Statistisk sentralbyrå
Statistics Norway

Decomposing the Output Gap

Robust Univariate and Multivariate Hodrick–Prescott Filtering with Extreme Observations

Håvard Hungnes

TALL SOM FORTELLER
DISCUSSION PAPERS
1031

Discussion Papers: comprise research papers intended for international journals or books. A preprint of a Discussion Paper may be longer and more elaborate than a standard journal article, as it may include intermediate calculations and background material etc.

The Discussion Papers series presents results from ongoing research projects and other research and analysis by SSB staff. The views and conclusions in this document are those of the authors.

Published: December 2025

Abstracts with downloadable Discussion Papers
in PDF are available on the Internet:

https://www.ssb.no/discussion-papers
http://ideas.repec.org/s/ssb/dispap.html

ISSN 1892-753X (electronic)

Abstract

This paper introduces two methodological improvements to the Hodrick– Prescott (HP) filter for decomposing GDP into trend and cycle components. First, we propose a robust univariate filter that accounts for extreme observations — such as the COVID–19 pandemic — by treating them as additive outliers. Second, we develop a multivariate HP filter that incorporates time–varying, import–adjusted budget shares of GDP sub–components. This adaptive weighting minimizes cyclical variance and yields a more stable trend estimate. Applying the framework to U.S. data, we find that private investment is the dominant source of cyclical fluctuations, while government expenditure exhibits a persistent counter–cyclical pattern. The proposed approach enhances real–time policy analysis by reducing endpoint bias and improving the identification of cyclical dynamics.

Keywords: output gap; Hodrick–Prescott filter; robust filtering; multivariate decomposition; additive outliers; time–varying budget shares; business cycle analysis

JEL classification: E32, C22, E37, C43, C51

Acknowledgements: I thank Thomas von Brasch, Terje Skjerpen, and participants at the 27th Dynamic Econometrics Conference, held at Bayes Business School in London, September 2025, for valuable feedback. I gratefully acknowledge financial support from Norges Bank in the form of a travel grant. The usual disclaimer applies.

Address: Håvard Hungnes, Statistics Norway, Research Department. E-mail: hhu@ssb.no

Sammendrag

Dette manuskriptet presenterer en utvidelse av den velkjente Hodrick–Prescott (HP)-filtermetoden for å dekomponere bruttonasjonalproduktet (BNP) i trend- og konjunkturkomponenter. Analysen forbedrer metoden på to måter. For det første innføres en robust behandling av ekstreme observasjoner, som for eksempel sjokket under koronapandemien. I stedet for at slike ekstreme utslag får påvirke trendestimatet for hele tidsserien, modelleres de som additive uteliggere i filterets optimeringsproblem. På denne måten kan man estimere et underliggende trendforløp som reflekterer utviklingen i økonomien uten at pandemien får en varig effekt på potensielt BNP.

For det andre utvikles et multivariat HP-filter som utnytter informasjonen i BNPs underkomponenter — privat konsum, privat investering, offentlig etterspørsel og eksport — justert for importinnhold. Mens tidligere multivariate tilnærminger har brukt faste vekter, foreslås her en metode med tidsvarierende, importjusterte budsjettandeler. Disse andelene glattes over tid og brukes til å re-vekte komponentene slik at den felles trenden blir mest mulig stabil. Dette gjør det mulig å redusere den sykliske variansen og samtidig følge endringer i BNPs sammensetning over tid.

Metoden anvendes på data for USA i perioden 1995–2024. Resultatene viser at privat investering er den komponenten som har bidratt mest til konjunktursvingninger, mens offentlig konsum og investeringer har virket klart motkonjunkturelt og har hatt en stabiliserende effekt på økonomien. Den nye multivariate metoden gir et trendestimat som ligger nær det som oppnås med et univariat filter, men som samtidig gir et mer nyansert bilde av hvilke komponenter som driver svingningene.

Tilnærmingen er spesielt egnet for sanntidsanalyser, der tidlige BNP-tall ofte er usikre og utsatt for revisjoner. Ved å legge større vekt på de mest stabile komponentene – og lavere vekt på mer volatile serier som investeringene – reduseres usikkerheten i trendestimatet, særlig mot slutten av observasjonsperioden. Dermed gir metoden et mer robust grunnlag for løpende vurderinger av konjunktursituasjonen og produksjonsgapet i økonomien.

1 Introduction

Decomposing the Gross Domestic Product (GDP) into trend and cycle components is a fundamental practice in macroeconometric analysis. Such decomposition a is particularly useful for identifying business cycles, which in turn informs both monetary and fiscal policy decisions.

Several methods exist for decomposing GDP, including band-pass filters (e.g., Christiano and Fitzgerald, 2003), the Beveridge–Nelson decomposition (Beveridge and Nelson, 1981), and structural time series models (e.g., Harvey, 1985). Among these, the Hodrick–Prescott (HP) filter remains widely used, despite facing severe criticism, most notably from Hamilton (2018). He argues that the filter introduces spurious dynamic relations and advocates replacing it with a regression-based alternative (the Hamilton regression filter). However, as pointed out by Harvey and Delle Monache (2009), even though the statistical model underlying the HP filter is technically misspecified — particularly due to its assumption of a white noise cyclical component — it nevertheless remains a useful and robust tool for practical trend-cycle decomposition. This view is supported by recent re-evaluations demonstrating that the HP filter often outperforms alternatives like the Hamilton regression filter in identifying cyclical patterns (Franke *et al.*, 2024; Schüler, 2024).

This paper proposes two methodological enhancements to the HP filter. First, we introduce a robust univariate filter that accounts for extreme observations. The standard HP filter is highly sensitive to large shocks, such as those experienced during the COVID-19 pandemic, which can distort the trend estimate. We propose a robust modification by treating such events as additive outliers within the filter's optimization problem, effectively isolating their impact without discarding valuable data. This is a deliberate modeling choice, distinct from a structural break, designed to estimate a counterfactual trend that reflects the pre-pandemic trajectory.

Second, we develop a multivariate HP filter that incorporates information from GDP sub-components. We refer to the weights of these components as 'budget shares', defined as the import-adjusted share of each final demand category (consumption, investment, government expenditure, and exports) in GDP. While existing multivariate methods typically assume constant budget shares (e.g., Kozicki, 1999; Harvey, 2006), our primary innovation is the use of time-varying budget shares. This is a crucial innovation, since a fixed-shares filter cannot account for shifts in the composition of GDP, such as a long-term change in the investment share of the economy.

Our contribution relates closely to recent developments in the literature addressing the challenge of large shocks and structural instability, such as during the COVID-19 pandemic. For instance, Kamber *et al.* (2025) propose refining the Beveridge-Nelson filter by optimizing the smoothing parameter to minimize changes in the trend during large shocks. Similarly, Morley *et al.* (2023) introduce variance

scale factors in a Bayesian VAR framework to effectively down-weight outlier observations. For the HP filter, Maranzano and Pelagatti (2025) propose a method for automatic detection of structural breaks.

While these approaches handle data instability stochastically, by parameter optimization, or by estimating break points, our method contributes by introducing a deterministic selection mechanism directly into the filtering problem. By treating extreme outliers as effectively uninformative, our framework provides a transparent algebraic derivation that allows us to explicitly quantify the 'price' of such interventions in terms of increased estimation uncertainty, complementing the numerical findings in the existing literature.

Our approach allows us to quantify how the importance of specific components (e.g., private investment) as drivers of the business cycle has evolved. This is particularly relevant for real-time analysis, where the ability to systematically down-weight noisy components is critical for reliable policy advice. Failure to account for time-varying budget shares can therefore lead to a biased decomposition of the business cycle and potentially misguided policy conclusions.

Applying this framework to U.S. data over the past 30 years, we find that private investment is the largest single contributor to cyclical fluctuations. In contrast, government expenditure exhibits a strong and persistent counter-cyclical pattern, consistent with a stabilizing role in the economy. By decomposing the aggregate business cycle into contributions from its underlying components, our method provides policymakers, such as central banks and fiscal authorities, with a more granular tool to assess economic conditions.

The remainder of the paper is structured as follows: Section 2 presents the univariate HP filter and our proposed correction for extreme observations. Section 3 develops the multivariate filter with time-varying budget shares. Section 4 concludes.

2 The univariate filter and extreme observations

This section lays the foundation for our analysis by first reviewing the standard univariate HP filter. In Section 2.1, we present its conventional formulation, derive its matrix solution, and discuss its statistical interpretation via a state-space model. In Section 2.2, we introduce our first methodological contribution: a robust modification of the filter that accounts for extreme observations with a specific application to the COVID-19 shock. In Section 2.3, we derive the variance-covariance matrices for the estimation errors of both the robust and generalized time-varying filters, extending the results of Schlicht (2005). Finally, in Section 2.4, we generalize the framework to allow for soft weighting and discuss its relationship to recent literature.

2.1 The standard HP filter and its state-space formulation

This section presents the standard Hodrick–Prescott (HP) filter and its statistical interpretation. We begin by defining the decomposition, then derive the matrix solution, and finally present the state-space representation.

The HP filter decomposes a time series y_t, typically log-transformed GDP, into the two unobserved components trend τ_t and cyclical ε_t:

$$y_t = \tau_t + \varepsilon_t. \tag{1}$$

The central challenge is to identify and separate these two unobserved components from the single observed series y_t. The HP filter, originally proposed by Hodrick and Prescott (1981), is widely used. It addresses this problem by defining the trend as the component that best fits the data, subject to a penalty on changes in its growth rate.

This trade-off between data fidelity and trend smoothness is formalized through the objective function below. Specifically, the HP filter finds the trend vector $\boldsymbol{\tau} = (\tau_1, \dots \tau_T)'$ that minimizes an objective function composed of two parts. The first is the "fit term" (or "goodness-of-fit term"), $\sum_{t=1}^{T}(y_t - \tau_t)^2$, which measures the deviation from the data. The second is the "penalty term", $\lambda \sum_{t=3}^{T} \left(\Delta^2 \tau_t\right)^2$, which penalizes the lack of smoothness. This penalty uses the standard second-difference operator, defined as $\Delta^2 \tau_t \equiv (\tau_t - \tau_{t-1}) - (\tau_{t-1} - \tau_{t-2}) = \tau_t - 2\tau_{t-1} + \tau_{t-2}$, for $t = 3, \dots, T$. The full objective function is:

$$\begin{aligned} L &= \sum_{t=1}^{T}(y_t - \tau_t)^2 + \lambda \sum_{t=3}^{T} \left(\Delta^2 \tau_t\right)^2 \\ &= (\boldsymbol{y} - \boldsymbol{\tau})'(\boldsymbol{y} - \boldsymbol{\tau}) + \lambda \boldsymbol{\tau}' \boldsymbol{D}_2' \boldsymbol{D}_2 \boldsymbol{\tau}, \end{aligned} \tag{2}$$

where $\boldsymbol{y} = (y_1, \dots, y_T)'$ is the vector of observations, and $\boldsymbol{D}_2$ is the $(T-2) \times T$ second-order differencing matrix that maps the vector $\boldsymbol{\tau}$ to the vector of its second differences $(\Delta^2 \tau_3, \dots, \Delta^2 \tau_T)'$:

$$\boldsymbol{D}_2 = \begin{bmatrix} 1 & -2 & 1 & 0 & 0 & \dots & 0 \\ 0 & 1 & -2 & 1 & 0 & \dots & 0 \\ \vdots & \ddots & \ddots & \ddots & \ddots & \ddots & \vdots \\ 0 & \dots & 0 & 1 & -2 & 1 & 0 \\ 0 & \dots & 0 & 0 & 1 & -2 & 1 \end{bmatrix}$$

The optimal trend component can be obtained by minimizing the loss function (2) with respect to $\boldsymbol{\tau}$, as shown in Danthine and Girardin (1989):

$$\hat{\boldsymbol{\tau}} = (\boldsymbol{I}_T + \lambda \boldsymbol{D}_2' \boldsymbol{D}_2)^{-1} \boldsymbol{y}, \tag{3}$$

where $\boldsymbol{I}_T$ is the identity matrix of dimension $T \times T$.

A formal statistical foundation for the HP filter is provided by its state-space representation, as developed by Akaike (1980) and Harvey and Trimbur (2008). In this framework, the components of the decomposition are given a statistical interpretation. The trend τ_t is modeled as a second-order integrated process (a smooth trend, see Harvey and Jaeger, 1993), following the process:

$$\tau_{t+1} = \tau_t + \beta_t, \quad \beta_{t+1} = \beta_t + \zeta_t, \quad \zeta_t \sim IID(0, \sigma_\zeta^2).$$

From this formulation of the trend, we can derive the relationship to the second-difference matrix $\boldsymbol{D}_2$. By substituting the equations, we find the second-difference of the trend:

$$\begin{aligned} \Delta^2 \tau_t \equiv \tau_t - 2\tau_{t-1} + \tau_{t-2} &= (\tau_t - \tau_{t-1}) - (\tau_{t-1} - \tau_{t-2}) \\ &= (\beta_{t-1}) - (\beta_{t-2}) \\ &= \zeta_{t-2} \end{aligned}$$

This key identity shows that the second-difference of the trend at time t is equal to the trend innovation from time $t-2$. This allows us to express the state-space model in matrix notation. The $(T-2)\times 1$ vector $\boldsymbol{D}_2\boldsymbol{\tau}$, which starts with the element $(\tau_3 - 2\tau_2 + \tau_1)$, corresponds exactly to the vector of innovations $\boldsymbol{\zeta} = (\zeta_1, \ldots, \zeta_{T-2})'$. The full model in (1) and the equations above can be written compactly as:

$$\boldsymbol{y} = \boldsymbol{\tau} + \boldsymbol{\varepsilon} \tag{4}$$

$$\boldsymbol{D}_2\boldsymbol{\tau} = \boldsymbol{\zeta} \tag{5}$$

where $\boldsymbol{y}$, $\boldsymbol{\tau}$, and $\boldsymbol{\varepsilon}$ are $T \times 1$ vectors, and $\boldsymbol{\zeta}$ is the $(T-2) \times 1$ vector of trend innovations.

The cyclical component ε_t is formally treated as a white noise process in the state-space representation. This statistical model clarifies the role of the smoothing parameter λ. For predictive purposes (e.g., using the Kalman filter), the optimal choice is $\lambda = \sigma_\varepsilon^2/\sigma_\zeta^2$, where σ_ε^2 is the variance of ε_t. This sets the parameter as the inverse of the signal-to-noise ratio, balancing the observation noise against the trend's growth rate noise.

However, this white noise assumption is unrealistic for persistent macroeconomic cycles. This simplification is the primary reason why a specific value for λ is typically imposed rather than estimated. As pointed out by Harvey and Jaeger (1993, p. 233):

> If [(1)] was believed to be the true model, [λ] could be estimated by maximum likelihood [...]. However, the whole reason for applying the HP filter is the belief that detrended data consist of something more than white noise. Thus, a value of [λ] is imposed, rather than estimated.

Although this implies that the filter's underlying statistical model is technically misspecified when applied to persistent cycles, the HP filter remains a valuable and robust tool for practical trend-cycle decomposition, as demonstrated by Harvey and Delle Monache (2009).

When applying the Hodrick–Prescott filter to quarterly, seasonally-adjusted U.S. GDP, we follow the convention of setting the smoothing parameter to $\lambda = 1600$ (Hodrick and Prescott, 1981, 1997). While this value is a widely adopted standard, its theoretical foundation has been debated and alternative selection methods have been proposed (e.g., Ravn and Uhlig, 2002; de Jong and Sakarya, 2016; Hamilton, 2018; Phillips and Shi, 2021).

In Figure 1, the blue line in the upper panel shows the standard HP trend, plotted together with the actual GDP in black. In the lower panel, the resulting cycle is shown in blue. The COVID-19 pandemic clearly affects the trend and cycle. With the standard HP filter, the economy is estimated to be almost 2 percent above trend in the quarters before the start of the pandemic, as the trend is pulled down in anticipation of the large negative shock. In the next section, we will therefore look at a better way to handle such extreme events.

