EUROPEAN CENTRAL BANK

EUROSYSTEM

Working Paper Series

Luca Fosso

Decomposing US economic fluctuations: a trend-cycle approach

No 3138

Abstract: This paper proposes a unified framework to study the permanent and transitory origins of US economic fluctuations. The model provides a reasonable account of the evolution of the economy in the post-war period and of the recent inflation episode. Overall, it constitutes a comprehensive framework to offer policy guidance and a flexible empirical counterpart to more heavily-parametrized structural models.

Keywords: *Trend-Cycle, business cycle fluctuations, stochastic trends.*

JEL Classification: *E32, C32, E52*

NON-TECHNICAL SUMMARY

This paper proposes a unified framework for modelling the long-term and the business cycle dynamics using a trend-cycle Vector Autoregression (TC-VAR). The approach is motivated by the fact that most macroeconomic time series can be viewed as the sum of a long-term and a business cycle component (Beaudry et al., 2020). Long-term dynamics (or trends) – intended as those fluctuations that occur beyond business cycle horizons – contain valuable information for policymakers. Trend inflation, for example, provides central bankers with an assessment of the ability of monetary policy in maintaining inflation within confines that are consistent with the objective target. Business cycle dynamics are the transitory cyclical perturbations from the long-run trajectory of the macroeconomy that policymakers seek to minimize.

The concomitant analysis of these two components has proven problematic in standard SVARs – see e.g., Sims, 2000, Fernald, 2007, Canova et al., 2010, 2013, Giannone et al., 2019, Bergholt et al., 2024. These studies document the limitations of conventional SVARs specifications in addressing the low-frequency component of the data and warn against indiscriminately imposing a deterministic structure on the long-run component of the data.

Building on Stock and Watson (1988, 2007) and Villani (2009), the modelling strategy presented in this paper deviates from standard VARs along two key dimensions: (i) trends are allowed to vary over time; (ii) low-frequency correlations among macro variables are derived from new-Keynesian theory.

There are three principal advantages to preferring a TC-VAR over a VAR with a deterministic constant. First, TC-VAR enables to study both the long-term trend and the business cycle dynamics of the data in question, thereby facilitating a comprehensive analysis. Second, as trends are allowed to vary over time and assumptions on the long-run comovement in the macroeconomic time series is derived from economic theory, the inference on the cyclical properties of the data is robust to breaks and structural changes in the mean. Consequently, the researcher is not required to de-mean, de-trend or detect possible structural breaks prior to estimation. Third, shock identification is robust to changes in the trends. Therefore, the exclusion of subsamples is no longer necessary, and valuable information is not lost.

The model reasonably represents some stylised facts of the US economy post-WWII. These facts include low trend growth since the 2000s (Fernald 2004, 2023) and low/stable inflation (Ascari & Fosso 2024; Hasenzagl et al. 2022). The model also enables the estimation of the real natural rate of interest, which captures the secular decline observed since the mid-1990s (Del Negro et al., 2017; Holston et al., 2017).

The post-pandemic inflation surge provides a useful laboratory to demonstrate the model's efficacy in evaluating the cyclical sources of key macro fluctuations around historical episodes. Three key results stand out. First, inflation was mostly a business cycle phenomenon. The model finds no evidence that trend inflation drifted away from the target throughout the post-pandemic recovery, despite the Federal Reserve's monetary policy was very accommodative following the pandemic. Second, the rise in inflation was mostly due to demand shocks. Supply shocks were prominent but mostly confined around the early stages of the recovery and the Russian invasion of Ukraine in 2022. This indicates that, according to the model's estimates, supply did not generate self-fulfilling inflationary pressures.

Overall, the model proposed in this paper constitutes a comprehensive framework for the joint monitoring of the evolution of key policy variables and the evaluation of the contribution of structural shocks to business cycle fluctuations. It thus offers a valuable platform for the provision of policy advice.

1 Introduction

Motivation Structural Vector Autoregressive (SVAR) models are arguably the most standard tool for the analysis of macroeconomic time series. Compared to Dynamic Stochastic General Equilibrium (DSGE) models, they permit more general forms of auto- and cross-correlations, thereby offering a more flexible yet structural representation of the data. Today, SVARs form the bedrock of the policy economist's toolkit, providing policy advice and informing monetary policy decisions.

One common assumption in standard VARs is to add an exogenous constant term to the autoregressive law of motion and impose a deterministic structure on the low-frequency component of the data. This presupposes that economic time series fluctuate around a stable long-run equilibrium. However, such an assumption stands in stark contrast with a large literature documenting substantial low frequency variation in macroeconomic time series (Lubik et al. (2019), Beaudry et al. (2020), etc.). The data plotted

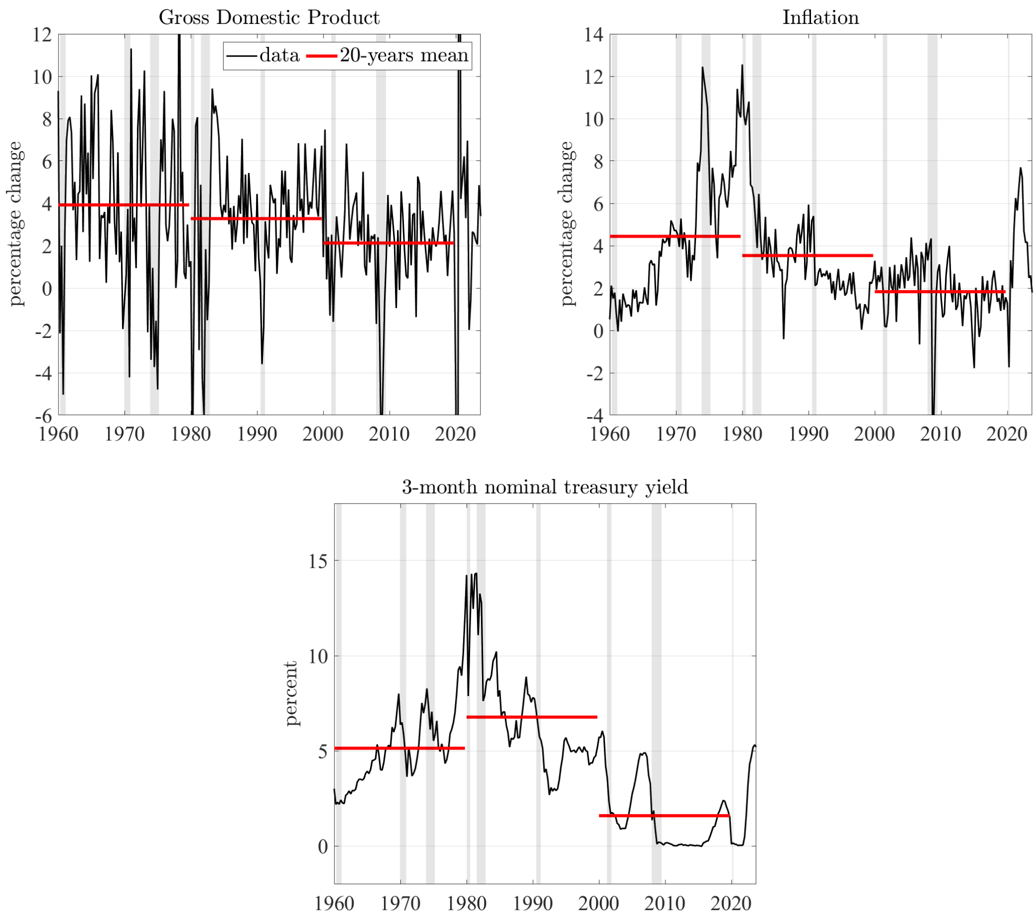

Figure 1: Data (black solid), 20-years sub-sample means (red solid), NBER recession dates (grey shaded areas).

in figure 1 do not seem to support the "stable mean" hypothesis. For example, while the average growth rate of real GDP ranged between 3.5% and 4% over the 1960-2000 period, it has fallen to 2% in the last 20 years. The variation in the sub-sample averages of inflation and interest rate is even more pronounced. Moreover, table 1 shows that their HP-filtered low-frequency components also exhibit some degree of co-movement, with the exception of inflation and output.

