Australian National University

Crawford School of Public Policy

CAMA

Centre for Applied Macroeconomic Analysis

Information Sets and Output Gap Estimates from Beveridge-Nelson and Unobserved Components Models

CAMA Working Paper 73/2026
August 2026

Tino Berger
University of Göttingen
Centre for Applied Macroeconomic Analysis, ANU

Benjamin Wong
Monash University
Centre for Applied Macroeconomic Analysis, ANU

Abstract

Morley, Nelson and Zivot (2003) show that the Beveridge-Nelson decomposition and an unrestricted unobserved-components model deliver the same cycle for U.S. real GDP once the covariance between trend and cycle innovations is freely estimated. We revisit this result in a narrow replication and then extend it in two directions. First, we repeat the univariate exercise over a longer sample ending in 2026. Second, we add labor-market and real-activity indicators that should be directly informative about the cycle. The narrow replication reproduces the equivalence between the Beveridge-Nelson and unrestricted unobserved-components decompositions in both samples. In the multivariate extension, the two gaps remain closely aligned, with a correlation of 0.995, and the estimated cycle displays persistent contractions around major U.S. recessions. We interpret these results as evidence that the volatile cycle in the unrestricted univariate model mainly reflects the limited cyclical information contained in GDP alone.

Keywords

output gap, Beveridge-Nelson decomposition, unobserved components, information set, trend-cycle decomposition

JEL Classification

C32, E32

Address for correspondence:

(E) cama.admin@anu.edu.au

ISSN 2206-0332

The Centre for Applied Macroeconomic Analysis in the Crawford School of Public Policy has been established to build strong links between professional macroeconomists. It provides a forum for quality macroeconomic research and discussion of policy issues between academia, government and the private sector.

The Crawford School of Public Policy is the Australian National University's public policy school, serving and influencing Australia, Asia and the Pacific through advanced policy research, graduate and executive education, and policy impact.

Information Sets and Output Gap Estimates from Beveridge–Nelson and Unobserved Components Models*†

Tino Berger[1,3] and Benjamin Wong[2,3]

[1]University of Göttingen
[2]Monash University, Australia
[3]Centre for Applied Macroeconomic Analysis, ANU

August 2026

Abstract

Morley, Nelson and Zivot (2003) show that the Beveridge-Nelson decomposition and an unrestricted unobserved-components model deliver the same cycle for U.S. real GDP once the covariance between trend and cycle innovations is freely estimated. We revisit this result in a narrow replication and then extend it in two directions. First, we repeat the univariate exercise over a longer sample ending in 2026. Second, we add labor-market and real-activity indicators that should be directly informative about the cycle. The narrow replication reproduces the equivalence between the Beveridge-Nelson and unrestricted unobserved-components decompositions in both samples. In the multivariate extension, the two gaps remain closely aligned, with a correlation of 0.995, and the estimated cycle displays persistent contractions around major U.S. recessions. We interpret these results as evidence that the volatile cycle in the unrestricted univariate model mainly reflects the limited cyclical information contained in GDP alone.

Keywords: output gap; Beveridge-Nelson decomposition; unobserved components; information set; trend-cycle decomposition.

1 Introduction

In an influential paper, Morley, Nelson, and Zivot (2003, MNZ henceforth) examine why two standard decompositions of U.S. real GDP give such different pictures of the business cycle. In the conventional unobserved-components (UC) approach, output is decomposed into a smooth stochastic trend and a persistent cyclical component. In the Beveridge-Nelson (BN) approach, the trend is defined as the model-implied long-run conditional forecast of output. The cycle is the deviation of current output from this forecast and is typically much smaller and less persistent.

MNZ show that this difference reflects a restriction that is usually imposed in the UC model, where the standard UC model sets the covariance between trend and cycle innovations to zero. Once this restriction is relaxed, the unrestricted UC decomposition lines up with the BN representation implied by the corresponding ARIMA model. In their empirical application with U.S. real GDP, the unrestricted UC and BN-ARMA

*Contact: Berger: tino.berger@wiwi.uni-goettingen.de Wong: benjamin.wong@monash.edu
†We thank Elmar Mertens and James Morley for helpful discussions and comments.

models have the same likelihood, which leads to identical cycles, and the estimated correlation between trend and cycle innovations in the UC model is strongly negative.

