CAHIER D'ÉTUDES
WORKING PAPER

N° 162

WHY YOU SHOULD NEVER USE THE HODRICK-PRESCOTT FILTER: COMMENT

ALBAN MOURA

AUGUST 2022

BANQUE CENTRALE DU LUXEMBOURG

EUROSYSTÈME

WHY YOU SHOULD NEVER USE THE HODRICK-PRESCOTT FILTER: COMMENT

ALBAN MOURA

Abstract. Hamilton (2018) argues that one should never use the Hodrick-Prescott (HP) filter, given its drawbacks and the existence of a better alternative. This comment shows that the main drawback Hamilton finds in the HP filter, the presence of filter-induced dynamics in the estimate of the cyclical component, is also a key feature of the alternative filter proposed by Hamilton. As with the HP filter, the Hamilton filter applied to a random walk extracts a cyclical component that is highly predictable, that can predict other variables, and whose properties reflect as much the filter as the underlying data-generating process. In addition, the Hamilton trend lags the data by construction and there is some arbitrariness in the choice of a key parameter defining the filter. Therefore, a more balanced assessment is that the HP and Hamilton filters provide different ways to look at the data, with neither being clearly superior from a practical perspective.

JEL Codes: B41, C22, E32.

Keywords: HP filter; Hamilton filter; business cycles; detrending; filtering.

July 2022. Banque centrale du Luxembourg, Département Économie et Recherche, 2 boulevard Royal, L-2983 Luxembourg (e-mail: alban.moura@bcl.lu). I thank James Hamilton and Robert Hodrick for very helpful feedback. I also thank Patrick Fève, Paolo Guarda, Olivier Pierrard, and colleagues at the Banque centrale du Luxembourg for useful comments. The views presented in this paper are personal and should not be reported as those of the Banque centrale du Luxembourg or the Eurosystem.

Résumé Non Technique

Les principales variables macro-économiques, telles que le produit intérieur brut (PIB) ou le taux de chômage, sont souvent décomposées en séparant leur tendance de long terme d'une composante cyclique. Que l'objectif soit d'estimer une composante cyclique d'un intérêt particulier (par exemple l'écart de production, ou *output gap*) ou de comparer les propriétés cycliques d'un modèle économique à celles des données, la difficulté reste la même : il existe une infinité de décompositions possibles et il appartient au chercheur de sélectionner la plus adaptée.

En 1981, Hodrick et Prescott ont proposé une méthode de décomposition, connue depuis sous le nom de filtre HP, qui est largement employée dans la littérature académique, dans les banques centrales et dans l'industrie financière. Néanmoins, cette méthode est sujette à d'importantes limitations qui compliquent son emploi. En particulier, le filtre HP peut déformer les propriétés de la composante cyclique et souffre d'un biais en fin d'échantillon, ce qui ajoute à l'incertitude des analyses en temps réel. Finalement, la décomposition calculée par le filtre HP dépend d'un paramètre dont la valeur doit être fixée a priori, de manière assez arbitraire.

Dans une contribution récente, Hamilton (2018) énumère ces limitations pour arriver à la conclusion que les économistes ne devraient jamais utiliser le filtre HP. En complément des critiques précédentes, Hamilton propose une nouvelle technique pour décomposer une série temporelle en tendance et cycle. Ce filtre de Hamilton constituerait une alternative supérieure au filtre HP en résolvant tous les problèmes de ce dernier.

Le présent travail montre que le filtre de Hamilton n'apporte qu'une amélioration très limitée par rapport au filtre HP. En se basant à la fois sur un exemple empirique étudié par Hamilton lui-même et sur une discussion analytique, il est démontré que la composante cyclique extraite par le filtre de Hamilton présente des caractéristiques statistiques voisines de celles du cycle obtenu par le filtre HP. Ceci suggère que le filtre de Hamilton ne permet pas de résoudre le problème de déformation de la composante cyclique. De plus, la tendance estimée par le filtre de Hamilton présente un décalage mécanique de plusieurs observations par rapport aux données. Enfin, les résultats du filtre de Hamilton dépendent eux aussi crucialement d'un paramètre dont la valeur est choisie de manière arbitraire. En conséquence, le filtre de Hamilton devrait être considéré comme un outil complémentaire au filtre HP, plutôt que comme une alternative clairement supérieure.

