ACI Research Paper #16-2026

GDP Trends and Cycles in Emerging Market Economies

Tingbin BIAN

Qu FENG

Benjamin WONG

August 2026

Please cite this article as:

Bian, Tingbin, Qu Feng, and Benjamin Wong, "GDP Trends and Cycles in Emerging Market Economies", Research Paper #16-2026, *Asia Competitiveness Institute Research Paper Series (August 2026)*

GDP Trends and Cycles in Emerging Market Economies

Tingbin Bian[*], Qu Feng[†], Benjamin Wong[‡]

August 5, 2026

Abstract

We propose using the permanent income hypothesis to decompose GDP into its permanent (i.e., trend) and transitory (i.e., cyclical) components. While the method is inherently non-parametric, we show that it aligns with a typical unobserved components model used to detrend real GDP and will thus recover the true permanent component as long as the permanent income hypothesis holds. The permanent income hypothesis holding has the implication that consumption and GDP share the same permanent component, including but not limited to, features such as a common stochastic permanent component. While the broader literature is replete with methods to perform trend-cycle decomposition, we show that the proposed method is perhaps most useful when considering features that one naturally associates with emerging market economies, such as more volatile permanent shocks and/or frequent breaks in trend growth rates. In cross-country applications, it produces more plausible cyclical measures and improves out-of-sample forecasting performance relative to standard filters. In a case study of China, it yields a more reliable cycle that aligns closely with major historical events.

Keywords: Trends and Cycles, Emerging Market Economies, Consumption, Permanent Income Hypothesis
JEL classification codes: C32, E21, E32

[*]Asia Competitiveness Institute, Lee Kuan Yew School of Public Policy, National University of Singapore, Singapore. Email: tingbin.bian@nus.edu.sg

[†]Economics Division, School of Social Science, Nanyang Technological University, Singapore. Email: qfeng@ntu.edu.sg

[‡]Department of Econometrics and Business Statistics, Monash Business School, Monash University, Australia. Email: benjamin.wong@monash.edu

1 Introduction

The permanent component of GDP measures long-run (or potential) output and shares the same order of integration as GDP. Because the unobserved components model defines the permanent component of GDP as a random walk process, shocks to the trend have permanent effects on output (Watson, 1986). As emphasized by Aguiar and Gopinath (2007), emerging market economies differ from advanced economies in that they are subject to larger permanent shocks, which correspond to a larger variance at frequency zero in the spectral density (Cochrane, 1988). Figure 1 compares the spectral densities of GDP growth for the two groups. The top panels report the normalized spectral densities: emerging market economies exhibit substantially greater spectral mass at or near frequency zero, whereas advanced economies concentrate their spectral mass in a relatively higher frequency band (Levy & Dezhbakhsh, 2003).[1] The bottom panels report cumulative spectral densities, which show that the overall volatility of GDP growth is markedly higher in emerging market economies.

A second distinguishing feature of emerging market economies is the prevalence of structural breaks in long-run (trend) growth. Figure 2 plots quarterly GDP growth, sub-sample mean growth rates, and estimated break dates for both groups. For advanced economies, changes in mean growth across breaks are relatively small; for emerging market economies, the shifts in mean growth following breaks are considerably larger. Taken together, these two features, greater volatility at frequency zero or low frequency and more pronounced structural breaks, challenge the definition of trend in the standard trend-cycle decomposition methods, e.g., the unobserved components model and the HP filter. As a result, the estimated permanent components, and analogously the cyclical component, can be distorted.

[1] The 6–32 quarter frequency band is often referred to as the business-cycle band in the context of band-pass filters.

Figure 1: **Comparison of Cumulative spectral Density**

(a) **Normalized spectral Density**

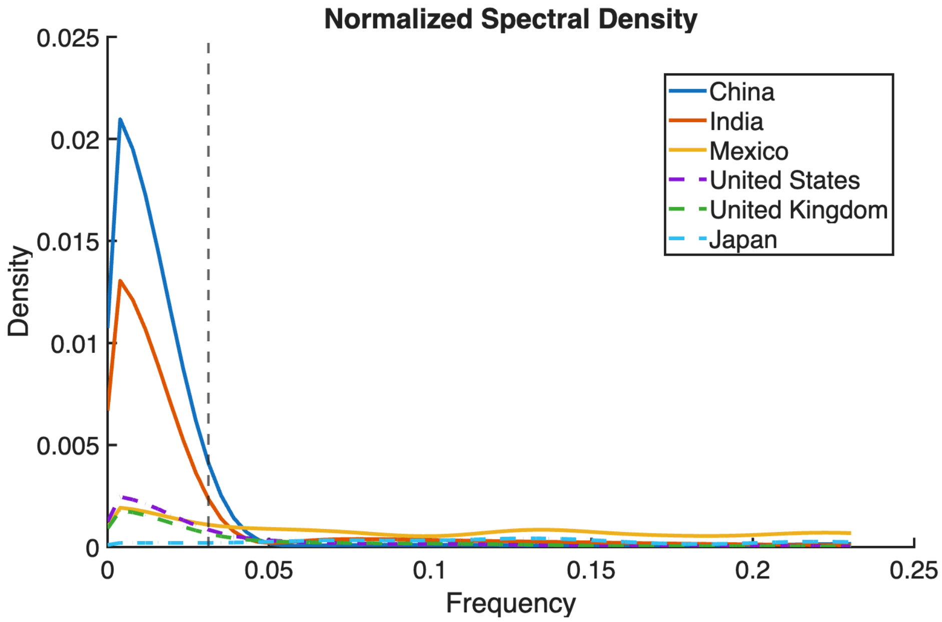

(b) **Non-Normalized Cumulated spectral Density**

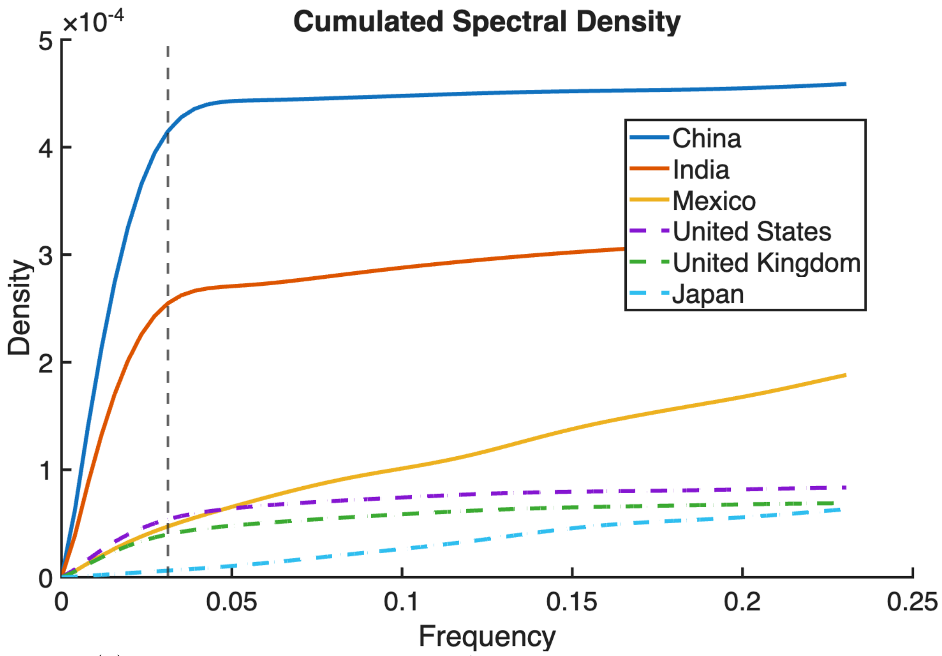

Note: (1). We exclude relatively high-frequency bands, as our primary interest lies in low-frequency fluctuations. The vertical dashed line indicates the 8-year frequency threshold. Variations to the left represent fluctuations with periods longer than 8 years, while those to the right correspond to shorter-term fluctuations. (2). Emerging market economies exhibit substantially greater spectral mass at frequency zero or low-frequency band, while advanced economies display a concentration of spectral mass in a relatively higher frequency band. (3). The overall volatility of GDP growth is markedly higher in emerging market economies.

Figure 2: **Comparison of Annualized Quarterly GDP Growth between Emerging Market Economies and Advanced Economies**

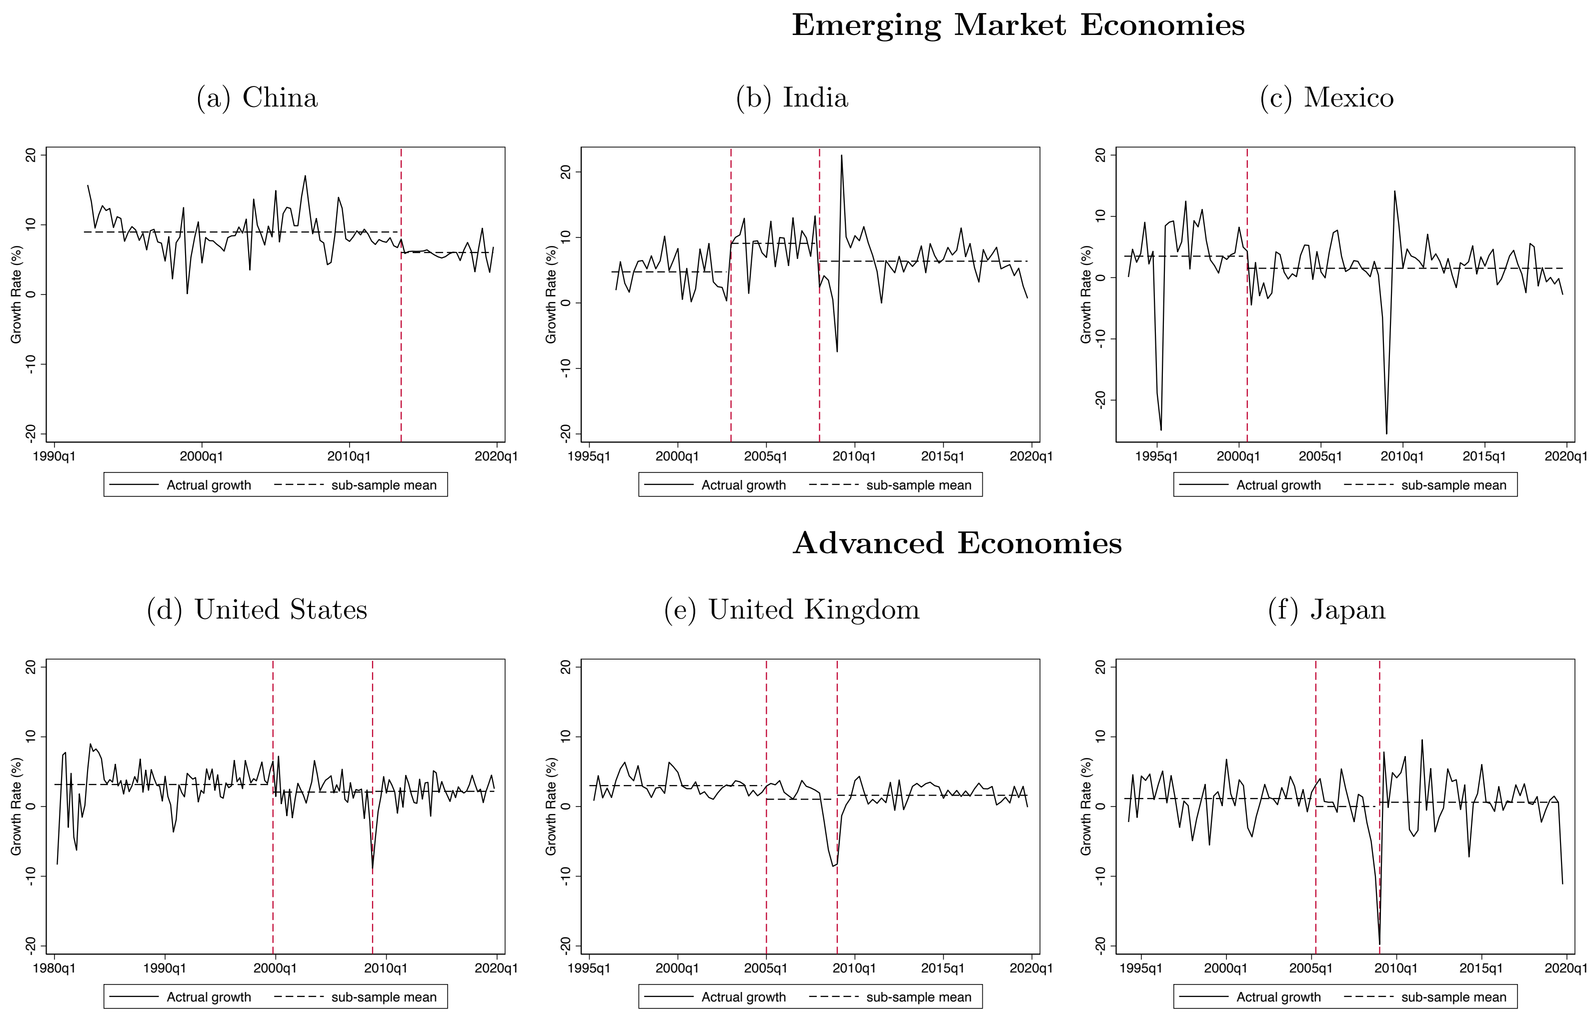

Note: We use Bai and Perron (1998) to identify structural breaks in the constant term of GDP growth in AR model. When the test detects more than two breakpoints, we restrict the number of breaks to two.

Our main contribution in this paper is to show that the permanent income hypothesis could help to identify GDP trends and cycles, particularly in emerging market economies. Under the standard permanent income hypothesis benchmark, consumption responds only to permanent income shocks; consequently, consumption and income share a common stochastic trend and consumption behaves as a random walk (Cochrane, 1994; Hall, 1978; King et al., 1991). These restrictions imply that consumption can be used to pin down the trends of GDP and, in turn, its cyclical component. Specifically, we show a simple non-parametric expression of GDP trends and cycles: the permanent component of output is a linear function of consumption, and the cyclical component can be identified from the gap between output and consumption, up to a steady-state wedge that reflects long-run saving behavior. Although the method is inherently non-parametric, it aligns with a canonical unobserved components model and can equivalently be interpreted as a Beveridge–Nelson decomposition (see Section 2). When the permanent income hypothesis holds, the approach recovers the true permanent and cyclical components; more importantly, relative to parametric models, it reduces reliance on correctly specifying the trend process, a central concern in emerging market economies, where trend shocks are large and trend growth may shift.

The observation that consumption can serve as a direct measure of the GDP trend is not itself new (see Crucini & Shintani, 2015). Our contribution is fourfold. First, we show why this approach is especially well suited to the features of emerging market economies. Second, we characterize, through Monte Carlo experiments, the conditions under which widely used alternatives fail. Third, we generalize the method to accommodate empirically relevant departures from the strict permanent income hypothesis. Fourth, we validate it in a broad cross-country sample and in a detailed case study of China. A natural concern is that the assumptions underlying the permanent income hypothesis, the absence of liquidity constraints and precautionary saving, may appear least tenable precisely in emerging market economies. We address this concern directly: our generalizations and robustness exercises show that the method continues to deliver reliable cycle estimates even when these assumptions are relaxed.

Our analysis proceeds in two steps. First, we use Monte Carlo experiments to illustrate how alternative decomposition methods can fail in environments that resemble emerging market economies. When permanent shocks are large or trend growth contains breaks, these methods tend to confound trend variation with cyclical movements, whereas the consumption-based method remains accurate as long as the permanent income hypothesis holds. Second, we extend the baseline framework to allow for empirically relevant departures from the strict permanent income hypothesis, including a transitory component in consumption, time variation in the steady-state wedge between consumption and income, and outliers such as the COVID-19 shock. Returning to the Monte Carlo design, we show that the modified procedure continues to provide reliable cycle estimates and often outperforms competing methods precisely in the environments that resemble emerging market economies.

On the empirical side, we first examine "ideal" cases using pre-COVID data, highlighting when the baseline method is expected to work well. Next, we apply the generalized specifications to a broad set of countries and evaluate the resulting cycles through their pseudo-out-of-sample forecasting performance. Finally, we study China as a case study: we compare alternative decompositions, relate the estimated cycles to major historical events, and discuss the applications for interpreting China's trend growth and cyclical fluctuations.

The rest of the paper is organized as follows. Section 2 introduces the consumption-based trend-cycle decomposition method. Section 3 presents the generalizations of the baseline framework and uses Monte Carlo simulations both to illustrate why commonly used filters can fail and to demonstrate the performance of the consumption-based method. Section 4 evaluates the implied cycle measures using pseudo-out-of-sample forecasting. Section 5 applies the framework to China. Section 6 concludes.

