University of Victoria | Economics

Department Discussion Paper DDP2401

ISSN 1914-2838

AN AGGREGATION-CONSISTENT IMPLEMENTATION OF THE HAMILTON FILTER

Marco Cozzi
Department of Economics, University of Victoria
Victoria, B.C., Canada V8W 2Y2

May, 2024

Abstract

I propose a modified implementation of the popular Hamilton filter, to make the cyclical component extracted from an aggregate variable consistent with the aggregation of the cyclical components extracted from its underlying variables. This procedure is helpful in many circumstances, for instance when dealing with a variable that comes from a definition or when the empirical relationship is based on an equilibrium condition of a growth model. The procedure consists of the following steps: 1) build the aggregate variable, 2) run the Hamilton filter regression on the aggregate variable and store the related OLS estimates, 3) use these estimated parameters to predict the trends of all the underlying variables, 4) rescale the constant terms to obtain mean-zero cyclical components that are aggregation-consistent. I consider two applications, exploiting U.S. and Canadian data. The former is based on the GDP expenditure components, while the latter on the GDP of its Provinces and Territories. I find sizable differences between the cyclical components of aggregate GDP computed with and without the adjustment, making it a valuable procedure for both assessing the output gap and validating empirically DSGE models.

Keywords: Business cycles, Filtering, Hamilton filter, Output gap, Trend-cycle decomposition.

JEL Classifications: C22, E30, E32.

Author Contact:
Marco Cozzi, Dept. of Economics, University of Victoria, P.O. Box 1700, STN CSC, Victoria, B.C., Canada V8W 2Y2; E-mail: mcozzi@uvic.ca; Tel: (250) 721-6535

An aggregation-consistent implementation of the Hamilton filter

Marco Cozzi, University of Victoria

May 2024

Abstract

I propose a modified implementation of the popular Hamilton filter, to make the cyclical component extracted from an aggregate variable consistent with the aggregation of the cyclical components extracted from its underlying variables. This procedure is helpful in many circumstances, for instance when dealing with a variable that comes from a definition or when the empirical relationship is based on an equilibrium condition of a growth model. The procedure consists of the following steps: 1) build the aggregate variable, 2) run the Hamilton filter regression on the aggregate variable and store the related OLS estimates, 3) use these estimated parameters to predict the trends of all the underlying variables, 4) rescale the constant terms to obtain mean-zero cyclical components that are aggregation-consistent. I consider two applications, exploiting U.S. and Canadian data. The former is based on the GDP expenditure components, while the latter on the GDP of its Provinces and Territories. I find sizable differences between the cyclical components of aggregate GDP computed with and without the adjustment, making it a valuable procedure for both assessing the output gap and validating empirically DSGE models.

JEL Classification Codes: C22, E30, E32.

Keywords: Business cycles, Filtering, Hamilton filter, Output gap, Trend-cycle decomposition.

Contact details: *Department of Economics, University of Victoria, 3800 Finnerty Road, Victoria, BC V8P 5C2, Canada. Tel: 1-250-721-6535, E-mail:* `mcozzi@uvic.ca`

1 Introduction

Decomposing a time series into trend and cyclical components is a ubiquitous task in applied Macroeconomics. Since the '80s, the HP filter of Hodrick and Prescott (1997) has been the most commonly used filtering tool.[1] Hamilton (2018) proposed an alternative regression filter (H, hereafter) that has quickly gained popularity.[2] The H filter derives the cycle (z_t^x) as the deviation of the actual value of a series (x_t) from its estimated trend ($z_t^x = x_t - g_t^x$), which is obtained by regressing the h-period-ahead value of the series (x_{t+h}) on a constant and p lags of the series ($\{x_{t-s}\}_{s=0}^{p-1}$). For quarterly data, Hamilton suggests setting $h = 8$ and $p = 4$, and the regression model is

$$x_{t+8} = \beta_0^x + \beta_1^x x_t + \beta_2^x x_{t-1} + \beta_3^x x_{t-2} + \beta_4^x x_{t-3} + \varepsilon_{t+8}^x$$

where $\left\{\beta_j^x\right\}_{j=0}^{p}$ are parameters to be estimated, and ε_t^x is the error term.

Schüler (2018) argues that the H and HP filters share some of the same pitfalls. Jönsson (2020) and Biolsi (2023) find that, in terms of real-time revision properties, the H filter outperforms both the HP and Beveridge-Nelson filters.[3] Quast and Wolters (2022) modify the H filter, to obtain more reliable real-time output gap estimates.