Table 1: HP-filtered trend correlations.

	output	inflation	interest rate
output	–		
inflation	-0.05	–	
interest rate	0.28	0.78	–

Notes: Output is defined as the quarterly annualized percentage change of real GDP. Inflation is the quarterly annualized percentage change of the PCE index. The nominal interest rate is defined as the three-month constant maturity treasury yield. Data spans from 1960Q1 to 2019Q4.

Low-frequency dynamics (or trends) – intended as those fluctuations that occur beyond business cycle horizons – have been shown to contain valuable information for policymakers. Trend inflation, for example, provides central bankers with an assessment of the efficacy of monetary policy in maintaining inflation within confines that are consistent with the objective target. In accordance with modern monetary theory (Woodford (2003)), the natural rate of interest is regarded as the yardstick for assessing the stance of monetary policy.

The estimation of these low-frequency components has nonetheless proved to be problematic in standard VARs (Sims (2000), Giannone et al. (2019), Bergholt et al. (2024a)). In particular, if data embed persistent transitional dynamics and/or breaks in the mean, the trend component is misspecified. The bias arising from misspecification may interact in unpredictable ways with the model's residuals and give rise to a nuanced parameter problem. In such a case, enforcing a deterministic trend may have adverse effects on the identification of the structural shocks. Fernald (2007), Canova et al. (2010, 2013) discuss at length the exceptional sensitivity of long-run restrictions to low-frequency correlations between hours worked and labor productivity when studying the effects of productivity shocks on hours. More recently, Bianchi et al. (2023b) highlight the limitations of the max-share template used by Angeletos et al. (2020) in identifying the "main business cycle shock". They demonstrate that when the data are characterized by structural shifts, the max-share approach is unable to recover the autocovariance structure of the data generating process at business cycle frequencies. These results emphasize once more the limitations of conventional SVARs specifications in addressing the low-frequency component of the data and warn against indiscriminately imposing a deterministic structure on the long-run component of the data.

CONTRIBUTION This paper designs a flexible time series model for the U.S. that expresses variables in deviation from their trend.[1] This modelling strategy embeds two key features: (i) trends are stochastic; (ii) low-frequency correlations are guided by economic theory, and consistent with a wide class of models. To fix ideas, the long-run co-movement between trend output growth and the nominal interest rate is derived from the optimal intertemporal consumption plan of the representative household in a standard three-equations New-Keynesian DSGE model à la Galí (2008). A Fisher equation postulates the long-run co-movement between nominal interest rate and inflation.

These two ingredients make inference on the cyclical component of the data robust to slow-moving fluctuations arising from persistent transitional dynamics and/or structural breaks in the unconditional mean. Therefore, differently from standard SVAR analysis, a more rigorous approach to the modelling of the long-run component of the data is taken. Importantly, neither preliminary filtering of the data nor reduction of the sample utilized for estimation are required.[2] Furthermore, no prior knowledge of the dates of structural breaks is necessary. This is particularly convenient as dates at which breaks occur are typically unknown *ex ante*. The identification of business cycle shocks through standard identification schemes is straightforward to implement and robust to breaks in the data. Thus, the model provides a reliable assessment of the relative contribution of the structural shocks driving the business cycle around key historical events.

Finally, while this is not explicitly addressed in the paper, the empirical setup proposed below is versatile enough to investigate the role of common factors, both trend and cycles, among countries. For instance, it can be designed to encompass both US and EA macroeconomic aggregates, thereby facilitating the study of the factors driving the long-term convergence/divergence between the two economies, as well as the sources of business cycle synchronization.[3]

RESULTS Based on an application to the US economy, the model provides a reasonable representation of some stylized facts about the US economy in the post-WWII period. These include the low trend growth since 2000s – consistent with the slowdown in productivity found by Fernald and Ramnath (2004) and Fernald et al. (2023) –, as well as the low and stable inflation trend over the twenty years preceding

[1]The modelling strategy builds on the unobserved component (UC) models literature (Watson (1986), Stock and Watson (1988, 2007), Villani (2009)).

[2]These shortcuts are inefficient and may introduce bias in other ways. For example, small-sample biases may interact in an unpredictable way with measurement errors, making evidence uninterpretable (Canova et al. (2010)).

[3]For a comprehensive discussion, the reader can refer to Bordo and Helbling (2010).

the Covid-19 pandemic (Hasenzagl et al. (2022), Ascari and Fosso (2024)). The model also allows to obtain an estimate of the real natural rate of interest, that captures the secular decline observed since the mid-1990s – Del Negro et al. (2017), Holston et al. (2017). The decline is attributed to slower trend growth as well as to structural changes in agents' preferences, possibly due for example to demographics, globalisation, etc. – Cesa-Bianchi et al. (2022).

The model is designed to obtain estimates of policy-relevant business cycle indicators, such as the output gap, the inflation gap and an interest rate gap. Since the model is instructed with data on global commodity prices, it also allows to extract a proxy of the "global commodity cycle", that turns out to be highly correlated with the estimated domestic gaps. Importantly, the strong co-movement between the price of energy commodities and domestic gaps is neither imposed nor informed with external proxies, but it genuinely reflects the underlying co-movement in the data captured by the model. Altogether, the estimated states give a comprehensive assessment of the state of the domestic cycle at any given point in time, also in relation to the conditions in the international market for commodities – i.e., a key benchmark for global economic conditions (Baumeister et al. (2022), Delle Chiaie et al. (2022)).

The model provides a reasonable account of the business cycle co-movement of key macro aggregates around selected historical episodes, as well. I take the post-pandemic inflation surge as a laboratory. Three main results emerge. First, inflation was mainly a business cycle phenomenon. According to the model, in fact, there is no evidence that trend inflation deviated significantly away from the target throughout the post-pandemic recovery. While not conclusive, this would suggests that the credibility gained by the Fed as a successful inflation targeter helped to prevent inflation expectations from becoming entrenched.[4] Importantly, the results on trend inflation are robust to the so-called end-point problem typical of state space models with unobserved states, suggesting that the model contains valuable information that can be used to monitor trend inflation in *real-time*. Second, the positive inflation gap was primarily due to demand shocks. Although supply shocks played a non-negligible role during the early phases of the recovery and following the Russian invasion of Ukraine in 2022, these shocks were short-lived. Consequently, according to the model's estimates, supply did not give rise to self-fulfilling inflationary pressures. Third, the model reconciles the re-acceleration of economic activity and the consistent easing of inflation in 2023 with improving of domestic supply conditions and slowly fading demand pressures.

[4] See Hajdini et al. (2025) for a discussion.

RELATED LITERATURE This paper is related to a growing body of literature that estimates VARs *with common stochastic trends* and studies the drivers of the secular trends in macroeconomic data – Del Negro et al. (2017), Crump et al. (2019), Ascari and Fosso (2024), Bergholt et al. (2024b), Bjørnland et al. (2025). In addition to the analysis of secular trends in the data, this paper also underscores the importance of an appropriate accounting of the slow-moving dynamics for the measurement and the assessment of business cycle fluctuations (Perron (1989)).

This paper is also related to Bergholt et al. (2024a), who discuss the poor performance of *correctly-specified* VARs in underpinning the long-run dynamics in the data through the deterministic component. This pathology becomes even more acute in the presence of large transitional dynamics and structural breaks, namely when the trend component is *misspecified*. While addressing the theoretical underpinnings of this issue is beyond the scope of this paper, the model outlined below helps mitigate this pathology through time variation in the trends and theory-robust restrictions to the long-run comovement of the data.

Finally, this work contributes to the literature studying the post-pandemic inflation episode – Ball et al. (2022), Bańbura et al. (2023), Eickmeier and Hofmann (2023), Shapiro (2024), Ascari et al. (2023, 2024), Giannone and Primiceri (2024), Mori (2024), Pinilla-Torremocha (2025). Similar to some of these findings, I document a major role of demand factors in driving the post-pandemic inflation episode in the US. Differently from previous studies, the analysis presented below also finds that trend inflation remained anchored at 2% throughout the post-pandemic recovery, providing a comprehensive assessment of the inflation outlook and Fed's monetary policy.

OUTLINE The remainder of the paper is organized as follows. The empirical methodology is exposed in section 2. Section 3 presents and discusses the results. Section 3.4 shows the importance of inflation expectations for the monitoring inflation in real-time. Section 4 discusses the implications of alternative long-run assumptions for the analysis of the cycle. Section 5 concludes.

