Minerva Research Institute

Empirical Analysis

DECONSTRUCTING DUTCH GDP – A COMPARISON OF METHODOLOGIES

28th November 2023

Authors:

Max Nuboer

Paul Haimerl

Supervisor:

Lenard Lieb

Abstract

Trend-cycle decomposition of GDP is subject to a rich set of different approaches. We demonstrate this heterogeneity in the literature by applying a selected list of commonly used techniques to Dutch real GDP.

Keywords: GDP, trend-cycle, business cycles

Table of Contents

1. Introduction

It is well appreciated that economic development does not evolve in a steady fashion but is subject to recurring boom-and-bust cycles. An accurate understanding of this business cycle is crucial. The private sector economy frequently relies on economic forecasts for investment decisions. Fiscal as well as monetary policy must adjust in accordance with the current economic regime to either cut recessions short or to leverage the benefits of an economic upturn. The visualisation of economic fluctuations is an important tool for public communication. Understanding the nature and potential driving factors of past business cycles holds many learnings for future policy measures. It is not hard to imagine countless similar motivations.

Consequently, ever since the seminal work by (Burns & Mitchell, 1946), a large literature strand has emerged that using quantitative methods to identify such cycles in the real gross domestic product (GDP). These techniques decompose a single time series, like GDP, into a short-run cycle and a long-term trend component. Therefore, they are often referred to as trend-cycle decompositions.

However, despite the great number of proposed methodologies in recent years and an associated, seemingly never-ending discussion, the literature is yet to converge on an approach of disentangling the short-run cyclical component from the long-run trend of macroeconomic time series (see (Kim & Kim, 2020) and (Hodrick R. , 2020) for recent discussions). As a result, applied researchers are exposed to a wide set of different trend-cycle decomposition techniques, that however, somewhat concerningly, are in part based on contradicting assumptions and produce vastly diverging inference (González-Astudillo & Roberts, 2021).

The heterogeneity in the literature also extends beyond the quantitative techniques to the qualitative space. When examining many contributions, it is striking that the nature of the trend and cycle are defined through the respective quantitative approach, instead of tailoring a technique to capture dynamics that correspond to some initial qualitative consideration. In short, the trend in GDP describes the equilibrium state of an economy. The cycle therefore displays deviations from this equilibrium, giving rise to the boom-and-bust notion (Burns & Mitchell, 1946). However, the concrete distinction on which kind of movements should be attributed to the trend or the cycle is not conclusive in the discipline. Some schools of thought favour a trend that is as smooth as possible with the cycle exhibiting large peaks and throughs, whereas others see the cycle as a residual capturing only minor short-term dynamics (see (Lucas, 1980); (Rebelo, 2005)).

This paper estimates the Dutch business cycle and in turn highlights the methodological heterogeneity by applying a selected list of quantitative trend-cycle decomposition approaches to a time series of Dutch real GDP. In particular, we employ the Beveridge-Nelson (BN) decomposition (Beveridge & Nelson, 1981), a boosted Hodrick-Prescott (HP) filter (Phillips & Shi, 2020) as well as three versions of an unobserved components (UC) model: (i) without trend-cycle correlation (Clark, 1987); (ii) allowing for a nonzero trend-cycle correlation (Morley, Nelson, & Zivot, 2003); (iii) specifying a fractionally integrated trend component. In composite, these methods cover most applied research and further include interesting relationships among themselves. Moreover, they each pertain to slightly different qualitative interpretations of the business cycle and the associated trend component. The previous applied literature almost exclusively focuses on the U.S. economy. We validate the common findings of contradicting results when analysing the Dutch economy.

The plan of this paper is as follows: Section 2 briefly introduces the methodology. Section 3 presents the data and the empirical results. Section 4 offers a discussion and concludes.[1]

2. Methodology

2.1. Preliminaries

Let $\{y_t\}_{t=1}^{T}$ denote the scalar time series of logarithmic Dutch real GDP. A logarithmic transformation accounts for exponential growth and is standard in the literature (Perron & Wada, 2009). Further assume that y_t admits to the decomposition into a long-run trend component τ_t and a short-run cyclical component c_t

$$y_t = \tau_t + c_t, \quad t = 1, \dots, T. \tag{1}$$

The trend τ_t, in stark contrast to the cycle c_t, captures the low frequency, slow moving patterns as an economy evolves over time. In consequence, the trend is specified as a smooth and persistent series. This implies that the innovations to the trend do not decay and are still observable even long after they occurred. Examples for such innovations are crucial technological changes and other long-lasting fundamental developments. The cyclical component on the other hand assumes non-persistent transitive dynamics. Such dependencies emerge when shocks affect the economy in the short-run only and quickly fade out in the subsequent periods. This gives a stationary, weakly dependent series c_t that always reverts to a long run mean of zero. In a nutshell, the cycle is interpreted as the boom-and-bust behaviour of an economy that occurs around a long-term GDP trend. Examples for cyclical shocks are energy price changes, wage rigidities or unanticipated policy measures.