In this note, I focus on a less studied characteristic of these filters, which I will refer to as their aggregation property. Namely, whenever dealing with a variable consisting of several components (say, GDP), the related cycle is robust irrespective of whether it is obtained by applying the filter to the aggregate variable, or by aggregating the individual cycles obtained from filtering the underlying variables. The HP filter satisfies this property, while the H filter does not. This is problematic when assessing the output gap and validating empirically DSGE models, because there are multiple estimates for the same cyclical component. The intuition for the HP filter's aggregation result can be drawn from its closed-form solution of the trend component ($g_t^{x,HP}$) when the sample is infinite, which is $g_t^{x,HP} = \sum_{j=-\infty}^{\infty} w_j^{\infty} x_{t+j}$, where $w_j^{\infty} = 0.8941^j(0.056168\cos(0.11168j) + 0.055833\sin(0.11168j))$.[4] These formulas show that the weights w_j^{∞} assigned to each observation do not depend on the observations themselves. For a given smoothing parameter, different series will assign the same weight to each time period. Therefore, aggregating the HP cycles of the various components leads to the same HP cycle derived from applying the filter to the aggregate variable. Without loss of generality, consider the case of GDP (y_t), using its most basic definition as the sum of consumption and investment ($y_t = c_t + i_t$). It follows that the aggregation property holds because

$$g_t^{y,HP} = \sum_{j=-\infty}^{\infty} w_j^{\infty} y_{t+j} = \sum_{j=-\infty}^{\infty} w_j^{\infty} c_{t+j} + \sum_{j=-\infty}^{\infty} w_j^{\infty} i_{t+j} = g_t^{c,HP} + g_t^{i,HP}$$

Differently, the weights used by the H filter are estimated OLS coefficients ($\left\{\widehat{\beta_j^x}\right\}_{j=0}^{p}$), which depend on the

[1] Harvey and Jaeger (1993), Cogley and Nason (1995), and Pedersen (2001) analyse the statistical issues of the HP filter.

[2] Contributions in virtually every field of economics and finance now use the H filter, which has already garnered 1743 citations. Recent influential papers using it are Caramp and Singh (2023), Bianchi, Ludvigson, and Ma (2022), and Nguyen (2021).

[3] The Beveridge-Nelson filter is presented in Kamber, Morley, and Wong (2018).

[4] For more details, see Appendix A in Hodrick (2020).

series (x_t) that is filtered. To better understand the essence of the problem, focus on the case with $h = 1$ and $p = 1$, such that

$$\widehat{\beta_1^x} = \frac{\sum_{t=2}^{T}(x_{t-1} - \bar{x})(x_t - \bar{x})}{\sum_{t=2}^{T}(x_{t-1} - \bar{x})^2}$$

where $\bar{x}$ denotes the sample average and T is the sample size.

Since $g_t^{y,H} = \widehat{\beta_0^y} + \widehat{\beta_1^y} y_{t-1}$ and $g_t^{c,H} + g_t^{i,H} = \widehat{\beta_0^c} + \widehat{\beta_1^c} c_{t-1} + \widehat{\beta_0^i} + \widehat{\beta_1^i} i_{t-1}$, the aggregation property $g_t^{y,H} = g_t^{c,H} + g_t^{i,H}$ holds when $\widehat{\beta_0^y} = \widehat{\beta_0^c} + \widehat{\beta_0^i}$ and $\widehat{\beta_1^y} = \widehat{\beta_1^c} = \widehat{\beta_1^i}$.

In the next section, I propose an alternative implementation of the Hamilton filter to make it aggregation-consistent.[5] An advantage of my procedure is that, unlike Hamilton's standard formulation, it leads to a unique assessment of the output gap, irrespective of the income definition. In particular, one could consider the expenditure, product or income approaches to measuring GDP (or the GDP in different sectors or regions), which would aggregate different components, leading to different aggregate cycles. Another key advantage is to provide a coherent setup to estimate DSGE models with moment-based approaches. Whenever in the set of moments used in estimation there are the persistence and/or volatility of GDP and its components, the moments filtered with Hamilton's standard procedure will not be unique, making the estimation somewhat arbitrary. To bypass these inconsistencies, I use the slope parameters estimated on the aggregate variable to predict the trends of all its components.

2 Adjusting the Hamilton Filter

The aggregation-consistent procedure of the adjusted Hamilton filter ($\widetilde{H}$) consists of the following steps.

First, the aggregate variable (y) is obtained from its components.[6]

Second, the Hamilton filter regression is run on the aggregate variable, and the related OLS estimates ($\left\{\widehat{\beta_j^y}\right\}_{j=0}^{p}$) are stored.