2 A TREND-CYCLE VAR FOR THE US ECONOMY

The trend-cycle VAR (TC-VAR henceforth) outlined below decomposes the data into two unobservable states: (i) a permanent component (or trend) that captures slow-moving changes in the data; (ii) a transitory component (or cycle) that describes the temporary deviations of the data from the trend. To fix ideas, consider an $n \times 1$ vector of data y_t, which is the sum of two unobserved states:

$$y_t = \Lambda^* y_t^* + \tilde{y}_t \tag{1}$$

y_t^* and $\tilde{y}_t$ constitute the stochastic permanent and transitory components of the data, respectively. An important object of interest is matrix Λ^*, which controls for the possible presence of $r \leq n$ common trends in the data.[5] Assumptions on Λ^* are discussed in the next sub-sections.[6] Trends are assumed to follow a random walk:

$$y_t^* = y_{t-1}^* + u_t^* \qquad u_t^* \sim N(0, \Sigma^*) \tag{2}$$

The fluctuations originating from transitory deviations of the data from y_t^* follow an autoregressive law of motion in reduced-form:

$$y_t - \Lambda^* y_t^* = \tilde{\Phi}(L)(y_{t-p} - \Lambda^* y_{t-p}^*) + \tilde{u}_t, \qquad \tilde{u}_t \sim N(0, \tilde{\Sigma}) \tag{3}$$

where $\tilde{y}_t = y_t - \Lambda^* y_t^*$. $\tilde{\Phi}(L)$ is the companion matrix of the autoregressive coefficients, whose roots are assumed to lie within the unit circle to ensure stationarity. Furthermore, permanent and transitory innovations are mutually uncorrelated, i.e., $corr(u_t^*, \tilde{u}_t) = 0$.[7]

Equations (1)-(3) constitute the model I confront with data. By combining relevant macroeconomic aggregates with reasonable assumptions on Λ^*, this model provides a unified framework to decompose the sources of US economic fluctuations and to monitor the slow-moving changes in the trends.

[5] Λ^* can be viewed as the matrix of loadings in an empirical time series model with a factor structure in the low-frequency component of the data. In this sense, the model presented here can be related to the dynamic factor model proposed by Antolin-Diaz et al. (2017), which includes a common drifting constant between GDP and consumption growth to model the secular slowdown in US trend growth.

[6] Λ^* can be calibrated, estimated, or both. Assumptions on Λ^* can also be formulated to identify the *structural* trends driving the data. The reader can refer to Bergholt et al. (2024b) for an application.

[7] This implies that the cycle does not affect the trend (and vice versa) by construction. This assumption may appear in contrast for example with the emerging evidence of temporary demand shocks having scarring effect on potential output – the so-called hysteresis effects (see Cerra et al. (2023) for a literature review). However, while such an assumption may constitute a potential limitation to the use of this empirical strategy, it ensures a clear separation between the permanent and the transitory components. It, therefore, enables to get an estimate of the cycle that exhibits meaningful fluctuations that are comparable to the ones published by official agencies, such as the Congressional Budget Office and the National Bureau of Economic Research. It is possible to relax this assumption, which would make the TC-VAR observationally equivalent to performing a Beveridge-Nelson decomposition on the data (Morley et al. (2003)). However, this comes with costs. As shown in Morley et al. (2003) and Grant and Chan (2017), when the correlation between trend and cycle is switched on, most of the variation in the data is soaked up by the trend. As a consequence, the estimated cycle is so small and noisy that its fluctuations would become irrelevant from a policy perspective. Finally, while having non-negligible implications for the estimation of TC-VARs, relaxing the uncorrelation assumption goes beyond the scope of this paper. Thus, I leave this methodological aspect to future research.

DATA The input vector y_t includes data on real gross domestic product, personal consumption expenditure, three-month constant maturity treasury yield, one-year ahead inflation expectations and the commodity price index released by the Commodity Research Bureau.[8] All input variables are expressed in annualized quarter-on-quarter percentage changes with the exception of the treasury yield, which enters the model in annualized levels. The estimation sample spans 1960Q1 – 2023Q4.

PRIORS The initial conditions of the steady states are distributed according to $y_0^* \sim \mathcal{N}(\underline{y_0^*}, I_q)$. The prior means $\underline{y_0^*}$ are defined using averages of input data over the pre-sample period, which spans 1954Q1-1959Q4. The initial conditions of the cycles are distributed according to $\tilde{y}_0 \sim \mathcal{N}(0_n, I_n)$. This assumption implies that cycles fluctuate symmetrically around a zero mean. Finally, the priors for the remainder model's coefficients are distributed according to:

$$\Sigma^* \sim \mathcal{IW}(\kappa^*, (\kappa^* + n + 1)\underline{\Sigma}^*) \tag{4}$$

$$\tilde{\Sigma} \sim \mathcal{IW}(\tilde{\kappa}, (\tilde{\kappa} + n + 1)\underline{\tilde{\Sigma}}) \tag{5}$$

$$vec(\tilde{\Phi})|\tilde{\Sigma} \sim \mathcal{N}(vec(\underline{\tilde{\Phi}}), \tilde{\Sigma} \otimes \underline{\Omega})\mathcal{I}(\tilde{\Phi}), \tag{6}$$

where and $\mathcal{I}(\tilde{\Phi})$ is an indicator function that is equal to one, when the VAR of the cycle block is stationary, zero otherwise. $\mathcal{IW}$ is the Inverse-Wishart distribution with κ degrees of freedom and mode $\underline{\Sigma}$. The prior mode of trend innovations $\underline{\Sigma}^*$ is assumed to be diagonal.[9] The priors on the main diagonal entries are conservative in limiting the amount of variance attributable to the trends. In combination with the aforementioned trend-cycle uncorrelation assumption, this prior assumption allows to clearly disentangle trends from cycles. Consistently, and similar to Del Negro et al. (2017), all prior volatilities are elicited to imply that the standard deviation of the expected change in the trends over a period of six decades is only one percentage point. The degrees of freedom $\kappa^* = 100$ imply a rather tight prior around the mode.

Moving to the cycle block, the priors for the lag coefficients are standard Minnesota with an overall tightness hyperparameter equal to 0.2, as in Giannone et al. (2015), and the own-lag hyperparameters centered around zero, instead of one, since this is the transitory block. The prior mode of the transitory

[8] I use the inflation expectations series available on the Cleveland Fed website. The series is calculated from a model that combines data on treasury yields, inflation, inflation swaps and survey-based inflation expectations. It starts from 1982 and is the longest publicly available series on US inflation expectations.

[9] While the prior excludes any correlation among trends a priori, it does not prevent the posterior from deviating from this assumption if the data indicates a correlation among trends. This represents a major difference to the approach taken in Bergholt et al. (2024b). In their study, the covariance structure of the trends is *imposed* to be diagonal. This is because their model is elicited to identify the *structural* trends driving US trend growth (e.g., technology-specific and labor-specific factors), which are orthogonal by construction. Instead, the purpose of this paper is to obtain a reasonable representation of these secular empirical trends, which is instrumental for a more robust evaluation of the business cycle dynamics.

innovations $\tilde{\Sigma}$ is assumed to be an identity matrix and rather uninformative, as the degrees of freedom are $\tilde{\kappa} = n + 2$.

The model is estimated with Kalman smoothing techniques, which allow to retrieve the latent states also in the case of an unbalanced dataset. This allows, for example, to treat the zero lower bound periods in 2008-2015 and 2020-2021 as missing observations. Section A in the Appendix provides a detailed exposition of the state space representation and the steps for the estimation with Gibbs sampling.