2 Linking consumption to the permanent component of output

At a high level, the method that we propose builds on the insight that consumption and GDP share a common stochastic trend. In turn, consumption becomes additional multivariate information to pin down the trend, and analogously the cycle, in real GDP. We note that this insight has previously been used by, among others, Stock and Watson (1988), King et al. (1991), Cochrane (1994), Morley (2007), and Crucini and Shintani (2015) to perform trend-cycle decompositions of real GDP for the U.S. or advanced economies.

We begin with a description of the problem, where the natural log of real GDP, y_t, can be decomposed into a permanent (trend) component τ_t and a stationary cyclical component $\tilde{y}_t$.:

$$y_t = \tau_t + \tilde{y}_t \tag{1}$$

The goal of trend-cycle decomposition is to accurately obtain these permanent and transitory components. To fix ideas, we consider a fairly canonical unobserved component model (e.g., see Watson (1986)) for the permanent and cyclical components:

$$\tau_t = \mu + \tau_{t-1} + e_{\tau,t},\ e_{\tau,t} \sim (0, \sigma^2_{e,\tau}) \tag{2}$$

$$\phi(L)\tilde{y}_t = e_{\tilde{y},t},\ e_{\tilde{y},t} \sim (0, \sigma^2_{e,\tilde{y}}) \tag{3}$$

where μ is the drift, which serves as the mean growth rate of real GDP, $\phi(L) = I - \phi_1 L - \ldots \phi_p L^p$ is a lag polynomial whose roots are outside the unit circle;[2] and $\sigma^2_{e,\tau}$ and $\sigma^2_{e,\tilde{y}}$ are the variances of the innovations in the permanent and cyclical components, respectively. The permanent component is thus the random walk component. The model above can be cast into state-space form, where one can obtain estimates of trend and cycles via the Kalman filter. Moreover, the above is connected to the Beveridge–Nelson decomposition, where the Kalman-filtered trend is equivalent to the Beveridge–Nelson decomposition of the ARIMA representation of y_t (see Morley et al., 2003).

Suppose that now one has information about the permanent component. Under balanced growth in a canonical real business cycle model (e.g., see King et al., 1991), GDP and consumption share a common stochastic trend. Together, these imply (e.g., see Morley, 2007):

[2] We describe the cyclical component as being a stationary autoregressive process, though this does not preclude any stationary invertible ARMA process.

$$c_t = \bar{c} + \gamma\tau_t + \tilde{c}_t \tag{4}$$

where c_t is the log of consumption, γ is the marginal propensity to consume out of permanent income, and $\tilde{c}_t$ is the cyclical component of consumption, which is stationary. The constant $\bar{c}$ captures the long-run impact of taxes and private saving on consumption. The permanent income hypothesis allows us to simplify Equation (4) further. Because the hypothesis implies that consumption responds only to changes in permanent income, consumption growth is unpredictable and consumption itself follows a random walk. Equation (4) then reduces to:

$$c_t = \bar{c} + \tau_t \qquad \text{or} \qquad \tau_t = c_t - \bar{c} \tag{5}$$

where GDP and consumption will cointegrate $[1 - 1]$ and $\gamma = 1$. Because consumption is unpredictable and follows a random walk, this also implies that consumption has no cyclical component. Together, these observations imply that if the permanent income hypothesis holds, consumption serves as a direct measure of the permanent component of real GDP. Relative to a standard unobserved components model, this amounts to excluding the transitory consumption component from Equation (4). This restriction follows directly from the standard permanent income hypothesis, under an intertemporal Euler equation with uncertainty and quadratic utility; we treat it as a benchmark and relax it in the generalizations of Section 3, which is especially relevant for emerging market economies. Empirically, the random-walk hypothesis can be tested following Hall (1978).

Combining (5) with (1), this implies that

$$\tilde{y}_t = y_t - c_t + \bar{c} \tag{6}$$

where $c_t - \bar{c}$ is the permanent component of real GDP. We can thus obtain the output gap (e.g. the cycle of real GDP) as $\tilde{y}_t = y_t - c_t + \bar{c}$. Since y_t and c_t cointegrate, $\bar{c}$ is the wedge between consumption and income. We do not place any structure on $\bar{c}$ for now, but note that, under the permanent income hypothesis, this can be interpreted as the steady-state level of savings (see Crucini & Shintani, 2015). To see this, note that if income and consumption cointegrate, savings will be stationary. Under this interpretation, $-\bar{c}$ is the steady-state level of savings, and the level of savings is $\tilde{y}_t - \bar{c}$. Thus, the level of savings is this steady-state level of savings together with fluctuations of the output gap.

As a starting point, we note that the discussion above implies that one can estimate the trend-cycle decomposition of GDP above in a non-parametric fashion. That is, one can calculate $\tilde{y}_t = y_t - c_t + \bar{c}$ without needing to specify the underlying unobserved components model implied by Equations (2) and (3). We note that this observation has been made by Crucini and Shintani (2015) generally, but our argument is that such an approach may be of utility, especially for emerging market economies for two key reasons.

First, as discussed in the introduction, emerging market economies experience both volatile breaks in trend growth and shocks to trend. That is, in Equation (2), μ may not be a constant and $\sigma^2_{e,\tau}$ is larger than in advanced economies. Breaks in μ can be somewhat problematic as such features can require one to be explicit about the form and location of such breaks in a parametric model. If the permanent income hypothesis holds, μ is part of permanent income, and therefore, consumption will provide sufficient information about the trend (and thus cycle) of output. More volatile shocks to the trend (i.e. $\sigma^2_{e,\tau}$ being larger) may also make the non-parametric approach more appealing. Suppose one ignores information on consumption, a common approach is simply to use the HP filter to obtain the trend and cycle of output. The HP filter has a tuning parameter which relies on the ratio of shocks to trend relative to cycle. A common tuning parameter is 1600 for the U.S. If emerging market economies experience more volatile shocks to the trend, it is unclear whether such a tuning parameter is appropriate. In principle, the non-parametric method should not require such knowledge, as any volatility in trend is ultimately reflected in consumption. Second, emerging market economies also have a shorter span of data due to data availability issues. The non-parametric method, in principle, should mitigate such concerns. In particular, it is known that short time series can pose specific challenges for estimating a parametric unobserved components model, one that the non-parametric model may mitigate.

Under the restriction $\gamma = 1$, household consumption serves as a direct proxy for the GDP trend: $\tau_t = -\bar{c} + c_t$. If the permanent income hypothesis holds, the non-parametric method will exactly estimate the true output gap. We note that it is possible that the non-parametric method may still be useful even with small departures from the pure permanent income hypothesis (i.e., γ is approximately, but not exactly 1, and a small cyclical component in consumption in Equation (4)). This is ultimately an empirical question that we will explore.

Moreover, this expression for the GDP trend could also be interpreted as an implementation of Beveridge–Nelson (BN) decomposition. The general reason is that the BN decomposition identifies the stochastic trend as the permanent component, and unobserved components models are intrinsically related to the BN decomposition (Morley et al., 2003). In our settings, under the standard permanent income hypothesis, consumption equals the present value of expected lifetime income, so only permanent shocks to income affect current consumption while transitory shocks have no impact. Consequently, consumption can be interpreted as the conditional expectation of the infinite-horizon path of income, which corresponds exactly to the definition of the BN permanent component.

An alternative way to see this equivalence is through a bivariate VECM representation of consumption and GDP growth, following Cochrane (1994).[3] As shown by Garratt et al. (2006), the BN trend can be interpreted as the conditional cointegrating equilibrium value, with deviations from this equilibrium being transitory and eventually reverting. In our framework, this implies that a linear combination of the BN trends of GDP and consumption is constant. When consumption follows a random walk (i.e., no transitory consumption component, and the BN trend of consumption equals consumption itself), this directly indicates that the BN trend of GDP is a linear function of consumption, same as in our previous analysis.

In what follows, we set up Monte Carlo exercises to explore the utility of the proposed non-parametric method in confronting EME features such as breaks in trend growth rates, larger volatility to trend shocks, and departures from the pure form of the permanent income hypothesis. We compare to standard trend-cycle decomposition methods such as the HP filter (tuning parameter 1600 and the Boosted HP filter by

[3] Morley (2007) shows the conversion between a UCM with a stochastic trend and a VARMA, and the VARMA can in turn be transformed into a VECM.

Phillips and Shi (2021)) and the Hamilton filter.

3 Baseline Model Performance and Generalizations

This section establishes two properties of the consumption-based method through Monte Carlo simulation. The method recovers the cycle exactly when the permanent income hypothesis holds, but its practical value rests on two things the hypothesis itself does not guarantee: that it remains accurate when the hypothesis holds only approximately, and that it outperforms standard alternatives precisely in the configurations, large permanent shocks and breaks in trend growth, that characterize emerging market economies. Taking the strict permanent income hypothesis as a benchmark, we first show that the cycle estimates from widely used univariate filters degrade exactly as the data acquire these emerging-market features, and that even a flexible bivariate Beveridge–Nelson VECM mismeasures the cycle once trend growth contains breaks, whereas the consumption-based method does not. We then relax the hypothesis along three empirically relevant dimensions—a transitory component in consumption, a time-varying consumption–income relationship, and outliers such as the COVID-19 shock—and show that the method continues to deliver reliable cycles, often still outperforming the alternatives in the very environments in which those alternatives are least reliable.

3.1 The Ideal Case: When the Permanent Income Hypothesis Holds Exactly

We begin with the ideal case in which the permanent income hypothesis holds exactly, so that the permanent component of output is, by construction, a linear function of consumption, regardless of the true trend specification. In this benchmark world, the consumption-based method recovers the true cycle exactly. The point of this subsection is therefore not that our method "wins" a horse race; perfection by construction would be uninformative on its own. The point is the pairing: if one finds this data-generating process plausible for emerging market economies, and the spectral-density and trend-break evidence of Section 1 suggests one should, then the standard tools mismeasure the cycle precisely in the configurations that characterize these economies, whereas a consumption-based reading does not. The pairing matters because the two features we impose are properties of the trend, whereas the permanent income hypothesis restricts only how consumption relates to that trend. A trend that is volatile or break-prone is therefore fully consistent with the hypothesis holding, so taking the data-generating process to be of this form is not a special assumption but the natural reading of an emerging-market trend through the lens of the hypothesis.

To make this concrete, we use Monte Carlo simulations calibrated to environments commonly observed in both advanced and emerging market economies, and we examine how each method behaves as the DGP acquires features characteristic of emerging market economies, larger trend shocks and structural breaks. The alternatives we compare are the HP filter, the boosted HP filter of Phillips and Shi (2021), and the Hamilton filter (Hamilton, 2018).[4] For a multivariate comparison, we also consider

[4] We report results for the boosted HP filter using BIC as the stopping rule. The results are very

the Beveridge–Nelson decomposition with the bivariate VECM of Cochrane (1994).[5]

Table 1: **Monte Carlo Simulation: Permanent Income Hypothesis Holds**

	DGP 1		DGP 2		DGP 3	
	Correlation	RMSE	Correlation	RMSE	Correlation	RMSE
HP Filter	0.733	1.809	0.559	2.429	0.723	1.838
Boosted HP-BIC	0.492	2.338	0.374	2.577	0.494	2.335
Hamilton Filter	0.624	2.974	0.463	4.490	0.556	3.479
BN-VECM	0.961	0.721	0.943	1.049	0.738	11.435
Consumption-based Method	1.000	0.000	1.000	0.000	1.000	0.000

Notes:
1. The table reports the correlation and root mean squared error (RMSE) between the estimated and true cycles.
2. In the baseline DGP, $y_t = \tau^t + \tilde{y}_t$, $c_t = \bar{c} + \tau^t$, $\tau^t = \mu + \tau_{t-1} + e_{\tau,t}$, $\tilde{y}_t = 1.53\tilde{y}_{t-1} - 0.61\tilde{y}_{t-2} + e_{\tilde{y},t}$, where $e_{\tau,t} \sim$ i.i.d. $N(0, \sigma^2_{e,\tau})$ and $e_{\tilde{y},t} \sim$ i.i.d. $N(0, \sigma^2_{e,\tilde{y}})$. In DGP 1, we use these parameter values: $\bar{c} = -11.34$, $\mu = 0.81$, $\sigma_{e,\tau} = 0.62$, $\sigma_{e,\tilde{y}} = 0.69$.
3. In the DGP 2, we allow for larger Trend Shocks, $\sigma_{e,\tau} = \sqrt{2}$.
4. In DGP 3, we allow for breaks in the mean growth rate μ, with the timing and magnitude drawn randomly.
5. The results are produced using seeds rng(123) and 1,000 iterations.

Across all DGPs in this subsection, we assume that the permanent income hypothesis holds exactly. That is, in Equation (4), $\gamma = 1$ and consumption contains no transitory component. Our primary focus is on the permanent component τ_t in Equation (2), examining how different decomposition methods perform when the underlying DGP features larger permanent shocks $\sigma_{e,\tau}$ and structural breaks in mean growth μ. We consider three DGPs chosen to reflect features representative of both advanced economies and emerging market economies:

1. **DGP 1 (Baseline):** A low signal-to-noise ratio in which permanent shocks are relatively small ($\sigma_{e,\tau} < \sigma_{e,\tilde{y}}$).

2. **DGP 2 (Larger Trend Shocks):** Permanent shocks are larger, capturing the higher volatility observed in emerging market economies.

3. **DGP 3 (Breaks in Trend Growth):** The mean growth rate μ experiences structural breaks, consistent with the large shifts documented in EME growth trends.

Table 1 summarizes the simulation results. As expected, the consumption-based method perfectly recovers the true cycle in all DGPs. DGP 1 serves as the natural reference point, featuring an advanced-economy calibration with a smooth trend and small permanent shocks. In this environment, both the HP and Hamilton filters perform reasonably well, yielding relatively high correlations with the true cycle and low RMSEs. This is consistent with their being designed for, and routinely applied to, such environments. The key pattern is what happens as the DGP takes on the features of emerging market economies. Under DGP 2, as the variance of permanent

similar when the ADF test is used as the stopping criterion.

[5] We also tried the bivariate UCM of Morley (2007). Indeed, our baseline DGP is nested within Morley's framework. However, that model is estimated using a Kalman filter, which requires many parameters to be identified and often suffers from an initial-value problem.

shocks rises, the HP and Hamilton filters increasingly misclassify trend movements as cyclical fluctuations, and their performance degrades relative to DGP 1: correlations fall and RMSE rises. Under DGP 3, mean growth exhibits a one-time break; because both the HP objective function and the Hamilton regression are local in nature, the break mainly affects estimates near the break date, so the deterioration relative to DGP 1, while present, is more modest than under DGP 2.

The Beveridge–Nelson decomposition based on a bivariate VECM performs well in DGP 1 and DGP 2, but for an instructive reason: when consumption is restricted to a random walk, the BN–VECM and the consumption-based method coincide, so its success in these columns is not an independent result but our own method in a different guise.[6] This is also why moving from DGP 1 to DGP 2 leaves it unaffected: the change in the signal-to-noise ratio does not alter a Blanchard and Quah (1989) decomposition. The two methods part company only under DGP 3, where the VECM does not model the break in trend growth and its cycle estimate deteriorates sharply, with a substantially larger RMSE, and this is precisely the configuration most relevant for emerging market economies.

Three conclusions follow. First, univariate filters are reliable only when the trend evolves smoothly, an assumption violated more often in emerging market economies. Second, adding consumption through a flexible multivariate model is not enough on its own: without restrictions that discipline the trend, the model can still mismeasure the cycle, as the VECM does once trend growth breaks. Third, when the permanent income hypothesis holds, the consumption-based method is exact regardless of how volatile or break-prone the trend is, precisely the conditions under which the standard tools break down. The force of the exercise lies in the pairing of the second and third points: to the extent one reads emerging-market output as having a volatile, break-prone trend, the cycle that standard filters report is mismeasured in exactly that case, while the consumption-based reading is not, a reading we test directly in Section 4. This is what motivates the closer examination of the permanent income hypothesis, and the empirical application, in the sections that follow.

3.2 Generalizing the Baseline Model

The baseline model illustrates the usefulness of the consumption-based approach. We now extend it to remain consistent with the spirit of the permanent income hypothesis while allowing modest departures from its strict assumptions. In empirical application, consumption may contain transitory components, the long-run relationship between consumption and income may change over time, and the data may be distorted by outliers such as the COVID-19 shock. We address each of these in turn, preserving the core intuition that consumption carries the permanent component of output.

Transitory Consumption

In the baseline model, we exclude the transitory consumption component, $\tilde{c}_t$, in line with the permanent income hypothesis implication that consumption follows a random walk. Empirically, however, the random-walk assumption may be rejected.

[6] We use two lags to estimate the bivariate VECM; it performs better than specifications with more lags, with additional results reported in the Appendix.

From an econometric standpoint, rejection may stem from structural breaks in the mean growth rate or from the limited length of the available sample. From the standpoint of economic theory, a transitory consumption component can be rationalized by liquidity constraints faced by "hand-to-mouth" consumers, or by habit formation and precautionary saving. We do not attempt to adjudicate among these explanations or to formally test for the presence of a transitory component. Our focus is the more practical question: if transitory consumption is present, can consumption still serve as a useful proxy for the trend, and to what extent does the method continue to outperform alternative decompositions?