Third, the estimated parameters are used to predict the trends of all the variables underlying the aggregate one, by using the slope parameters already obtained ($\left\{\widehat{\beta_j^y}\right\}_{j=1}^{p}$) and by setting the constant terms ($\left\{\widehat{\beta_0^c}, \widehat{\beta_0^i}\right\}$) to an educated guess (e.g., splitting $\widehat{\beta_0^y}$ according to the shares of the GDP components: $\widehat{\beta_0^x} = (\bar{x}/\bar{y})\widehat{\beta_0^y}, x = \{c, i\}$)

$$g_t^{x,\widetilde{H}} = \left(\frac{\bar{x}}{\bar{y}}\right)\widehat{\beta_0^y} + \sum_{j=1}^{p} \widehat{\beta_j^y} x_{t-j+1-h}, x = \{c, i\}$$

Fourth, the cycles (z_t^x) are computed using the usual definition

$$z_t^x = x_t - g_t^{x,\widetilde{H}}, x = \{c, i\}$$

[5] The analysis deals with the cyclical components in levels, as I do not perform the log transformation of the variables. Obtaining the percentage deviations is straightforward. However, in this context log transformations are not informative, as the resulting cyclical components would not be additive by assumption.

[6] For the sake of notation, I focus on a two-variable case, c and i. The procedure is trivially generalized to an arbitrary number of variables.

but they are rescaled by subtracting their averages to make them mean-zero, which is a meaningful step because according to Hamilton's standard procedure they are regression residuals

$$z_t^{x,\widetilde{H}} = z_t^x - \overline{z^x}, x = \{c, i\}$$

3 Two Applications of the Adjusted Hamilton Filter

I now present two applications, exploiting U.S. and Canadian quarterly data.[7]

In the first application, I use U.S. macroeconomic data in the 1947Q1-2023Q1 period. I compute aggregate GDP applying its expenditure approach definition for an open economy. Fig. (1) displays the differences between the H and $\widetilde{H}$ filters, when they are implemented to compute the U.S. business cycles. As for GDP, the assessment of the cyclical component can be off by up to ± 200 billion, which is sizable. As for Consumption and Import, their plots reveal a systematic difference between the two approaches. The cyclical component of Consumption (Import) is overstated by the H filter until the early '90s (early '00s), and understated afterwards. This is largely due to the persistent changes that these variables experienced in their GDP shares.[8]
In terms of differences between the two assessments of the cyclical components, the median percentage gaps are -9.5% for Consumption, -3.1% for Investment, -29.5% for Government Purchases, -3.1% for Export, and -32.5% for Import. These discrepancies are quantitatively important, affecting non-trivially the assessment of the output gap. Fig. (2) displays the cycles for U.S. GDP and its expenditure components, computed with the $\widetilde{H}$ filter. Since the main features of the U.S. macroeconomic aggregates are common to many developed economies (e.g., the values and trends of the various GDP components shares, a high volatility of Investment, the smoothness of Consumption), I expect the discrepancies between the two implementations to display similar behaviors for a number of countries.

In the second application, I use GDP data for the Canadian Provinces and Territories (regions, hereafter) in the 1947Q1-2019Q4 period. Aggregate GDP for Canada is obtained by adding up the GDPs of all eleven regions. Fig. (3) displays the differences in the cyclical component of GDP. As for Canadian GDP, the assessment of cyclical income can be off by up to ± 20 billion, which is consistent with the findings for the U.S.[9] The plots for British Columbia's and Ontario's discrepancies in their filtered GDP show a systematic difference between the two approaches, as the gap for British Columbia (Ontario) is decreasing (increasing) over the sample period.
In terms of differences between the two assessments of the cyclical components, the median percentage gaps are between -1.2% (Quebec) and -44.1% (Northern Territories), and are substantial for most regions. Fig. (4) displays the cycles for the GDP of Canada and its regions, computed with the $\widetilde{H}$ filter. Most notably, larger regions show cycles that are more persistent, unlike smaller ones whose GDP cycles tend to be very volatile.

[7] For the U.S., I use NIPA data. For Canada, I use CANSIM, Statistics Canada, and Conference Board of Canada data.

[8] I uncovered the importance of these shares analyzing some realistic data generating processes with theoretical and Monte Carlo methods.

[9] Recall that total income in the U.S. is approximately ten times larger than Canadian GDP.

4 Conclusion

I proposed an alternative implementation of the Hamilton filter to render it aggregation-consistent, which avoids the pitfall of having multiple assessments of the same cyclical component. I considered two applications, using U.S. and Canadian data. I showed quantitatively important differences between the cyclical component of aggregate GDP, making the adjustment a valuable procedure for assessing the output gap.

To conclude with, the applicability of my analysis goes beyond the Hamilton filter, because any regression-based filter is not aggregation-consistent.

Disclosure Statement: there are no relevant financial or non-financial competing interests to report.