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 cyclical estimates at the boundaries of the sample are not reliable. (c) Common choices for the smoothing parameter are arbitrary.[1] 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. Hamilton's dismissal of the HP filter runs counter to widespread practice in academia, policy-making institutions, and the private sector.

Given the attention received by Hamilton's paper (more than 1,100 citations on Google Scholar as of May 2022), it seems important to fully understand how his alternative filter works and how it compares to the HP filter. This comment focuses on filter-induced dynamics in the estimated cycles, showing that the Hamilton filter suffers from a similar drawback as the HP filter. The issue is especially clear for the detrending of random walks, which happens to be Hamilton's focus. In this case, one can show analytically that the Hamilton cycle exhibits strong serial correlation, is highly predictable from its past, and can predict other variables. The same empirical example Hamilton uses to dismiss the HP filter, the filtering of consumption and stock prices, confirms the existence of such dynamics: strikingly, the Hamilton filter extracts cyclical components whose persistence and comovements mirror those found in HP cycles. The analytical discussion confirms that these properties originate from the filter: persistence follows mechanically from Hamilton defining the cyclical component as a multi-step-ahead forecast error in a linear regression, a variable almost necessarily autocorrelated. Therefore, focusing on a backward-looking filter like the Hamilton regression is not enough to avoid creating seemingly spurious persistence and predictability patterns in the estimated cyclical component. Furthermore, the trend estimated by the Hamilton approach lags the data by construction.

The analytical discussion also shows how the choice of the forecast horizon h in the Hamilton regression shapes the properties of the estimated cycles. With larger values of h, the Hamilton filter extracts more volatile and more persistent cyclical components from the same DGP. Yet, there is no clear criterion to pick up the forecast horizon. Hamilton (2018) suggests using $h = 8$ for quarterly data, on the grounds that a two-year horizon should be the benchmark for business-cycle analysis and that a multiple of 4 helps dealing with seasonal patterns. But there is some arbitrariness in this choice. For instance, one could invoke Angeletos, Collard, and Dellas' (2020) finding that a shock dominating at business-cycle

[1] These limitations of the HP filter have been known for some time in the literature; see, e.g., Nelson and Plosser (1982), Harvey and Jaeger (1993), and Cogley and Nason (1995).

frequencies (6–32 quarters) has a footprint in the time domain that peaks within a year to motivate using $h = 4$ instead. For macroeconomic time series behaving like random walks, this change would roughly halve the volatility and the persistence of estimated Hamilton cycles.

Overall, these results make it difficult to argue that the Hamilton filter will systematically outperform the HP filter in macroeconomic applications. On the one hand, the Hamilton filter improves on the HP filter by yielding consistent estimates at the boundaries of the sample, even though the empirical results reported in Hall and Thomson (2021) suggest that the gain can be marginal in practice. On the other hand, Hamilton cycles present filter-induced dynamics that resemble those found in HP cycles and the Hamilton filter relies on a rather arbitrary choice for the forecast horizon just as the HP filter depends on the smoothing parameter. Therefore, a more balanced assessment is that the two filters provide different views of the data, and that whether one of the two views is more interesting remains an open question. (This idea is borrowed from Burnside, 1998.) Until this question is answered, the HP and Hamilton filters should be regarded as complementary tools for business-cycle analysis.

Previous studies offer a critical evaluation of the Hamilton filter. For instance, Schüler (2018) uses spectral methods to show that the Hamilton filter emphasizes cycles with longer duration than typical business cycles and that it mutes shorter cycles, leading to a failure to reproduce the chronology of U.S. business cycles. Hodrick (2020) applies simulation methods to compare the Hamilton filter with alternative detrending strategies, including the HP filter, and finds that the Hamilton filter yields better cyclical estimates for simple models, while the HP filter performs better for complex models. Compared to these studies, this comment focuses more narrowly on the mechanical impact the Hamilton filter has on estimated cycles, which is of particular interest for applied economists.