Recall that the estimated cycle is $\hat{\tilde{y}}_t = y_t - \frac{1}{\gamma}c_t + \frac{1}{\gamma}\hat{\bar{c}}$. For simplicity, suppose the cointegration vector is $[1, -1]$ and that the estimated steady-state level of savings, $\hat{\bar{c}} = \overline{(y_t - c_t)}$, equals the true value. The estimated cycle then simplifies to $\hat{\tilde{y}}_t = \tilde{y}_t - \tilde{c}_t$, so the estimation error is $\hat{\tilde{y}}_t - \tilde{y}_t = -\tilde{c}_t$. Assume further that $\mathrm{std}(\tilde{c}_t) = \beta\,\mathrm{std}(\tilde{y}_t)$ and $\mathrm{cor}(\tilde{c}_t, \tilde{y}_t) = \rho$. Because the error $-\tilde{c}_t$ is mean zero, the root mean squared error between the estimated and true cycles equals $\mathrm{std}(\tilde{c}_t) = \beta\,\mathrm{std}(\tilde{y}_t)$, and the correlation between $\hat{\tilde{y}}_t = \tilde{y}_t - \tilde{c}_t$ and $\tilde{y}_t$ is $\frac{1-\rho\beta}{\sqrt{1+\beta^2-2\rho\beta}}$.

The case $\beta = 0$ recovers the baseline model, in which no transitory consumption is present. Even when β is sizable, for instance, $\beta = 0.67$, the correlation remains above 0.73 for any value of ρ.[7] This matches the HP filter's correlation of 0.73 in the smooth-trend benchmark (DGP 1 of the previous subsection) and, more tellingly, exceeds that of every univariate filter once trend shocks are large or trend growth contains breaks, and of the BN–VECM once breaks are present, the very environments in which a transitory consumption component is itself most likely. Note that this comparison is deliberately conservative: it pits the consumption method with a transitory component against the alternatives in their best case. Even so, the consumption-based approach continues to provide a reliable measure of GDP trends and cycles. The next subsection presents additional Monte Carlo evidence on this robustness.

Time-Varying Steady-State Savings

We show that several realistic departures from the strict permanent income hypothesis make the steady-state savings ratio $\bar{c}$ time-varying, and we adopt a rolling-window estimator that accommodates them without having to date breaks. Under the standard permanent income hypothesis, a cointegration coefficient of one implies that the log (C/Y) ratio is stationary in the long run. In the baseline framework, $\bar{c}$ corresponds to the sample mean of the log C/Y ratio, and $-\bar{c}$ represents the steady-state level of savings (Morley, 2007). When $\bar{c}$ is constant, the steady-state level of savings remains unchanged, implying a stable long-run C/Y ratio.

However, the cointegration relationship between GDP and consumption, and hence the steady-state level of savings, could instead be time-varying, implying that $\bar{c}$ is not constant. For example, Campbell and Mankiw (1989) shows that, once the real interest rate is allowed to affect consumption growth, the permanent income hypothesis implies

[7] When restricting $\rho = 0$ (the two transitory components are uncorrelated), the correlation exceeds 0.73 for $\beta \leq 0.94$.

the following dynamics for the log C/Y ratio:

$$c_t - y_t = E_t \sum_{j=1}^{\infty} \rho^j (\Delta y_{t+j} - \sigma r_{t+j}) - \rho\mu/(1-\rho)$$

where r_t is the real interest rate, ρ is the discount factor, and σ measures the effect of the real-interest-rate changes on consumption growth. When $\sigma = 0$, the expression collapses to the standard permanent income hypothesis result, the left-hand side (log C/Y ratio) is stationary, and the right-hand side captures the classic idea of "saving for a rainy day". If income is difference stationary, then the left-hand side is then also stationary and the steady-state level of savings is unchanged.

However, as argued by Campbell and Mankiw (1989) and Summers (1981), a higher real interest rate reduces the market value of wealth and therefore lowers consumption even when future income is unchanged, that is, $\sigma \neq 0$. In this case, if the real interest rate is $I(1)$, then the left-hand side, log C/Y ratio, is also $I(1)$. To see this, suppose for simplicity that r_t is a driftless random walk; then the corresponding term on the right-hand side satisfies $E_t \sum_{j=1}^{\infty} \rho^j (\sigma r_{t+j}) = \frac{\sigma}{1-\rho} r_t$. A substantial literature documents that the real interest rate is non-stationary (Del Negro et al., 2017; Hamilton et al., 2016; Morley et al., 2024), which implies that the steady-state level of savings is time-varying. Re-arrange the equation above:

$$y_t = \rho\mu/(1-\rho) + \frac{\sigma}{1-\rho} r_t + c_t - E_t \sum_{j=1}^{\infty} \rho^j (\Delta y_{t+j})$$

This is useful because it suggests that the real interest rate can be used as an additional variable when testing the cointegration relationship implied by permanent income hypothesis. Empirically, we confirm this mechanism for the United States and Mexico: using the Johansen test, we find a cointegration relationship between the log C/Y ratio $(c_t - y_t)$ and the real interest rate at annual frequency (see Appendix).

Combining the previous expression with the decomposition in Equation (1):

$$c_t = (-\rho\mu/(1-\rho) - \frac{\sigma}{1-\rho} r_t) + \tau_t + (\tilde{y}_t + E_t \sum_{j=1}^{\infty} \rho^j (\Delta y_{t+j}))$$

so that $\bar{c}$ in Equation (4) inherits a time-varying term in r_t and should be treated as time-varying.

Beyond the omission of the real interest rate, several other deviations from the standard permanent income hypothesis can also generate a time-varying log C/Y ratio, for example, GDP being $I(2)$, omitted risk (Hahm & Steigerwald, 1999), omitted wealth, and other missing variables. To address such concerns, Crucini and Shintani (2015) propose augmenting the permanent income hypothesis framework with an additional linear trend to capture the omitted variables or other departures from the benchmark model. A deterministic linear trend, however, implicitly assumes that the C/Y ratio evolves monotonically, which is unlikely in practice; their alternative of kinked trends requires identifying structural breakpoints, which is especially difficult for emerging market economies with short samples. To avoid this limitation, we adapt Kamber et al. (2018) and use a rolling-window approach that allows $\bar{c}$ to vary over time in Equation (4).

$$\bar{c}_t = \frac{1}{39} \sum_{i=-19}^{19} (y_{t+i} - c_{t+i}). \tag{7}$$

The symmetric rolling window nests the deterministic-trend case without imposing it: when the underlying C/Y ratio does contain a linear trend, the estimated $\bar{c}_t$ closely approximates that trend, but when it does not, the window does not force one. This is why it improves on the fixed linear or kinked trends of Crucini and Shintani (2015) rather than merely reproducing them. It also mitigates problems that arise when $\gamma \neq 1$. **campbell1987trend<empty citation>** and Attanasio (1999) show that consumption is proportional to lifetime income, with the proportionality factor determined by the subjective discount rate and the real interest rate. If the true γ differs from unity but is imposed to be one, the log C/Y ratio will mechanically exhibit a deterministic trend.[8]

Outliers: COVID-19

A further empirical concern is the presence of outliers in the data, which can distort the trend-cycle decomposition. We focus on outliers in consumption data that arise when household behavior is constrained by extraordinary events such as natural disasters, wars, or pandemics. A prominent recent example is the COVID-19 shock. During the COVID-19 pandemic, strict lockdowns produced an unprecedented collapse in consumption, followed by a rapid rebound once restrictions were lifted.

These dynamics were driven primarily by temporary constraints rather than permanent income shocks, and therefore introduce a large transitory component in consumption that is inconsistent with the permanent income hypothesis framework. To fix ideas, one can think of consumption growth during the lockdown as

$$\Delta c_t = \mu + \theta D_t + \epsilon_t,$$

where $D_t = 1$ in the lockdown quarters, $D_t = -1$ in the rebound that follows, and $D_t = 0$ otherwise, with $\theta < 0$ capturing the negative effect of the lockdown on consumption. This representation is only illustrative: it makes explicit that the lockdown movements are large but transitory and offsetting. We do not implement it directly, because the timing of lockdowns differs across countries and several countries experienced more than one lockdown, so no single dated dummy applies across the panel. Instead, we smooth consumption growth over a common, longer window: we average consumption growth over 2020Q2–2021Q3 and assign the same quarterly growth rate to each of these quarters.[9]

Treating the lockdown shock as transitory is consistent with the logic of the Beveridge–Nelson decomposition, which defines the trend as the long-horizon conditional expectation of the series. During the lockdown, households expect the impact on consumption to dissipate once restrictions are lifted, while their long-run income expectations are

[8] It can be shown as,

$$c_t - y_t = (c_{t-1} - y_{t-1}) + (\Delta c_t - \Delta y_t) = (c_{t-1} - y_{t-1}) + (\gamma - 1)\mu + (\gamma - 1)e_{\tau,t} - (\Delta \tilde{y}_t)$$

[9] For China, we average window over 2020Q1–2020Q2.

largely unaffected. Consumption movements over this period therefore reflect mainly transitory shocks and carry little information about the permanent component of output.

If left unadjusted, such outliers can bias the estimated cycle by overstating short-run volatility and misclassifying temporary disruptions as permanent shocks, particularly when the magnitude of the transitory shock (θD_t) is much larger than that of the permanent shock (ϵ_t). The averaging adjustment removes this contamination: it requires only that the transitory deviations average out over the window, so that average growth over the window reflects the trend, leaving the estimated trends and cycles for these quarters undistorted by the unusually large transitory shocks. For subsequent periods, consumption no longer exhibits outlier behavior and primarily reflects permanent shocks; from the perspective of the Beveridge–Nelson decomposition, the transitory lockdown shock has dissipated.

3.3 Does the Generalized Consumption-Based Method Still Work? Monte Carlo Evidence

We now return to Monte Carlo simulations to assess the robustness of the consumption-based approach once the baseline permanent income hypothesis assumptions are relaxed. We deliberately introduce two departures: (i) an additional transitory component in consumption, and (ii) time variation in the steady-state level of savings.

In DGP 4, we add a transitory component to consumption, modeled as an AR(2) process calibrated to Morley (2007), with the volatility of the transitory consumption shock set to one-half of GDP's volatility. Appendix considers an alternative specification, using the same volatility. Unsurprisingly, the consumption-based method performs better when the transitory consumption is relatively small.

DGP 5 considers changing steady-state level of savings. Motivated by the discussion of missing variable above, we allow $\bar{c}_t$ in Equation (4) to vary over time, modeling it as a random walk with drift. The drift term governs the average quarterly change in the steady-state level of savings and can be interpreted as a deterministic shift in the C/Y ratio, as in Crucini and Shintani (2015); we additionally allow for a one-time shock to $\bar{c}_t$. When the long-run C/Y ratio is unstable, applying the consumption-based method under a constant-C/Y assumption yields biased estimates. We show that a rolling-window implementation mitigates this problem.

Table 2: **Monte Carlo Simulation: Permanent Income Hypothesis Does Not Hold**

	DGP 4		DGP 5	
	Correlation	RMSE	Correlation	RMSE
Consumption-based Method	0.961	0.754	0.628	3.373
C-based with Rolling mean	0.764	1.736	0.715	1.871
BN-VECM	0.908	1.304	0.725	2.308

Notes:
1. Baseline settings, seeds, and iteration choices are the same as in Table 1.
2. In DGP 4, we allow for a transitory consumption. $\sigma_c = \frac{1}{2}\sigma_y = 0.345$.
3. In DGP 5, we allow $\bar{c}_t$ to be time-varying, specifically, $\bar{c}_t = 0.03 + D_t + \bar{c}_{t-1} + v_t$, with $\sigma_{e,v} = 0.05$ and D_t being a one-time shock.

Table 2 reports only the bivariate methods and omits the univariate filters from Table 1. The reason is that the univariate filters, the HP filter, the boosted HP filter, and the Hamilton filter, take only y_t as input, so their performance depends solely on the DGP for y_t. Because the experiments in this subsection alter only the consumption process and leave the y_t process unchanged, these filters would reproduce their Table 1 numbers exactly, and reprinting them would add nothing. This same property makes the comparison across tables especially clean: since the univariate filters depend only on y_t, each entry in Table 2 can be compared directly against all three DGP columns of Table 1 for those methods, with no ambiguity about which column is the relevant benchmark. The BN-VECM is the exception, and is therefore reprinted here: it uses the joint dynamics of $\{y_t, c_t\}$ rather than y_t alone, so changing the consumption process does change its results.

In both DGP 4 and DGP 5, the consumption-based method no longer recovers the cycle perfectly, since these DGPs are deliberately constructed to violate the permanent income hypothesis assumptions. Its performance nonetheless remains strong and, in every case, exceeds that of the HP filter benchmark reported in Table 1. The advantage is particularly pronounced when GDP is subject to large trend shocks: in that environment (DGP 2 of Table 1) the HP-filter correlation with the true cycle falls to 0.56, whereas the modified consumption-based method with a rolling-window mean maintains correlations above 0.7 even when the permanent income hypothesis does not strictly hold.

The Beveridge–Nelson decomposition based on a bivariate VECM performs well in both DGPs, which is expected because both departures can be captured by a VECM. For DGP 4, consumption growth is regressed on its own lag, which implicitly allows for a transitory component in consumption; for DGP 5, the lagged cointegration errors in the VECM help capture deviations in the steady-state level of savings, so its performance is slightly better than the rolling-window specification here. However, when additional breaks in trend growth are present, the BN–VECM performs worse, as illustrated in DGP 3.

Finally, time variation in the steady-state level of savings may also arise from structural breaks in $\bar{c}$ in Equation (4), or from cointegration between consumption and income with a coefficient different from unity. Appendix considers these additional DGPs, along with hybrid cases that combine multiple departures from the baseline, such as allowing for transitory consumption together with varying steady states. The conclusions are unchanged.

4 Cross-Country Results and Evaluation

4.1 Data

We collect quarterly, seasonally adjusted data on real GDP and real private consumption from the OECD database, accessed via CEIC. The sample includes both advanced economies and emerging market economies, classified according to the IMF's World Economic Outlook groups, with one exception: we assign all Eastern European economies to the emerging market group, because, with data beginning only in the 1990s, these transition economies are better grouped with emerging markets than with mature advanced economies over our sample.[10] The dataset extends through 2023Q4. For emerging market economies, sample start dates vary from the early 1990s to the early 2000s, depending on data availability. For advanced economies, observations prior to 1980 are excluded to maintain comparability across countries. For China, where expenditure-side GDP is available only at an annual frequency, we rely on the quarterly real GDP and consumption estimates constructed by Chang et al. (2016), spanning 1992Q1–2022Q4.[11] A summary of the sample span and country classification is provided in the Appendix.

Although the permanent income hypothesis is fundamentally about nondurable consumption, such disaggregated data are unavailable for most countries in our sample. Following Aguiar and Gopinath (2007), we therefore use aggregate real private consumption and real GDP as proxies for consumption and income. We use aggregate rather than per-capita series, consistent with the concepts of business cycles and output gaps. Using per-capita data is innocuous for describing cycles because population enters only the long-run trend: subtracting log-population from both sides of Equation (6) shifts the estimated trend but leaves the cyclical component unchanged.

4.2 Results: Ideal Cases using Baseline Model

In this subsection, we ask where the strict permanent income hypothesis (i) consumption follows a random walk and (ii) consumption and income cointegrate with a unit coefficient, holds closely enough for the baseline model to apply directly. Using the pre-COVID cross-country panel data, we test the random walk hypothesis following Hall (1978), examine the long-run consumption-income relationship using both the Engle–Granger and Johansen procedures, and estimate the long-run coefficient for each country using Dynamic OLS (DOLS). Details of these tests are provided in the Appendix.

The strict conditions are met almost nowhere: Japan is the only country in our sample that satisfies both the random walk and the unit-coefficient cointegration assumptions. Figure 3a shows the baseline consumption-based cycle for Japan alongside the HP-filtered cycle. The two series are closely aligned, especially before 2010: the correlation is 0.52 and 70% of observations share the same sign. We also report Colombia, a near-miss, in Figure 3b, where the correlation is 0.47 and 68% of observations share the same sign.[12] We report these correlations as descriptive summaries only. Because the HP filter is itself an imperfect benchmark, agreement between the two measures is not in itself evidence that either is accurate; we return to accuracy directly through the forecasting evaluation.

4.3 Results: Generalized Model

Having shown through simulation that the modified consumption-based method remains informative under departures from the random walk and cointegration assump-

[10] IMF classification available at: https://www.imf.org/en/Publications/WEO/weo-database/2023/April/groups-and-aggregates.

[11] The data is updated annually by the Federal Reserve Bank of Atlanta and can be accessed at: https://www.atlantafed.org/cqer/research/chinamacroeconomy.