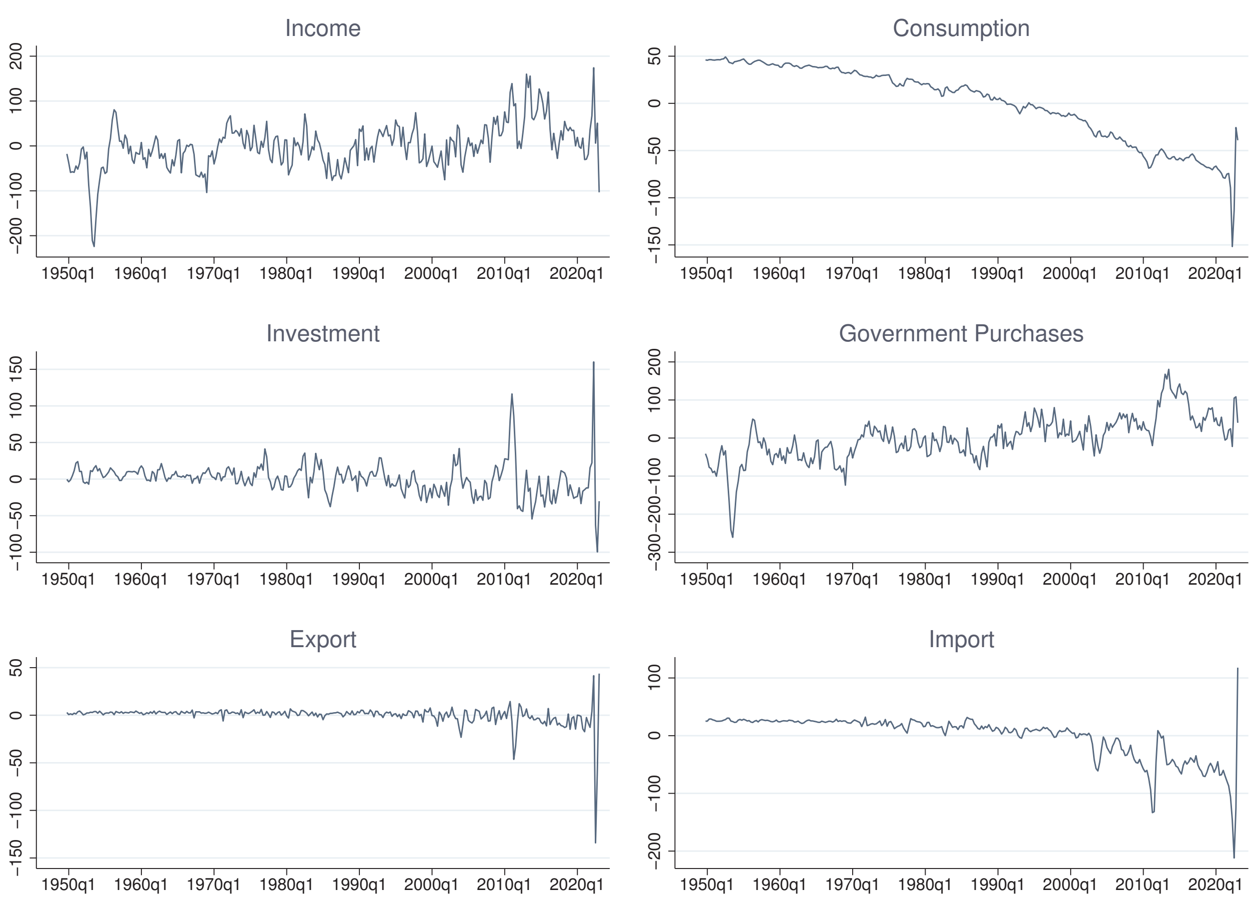

Figure 1: Differences in the cyclical parts of U.S. GDP and its expenditure components ($H - \widetilde{H}$), Billion of USD.