2. The HP and Hamilton Filters

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

Both the HP and the Hamilton filter 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 series that does not differ much from the observed series. This objective can be formalized by choosing g_t as the solution to the following program:

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

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

On the other hand, the Hamilton filter 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 the observed variable x_t on a constant, the realization h periods ago x_{t-h}, and $p-1$ additional lags $x_{t-h-1}, \ldots, x_{t-h-p+1}$. For quarterly time series, Hamilton (2018) suggests using $h = 8$ quarters and $p = 4$ lags, so that the regression has the following form:

$$x_t = b_0 + b_1 x_{t-8} + b_2 x_{t-9} + b_3 x_{t-10} + b_4 x_{t-11} + u_t. \tag{2}$$

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

3. Cyclical Dynamics of Stock Prices and Consumption

Section III.A in Hamilton (2018) questions the appropriateness of applying the HP filter to detrend typical economic time series. (Unless otherwise specified, all quotes reported in this section are from Hamilton's Section III.A. p. 833.) Hamilton argues that many such series resemble random walks and shows that detrending a random walk with the HP filter generates spurious dynamics, in the sense that the extracted cycle features high persistence, in contrast to the serially uncorrelated innovations of the underlying process. He provides 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 to 100 times the natural log of the end-of-quarter value for the S&P composite stock price index published 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 series follow 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 $\lambda = 1,600$. As emphasized by Hamilton, the cyclical components of stock prices and real consumption display strong persistence, so that they are predictable from their past values. Furthermore, the cross-correlograms exhibit rich autoregressive structures with wave-like patterns, indicating that the cyclical component from one series can be forecast from the cyclical component of the other.

The discrepancy between the dynamic properties of the first differences of the data and those of 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 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 cyclical 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 instead of adapting to the specific time series at hand. To overcome these deficiencies, Hamilton designs his detrending approach as an estimated backward-looking regression. Because the coefficients $b_0, \ldots, b_4$ in equation (2) are estimated from the data, the filter has the flexibility to adapt to the underlying DGP. Because the regression uses only past information, the estimated cyclical component 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 seem like 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 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-filtered series. Focusing on the absolute size of the correlations, there appears to be even *more persistence* and *more cross-variable predictability* in Hamilton cycles than in HP cycles. This is confirmed by the 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 respectively).

Of course, the HP and Hamilton cycles extracted from stock prices and real consumption are different, as shown in Figure 4 and Table 1. For instance, the Hamilton cycles are about twice as volatile and exhibit stronger persistence than the HP cycles. Nevertheless, the contemporaneous correlation between the HP and Hamilton cycles extracted from a given variable is high: 0.71 for stock prices and 0.66 for real consumption. Although Table 1 does not report these statistics, the correlation between the HP and Hamilton cycles is maximal contemporaneously, so that neither series seem to lead or lag the other. Finally, comparison of the autocorrelation functions highlights their similar dynamics.

These are surprising results, which weaken Hamilton's case for his alternative to the HP filter. If one accepts Hamilton's view that the serial correlation and predictability found in the HP cycles are artificial, then it is difficult not to draw the same conclusion regarding the same features in the Hamilton cycles. Alternatively, if one is willing to accept the cyclical component from Hamilton's filter, then comparing Figure 2 to 3 would indicate that the HP filter does at least a satisfactory job estimating the cyclical properties of the data. In either case, based on this empirical example, it is unclear why one would choose the Hamilton filter over the HP filter.

Another important property appears in Figure 5, 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 due to the smoothness of the series.) Unsurprisingly, the HP trend is a smooth variable that lies well inside the path of the actual time series. As Hamilton argues, the HP filter generates this smooth trend by making use of both past, present, and future observations of the variable. If Hamilton regards this behavior as a drawback ("HP-filtered series exhibit the visual properties that they do, precisely because they impose patterns that are not a feature of the data-generating process and could not be recognized in real time", p. 835), any smoothed estimate would also make use of all sample information, including future observations not available in real time. For instance, output-gap estimates based on DSGE models (e.g., Edge, Kiley, and Laforte, 2008; Justiniano, Primiceri, and Tambalotti, 2013) or trend-cycle decompositions based on unobserved component models (e.g., Harvey and Trimbur, 2003; Harvey, Trimbur, and Van Dijk, 2007) also exploit all available information without raising particular suspicion.

As expected, the Hamilton trend is not as smooth as the HP trend because it does not make use of future information. However, a direct consequence of using an 8-quarter-ahead forecast is that the Hamilton trend reacts to economic developments with a mechanical two-year delay. This is especially easy to spot in the later part of the sample: the 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. To many economists, this pattern would appear as a drawback and few people would consider the red line in Figure 4 as the best possible estimate of the trend in stock prices. Therefore, while one can sympathize with Hamilton's objective to avoid using future information in the trend-cycle decomposition, one should also realize that his filter generates peculiar timing implications.