[12] Colombia violates both the random walk hypothesis and $\gamma = 1$, but in each case the rejection is only marginal.

Figure 3: **Cycles by Baseline Method before COVID-19**

(a) **Japan**

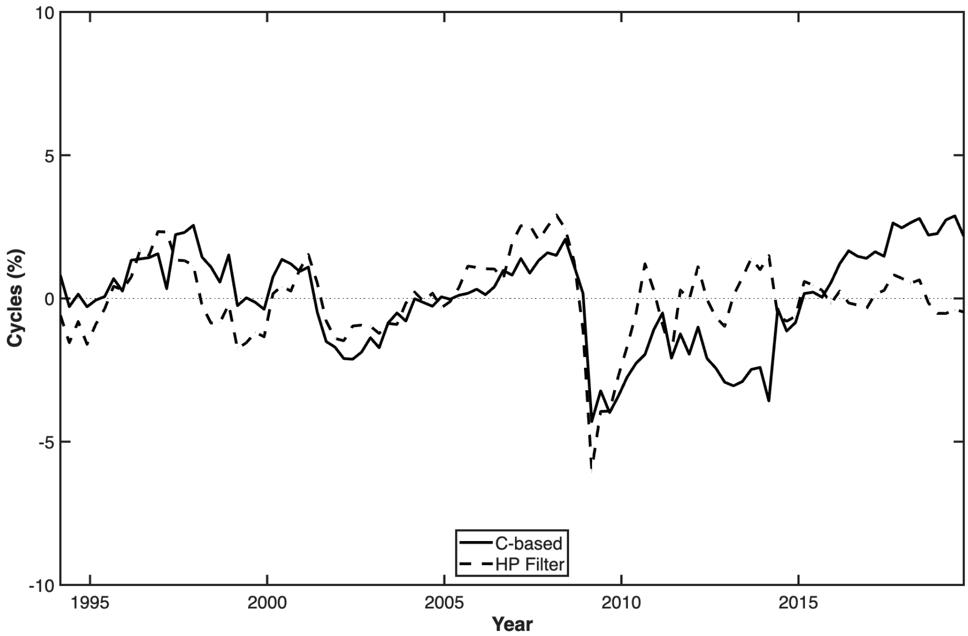

(b) **Colombia**

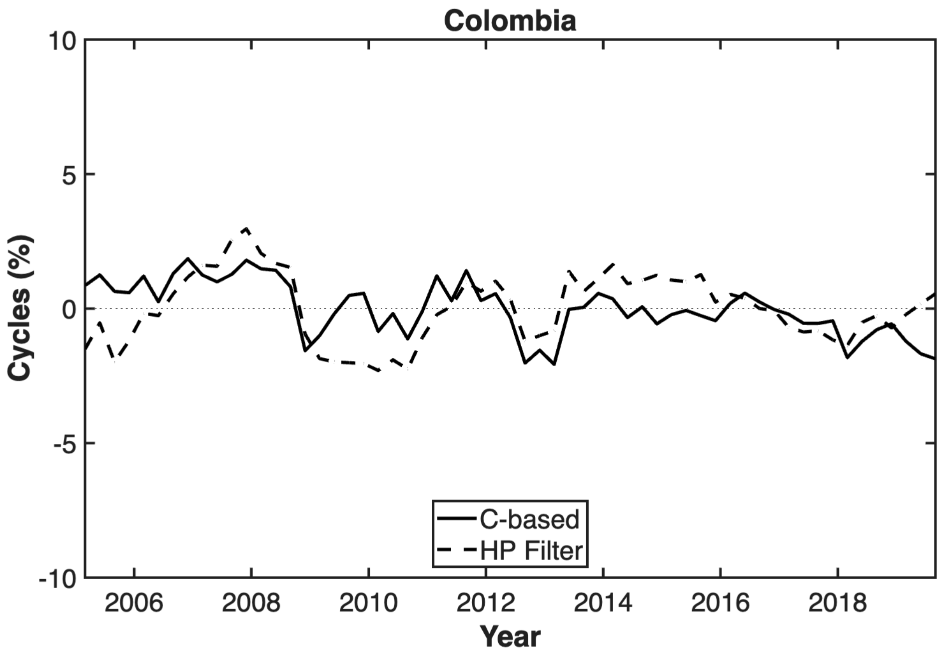

tions, and having illustrated the baseline cycle for representative countries, we now apply the generalized method to a broader set of countries over the full sample, including the COVID-19 period, using the outlier adjustment described above. We begin with the United States as a benchmark and then turn to other advanced and emerging market economies. For the United States, the consumption-based method produces trend and cycle estimates closely aligned with the standard HP filter, which reassures us that, for economies with relatively smaller trend shocks, the method does not depart substantially from established practice. We then examine how the alignment between the two approaches varies with economic structure.

Figures 4 and 5 plot the estimated cycles. For illustration, we take the United States and shade NBER-dated recessions. The cointegration coefficient γ is estimated at 0.92, statistically different from the unit coefficient implied by the standard permanent income hypothesis, so, following the Monte Carlo evidence, we use a rolling-window procedure to estimate a time-varying $\bar{c}$. The resulting consumption-based and HP-filtered cycles display broadly similar dynamics: both show pronounced declines in each of the six recession episodes, and their full sample correlation is high (0.68).

We next plot the cycles for the United Kingdom, India, and Mexico. All show a pronounced decline during the COVID-19 period, when most of the volatility in output appears to reflect transitory rather than permanent shocks. This is precisely the configuration in which our Monte Carlo evidence indicates the HP filter recovers the cyclical component well, so over the COVID window the HP filter is a credible reference. The COVID-adjusted consumption-based cycle tracks it closely over the same window, which indicates that the outlier adjustment is behaving as intended.[13] More generally, the two methods produce similar cyclical patterns with largely consistent turning points, though the consumption-based estimates display greater short-run variation and the HP filter produces smoother cycles.

To assess how closely the consumption-based method aligns with the HP filter across advanced economies and emerging market economies, Figure 6 reports the correlation between the two cycle measures; the Appendix reports the share of observations in which they have the same cyclical sign. Most advanced economies (the filled bars) show both higher correlations and greater sign consistency, while the pattern is largely reversed for emerging market economies. This is consistent with our Monte Carlo evidence and with Aguiar and Gopinath (2007): in advanced economies, the consumption-based and HP cycles are broadly aligned, whereas in emerging market economies they diverge more often, owing to the prevalence of large permanent shocks. The correlation gap alone, however, cannot say which measure is the more accurate; divergence is equally consistent with HP mismeasuring the EME cycle or with the consumption method doing so. We resolve this directly in the next subsection, where forecasting performance provides an external criterion for accuracy.

4.4 Pseudo-out-of-sample Forecast

Furlanetto et al. (2023) argue that a useful criterion for assessing trend cycle estimates is whether they improve forecasting performance. This follows directly from the BN decomposition: as Nelson (2008) notes, when output lies below trend, subsequent

[13]For India, an additional sharp decline and rapid rebound in the late 1990s suggests a further episode that may warrant outlier adjustment.

Figure 4: **Comparison of Cycles in Advanced Economies**

(a) **United States**

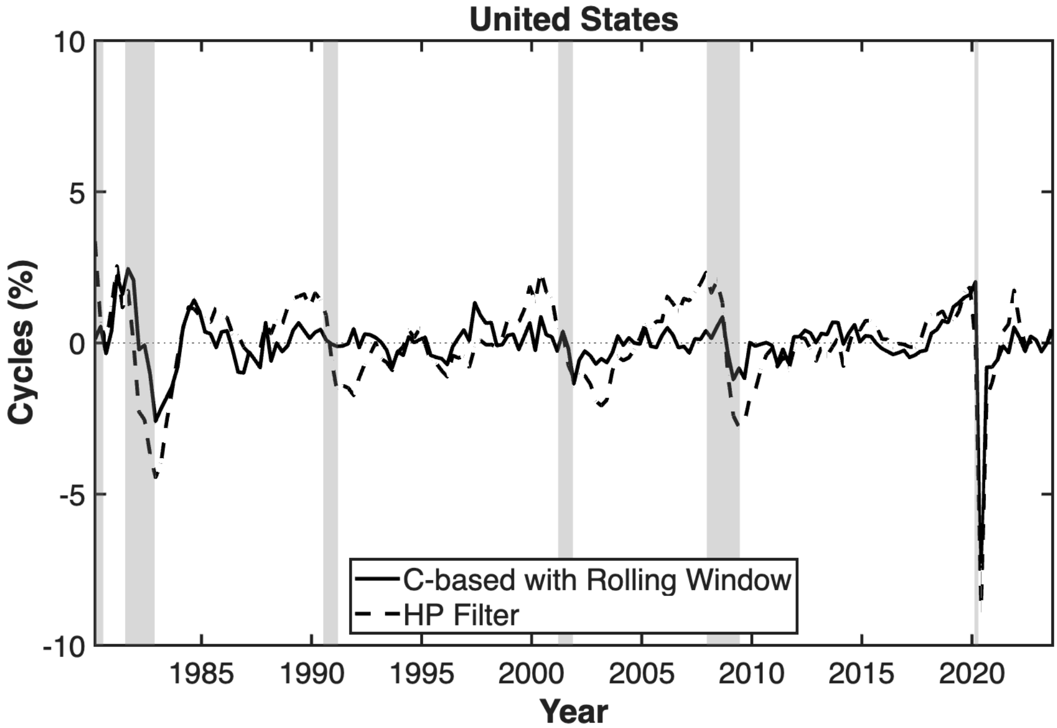

(b) **United Kingdom**

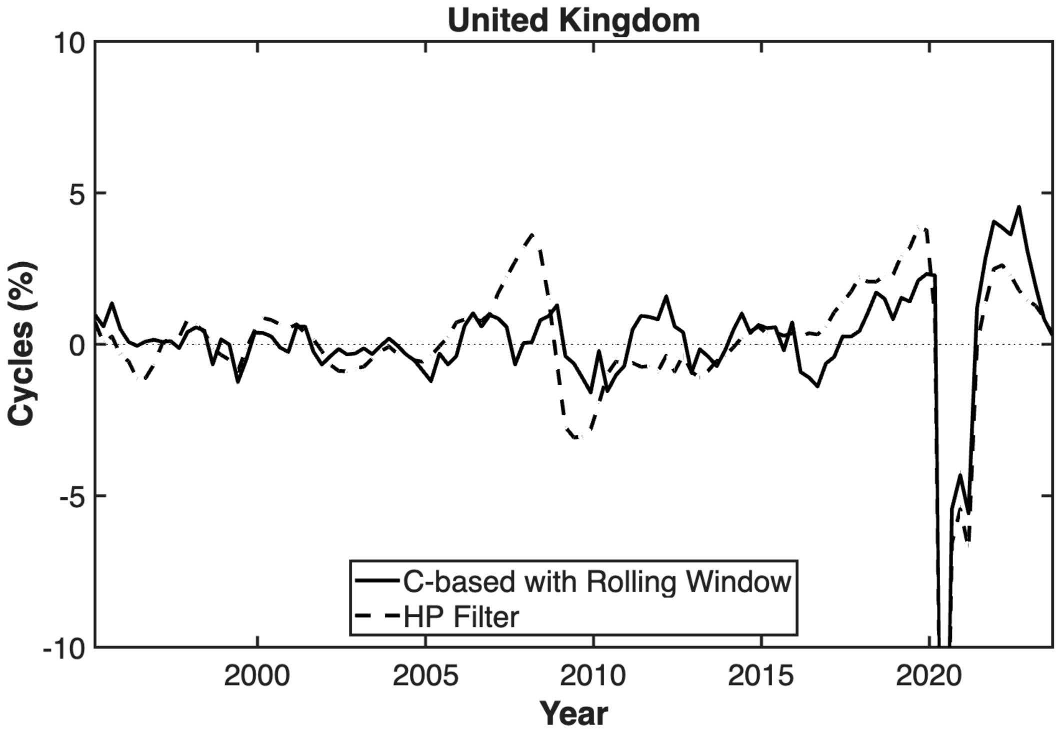

Figure 5: **Comparison of Cycles in Emerging Market Economies**

(a) **India**

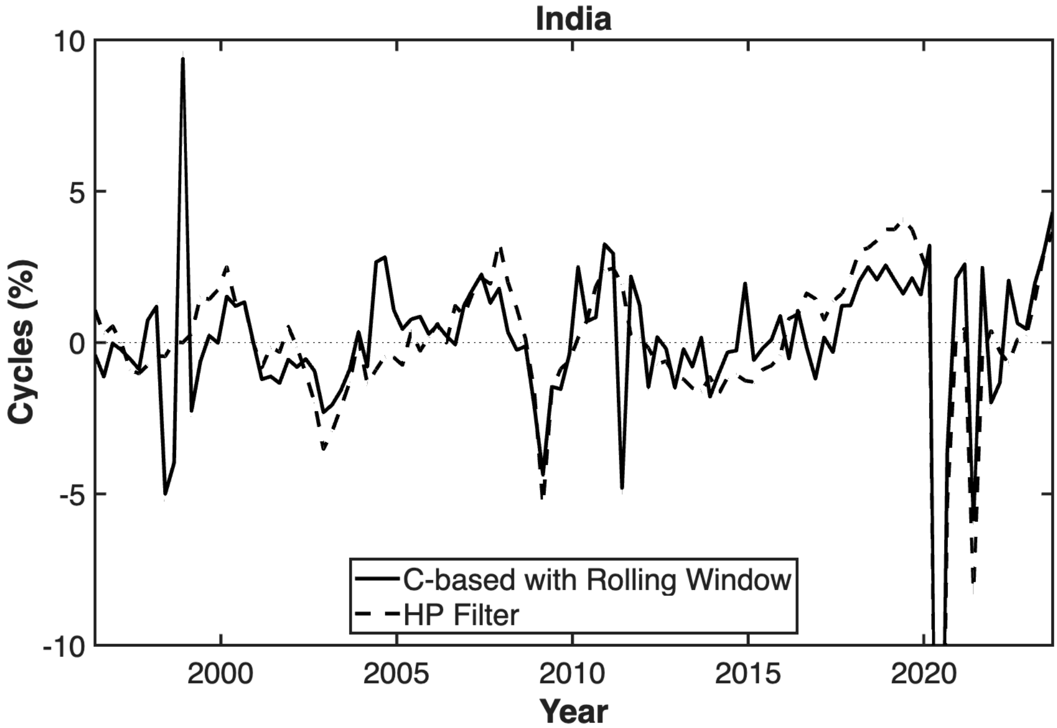

(b) **Mexico**

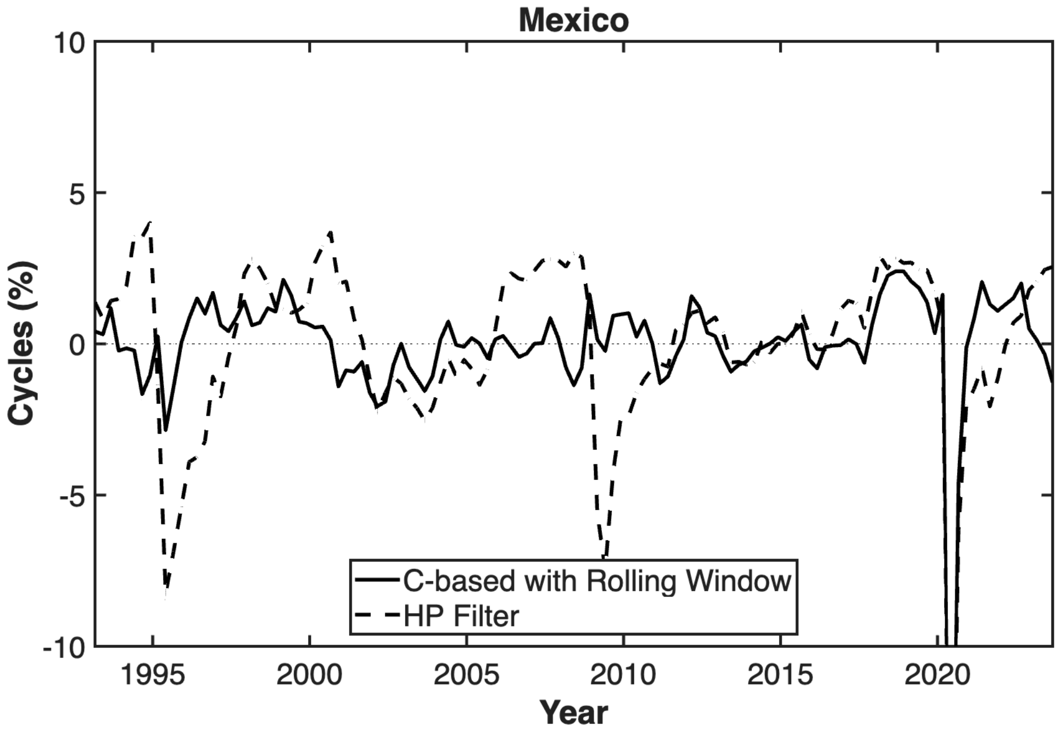

Figure 6: **Correlation between C-based and HP filtered Cycles**

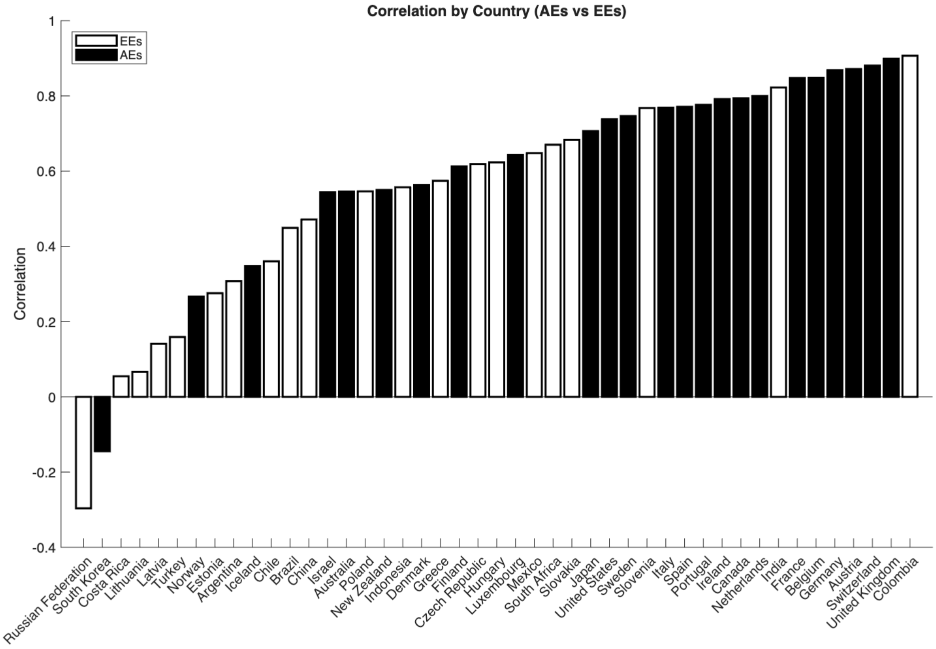

growth should be higher, because transitory shocks eventually dissipate and output reverts toward trend. Following Nelson (2008) and Kamber et al. (2018), we evaluate forecasting performance using the regression