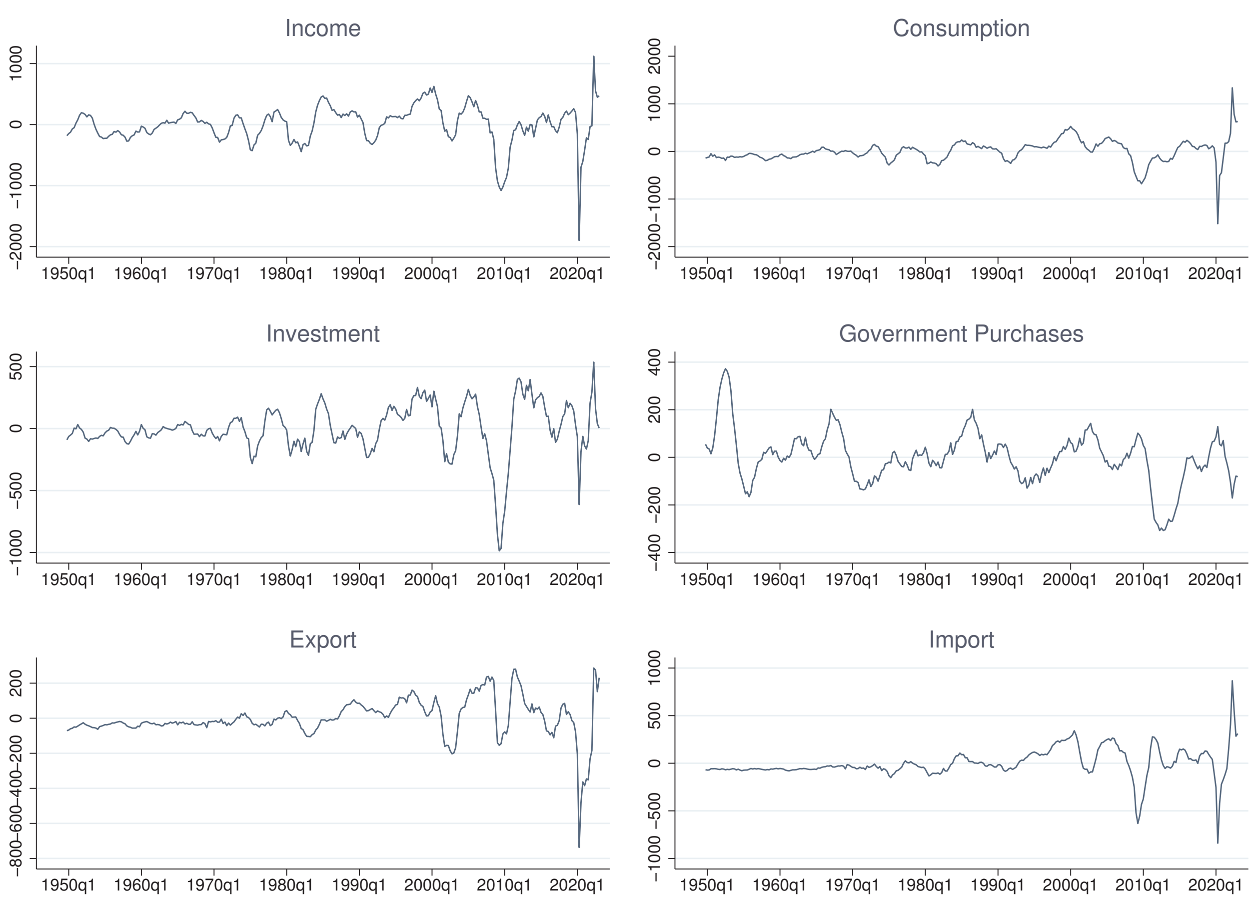

Figure 2: Adjusted (aggregation-consistent) cycles for the U.S. GDP and its expenditure components ($\widetilde{H}$), Billion of USD.