2.2 A modification for extreme observations

This section introduces a robust extension of the HP filter designed to handle extreme observations, such as the COVID-19 pandemic. Extreme shocks pose a challenge for trend–cycle decomposition. The standard HP filter is sensitive to large deviations, which can distort the estimated trend across the entire sample. To address this, we propose a modification that treats extreme events as additive outliers. This approach isolates their impact.

Consider the following modified specification:

$$y_t = \delta_t + \tau_t^u + \varepsilon_t^u \tag{6}$$

where δ_t represents a process of impulse dummies for the COVID-19 period, which are included as additive outliers (Fox, 1972). Furthermore, the superscript 'u' denotes the underlying (uncontaminated) process, excluding the direct impact of these extreme shocks.

In the U.S., the effect of COVID-19 on GDP was almost only noticeable in 2020. Since we are studying U.S. data here, we therefore use additive outliers only for the four quarters in 2020, as this period captures the vast majority of the initial economic decline and subsequent rebound. Thus, for all other quarters we have $\delta_t = 0$.

Two possible interpretations arise regarding the relationship between the components we want to identify (τ_t, ε_t) and the underlying components $(\tau_t^u, \varepsilon_t^u)$:

- If we want the trend during the pandemic to be the potential GDP if there had not been a pandemic; $\tau_t = \tau_t^u$, and $\varepsilon_t = \varepsilon_t^u + \delta_t$.
- If we want the trend during the pandemic to be the potential GDP given that there is a pandemic, and we assume that every period during the pandemic is different (i.e., special) such that we can only use one observation to estimate the trend; $\tau_t = \tau_t^u + \delta_t$, and $\varepsilon_t = \varepsilon_t^u$.

For both cases, the objective function becomes:

$$\begin{aligned} L &= \sum_{t=1}^{T} 1_{t \neq COVID19}(y_t - \tau_t^u)^2 + \lambda \sum_{t=2}^{T-1} \left[(\tau_{t+1}^u - \tau_t^u) - (\tau_t^u - \tau_{t-1}^u)\right]^2 \\ &= (\boldsymbol{y} - \boldsymbol{\tau}^u)' \boldsymbol{H} (\boldsymbol{y} - \boldsymbol{\tau}^u) + \lambda \boldsymbol{\tau}^{u\prime} \boldsymbol{D}_2' \boldsymbol{D}_2 \boldsymbol{\tau}^u \end{aligned}$$

where $1_{t \neq COVID19}$ is an indicator that is 1 if the period is not part of the COVID-19 pandemic, and zero otherwise; and $\boldsymbol{H}$ is a diagonal matrix with those indicators along its main diagonal.

For the remainder of this paper, we follow the first interpretation presented in the bullet point above, assuming that the potential trend $\boldsymbol{\tau}$ is equal to the underlying trend ($\boldsymbol{\tau} = \boldsymbol{\tau}^u$), where the estimated trend $\hat{\boldsymbol{\tau}}$ is given by the solution to this minimization problem.

The estimated trend then becomes

$$\hat{\boldsymbol{\tau}} = (\boldsymbol{H} + \lambda \boldsymbol{D}_2' \boldsymbol{D}_2)^{-1} \boldsymbol{H} \boldsymbol{y}, \tag{7}$$

see Dermoune *et al.* (2009, Prop. 2.1). The matrix $(\boldsymbol{H} + \lambda \boldsymbol{D}_2' \boldsymbol{D}_2)$ is full rank as long as there are at least two distinct time points with non-zero diagonal elements in $\boldsymbol{H}$, that is, at least two regular observations outside the excluded period. This condition ensures that both the level and the slope of the trend are identified. In practice, this requirement is easily satisfied, since the excluded (extreme) observations typically form a short block within the interior of the sample. A formal proof of the rank condition is provided in Appendix C.1.

Both interpretations above imply that the trend after the pandemic is a projection based on the trend before the pandemic, effectively assuming that the shock had no lasting impact on the trend's level or slope. If one wishes to allow for a permanent effect or simply adapt to extreme volatility, the standard fixed-parameter framework is insufficient. Recent literature highlights the need for such adaptation. For instance, Mei *et al.* (2024) introduce a boosted HP filter that uses iterative machine-learning techniques to capture complex trend dynamics. Similarly, Maranzano and Pelagatti (2025) propose a modification that automatically detects and incorporates structural breaks. Conversely, to handle the massive but largely transitory COVID-19 shock, Kamber *et al.* (2025) demonstrate that optimizing the

Figure 1: Univariate HP trend and cycle, with and without COVID-19 correction

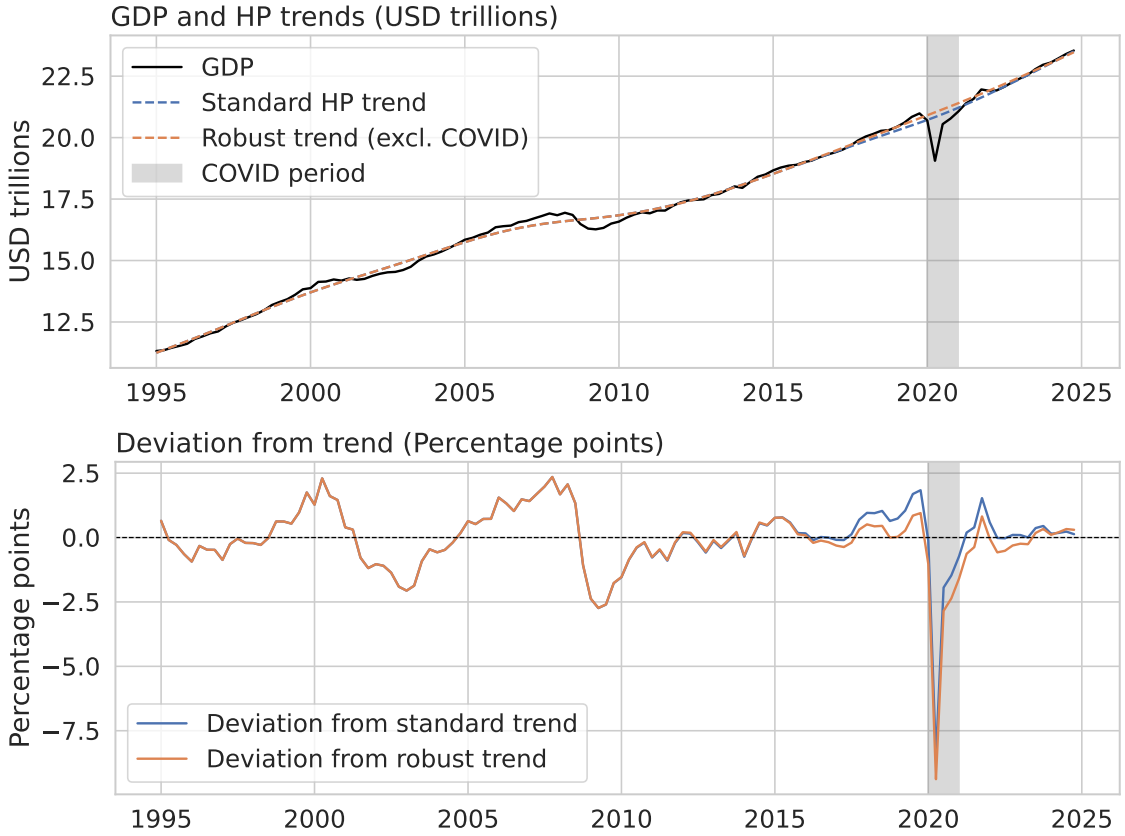

Notes: Upper panel: Actual GDP ($\exp(y_t)$, black solid line), standard HP trend ($\exp(\hat{\tau}_t)$ from (3), blue dotted line), and robust HP trend ($\exp(\hat{\tau}_t^u)$ from (7), red dotted line), all shown in levels (USD Trillions). Lower panel: Implied cyclical component, shown as percentage deviation of actual GDP from each trend ($100 \times (\exp(y_t) - \exp(\hat{\tau}_t))/\exp(\hat{\tau}_t)$), shown in percentage points. Shaded area indicates the COVID-19 period (2020).

signal-to-noise ratio in that period to minimize trend changes is necessary to prevent the output collapse from being misinterpreted as a trend shift.

In our context, such flexibility can be implemented by applying a time-varying smoothing parameter λ_t. The objective function becomes:

$$L = (\boldsymbol{y} - \boldsymbol{\tau})'(\boldsymbol{y} - \boldsymbol{\tau}) + \boldsymbol{\tau}' \boldsymbol{D}_2' \boldsymbol{\Lambda} \boldsymbol{D}_2 \boldsymbol{\tau},$$

with solution

$$\hat{\boldsymbol{\tau}} = (\boldsymbol{I}_T + \boldsymbol{D}_2' \boldsymbol{\Lambda} \boldsymbol{D}_2)^{-1} \boldsymbol{y}, \tag{8}$$

where $\boldsymbol{\Lambda} = \text{diag}(\lambda_1, \ldots, \lambda_{T-2})$ is a $(T-2) \times (T-2)$ diagonal matrix with the (assumed known) time-varying weights.

Each diagonal element λ_k in $\boldsymbol{\Lambda}$ (for $k = 1, \ldots, T-2$) corresponds to the penalty on the k-th trend innovation, ζ_k. Setting a specific element $\lambda_k = 0$ eliminates the smoothness penalty associated with the shock ζ_k. Based on our state-space formulation, this shock ζ_k introduces a 'kink' in the trend's growth rate that begins to affect the trend's level at time $t = k + 2$.

Allowing for a full structural break or level shift in the trend would require setting at least two consecutive elements, λ_k and λ_{k+1}, to zero. This action effectively breaks the smoothness constraint for the trend within the affected window. Furthermore, to force the trend to exactly equal the actual data for a block of periods, say from $t = p$ to $t = q$, requires setting $\lambda_k = 0$ for an even wider window, specifically for indices k from $p - 2$ through q.

This action of setting $\lambda_k = 0$ for one or more periods effectively breaks the smoothness constraint for the trend. Specifically, setting $\lambda_k = 0$ for a single period breaks the slope of the trend, whereas setting it equal to zero for two or more consecutive periods breaks both the slope and the level. Consequently, in the latter case, the estimation becomes equivalent to applying the HP filter separately to the series segments before and after the periods with extreme observations. While this formally allows for modeling structural breaks, it comes at a significant cost. In particular, the trend estimate for the post-pandemic period would be based only on the few observations available after the shock. This would make the resulting trend and cycle decomposition highly uncertain, particularly at the end of the sample where estimates are already less reliable. Given this trade-off, we opt for the additive outlier model, which leverages the full sample history to provide a more stable, albeit counterfactual, post-pandemic trend.

Figure 1 also shows the trend (top panel) and the cycle (bottom panel) when controlling for the COVID-19 pandemic. We see that the trend is estimated to be higher before and after the pandemic year 2020 when we control for the pandemic. Thus, we find that the economy was closer to being cyclically neutral before the pandemic. This is because we do not include the observations in 2020 when minimizing the deviation between actual GDP and trend.

2.3 Variance of the trend estimation error

We now derive the variance–covariance matrix of the trend estimation error. First, we consider the robust HP filter with additive outlier correction. Then, we generalize the result to the case with time-varying smoothing parameters. The standard case, as shown in Schlicht (2005), is nested within both of our results: it is recovered by setting $\boldsymbol{H} = \boldsymbol{I}$ in Proposition 2.1, or by setting $\boldsymbol{\Lambda} = \lambda \boldsymbol{I}$ in Proposition 2.2. We begin with the proposition for the outlier-correction filter.

Proposition 2.1 (Variance with outlier correction)**.** *Consider the robust trend estimator $\hat{\boldsymbol{\tau}}$ defined in (7), which excludes extreme observations from the fit term:*

$$\hat{\boldsymbol{\tau}} = \left(\boldsymbol{H} + \lambda \boldsymbol{D}_2' \boldsymbol{D}_2\right)^{-1} \boldsymbol{H} \boldsymbol{y}$$

where $\boldsymbol{H} = \operatorname{diag}(h_1, \ldots, h_T)$ with $h_t \in \{0, 1\}$ and $\lambda > 0$. Assume the underlying data

are generated by the state-space model in (4)–(5), with the following assumptions on the noise terms:

1. *The observation noise* $\boldsymbol{\varepsilon}$ *is IID:* $Var(\boldsymbol{\varepsilon}) = \sigma_\varepsilon^2 \boldsymbol{I}_T$.
2. *The trend noise vector* $\boldsymbol{\zeta}$ *is a* $(T-2) \times 1$ *vector* $\boldsymbol{\zeta} = (\zeta_1, \ldots, \zeta_{T-2})'$ *which is IID with* $Var(\boldsymbol{\zeta}) = (\sigma_\varepsilon^2/\lambda)\boldsymbol{I}_{T-2}$.

Under these assumptions, the variance–covariance matrix of the estimation error is given by:

$$Var(\hat{\boldsymbol{\tau}} - \boldsymbol{\tau}) = \sigma_\varepsilon^2 \left(\boldsymbol{H} + \lambda \boldsymbol{D}_2' \boldsymbol{D}_2\right)^{-1} \tag{9}$$

Proof. We define the filter matrix $\boldsymbol{A} = (\boldsymbol{H} + \lambda \boldsymbol{D}_2' \boldsymbol{D}_2)^{-1}$. The estimator for the trend is then $\hat{\boldsymbol{\tau}} = \boldsymbol{AHy} = \boldsymbol{AH}(\boldsymbol{\tau} + \boldsymbol{\varepsilon})$, where in the last equality we use (4). The estimation error then becomes $\hat{\boldsymbol{\tau}} - \boldsymbol{\tau} = (\boldsymbol{AH} - \boldsymbol{I})\boldsymbol{\tau} + \boldsymbol{AH\varepsilon}$. Using the matrix identity $\boldsymbol{A}(\boldsymbol{H} + \lambda \boldsymbol{D}_2' \boldsymbol{D}_2) = \boldsymbol{I}$, we have that $(\boldsymbol{AH} - \boldsymbol{I}) = -\lambda \boldsymbol{AD}_2' \boldsymbol{D}_2$. This allows us to rewrite the error using (5) into its fundamental noise components:

$$\hat{\boldsymbol{\tau}} - \boldsymbol{\tau} = -\lambda \boldsymbol{AD}_2' \boldsymbol{\zeta} + \boldsymbol{AH\varepsilon}$$

Since $\boldsymbol{\zeta}$ and $\boldsymbol{\varepsilon}$ are independent, with variances as defined in the proposition, the full variance calculation (see Appendix C.2) shows that the expression collapses neatly back to $\sigma_\varepsilon^2 \boldsymbol{A} = \sigma_\varepsilon^2 \left(\boldsymbol{H} + \lambda \boldsymbol{D}_2' \boldsymbol{D}_2\right)^{-1}$. □

For the variance expression to be well-defined, the matrix $(\boldsymbol{H} + \lambda \boldsymbol{D}_2' \boldsymbol{D}_2)$ must be invertible. As shown in Appendix C.1, this condition holds as long as $\boldsymbol{H}$ contains at least two non-zero diagonal elements (i.e., at least two observations are included in the estimation).