$$\Delta y_{t+h} = \alpha + \beta \tilde{y}_t + \epsilon_{t+h}, \tag{8}$$

where Δy_{t+h} denotes h-quarter-ahead GDP growth and $\tilde{y}_t$ is the cycles (or output gap).

We conduct the forecast evaluation on pre-COVID data, with the pseudo-out-of-sample period starting in 2012Q1 for all countries. Because sample availability differs across countries, the in-sample and out-of-sample lengths vary across the panel. This is a deliberate separation from the cycle estimation above: the cycles are estimated on the full sample with the COVID adjustment, whereas the forecast comparison excludes the COVID period entirely, so that the smoothed pandemic quarters play no role in the forecasting results.

We take the consumption-based method with a symmetric rolling window as the benchmark and report relative root mean squared errors (RRMSE). By construction, an RRMSE value above one means a given measure has a larger forecast RMSE than the benchmark, that is, values above one favor the consumption-based method. We also report Diebold–Mariano confidence bands following Diebold and Mariano (1995).[14]

Figure 7 reports the out-of-sample results for the United States, the United Kingdom, India, and Mexico, comparing four output-gap estimates: the consumption-based method with a constant mean, the consumption-based method with a symmetric rolling window, the HP filter, and the Hamilton filter. Across these countries, the consumption-based approach delivers substantially lower forecast errors than both the

[14] The bands show the confidence interval for the mean loss differential in the Diebold–Mariano test.

Figure 7: **OOS Forecast**

(a) **US**

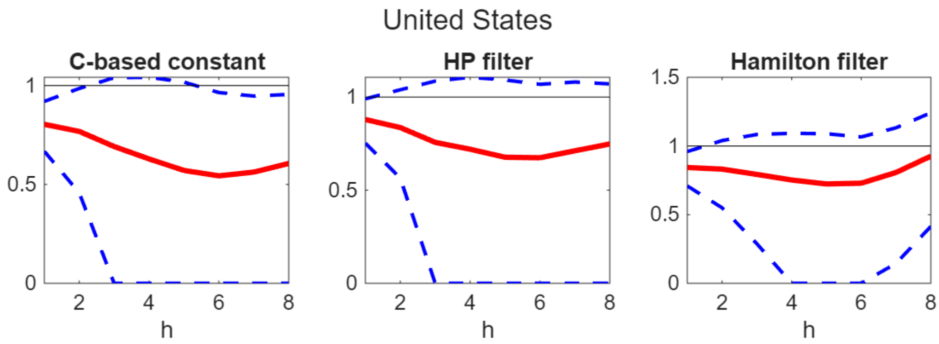

(b) **UK**

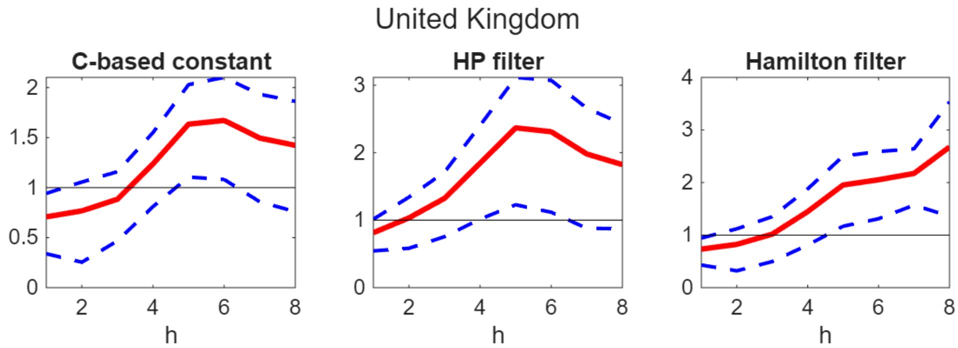

(c) **India**

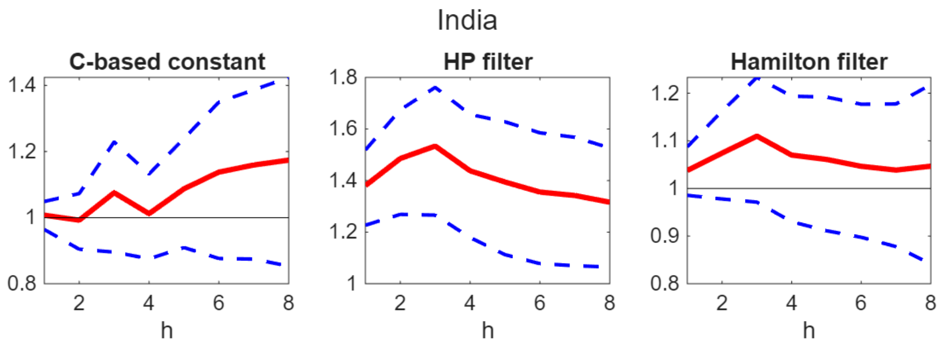

(d) **Mexico**

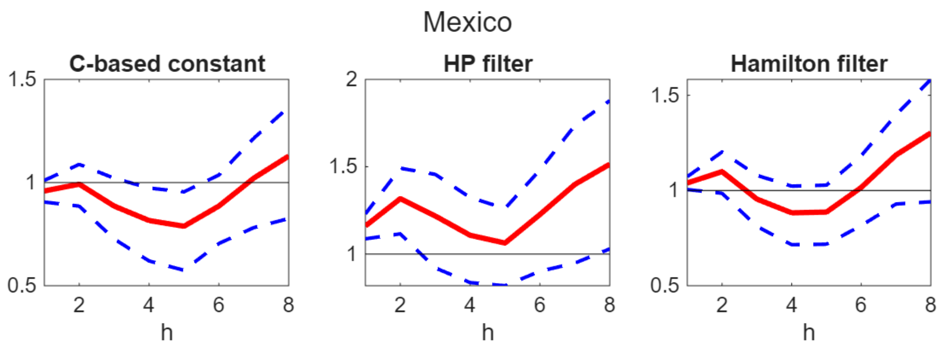

Notes: The graphs plot out-of-sample RRMSEs compared to forecasts based on the consumption-based estimated output gap with symmetric rolling window. Out-of-sample evaluation begins in 2012Q1. The bands depict 90% confidence intervals from a two-sided Diebold and Mariano (1995) test of equal forecast accuracy.

Figure 8: **RRMSE: HP Filter relative to C-based**

(a) **Emerging Market Economies**

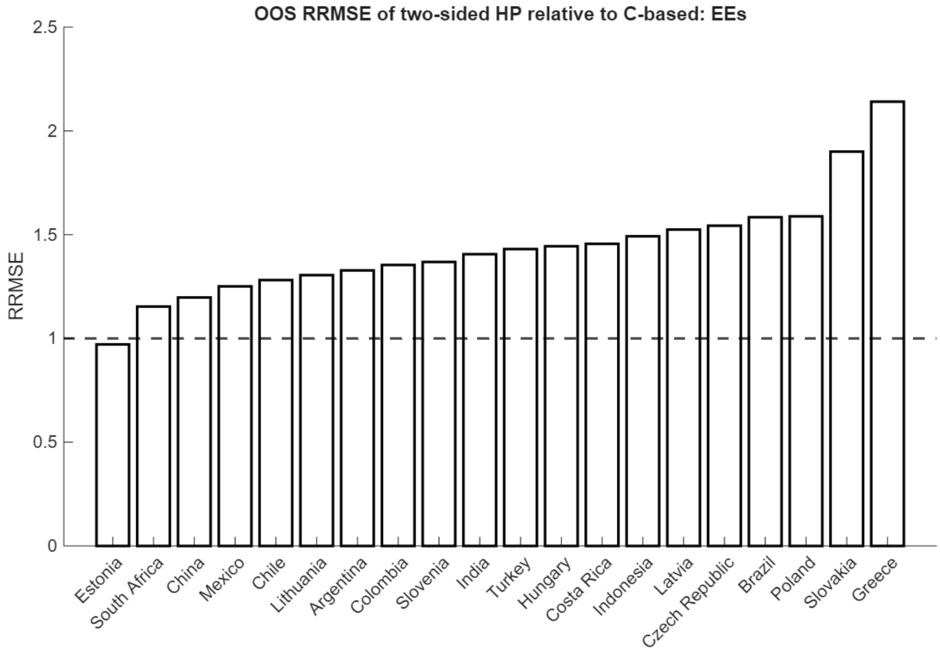

(b) **Advanced Economies**

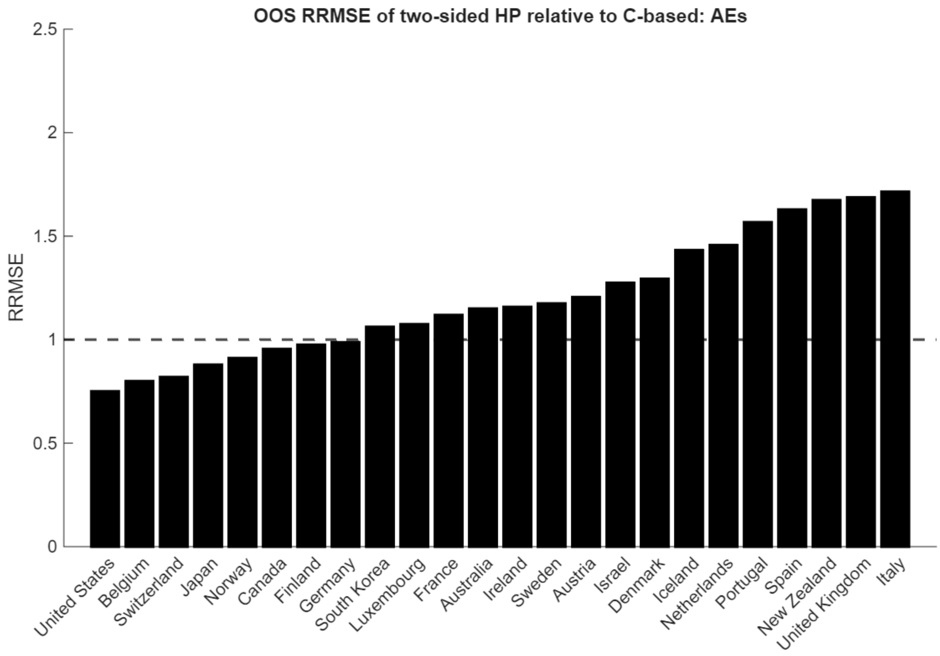

HP and Hamilton filters. For the United States, the consumption-based method with a constant mean is preferred (it attains the lowest forecast RMSE).

Beyond the four-country comparisons, Figure 8 summarizes the full sample by plotting, for each country, the RRMSE of the HP filter relative to the consumption-based benchmark, averaged across horizons $h = 1, \dots, 8$; values above one again favor the consumption-based method. The forecasting gains are larger in emerging market economies: for nearly all of them (Estonia is the exception), the consumption-based method outperforms the HP filter by a wide margin. Among advanced economies, it also generally improves accuracy, though the average gains are more modest.

5 Case Study: China's Trends and Cycles

This section uses China as a case study to assess the applicability of the consumption-based method and to examine whether the estimated patterns provide a plausible characterization of the economy. China is a prominent emerging market economy that has experienced substantial trend shocks and notable growth slowdowns. As the cross-country evidence of Section 1 would lead one to expect, its output fluctuations are heavily concentrated at low frequencies, roughly 83% of the variance of output growth, and its quarterly GDP growth fell markedly after 2013Q3, with the annualized rate declining from about 9% to 5.9% before COVID.

These features make China a natural environment for evaluating the consumption-based approach. We begin by presenting the consumption-based estimates alongside those from alternative decomposition methods, highlighting the limitations of standard filters. We then relate the consumption-based cycles to major historical events to assess whether their timing and direction are plausible. Finally, we discuss the broader patterns that emerge from China's estimated trends and cycles.

5.1 Trends and Cycles Estimations

Figure 9a plots the log of actual real GDP alongside the estimated trends.[15] The difference between the two lines is the estimated cycle, shown in Figure 9b.[16] The cycle measures the percentage deviation of actual GDP from its potential level: a positive value indicates that actual GDP exceeds potential, and a negative value that it falls short.

Comparison with Filtering Methods. Figure 10 compares the consumption-based cycles with those from the HP filter and the OECD Composite Leading Indicator (CLI).[17] Both the HP-filtered cycles and the OECD CLI exhibit limited volatility, with maximum amplitudes around 2.5%, and all three methods display trough-to-trough cycles longer than eight years, well beyond the duration typically found for the United

[15] The shaded area corresponds to 2020Q1–2020Q2, for which we apply the COVID-19 consumption adjustment described in Section 3.

[16] The 95% confidence interval for the estimated cycle is derived from the estimated γ, obtained via a DOLS regression of $(y - c)$ on its lags using HAC-robust standard errors.

[17] The OECD CLI is constructed as the common component of HP-filtered macroeconomic time series and is available at https://www.oecd.org/en/data/indicators/composite-leading-indicator-cli.html.