4. Simple Properties of the Hamilton Filter

This section elaborates on the previous example by showing how the Hamilton filter mechanically shapes the statistical properties of the cyclical component it estimates. The discussion focuses on the random-walk case for simplicity, but the generalization is obvious.

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 power for each other, but a common shock might induce contemporaneous comovement.[2]

Section IV.B in Hamilton (2018) shows that, in population, the cyclical components obtained by applying the Hamilton filter to x_t and y_t are equal to the forecast errors at horizon h:

$$v_t^x = x_t - x_{t-h} = \sum_{j=0}^{h-1} \epsilon_{t-j}, \qquad v_t^y = y_t - y_{t-h} = \sum_{j=0}^{h-1} \eta_{t-j}. \tag{3}$$

Since ϵ_t and η_t are white noise processes, it is straightforward to compute the second moments of these random variables:

$$\mathrm{Var}(v_t^x) = h\sigma_\epsilon^2, \qquad \mathrm{Var}(v_t^y) = h\sigma_\eta^2,$$

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

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

These moments highlight three key properties of the Hamilton filter. First, it extracts a persistent cycle out of a random walk. Second, it extracts interrelated cycles out of correlated random walks. Third, the variances, the persistence, and the joint dynamics of the cycles are mechanically determined by the forecast horizon h used by the filter. All three properties

[2]This 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:

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

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

follow from Hamilton's definition of the cycle as a forecast error: as shown by equation (3), two realizations of v_t^x and v_t^y separated by j periods share $h-j$ common innovations when $j \leq h$, which necessarily leads to serial correlation and comovement. While the random-walk setting makes this feature especially apparent given the permanent effect of shocks, similar persistence and predictability would arise from applying the Hamilton filter to more general ARIMA processes.

These properties explain the dynamics found in Hamilton-filtered stock prices and real consumption. Thus, the empirical example in Section 3 represents the normal behavior of the filter and demonstrates that filter-induced dynamics are as present in Hamilton cycles as in HP cycles. This finding challenges Hamilton's claim that his alternative detrending approach can avoid all drawbacks of the HP filter.

The choice of the forecast horizon h provides another illustration of filter-induced dynamics. Hamilton (2018, Section IV.C, p. 838) motivates his recommendation of $h = 8$ for quarterly data from two arguments: (i) having h be a multiple of 4 is useful to purge the estimated cycles from potential seasonal patterns, and (ii) the notion that "a two-year horizon should be the standard benchmark" for business-cycle analysis. If the first argument is objective, the second one seems less empirically or theoretically grounded. For instance, the frequency band typically associated with business cycles ranges from six quarters to eight years (e.g., Stock and Watson, 1999), so that it is not clear what makes $h = 8$ the benchmark choice.

Alternatively, Angeletos, Collard, and Dellas (2020) show that the shock which contributes most to the variance of key macroeconomic variables over the standard business-cycle frequency band (6-32 quarters) has a footprint in the time domain that peaks within a year. This result suggests that setting $h = 4$ might constitute an interesting alternative for quarterly data. For the many macroeconomic time series featuring random-walk like behavior, this change would roughly halve the variance and the persistence of the estimated cycles compared to Hamilton's benchmark of $h = 8$. More generally, while the choice of the forecast horizon in the Hamilton filter has important consequences for the estimated trend-cycle decomposition, there is no conclusive argument to settle on the best value. Thus, the arbitrariness in choosing h for the Hamilton filter mirrors that in choosing λ for the HP filter.

Overall, both the empirical example and analytical discussion emphasize that the population characteristic estimated by the Hamilton filter, the forecast error in the linear regression of the variable at $t+h$ on a constant and p lags, corresponds to a very particular view of the trend-cycle decomposition. In particular, the Hamilton approach has mechanical effects on the timing of the estimates relative to the data, as well as on the persistence and volatility of the estimated cyclical component. In the important random-walk case, these filter-induced dynamics resemble those found in HP cycles when it comes to "artificial" serial correlation and "spurious" predictability. Highlighting these properties of the Hamilton filter to the audience of applied economists is the main purpose of this comment.

5. Conclusion

Hamilton (2018) argues that economists should stop using the HP filter. However, the alternative filter proposed by Hamilton can be criticized using essentially the same arguments invoked against the HP filter, namely the presence of filter-induced dynamics in the estimated cycles and the relative arbitrariness of a key parameter choice. Furthermore, the trends estimated by the Hamilton approach will lag the data by construction. These results cast doubts on Hamilton's claim that his filter will always outperform the HP filter in practice. A more balanced assessment is that the two filters provide different ways to look at the cyclical properties of the data, with neither appearing to be clearly superior.