The resulting output gap, however, often lacks persistence and amplitude, and is difficult to reconcile with the more persistent movements one usually associates with the business cycle. Consequently, this has also led to a perception that one *needs* the zero correlation between trend and cycle innovations when specifying a UC model in order to obtain these more persistent and large amplitude cycles.[1]

We first replicate MNZ in a narrow sense, reproducing their results with their sample, and also extending through to 2026. The narrow replication of MNZ also confirms their point that the restriction on the correlation between trend and cycle innovations is the source of difference between the UC and BN decomposition, and that freeing up the restriction leads to an equivalence between the UC and BN output gap. In our wider replication, we ask whether one indeed requires the zero correlation that one imposes between trend and cycle innovations in a UC model, as commonly believed and practised by the UC literature, in order to obtain large and persistent cycles. Our wider replication suggests this is not necessarily true. In particular, we show that the information set that one uses to apply the BN decomposition or UC model matters. When we condition both the UC model and BN decomposition on a standard set of cyclical variables — U-2 unemployment rate, industrial production growth, and hours growth — two results stand out. First, we can obtain both large and persistent output gaps *and* a non-zero correlation between trend and cycle, suggesting that one does not necessarily require zero correlation between trend and cycle in order to obtain the persistence and large cycle as one expects from an output gap. Second, when we also allow for the same information set using a BN decomposition that allows for multivariate information in the spirit of Morley and Wong (2020), the obtained output gap is almost identical to that of the UC model with cyclical indicators. This second result is consistent with MNZ's intuition that there is an equivalence between the BN decomposition and UC model, with our added caveat that this holds as long as one conditions on the same information set in both settings.

The rest of the paper is organized as follows. Section 2 presents the narrow replication. Section 3 presents the wider replication where we consider the role of enlarging the information set and how to think about the equivalence between the Kalman filtered trend from the UC model and the BN decomposition. Section 4 discusses reconciling output gap estimates and the associated information sets. Section 5 offers concluding remarks.

2 Narrow Replication

MNZ start from the standard trend-cycle representation of log real GDP. Let y_t denote 100 times log real GDP. The UC model writes

$$y_t = \tau_t + c_t, \tag{1}$$

where τ_t is the stochastic trend and c_t is the cycle, or output gap. The trend follows a random walk with drift,

$$\tau_t = \mu + \tau_{t-1} + \eta_t, \qquad \eta_t \sim N(0, \sigma_\eta^2), \tag{2}$$

[1]We note that specifying zero correlation between trend and cycle innovations remains a prevalent practice, especially when one considers extensions of the base UC model to the related trend-cycle VAR literature. Recent examples include Del Negro, Giannone, Giannoni, and Tambalotti (2017) and Maffei-Faccioli (2025).

where μ is the drift. The cycle follows a stationary AR(2) process,

$$c_t = \phi_1 c_{t-1} + \phi_2 c_{t-2} + \varepsilon_t, \qquad \varepsilon_t \sim N(0, \sigma_\varepsilon^2). \tag{3}$$

The key restriction in the conventional UC model is $\text{cov}(\eta_t, \varepsilon_t) = 0$. We refer to this restricted model as UC-0. The unrestricted model, UC-UR, estimates this covariance, or equivalently the correlation of trend and cycle innovations, $\rho_{\eta\varepsilon}$.

The BN decomposition defines trend output as the long-horizon conditional forecast of output, adjusted for its drift,

$$\tau_t^{BN}(\Omega_t) \equiv \lim_{h\to\infty} \text{E}(y_{t+h} - h\mu \mid \Omega_t)\,, \tag{4}$$

where h is the forecast horizon, Ω_t is the information set available at time t, and $\text{E}(\cdot \mid \Omega_t)$ denotes the model-implied conditional expectation. The corresponding BN cycle is

$$c_t^{BN}(\Omega_t) \equiv y_t - \tau_t^{BN}(\Omega_t) = -\sum_{j=1}^{\infty} \text{E}(\Delta y_{t+j} - \mu \mid \Omega_t)\,. \tag{5}$$

In the narrow replication, Ω_t contains current and past observations of y_t. Section 3 enlarges this information set to include the additional labor-market and real-activity indicators.