The following subsections offer a brief description of some commonly employed quantitative techniques to estimate the latent τ_t and c_t based on the observed y_t. Those methods all separate the trend from the

---

[1] All figures, tables and estimates can be conveniently replicated by the executing R-Notebook available at github.com/Paul-Haimerl/PES-Dutch-Business-Cycle. Furthermore we would like to thank the Data Science Research Infrastructure (DSRI) hosted at Maastricht University for their computational budget.

cyclical component by exploiting the different levels of persistence. However, the particular way in which this structural differentiation is imposed as well as some underlying assumptions differentiate the various approaches.

2.2. Beveridge Nelson Decomposition

Recall that innovations to the trend component do not decay over time. In this fashion, the BN decomposition expresses the trend as

$$\tau_t = \tau_{t-1} + \mu + \epsilon_t, \qquad \epsilon_t \sim iid\ N(0, \sigma_\epsilon^2). \tag{2}$$

As a consequence, the trend at any given time period is simply the sum of a linear deterministic time trend and an unweighted sum of all past innovations $\{\epsilon_t\}_{t=1}^{T}$, a process commonly known as Random Walk with drift (a property also termed $I(1)$, since $\{\tau_t\}_{t=1}^{T}$ has to be difference once in order to obtain a stationary, i.e. mean-reverting, time series). A crucial characteristic of such a process is that, conditional on τ_t and controlling for the deterministic drift μ, future observations are only a function of the stochastic residual ϵ_{t+k} with $k = 1,2, \ldots$ and are therefore not predictable. In contrast, the stationary cyclical component always returns to its long-term average of zero after a shock occurred (i.e. $I(0)$) (Beveridge & Nelson, 1981); (Morley, 2010).

The basic idea of the BN decomposition is thus to separate the predictable (cycle) from the unpredictable (trend) part of a time series. The cycle is thus constituted by the sum of all forecastable changes in the series of GDP, beyond the drift μ. The trend, on the other hand, is everything that is left when the mean-reverting cyclical impulses have dissipated, i.e. the k-step forecast $E(y_{t+k}|y_t, y_{t-1}, \ldots)$ as $k \to \infty$. Given these two insights, it is also straightforward to show that the residuals of the trend and those of the cycle are perfectly negatively correlated. As a result, a unit innovation enters the trend with a positive and the cycle with a negative sign or vice-versa (see (Beveridge & Nelson, 1981) sec. 2, for further details).

In practice, a researcher must select an appropriately large, yet finite value for k to obtain a computationally feasible solution. Furthermore, in order to produce forecasts in the first place, a parametric ARIMA model with adequately selected autoregressive as well as moving average lags needs to be specified, which describes the observed series y_t.

The BN-decomposition typically results in a very shallow cyclical pattern with the trend capturing most of the variation in the GDP series. Subsequently, this methodology may pertain more to a qualitative interpretation that attributes only minor importance to the cycle in favour of the trend component.

2.3. Hodrick-Prescott Filter

In contrast to the BN´s limiting forecast technique, the HP-filter, as noted by the authors, traces out the trend component of a series only by imposing smoothness (Hodrick & Prescott, 1997). Essentially, the HP-technique maximizes the part of the observed series that is captured by the trend captures whilst also preventing the estimated trend from being overly erratic. This results in a reasonably smooth yet fitting trend component. However, the implications of how this is achieved are a frequently cited weakness of the technique (see (Hamilton, 2018) and (Phillips & Jin, 2021)).