More generally, there is really nothing new or wrong in recognizing that detrending data affects its statistical properties in a way that depends on the chosen approach: using a polynomial time trend, the HP filter, a band-pass filter, or the Hamilton filter to separate the trend from the cycle will necessarily lead to different estimates of the cyclical component. Canova (1998) illustrated this point nicely twenty years ago. As stressed by Burnside (1998), this is not a major issue when the goal is to relate stationary economic models to non-stationary data, since it is always possible to compare filtered real-world data with filtered series from the model. For instance, the widespread software package Dynare (2011) automatically computes moments for HP- and band-pass filtered series simulated from DSGE models, allowing for straightforward comparison between theory and data. It would be useful to also implement the Hamilton filter, providing economists yet another window through which they can compare their models to reality.

References

Adjemian, S., H. Bastani, M. Juillard, F. Karamé, J. Maih, F. Mihoubi, W. Mutschler, G. Perendia, J. Pfeifer, M. Ratto, and S. Villemot (2011): "Dynare: Reference Manual Version 4," Dynare Working Papers 1, CEPREMAP.

Angeletos, G.-M., F. Collard, and H. Dellas (2020): "Business-Cycle Anatomy," *American Economic Review*, 110(10), 3030–3070.

Burnside, C. (1998): "Detrending and Business Cycle Facts: A Comment," *Journal of Monetary Economics*, 41(3), 513–532.

Canova, F. (1998): "Detrending and Business Cycle Facts," *Journal of Monetary Economics*, 41(3), 475–512.

Cogley, T., and J. M. Nason (1995): "Effects of the Hodrick-Prescott Filter on Trend and Difference Stationary Time Series: Implications for Business Cycle Research," *Journal of Economic Dynamics and Control*, 19(1-2), 253–278.

Edge, R. M., M. T. Kiley, and J.-P. Laforte (2008): "Natural Rate Measures in an Estimated DSGE Model of the U.S. Economy," *Journal of Economic Dynamics and Control*, 32(8), 2512–2535.

Hall, V. B., and P. Thomson (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.

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

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

Harvey, A. C., and T. M. Trimbur (2003): "General Model-Based Filters for Extracting Cycles and Trends in Economic Time Series," *The Review of Economics and Statistics*, 85(2), 244–255.

Harvey, A. C., T. M. Trimbur, and H. K. Van Dijk (2007): "Trends and Cycles in Economic Time Series: A Bayesian Approach," *Journal of Econometrics*, 140(2), 618–649.

Hodrick, R. J. (2020): "An Exploration of Trend-Cycle Decomposition Methodologies in Simulated Data," NBER Working Papers 26750, National Bureau of Economic Research, Inc.

Hodrick, R. J., and E. C. Prescott (1981): "Post-War U.S. Business Cycles: An Empirical Investigation," Discussion Papers 451, Northwestern University.

——— (1997): "Postwar U.S. Business Cycles: An Empirical Investigation," *Journal of Money, Credit and Banking*, 29(1), 1–16.

Justiniano, A., G. E. Primiceri, and A. Tambalotti (2013): "Is There a Trade-Off between Inflation and Output Stabilization?," *American Economic Journal: Macroeconomics*, 5(2), 1–31.

Nelson, C. R., and C. I. Plosser (1982): "Trends and Random Walks in Macroeconmic Time Series : Some Evidence and Implications," *Journal of Monetary Economics*, 10(2), 139–162.

Schüler, Y. (2018): "On the Cyclical Properties of Hamilton's Regression Filter," Discussion Papers 03/2018, Deutsche Bundesbank.

Stock, J. H., and M. W. Watson (1999): "Business Cycle Fluctuations in US Macroeconomic Time Series," in *Handbook of Macroeconomics*, ed. by J. B. Taylor, and M. Woodford, vol. 1 of *Handbook of Macroeconomics*, chap. 1, pp. 3–64. Elsevier.