A crucial implication of this formula is that the estimation error variance increases when observations are excluded. Formally, excluding observations means we are in a case where $\boldsymbol{H} \neq \boldsymbol{I}$. The variance matrices can be compared in the Löwner order (the positive semi-definite ordering). Since $\boldsymbol{I}$ and $\boldsymbol{H}$ are diagonal with $h_t \in \{0, 1\}$, the matrix $(\boldsymbol{I} - \boldsymbol{H})$ is positive semi-definite, which we denote as $\boldsymbol{I} \succeq \boldsymbol{H}$. This implies that $(\boldsymbol{I} + \lambda \boldsymbol{D}_2' \boldsymbol{D}_2) \succeq (\boldsymbol{H} + \lambda \boldsymbol{D}_2' \boldsymbol{D}_2)$. For positive definite matrices, this ordering is reversed upon inversion (see, e.g., Horn and Johnson, 2013, Corollary 7.7.4):

$$\left(\boldsymbol{H} + \lambda \boldsymbol{D}_2' \boldsymbol{D}_2\right)^{-1} \succeq \left(\boldsymbol{I} + \lambda \boldsymbol{D}_2' \boldsymbol{D}_2\right)^{-1} \tag{10}$$

Given that $\boldsymbol{H} \neq \boldsymbol{I}$, the matrices are not equal. This confirms the intuition that excluding observations reduces informational content and increases the uncertainty of the trend estimate. This result highlights the method's central trade-off: While failing to exclude extreme observations may introduce bias into the trend estimate,

our derivation shows that the price for excluding observations is an unavoidable increase in the estimate's variance (uncertainty).

Finally, we generalize this variance derivation to the case of a time-varying smoothing parameter.

Proposition 2.2 (Variance with time-varying smoothing parameter)**.** *Consider the generalized trend estimator $\hat{\boldsymbol{\tau}}$ defined in (8), which allows for time-varying smoothness penalties:*

$$\hat{\boldsymbol{\tau}} = \left(\boldsymbol{I}_T + \boldsymbol{D}_2'\boldsymbol{\Lambda}\boldsymbol{D}_2\right)^{-1}\boldsymbol{y}$$

where $\boldsymbol{\Lambda} = \operatorname{diag}(\lambda_1, \ldots, \lambda_{T-2})$ is a diagonal matrix of known, time-varying weights. Assume the underlying data are generated by the state-space model in (4)–(5), with the following assumptions on the noise terms:

1. *The observation noise $\boldsymbol{\varepsilon}$ is IID: $Var(\boldsymbol{\varepsilon}) = \sigma_\varepsilon^2 \boldsymbol{I}_T$.*
2. *The trend noise vector $\boldsymbol{\zeta} = (\zeta_1, \ldots, \zeta_{T-2})'$ is independent but not identically distributed, with a variance structure given by $Var(\boldsymbol{\zeta}) = \sigma_\varepsilon^2 \boldsymbol{\Lambda}^{-1}$.*

Under these assumptions, the variance–covariance matrix of the estimation error is given by:

$$Var(\hat{\boldsymbol{\tau}} - \boldsymbol{\tau}) = \sigma_\varepsilon^2 \left(\boldsymbol{I}_T + \boldsymbol{D}_2'\boldsymbol{\Lambda}\boldsymbol{D}_2\right)^{-1} \tag{11}$$

Proof. The proof follows the same algebraic structure as that in Proposition 2.1, replacing the scalar λ with the diagonal matrix $\boldsymbol{\Lambda}$ to accommodate time-varying penalties. The full derivation is provided in Appendix C.3. □

2.4 Generalization to soft weighting

Although we have defined $\boldsymbol{H}$ as a binary selection matrix ($h_t \in \{0, 1\}$), the framework naturally extends to allow for continuous weights $0 \leq h_t \leq 1$. In a Weighted Least Squares (WLS) context, a weight of $h_t < 1$ corresponds to assuming a higher variance for the observation at time t (specifically, inflating the variance by a factor of $1/h_t$).

This provides a direct theoretical link to the recent work by Morley *et al.* (2023), who handle COVID-19 outliers by estimating variance scale factors. Their approach of "down-weighting" observations via variance inflation is mathematically equivalent to setting diagonal elements of $\boldsymbol{H}$ to values strictly between 0 and 1. Our binary case ($h_t = 0$) represents the limiting case where the variance of the outlier is treated as infinite (zero information).

Consequently, the variance inequality derived in Section 2.3 holds for this general case as well. Any down-weighting of data ($h_t < 1$) implies $\boldsymbol{H} \preceq \boldsymbol{I}$ in the positive

semi-definite sense. This leads to the same conclusion regarding uncertainty as in (10). This shows that there is a continuum between fully trusting the data (Standard HP), partially trusting the data (as in variance scaling approaches), and disregarding the data (our binary exclusion), where lower trust invariably leads to higher posterior uncertainty about the trend estimate.

3 The multivariate filter and its decomposition

The previous section introduced a robust univariate HP filter that accounts for extreme observations. While this improves the trend estimate for aggregate GDP, it does not utilize the information embedded in the composition of GDP. To address this limitation, we extend the framework to a multivariate setting, allowing GDP subcomponents to jointly inform a common underlying trend. The goal is to extract a smooth aggregate trend by optimally combining information from import-adjusted expenditure categories: private consumption, private investment, government expenditure, and exports. Each component may exhibit distinct cyclical dynamics.

3.1 Multivariate model framework

Kozicki (1999) proposes a method for extracting a common trend from multiple economic series, which Harvey (2006, Ch. 7.3.3) develops further. While recent studies explore multivariate decomposition using complex Bayesian unobserved component models (Jahan-Parvar *et al.*, 2024), our approach retains the transparency of the HP framework while integrating time-varying composition effects. We apply this framework to decompose GDP into its main final demand components: private consumption, private investment, government expenditure (consumption and investment), and exports. Since GDP is defined as total final expenditure minus imports, each component must be adjusted for its import content to ensure consistency.

Suppose we have information on the import share associated with each final use. Let ι_C denote the import share of private consumption (C), ι_I for (private and government) investment ($Ip + Ig$), ι_G for government consumption ($G - Ig$), and ι_A for exports (A). Then total imports (Imp) are given by

$$Imp = \iota_C C + \iota_I(Ip + Ig) + \iota_G(G - Ig) + \iota_A A,$$

and GDP is correspondingly given by

$$Y = (1 - \iota_C)C + (1 - \iota_I)(Ip + Ig) + (1 - \iota_G)(G - Ig) + (1 - \iota_A)A.$$

Since the trend is estimated using log-transformed quantities, we apply a log transformation. Applying a first-order Taylor approximation of the log-transformed

Figure 2: Private consumption, government consumption and investments, private investments, and exports, all in shares of GDP

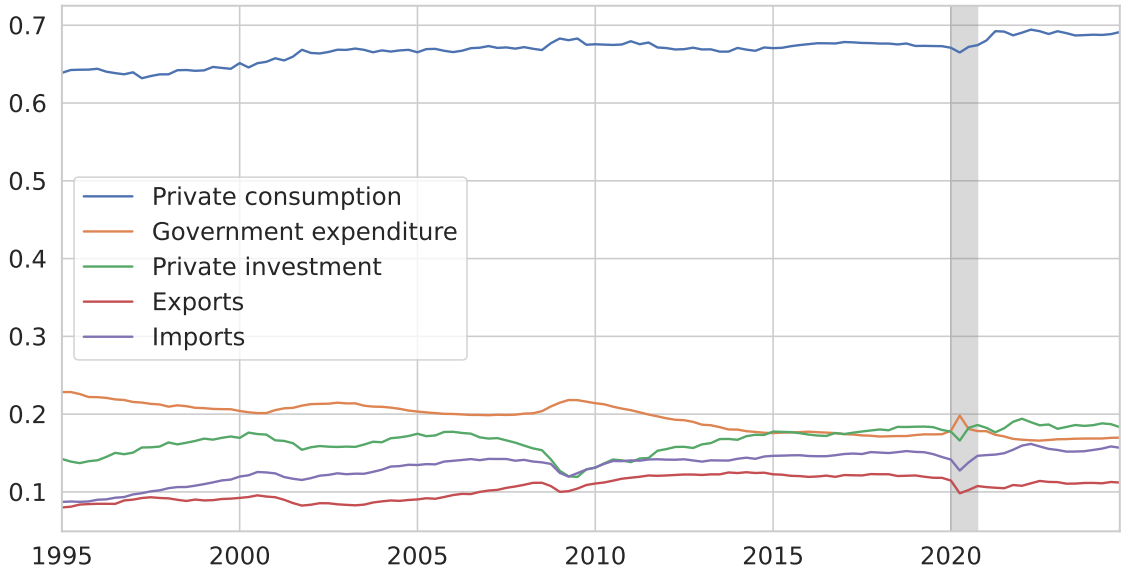

Notes: National account figures given as share of GDP; Imp_t/Y_t (blue), C_t/Y_t (red), G_t/Y_t) (green), Ip_t/Y_t (red), and A_t/Y_t (purple). For more details, see Appendix A. Source: The U.S. Bureau of Economic Analysis.

GDP identity, evaluated around the components' average budget shares over the sample period, gives

$$y \approx \mu_C c + \mu_{Ip} ip + \mu_G g + \mu_A a, \tag{12}$$

where $y = \log(Y)$, $\mu_C = \frac{1}{T}\sum_{t=1}^{T}(1-\iota_{C,t})C_t/Y_t$ denotes the average import-adjusted budget share for consumption over the sample period, and $c = \log((1-\iota_C)C) - \log(\mu_{C,t})$ is the import-adjusted log component of private consumption, rescaled by its own budget share $\mu_{C,t}$ to be consistent with the level of GDP. This budget share used for rescaling may be time-varying (hence the subscript t). Analogous definitions apply to the average shares and the rescaled log components for private investments and exports. For government expenditure (G), the calculation of the average share μ_G accounts for the fact that government consumption $(G - I_g)$ and government investment (I_g) have different import intensities, using their respective import shares (ι_G and ι_I): $\mu_G = \frac{1}{T}\sum_{t=1}^{T}\left[(1-\iota_{G,t})(G_t - I_g)/Y_t + (1-\iota_{I,t})I_{g,t}/Y_t\right]$. The import-adjusted log component of government expenditure is given by $g = \log((1-\iota_G)G) + \log((1-\iota_I)Ig) - \log(\mu_{C,t})$. Note that since the average budget shares are calculated from level data before averaging, they sum to exactly one: $\sum_i \mu_i = 1$.

Following Kozicki (1999) and Harvey (2006, Chapt. 7.3.3), we assume that these log-transformed components follow a common trend:

$$\boldsymbol{y}_t = \boldsymbol{z}\tau_t + \boldsymbol{\varepsilon}_t, \tag{13}$$

where $\boldsymbol{y}_t = (c_t, ip_t, g_t, a_t)'$ is the vector of import-adjusted, rescaled log final use categories, $\boldsymbol{z} = (1, 1, 1, 1)'$ implies the same loadings on the common trend τ_t, and $\boldsymbol{\varepsilon}_t$ is a vector of stationary deviations with zero mean and covariance matrix $\boldsymbol{\Sigma}_\varepsilon$. No assumption of normality or absence of autocorrelation (white noise) is made for $\boldsymbol{\varepsilon}_t$ at this stage.

3.2 Estimation of the common trend

Since all components are assumed to share a common trend, any linear combination with constant weights will also follow that trend. Weighting each component by its average import-adjusted budget share approximately reproduces the univariate GDP decomposition, see (12).

Alternatively, we can choose weights that minimize the variance of the cyclical component, following the approach in Harvey (2006, Ch. 7.3.3). The objective is to extract the smoothest and most persistent trend that is common across all GDP components. By constructing a weighted average of the components that minimizes the variance of the weighted cyclical deviations, we are by definition isolating this common, stable underlying trend. We do this by using the weights:

$$\hat{\boldsymbol{\omega}} = \left(\boldsymbol{z}'\hat{\boldsymbol{\Sigma}}_\varepsilon^{-1}\boldsymbol{z}\right)^{-1}\boldsymbol{z}'\hat{\boldsymbol{\Sigma}}_\varepsilon^{-1}, \tag{14}$$

where $\hat{\boldsymbol{\Sigma}}_\varepsilon$ is an estimate of the variance-covariance matrix $\boldsymbol{\Sigma}_\varepsilon$.

This weighting scheme assigns higher weights to components with lower volatility or negative correlation with other components, thereby emphasizing stable contributors to the common trend. This weighting scheme corresponds to the Generalized Least Squares (GLS) solution for estimating the common factor τ_t in the model (13).

Let

$$\tilde{y}_t = \hat{\boldsymbol{\omega}}'\boldsymbol{y}_t \tag{15}$$

be the weighted average of the components that minimizes the cyclical part of (13). The trend and cycle are estimated from the weighted series $\tilde{y}_t$, which follows the same structure as the univariate model:

$$\tilde{y}_t = \tau_t + \tilde{\varepsilon}_t, \quad \tilde{\varepsilon}_t = \hat{\boldsymbol{\omega}}'\boldsymbol{\varepsilon}_t, \tag{16}$$

where $\tilde{\varepsilon}_t = \hat{\boldsymbol{\omega}}\boldsymbol{\varepsilon}_t$ is the resulting weighted-average (filtered) cyclical component. The trend τ_t and filtered cycle $\hat{\varepsilon}_t$ can then be estimated by using the same methods as in the univariate case. The estimation proceeds iteratively, starting from an initial trend estimate (e.g., based on GDP alone as in Section 2). We alternate between (i) estimating the weight vector $\hat{\boldsymbol{\omega}}$ using (14), and (ii) updating the trend from the weighted series $\tilde{y}_t$ using the HP filter. This process is repeated until convergence.

Figure 3: Import shares by expenditure category, United States

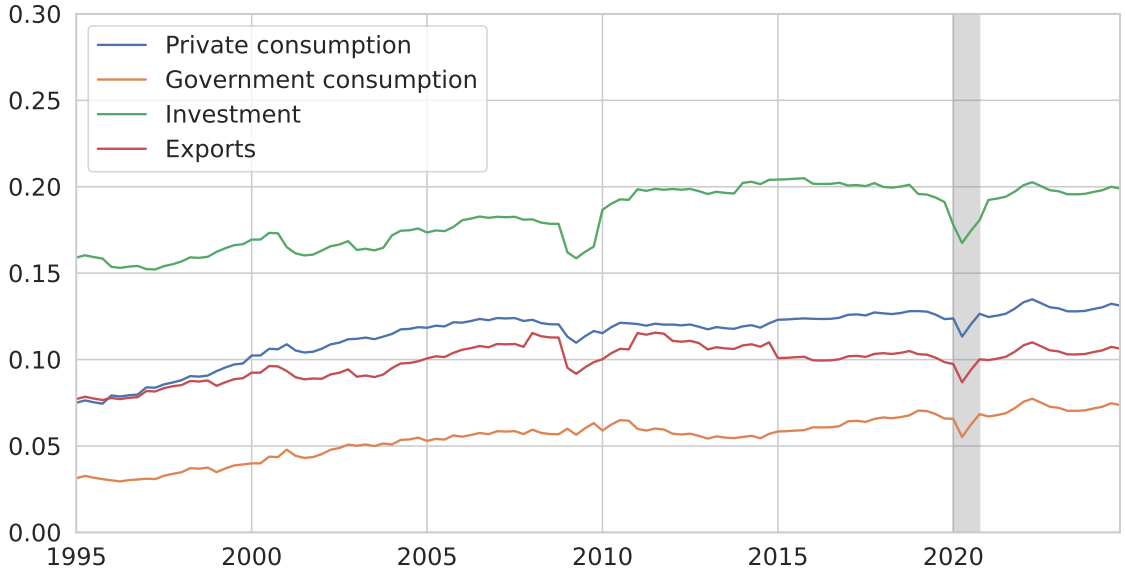

Notes: Import shares, derived in Appendix B, based on the input-output system (industry-industry) from OECD.

Convergence is defined as the point where the Euclidean norm of the change in weights between iterations satisfies

$$||\hat{\boldsymbol{\omega}}^{(k)} - \hat{\boldsymbol{\omega}}^{(k-1)}|| < \epsilon,$$

with a tolerance threshold $\epsilon = 10^{-6}$.

3.3 Import-adjusted components and budget shares

Figure 2 shows the final demand components and imports as shares of GDP. Although the shares are relatively stable over time, they seem to have a slightly positive or negative underlying trend.

The import shares (ι_C, ι_I, ι_G, and ι_A) are derived from the input-output tables in the national accounts, as detailed in Appendix B. Figure 3 shows that import intensity varies substantially across components: investment (which includes both private and public) has the highest import share, while government consumption has the lowest. These patterns are consistent with international evidence on import content across expenditure categories (e.g., Bussière *et al.*, 2013).