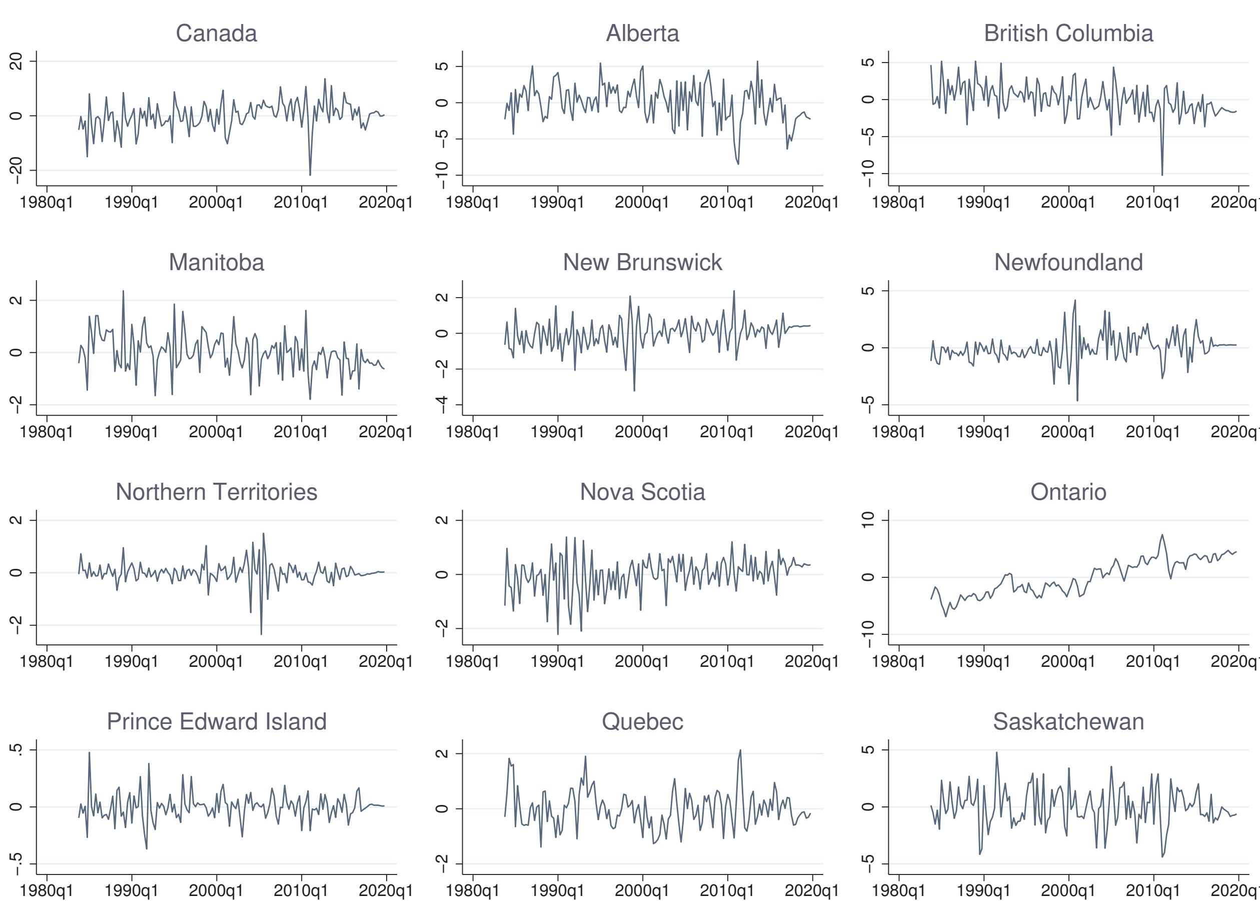

Figure 3: Differences in the cyclical components of GDP for Canada and its Provinces/Territories ($H - \widetilde{H}$), Billion of CAD.