Starting again from equation (1), but now let the trend estimate be the series τ_t that minimizes the following criterion

$$\arg\min_{\tau} \left\{ \sum_{t=1}^{T} (y_t - \tau_t)^2 + \lambda \sum_{t=1}^{T} [(g_t - g_{t-1}) - (g_{t-1} - g_{t-2})]^2 \right\}. \tag{3}$$

The cycle follows simply as the residual $c_t = y_t - \hat{\tau}_t$.[2] The first term in (3) evaluates the fitness of the trend. In other words, how much of the variation in the original data does the trend assume. The second part penalizes changes in the growth rate of the trend. As a consequence, the final estimated trend balances the fit to the actual GDP series with how flexible it is. The parameter λ governs how much weight is placed on the penalty term. For $\lambda = 0$ the estimated trend would simply equal the observed series and for $\lambda = \infty$ the trend equals a linear function (Hodrick R. , 2020).

Choosing an appropriate value for λ is not straightforward. Hodrick and Prescott (1997) recommend $\lambda = 1600$ for quarterly macroeconomic data, based on prior economic considerations of how much percentage points the cycle should deviate from the trend. Cross-validation is another technique to tune the penalty parameter. Nonetheless, the vast majority of applied work employs $\lambda = 1600$ (Hamilton, 2018).

The HP-filter is not without its fair share of criticism. First and foremost, it is straightforward to show that at its core it requires y_t to be of $I(2)$ (has to be differenced twice to produce a stationary series). However, this property is frequently not satisfied in practice. A misspecification can introduce complex spurious dynamics to the estimates. Furthermore, the trend and cycle estimates are sensitive to the choice of λ. Setting the penalty parameter to low produces a choppy trend that precludes a sensible analysis of the long-term patterns. On the other hand, specifying the parameter to low leaves trend residuals that pollute the cycle. Moreover, as is the case with any two-sided filter, the HP-filter effectively averages future and past values. This leads to strong revisions of old estimates as new data becomes available and complicates the comparison of time periods in the middle with periods towards the beginning and end of the

[2] Note that $\hat{x}$ denotes an estimate of x.

observational horizon. The weights placed on adjacent observations shift when estimating very early or late observations.

However, in a recent contribution (Phillips & Jin, 2021) show that by boosting, applying the HP-filter to the estimated cycle over and over, many problems of the original approach, in particular spurious cycles, are remedied. The authors argue that a singular application of the HP-filter simply does not suffice do trace out all long-term dynamics. This is a reason for why the HP-filter typically produces large cycle variation, as compared to other techniques. Subsequently, running the filter sequentially repeatedly on the estimated cycle of the previous iteration solves this shortcoming. In the interest of brevity we refer to (Phillips & Shi, 2020) and (Mei, Phillips, & Shi, 2022) for further details. Due to the dramatically improved performance, we will consider this boosted HP-filter in the following.

2.4. Unobserved Components Models

UC models provide the most general framework within the commonly used trend-cycle decomposition methods. In contrast to the approaches presented above it requires a specific functional form for both the trend as well as the cyclical component. Given a certain specification, the framework of UC models is also general enough to nest the BN decomposition and the HP-filter (see (Morley, Nelson, & Zivot, 2003) and (Grant & Chan, 2017) for the transformations).

An explicit parametric form for both components allows the system to be cast in state space form and then solved efficiently via the Kalman filter recursions. Parameter estimation is performed through (Quasi-) Maximum Likelihood or the computationally less expensive Conditional Sum of Squares (CSS) optimization (Harvey, 1990); (Hartl, 2022).

We consider three popular variants of UC models. The first, termed *UC-0*, posits a trend function as in (2) and sets up the cyclical component as a second order autoregressive process, as is the standard in the literature

$$c_t = \phi_1 c_{t-1} + \phi_2 c_{t-2} + \eta_t, \qquad \eta_t \sim iid\left(0, \sigma_\eta^2\right). \tag{4}$$

In order to ensure stationarity of the cycle, all roots of the lag polynomial in (4) must lie outside the unit circle. Moreover, the core characteristic of the *UC-0* model is that cycle and trend components are specified to develop in isolation of each other, i.e. an innovation to the cycle does not correlate with the trend in any way. This gives $\sigma_{\epsilon\eta} = 0$ and is in stark contrast to the BN decomposition, where $\sigma_{\epsilon\eta} = -1$ by construction.

However, it is not hard to challenge this zero correlation assumption for many practical applications. For example, important technological advances may hold a positive effect on the GDP trend but trigger a downturn in the cycle as workers are substituted by the new technology and not yet retrained. Similarly, a national bank increasing the funds rate can depress an economy in the short-run but pave the way for long-term success. To incorporate such dynamics, given certain requirements that are however met by the model in (1), (2), (3) and (4), it is possible to relax the zero trend-cycle correlation constraint and let $\sigma_{\epsilon\eta}$ to be estimated together with the other system parameters (Morley, Nelson, & Zivot, 2003). This constitutes the second *UC-Corr* model.