3.3.1 Constant budget shares

Figure 4 shows the estimated trend and cycle for U.S. GDP using the multivariate HP filter with fixed import-adjusted budget shares, i.e $\mu_{C,t} = \mu_C$ for c, and analogous for ip, g, and a. Under the assumption of constant shares, the estimated trend suggests that GDP remained above potential from 2012 onward, excluding the COVID-19 period. This pattern likely reflects misspecification due to the as-

Figure 4: Decomposition of U.S. GDP using the multivariate filter with fixed budget shares

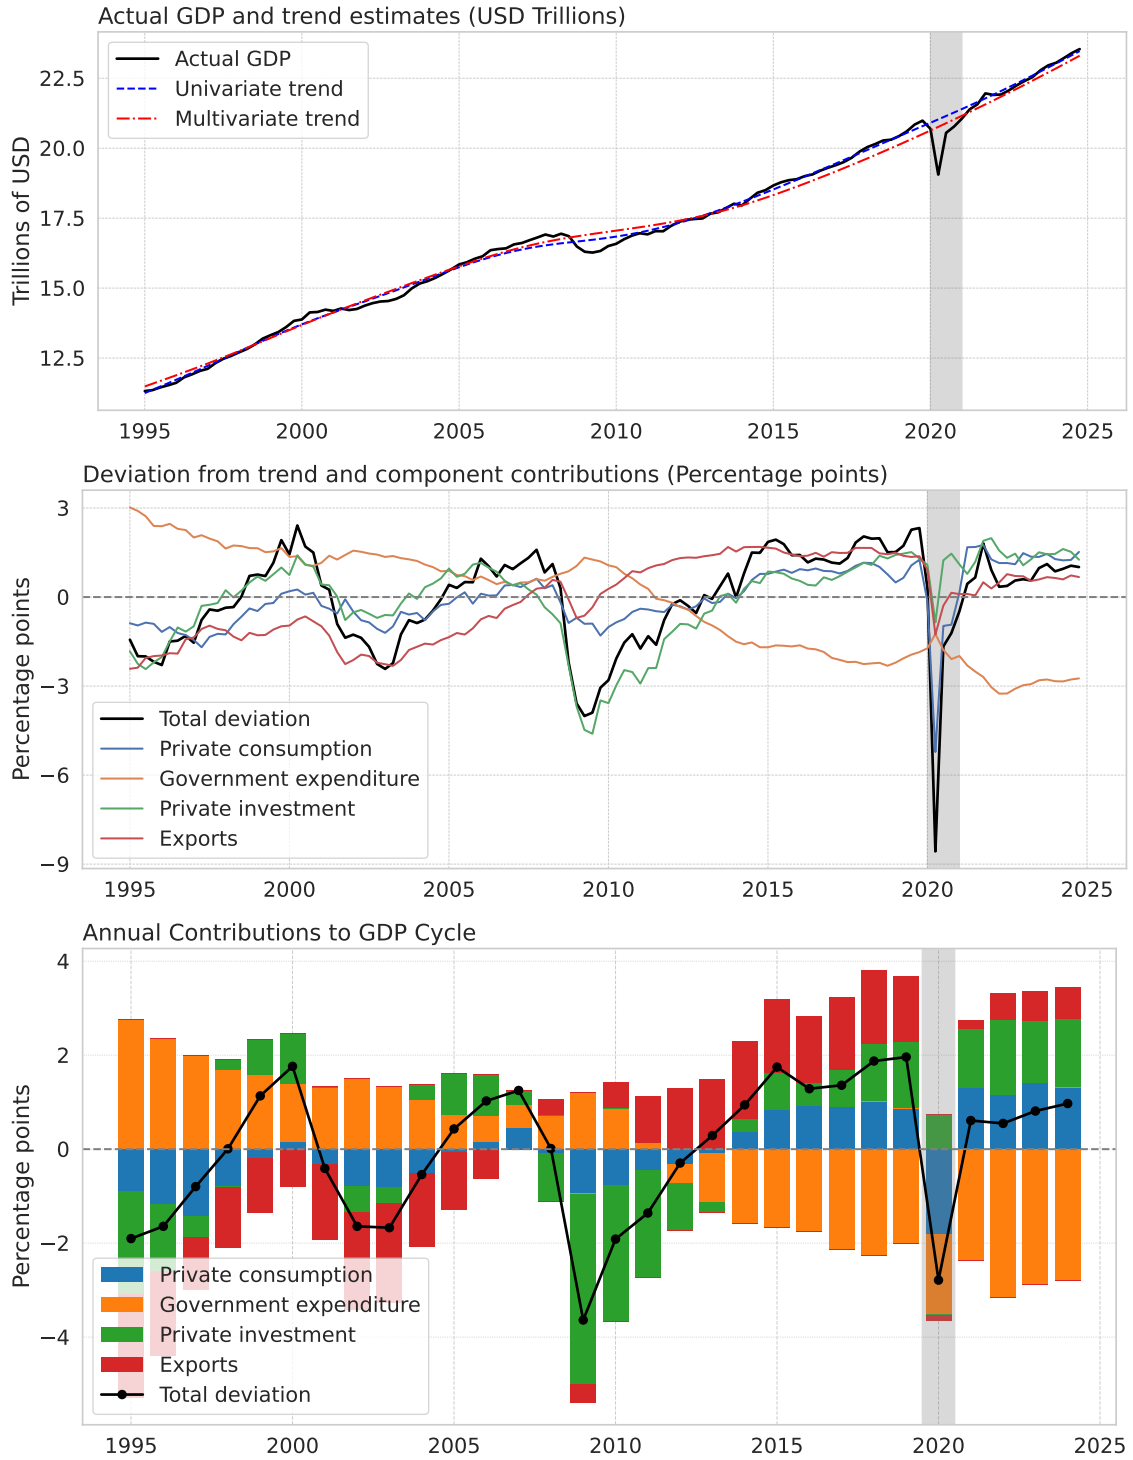

Notes: Upper panel: Actual GDP ($\exp(y_t)$, black solid line), robust univariate HP trend ($\exp(\hat{\tau}_t)$ from (7), blue dotted line), and multivariate trend with constant shares (red dotted line), all shown in levels (USD Trillions). Lower panel: Implied cyclical component, shown as percentage deviation of actual GDP from each trend ($100 \times (\exp(y_t) - \exp(\hat{\tau}_t))/\exp(\hat{\tau}_t)$), shown in percentage points. (Formula for consumption contribution: $100 \cdot \exp(\mu_C)(\exp(c_t) - \exp(\hat{\tau}_t))/\exp(\hat{\tau}_t)$.) Shaded area indicates the COVID-19 period (2020).

sumption of constant shares, rather than genuine cyclical dynamics. For example, the share of government expenditure in GDP declined during this period. As shown in the middle and lower panels of Figure 4, its contribution to the cycle is positive until 2011 and negative thereafter. For exports, we have the opposite picture, as it contributes negatively to the cycle each year until 2006 and positively each year from 2010 when we exclude the year 2020.

3.3.2 Time-varying budget shares

Figure 5 shows the results of estimating the multivariate HP filter with time-varying budget shares. The estimation procedure is designed to extract a smooth common trend from the GDP components with time-varying import-adjusted shares. The process proceeds in two steps.

Step 1: Estimate the time-varying shares ($\mu_{i,t}$). Let $s_{c,t} = C_t(1 - \iota_{C,t})/Y_t$ be the observed import-adjusted budget share for private consumption, and $s_{ip,t}$, $s_{g,t}$, and $s_{a,t}$ be the observed import-adjusted shares of private investment, government expenditure, and exports, respectively. We model the observed import-adjusted share $s_{i,t}$ for each component $i \in c, ip, g, a$ as a noisy signal around a latent smooth process $\mu_{i,t}$ assumed to follow a random walk without drift:

$$\begin{aligned} s_{i,t} &= \mu_{i,t} + \nu_{i,t}, \quad \nu_{i,t} \sim N(0, \sigma^2_{\nu,i}), \\ \mu_{i,t} &= \mu_{i,t-1} + \eta_{i,t}, \quad \eta_{i,t} \sim N(0, \sigma^2_{\eta,i}), \end{aligned}$$

where the signal-to-noise ratio $\lambda_s = \sigma^2_{\nu,i}/\sigma^2_{\eta,i}$ is assumed to be common across all components. We estimate the vector $\boldsymbol{\mu}_i$ using a smoothness-penalized least squares approach, applying the COVID-19 correction matrix $\boldsymbol{H}$ (defined in Section 2.2):

$$\hat{\boldsymbol{\mu}}_i = (\boldsymbol{H} + \lambda_s \boldsymbol{D}_1' \boldsymbol{D}_1)^{-1} \boldsymbol{H} \boldsymbol{s}_i, \quad i \in \{c, ip, g, a\}, \tag{17}$$

where $\boldsymbol{s}_i = (s_{i,1}, s_{i,2}, \ldots, s_{i,T})'$ is the time series of the observed budget shares of component i, and $\boldsymbol{D}_1$ is the first-order difference matrix of dimension $(T-1) \times T$:

$$\boldsymbol{D}_1 = \begin{bmatrix} -1 & 1 & 0 & \ldots & 0 \\ 0 & -1 & 1 & \ldots & 0 \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ 0 & \ldots & 0 & -1 & 1 \end{bmatrix}.$$

This estimation procedure ensures that the smoothed shares maintain the adding-up constraint ($\sum_i \hat{\boldsymbol{\mu}}_i = \mathbf{1}_T$) by construction, as stated in the following proposition.

Proposition 3.1 (Adding-up constraint)**.** *Given that the observed import-adjusted shares satisfy the adding-up condition ($\sum_i \boldsymbol{s}_i = \mathbf{1}_T$), and assuming the matrix ($\boldsymbol{H} + \lambda_s \boldsymbol{D}_1' \boldsymbol{D}_1$) is invertible, the estimated smoothed shares from (17) also satisfy the*

Figure 5: Decomposition of U.S. GDP using the multivariate filter with time-varying budget shares

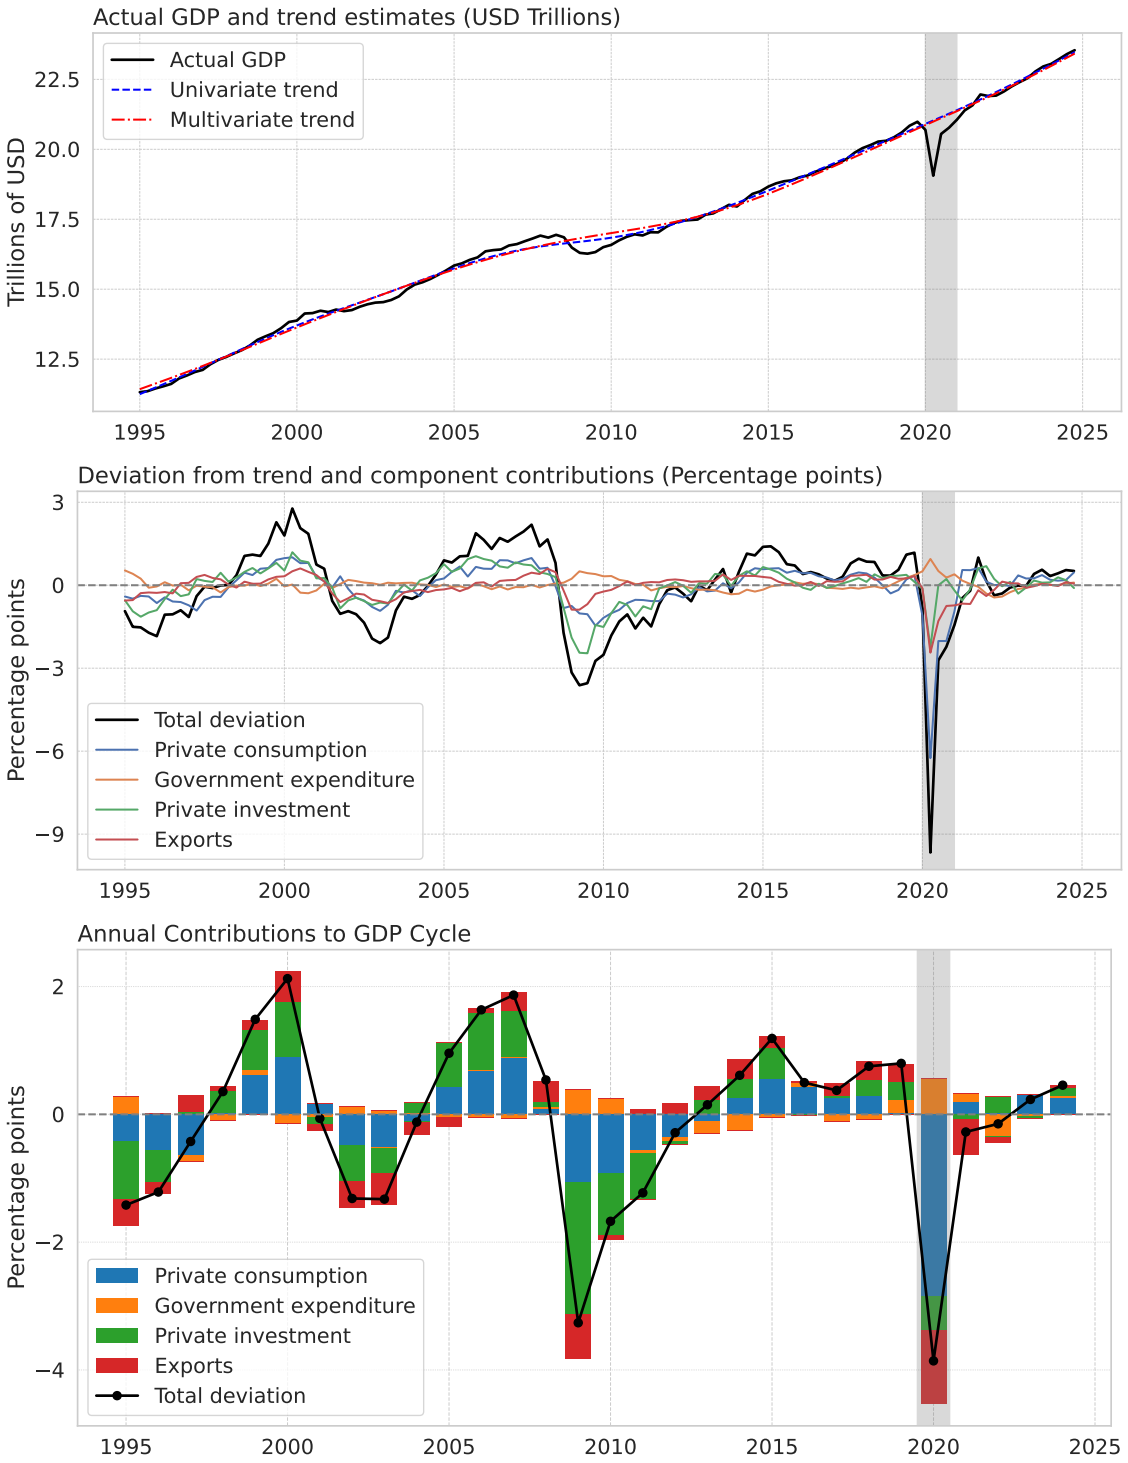

Notes: Upper panel: Actual GDP ($\exp(y_t)$, black solid line), robust univariate HP trend ($\exp(\hat{\tau}_t^u)$ from (7), blue dotted line), and multivariate trend with time-varying shares (red dotted line), all shown in levels (USD Trillions). Lower panel: Implied cyclical component, shown as percentage deviation of actual GDP from each trend ($100 \times (\exp(y_t) - \exp(\hat{\tau}_t))/\exp(\hat{\tau}_t)$), shown in percentage points. (Formula for consumption contribution: $100 \cdot \exp(\mu_{C,t})(\exp(c_t) - \exp(\hat{\tau}_t))/\exp(\hat{\tau}_t)$.) Shaded area indicates the COVID-19 period (2020).