Figure 9: **China's Trends and Cycles**

(a) **Actual GDP and Trends**

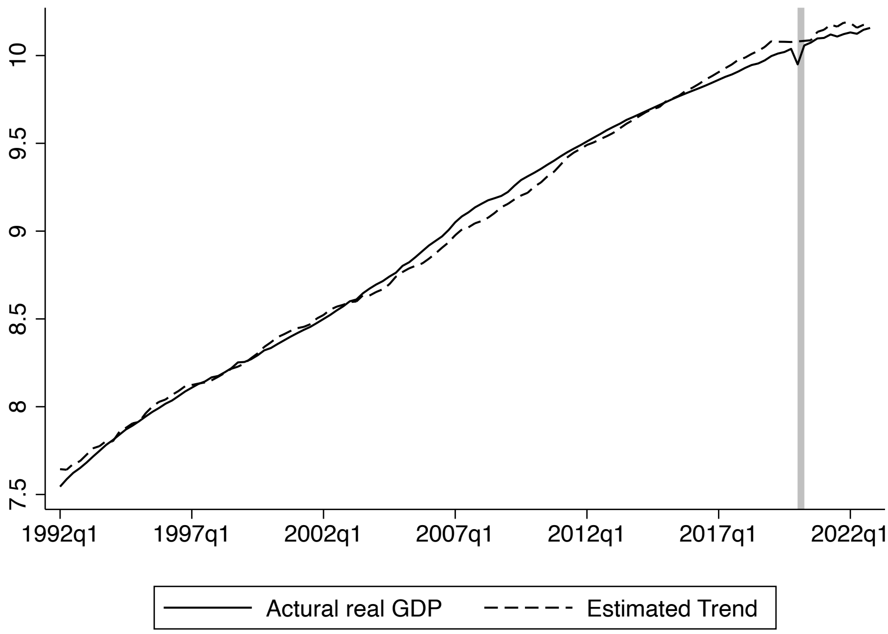

(b) **Cycles**

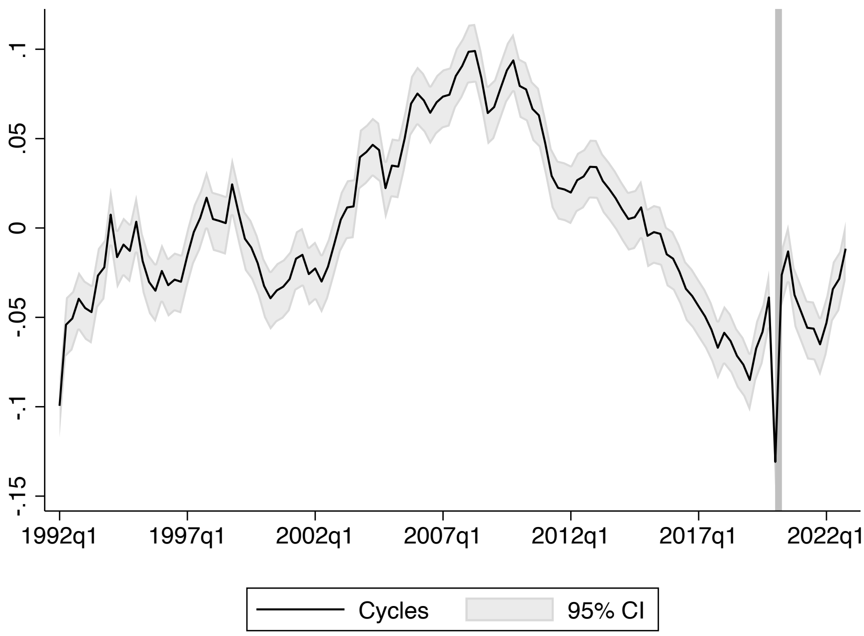

Figure 10: **Estimated Cycles and Filtering Methods**

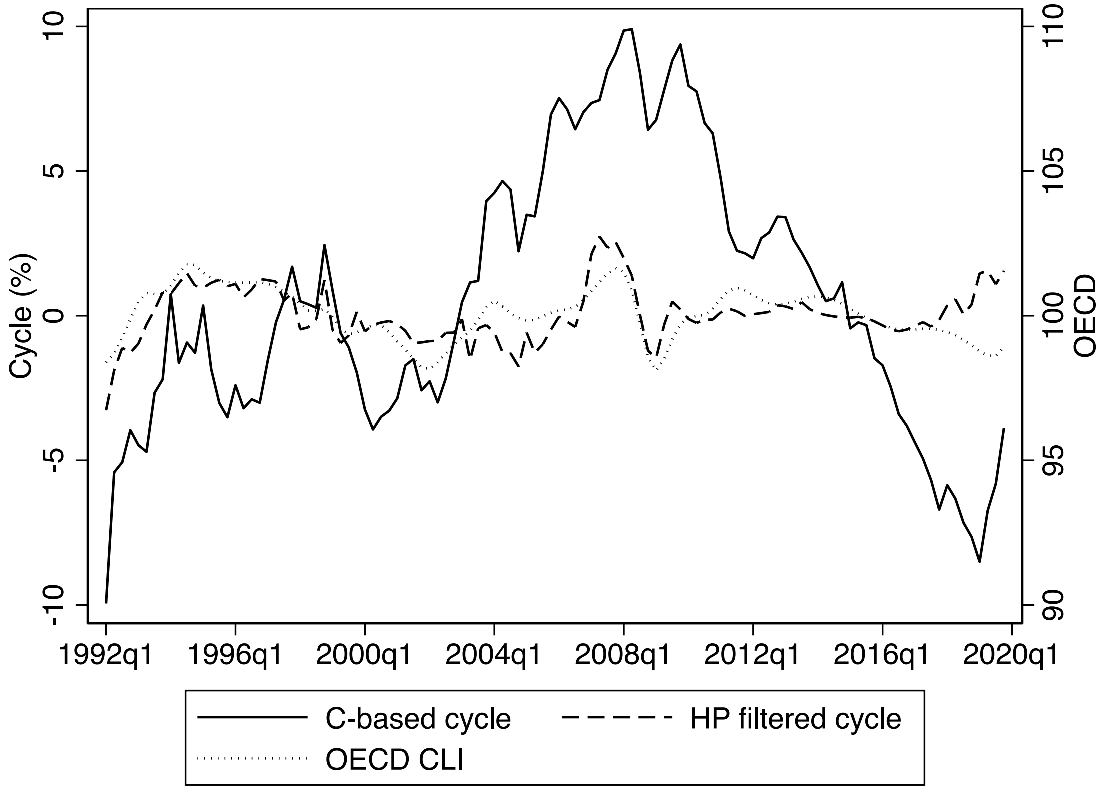

Notes: The OECD Composite Leading Indicator (CLI) measures the common component of HP-filtered cycles across multiple macroeconomic time series and is available at https://www.oecd.org/en/data/indicators/composite-leading-indicator-cli.html. The HP filter is estimated using the standard smoothing parameter of 1600. Estimates using alternative parameter values of 400 and 800 are highly correlated with the baseline, with correlation coefficients of 0.95 and 0.98, respectively.

States (Stock & Watson, 1999) and beyond the eight-year upper bound conventionally used to delimit business cycles.

Figure 11 reports the spectral densities of the estimated cycles from the three methods, and the striking feature is what they share. In all three panels, including the HP filter and the CLI, which are expressly designed to isolate fluctuations in the 6–32-quarter band, the spectral mass is concentrated at the lowest frequencies and decays as frequency rises; none displays the interior, hump-shaped peak that characterizes U.S. business cycles. In other words, even methods built to suppress low-frequency variation cannot remove it here: China's cyclical movements genuinely reside at low frequencies, consistent with the cross-country spectral evidence of Section 1. This is precisely the setting in which imposing a fixed business-cycle pass-band is least appropriate, and in which a method that does not pre-commit to one, such as the consumption-based decomposition, has the most to offer.

Table 3: **Cycles and China's Major Events**

Date	Event	Cycle Position	Cycle Direction
1989Q2	Tiananmen Square Events	$-$	$\downarrow$
1992Q1	Deng's Southern Tour	$-$	$\uparrow$
1997Q3	Asian Financial Crisis	$+$	$\downarrow$
1998Q3	100bn Fiscal Expansion	$-$	$\uparrow$
2001Q4	WTO Accession	$-$	$\uparrow$
2008Q3	Global Financial Crisis	$+$	$\downarrow$
2009–2010	4tn Fiscal Expansion	$+$	$\uparrow$
2010–2019	Declining FDI and Net Exports	$+/-$	$\downarrow$

Notes:
1. The column "Cycle Position" reports whether the economy is above (+) or below (−) potential at the time of the event, based on the estimated cycle.
2. The column "Cycle Direction" indicates whether the economy is moving toward a peak (↑) or toward a trough (↓).
3. Events are selected to illustrate key episodes in which external shocks, policy interventions, or institutional reforms align with turning points in the estimated cycle.
4. 1989Q2 predates the estimation sample (which begins 1992Q1); the reported position and direction reflect the depressed external-demand conditions inherited at the start of the sample, falling exports and FDI, not a reading of the estimated cycle at that date.

Relating Cycles to Major Events. The estimated cycles can be read against China's economic history. We stress at the outset that this is an illustrative exercise, not a test: the events in Table 3 and Figure 12 are selected ex post, after the cycle has been estimated, and are then matched to its turning points. The exercise therefore speaks to whether the estimated cycle is interpretable and narratively coherent, not to whether it is correct; accuracy is assessed separately through the forecasting evaluation of Section 4. With that caveat, the cycle lines up with major episodes in a plausible way. Figure 12 plots the cycle together with these milestones, and Table 3 lists a representative subset with their cycle positions and directions; the discussion below covers the episodes more fully than the table.

Our sample begins in 1992, so events before that date enter only as context for

Figure 11: **Comparison of Spectral Density**

(a) **C-based Cycles in China**

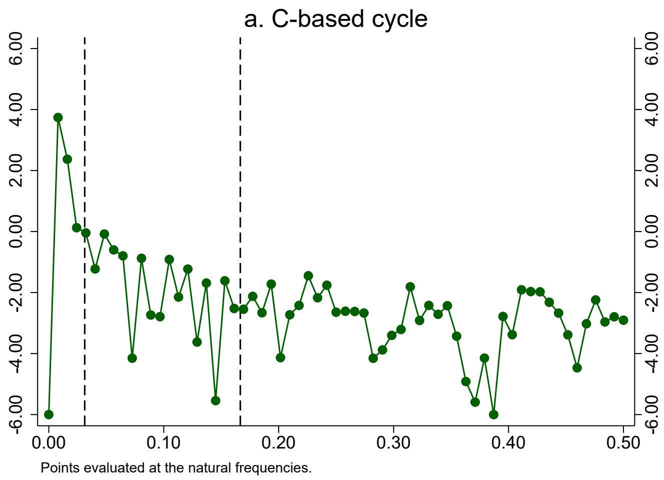

(b) **HP Cycles in China**

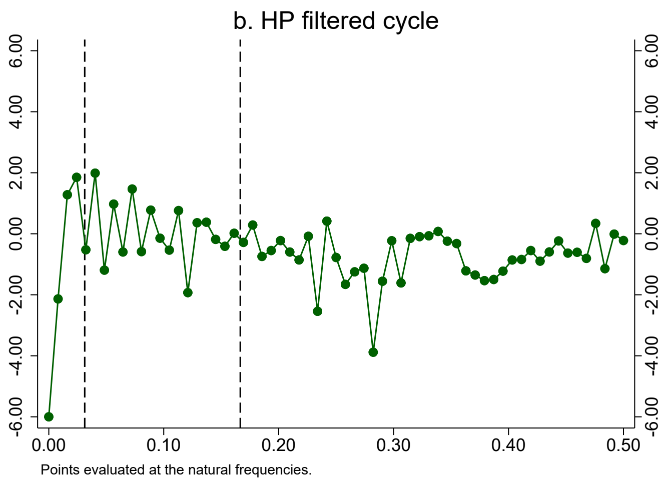

(c) **OECD CLI in China**

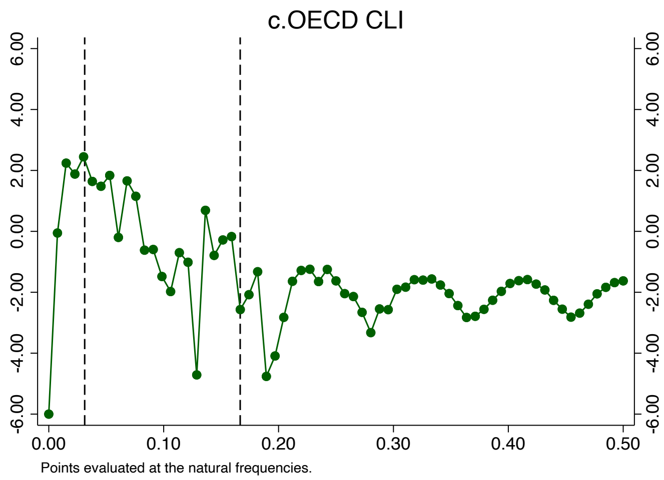

Figure 12: **Estimated Cycles and Major Events**

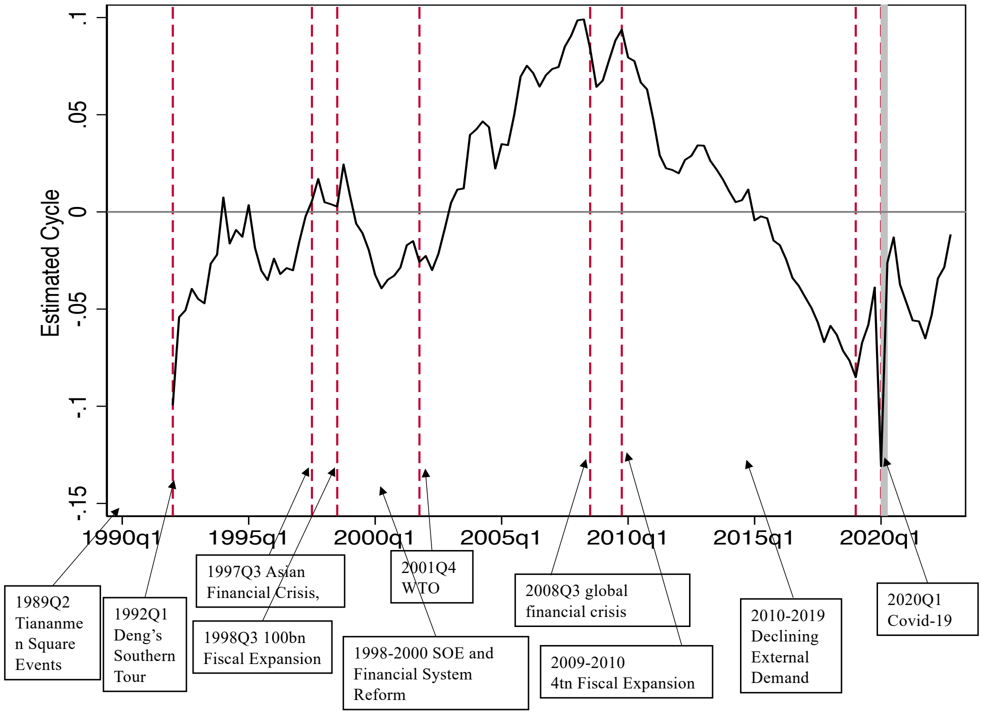

the starting conditions rather than as dated turning points in the estimated cycle.[18] The first cycle peaks early in the sample. A pivotal turning point comes in 1992Q1, when Deng Xiaoping's southern tour, followed by the 14th National Congress's goal of a "socialist market economy," relaxes bank-credit controls and triggers an investment boom, inflation, and property-market pressures. The government's "16-point" plan and accompanying monetary measures cool the economy, bringing this cycle to an end around 1995.

The next cycle is shaped by the 1997 Asian Financial Crisis, which sharply reduces external demand. The government responds with expansionary fiscal and monetary policy, including a 100-billion-yuan bond issue for infrastructure in 1998, alongside major reforms of state-owned enterprises and the financial system over 1998–2000; the economy moves from peak to trough over this period. In late 2001, WTO accession initiates a recovery, raising net exports and FDI until the 2008 Global Financial Crisis. As external demand collapses, the 4-trillion-yuan stimulus of 2009–2010 pushes the economy toward another peak, after which the subsequent decline in external demand, through the 2010s, as FDI and net export growth fade, draws it back down.