Lastly, as previously mentioned, it is fair to say that the concrete value of the integration order d of GDP is somewhat of an open question. A natural solution to this model specification problem is to simply let the integration order of the trend be a parameter to be estimated (Hartl, 2022). This, in combination with a nonzero trend-cycle correlation gives the *UC-Frac* model.[3]

3. Data and Empirical Results

3.1. Dutch Real GDP

Figure 1 sketches the observed series of Dutch real GDP and its logarithmic transformation for the available observational horizon. As common in the literature, we work with data that has been seasonally adjusted. GDP is expressed in chained Euro to allow for a convenient comparison across time periods.

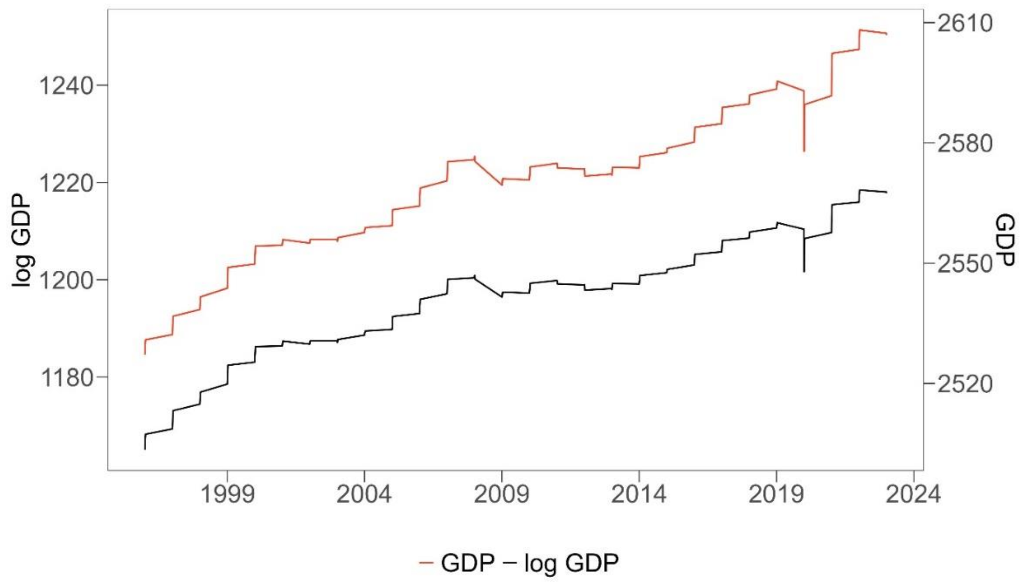

[3] Note that the cyclical component does not change relative to *UC-0* or *UC-Corr*.

Figure 1 Seasonally adjusted quarterly Dutch GDP from 1996:1 until 2023:3 in Bn. of chained 2010 Euros (orange, right scale). Logarithmic series times 100 to allow for an easy interpretation in percentage points (black, left scale).

A brief list and description of the economic regimes present during the observed time periods is provided in the following.

Q1 1996 – Q4 2000	From the first quarter of 1996 to the last of 2000, there is a gradual rise in Dutch GDP. The Netherlands even exceeded the average of the EU during this period (Albers & Langedijk, 2004). Furthermore, the beginning of the European Union coincides with this time span. The gradual economic integration of the Netherlands within the EU correlates with the experienced growth.
Q4 2000 – Q1 2004	Also affected by the Dot-Com Bubble, the Dutch economy became stagnant. Whilst an episode of low growth was also felt in other EU countries, the Netherlands was affected more intensely than most peers (Albers & Langedijk, 2004).
Q1 2004 – Q3 2008	In this time the Dutch economy made a recovery from the previous economic slump. A gradual upturn can be seen throughout the later quarters of 2004 and 2005 until becoming steeper up until the third quarter of 2008, when the global financial crisis perturbed the world economy.
Q3 2008 – Q1 2009	Economic growth rapidly declined when the global financial crisis struck, a recession on a scale not seen since the great depression. Whilst the Dutch economy was able to maintain growth until the third quarter of 2008, there was a subsequent rapid decline. It was thought that the Dutch economy was in a strong position to deal with the oncoming crisis as a result of the surplus in the current account that had been running along with a low figure for the unemployment rate (Masselink & van den Noord, 2009).
Q1 2011 – Q1 2013	Following a measured recovery and increase in growth from the 2008 financial crisis, the Dutch economy once again a period of stagnation. Whilst the fall in growth rates was not as significant as in 2008, it lasted far longer for around 2 years. Especially in other EU members such as Spain or Greece, the effects of the crisis were still ongoing. Household income in the Netherlands was falling and in composite with low consumer confidence, consumption fell drastically.
Q1 2013 – Q4 2019	From the first quarter of 2013 to the final quarter of 2019 the Netherlands saw a period of sustained economic growth. This marked the full recovery from the financial crisis. Consumer confidence began to rise again and aggregate spending in the economy increased along with this.
Q4 2019 – Q2 2022	In the beginning of 2020, the Dutch economy experienced the greatest ever one-quarter fall in economic growth. This recessionary period was the result of