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

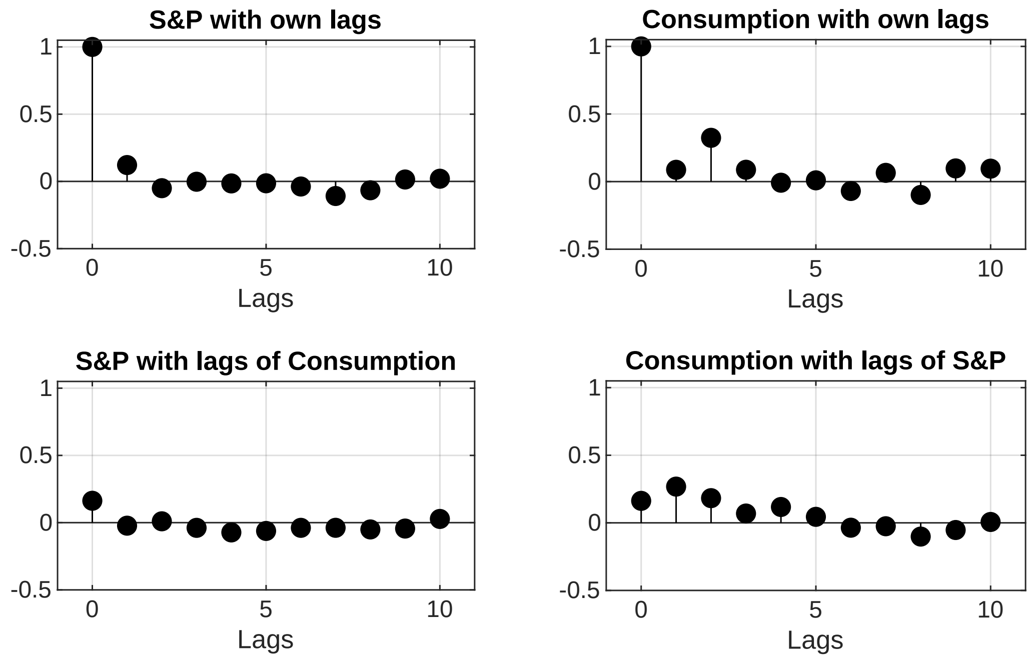

Upper left: Autocorrelations of the first difference of end-of-quarter value for S&P composite. Upper right: Autocorrelations of the first difference of real consumption. Lower panels: Cross-correlations.

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

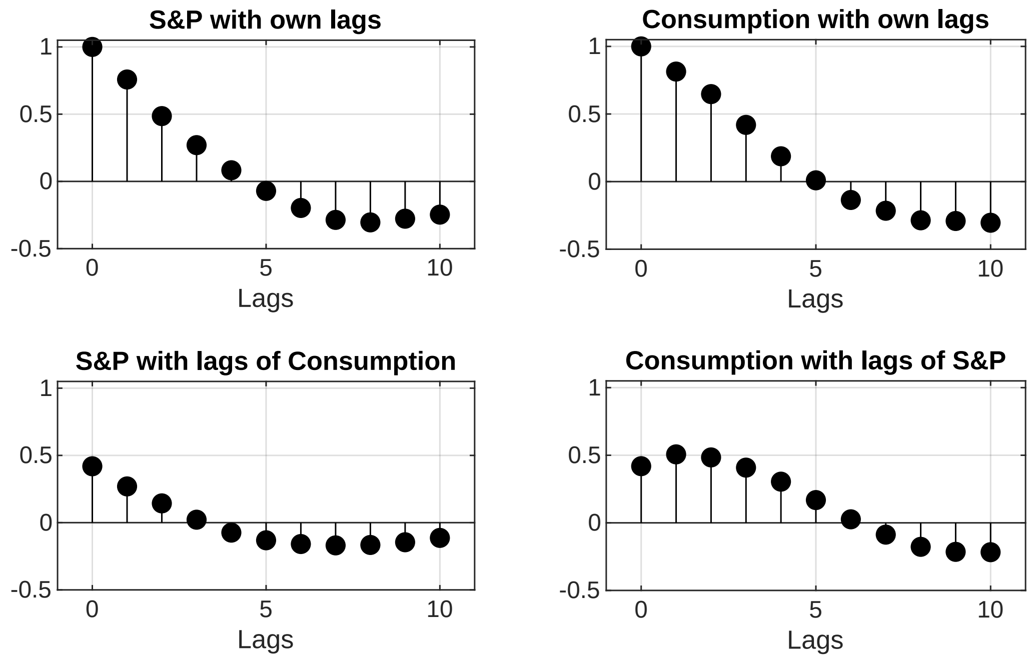

Upper left: Autocorrelations of HP-filtered end-of-quarter value for log S&P composite. Upper right: Autocorrelations of HP-filtered log real consumption. Lower panels: Cross-correlations. Smoothing parameter: $\lambda = 1,600$.