Figure 13: **Actual and Trend GDP Growth over Subperiods**

(a) **Actual GDP Growth**

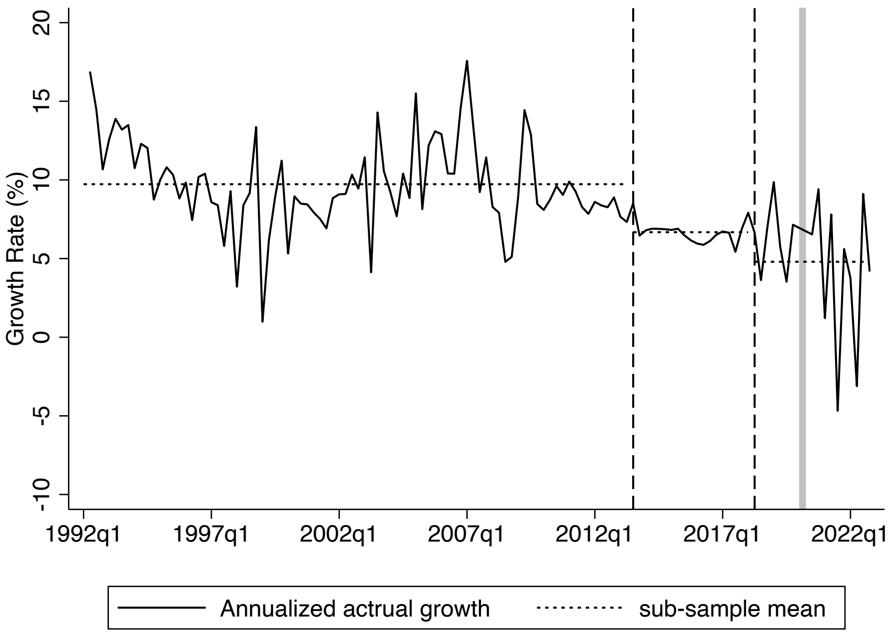

(b) **Trend Growth**

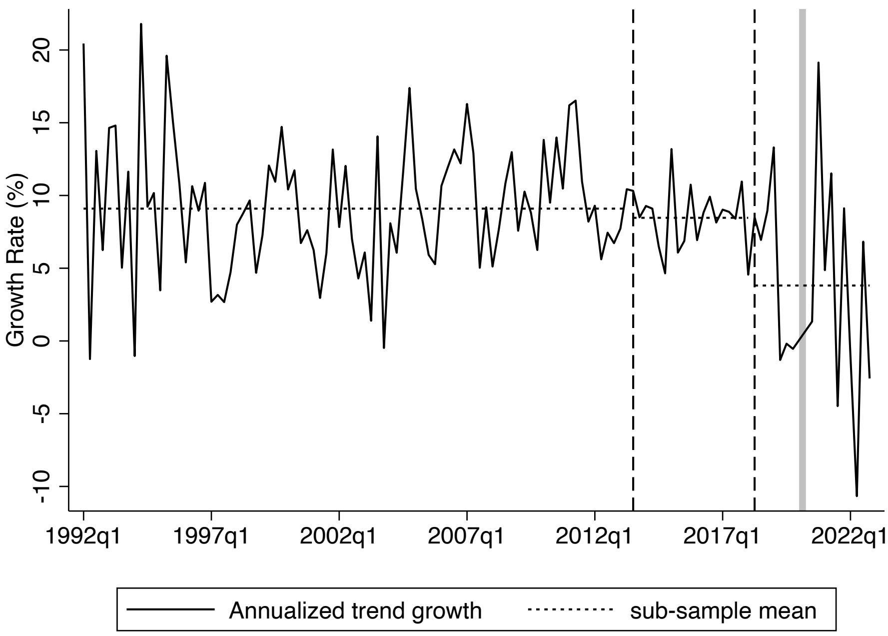

5.2 Trends and Cycles Patterns

Trends. Figures 13a and 13b plot the annualized growth rates of actual real GDP and of the estimated trend using the benchmark method. Actual GDP growth begins to slow in 2013Q3, whereas trend growth remains relatively stable until 2017Q4, when it starts to decline. These dynamics suggest three subperiods. Before 2013Q3, average actual growth rate is 9.7%, slightly above trend growth of 9.1%. Between 2013Q3 and 2018Q1, average actual growth rate falls to 6.7% while trend growth declines more modestly, to 8.5%. After 2018Q2, both fall substantially, indicating a combination of cyclical and trend slowdown. The early slowdown in actual GDP growth thus appears primarily cyclical, whereas the post-2018 period reflects both a structural slowdown in trend growth and continued cyclical weakness.[19]

Figure 14: **Estimated Cycles and Turning Points**

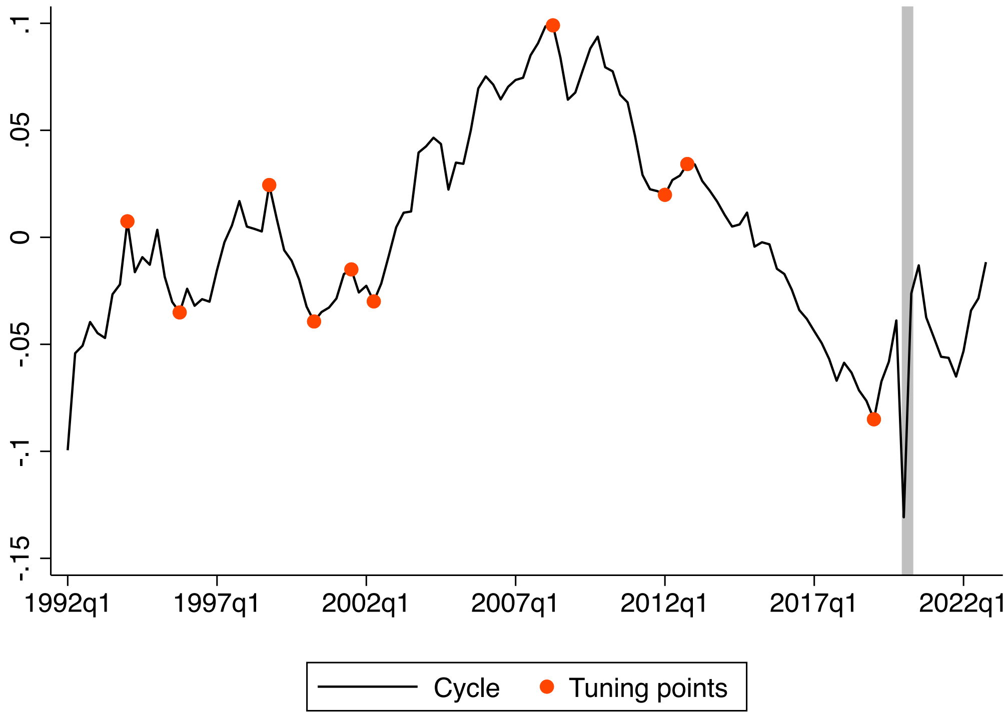

Notes: The figure plots the estimated cycles together with peaks and troughs identified using the Bry and Boschan (1971) algorithm under the 3/3/7 rule. The shaded area denotes the first two quarters of the COVID-19 pandemic.

Business Cycles. We date turning points using the Bry and Boschan (1971) algorithm.[20] Following the NBER convention, recessions are dated from a peak to the

[18]In particular, the 1989 Tiananmen Square events predate our sample; we cannot locate them on the estimated cycle and treat them only as background to the depressed external demand, falling exports and FDI, observed at the start of the sample.

[19]The shaded area again denotes 2020Q1–Q2, corresponding to the onset of the COVID-19 pandemic. These quarters are excluded from the figure due to their large scale relative to normal periods.

[20]The default algorithm applies a 2/2/5 rule that requires a minimum of two quarters in each of upswings and downswings and a minimum cycle length of five quarters. To avoid capturing minor

subsequent trough and expansions from a trough to the next peak. Over the period 1992Q1–2019Q4, we identify five peaks and five troughs, as shown in Figure 14.[21] The formal dating confirms what visual inspection of Figure 10 already suggested: China's cycles are unusually long. Their durations range from about two years to nearly seventeen, far exceeding both typical U.S. business cycles and the eight-year upper bound conventionally used to delimit them. As Canova (1994) emphasizes, the choice of decomposition method is itself decisive for the timing of turning points.

Table 4: **Correlation between Estimated Cycles and Major Macro Time Series**

Series	Correlation with the estimated cycle			
	1992Q1–1995Q4	1996Q1–2000Q2	2002Q2–2019Q1	Full Sample
Internal				
M2 (YoY growth)	0.29	-0.42	0.81	0.18
Fiscal balance (% of nominal GDP)	No data	0.39	0.80	0.67
Fixed asset investment (YoY growth)	-0.05	-0.04	0.77	0.52
Gross fixed capital formation (YoY growth)	-0.55	0.11	0.54	0.14
Domestic credit (YoY growth)	0.06	-0.06	0.05	0.03
Required reserve ratio (%)	No data	-0.04	-0.07	0.19
External				
Nonreserve liabilities (YoY growth)	No data	-0.41	0.25	0.24
Foreign direct investment (YoY growth)	No data	-0.19	0.22	0.25
Net exports (YoY growth)	0.54	-0.20	0.17	0.08
Nonreserve liabilities (log level)	No data	-0.53	-0.05	0.15
Foreign direct investment (log level)	No data	0.39	0.01	0.22
Net exports (log level)	0.03	0.33	0.45	0.37

To gauge the roles of internal and external factors, we compute correlations between the estimated cycles and a set of macroeconomic time series over three business-cycle periods (Table 4). Because these sub-periods are themselves delimited by the estimated turning points, the within-period correlations should be read as descriptive co-movements rather than as independent evidence. For external factors, log net exports and FDI growth, the co-movement strengthens noticeably after WTO accession.

Among internal factors, the fiscal balance and M2 growth co-move most strongly with the cycle, consistent with a prominent role for fiscal and monetary conditions in China's fluctuations (Chen & Zha, 2018; Chen et al., 2018); we emphasize that these are correlations and are not evidence of a causal channel, particularly as both variables respond endogenously to the cycle. The M2 correlation is 0.81 over the pre-COVID cycle period 2002Q2–2019Q1 (Table 4) and a similar 0.82 over the longer 2000Q1–2022Q4 window (Figure 15), so the co-movement is stable across windows. By contrast, two other instruments, domestic credit and the required reserve ratio, co-move only weakly with the cycle in all periods, which highlights M2 growth as the more informative summary of China's monetary stance. Figure 15 plots the estimated cycle against M2 growth over the recent cycle periods.

fluctuations, we adopt the 3/3/7 rule of Kulish and Pagan (2021).

[21] We exclude the COVID-19 period, as consumption patterns were heavily constrained by lockdowns. The brief expansion in 2012Q1–2012Q4 is also omitted, as it appears to be a temporary reversal within a longer recession from 2008Q2 to 2019Q1.

Figure 15: **Estimated Cycle and M2 Growth**

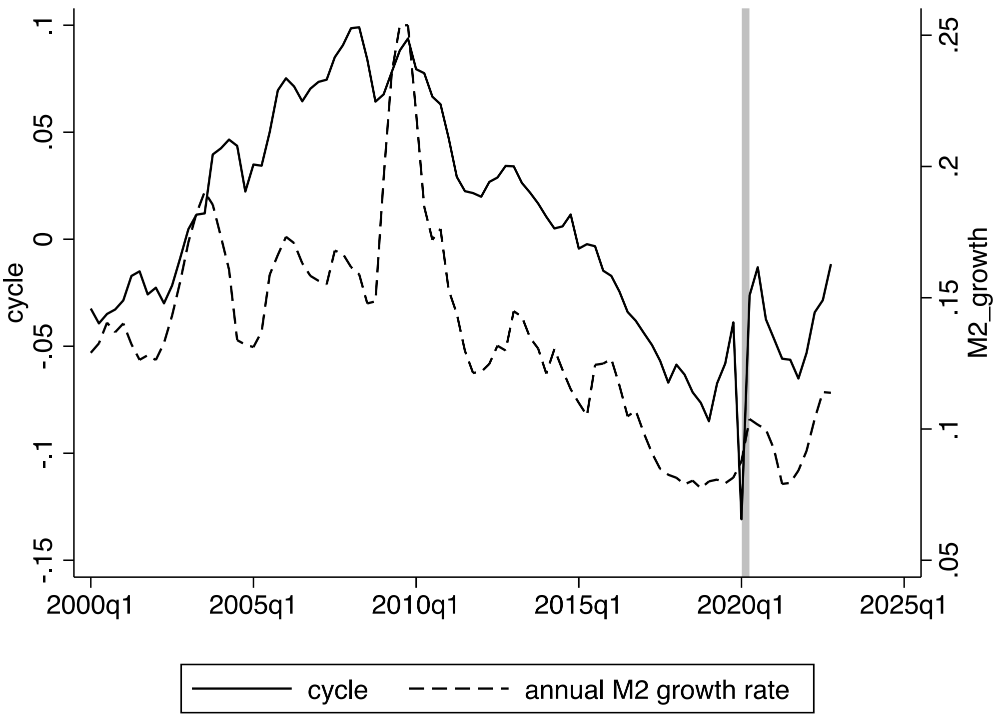

Notes: The correlation between the estimated cycle and the annual M2 growth rate is 0.82 over the period 2000Q1–2022Q4.

6 Conclusion

This paper develops a consumption-based trend-cycle decomposition for emerging market economies and evaluates it against standard filtering methods. Motivated by the permanent income hypothesis, the method exploits the long-run relationship between consumption and income to identify the trend in GDP. The observation that consumption can serve as a direct measure of the GDP trend is not itself new (Crucini & Shintani, 2015); our contribution is to show why this approach is especially suited to emerging market economies, to characterize through Monte Carlo simulation the conditions under which standard filters fail, to generalize the method to empirically relevant departures from the strict permanent income hypothesis, and to validate it across a broad cross-country sample and a detailed study of China. Both the simulations and the empirical results indicate that the method delivers cycles consistent with the HP filter for advanced economies, while for emerging market economies, it yields better out-of-sample forecasting performance, the environment in which standard filters are least reliable. Overall, the consumption-based method offers a robust and practical alternative for decomposing GDP, particularly in economies subject to large permanent shocks and structural breaks.

From a policy perspective, reliable identification of the GDP trend is essential for assessing potential output and setting macroeconomic policy. By remaining robust under structural change and large permanent shocks, the consumption-based approach offers policymakers in emerging market economies a more dependable tool for monitoring economic performance and informing policy decisions.

References

Aguiar, M., & Gopinath, G. (2007). Emerging market business cycles: The cycle is the trend. *Journal of Political Economy*, *115*(1), 69–102.

Attanasio, O. P. (1999). Consumption. *Handbook of macroeconomics*, *1*, 741–812.

Bai, J., & Perron, P. (1998). Estimating and testing linear models with multiple structural changes. *Econometrica*, 47–78.

Blanchard, O. J., & Quah, D. (1989). The dynamic effects of aggregate demand and aggregate supply. *American Economic Review*, *79*(4), 655–673.

Bry, G., & Boschan, C. (1971). Front matter to" cyclical analysis of time series: Selected procedures and computer programs". In *Cyclical analysis of time series: Selected procedures and computer programs* (pp. 13–2).

Campbell, J. Y., & Mankiw, N. G. (1989). Consumption, income, and interest rates: Reinterpreting the time series evidence. *NBER Macroeconomics Annual*, *4*, 185–216.

Canova, F. (1994). Detrending and turning points. *European Economic Review*, *38*(3-4), 614–623.

Chang, C., Chen, K., Waggoner, D. F., & Zha, T. (2016). Trends and cycles in china's macroeconomy. *NBER Macroeconomics Annual*, *30*(1), 1–84.

Chen, K., Ren, J., & Zha, T. (2018). The nexus of monetary policy and shadow banking in china. *American Economic Review*, *108*(12), 3891–3936.

Chen, K., & Zha, T. (2018). Macroeconomic effects of china's financial policies. *Handbook of China's Financial Markets*, *1*, 151–182.

Cochrane, J. H. (1988). How big is the random walk in gnp? *Journal of Political Economy*, *96*(5), 893–920.

Cochrane, J. H. (1994). Permanent and transitory components of gnp and stock prices. *Quarterly Journal of Economics*, *109*(1), 241–265.

Crucini, M. J., & Shintani, M. (2015). Measuring international business cycles by saving for a rainy day. *Canadian Journal of Economics*, *48*(4), 1266–1290.

Del Negro, M., Giannone, D., Giannoni, M. P., & Tambalotti, A. (2017). Safety, liquidity, and the natural rate of interest. *Brookings Papers on Economic Activity*, *2017*(1), 235–316.

Diebold, F. X., & Mariano, R. S. (1995). Comparing predictive accuracy. *Journal of Business & Economic Statistics*, *13*(3), 253–263.

Furlanetto, F., Hagelund, K., Hansen, F., & Robstad, Ø. (2023). Norges bank output gap estimates: Forecasting properties, reliability, cyclical sensitivity and hysteresis. *Oxford Bulletin of Economics and Statistics*, *85*(1), 238–267.

Garratt, A., Robertson, D., & Wright, S. (2006). Permanent vs transitory components and economic fundamentals. *Journal of Applied Econometrics*, *21*(4), 521–542.

Hahm, J.-H., & Steigerwald, D. G. (1999). Consumption adjustment under time-varying income uncertainty. *Review of Economics and Statistics*, *81*(1), 32–40.

Hall, R. E. (1978). Stochastic implications of the life cycle-permanent income hypothesis: Theory and evidence. *Journal of Political Economy*, *86*(6), 971–987.

Hamilton, J. D. (2018). Why you should never use the hodrick-prescott filter. *Review of Economics and Statistics*, *100*(5), 831–843.

Hamilton, J. D., Harris, E. S., Hatzius, J., & West, K. D. (2016). The equilibrium real funds rate: Past, present, and future. *IMF Economic Review*, *64*(4), 660–707.

Kamber, G., Morley, J., & Wong, B. (2018). Intuitive and reliable estimates of the output gap from a beveridge-nelson filter. *Review of Economics and Statistics*, *100*(3), 550–566.

King, R. G., Plosser, C. I., Stock, J. H., & Watson, M. W. (1991). Stochastic trends and economic fluctuations. *American Economic Review*, *81*(4), 819.

Kulish, M., & Pagan, A. (2021). Turning point and oscillatory cycles: Concepts, measurement, and use. *Journal of Economic Surveys*, *35*(4), 977–1006.

Levy, D., & Dezhbakhsh, H. (2003). International evidence on output fluctuation and shock persistence. *Journal of Monetary Economics*, *50*(7), 1499–1530.

Morley, J. (2007). The slow adjustment of aggregate consumption to permanent income. *Journal of Money, Credit and Banking*, *39*(2-3), 615–638.

Morley, J., 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*(2), 235–243.