	the COVID-19 pandemic. The pandemic heavily restricted the international flow of people, goods and services alike, decimating world trade, an economic sector of particular importance to the Netherlands.
Q2 2020 – Q1 2023	In the past three years, the Dutch economy was able to make a recovery following the COVID-19 pandemic. However, towards the end of the observational period the data suggests another reduction of growth rates. Several key interest rate hikes by central banks around the world have introduced an economic regime not present since before the financial crisis, the effects of which are still to be determined.

3.2. Estimating the Trend and Cycle

The following subsection implements the various estimation strategies laid out above. The BN decomposition is solved as to imply a cyclical with two autoregressive lags, in accordance with the UC models. Furthermore, we specify an appropriately long forecasting horizon in order to capture the permanent component. The boosted HP-filter is executed as specified in (Phillips & Shi, 2020) with $\lambda = 1600$ and includes eight boosting iterations.[4] As common in the literature, the UC models are estimated with a numerical optimization routine. For each model variant we construct a large grid of starting values with appropriate support and pick the combination of initial values yielding the best overall fit. The optimization run based on these initial values serves as the final estimates. The system parameters for the considered UC models are displayed in Table 1.

	UC–0	UC-Corr	UC-Frac
d	1 -	1 -	0.474 (0.026)
σ_ϵ	1 -	1 -	1 -
σ_η	0.224 (0.001)	0.224 (0.001)	0.4 (0.001)
ρ	0 -	-0.918 (0.008)	-0.999 (0.005)
ϕ_1	0.607 (0.015)	-1.787 (0.015)	-1.761 (0.035)
ϕ_2	0.555 (0.017)	0.859 (0.017)	0.812 (0.036)
μ	0.006	0.005	5.197
CSS	-176.003	-170.353	-154.588

[4] As laid out in (Phillips & Shi, 2020), the boosted HP-filter is largely invariant to changes in the penalty parameter λ. As a consequence we resort to the standard value in the interest of simplicity.

Table 1 CSS estimates for the three UC models. Standard errors are in parenthesis. A lack of standard errors denotes a non-estimated quantity. Note that with the CSS-optimization, the variance-covariance-matrix of system innovations is only defined up to a nonzero multiplicative constant. As a consequence, we normalize σ_ϵ *to 1. A computationally more expensive quasi-maximum likelihood estimator however could uncover the level of* σ_ϵ.

Briefly turning to the estimation results of the UC models, it is striking that the *UC-Corr* as well as the *UC-Frac* variants tend towards a corner solution of -1 for the trend-cycle correlation. This replicates previous contributions such as (Morley, Nelson, & Zivot, 2003). Furthermore, the almost perfect negative correlation has a strong knock-on effect on the cycle coefficients, as demonstrated by the difference between *UC-0* and *UC-Corr* with regard to $\widehat{\phi_1}$ and $\widehat{\phi_2}$. However, most interestingly, when controlling for a drift, the *UC-Frac* model specifies a GDP trend that exhibits long-memory characteristics while still being stationary. Previous empirical studies of macroeconomic indicators, in particular US GDP, usually find a level of persistence in between 1 and 1.5 (Hartl, Tschernig, & Weber, 2020). This is in stark contrast to our estimates for the Dutch GDP series. Moreover, this indicates that the remaining techniques presented here severely over-difference the permanent component and thus eliminate important information. This dynamic also explains the large deviations in the drift parameter μ among the various specifications in Table 1.[5]

Figure 2 presents the respective estimates of the cycle and Figure 3 of the trend component. It is apparent that, albeit for the *Frac-UC* model, all techniques are roughly in accordance. Almost all variation of GDP is attributed to the trend components and the cycles appear to move laterally around zero with little variability. Merely the time periods subject to the COVID-19 pandemic result in a large negative spike of the cycle. In contrast, the *Frac-UC* cycle seems to be the only technique that replicates the business cycle chronology laid out in the previous sub-section. The trend as estimated by the *Frac-UC* model is far smoother and subsequently more movement is allocated to the cycle.