MNZ show that the UC-ARMA model and the BN-ARIMA representation have the same reduced form once the covariance between trend and cycle innovations is treated as an unrestricted parameter. In their empirical application, the BN decomposition is computed from an ARIMA$(2, 1, 2)$ model for y_t, while the corresponding unrestricted UC model is estimated in state-space form. In their estimates for U.S. real GDP, MNZ find the empirical counterpart of this equivalence, with the two models having the same likelihood and delivering the same cycle.

Table 1 summarizes the narrow replication for the original MNZ sample and for the extended sample ending in 2026. Extending the sample leaves the estimated dynamics largely intact. The main visible change is a lower estimated drift of trend output, consistent with lower average growth over the longer postwar sample. The likelihood comparison is central to the narrow replication. In both samples, the BN-ARMA and unrestricted UC specifications attain the same log likelihood. The updated sample therefore preserves the central empirical conclusion of MNZ.[2]

Figure 1 compares the output gaps from the unrestricted and zero covariance specifications. Panel (a) reports the BN and UC-UR gaps. Since these gaps are numerically identical within each sample, the panel plots one line per sample. Panel (b) reports the corresponding UC-0 gaps. In both panels, the estimate for the original MNZ sample is shown alongside the estimate for the sample extended through 2026.

[2]We note that the estimate of the correlation is -1 with a standard error of 0.00 when we extend the UC-UR to 2026. This is due to the effect of the Covid-19 pandemic. When we estimate the model up till 2019Q4, we estimate a correlation of -0.92 with a standard error of 0.08. When we estimate the UC-UR by treating the two outlier quarters of 2020Q1 and 2020Q2 as being missing (i.e. setting the Kalman gain to zero), we obtain an estimate of -0.91 with a standard error of 0.07. We also note that while the extended sample changes the filtered estimate of the trend relative to the original MNZ paper, either stopping estimation before the pandemic or treating the outlier quarters as missing observations yields estimates of the cycle from the UC-UR that are very similar to those of MNZ. We note that a more explicit treatment of the pandemic would require one to model some form of heteroskedasticity in the UC innovations, which would make maximum likelihood estimation challenging as these likely involve Bayesian methods. Moreover, heteroskedasticity in the innovations implies that the reduced form becomes an ARMA with time-varying coefficients (e.g., see Johannsen and Mertens, 2021), which will once again likely involve Bayesian methods. For these reasons, we have chosen not to consider an explicit treatment of the Covid outliers given it complicates comparisons relative to the original MNZ paper, and is also somewhat secondary to the point we will make about multivariate information later.

Table 1: Narrow Replication

		Original sample 1947–1998	Extended sample 1947–2026
Restricted UC model (UC-0)			
Trend process	Drift μ	0.81 (0.05)	0.76 (0.05)
	Innovation σ_η	0.69 (0.12)	0.97 (0.08)
Cycle process	ϕ_1	1.53 (0.11)	1.49 (0.16)
	ϕ_2	-0.61 (0.14)	-0.59 (0.17)
	Innovation σ_ε	0.62 (0.15)	0.42 (0.16)
	Implied period	7.80 years	6.30 years
	Log likelihood	-286.61	-475.96
BN reduced form: ARIMA(2,1,2)			
	Drift μ	0.82 (0.09)	0.76 (0.21)
	ϕ_1	1.34 (0.14)	1.49 (0.30)
	ϕ_2	-0.71 (0.15)	-0.98 (0.38)
	θ_1	-1.05 (0.18)	-1.44 (0.57)
	θ_2	0.52 (0.19)	0.93 (0.60)
	SE of regression	0.97 (0.05)	1.07 (0.08)
	$\kappa(1)$	1.28	1.03
	Implied period	2.40 years	2.20 years
	Log likelihood	-284.65	-471.03
Unrestricted UC model (UC-UR)			
Trend process	Drift μ	0.82 (0.09)	0.76 (0.06)
	Innovation σ_η	1.24 (0.16)	1.10 (0.00)
Cycle process	ϕ_1	1.34 (0.15)	1.49 (0.00)
	ϕ_2	-0.71 (0.15)	-0.98 (0.00)
	Innovation σ_ε	0.75 (0.29)	0.10 (0.00)
	Implied period	2.40 years	2.20 years
	Covariance $\sigma_{\eta\varepsilon}$	-0.84 (0.40)	-0.11 (0.00)
	Correlation $\rho_{\eta\varepsilon}$	-0.91 (0.07)	-1.00 (0.00)
	Log likelihood	-284.65	-471.03

Notes: Standard errors are in parentheses. Some extended-sample UC-UR standard errors round to 0.00 because the correlation of trend and cycle innovations is estimated at the admissible boundary. $\kappa(1)$ denotes the long-run multiplier used in the BN decomposition.