2.1 Theory-based Assumptions for the Long-Run

An important focus of the model is matrix Λ^*, which requires additional assumptions to pin down the long-run component in the data and, consequently, define the transitory fluctuations around it. This sub-section exposes the assumptions on Λ^* and their link with economic theory. To fix ideas, let us start from the optimal intertemporal consumption plan implied by a wide class of model consistent with new-Keynesian theory – see, e.g., Galí (2008):

$$1 + i_t = \beta_t^{-1} \mathbb{E}_t \left[\left(\frac{C_t}{C_{t+1}} \right)^{-\sigma} \frac{P_{t+1}}{P_t} \right] \tag{7}$$

β_t is a stochastic discount factor. The household willingness to substitute consumption between periods is pinned down by σ. In the long-run, the log-linear version of eq. (7) is given by:[10]

$$i_t^* = -log(\beta_t^*) + \sigma \Delta c_t^* + \pi_t^* \tag{8}$$

which leaves us with a decomposition of the nominal natural rate of interest into three underlying long-run drivers. The first term in the right-hand side of the equation – $z_t^* = -log(\beta_t^*)$ – captures slow-moving shifts in households savings-consumption preferences due to, among other things, demographics and globalisation. The second term defines the long-run dynamics of real output, which is set to grow in the long-run on at the growth rate of technological progress Δg_t^*. π_t^* refers to positive trend inflation. Eq. (8) can therefore be rewritten as:

$$i_t^* = z_t^* + \sigma \Delta g_t^* + \pi_t^* \tag{9}$$

[10]The $*$ superscript denotes steady state variables. For consistency with the assumptions of time varying trends in the empirical model, steady states maintain the time subscripts.

Moreover, as in Del Negro et al. (2017), an analytical solution for the long-run real natural rate of interest can derived by postulating a Fisher equation in (9):

$$r_t^* = z_t^* + \sigma \Delta g_t^* \tag{10}$$

Equipped with the solution in Eq. (9), a set of theory-based long-run restrictions on the elements of Λ^* are available. This can be seen by unfolding the second term in the right-hand side of eq. (1):

$$\begin{bmatrix} \Delta c_t^{\mathcal{T}} \\ \pi_t^{\mathcal{T}} \\ \pi_{e,t}^{\mathcal{T}} \\ i_t^{\mathcal{T}} \\ \pi_{c,t}^{\mathcal{T}} \end{bmatrix} = \underbrace{\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ \sigma & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}}_{\Lambda^*} \begin{bmatrix} \Delta g_t^* \\ \pi_t^* \\ z_t^* \\ \pi_{c,t}^* \end{bmatrix} \tag{11}$$

Δg_t^* gives raise to a long-run comovement between the growth rate of real output and nominal interest rates. The long-run loading of Δg_t^* on nominal interest rates is given by σ, the intertemporal elasticity of substitution in the theoretical model. σ is calibrated to unity, consistent with a representative household maximizing a lifetime logarithmic utility function.[11] π_t^* is the the time-varying nominal anchor that jointly contributes to the long-run dynamics of inflation, one-year ahead inflation expectations and the nominal rate of interest.[12] z_t^* captures slow-moving shifts in agents' preferences affecting interest rates in the long-run. Finally, $\pi_{c,t}^*$ is a trend in internationally traded commodity prices which is assumed to be exogenous to the long-run dynamics of the domestic economy. According to eq. (11), the long-run is therefore characterized by four time-varying trends.

Before moving to the next section, a few comments are in place. The use of economic theory to make inference in empirical time series models relates this paper to the DSGE-VAR methodology pioneered by Del Negro et al. (2007). In contrast to their approach, however, the purpose of this paper is not to infer on the full covariance structure implied by theory. Rather, it is to conduct a robust analysis of the cyclical component, with a particular emphasis on addressing the misspecification of the trend component.[13]

[11] It is otherwise possible to formulate a prior and estimate it. In a robustness check, I relax this assumption and show that the posterior density of σ concentrates around unity.

[12] The common trend between inflation and one-year ahead inflation expectations is motivated by the strong signal embedded in survey expectations measures, which helps reducing the uncertainty around trend inflation estimates. I will discuss the importance of this assumption in Section 3.4.

[13] Bergholt et al. (2024b) use economic theory to inform the estimation of a BVAR with a specific focus on the identification of structural trends. They formulate priors on specific structural elasticities to study the permanent impact of gender-specific

2.2 IDENTIFYING BUSINESS CYCLE SHOCKS

The theory-robust assumptions outlined in the previous Section pin down the long-run and allow to express variables in deviations from their time-varying long-run components. The next step is, therefore, to identify the shocks that are responsible for the transitory fluctuations around it – i.e., the business cycle shocks. Before discussing the identification scheme, let me briefly motivate one important modelling strategy that will become particularly beneficial for the interpretation of the results. While all the input variables enter the model in growth rates, the state space estimates the cycle of the log-levels of real GDP and commodity prices as opposite to the cycle of their growth rates.[14] This is particularly convenient because it allows to retrieve a measure of the output gap and of the commodity price cycle.[15]

The vector $\tilde{y}_t = y_t - \Lambda^* y_t^*$ thus includes: output gap $\tilde{g}_t$, inflation gap $\tilde{\pi}_t$, inflation expectations gap $\tilde{\pi}_{e,t}$, nominal interest rate gap $\tilde{i}_t$, commodity price gap $\tilde{\pi}_{c,t}$. In a more compact form, eq. (3) becomes:

$$\tilde{y}_t = \tilde{\Phi}(L)\tilde{y}_{t-p} + \tilde{u}_t, \qquad \tilde{u}_t \sim N(0, \tilde{\Sigma}) \tag{12}$$

where $\tilde{u}_t = \tilde{P}^{-1}\tilde{\varepsilon}_t$ and $\tilde{\Sigma} = (\tilde{P}'\tilde{P})^{-1}$. $\tilde{\varepsilon}_t$ is the vector of structural disturbances. The reduced-form residuals can therefore be mapped into economic meaningful shocks through the usual identification schemes employed in the VAR literature. I opt for a combination of zero and sign restrictions following the QR decomposition algorithm proposed by Arias et al. (2018).[16] Restrictions are imposed directly on impulse response functions and are summarized in Table 2. Four main sources of business cycle fluctuations are identified.[17] Demand and monetary policy shocks are responsible for the same co-movement between output and inflation gaps and are mutually exclusive via the opposite sign restriction on the nominal interest rate gap. Domestic and foreign supply shocks capture exogenous shifts in aggregate supply and are therefore responsible for the opposite co-movement between output and inflation gaps. Domestic supply disturbances are disentangled from imported ones through a zero impact restriction on the commodity price gap.

Overall, the model presented in this Section provides a parsimonious representation of the busi-

shocks on the US macroeconomy.

[14] This is accomplished via an appropriate formulation of the state space. See section 5 in the Appendix.

[15] Consistent with the literature on commodity prices, shocks to commodity prices can then be interpreted as transitory deviations from the level (Hamilton (2009), Kilian (2009)).

[16] Given the explicit link with structural models, set identification represents a convenient choice because the estimation uncertainty surrounding, for example, the IRFs can be directly interpreted as the parameter uncertainty implied by alternative calibrations of the underlying structural model.

[17] This makes the system partially identified. The residual shock thus soaks up all the remaining fluctuations that do not accrue from the identified shocks.

ness cycle. Moreover, while deeply parameterized structural models are firmly grounded on tight cross-equation restrictions, the model exposed above does not commit to one specific parametrization and constitutes a flexible "model-free" counterpart.

Table 2: Sign restrictions

	Business cycle	Monetary Policy	Domestic cost-push	imported cost-push
$\tilde{g}_t$	–	–	–	–
$\tilde{\pi}_t$	–	–	+	+
$\tilde{\pi}_{e,t}$				
$\tilde{i}_t$	–	+		
$\tilde{\pi}_{c,t}$			**0**	+

Notes: Restrictions are imposed on impact.

3 Results

This section presents the results and it is organized as follows. Sub-sections 3.1-3.2 discuss the model-based estimates of stars and gaps. Sub-section 3.3 investigates the drivers of the business cycle around selected historical episodes.

3.1 Slow-moving stars

Figure 2 plots the estimates of Δg_t^*, π_t^* and i_t^* together with observable data on real gross domestic product, inflation and inflation expectations, nominal interest rate. The results confirm some stylized facts about the evolution of the US economy over the past decades, corroborating the restrictions imposed to pin down the long-run.

Consistent with the facts documented by Fernald and Ramnath (2004) and Fernald (2015), trend growth accelerated in the mid-1990s (see the left panel), supported by advances in information technology.[18] During the first decade of 2000s, it declined by almost two percentage points and went even below 2% after the Great Recession, providing additional evidence in favour of long lasting effects of recessions on the growth rate of output – see, e.g., Summers (2015), Blanchard et al. (2015) and Maffei-Faccioli (2021).