When recalling that $\hat{d} \approx 0.5$ this may seem a bit puzzling at first. Lower integration orders are typically associated with higher, variability. At the other side of the spectrum, over-differencing eliminates more dynamics than necessary to obtain a stationary series and thus produces an overly smooth trend (Hassler, 2018). Nevertheless, in this application the trend of the *Frac-UC* model is integrated of order ≈ 0.5 whereas the other models involve $d = 1{,}2$ and still features a lower variance. However, also note that the variance parameters in Table 1 are only identified up to a multiplicative constant, a feature of the employed CSS estimator. Therefore, an interpretation of these parameters can be easily misleading. We suspect a different dynamic at play here. Over-differencing the data could make the true trend-dynamics simply unidentifiable. Subsequently, cleanly disentangling the trend from the cycle becomes impossible and the cycle assumes part of the innovations actually pertaining to the trend and vice-versa.

---

[5] Given a Maximum Likelihood optimization, a Likelihood-ratio test that performs model selection would have been interesting to see in this context.

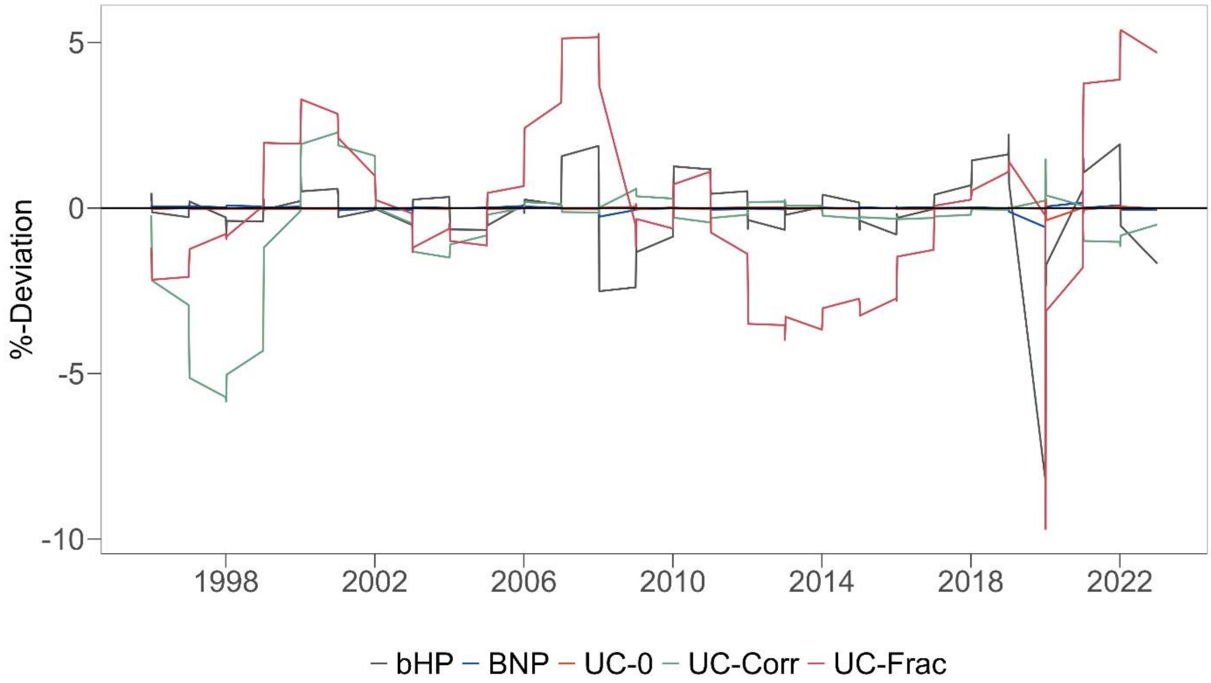

Figure 4 Cycle estimates in percentage deviation from the trend. Boosted HP-filter (grey), BN decomposition (blue), UC-0 (orange, solid), UC-Corr (orange, green), UC-Frac (red).