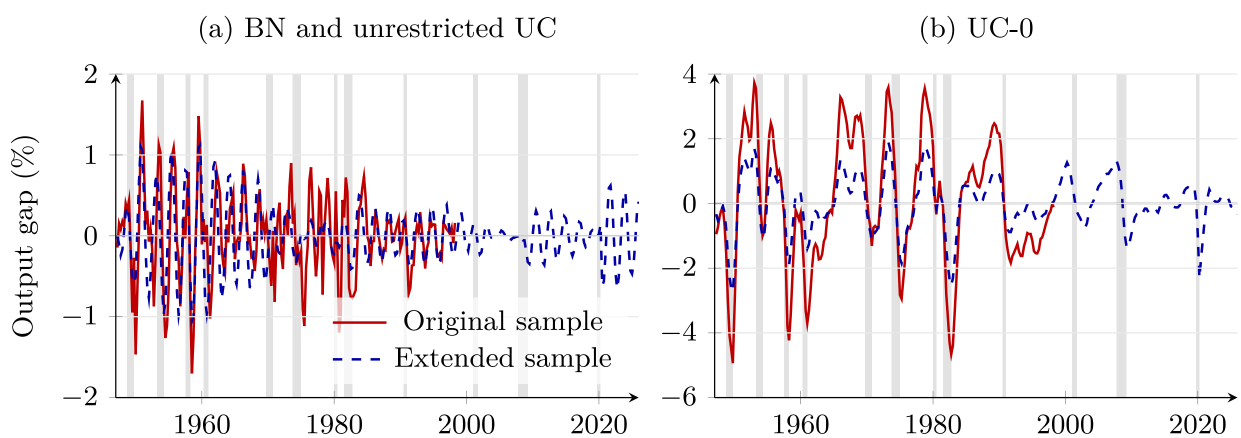

Figure 1: Narrow replication output gaps. Panel (a) shows the coincident BN and unrestricted-UC gaps, and Panel (b) shows the UC-0 gaps. NBER recessions are shaded.

The narrow replication therefore confirms the main empirical result in MNZ. Once the covariance restriction is removed, the univariate BN and unrestricted UC decompositions deliver the same cyclical component. Figure 1 also shows how the zero covariance restriction affects the properties of the estimated cycle. The UC-0 gaps are substantially larger and more persistent, whereas the unrestricted gaps display considerably more high-frequency variation. This contrast helps explain the prevalence of the zero covariance restriction in applied UC models. It also motivates the multivariate extension, where we examine whether additional cyclical information can yield persistent output gaps while leaving the covariance between trend and cycle innovations unrestricted.

3 Enlarging the Information Set

We consider an enlarged information set by considering three variables in addition to real GDP: U-2 unemployment rate, industrial production growth, and hours growth. The set of multivariate information is small by design, so that we can continue to use the maximum likelihood approach used by MNZ and avoid Bayesian regularization or factor-based dimension reduction that is often useful in large-dataset output-gap models. The choice of U-2 unemployment rate, defined as job losers as a percentage of the labor force, follows the argument in Berger, Boll, Morley, and Wong (2022) and has a clear economic motivation. Job loss is closely tied to cyclical contractions in labor demand, whereas the headline unemployment rate also reflects more persistent movements in labor force participation and other low-frequency labor-market developments. U-2 is therefore a more direct measure of cyclical labor-market slack. We add industrial production growth as a timely real-activity indicator that tends to move strongly over the business cycle. Hours growth captures variation in labor input on the intensive margin.