adding-up condition:

$$\sum_i \hat{\boldsymbol{\mu}}_i = \mathbf{1}_T$$

Proof. Let $\boldsymbol{A} = (\boldsymbol{H} + \lambda_s \boldsymbol{D}_1' \boldsymbol{D}_1)^{-1}$. Summing the estimated shares from (17) across all components i gives:

$$\sum_i \hat{\boldsymbol{\mu}}_i = \sum_i (\boldsymbol{A}\boldsymbol{H}\boldsymbol{s}_i) = \boldsymbol{A}\boldsymbol{H}\left(\sum_i \boldsymbol{s}_i\right) = \boldsymbol{A}\boldsymbol{H}\mathbf{1}_T,$$

where we use the assumption that $\sum_i \boldsymbol{s}_i = \mathbf{1}_T$. We now show that $\boldsymbol{A}\boldsymbol{H}\mathbf{1}_T = \mathbf{1}_T$. When $\boldsymbol{A}$ is invertible, this is equivalent to showing $\boldsymbol{H}\mathbf{1}_T = \boldsymbol{A}^{-1}\mathbf{1}_T$. Substituting the definition of $\boldsymbol{A}^{-1}$ gives:

$$\boldsymbol{A}^{-1}\mathbf{1}_T = (\boldsymbol{H} + \lambda_s \boldsymbol{D}_1' \boldsymbol{D}_1)\mathbf{1}_T = \boldsymbol{H}\mathbf{1}_T + \lambda_s \boldsymbol{D}_1'(\boldsymbol{D}_1\mathbf{1}_T)$$

The penalty vanishes because the first-difference of a constant vector $\mathbf{1}_T$ is zero: $\boldsymbol{D}_1\mathbf{1}_T = \mathbf{0}$. Thus, $\boldsymbol{A}^{-1}\mathbf{1}_T = \boldsymbol{H}\mathbf{1}_T$. This proves the identity, as it implies:

$$\sum_i \hat{\boldsymbol{\mu}}_i = \boldsymbol{A}(\boldsymbol{H}\mathbf{1}_T) = \boldsymbol{A}(\boldsymbol{A}^{-1}\mathbf{1}_T) = (\boldsymbol{A}\boldsymbol{A}^{-1})\mathbf{1}_T = \mathbf{1}_T \qquad \square$$

Step 2: Estimate the common trend (τ_t). Using the smoothed budget shares $\hat{\boldsymbol{\mu}}_{i,t}$, we define the rescaled log-components as $c_t = \log((1-\iota_C)C) - \log(\hat{\mu}_{C,t})$, and similarly for ip_t, g_t, a_t. This forms a new vector $\boldsymbol{y}_t = (c_t, ip_t, g_t, a_t)'$ where the components' rescaling weights (the smoothed budget shares) are time-varying. We then apply the same iterative multivariate procedure described in (13)—(16) to this new $\boldsymbol{y}_t$ vector to estimate the single, common trend τ_t.

Incorporating time-varying budget shares leads to a trend estimate that closely aligns with the univariate GDP trend, while preserving the ability to decompose cyclical contributions from individual components. We set the smoothing parameter for estimating the process of the smoothed budget shares to $\lambda_s = 40$. This choice is theoretically consistent with the conventional HP parameter for quarterly GDP ($\lambda = 1600$). Specifically, since the shares are smoothed with a first-order filter appropriate for an $I(1)$ process, we follow the methodology used in Harvey and Delle Monache (2009) for comparing filters of different orders. This involves setting the signal-to-noise ratio for our filter equal to the square root of the ratio for the second-order HP filter, which implies $\lambda_s = \sqrt{1600} = 40$ (see Harvey and Delle Monache, 2009, Sec. 4.3). In Appendix D, we conduct the same analysis with other values of λ_s and show that the qualitative results do not rely on this specific value.

The results in Figure 5 support the interpretation of government expenditure as countercyclical. Following the 2008 financial crisis and again in 2020, government expenditure remained above trend, while GDP and the other components fell below trend. (It should be noted that observations from 2020 are excluded from the

Table 1: Average budget shares and estimated multivariate weights (percent)

	Univariate	Multivariate	
	Budget share	Time-invariant	Time-varying
Private consumption	59.1	88.3	42.0
Government expenditure	18.0	11.0	56.0
Private investment	13.5	-1.5	-0.8
Exports	9.4	1.2	2.7

Notes: 'Univariate Budget share' refers to the sample average of the import-adjusted shares (μ_i). The 'Weights' ($\boldsymbol{\omega}$) refer to the estimated optimal weights from (14). The near-zero or negative weights for private investment and exports in the time-varying model reflect their high cyclical volatility, which the filter optimally down-weights to identify the smoothest common trend.

estimation of the trend.)

The next section compares the performance of the multivariate filter under fixed and time-varying budget shares, highlighting the empirical implications of our methodological choices.

3.4 Empirical results

Table 1 shows the estimated optimal weights, $\hat{\boldsymbol{\omega}}$, from (14) for the multivariate specifications, compared with the average budget shares in the first column. In the model with time-varying budget shares, government expenditure receives the highest weight of 56 percent, while private consumption is given a weight of 42 percent. Private investment and exports receive near-zero or slightly negative weights. This is an important result of the optimization: components with high cyclical volatility, such as private investment, are given a low weight to identify the smoothest common trend. A negative weight can arise if a component is highly volatile and/or negatively correlated with the smoother components that dominate the common trend.

From Figure 5, we also see that the largest contribution to cyclical fluctuations comes from private investment. Fluctuations in household consumption also contribute, while exports contribute to a much lesser extent.

To evaluate the empirical performance of the proposed filtering methods, we compare the statistical properties of the trend and cycle components across the four filter specifications shown in Figure 6. The upper panel shows the smoothness of the trend, measured as the variance of the trend's second difference ($Var(\Delta^2\hat{\tau}_t)$). The lower panel shows the variance of the implied cyclical component (the gap between actual GDP and the estimated trend, $Var(y_t - \hat{\tau}_t)$). All results are shown relative to the robust univariate filter (the second bar in Figure 6), which serves as our benchmark.

The upper panel illustrates the primary benefit of our approach. Correcting for the COVID-19 shock (moving from the standard univariate filter to the robust univariate filter) reduces the trend variance substantially. Furthermore, incorporating

Figure 6: Relative trend smoothness and cyclical variance

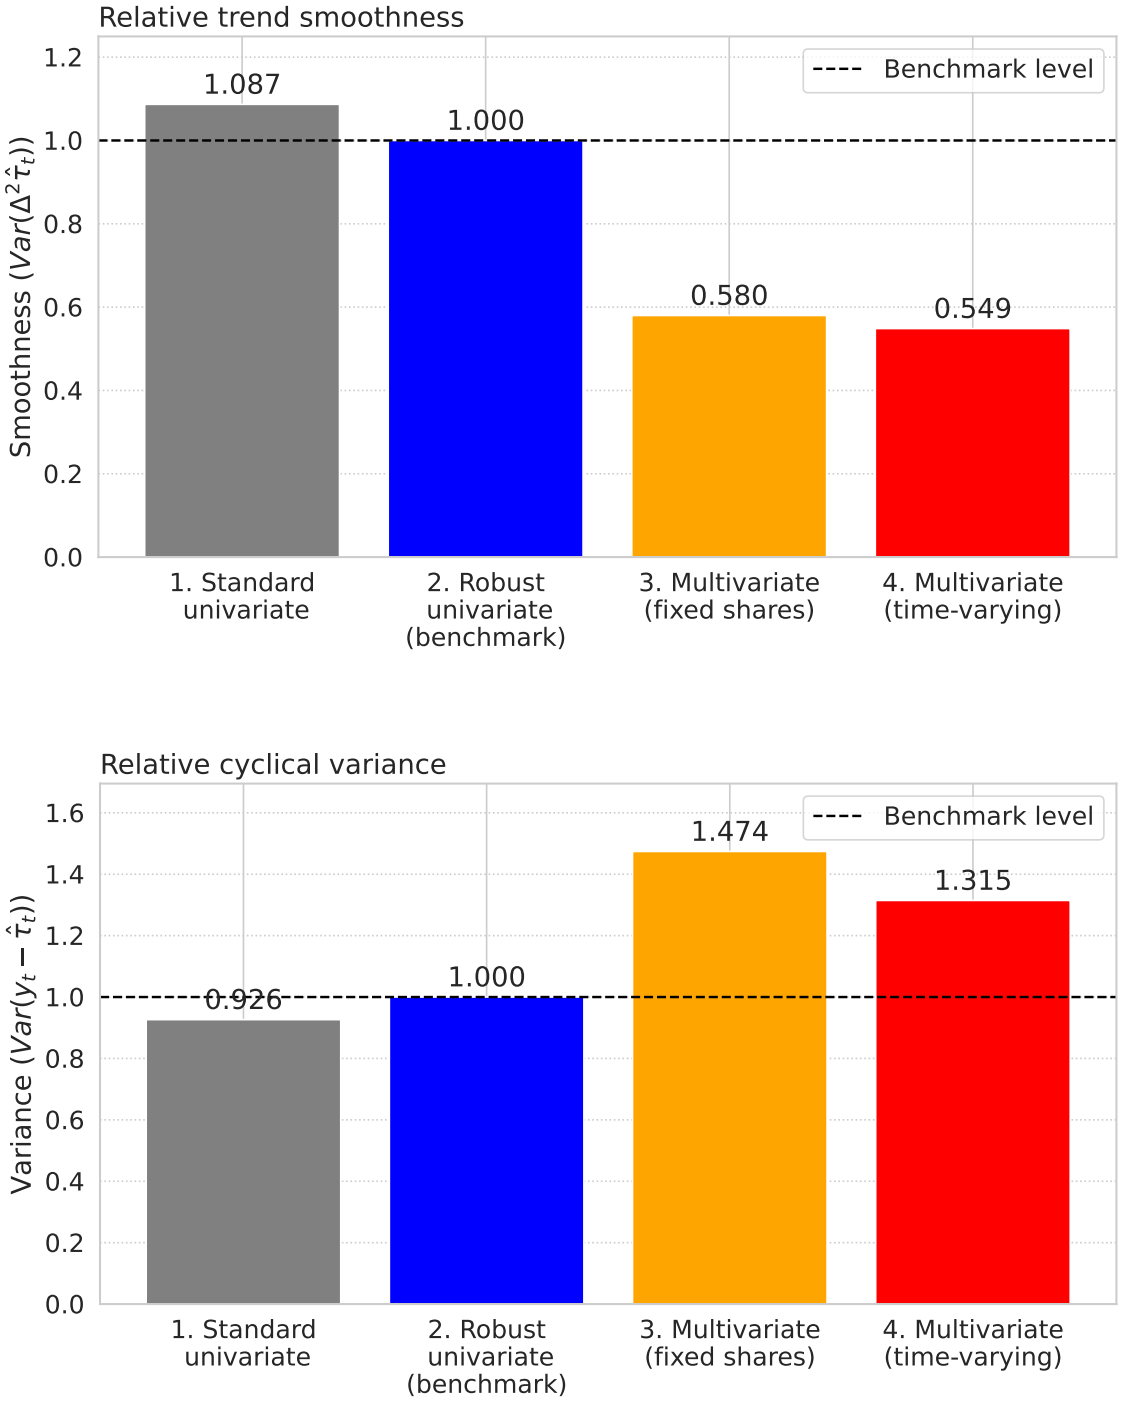

Notes: The figure displays the variance of the trend's second difference ($Var(\Delta^2\hat{\tau})$, upper panel) and the variance of the implied cyclical component ($Var(y_t - \hat{\tau}_t)$, lower panel), each relative to the robust univariate filter (bar 2), which is set to 1. In the upper panel, a lower value implies a smoother trend. The four specifications are as follows: (1) standard univariate HP, (2) robust univariate HP (used as benchmark), (3) multivariate (fixed shares), and (4) multivariate (time-varying).

component information (the multivariate filters with both fixed and time-varying budget shares) yields an even smoother trend, with our final time-varying model producing the most stable trend overall.

The lower panel, which shows the variance of the implied cycle $(y_t - \hat{\tau}_t)$, reveals a more nuanced picture. The robust and multivariate filters produce a cycle with a higher variance than the standard HP filter. This is an expected result. The standard filter (bar 1) artificially lowers its cyclical variance by distorting the trend to partially absorb the extreme COVID-19 shock. Our robust methods (bars 2–4) correctly classify this shock as part of the cycle, which naturally increases its overall measured variance.

More importantly, the multivariate filters also show a higher implied cyclical variance (bars 3 and 4) than the robust univariate benchmark (bar 2). A higher cyclical variance in the multivariate specifications does not indicate inferior performance. Rather, it reflects the improved separation between trend and cycle: a smoother trend implies that more short-term fluctuations are correctly attributed to the cycle. Because the multivariate trend is smoother and tracks the underlying long-run path more accurately, it no longer absorbs as much of the short-term fluctuations from the original GDP series (y_t). Instead, these fluctuations are correctly allocated to the implied cyclical component $(y_t - \hat{\tau}_t)$, resulting in a more precise, albeit more volatile, measure of the business cycle.

It is important to distinguish between the implied cycle $(y_t - \hat{\tau}_t)$, which reflects deviations from the official GDP series, and the filtered cycle $(\tilde{y}_t - \hat{\tau}_t)$, which is derived from the component-weighted series. The latter is smoother by construction and better suited for structural analysis. Since the smoothing parameter $\lambda = 1600$ is used for all specifications, the filter imposes a fixed ratio between the filtered cycle variance and the trend's second-difference variance $(Var(\tilde{y}_t - \hat{\tau}_t)/Var(\Delta^2\hat{\tau}_t) = \lambda)$ for the specific series being filtered (i.e., y_t or $\tilde{y}_t$). As the upper panel shows, our multivariate models produce a substantially smoother trend (a lower $Var(\Delta^2\hat{\tau}_t)$). Therefore. it follows directly from the fixed λ that the variance of the filtered cyclical component $\tilde{y}_t - \hat{\tau}_t$ is also progressively reduced. Thus, the gain from the multivariate approach is twofold: it yields a smoother aggregate trend $(\hat{\tau}_t)$ and simultaneously identifies a less volatile underlying cyclical component for the component-weighted series $(\tilde{y}_t)$, even though it more accurately captures the full volatility of the business cycle in the original GDP data.

While our analysis holds $\lambda = 1600$ constant for a consistent comparison, the results support the view that the choice of λ is context-dependent (see, e.g., Harvey and Jaeger, 1993). By construction, the multivariate input series $\tilde{y}_t$ is smoother than the original GDP series y_t. Consequently, applying the same fixed $\lambda = 1600$ results in a multivariate trend with a substantially lower smoothness variance $(Var(\Delta^2\hat{\tau}_t))$, which is approximately 55-58 percent of the robust univariate benchmark. To en-

force the same degree of trend smoothness (i.e., the same $Var(\Delta^2\hat{\tau}_t)$) across specifications, λ should be reduced to approximately 900 for the multivariate filters. This observation reinforces the argument that the choice of smoothing parameter is context-dependent. Nevertheless, we maintain the conventional fixed-λ approach in our main analysis to ensure transparency and comparability with the standard univariate filter.

4 Conclusions

This paper has demonstrated two methodological enhancements to the classical Hodrick–Prescott filter for decomposing GDP into trend and cycle: (i) a robust univariate filter that accounts for extreme observations such as the COVID-19 shock, and (ii) a multivariate filter that leverages time-varying, import-adjusted budget shares of GDP components. Applying this framework to U.S. data, we find that private investment is the largest single contributor to cyclical fluctuations, while government expenditure exhibits a strong and persistent counter-cyclical pattern. By enabling a more granular and robust decomposition of the output gap, the proposed framework provides policymakers with a data-driven tool to identify the sources of cyclical fluctuations and design targeted stabilization measures.