The secular dynamics of trend inflation can be divided into a pre- and post-Volcker period. Before Volcker's administration, trend inflation was running at levels that were well above 2%, especially around

[18]For an extensive discussion on the log-term dynamics of productivity growth in the US, see also Byrne et al. (2016) and Cette et al. (2016). For a cross-country perspective, instead, the reader can refer to Fernald et al. (2023).

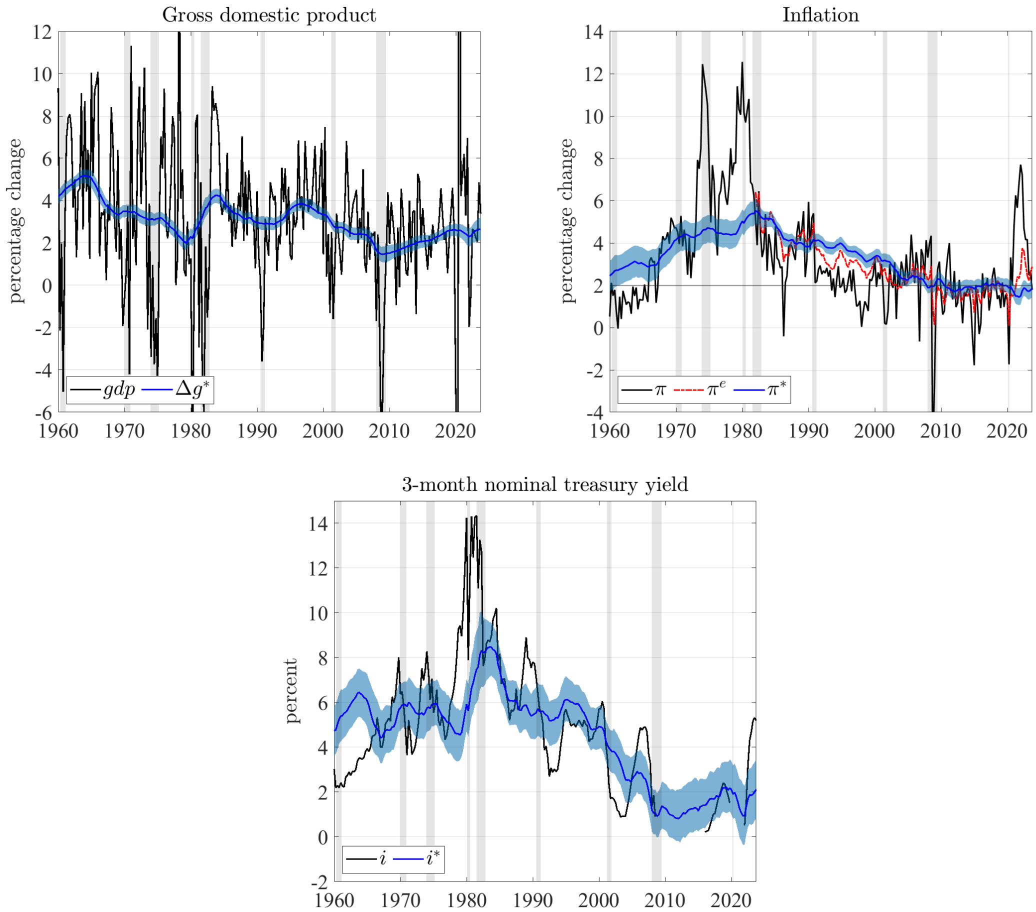

Figure 2: Data (black dashed), median estimate (blue solid), 68% coverage bands (blue shaded areas), NBER recession dates (grey shaded areas).

the two oil crises in the 1970s (see the middle panel), reflecting among other things the limited ability of monetary policy to maintain inflation expectations under control (Goodfriend and King (2005)). The Volcker administration represented a turning point and, starting from early 1980s, trend inflation was set on a downward path and settled close to the 2% target level in the early 2000s, consistent with well-anchored inflation expectations (red dashed line in the panel). Finally, and importantly, while the post-pandemic inflation surge led to the rise of one-year ahead inflation expectations, trend inflation remained anchored at the 2% target, providing evidence against a de-anchoring of medium to long-term expectations.

The nominal natural interest rate fluctuated within the 4%-6% window for most of the sample until mid-1990s, when it started to decline. In the aftermath of the Great Recession, when the policy rate hit the zero lower bound and inflation was low and stable, it reached its lowest level. Since the beginning of the post-pandemic tightening cycle, the nominal natural interest rate has slightly risen again, albeit remaining at historically low levels.

While not being the main focus of this paper, the model also enables to retrieve an estimate of the real natural rate of interest (r^*). Figure 3 plots the estimate of r^* based on eq. (10). Similar to Del Negro et al. (2017) and Holston et al. (2017), the estimate exhibits a persistent decline that started in the late 1990s. Since the Great Recession, it stabilizes and fluctuates around zero until the end of the sample. Table 3 breaks down the historical decline over the last twenty years into the contribution of trend growth and agents' preferences. According to the model's results, the decline of r^* can be attributed to both lower growth (Δg^*) and changes in preferences (z^*), with the latter partly reflecting the lower demand for savings associated with the increased population aging, as documented by Cesa-Bianchi et al. (2022), who show that this is a common feature of all major advance economies.

Overall, this subsection shows that model-based estimates of the trends provide a reasonable account of some stylized facts about the US economy in the post-war period, which would have been otherwise impossible to capture in a standard VAR setting. It highlights once more the importance of taking a more rigorous approach to the long-run component of the data especially when the goal is to study the sources of transitory deviations from the trend. The next subsection discusses the estimation of the gaps and the analysis of their main drivers with an special focus on the sources of the post-pandemic inflation surge.

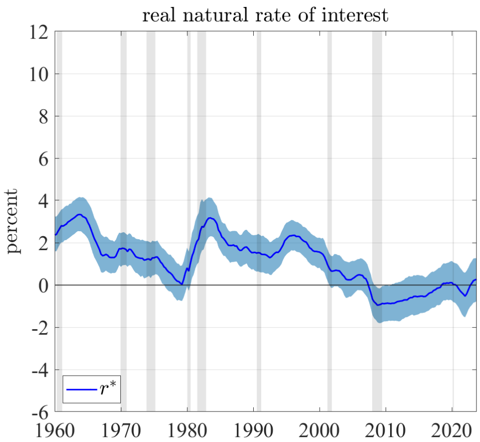

Figure 3: Median estimate (blue solid), 68% coverage bands (blue shaded areas), NBER recession dates (grey shaded areas).

r_t^*	-2.23
	(-2.70, -1.11)
Δg_t^*	-1.18
	(-1.64, -0.65)
z_t^*	-1.05
	(-1.16, -0.46)

Table 3: changes in trend, 1996-2023. 95% coverage bands in parentheses.

3.2 UNDERSTANDING THE GAPS

Consistent with the representation in eq. (12), figure 4 reports the variables in deviations from their steady state. The top-left panel displays the percentage deviations of output from its log-level trend.[19] This leaves us with a model-based estimate of the output gap, which is closely comparable to other reference measures, such as the one released by the Congressional Budget Office (CBO). The correlation between the CBO and the model-based estimate of the output gap is in fact beyond 0.8. Intuitively, the gap is positive during expansions and turns quickly negative when the economy is hit by a recession - see the grey shaded areas reporting the NBER recessions dates. The top-middle panel plots the inflation gap, which tracks key historical events closely. For instance, the gap is largely positive around the two oil price crises in the 1970s and negative around the Great Recession episode. As new-Keynesian theory would suggests, the cyclical behaviour of one-year ahead inflation expectations is tightly linked to the one of inflation but much less volatile. The bottom-left panel provides an estimate of the nominal interest rate in percentage deviations from its natural rate. Positive (negative) values can be therefore interpreted as periods during which monetary policy was restrictive (loose). To see this, focus for example on the period following the Great Recession: the interest rate gap is negative, reflecting the large accommodating stance until 2015. Furthermore, the estimate of the interest rate gap seems to track the Fed's stance in the most recent period quite well: loose at the onset of the Covid-19 pandemic but rapidly turning restrictive in response to the subsequent inflation upturn. Finally, the bottom-right panel plots the commodity price cycle, which is defined in percentage deviations from the trend of commodity prices in log-levels and is closely correlated with the inflation gap, especially in the pre-Great Moderation period. Intuitively, following temporary increases in commodity prices, the global market gets overheated and the gap becomes positive – see e.g., the positive gap during the 1970s and since 2020.