As far as we are able to, we implement the enlarged information set extension as symmetrically as we can across the two methods. For the BN decomposition, we add these three variables to real GDP growth to estimate a four-variable VAR and obtain the output gap as the BN cycle of real GDP, similar to what is done by Morley and Wong (2020).[3] For the UC model, in addition to equations (1)–(3), we append additional observation equations that link the additional indicators to the output cycle:

$$x_t = \gamma + \alpha c_t + u_t, \qquad u_t \sim N(0, \Sigma_u). \tag{6}$$

Here $x_t = (u2_t, ip_t, h_t)'$ collects U-2 unemployment, industrial production growth, and hours growth. The vector γ contains intercepts, α contains the loadings of the indicators on the output cycle c_t, and u_t collects idiosyncratic measurement disturbances with diagonal covariance matrix Σ_u. The covariance between the GDP trend and cycle innovations remains unrestricted.[4] Thus the exercise changes the information set while keeping the central MNZ structure in place.[5]

Table 2 compares the parameter estimates from the unrestricted univariate UC model, UC-0, and the UC model with cyclical indicators. The sum of the cycle autoregressive coefficients is 0.87 in the model with cyclical indicators, close to the UC-0 estimate of 0.90 and compared with 0.51 in the unrestricted univariate

[3]The VAR is estimated with two lags by least squares and the BN cycle of real GDP is calculated based on the VAR coefficients (see Morley, 2002, for details). We also choose a low order VAR mostly to obviate some of the overfitting issues discussed by Morley and Wong (2020). The choice of two lags also reflects that in the hypothetical case where the cycle is observed, the UC model has a restricted VAR(2) representation of the observed cycle and GDP growth (see Kamber, Morley, and Wong, 2018) which follows from the cycle being an AR(2). In general, given the interpretation that additional information is more likely to be the true cycle with some measurement error, this would induce possible MA dynamics and the reduced form is likely to be VARMA rather than a finite VAR. That said, we see using a VAR(2) as a reasonable compromise for the sake of our example since VARMAs are typically more challenging to estimate in finite samples, particularly with maximum likelihood methods.

[4]A closely related result is reported by Basistha (2007). Using a bivariate correlated UC model for Canadian output and inflation, he obtains a persistent cycle together with a strongly negative correlation between trend and cycle innovations and shows in Monte Carlo experiments that the bivariate specification improves the precision of the estimated trend-cycle parameters.

[5]We note that this differs from multivariate UC models in the broader literature, which often apply a UC decomposition to all observed variables, for example by estimating both an output gap and an unemployment gap, as in Sinclair (2009), or by specifying a common cyclical component across multiple indicators, as in Basistha and Startz (2008). In our model, the additional series enter only as part of the information set used to estimate the output cycle, rather than receiving their own trend-cycle decompositions. This keeps the exercise symmetric with the VAR, which also includes the additional variables as information while remaining focused on a single cycle.

Table 2: UC parameter estimates

		Univariate UC-UR 1947–2026	Univariate UC-0 1947–2026	UC with cyclical indicators 1967–2026
Trend process	Drift μ	0.76 (0.06)	0.76 (0.05)	0.68 (0.05)
	Innovation σ_η	1.10 (0.00)	0.97 (0.08)	0.70 (0.15)
Cycle process	ϕ_1	1.49 (0.00)	1.49 (0.16)	0.76 (0.06)
	ϕ_2	-0.98 (0.00)	-0.59 (0.17)	0.11 (0.07)
	Innovation σ_ε	0.10 (0.00)	0.42 (0.16)	1.07 (0.22)
Trend-cycle innovations	Covariance $\sigma_{\eta\varepsilon}$	-0.11 (0.00)	0 (restricted)	-0.31 (0.33)
	Correlation $\rho_{\eta\varepsilon}$	-1.00 (0.00)	0 (restricted)	-0.41 (0.28)
Indicator loadings	U-2	—	—	-0.65 (0.11)
	Industrial production	—	—	0.17 (0.06)
	Hours	—	—	0.19 (0.05)

Notes: Standard errors are in parentheses. The sample for the UC model with cyclical indicators starts in 1967, when U-2 data become available.

model. At the same time, the covariance between trend and cycle innovations remains unrestricted, with an estimated correlation of -0.41. This pattern indicates that cyclical information can generate persistence comparable to UC-0 within a model that estimates the trend-cycle covariance. Figure 2 reports the central result for the enlarged information set. The BN and unrestricted-UC gaps are closely aligned, with a correlation of 0.995.