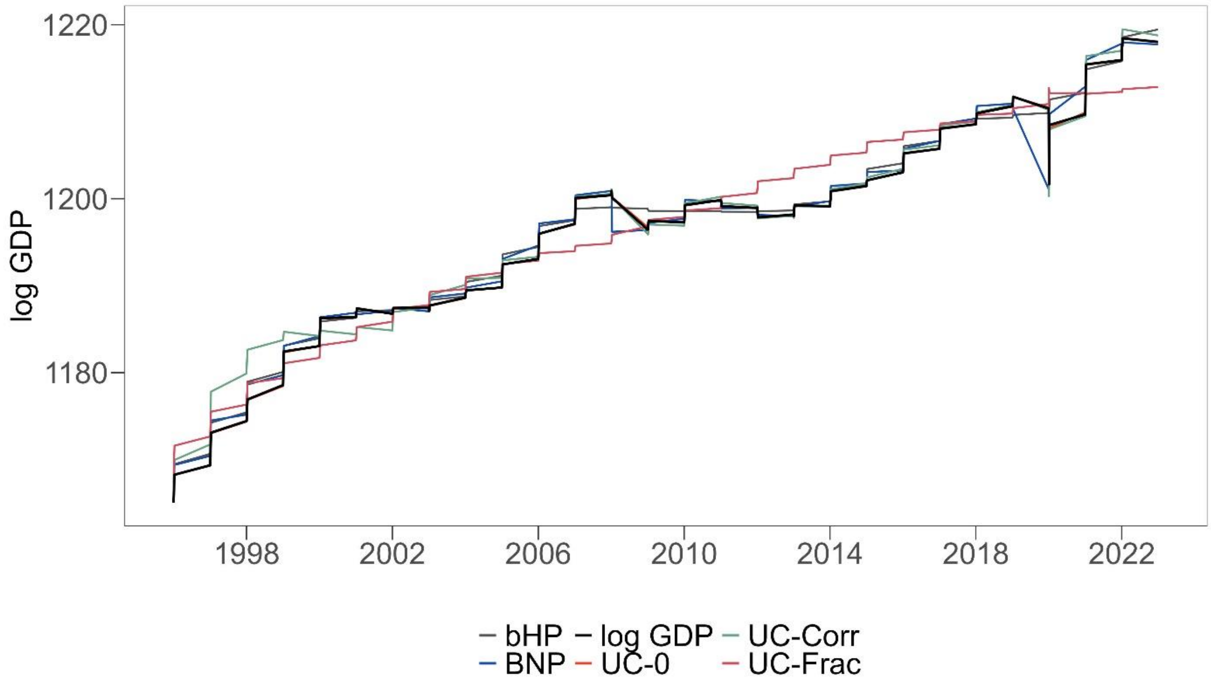

Figure 5 Estimates of the trend in logarithmic GDP. Log GDP (black), boosted HP-filter (grey), BN decomposition (blue), UC-0 (orange), UC-Corr (green), UC-Frac (red).

4. Conclusion

This paper estimates the Dutch Business cycle using a variety of commonly employed techniques. We replicate prior results of different trend-cycle decompositions resulting in diverging estimates. Overall, a UC model with a nonzero trend-cycle correlation and a fractionally integrated trend component seems to fit the data generating process of Dutch real GDP best. Nevertheless, as already pointed out, this evaluation also depends on the qualitative interpretation of the trend-cycle decomposition to some extent. Therefore, it is not straightforward to name any one winner.

Furthermore, it is important to highlight that most similar empirical studies are based on a vastly more extensive observational horizon, spanning 60 years or more. Due to data availability, our application only includes just shy of 30 years worth of quarterly observations. The short time span can limit the robustness of results and lead to unstable estimates, in particular when relying on numerical optimization. There do simply do not exist many observations that describe a relationship over a long time span such as 20 years or more. As a consequence, out inference is in part based on few usable data points which hurts the robustness.

In a similar vein, a shorter observed duration is also limiting from an economic perspective. Economic regimes such as the stagflation period of the 1970s are not reflected in our data. This may particularly impact the persistence parameter of the *UC-Frac* model, as an overall less monotonic trajectory of the data is likely to increase the estimated integration order. It is common for many macroeconomic variables that they appear to follow a linear trend in the short to medium term, as confirmed by the *UC-Frac* model. Only when observing a considerate amount of time periods emerge the true underlying patterns. In consequence, trend and cycle characteristics may differ when extending the observational horizon.