While the empirical application relies on final, revised data, the multivariate framework is also well-suited for real-time analysis, where preliminary data are often subject to revision. In such applications, initial data releases are often noisy, and certain components are subject to large revisions. Private investment is a prime example, being both highly volatile and frequently revised (see e.g. Helliesen *et al.*, 2022). A powerful feature of our methodology is that it endogenously adapts to these data characteristics. By assigning lower weights to volatile components, such as investment, while giving greater weight to more stable components such as government expenditure, the filter automatically produces a more reliable initial trend estimate. By assigning lower weights to volatile and revision-prone components, the filter endogenously stabilizes the trend estimate near the sample endpoint, thereby mitigating the well-known "end-point problem" of the HP filter (see, e.g., Orphanides and Norden, 2002). While other solutions address this by extending the series with model-based forecasts (Kaiser and Maravall, 1999) or adjusting the one-sided filter directly (Wolf *et al.*, 2024), our framework offers an alternative, endogenous stabilization by automatically down-weighting the most volatile and revision-prone data components.

The findings carry important implications for macroeconomic monitoring. The dominance of private investment in the business cycle underscores the importance of tracking financial conditions and corporate sentiment. Furthermore, the consistently counter-cyclical pattern of government expenditure suggests that automatic fiscal

stabilizers have played a significant role in the U.S. economy during the sample period.

References

Akaike H (1980). "Seasonal adjustment by a bayesian modeling." *Journal of Time Series Analysis*, **1**(1), 1–13. doi:10.1111/j.1467-9892.1980.tb00296.x.

Beveridge S, Nelson CR (1981). "A new approach to decomposition of economic time series into permanent and transitory components with particular attention to measurement of the 'business cycle'." *Journal of Monetary Economics*, **7**(2), 151–174. doi:10.1016/0304-3932(81)90040-4.

Bussière M, Callegari G, Ghironi F, Sestieri G, Yamano N (2013). "Estimating trade elasticities: Demand composition and the trade collapse of 2008-2009." *American Economic Journal: Macroeconomics*, **5**(3), 118–151. doi:10.1257/mac.5.3.118.

Christiano LJ, Fitzgerald TJ (2003). "The Band Pass Filter*." *International Economic Review*, **44**(2), 435–465. doi:10.1111/1468-2354.t01-1-00076.

Danthine JP, Girardin M (1989). "Business cycles in Switzerland." *European Economic Review*, **33**(1), 31–50. doi:10.1016/0014-2921(89)90035-4.

de Jong RM, Sakarya N (2016). "The econometrics of the Hodrick-Prescott filter." *Review of Economics and Statistics*, **98**(2), 310–317. doi:10.1162/rest_a_00523.

Dermoune A, Djehiche B, Rahmania N (2009). "Multivariate Extension of the Hodrick-Prescott Filter-Optimality and Characterization." *Studies in Nonlinear Dynamics & Econometrics*, **13**(3). doi:10.2202/1558-3708.1656.

Fox AJ (1972). "Outliers in Time Series." *Journal of the Royal Statistical Society Series B: Statistical Methodology*, **34**(3), 350–363. doi:10.1111/j.2517-6161.1972.tb00912.x.

Franke R, Kukacka J, Sacht S (2024). "Is the Hamilton regression filter really superior to Hodrick–Prescott detrending?" *Macroeconomic Dynamics*, **29**. ISSN 14698056. doi:10.1017/S136510052400018X.

Hamilton JD (2018). "Why You Should Never Use the Hodrick-Prescott Filter." *The Review of Economics and Statistics*, **100**(5), 831–843. doi:10.1162/rest_a_00706.

Harvey AC (1985). "Trends and Cycles in Macroeconomic Time Series." *Journal of Business & Economic Statistics*, **3**(3), 216. doi:10.2307/1391592.

Harvey AC (2006). "Chapter 7 Forecasting with Unobserved Components Time Series Models." In G Elliott, C Granger, A Timmermann (eds.), "Handbook of Economic Forecasting," volume 1, pp. 327–412. doi:10.1016/S1574-0706(05)01007-4.

Harvey AC, Delle Monache D (2009). "Computing the mean square error of unobserved components extracted by misspecified time series models." *Journal of Economic Dynamics and Control*, **33**(2), 283–295. doi:10.1016/j.jedc.2008.05.007.

Harvey AC, Jaeger A (1993). "Detrending, stylized facts and the business cycle." *Journal of Applied Econometrics*, **8**(3), 231–247. doi:10.1002/jae.3950080302.

Harvey AC, Trimbur T (2008). "Trend Estimation and the Hodrick-Prescott Filter." *Journal of the Japan Statistical Society*, **38**(1), 41–49. doi:10.14490/jjss.38.41.

Helliesen MK, Hungnes H, Skjerpen T (2022). "Revisions in the Norwegian National Accounts: accuracy, unbiasedness and efficiency in preliminary figures." *Empirical Economics*, **62**(3), 1079–1121. doi:10.1007/s00181-021-02065-9.

Hodrick RJ, Prescott EC (1981). "Post-War U.S. Business Cycles: An Empirical Investigation." *Discussion Papers 451, from Northwestern University, Center for Mathematical Studies in Economics and Management Science*. URL https://econpapers.repec.org/paper/nwucmsems/451.htm.

Hodrick RJ, Prescott EC (1997). "Postwar U.S. Business Cycles: An Empirical Investigation." *Journal of Money, Credit and Banking*, **29**(1), 1. doi:10.2307/2953682.

Horn RA, Johnson CR (2013). *Matrix Analysis*. Cambridge iversity ress, 2nd edition. URL https://www.cambridge.org/9780521548236.

Jahan-Parvar MR, Knipp C, Szerszeń PJ (2024). "Trend-Cycle Decomposition and Forecasting Using Bayesian Multivariate Unobserved Components." *Finance and Economics Discussion Series*, (2024-100), 1–28. doi:10.17016/feds.2024.100.

Kaiser R, Maravall A (1999). "Estimation of the business cycle: A modified Hodrick-Prescott filter." *Spanish Economic Review*, **1**(2), 175–206. doi:10.1007/s101080050008.

Kamber G, Morley J, Wong B (2025). "Trend-cycle decomposition in the presence of large shocks." *Journal of Economic Dynamics and Control*, **173**. doi:10.1016/j.jedc.2025.105066.

Kozicki S (1999). "Multivariate detrending under common trend restrictions: Implications for business cycle research." *Journal of Economic Dynamics and Control*, **23**(7), 997–1028. doi:10.1016/S0165-1889(98)00048-7.

Maranzano P, Pelagatti M (2025). "A Hodrick–Prescott filter with automatically selected breaks." *Economic Modelling*, **150**, 107132. doi:10.1016/j.econmod.2025.107132.

Mei Z, Phillips PC, Shi Z (2024). "The boosted Hodrick-Prescott filter is more general than you might think." *Journal of Applied Econometrics*, **39**(7), 1260–1281. doi:10.1002/jae.3086.

Meyer C (2000). *Matrix Analysis and Applied Linear Algebra*. SIAM, Philadelphia. doi:10.1137/1.9780898719512.

Miller RE, Blair PD (2009). *Input-Output Analysis*. Cambridge University Press, New York. doi:10.1017/CBO9780511626982.

Morley J, Rodríguez-Palenzuela D, Sun Y, Wong B (2023). "Estimating the euro area output gap using multivariate information and addressing the COVID-19 pandemic." *European Economic Review*, **153**. doi:10.1016/j.euroecorev.2023.104385.

Orphanides A, Norden Sv (2002). "The Unreliability of Output-Gap Estimates in Real Time." *Review of Economics and Statistics*, **84**(4), 569–583. doi:10.1162/003465302760556422.

Phillips PCB, Shi Z (2021). "Boosting: Why you can use the HP filter." *International Economic Review*, **62**(2), 521–570. doi:10.1111/iere.12495.

Ravn MO, Uhlig H (2002). "On Adjusting the Hodrick-Prescott Filter for the Frequency of Observations." *Review of Economics and Statistics*, **84**(2), 371–376. doi:10.1162/003465302317411604.

Schlicht E (2005). "Estimating the Smoothing Parameter in the So-called Hodrick-Prescott Filter." *Journal of the Japan Statistical Society*, **35**(1), 99–119. doi:10.14490/jjss.35.99.

Schüler YS (2024). "Filtering economic time series: On the cyclical properties of Hamilton's regression filter and the Hodrick-Prescott filter." *Review of Economic Dynamics*, **54**. doi:10.1016/j.red.2024.101237.

Wolf E, Mokinski F, Schüler Y (2024). "On Adjusting the One-Sided Hodrick–Prescott Filter." *Journal of Money, Credit and Banking*. doi:10.1111/jmcb.13240.

A Data description

Table A.1: National accounts variables used in the analysis (chained 2017 dollars, seasonally adjusted)

Variable	Description	Macrobond code
Y	Gross Domestic Product	usnaac0169
Imp	Imports	usnaac0210
C	Private Consumption	usnaac0170
G	Government Expenditure	usnaac0213
A	Exports	usnaac0207
I_p	Private Investment ($I - I_g$)	usnaac0191
I_g	Government Investment	usnaac0217 + usnaac0220 + usnaac0223

The U.S. Bureau of Economic Analysis (BEA) https://www.bea.gov/itable/nipa-underlying-detail and Macrobond https://www.macrobond.com.

The seasonally adjusted national accounts series are sourced from Macrobond, based on data from the U.S. Bureau of Economic Analysis (BEA), specifically Table 1.5.6 in the National Income and Product Accounts (NIPA), Quarterly release. Chained dollar series are computed as the product of chain-type quantity indices and the corresponding current-dollar values (in 2017 dollars), divided by 100. Due to the chain-weighting procedure, these series are generally not additive across components.

Government investment (I_g) consists of gross investment in national defense (usnaac0217), federal nondefense (usnaac0220), and state and local government (usnaac0223). Data for these series are available from 2002 onwards. Prior to 2002, government investment is imputed as 19 percent of the sum of government consumption and investment, consistent with later observations. The residual is attributed to government consumption.

B Import shares

To ensure consistency in the multivariate decomposition of GDP, we adjust each final demand component for its import content. This appendix outlines the methodology used to derive these import shares, which are essential for constructing import-adjusted budget weights. First, we use the OECD's Input-Output (I-O) tables to calculate the direct and indirect import content for each final demand category. Second, since the I-O tables are not available for the most recent years, we explain how these shares are extrapolated. Finally, we describe an adjustment procedure to ensure that the sum of the import-adjusted components is consistent with the observed GDP data from the national accounts.

Table B.1: Input-Output table with domestic production and import

		Intermediate			Final demand			Total
		Ind. 1	...	Ind. S	Use 1	...	Use K	
Domestic prod.	Ind. 1 ⋮ Ind. S	$\boldsymbol{Z}^d$			$\boldsymbol{F}^d$			$\boldsymbol{x}$
Import	Ind. 1 ⋮ Ind. S	$\boldsymbol{Z}^m$			$\boldsymbol{F}^m$			$\boldsymbol{m}$

Notes: Illustration of input-output industry-industry table: $\boldsymbol{Z}^d$ and $\boldsymbol{Z}^m$ $(S \times S)$ are inter-industry transactions and imported-industry transactions; $\boldsymbol{F}^d$ and $\boldsymbol{F}^m$ $(S \times K)$ are final use matrices, $\boldsymbol{x}$ $(S \times 1)$ is the vector of total output; and $\boldsymbol{m}$ $(S \times 1)$ is the vector of import. S is the number of industries, and K is the number of final use categories.

B.1 Import shares based on the Input-Output system

We use data from the OECD Input-Output Database (https://www.oecd.org/en/data/datasets/input-output-tables.html), which provides detailed information on deliveries of input factors across industries and to final demand categories. At the time of writing, input-output tables for the United States are available from 1995 to 2020. These tables distinguish between inputs sourced domestically and those sourced through imports, as illustrated in Table B.1.

Let S denote the number of industries.[1] Define the matrix $\boldsymbol{Z}^d$ as an $S \times S$ matrix, where each element represents the value of domestically produced intermediate inputs from the industry in the row that is used by the industry in the column during the reference year. Similarly, $\boldsymbol{Z}^m$ is an $S \times S$ matrix where each entry indicates the value of imported intermediate inputs (to one industry (row) from one industry (column)).

The matrices $\boldsymbol{F}^d$ and $\boldsymbol{F}^m$ have dimensions $S \times K$, where K is the number of final demand categories (e.g., household consumption, government consumption, exports, gross fixed capital formation, changes in inventories, etc.). Each element of $\boldsymbol{F}^d$ represents the final demand for domestically produced goods and services, whereas the entries in $\boldsymbol{F}^m$ capture the corresponding final demand for imports.

From this, total industry output, $\boldsymbol{x}$, and total imports, $\boldsymbol{m}$, are defined as:

$$\boldsymbol{x} = \boldsymbol{Z}^d \mathbf{1}_S + \boldsymbol{F}^d \mathbf{1}_K, \tag{B.1}$$

$$\boldsymbol{m} = \boldsymbol{Z}^m \mathbf{1}_S + \boldsymbol{F}^m \mathbf{1}_K, \tag{B.2}$$

where $\mathbf{1}_N$ is a vector of N ones.

To analyze the production structure, we normalize the intermediate flows by

[1] In the current OECD input-output tables, $S = 45$, whereas Bussière *et al.* (2013) used 48 industries.

total output to obtain the technical coefficient matrices:

$$\boldsymbol{A}^d = \boldsymbol{Z}^d \left(\text{diag}(\boldsymbol{x})\right)^{-1}, \tag{B.3}$$

$$\boldsymbol{A}^m = \boldsymbol{Z}^m \left(\text{diag}(\boldsymbol{x})\right)^{-1}, \tag{B.4}$$

where $\text{diag}(\boldsymbol{x})$ is an $S \times S$ diagonal matrix with the elements of $\boldsymbol{x}$ along its diagonal. The coefficient a_{ij}^d in $\boldsymbol{A}^d$ represents the value of domestic inputs from industry i required per unit of output in industry j, whereas a_{ij}^m in $\boldsymbol{A}^m$ represents the corresponding imported input requirement. These technical coefficient matrices are assumed to be fixed if demand (and thus output) changes; see, e.g., Miller and Blair (2009).

By using the fixed technical coefficients from (B.4), total output in (B.1) can alternatively be expressed as:

$$\boldsymbol{x} = \left(\boldsymbol{I}_S - \boldsymbol{A}^d\right)^{-1} \boldsymbol{F}^d \boldsymbol{1}_K, \tag{B.5}$$

where $\boldsymbol{I}_S$ is the $S \times S$ identity matrix, and $\left(\boldsymbol{I}_S - \boldsymbol{A}^d\right)^{-1}$ is the Leontief inverse. The Leontief inverse is reflecting the total (direct and indirect) input requirements per unit of final demand when also taking into account that output in one industry is used as input in another industry (cf. Miller and Blair, 2009, p. 21).

By combining (B.2) and (B.4) with (B.5), we can express total imports as a function of final demand:

$$\boldsymbol{m} = \underbrace{\boldsymbol{F}^m \boldsymbol{1}_K}_{\text{Direct imports}} + \underbrace{\boldsymbol{A}^m \left(\boldsymbol{I}_S - \boldsymbol{A}^d\right)^{-1} \boldsymbol{F}^d \boldsymbol{1}_K}_{\text{Indirect imports}}. \tag{B.6}$$

Equation (B.6) decomposes total imports into two components: (i) direct imports, which are final goods and services purchased from abroad, and (ii) indirect imports, which are foreign inputs embedded in domestically produced goods and services.