Relative to the univariate gap in Figure 1, the estimated cycle is also more consistent with standard business-cycle chronology. It is negative around the mid-1970s and early-1980s recessions, declines markedly during the global financial crisis, and displays the large but short-lived contraction associated with the Covid-19 recession. Over the pre-Covid sample, this chronology is broadly consistent with the multivariate BN estimates in Morley and Wong (2020), where the output gap is positively associated with NBER reference cycles and is most clearly negative around recessions and the ends of expansions.

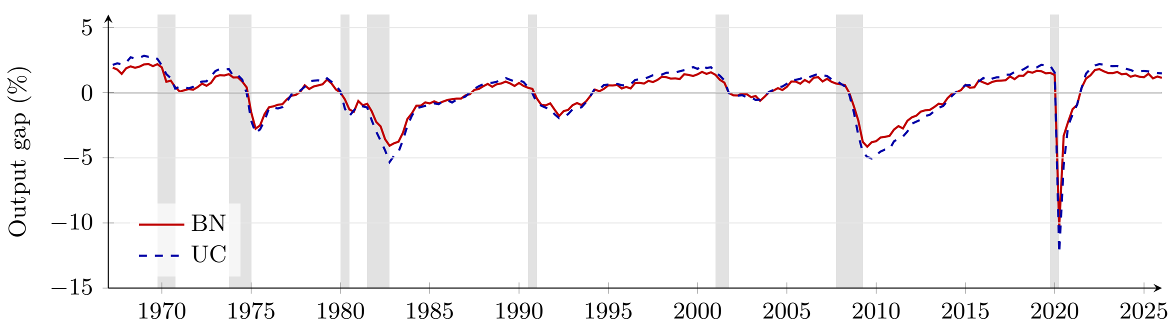

Figure 2: BN and unrestricted-UC gaps with cyclical indicators.

Using a BN decomposition of a VAR, we can summarize which innovations account for movements in the gap using the informational decomposition developed by Morley and Wong (2020). We decompose the BN cycle into the historical contributions of the VAR forecast errors. These contributions add up to the estimated output gap in each period. Figure 3 shows that all four variables are informative for the gap, even though their relative importance varies over time. The early-1980s recession is broad based, with negative contributions from all four variables. The global financial crisis is accounted for mainly by industrial production and hours, with U-2 also contributing negatively. The Covid-19 recession is different, with the large temporary fall in the gap dominated by U-2.

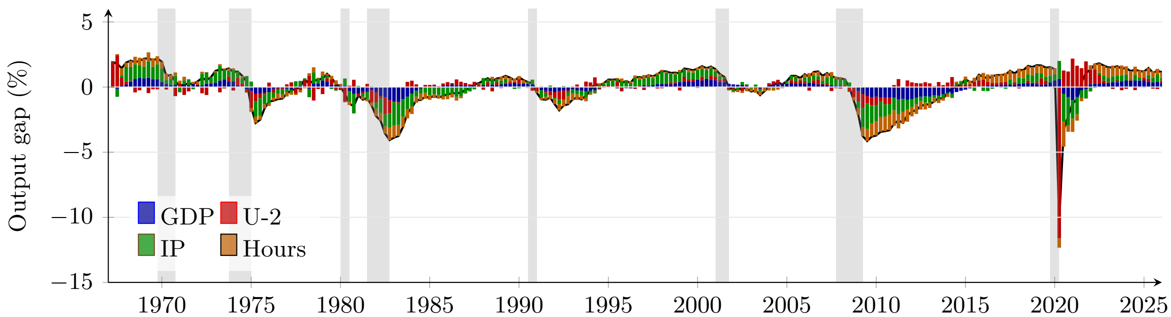

Figure 3: Historical decomposition of the multivariate BN gap. The bars show the accumulated contribution of each VAR innovation to the BN cycle; the black line is the BN gap.

Table 3: BN informational decomposition

Variable	Standard deviation of informational contribution
GDP	0.40
U-2	0.90
IP	0.65
Hours	0.47

Notes: Entries are the standard deviations of the informational contributions of each variable to the BN gap, following Morley and Wong (2020).

Table 3 summarizes this decomposition by reporting the standard deviation of each variable's informational contribution to the output gap, following Morley and Wong (2020). The largest entry is for U-2, so the decomposition points to job-loss unemployment as the most relevant conditioning variable in the BN system. Industrial production is the next most relevant source of information, while GDP and hours contribute less by this metric.