Another interesting extension would be an approach that includes a structural break in the drift component, as proposed in (Perron & Wada, 2009). However, this naturally requires an actually observed structural break, a feature that is arguably not reflected in the data available to us.

References

Albers, R., & Langedijk, S. (2004). *The Netherlands: From Riches to Rags.* European Commission, Diretorate-General for Economic and Financial Affairs.

Beveridge, S., & Nelson, C. R. (1981). A new approach to decomposition of economic time series into permanent and transitory components with particular attention to measurement of the 'business cycle'. *Journal of Monetary economics, 7*, 151–174.

Burns, A. F., & Mitchell, W. C. (1946). *Measuring business cycles.* National bureau of economic research.

Clark, P. K. (1987, November). The Cyclical Component of U. S. Economic Activity. *The Quarterly Journal of Economics, 102*, 797. doi:10.2307/1884282

González-Astudillo, M., & Roberts, J. M. (2021, August). When are trend–cycle decompositions of GDP reliable? *Empirical Economics, 62*, 2417–2460. doi:10.1007/s00181-021-02105-4

Grant, A. L., & Chan, J. C. (2017, February). Reconciling output gaps: Unobserved components model and Hodrick–Prescott filter. *Journal of Economic Dynamics and Control, 75*, 114–121. doi:10.1016/j.jedc.2016.12.004

Hamilton, J. D. (2018, December). Why You Should Never Use the Hodrick-Prescott Filter. *The Review of Economics and Statistics, 100*, 831–843. doi:10.1162/rest_a_00706

Hartl, T. (2022). The fractional unobserved components model: a generalization of trend-cycle decompositions to data of unknown persistence.

Hartl, T., Tschernig, R., & Weber, E. (2020). Fractional trends and cycles in macroeconomic time series. *arXiv preprint arXiv:2005.05266.*

Harvey, A. C. (1990). Forecasting, structural time series models and the Kalman filter.

Hassler, U. (2018). *Time series analysis with long memory in view.* John Wiley & Sons.

Hodrick, R. (2020, February). *An Exploration of Trend-Cycle Decomposition Methodologies in Simulated Data.* National Bureau of Economic Research. doi:10.3386/w26750

Hodrick, R. J., & Prescott, E. C. (1997, February). Postwar U.S. Business Cycles: An Empirical Investigation. *Journal of Money, Credit and Banking, 29*, 1. doi:10.2307/2953682

Kim, C.-J., & Kim, J. (2020, June). TREND-CYCLE DECOMPOSITIONS OF REAL GDP REVISITED: CLASSICAL AND BAYESIAN PERSPECTIVES ON AN UNSOLVED PUZZLE. *Macroeconomic Dynamics, 26*, 394–418. doi:10.1017/s1365100520000218

Masselink, M., & van den Noord, P. (2009). The global financial crisis and its effects on the Netherlands. *ECFIN Country Focus, 6*(10), 1-7.

Mei, Z., Phillips, P. C., & Shi, Z. (2022). The boosted HP filter is more general than you might think. *arXiv preprint arXiv:2209.09810.*

Morley, J. C. (2010, June). THE TWO INTERPRETATIONS OF THE BEVERIDGE–NELSON DECOMPOSITION. *Macroeconomic Dynamics, 15*, 419–439. doi:10.1017/s1365100510000118

Morley, J. C., Nelson, C. R., & Zivot, E. (2003). Why are the Beveridge-Nelson and unobserved-components decompositions of GDP so different? *Review of Economics and Statistics, 85*, 235–243.

Netherlands, Statistics. (2021). *Economic impact of COVID-19.* Retrieved from https://www.cbs.nl/en-gb/dossier/coronavirus-crisis-cbs-figures/economic-impact-of-covid-19#:~:text=In%20Q2%202020%2C%20the%20Dutch,domestic%20product%20(GDP)%20rebounded.

Perron, P., & Wada, T. (2009, September). Let\textquotesingles take a break: Trends and cycles in US real GDP. *Journal of Monetary Economics, 56*, 749–765. doi:10.1016/j.jmoneco.2009.08.001

Phillips, P. C., & Jin, S. (2021, January). BUSINESS CYCLES, TREND ELIMINATION, AND THE HP FILTER. *International Economic Review, 62*, 469–520. doi:10.1111/iere.12494

Phillips, P. C., & Shi, Z. (2020, December). Boosting: Why You Can Use the HP Filter. *International Economic Review, 62*, 521–570. doi:10.1111/iere.12495