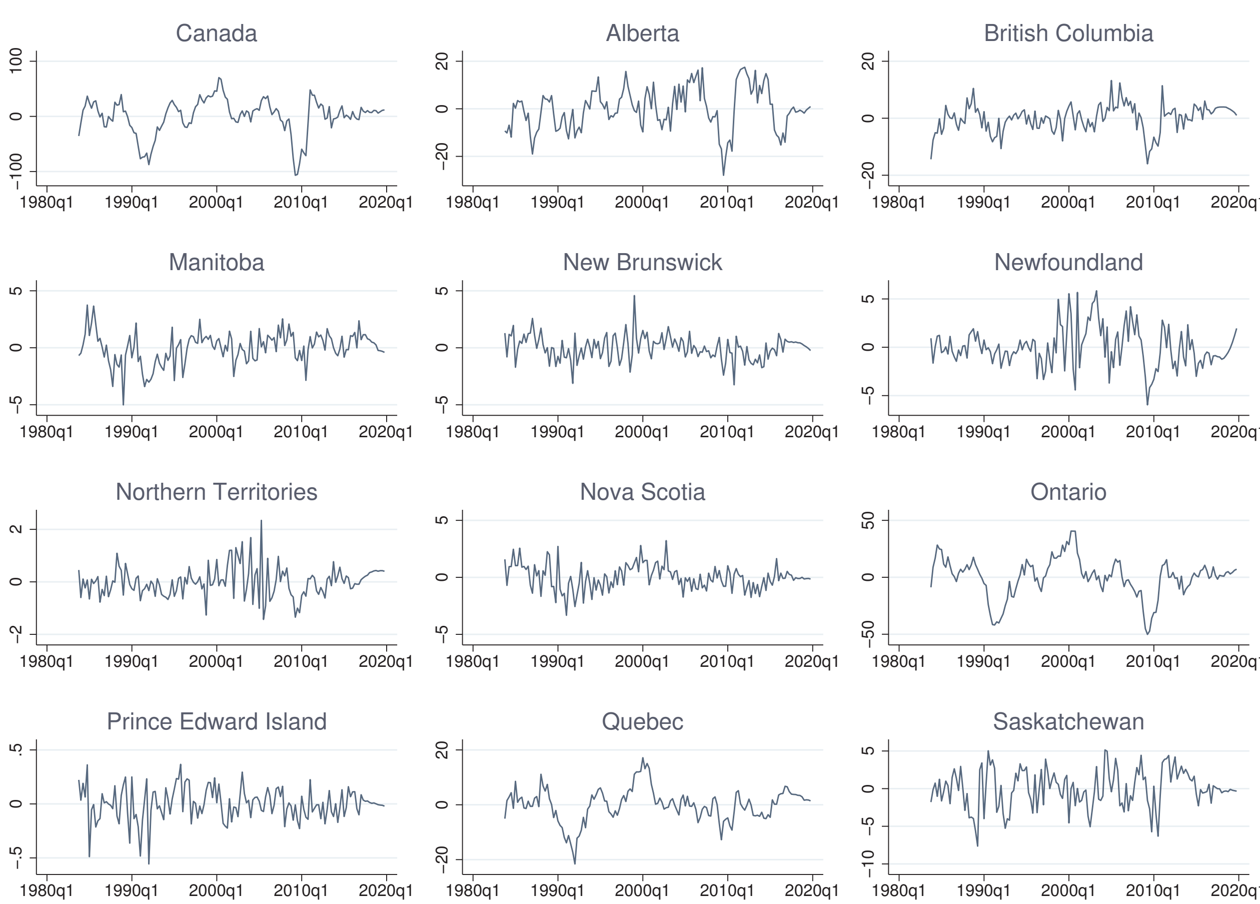

Figure 4: Adjusted (aggregation-consistent) cycles for the GDP of Canada and its Provinces/Territories ($\widetilde{H}$), Billion of CAD.

References

Bianchi, F., Ludvigson, C., and Ma, S. (2022). "Belief Distortions and Macroeconomic Fluctuations," *American Economic Review*, Vol. 112, 2269-2315.

Biolsi, C. (2023). "Do the Hamilton and Beveridge–Nelson filters provide the same information about output gaps? An empirical comparison for practitioners," *Journal of Macroeconomics*, Vol. 75, 103496.

Caramp, N., and Singh, S. (2023). "Bond Premium Cyclicality and Liquidity Traps," *Review of Economic Studies*, Vol. 90, 2822-2879.

Cogley, T., and Nason, J. (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*, Vol. 19, 253-278.

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

Harvey, A., and Jaeger, A. (1993). "Detrending, stylized facts and the business cycle," *Journal of Applied Econometrics*, Vol. 8, 231-247.

Hodrick, R. (2020). "An Exploration of Trend-Cycle Decomposition Methodologies in Simulated Data," *NBER*, Working Paper No. 26750.

Hodrick, R., and Prescott, E. (1997). "Postwar US business cycles: an empirical investigation," *Journal of Money, Credit, and Banking*, Vol. 29, 1-16.

Jönsson, K. (2020). "Real-time US GDP gap properties using Hamilton's regression-based filter," *Empirical Economics*, Vol. 59, 307-314.

Kamber, G., Morley, J., and Wong, B. (2018). "Intuitive and reliable estimates of the output gap from a Beveridge-Nelson filter," *Review of Economics and Statistics*, Vol. 100, 550-566.

Nguyen, T. (2021). "Public Debt, Consumption Growth, and the Slope of the Term Structure," *Review of Financial Studies*, Vol. 35, 3742-3776.

Pedersen, T. (2001). "The Hodrick-Prescott filter, the Slutzky effect, and the distortionary effect of filters," *Journal of Economic Dynamics and Control*, Vol. 25, 1081-1101.

Quast, J., and Wolters, M. (2022). "Reliable Real-Time Output Gap Estimates Based on a Modified Hamilton Filter," *Journal of Business & Economic Statistics*, Vol. 40, 152-168.

Schüler, Y. (2018). "On the cyclical properties of Hamilton's regression filter," *Deutsche Bundesbank*, Discussion Paper No. 03/2018.

Appendix A - Supplementary Analysis

In principle, one would like to know whether my empirical results can be generalised to other data. The Hamilton multivariate regression framework is too complex to allow for clear-cut generalizations. From an analytical perspective, I tried to understand the aggregation properties of the Hamilton filter working with simple Data Generating Processes (DGPs), such as stationary autoregressive processes of order 1, and random walks.