Our goal is to find the total import share based on the input-output (IO) system, ι_k^{IO}, for each final demand category k. This share is the sum of the direct and indirect import content per unit of expenditure in that category. Following the methodology in Bussière *et al.* (2013), these are calculated as:

$$\iota_k^{IO} = \iota_k^{dir} + \iota_k^{ind}, \tag{B.7}$$

$$\text{where} \quad \iota_k^{dir} = \frac{\boldsymbol{1}_S' \boldsymbol{F}^m \boldsymbol{e}_k}{\boldsymbol{1}_S' \left(\boldsymbol{F}^d + \boldsymbol{F}^m\right) \boldsymbol{e}_k}, \tag{B.8}$$

$$\text{and} \quad \iota_k^{ind} = \frac{\boldsymbol{1}_S' \boldsymbol{A}^m \left(\boldsymbol{I}_S - \boldsymbol{A}^d\right)^{-1} \boldsymbol{F}^d \boldsymbol{e}_k}{\boldsymbol{1}_S' \left(\boldsymbol{F}^d + \boldsymbol{F}^m\right) \boldsymbol{e}_k}. \tag{B.9}$$

Here, $\boldsymbol{e}_k$ is a selection vector of length K with a 1 in position k and zeros elsewhere. The numerator in (B.8) sums the direct imports for category k, while the

Figure B.1: Import shares by expenditure category, United States

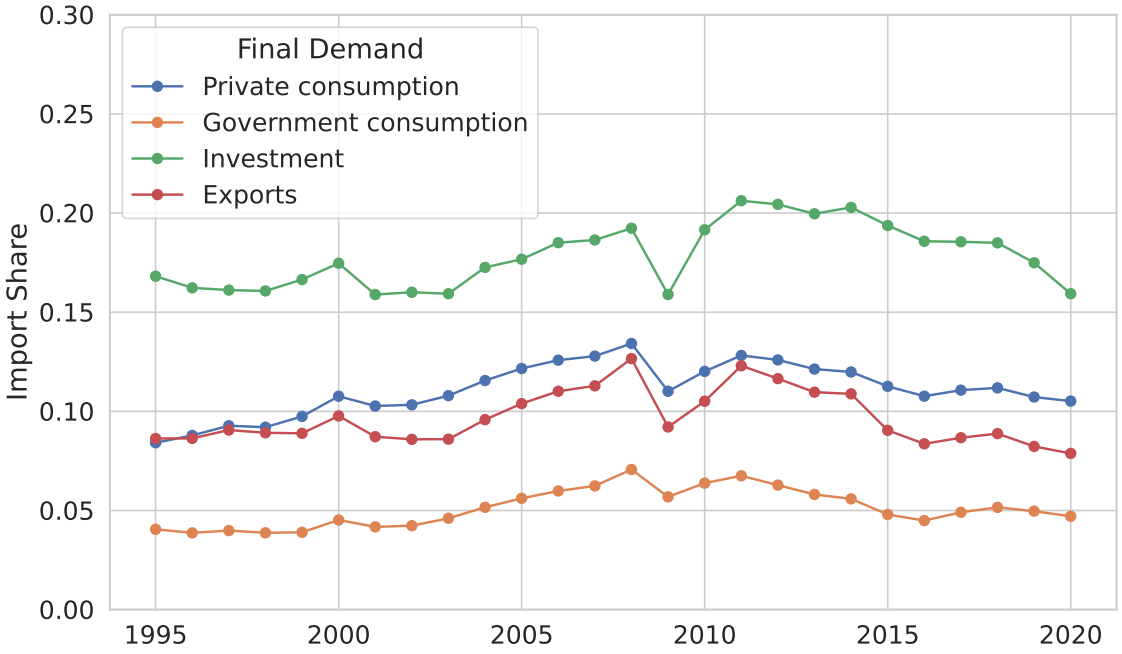

Notes: Import shares based on the input-output system; $\boldsymbol{\iota}_C^{IO}$ (blue), $\boldsymbol{\iota}_G^{IO}$ (orange), $\boldsymbol{\iota}_I^{IO}$ (green), and $\boldsymbol{\iota}_A^{IO}$ (red).

numerator in (B.9) sums the indirect imports. The common denominator is the total expenditure in category k.

Figure B.1 shows the import shares for household consumption (including NPISH, that is, the final consumption expenditure of non-profit institutions serving households), government consumption, investment (both private and government/public) and export. These are calculated for the years 1995-2020.

B.2 Import shares for 2021–2024

We need to extend the import shares to the years 2021–2024. Normally, it would be natural to assume that the import shares in these years were approximately equal to the last year of the calculated import shares. However, since this is from the year 2020, a year that was heavily affected by the COVID19 pandemic, we instead chose to set them equal to the import shares for 2019.

This approach ensures that import shares for years beyond 2020 are not influenced by the atypical trade patterns observed during the pandemic. By anchoring the post-2020 import shares to the 2019 values, we aim to provide a more stable and representative basis for subsequent analyses. While econometric methods could be used to forecast these shares, our approach provides a simple and transparent benchmark that avoids introducing additional uncertainty from model specification or forecast error.

B.3 Further adjustments of the import shares

To reconcile the discrepancies between the observed imports and those implied by the import-adjusted demand components, we introduce two alternative time-varying adjustment factors: $\delta_{1,t}$ aligns the model with actual imports, while $\delta_{2,t}$ ensures consistency with observed GDP. These factors serve to align model-generated imports and GDP with actual observed data.

Let the preliminary model-generated imports ($\widetilde{Imp}$) and model-generated GDP ($\widetilde{Y}$) be given by the following:

$$\begin{aligned}\widetilde{Imp} &= \iota_C^{IO} \cdot C + \iota_G^{IO} \cdot (G - I_g) + \iota_I^{IO} \cdot (I_p + I_g) + \iota_A^{IO} \cdot A \\ \widetilde{Y} &= (1 - \iota_C^{IO}) \cdot C + (1 - \iota_G^{IO}) \cdot (G - I_g) + (1 - \iota_I^{IO}) \cdot (I_p + I_g) + (1 - \iota_A^{IO}) \cdot A\end{aligned}$$

where

- C is private consumption,
- G is government expenditures (consumption $G - I_g$ plus investments I_g),
- I_p is private investment,
- I_g is public investment,
- A is exports, and
- ι_k ($k = C, G, I, A$) are the import shares for each category based on the imput-output system.

The adjustment factors are then defined as:

$$\delta_1 = \frac{Imp - \widetilde{Imp}}{C + (G - I_g) + (I_p + I_g) + A} \tag{B.10}$$

$$\delta_2 = -\frac{Y - \widetilde{Y}}{C + (G - I_g) + (I_p + I_g) + A} \tag{B.11}$$

where:

- Imp is the observed imports, and
- Y is the observed GDP.

By computing these adjustment factors, we can refine our import estimates to better align with observed data, enhancing the accuracy of our economic analyses.

If the general budget equation ($Y + Imp = C + G + Ip + A$) holds, then we would have $\delta_1 = \delta_2$. But due to chaining, the general budget equation will not hold exactly. Thus, we must choose between ensuring consistency with observed imports (using δ_1) or with observed GDP (using δ_2). Given that the primary objective of

Figure B.2: Adjustments factors for the import shares, percentage points

Notes: Adjustment vectors given by (B.10) in blue and (B.11) in orange, measured in percentage points. The shaded area indicates the COVID-19 period (2020).

this paper is to decompose the output gap of the official GDP series, ensuring that our import-adjusted components sum precisely to this observed GDP is the logical priority. We therefore choose the latter, and set $\delta = \delta_2$. The adjusted import shares are then:

$$\iota_k = \iota_k^{IO} + \delta, \quad k = C, G, I, A.$$

In Figure B.2, the two correction factors are plotted. We see that they are largely the same, even though they deviate somewhat in the first years of the sample. The notable upward trend in the factors, particularly after 2015, suggests that the import content of final demand may have been increasing at a rate not fully captured by the periodically updated input-output coefficients.

C Additional derivation

This appendix provides formal derivations and proofs for the filter matrices and variance expressions used in the paper. Appendix C.1 establishes the rank conditions required for the penalized least squares estimators to be well-defined. Appendix C.2 derives the variance–covariance matrix of the trend estimation error for the robust HP filter, detailing the proof of Proposition 2.1. Appendix C.3 presents the proof of Proposition 2.2, which generalizes the variance expression to the case of a time-varying smoothing parameter.

C.1 Rank conditions for penalized filter matrices

This appendix proves the rank conditions for the two penalized least squares filter matrices used in the paper. We present a general lemma and then show how it applies to the first-difference ($n = 1$) and second-difference ($n = 2$) cases.

Lemma C.1 (General Rank Condition for n-th Difference)**.** *Let $\boldsymbol{D}_n$ denote the $(T-n)\times T$ n-th difference matrix, and $\boldsymbol{H} = \text{diag}(h_1, \ldots, h_T)$ be a diagonal selection matrix with binary entries $h_t \in \{0, 1\}$ indicating which observations are included in the fit term. For any $\lambda > 0$, the matrix*

$$\boldsymbol{M}_n = \boldsymbol{H} + \lambda \boldsymbol{D}_n' \boldsymbol{D}_n$$

is positive definite (and hence invertible) if and only if $\boldsymbol{H}$ includes at least n distinct indices $t_1, \ldots, t_n$ such that $h_{t_1} = \cdots = h_{t_n} = 1$.

Proof. We show that for any non-zero vector $\boldsymbol{x} \in \mathbb{R}^T$, $\boldsymbol{x}'\boldsymbol{M}_n\boldsymbol{x} > 0$. The quadratic form is $\boldsymbol{x}'\boldsymbol{M}_n\boldsymbol{x} = \sum_{t=1}^T h_t x_t^2 + \lambda \|\boldsymbol{D}_n \boldsymbol{x}\|^2$. If $\boldsymbol{x}'\boldsymbol{M}_n\boldsymbol{x} = 0$, both non-negative terms must be zero: (i) $\sum_{t=1}^T h_t x_t^2 = 0$ and (ii) $\boldsymbol{D}_n\boldsymbol{x} = \boldsymbol{0}$.

From (ii), $\boldsymbol{x}$ must be in the null space of $\boldsymbol{D}_n$. The null space of the n-th difference matrix is the space of polynomials of degree $n-1$:

$$x_t = a_0 + a_1 t + \cdots + a_{n-1} t^{n-1}, \quad a_i \in \mathbb{R}.$$

This is a vector space with n coefficients (dimensions).

From (i), x_t must be zero for all t where $h_t = 1$. By the lemma's assumption, there are at least n distinct indices $t_1, \ldots, t_n$ where this holds. This gives a homogeneous system of n linear equations for the n unknown coefficients $a_0, \ldots, a_{n-1}$:

$$\begin{aligned} a_0 + a_1 t_1 + a_2 t_1^2 + \cdots + a_{n-1} t_1^{n-1} &= 0 \\ a_0 + a_1 t_2 + a_2 t_2^2 + \cdots + a_{n-1} t_2^{n-1} &= 0 \\ &\vdots \\ a_0 + a_1 t_n + a_2 t_n^2 + \cdots + a_{n-1} t_n^{n-1} &= 0 \end{aligned}$$

We can write this as $\boldsymbol{T}\boldsymbol{a} = \boldsymbol{0}$, where

$$\boldsymbol{T} = \begin{bmatrix} 1 & t_1 & t_1^2 & \ldots & t_1^{n-1} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & t_n & t_n^2 & \ldots & t_n^{n-1} \end{bmatrix} \text{ and } \boldsymbol{a} = \begin{bmatrix} a_0 \\ \vdots \\ a_{n-1} \end{bmatrix}.$$

A standard result in linear algebra is that this matrix $\boldsymbol{T}$ (a Vandermonde matrix) is invertible (has full rank) if and only if all t_i are distinct (see, e.g., Meyer, 2000, Ch. 6.1). This follows from the fact that the determinant of a Vandermonde matrix

with entries $t_1, \ldots, t_n \in \{1, \ldots, T\}$ is given by

$$\det(\boldsymbol{T}) = \prod_{1 \le i < j \le n} (t_j - t_i),$$

which is non-zero if and only if all t_i are distinct.

By the lemma's assumption, $t_1, \ldots, t_n$ are distinct, so $\boldsymbol{T}$ is invertible. Pre-multiplying by $\boldsymbol{T}^{-1}$ gives:

$$\boldsymbol{T}^{-1}\boldsymbol{T}\boldsymbol{a} = \boldsymbol{T}^{-1}\boldsymbol{0} \implies \boldsymbol{I}\boldsymbol{a} = \boldsymbol{0} \implies \boldsymbol{a} = \boldsymbol{0}$$

This implies $a_0 = \cdots = a_{n-1} = 0$, and thus $\boldsymbol{x} = \boldsymbol{0}$, which contradicts the initial assumption that $\boldsymbol{x} \neq \boldsymbol{0}$. It follows that $\boldsymbol{M}_n$ is positive definite and invertible. □

Remark C.1 (Case $n = 1$: First-difference filter for budget shares)**.** *For the filter $\boldsymbol{M}_1 = \boldsymbol{H} + \lambda_s \boldsymbol{D}_1' \boldsymbol{D}_1$ used to smooth the budget shares (Section 3.3.2), we have $n = 1$. The null space of $\boldsymbol{D}_1$ consists of constant vectors ($x_t = a_0$). Lemma C.1 implies that $\boldsymbol{M}_1$ is invertible if $\boldsymbol{H}$ contains at least $n = 1$ non-zero diagonal element. This condition is used in the proof of Proposition 3.1.*

Remark C.2 (Case $n = 2$: Second-difference (HP) filter)**.** *For the robust HP filter $\boldsymbol{M}_2 = \boldsymbol{H} + \lambda \boldsymbol{D}_2' \boldsymbol{D}_2$ (Section 2.2), we have $n = 2$. The null space of $\boldsymbol{D}_2$ consists of linear trends ($x_t = a_0 + a_1 t$). Lemma C.1 implies that $\boldsymbol{M}_2$ is invertible if $\boldsymbol{H}$ contains at least $n = 2$ distinct non-zero diagonal elements. This condition is referenced in Proposition 2.1 and the main text.*

C.2 Variance of the trend estimation error

Let the estimation error be

$$\hat{\boldsymbol{\tau}} - \boldsymbol{\tau} = -\lambda \boldsymbol{A}\boldsymbol{D}_2'\boldsymbol{\zeta} + \boldsymbol{A}\boldsymbol{H}\boldsymbol{\varepsilon}$$

The variance of the error, assuming independence between $\boldsymbol{\zeta}$ and $\boldsymbol{\varepsilon}$, and substitute $\mathrm{Var}(\boldsymbol{\zeta}) = \frac{1}{\lambda}\sigma_\varepsilon^2 \boldsymbol{I}$ and $\mathrm{Var}(\boldsymbol{\varepsilon}) = \sigma_\varepsilon^2 \boldsymbol{I}$:

$$\begin{aligned}
\mathrm{Var}(\hat{\boldsymbol{\tau}} - \boldsymbol{\tau}) &= \mathrm{Var}(-\lambda \boldsymbol{A}\boldsymbol{D}_2'\boldsymbol{\zeta}) + \mathrm{Var}(\boldsymbol{A}\boldsymbol{H}\boldsymbol{\varepsilon}) \\
&= (-\lambda \boldsymbol{A}\boldsymbol{D}_2')\,\mathrm{Var}(\boldsymbol{\zeta})\,(-\lambda \boldsymbol{A}\boldsymbol{D}_2')' + (\boldsymbol{A}\boldsymbol{H})\,\mathrm{Var}(\boldsymbol{\varepsilon})\,(\boldsymbol{A}\boldsymbol{H})' \\
&= (-\lambda \boldsymbol{A}\boldsymbol{D}_2') \left(\frac{1}{\lambda}\sigma_\varepsilon^2 \boldsymbol{I}\right) (-\lambda \boldsymbol{D}_2 \boldsymbol{A}') + (\boldsymbol{A}\boldsymbol{H}) \left(\sigma_\varepsilon^2 \boldsymbol{I}\right) (\boldsymbol{H}'\boldsymbol{A}')
\end{aligned}$$