The multivariate extension also clarifies how our exercise relates to the existing output-gap literature. The comparison remains between BN and unrestricted UC estimates constructed from a common information set, but that information set now includes additional indicators of cyclical conditions. In this respect, the exercise is closer to Morley and Wong (2020) and Barigozzi and Luciani (2023) than to an alternative univariate restriction. The additional variables help identify the cyclical state, and the BN and UC decompositions remain closely aligned once this information is included.

Taken together, the results suggest a different emphasis on the univariate finding. In the MNZ model, the covariance restriction is the formal distinction between UC-0 and the unrestricted UC/BN representation. The volatility of the unrestricted univariate cycle also reflects the limited information in GDP alone. Once the decomposition is conditioned on labor-market and real-activity indicators, BN and unrestricted UC remain closely aligned while the resulting gap becomes closer to a conventional business-cycle measure.[6]

[6]We note that the conditioning on cyclical variables does imply we interpret these cyclical variables as *only* containing information about the cycle. This is a subtle difference when one considers multivariate UC models that estimates trend-cycle decompositions for multiple variables. For example, Sinclair (2009) estimates the natural rate of unemployment and the output gap jointly in a bivariate UC model. In such a setting, the unemployment rate is also in the information set. But in such settings, the unemployment rate has information about both trend unemployment and cyclical unemployment, where only the latter is related to the cycle. A necessary, though perhaps not sufficient, condition for cyclical variables *only* containing information about the cycle in our context is that these variables have to be stationary, a condition easily satisfied with all three cyclical pieces of information we consider in our exercise.

4 Reconciling output gap estimates and information sets

MNZ have been influential in showing that output gaps obtained via a BN decomposition and a standard UC model are identical once one relaxes the covariance restriction between trend and cycle innovations. By adopting a wider replication and understanding the role of information sets, our work also potentially reconciles the role of the zero covariance restriction and how one thinks about alternative methods that use the BN decomposition in both univariate and multivariate environments, such as Kamber, Morley, and Wong (2018) and Morley and Wong (2020) respectively.

Our analysis potentially helps reconcile the often-used zero covariance restriction in UC models. As we have shown, one does not necessarily require a zero covariance restriction between trend and cycle to obtain a large and persistent output gap — a larger information set can deliver the same large and persistent output gap whether in the UC or BN setting.[7] The conclusion we thus draw from this exercise is that it may be the lack of inclusion of relevant information that is preventing us from estimating large and persistent cycles. It suggests that in a setting where this information is available, it is perhaps advisable for the analyst to build this information into either their UC model or a BN decomposition of a multivariate model like a VAR. However, lacking this information, because real GDP on its own is only weakly informative about the cycle, the zero restriction may be serving the role of embedding this information in the trend-cycle decomposition problem. While we use maximum likelihood to align with the original MNZ paper, we also note that UC models are now routinely estimated using Bayesian methods, and a common practice is to specify a prior where the variance of trend innovations is small. In some regards, one can also view the specification of such a prior as the analyst bringing some non-data information to bear as the data may be non-informative about the volatility of trend innovations. We should add that our discussion is not unique to the UC model. One can view the approach by Kamber, Morley, and Wong (2018), where they impose a signal-to-noise ratio restriction when they apply a BN decomposition to a long order AR model as an exercise in the same spirit as placing a prior on the variance on trend innovations in the UC model. Viewed from this perspective, one could view these as approaches for the analyst to circumvent the lack of useful multivariate forecasting information by appealing to extraneous information.

5 Conclusion

MNZ trace the large difference between standard UC and BN decompositions of GDP to the zero covariance restriction imposed in the standard UC model. Our narrow replication confirms this result in the original sample and in a longer sample ending in 2026. The BN-ARMA and unrestricted-UC likelihoods coincide, and the implied correlation between trend and cycle innovations is strongly negative.

The multivariate extension adds a small set of labor-market and real-activity indicators to the MNZ framework and estimates the BN and unrestricted-UC decompositions from this common information set. The resulting gaps are highly correlated and visually close. At the same time, the estimated cycle has the persistence and timing associated with a conventional business-cycle measure, including pronounced contractions around the early 1980s, the global financial crisis, and the Covid-19 recession.