3.3 THE DRIVERS OF THE POST-PANDEMIC INFLATION EPISODE

This sub-section employs the post-pandemic inflation episode as a laboratory to show how the model helps rationalizing the drivers of business cycle fluctuations around key historical events. The discussion is backed by the shock decompositions of key input variables for the Federal Reserve. Shocks are identified following the restrictions exposed in table 2. Figure 5 plots the contribution of structural shocks to output gap (top panel), inflation gap (middle panel) and inflation expectations gap (bottom panel). These results attempt to shed light on some of the facts that have been at the center of policy and academic

[19] As anticipated in Section 2, similarly to macroeconomic models, I exploit the state space representation to model the cyclical behaviour of the output in log-levels as opposite to output growth.

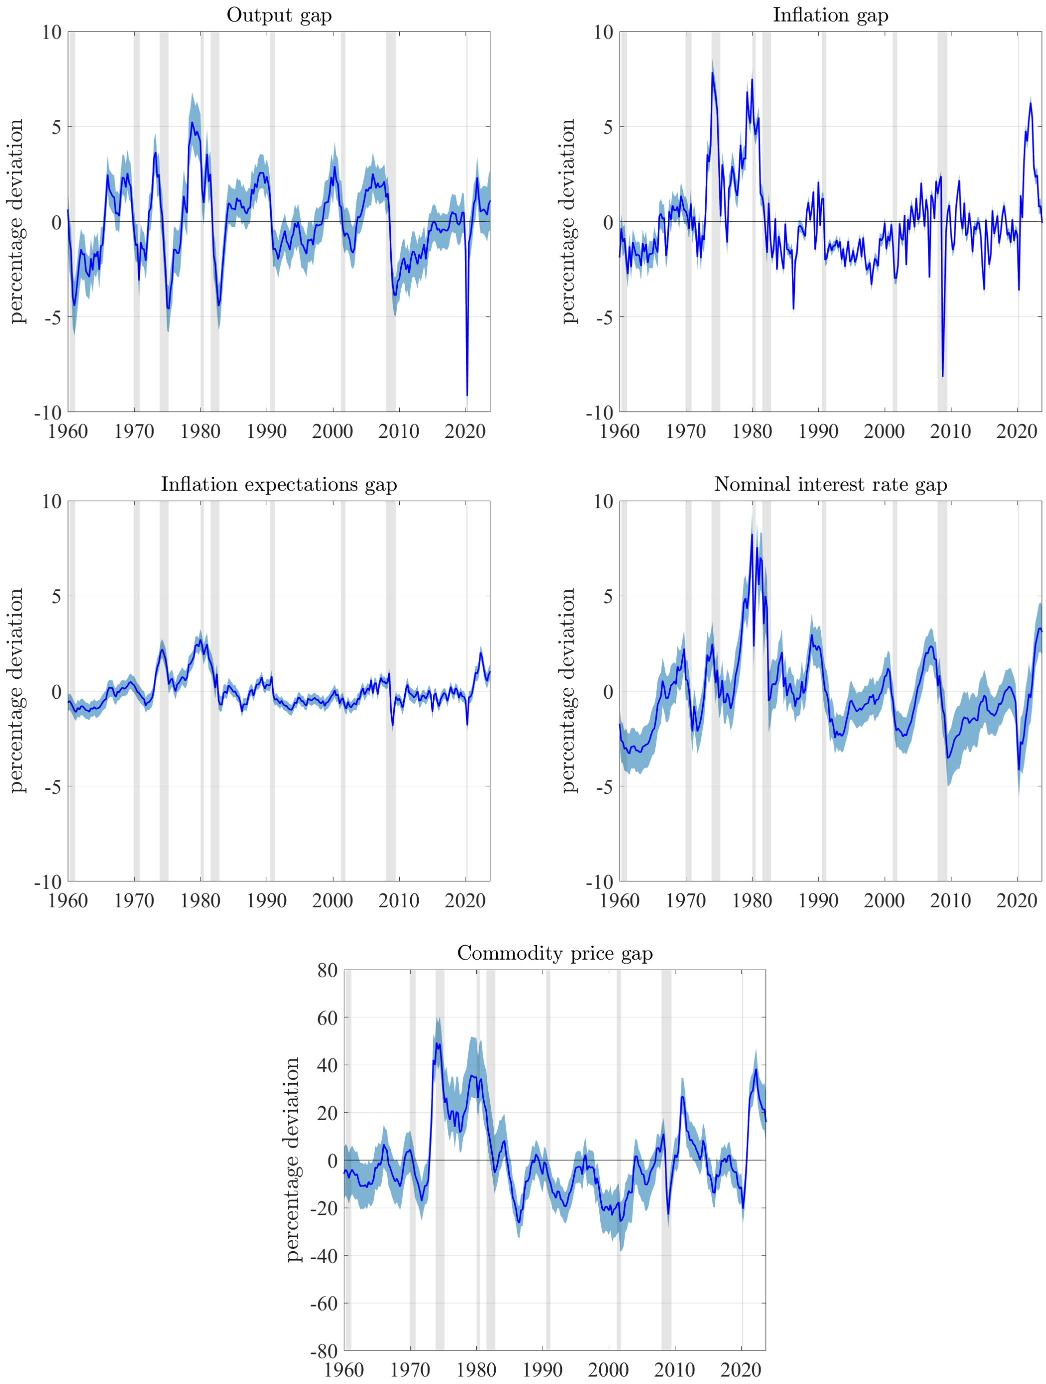

Figure 4: Median estimate (blue solid), 68% coverage bands (blue shaded areas), NBER recession dates (grey shaded areas).

debate throughout the post-pandemic recovery.

First, in the US the recent inflation episode was primarily *demand*-driven.[20] While supply cost-

[20]While findings on the US are consistent across studies – see e.g., by Bergholt et al. (2024a), Eickmeier and Hofmann (2023)

push shocks played an important role, especially in the early phases of the recovery and following the Russian invasion of Ukraine in 2022, business cycle demand and monetary policy were responsible for approximately two thirds of inflationary pressures. Monetary policy (light blue bars) was in fact very accommodating in the aftermath of the pandemic and accounted for a non-negligible portion of the inflation surge. However, with the beginning of the tightening cycle, monetary policy contributed about one percentage point to the disinflation process during 2022 – see the difference in the contribution between 2023Q1 and 2022Q1.

Second, business cycle demand (dark blue bars) was the *main* driver of the US recovery from the Covid-19 pandemic, suggesting, among other things, that the massive fiscal injection of liquidity by Biden's administration in 2020-21 was partly responsible for the overheating of the US economy. This is also corroborated in Mori (2024), who explicitly factors out the contribution of the fiscal stimulus from business cycle demand. Moreover, demand represented the main source of *persistence* of the current business cycle and, while slowly receding, it prevented inflation to return on target in the most recent quarters.

Third, after deviating about 2% from its long-run trend in 2021, the output gap steadily narrowed in 2022, as demand was slowly cooling and the economy seemed to be headed towards a soft landing. Since 2023, however, economic activity rebounded and the output gap widened again. According to the model, the rebound of output gap was primarily due to supply conditions rather than overheating demand. As a matter of fact, demand continued to slowdown in 2023, albeit very slowly. Accordingly, the rebound of economic activity should not be deemed as an additional source of upside risk to the inflation outlook.

Overall, the results suggest that the recent cycle mainly originated from demand, in line with the predictions of new-Keynesian business cycle theory. In addition, the disinflation process observed in the second half of 2022 and in 2023 was favoured by a credible and restrictive monetary policy and the re-balancing of supply conditions.

and Shapiro (2024) –, evidence on the Euro Area is more mixed. Bańbura et al. (2023), for example, document a relatively larger and more persistent contribution of supply, whereas Ascari et al. (2023, 2024) attribute a larger role to demand.

