.1016/j.jedc.2007.09.011](https://doi.org/10.1016/j.jedc.2007.09.011).
- Hall, V. B., & Thomson, R. G. (2021).** “Does Hamilton’s OLS Regression Provide a Better Alternative” to the Hodrick-Prescott Filter? A New Zealand Business Cycle Perspective.” *Journal of Business Cycle Research*, 17(2): 151-183. DOI: [10.1007/s41549-021-00059-1](https://doi.org/10.1007/s41549-021-00059-1).
- Hamilton, J. D. (2018).** “Why You Should Never Use the Hodrick-Prescott Filter.” *The Review of Economics and Statistics*, 100(5): 831-843. DOI: [10.1162/rest_a_00706](https://doi.org/10.1162/rest_a_00706).
- Hamilton, J. D., & Xi, J. (2024).** “Principal Component Analysis for Nonstationary Series.” Manuscript. URL: <https://www.nber.org/papers/w32068>
- Harvey, A. C. (1985).** “Trends and Cycles in Macroeconomic Time Series.” *Journal of Business & Economic Statistics*, 8(3): 231-247. DOI: [10.1080/07350015.1985.10509453](https://doi.org/10.1080/07350015.1985.10509453).
- Harvey, A. C., & Jaeger, A. (1993).** “Detrending, Stylized Facts and the Business Cycle.” *Journal of Applied Econometrics*, 8(3): 231–247. DOI: [10.1002/jae.3950080302](https://doi.org/10.1002/jae.3950080302).

Harvey, A. C., & Trimbur, T. M. (2003). “General Model-Based Filters for Extracting Cycles and Trends in Economic Time Series.” *The Review of Economics and Statistics*, 85(2): 244-255. DOI: [10.1162/003465303765299774](https://doi.org/10.1162/003465303765299774).

Hodrick, R. J. (2020). “An Exploration of Trend-Cycle Decomposition Methodologies in Simulated Data.” NBER Working Papers 26750, National Bureau of Economic Research, Inc. URL: <https://www.nber.org/papers/w26750>.

Hodrick, R. J., & Prescott, E. (1981). “Post-War U.S. Business Cycles: An Empirical Investigation.” Discussion Papers 451, Northwestern University, Center for Mathematical Studies in Economics and Management Science. PDF: <https://www.kellogg.northwestern.edu/research/math/papers/451.pdf>.

Hodrick, R. J., & Prescott, E. (1997). “Postwar U.S. Business Cycles: An Empirical Investigation.” *Journal of Money, Credit and Banking*, 29(1): 1-16. DOI: [10.2307/2953682](https://doi.org/10.2307/2953682).

Jonsson, K. (2020a). “Cyclical Dynamics and Trend/Cycle Definitions: Comparing the HP and Hamilton Filters” *Journal of Business Cycle Research*, 16(2): 151-162. DOI: [10.1007/s41549-020-00039-x](https://doi.org/10.1007/s41549-020-00039-x).

Jonsson, K. (2020b). “Real-Time US GDP Gap Properties using Hamilton’s Regression-Based Filter.” *Empirical Economics*, 59(1): 307-314. DOI: [10.1007/s00181-019-01631-6](https://doi.org/10.1007/s00181-019-01631-6).

Justiniano, A., Primiceri, G. E., & Tambalotti, A. (2013). “Is There a Trade-Off between Inflation and Output Stabilization?” *American Economic Journal: Macroeconomics*, 5(2): 1-31. DOI: [10.1257/mac.5.2.1](https://doi.org/10.1257/mac.5.2.1).

Mei, Z., Phillips, P. C. B., & Shi, Z. (2022). “The Boosted HP Filter is More General than You Might Think.” Papers 2209.09810, arXiv.org. URL: <https://arxiv.org/abs/2209.09810>.

Phillips, P. C. B., & Shi, Z. (2021). “Boosting: Why You Can Use The Hp Filter.” *International Economic Review*, 62(2): 521-570. DOI: [10.1111/iere.12495](https://doi.org/10.1111/iere.12495).

Quast, J., & Wolters, M. H. (2022). “Reliable Real-Time Output Gap Estimates Based on a Modified Hamilton Filter.” *Journal of Business & Economic Statistics*, 40(1): 152-168. DOI: [10.1080/07350015.2020.1784747](https://doi.org/10.1080/07350015.2020.1784747).

Ravn, M. O., & Uhlig, H. (2002). “On Adjusting the Hodrick-Prescott Filter for the Frequency of Observations.” *The Review of Economics and Statistics*, 84(2): 371-375. DOI: [10.1162/003465302317411604](https://doi.org/10.1162/003465302317411604).

Schuler, Y. S. (2021). “On the Cyclical Properties of Hamilton’s Regression Filter.” Discussion Papers 03/2018, Deutsche Bundesbank. PDF: <https://www.bundesbank.de/resource/blob/704152/7dc6c0968c9581dff4878e724ac6460b/mL/2018-02-14-dkp-03-data.pdf>

Stock, J. H., & Watson, M. W. (1999). “Business Cycle Fluctuations in US Macroeconomic

Time Series.” *The Review of Economics and Statistics*, 81(4): 575–593. DOI: [10.1016/S1574-0048\(99\)01004-6](https://doi.org/10.1016/S1574-0048(99)01004-6).

Stock, J. H., & Watson, M. W. (1999). “Business Cycle Fluctuations in US Macroeconomic Time Series” In: Taylor, J. B. & Woodford, M. (Eds.), *Handbook of Macroeconomics*, vol. 1, chap. 1, pp. 3–64. Elsevier. DOI: [10.1016/S1574-0048\(99\)01004-6](https://doi.org/10.1016/S1574-0048(99)01004-6).