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

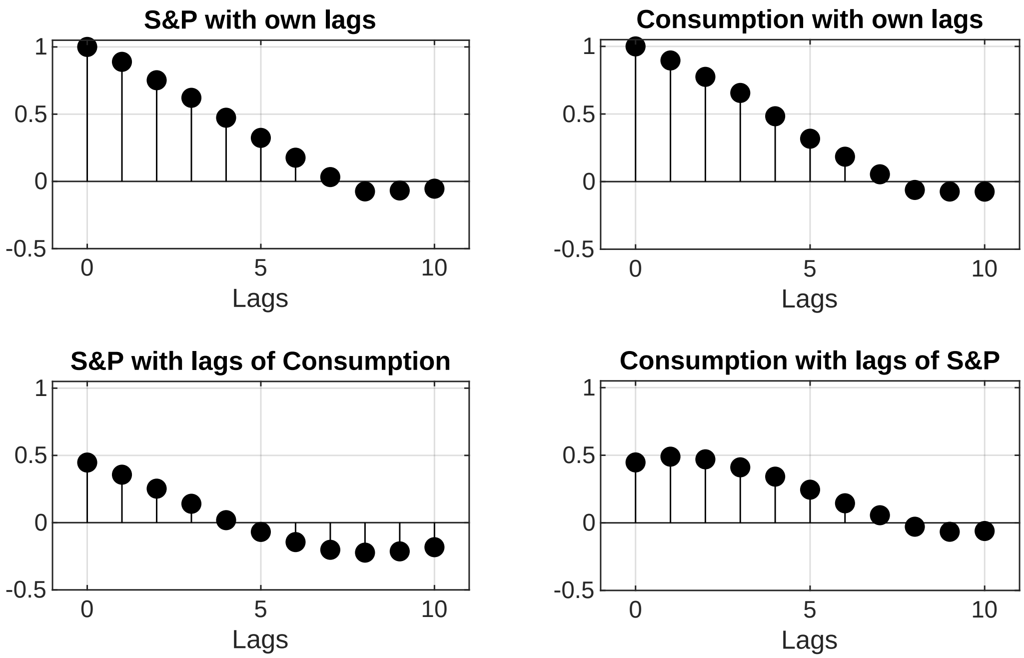

Upper left: Autocorrelations of Hamilton-filtered end-of-quarter value for log S&P composite. Upper right: Autocorrelations of Hamilton-filtered log real consumption. Lower panels: Cross-correlations. Regression parameters: $p = 4$ and $h = 8$.

Figure 4. HP- and Hamilton-filtered stock prices and real consumption.