We simplify the expression by factoring out the scalars. Since $\boldsymbol{H}$ (diagonal) and $\boldsymbol{D}_2'\boldsymbol{D}_2$ are symmetric, $\boldsymbol{A}$ is also symmetric. Thus $\boldsymbol{A}' = \boldsymbol{A}$ and $\boldsymbol{H}' = \boldsymbol{H}$. Furthermore, we use the property $\boldsymbol{H}^2 = \boldsymbol{H}$ (since $\boldsymbol{H}$ is a diagonal matrix of 0s and

1s):

$$\begin{aligned}\text{Var}(\hat{\boldsymbol{\tau}} - \boldsymbol{\tau}) &= \lambda\sigma_\varepsilon^2 \boldsymbol{A}\boldsymbol{D}_2'\boldsymbol{D}_2\boldsymbol{A} + \sigma_\varepsilon^2 \boldsymbol{A}\boldsymbol{H}\boldsymbol{H}\boldsymbol{A} \\ &= \lambda\sigma_\varepsilon^2 \boldsymbol{A}\boldsymbol{D}_2'\boldsymbol{D}_2\boldsymbol{A} + \sigma_\varepsilon^2 \boldsymbol{A}\boldsymbol{H}\boldsymbol{A}\end{aligned}$$

Factor out common terms $\sigma_\varepsilon^2\boldsymbol{A}$ and $\boldsymbol{A}$, and apply the definition $\boldsymbol{A}^{-1} = (\boldsymbol{H} + \lambda\boldsymbol{D}_2'\boldsymbol{D}_2)$:

$$\begin{aligned}\text{Var}(\hat{\boldsymbol{\tau}} - \boldsymbol{\tau}) &= \sigma_\varepsilon^2 \boldsymbol{A}\left(\lambda\boldsymbol{D}_2'\boldsymbol{D}_2 + \boldsymbol{H}\right)\boldsymbol{A} \\ &= \sigma_\varepsilon^2 \boldsymbol{A}\left(\boldsymbol{A}^{-1}\right)\boldsymbol{A} \\ &= \sigma_\varepsilon^2 \boldsymbol{A}\end{aligned}$$

C.3 Proof of Proposition 2.2: Variance with time-varying λ

We seek the variance of the estimation error $\hat{\boldsymbol{\tau}} - \boldsymbol{\tau}$. First, define the filter matrix $\boldsymbol{A}$ for notational simplicity:

$$\boldsymbol{A} = (\boldsymbol{I}_T + \boldsymbol{D}_2'\boldsymbol{\Lambda}\boldsymbol{D}_2)^{-1}$$

The trend estimate is $\hat{\boldsymbol{\tau}} = \boldsymbol{A}\boldsymbol{y}$. We can now write the estimation error by substituting the state-space model $\boldsymbol{y} = \boldsymbol{\tau} + \boldsymbol{\varepsilon}$:

$$\hat{\boldsymbol{\tau}} - \boldsymbol{\tau} = \boldsymbol{A}\boldsymbol{y} - \boldsymbol{\tau} = \boldsymbol{A}(\boldsymbol{\tau} + \boldsymbol{\varepsilon}) - \boldsymbol{\tau} = (\boldsymbol{A} - \boldsymbol{I}_T)\boldsymbol{\tau} + \boldsymbol{A}\boldsymbol{\varepsilon}$$

We use the matrix identity derived from the definition of $\boldsymbol{A}$. Since $\boldsymbol{A}(\boldsymbol{I}_T + \boldsymbol{D}_2'\boldsymbol{\Lambda}\boldsymbol{D}_2) = \boldsymbol{I}_T$, it follows that:

$$\boldsymbol{A} - \boldsymbol{I}_T = -\boldsymbol{A}\boldsymbol{D}_2'\boldsymbol{\Lambda}\boldsymbol{D}_2$$

Substituting this back into the error equation and use $\boldsymbol{D}_2\boldsymbol{\tau} = \boldsymbol{\zeta}$ to express the error in terms of the fundamental noise components $\boldsymbol{\varepsilon}$ and $\boldsymbol{\zeta}$:

$$\begin{aligned}\hat{\boldsymbol{\tau}} - \boldsymbol{\tau} &= (-\boldsymbol{A}\boldsymbol{D}_2'\boldsymbol{\Lambda}\boldsymbol{D}_2)\,\boldsymbol{\tau} + \boldsymbol{A}\boldsymbol{\varepsilon} \\ &= -\boldsymbol{A}\boldsymbol{D}_2'\boldsymbol{\Lambda}\boldsymbol{\zeta} + \boldsymbol{A}\boldsymbol{\varepsilon}\end{aligned}$$

We can now compute the variance. As $\boldsymbol{\zeta}$ and $\boldsymbol{\varepsilon}$ are assumed to be independent, with variances $\text{Var}(\boldsymbol{\zeta}) = \sigma_\varepsilon^2\boldsymbol{\Lambda}^{-1}$ and $\text{Var}(\boldsymbol{\varepsilon}) = \sigma_\varepsilon^2\boldsymbol{I}_T$:

$$\begin{aligned}\text{Var}(\hat{\boldsymbol{\tau}} - \boldsymbol{\tau}) &= \text{Var}(-\boldsymbol{A}\boldsymbol{D}_2'\boldsymbol{\Lambda}\boldsymbol{\zeta}) + \text{Var}(\boldsymbol{A}\boldsymbol{\varepsilon}) \\ &= (\boldsymbol{A}\boldsymbol{D}_2'\boldsymbol{\Lambda})\text{Var}(\boldsymbol{\zeta})(\boldsymbol{A}\boldsymbol{D}_2'\boldsymbol{\Lambda})' + \boldsymbol{A}\text{Var}(\boldsymbol{\varepsilon})\boldsymbol{A}' \\ &= (\boldsymbol{A}\boldsymbol{D}_2'\boldsymbol{\Lambda})(\sigma_\varepsilon^2\boldsymbol{\Lambda}^{-1})(\boldsymbol{\Lambda}'\boldsymbol{D}_2\boldsymbol{A}') + \boldsymbol{A}(\sigma_\varepsilon^2\boldsymbol{I}_T)\boldsymbol{A}'\end{aligned}$$

We simplify the expression and use symmetry. Since $\boldsymbol{\Lambda}$ is diagonal, $\boldsymbol{\Lambda}' = \boldsymbol{\Lambda}$. The

entire matrix $\boldsymbol{A}^{-1} = (\boldsymbol{I}_T + \boldsymbol{D}_2'\boldsymbol{\Lambda}\boldsymbol{D}_2)$ is symmetric, therefore its inverse $\boldsymbol{A}$ is also symmetric, $\boldsymbol{A}' = \boldsymbol{A}$. Finally, factor out $\sigma_\varepsilon^2\boldsymbol{A}$ and $\boldsymbol{A}$:

$$\begin{aligned}\text{Var}(\hat{\boldsymbol{\tau}} - \boldsymbol{\tau}) &= \sigma_\varepsilon^2\boldsymbol{A}\boldsymbol{D}_2'(\boldsymbol{\Lambda}\boldsymbol{\Lambda}^{-1})\boldsymbol{\Lambda}\boldsymbol{D}_2\boldsymbol{A} + \sigma_\varepsilon^2\boldsymbol{A}\boldsymbol{A} \\ &= \sigma_\varepsilon^2\boldsymbol{A}\boldsymbol{D}_2'(\boldsymbol{I}_{T-2})\boldsymbol{\Lambda}\boldsymbol{D}_2\boldsymbol{A} + \sigma_\varepsilon^2\boldsymbol{A}^2 \\ &= \sigma_\varepsilon^2\boldsymbol{A}\left(\boldsymbol{D}_2'\boldsymbol{\Lambda}\boldsymbol{D}_2 + \boldsymbol{I}_T\right)\boldsymbol{A}\end{aligned}$$

By definition, $\boldsymbol{A}^{-1} = (\boldsymbol{I}_T + \boldsymbol{D}_2'\boldsymbol{\Lambda}\boldsymbol{D}_2)$ and substituting back:

$$\begin{aligned}\text{Var}(\hat{\boldsymbol{\tau}} - \boldsymbol{\tau}) &= \sigma_\varepsilon^2\boldsymbol{A}\left(\boldsymbol{A}^{-1}\right)\boldsymbol{A} \\ &= \sigma_\varepsilon^2\boldsymbol{A} \\ &= \sigma_\varepsilon^2\left(\boldsymbol{I}_T + \boldsymbol{D}_2'\boldsymbol{\Lambda}\boldsymbol{D}_2\right)^{-1}\end{aligned}$$

D Alternative choices of smoothing parameter for budget shares

To assess the robustness of our main findings with respect to the choice of the smoothing parameter for the budget shares, λ_s, we re-estimate the component contributions to the business cycle using alternative values. Figure D.1 displays the decomposition for $\lambda_s = 20$, $\lambda_s = 40$ (our baseline value from the main text), and $\lambda_s = 80$. Table D.1 confirms that the weights do not change much with these alternative smoothing parameters.

While a lower value of λ_s allows for a more volatile underlying trend in the shares and a higher value implies a smoother one, the central conclusion of our analysis remains unchanged. Across all three specifications, private investment consistently emerges as the dominant contributor to business cycle fluctuations, while government expenditure exhibits a clear counter-cyclical pattern. The main findings of the paper are therefore robust to reasonable variations in this parameter.

Table D.1: Weights in the models with alternative smoothing parameter for budget shares, in percent

	Multivariate and time-varying		
	$\lambda_s = 20$	$\lambda_s = 40$	$\lambda_s = 80$
Private consumption	41.6	42.0	42.5
Government expenditure	59.8	56.0	53.3
Private investments	-5.0	-0.8	2.0
Export	3.7	2.7	2.2

Notes: Weights ($\hat{\boldsymbol{\omega}}$) alternative smoothing parameter for budget shares, where $\lambda_s = 40$ is used in the paper.

Figure D.1: Deviation from trend and contributions to these deviations, in percentage points. Multivariate HP trend and cycle with COVID-19 correction, based on time-varying budget shares with different smoothing parameter

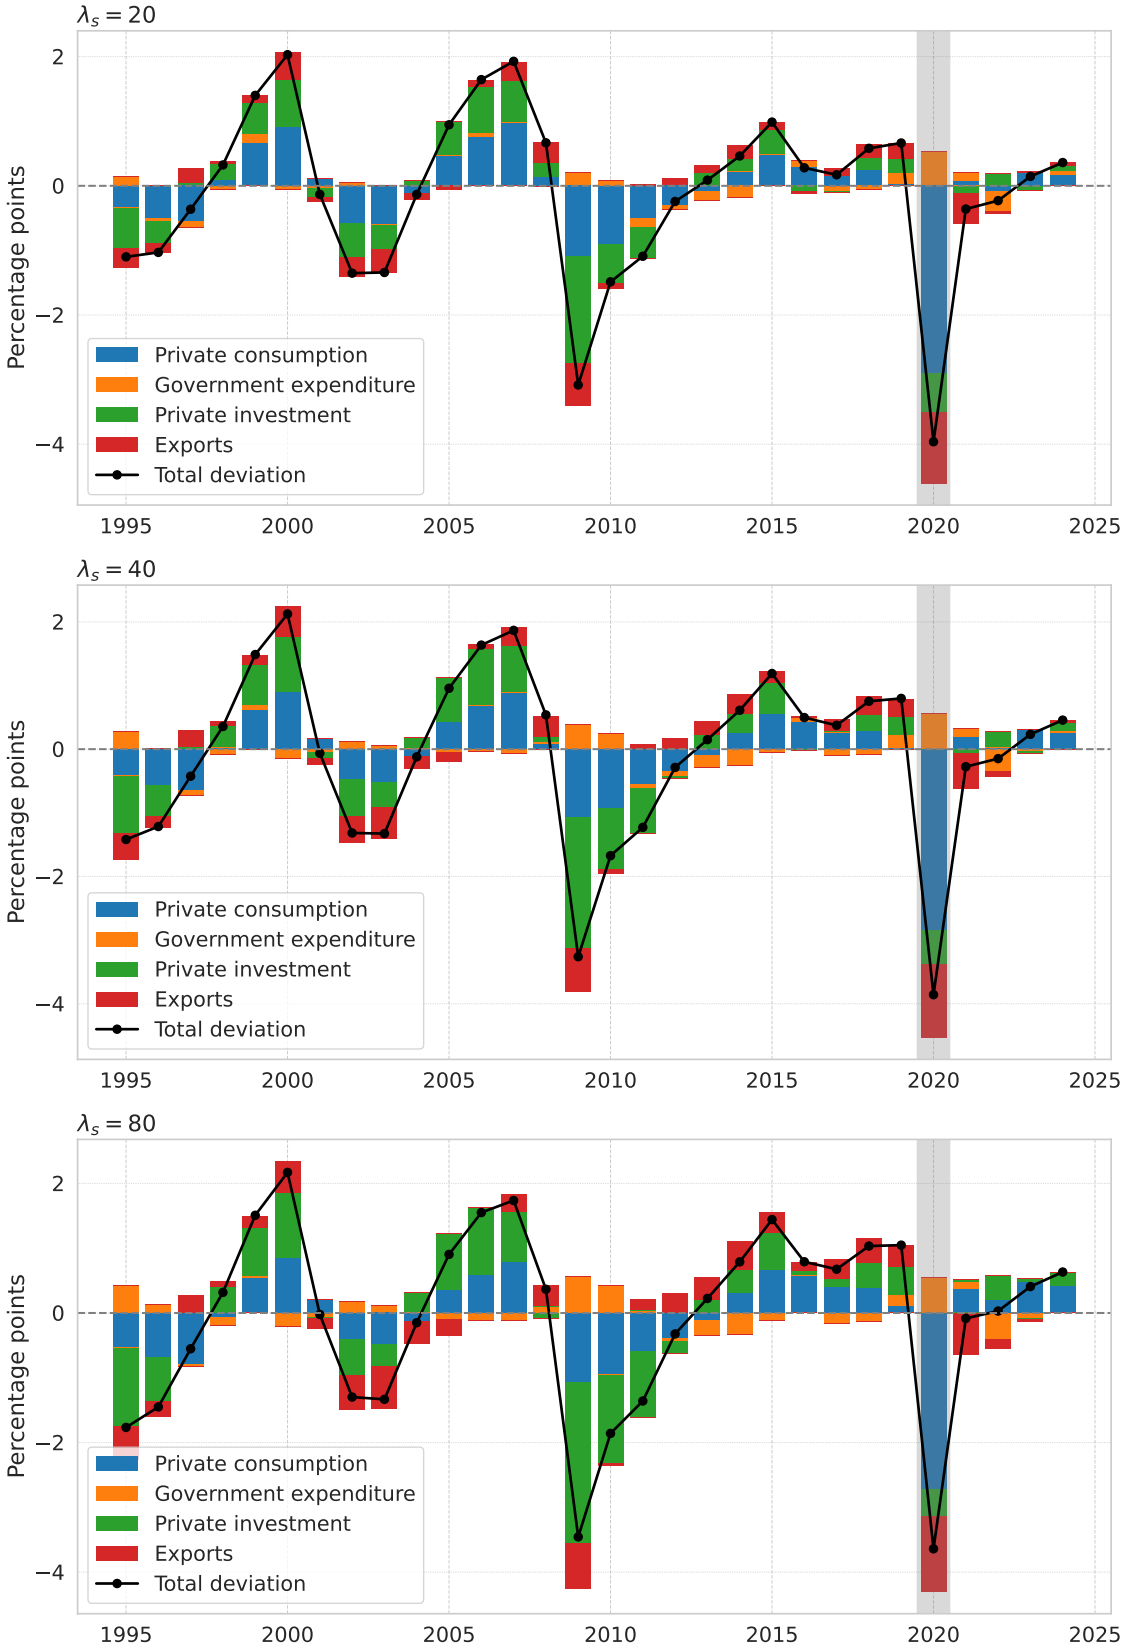

Notes: Upper panel: $\lambda_s = 20$. Middle panel: $\lambda_s = 40$. Lower panel: $\lambda_s = 80$. Shaded area indicates the COVID-19 period (2020).