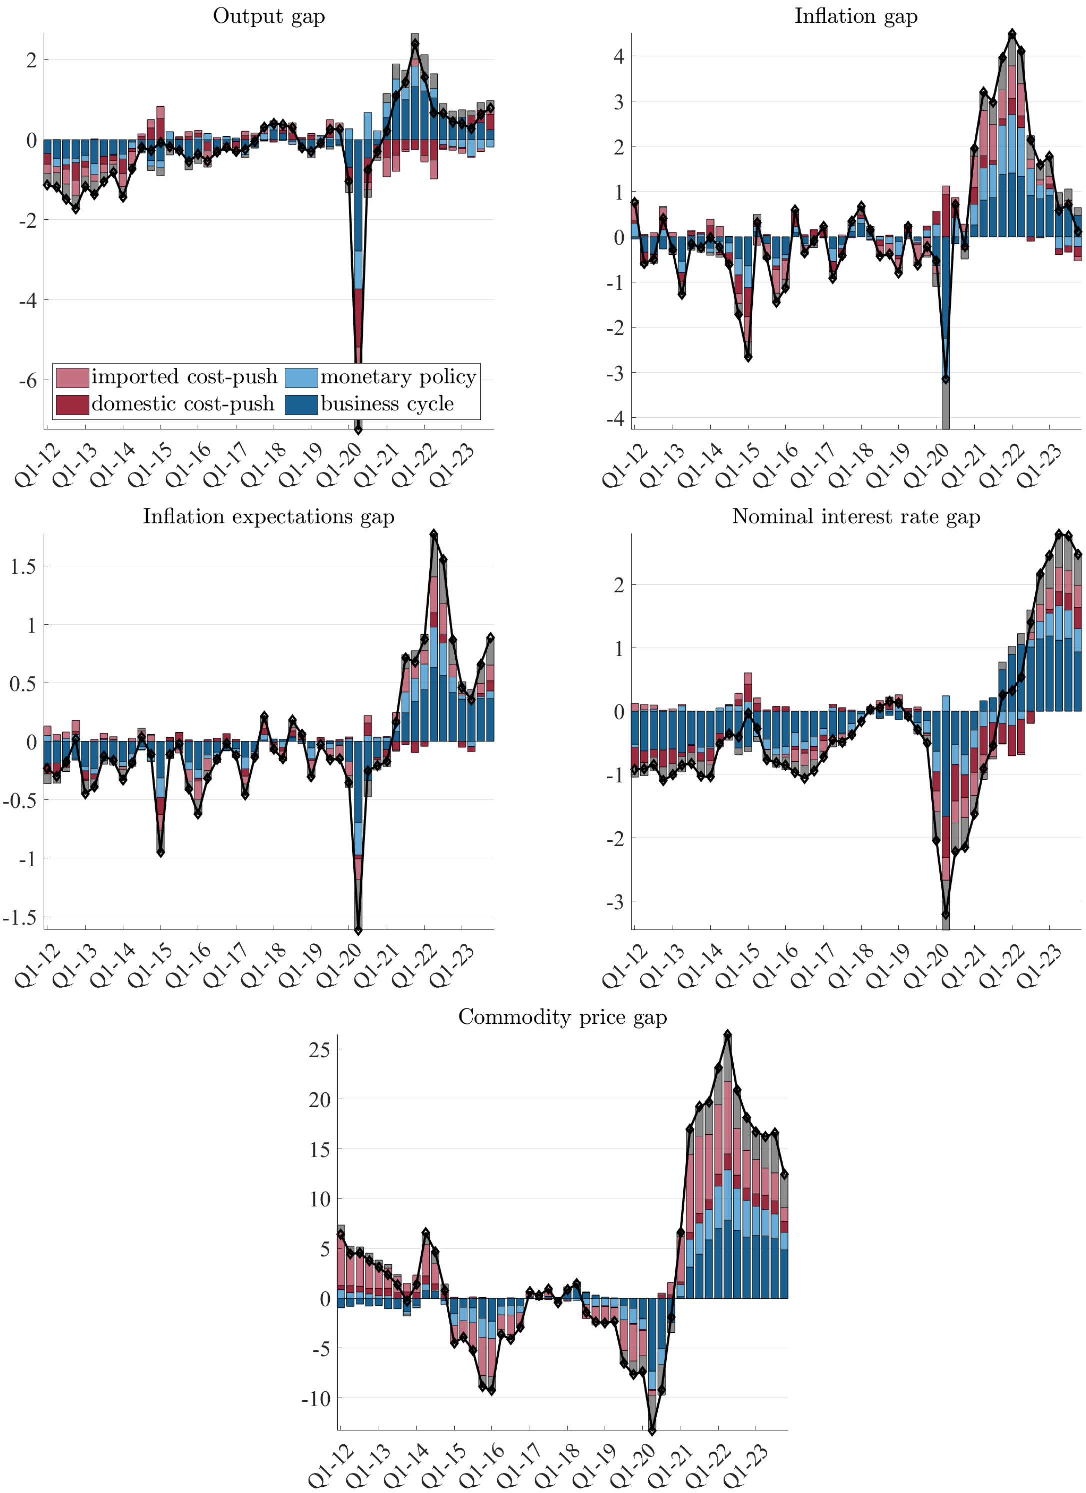

Figure 5: Historical decomposition, percentage deviation from the time-varying trend, median point estimate. Top panels: output gap ($\tilde{x}_t$) and inflation gap ($\tilde{\pi}_t$). Middle panels: inflation expectations gap ($\tilde{\pi}_{e,t}$) and nominal interest rate gap ($\tilde{i}_t$). Bottom panel: commodities price gap ($\tilde{\pi}_{c,t}$). Note: gray bars refer to the residual shock.

3.4 Trend Inflation in Real-Time

The results presented so far demonstrate how the model accounts for the most recent inflation events. However, a comprehensive evaluation of the current inflation dynamics requires to complement this *historical* analysis with a *real-time* assessment. From a policy perspective, the performance of the model in real-time is important, because it complements the historical information of the shocks decomposition with a monitoring of the risk of inflation becoming de-anchored. In previous sections, I argued that augmenting the model with a measure of inflation expectations was beneficial for the estimation of trend inflation. Because inflation is by definition a forward-looking variable, indicators and surveys that measure agents' beliefs about future inflation outcomes are generally deemed to embed valuable information on the low-frequency component of actual inflation – see e.g., Mertens (2016) and Zaman (2021).

This section briefly discusses the model's ability to track trend inflation in real-time and highlights the contribution of inflation expectations in (i) reducing estimation *uncertainty* and (ii) delivering stable *real-time* estimates. For this purpose, I compare the real-time estimate of trend inflation with its historical (or smoothed) estimate – see figure 6. First, note that both the real-time and the smoothed series become more precisely estimated, as inflation expectations become available in 1982. Second, while differences between the smoothed and the real-time estimates are large before 1982, they become negligible thereafter. Since then, in fact, historical revisions are never statistically significant, making real-time estimates incredibly stable. This result confirms and strengthens the previous finding that trend inflation remained anchored at the 2% target throughout the post-pandemic recovery. In addition, one can conclude that the model not only represents a useful apparatus for an historical analysis of the drivers of inflation but, thanks to the strong signal embedded in inflation expectations, also a credible device for real-time analysis.

Finally, while the inclusion of inflation expectations generates clear gains in terms of reduced estimation uncertainty and real-time analysis, a burgeoning body of literature documents as these measures often violate the full information rational expectations assumption – see e.g., Coibion and Gorodnichenko (2012, 2015), Bordalo et al. (2020), Bianchi et al. (2022), Bianchi et al. (2023a), among others. Mertens (2016), Nason and Smith (2021) and more recently Fisher et al. (2024), for example, show the benefits of modelling this feature of the data in a univariate time series setting for the inflation process. However, while this assumption is likely to further improve the predictive power of the model, its implementation in a multivariate framework is non-trivial and beyond the scope of this paper. For this reason, the model

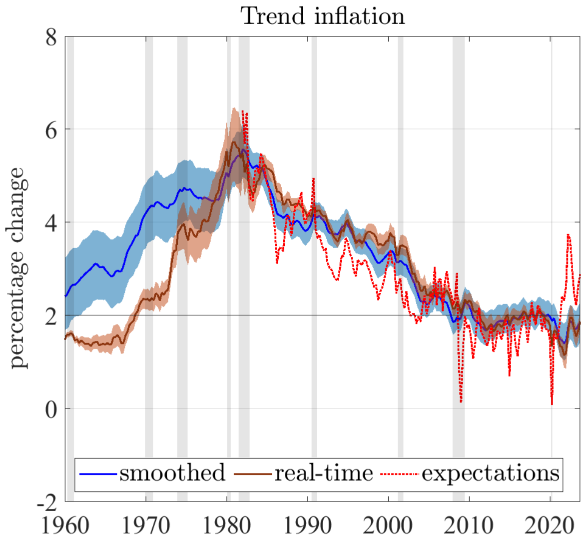

Figure 6: Median estimates (solid blue and brown lines), 68% coverage bands (shaded blue and brown areas), inflation expectations (red solid line), NBER recession dates (grey shaded areas).

preserves the full rational expectations assumption, which is captured by the common trend between actual inflation and one-year ahead inflation expectations.