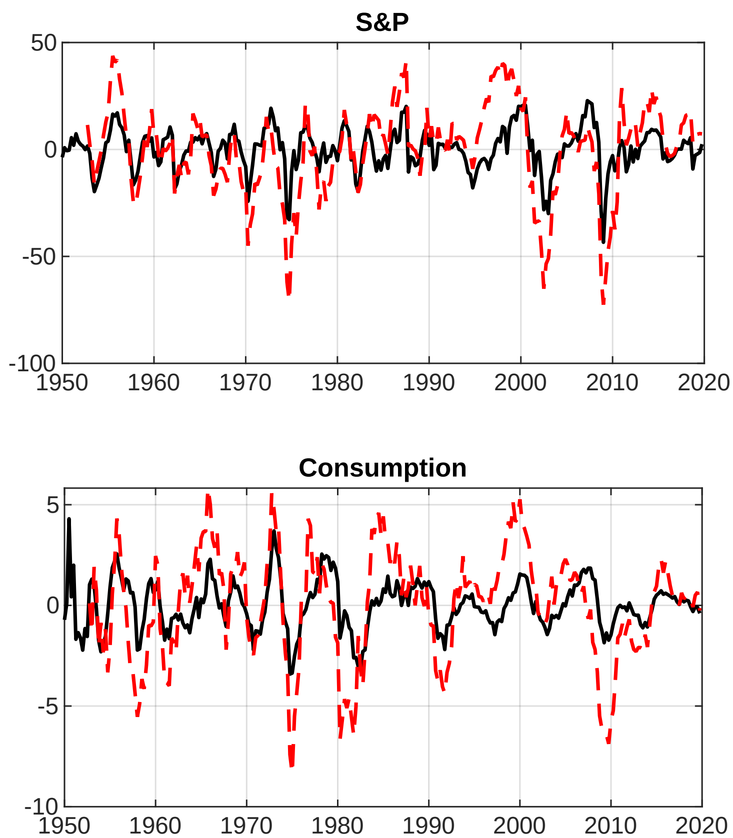

Solid black lines: HP-filtered series ($\lambda = 1,600$). Dashed red lines: Hamilton-filtered series ($p = 4$, $h = 8$).

FIGURE 5. HP and Hamilton trends for stock prices.

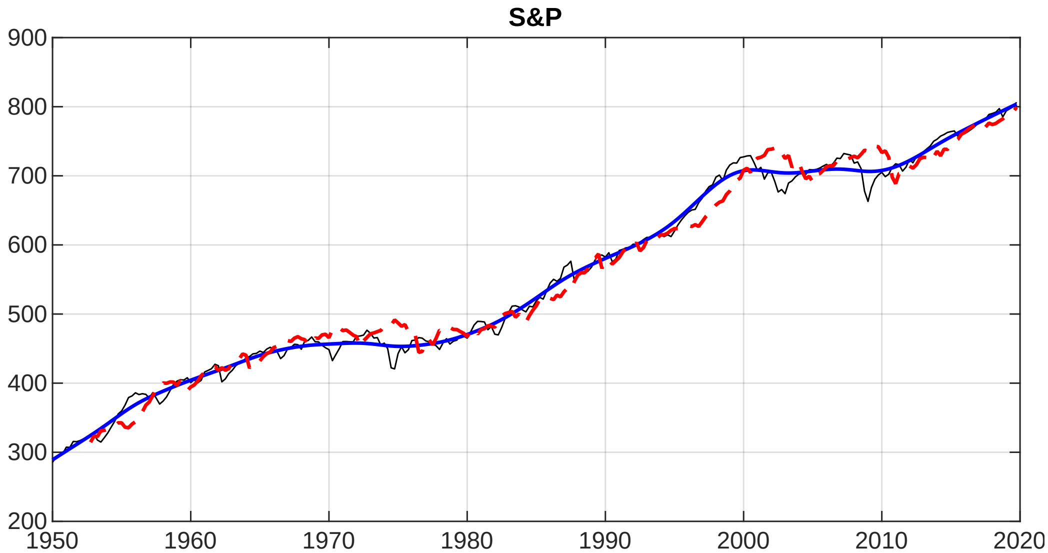

Thin black line: Data. Thick blue line: HP trend ($\lambda = 1,600$). Dashed red line: Hamilton trend ($p = 4$, $h = 8$).

TABLE 1. Business-cycle statistics.

Panel A - Volatility and persistence

	Standard deviation	Autocorrelation
Stock prices		
First difference	7.20	0.12
HP cycle	9.99	0.76
Hamilton cycle	20.95	0.89
Real consumption		
First difference	0.81	0.09
HP cycle	1.23	0.81
Hamilton cycle	2.73	0.90

Panel B - Contemporaneous correlation

	Stock prices			Real consumption		
	First diff.	HP filter	Hamilton filter	First diff.	HP filter	Hamilton filter
Stock prices						
First difference	1.00					
HP filter	0.34	1.00				
Hamilton filter	0.31	0.71	1.00			
Real consumption						
First difference	0.18	0.29	0.29	1.00		
HP filter	−.15	0.45	0.33	0.24	1.00	
Hamilton filter	−.03	0.36	0.45	0.40	0.66	1.00

Notes. Stock prices: 100 times the natural log of the end-of-quarter value for the S&P composite stock price index. Consumption: 100 times the natural log of real personal consumption expenditures from the U.S. NIPA. Sample: 1950Q1 to 2019Q4. HP cycles computed with $\lambda = 1,600$; Hamilton cycles computed with $p = 4$ and $h = 8$.

BANQUE CENTRALE DU LUXEMBOURG

EUROSYSTÈME

2, boulevard Royal
L-2983 Luxembourg

Tél.: +352 4774-1
Fax: +352 4774 4910

www.bcl.lu • info@bcl.lu