As for the AR(1) case, with arbitrary values for the Hamilton filter parameters h and p, the formulas get too involved. I therefore focused on the simplest possible case, setting $h = 1$ and $p = 1$. In particular, I assumed that the two components (c and i) of the aggregate variable y (say, GDP) are $x_t = \rho_x x_{t-1} + \varepsilon_t^x, x = \{c, i\}, \varepsilon_t^x \sim N(0, \sigma_x^2)$, where $-1 < \rho_x < 1$ denotes the two autocorrelations, and σ_x^2 the two variances of the shocks. Assuming weak exogeneity, using OLS in large samples to estimate the Hamilton regression on the two components separately (asymptotically) recovers the true autocorrelation parameters. It follows that the two trend components are $g_t^{x,H} = \rho_x x_{t-1}, x = \{c, i\}$ and the aggregate trend component is $g_t^{y,H} = \rho_c c_{t-1} + \rho_i i_{t-1}$. Manipulating the last formula (adding and subtracting $\rho_c i_{t-1}$, collecting terms, and multiplying and dividing the second term by y_{t-1}) leads to the following representation of the aggregate trend $g_t^{y,H} = [\rho_c + (\rho_i - \rho_c)\frac{i_{t-1}}{y_{t-1}}]y_{t-1}$. With the postulated DGPs, the aggregate variable is $y_t = c_t + i_t = \rho_c c_{t-1} + \rho_i i_{t-1} + \varepsilon_t^c + \varepsilon_t^i$. Manipulating this formula using the same steps as above, shows that the Hamilton regression $y_t = \beta^y y_{t-1} + e_t^y$ deals with an error term $e_t^y = (\rho_i - \rho_c) i_{t-1} + \varepsilon_t^c + \varepsilon_t^i$ (because, from the perspective of the aggregate variable, the Hamilton regression is misspecified as it is omitting the regressor i_{t-1}). Using standard steps to derive the consistency of an estimator, leads to the following representation of the aggregate trend $g_t^{y,H} = [\rho_c + (\rho_i - \rho_c)\frac{E[i_{t-1}y_{t-1}]}{E[i_t y_t]+E[c_t y_t]}]y_{t-1}$. Comparing the two expressions for the aggregate trend, one concludes that they are equal only when $\frac{i_{t-1}}{y_{t-1}} = \frac{E[i_{t-1}y_{t-1}]}{E[i_t y_t]+E[c_t y_t]}$. The main take-away of this analysis is that: 1) if the components c_t and i_t have exactly the same autocorrelation $\rho = \rho_c = \rho_i$, then there is no issue and the Hamilton filter aggregates; 2) If $\rho_c \neq \rho_i$, then the two following elements play a role in creating a wedge between the two trends: a) the share of the components on the aggregate variable $(\frac{x_{t-1}}{y_{t-1}})$, and b) the covariance between the components and the aggregate variable as a share of the sum of the covariances $(\frac{E[x_{t-1}y_{t-1}]}{E[i_t y_t]+E[c_t y_t]})$.

The previous analysis is quite informative, but deals with a scenario that is not realistic when working with macroeconomic data, which are non-stationary. The issue is that, for most macroeconomic applications, a credible DGP should display a unit root or near unit root behavior. I therefore analysed the issue of filter aggregation with Monte Carlo methods, simulating DGPs that feature some typical characteristics of macroeconomic data. I considered a number of DGPs, assuming that they combine two components: 1) a stationary autoregressive process of order 1, with high persistence (I considered values ranging from 0.90 to 0.99, so that unit root tests have a hard time distinguishing the DGP from a random walk), 2) a deterministic (exponential) growth term, parametrized to deliver an annual growth rate of 2%. I assumed mean-zero normally distributed shocks, selecting a value for the standard deviation that replicates the empirical values of the standard deviations for the (filtered) U.S. GDP components.

The Monte Carlo analysis confirms the finding above that the shares of the variables underlying the aggregate one play an important role in driving the discrepancy between the H filter and its adjusted version $\widetilde{H}$. When these ratios are flat, or display relatively small oscillations, the difference between the two implementations of the Hamilton filter is minimal. In the U.S. data, the consequences of the relatively stable shares of Investment and Export in total GDP can be clearly seen in the first part of the sample in Fig.1. In the related plots, there are small discrepancies between using H or $\widetilde{H}$. Differently, it is well-known that the shares of Consumption and Government purchases underwent persistent changes, which are reflected in the systematic trends in the difference between using H and $\widetilde{H}$.

As for the Random Walk case, I assumed that the two components (c and i) are generated by $x_t = x_{t-1} + \varepsilon_t^x, x = \{c, i\}, \varepsilon_t^x \sim N(0, \sigma_x^2)$. In large samples, it is possible to apply the same analysis as in Hamilton's paper, which implies that the Hamilton filter collapses to a difference filter, as $\widehat{\beta_1^x} = 1, x = \{c, i\}$. Since this result applies also when estimating the Hamilton regression on the aggregate variable, the Hamilton filter aggregates. In small samples, just like in the classical Dickey-Fuller analysis, there are no theoretical results based on closed-form formulas and standard distributions of test statistics. Without performing a thorough Monte Carlo analysis, which is beyond the scope of this paper, it is difficult to make general statements. Finally, in the Hamilton regressions considered in the empirical analysis of the paper (with $h = 8$ and $p = 4$), a statistical test on $\widehat{\beta_1^x} = 1$ does not reject the null at common levels of statistical significance only in a couple of cases (for instance, with U.S. GDP). The data I work with do not seem to support the theoretical property that all (but the first) estimated slope coefficients $\widehat{\beta_j^x}$ should not be statistically different from zero, while $\widehat{\beta_1^x}$ should be equal to one. However, this could be due to the relatively small sample sizes of the two datasets I rely on.