[7] We note that there are occurrences within the UC literature that incorporate extra information, such as using surveys together with the zero covariance restriction between trend and cycle (e.g. Mertens, 2016; Chan, Clark, and Koop, 2018; Zaman, forthcoming) and also versions of multivariate versions of UC models that do not necessarily restrict the covariance between trend and cycle innovations (e.g. Garnier, Mertens, and Nelson, 2015; Leiva-Leon, Sekkel, and Uzeda, 2026). Nonetheless, our observation that one can use multivariate information *and* free up the zero correlation, especially within the context of the output gap models, has, to our knowledge, never been made.

The conclusion is therefore deliberately narrow. The volatile and noisy cycle from the unrestricted univariate model reflects the limitations of using GDP alone to measure cyclical conditions. Conditioning the decomposition on additional cyclical indicators preserves the close connection between BN and unrestricted UC and yields a cycle that is more useful as a measure of cyclical conditions.

References

Barigozzi, M., and M. Luciani (2023): "Measuring the Output Gap Using Large Datasets," *The Review of Economics and Statistics*, 105(6), 1500–1514.

Basistha, A. (2007): "Trend-Cycle Correlation, Drift Break and the Estimation of Trend and Cycle in Canadian GDP," *Canadian Journal of Economics*, 40(2), 584–606.

Basistha, A., and R. Startz (2008): "Measuring the NAIRU with Reduced Uncertainty: A Multiple-Indicator Common-Cycle Approach," *The Review of Economics and Statistics*, 90(4), 805–811.

Berger, T., P. D. Boll, J. Morley, and B. Wong (2022): "Cyclical Signals from the Labor Market," *Oxford Open Economics*, 1(1), 1–16.

Chan, J. C. C., T. E. Clark, and G. Koop (2018): "A New Model of Inflation, Trend Inflation, and Long-Run Inflation Expectations," *Journal of Money, Credit and Banking*, 50(1), 5–53.

Del Negro, M., D. Giannone, M. P. Giannoni, and A. Tambalotti (2017): "Safety, Liquidity, and the Natural Rate of Interest," *Brookings Papers on Economic Activity*, 2017(1), 235–316.

Garnier, C., E. Mertens, and E. Nelson (2015): "Trend Inflation in Advanced Economies," *International Journal of Central Banking*, 11(S1), 65–136.

Johannsen, B. K., and E. Mertens (2021): "A Time-Series Model of Interest Rates with the Effective Lower Bound," *Journal of Money, Credit and Banking*, 53(5), 1005–1046.

Kamber, G., J. Morley, and B. Wong (2018): "Intuitive and Reliable Estimates of the Output Gap from a Beveridge-Nelson Filter," *The Review of Economics and Statistics*, 100(3), 550–566.

Leiva-Leon, D., R. Sekkel, and L. Uzeda (2026): "Do Monetary Policy Shocks Affect the Neutral Rate of Interest?," Staff Working Paper 2026-6, Bank of Canada.

Maffei-Faccioli, N. (2025): "Identifying the Sources of the Slowdown in Growth: Demand Versus Supply," *Journal of Applied Econometrics*, 40(2), 181–194.

Mertens, E. (2016): "Measuring the Level and Uncertainty of Trend Inflation," *The Review of Economics and Statistics*, 98(5), 950–967.

Morley, J., and B. Wong (2020): "Estimating and Accounting for the Output Gap with Large Bayesian Vector Autoregressions," *Journal of Applied Econometrics*, 35(1), 1–18.

Morley, J. C. (2002): "A State-Space Approach to Calculating the Beveridge-Nelson Decomposition," *Economics Letters*, 75(1), 123–127.

Morley, J. C., C. R. Nelson, and E. Zivot (2003): "Why Are the Beveridge-Nelson and Unobserved-Components Decompositions of GDP So Different?," *The Review of Economics and Statistics*, 85(2), 235–243.

Sinclair, T. M. (2009): "The Relationships between Permanent and Transitory Movements in U.S. Output and the Unemployment Rate," *Journal of Money, Credit and Banking*, 41(2–3), 529–542.

Zaman, S. (forthcoming): "A Unified Framework to Estimate Macroeconomic Stars," *The Review of Economics and Statistics*, pp. 1–45, Advance online publication.