Morley, J., Tran, T. D., & Wong, B. (2024). A simple correction for misspecification in trend-cycle decompositions with an application to estimating r. *Journal of Business & Economic Statistics*, *42*(2), 665–680.

Nelson, C. R. (2008). The beveridge–nelson decomposition in retrospect and prospect. *Journal of Econometrics*, *146*(2), 202–206.

Phillips, P. C., & Shi, Z. (2021). Boosting: Why you can use the hp filter. *International Economic Review*, *62*(2), 521–570.

Stock, J. H., & Watson, M. W. (1988). Variable trends in economic time series. *Journal of Economic Perspectives*, *2*(3), 147–174.

Stock, J. H., & Watson, M. W. (1999). Business cycle fluctuations in us macroeconomic time series. *Handbook of Macroeconomics*, *1*, 3–64.

Summers, L. H. (1981). Capital taxation and accumulation in a life cycle growth model. *American Economic Review*, *71*(4).

Watson, M. W. (1986). Univariate detrending methods with stochastic trends. *Journal of Monetary Economics*, *18*(1), 49–75.

A Data and Tests

A.1 Permanent Income Hypothesis Assumption

Random Walk

To test the random walk hypothesis, we follow Hall (1978) by regressing consumption growth on second to fourth lags of GDP and consumption growth. [22] We estimate the equation using heteroskedasticity- and autocorrelation-consistent (HAC) standard errors as follows:

$$\Delta c_t = \alpha + \beta_2 \Delta c_{t-2} + \beta_3 \Delta c_{t-3} + \beta_4 \Delta c_{t-4} + \gamma_2 \Delta y_{t-2} + \gamma_3 \Delta y_{t-3} + \gamma_4 \Delta y_{t-4} + \varepsilon_t \quad (9)$$

If all coefficients are insignificant, which suggests the serial independence of consumption growth, then the consumption follows a random walk.

Cointegration

To test the cointegration assumption, we employ both the Engle–Granger and Johansen tests. After confirming the presence of a cointegrating relationship between GDP and consumption, we estimate the long-run coefficients by regressing GDP on consumption using Dynamic Ordinary Least Squares (DOLS) with heteroskedasticity- and autocorrelation-consistent (HAC) standard errors.

An alternative approach is to test the stationarity of the logarithmic C/Y ratio directly using unit root tests such as the Augmented Dickey–Fuller (ADF) or Phillips–Perron (PP) tests. As noted by Cochrane (1994), the stationarity of the C/Y ratio is mathematically equivalent to the existence of a cointegration relationship between consumption and GDP.

[22] Hall (1978) and Campbell and Mankiw (1989) note that if the permanent income hypothesis holds in continuous time, and observed consumption represents a time average of a random walk. Consequently, changes in consumption exhibit a first-order serial correlation of approximately 0.25, which may lead to a rejection of the hypothesis even when it is true.

Table A1. **Test Results**

AEs	Random Walk	Cointegration	Coefficient = 1	EMEs	Random Walk	Cointegration	Coefficient = 1
Australia	Yes	No	No	Argentina	Yes	Yes	No
Austria	No	Yes	No	Brazil	Yes	Yes	No
Belgium	No	No	No	Chile	No	No	No
Canada	Yes	No	No	China	No	Yes	Yes
Denmark	Yes	No	No	Colombia	No	Yes	No
Finland	No	No	No	Costa Rica	No	No	No
France	No	No	No	Czech Republic	No	Yes	No
Germany	No	Yes	No	Estonia	No	Yes	No
Iceland	No	Yes	No	Greece	No	No	No
Ireland	No	No	No	Hungary	No	No	No
Israel	Yes	No	No	India	Yes	Yes	No
Italy	No	Yes	Yes	Indonesia	No	No	No
Japan	Yes	Yes	Yes	Latvia	No	Yes	Yes
Luxembourg	Yes	Yes	No	Lithuania	Yes	No	No
Netherlands	No	No	No	Mexico	Yes	No	No
New Zealand	No	No	No	Poland	No	Yes	No
Norway	No	Yes	No	Russian Federation	No	No	No
Portugal	No	Yes	No	Slovakia	No	Yes	No
South Korea	Yes	Yes	No	Slovenia	No	No	No
Spain	Yes	Yes	No	South Africa	No	Yes	No
Sweden	Yes	No	No	Turkey	Yes	Yes	No
Switzerland	No	No	No				
United Kingdom	Yes	Yes	No				
United States	No	Yes	No				

Notes:
The random walk hypothesis is tested following Hall (1978). The cointegration relationship between consumption and income is examined using the Engle–Granger and Johansen tests, while the restriction that the cointegration coefficient equals one is evaluated using DOLS regression.

A.2 Cointegration between Log(C/Y) and Real Interest Rate

The Johansen test suggests the presence of a cointegration relationship between the log C/Y ratio ($c_t - y_t$) and the real interest rate for both Mexico and the United States at the annual frequency. The p-values for the null of at most one cointegrating relationship are 0.001 for Mexico and 0.045 for the United States. The figure plots the log C/Y ratio, the real interest rate, and the estimated cointegration error.

Figure 16: Log C/Y Ratio and Real Interest Rate

(a) United States

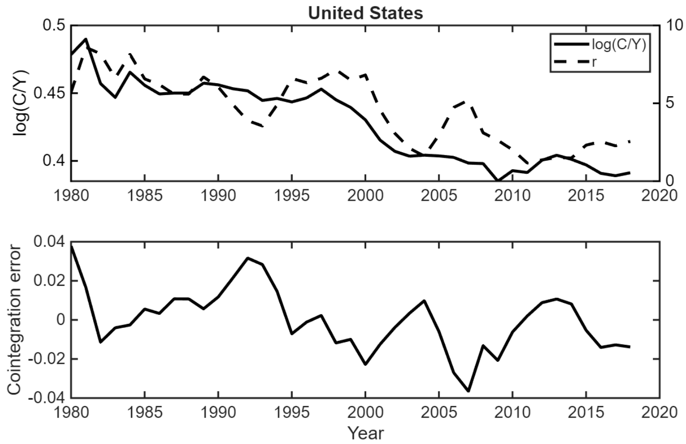

(b) Mexico

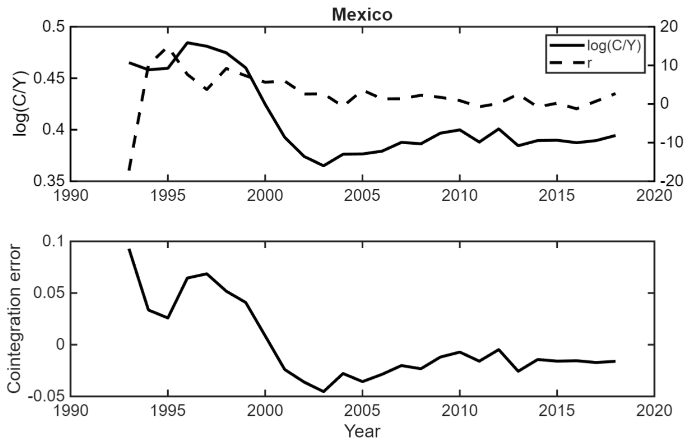

A.3 Data Span and Country Classification

Table 5: **Data Span and Country Classification**

Country	IMF_class	StartYQ	EndYQ
Argentina	0	2004Q1	2023Q4
Australia	1	1980Q1	2023Q3
Austria	1	1995Q1	2023Q4
Belgium	1	1995Q1	2023Q4
Brazil	0	1996Q1	2023Q3
Canada	1	1980Q1	2023Q4
Chile	0	1996Q1	2023Q3
China	0	1992Q1	2022Q4
Colombia	0	2005Q1	2023Q3
Costa Rica	0	1991Q1	2023Q4
Czech Republic	0	1995Q1	2023Q4
Denmark	1	1995Q1	2023Q3
Estonia	0	1995Q1	2023Q4
Finland	1	1990Q1	2023Q3
France	1	1980Q1	2023Q4
Germany	1	1991Q1	2023Q4
Greece	0	1995Q1	2023Q3
Hungary	0	1995Q1	2023Q3
Iceland	1	1995Q1	2023Q3
India	0	1996Q2	2023Q4
Indonesia	0	2000Q1	2023Q4
Ireland	1	1995Q1	2023Q4
Israel	1	1995Q1	2023Q3
Italy	1	1996Q1	2023Q4
Japan	1	1994Q1	2023Q3
Latvia	0	1995Q1	2023Q4
Lithuania	0	1995Q1	2023Q4
Luxembourg	1	1995Q1	2023Q3
Mexico	0	1993Q1	2023Q4
Netherlands	1	1996Q1	2023Q3
New Zealand	1	1987Q2	2023Q3
Norway	1	1980Q1	2023Q3
Poland	0	1995Q1	2023Q3
Portugal	1	1995Q1	2023Q4
Russian Federation	0	2003Q1	2021Q3
Slovakia	0	1995Q1	2023Q3
Slovenia	0	1995Q1	2023Q3
South Africa	0	1993Q1	2023Q4
South Korea	1	1980Q1	2023Q4
Spain	1	1995Q1	2023Q4
Sweden	1	1993Q1	2023Q4
Switzerland	1	1980Q1	2023Q3
Turkey	0	1998Q1	2023Q3
United Kingdom	1	1980Q1	2023Q3
United States	1	1980Q1	2023Q4

Notes:
1. We collect quarterly, seasonally adjusted data on real GDP and real private consumption from the OECD database.
2. IMF class equals 1 for advanced economies and 0 for emerging market economies. Some Eastern European countries are reclassified as emerging market economies.

B Additional Monte Carlo

This appendix provides additional Monte Carlo evidence that complements the simulations in the main text. We extend the baseline exercise along three dimensions. First, we consider a short-sample design with only 60 observations to assess short sample performance. Second, we examine hybrid DGPs that combine multiple departures from the baseline, such as large trend shocks together with transitory consumption components. Third, we increase the number of Monte Carlo replications to verify that the results are stable and not driven by simulation noise.

All DGPs introduce an additional transitory component in consumption, same as DGP 4 in Table 2. In DGP 6, the volatility of the transitory consumption shock is set equal to that of GDP, whereas in all other DGPs it is set to half of GDP's volatility. It shows that the consumption-based method continues to outperform both the HP filter and the Hamilton filter even in the presence of transitory consumption. Compared to DGP 4, performance improves when the transitory component is relatively small.

In DGP 7, we introduce a larger permanent shock to reflect features typical of emerging market economies. In DGP 8, we allow for a structural break in trend growth. In both cases, the consumption-based method still captures the trend dynamics effectively, as the gap between income and consumption $(y-c)$ remains informative about the permanent component. By contrast, the HP and Hamilton filters perform less reliably under these conditions, with lower correlations and higher RMSEs.

DGPs 9 and 10 assume breaks in $\bar{c}$ from Equation (4), implying that the long-run C/Y ratio varies over time. In these cases, the consumption-based method, when implemented under the assumption of a constant C/Y ratio, fails to accurately estimate the cyclical component. Given the short sample sizes typical of emerging market economies, identifying the exact break points is often infeasible. We therefore propose estimating a time-varying γ_t using a rolling-window approach to improve the robustness of the method.

Similarly, DGP 11 considers the case where the true cointegration coefficient between consumption and income is not unity. When the consumption-based method incorrectly imposes a cointegration coefficient of one, it fails to fully capture the true cycle, especially over longer horizons. However, this scenario can also be interpreted as involving a time-varying long-run C/Y ratio. The rolling-window approach can thus help improve performance in this environment as well.

DGP 12 combines the missing variable with large trend shocks. In the short-sample setting ($T = 60$), both the constant $\bar{c}$ specification and the rolling-window estimator outperform all alternative methods. Interestingly, the constant $\bar{c}$ performs better than the rolling-window approach. In the long-sample setting ($T = 160$), however, the relative performance reverses: the constant $\bar{c}$ estimator deteriorates, while the rolling-window method improves. This contrast primarily reflects the influence of the drift parameter η, whose cumulative effect becomes more pronounced in longer samples. When η is relatively large, the advantage of the rolling-window estimator becomes even more evident.

Table 6: **Monte Carlo Simulation: With Transitory Consumption**

	T=60		T=160	
	Correlation	RMSE	Correlation	RMSE
DGP 6. Same Transitory Consumption Shock Variance, $\sigma_c = \sigma_y$				
HP Filter	0.736	1.684	0.733	1.819
Boosted HP-BIC	0.481	2.232	0.492	2.354
Hamilton Filter	0.634	2.654	0.624	3.000
Consumption-based method	0.876	1.315	0.873	1.479
C-based with Rolling mean	0.737	1.709	0.709	1.899
DGP 7. Larger Permanent Shock Variance, $\sigma_v = \sqrt{2}$, $\sigma_c = 0.5\sigma_y$				
HP Filter	0.556	2.316	0.562	2.438
Boosted HP-BIC	0.357	2.469	0.376	2.592
Hamilton Filter	0.486	3.846	0.466	4.517
Consumption-based method	0.964	0.658	0.963	0.739
C-based with Rolling mean	0.799	1.527	0.764	1.748
DGP 8. Two Breaks on Trend Growth				
HP Filter	0.706	1.770	0.723	1.850
Boosted HP-BIC	0.482	2.230	0.493	2.353
Hamilton Filter	0.583	2.966	0.559	3.491
Consumption-based method	0.964	0.658	0.963	0.739
C-based with Rolling mean	0.799	1.527	0.764	1.748
DGP 9. Two Breaks on Mean of Log(C/Y), $\sigma_c = 0.5\sigma_y$				
HP Filter	0.736	1.684	0.733	1.819
Boosted HP-BIC	0.481	2.232	0.492	2.354
Hamilton Filter	0.634	2.654	0.624	3.000
Consumption-based method	0.660	2.713	0.699	2.732
C-based with Rolling mean	0.708	1.801	0.731	1.837
DGP 10. Two Breaks on Trend Growth and Mean of Log(C/Y)				
HP Filter	0.542	2.377	0.557	2.460
Boosted HP-BIC	0.358	2.473	0.377	2.593
Hamilton Filter	0.461	4.065	0.436	4.861
Consumption-based method	0.660	2.713	0.699	2.732
C-based with Rolling mean	0.708	1.801	0.731	1.837
DGP 11. True Cointegration Coefficient $= 0.90$, $\sigma_c = 0.5\sigma_y$				
HP Filter	0.736	1.684	0.733	1.819
Boosted HP-BIC	0.481	2.232	0.492	2.354
Hamilton Filter	0.634	2.654	0.624	3.000
Consumption-based method	0.840	1.535	0.561	3.829
C-based with Rolling mean	0.796	2.533	0.763	1.752
DGP 12. Missing variable, transitory consumption, large trend shocks, $\sigma_c = 0.5\sigma_y$, $\sigma_v = \sqrt{2}$				
HP Filter	0.556	2.316	0.562	2.438
Boosted HP-BIC	0.357	2.469	0.376	2.592
Hamilton Filter	0.486	3.846	0.466	4.517
Consumption-based Method	0.912	1.073	0.737	2.425
C-based Method with rolling window	0.798	1.531	0.764	1.750

Notes:
1. Baseline settings, seeds, and iteration choices are the same as in Table 1.
2. In DGP 6, $\sigma_c = \sigma_y = 0.69$, and in DGPs 7–12, $\sigma_c = 0.5\sigma_y = 0.345$.
3. In DGPs 8–10, we allow for breaks in $\bar{c}$ and μ, with the timing and magnitude drawn randomly.
4. In DGP 12, we allow $\bar{c}_t$ to be time-varying, specifically, $\bar{c}_t = 0.05 + \bar{c}_{t-1} + v_t$, with $\sigma_{e,v} = 0.05$.

B.1 Monte Carlo, BN-VECM with More Lags

In Tables 1 and 2, the BN–VECM is estimated with two lags, as specifications with more lags yield worse results (see Table 7 below).

Table 7: **Monte Carlo Simulation: BN-VECM with More Lags**

	4 Lags		6 Lags		8 Lags	
	Correlation	RMSE	Correlation	RMSE	Correlation	RMSE
DGP 1	0.956	0.785	0.949	0.894	0.935	1.173
DGP 2	0.915	1.220	0.880	1.407	0.840	1.945
DGP 3	0.700	28.993	0.644	7.535	0.591	18.955
DGP 4	0.876	1.444	0.842	1.588	0.802	1.899
DGP 5	0.699	2.324	0.663	2.498	0.630	4.762

Notes:
1. All DGPs are exactly the same as Tables 1 and 2, only the lag number of VECM is changed.

C Additional Results and Discussions

C.1 Comparison between C-based and HP Filtered Cycles

We show the share of observations in which they indicate the same cyclical sign. Same as the evidence in the correlation comparison in Figure 6, in advanced economies, the consumption-based method yields trend and cycle estimates that are broadly aligned with those obtained from the HP filter, whereas in emerging market economies the two measures are more likely to diverge.

Figure 17: **Same Sign Observation between C-based and HP filtered Cycles**