4 Alternative Assumptions on the Long-run

This section discusses how the baseline results would have changed, had an alternative specification been imposed a dogmatic prior in favor of a deterministic trend. To do so, I essentially estimate a version of the model in Section 2, that switches off the prior variance of the innovations of the unobserved "star" variables. Importantly, note that all the remaining assumptions – such as the theory-based long-run restrictions, the treatment of the zero lower bound periods, etc. – do not change. In this way, any differences between the two models arise solely from the dogmatic assumption against time-variation in the steady state. For conciseness, the discussion focuses only on output and inflation gap. The other results are available in Section B of the Appendix.

Figure 7 plots the model-based estimates of output and inflation gaps. Compared to the baseline estimates in figure 4, the uncertainty surrounding output gap is much larger. Strikingly, since the model is no longer able to accommodate changes in trend growth, it fails to capture the long-lasting effects of the Great Recession. As a consequence, the estimated output gap becomes persistently negative thereafter, suggesting that the economy has operated below potential over the last decade and in the aftermath of the Covid-19 pandemic. Moreover, this result is also in contrast with the strong and rapid rebound of the

US economy in the most recent years. Similarly, inflation gap is negative over most of the 1990-2023

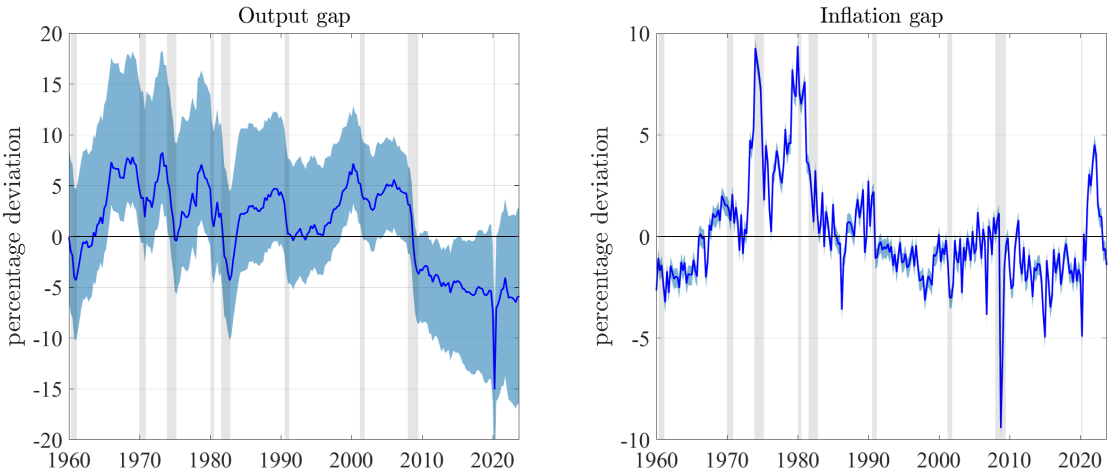

Figure 7: Median estimate (blue solid), 68% coverage bands (blue shaded areas), NBER recession dates (grey shaded areas).

period, with the exception of the post-pandemic recovery. Moreover, the estimated gap turns negative in the first half of 2023, suggesting that inflation was below trend, despite the Bureau of Labor Statistics registered a 2.5% level in the second quarter of 2023.

Figure 8 offers a breakdown of the shocks driving output and inflation gap during the post-pandemic inflation episode. Compared to the baseline results, three main differences emerge. First, the inflation surge is now mainly attributed by supply shocks, especially imported cost-push, which contribute to at least half of the post-pandemic rise. Second, while monetary policy (light blue bars) accounts for a similar share of the output and inflation gap fluctuations to the baseline model, the contribution of business cycle demand (dark blue) is much more marginal than in the baseline. According to the “dogmatic” model, in fact, it only exerts very mild upside pressures on inflation. Third, the negative output gap is almost entirely attributed to large and persistent domestic cost-push shocks.

Overall, these results draw very different policy conclusions, relative to the baseline model. For example, the limited contribution of business cycle demand (dark blue) would suggest that the largely accommodating fiscal stance of Biden’s administration was insufficient in stimulating private demand. Moreover, as the model imputes a relatively larger importance to supply-driven inflationary pressures, the Fed’s monetary tightening appears to be disproportionate and partly responsible for the negative inflation gap in 2023.

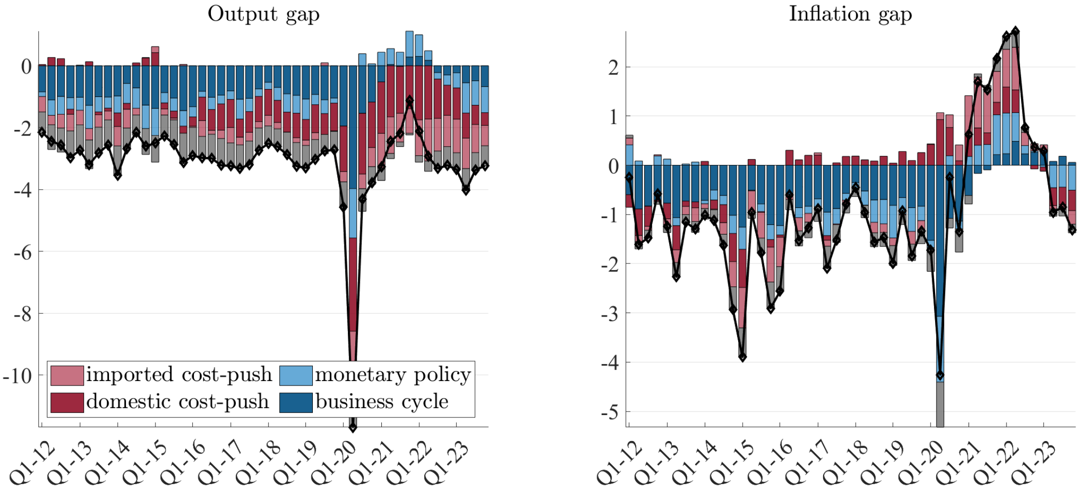

Figure 8: Historical decomposition, percentage deviation from the time-varying steady state, median point estimate. Top panel: output gap ($\tilde{x}_t$). Middle panel: inflation gap ($\tilde{\pi}_t$). Bottom panel: inflation expectations gap ($\tilde{\pi}_{e,t}$). Note: gray bars refer to the residual shock.

5 Conclusions

This paper proposes a time series model to jointly study trend and cyclical dynamics of key US macroeconomic aggregates. The model differs from standard VARs along two key dimensions: (i) trends are stochastic; (ii) long-run comovements are derived from economic theory. These two ingredients make inference on the cyclical properties of the data robust to slow-moving fluctuations arising from persistent transitional dynamics and/or structural breaks in the unconditional mean.

The results provide a reasonable account of the economic developments in the post-WWII period and of the most recent post-pandemic inflation episode. Importantly, its robustness to the so-called end-point problem, makes the model a valuable tool, not only for *historical* analysis, but also for *real-time* monitoring a trend inflation.

Overall, this model constitutes a comprehensive framework to offer policy guidance and a flexible empirical counterpart to heavily-parametrized DSGE models for the analysis of business cycle fluctuations.

References

G.-M. Angeletos, F. Collard, and H. Dellas. Business-cycle anatomy. American Economic Review, 110(10):3030–70, October 2020. doi: 10.1257/aer.20181174. URL https://www.aeaweb.org/articles?id=10.1257/aer.20181174.

J. Antolin-Diaz, T. Drechsel, and I. Petrella. Tracking the slowdown in long-run gdp growth. The Review of Economics and Statistics, 99(2):343–356, 05 2017. ISSN 0034-6535. doi: 10.1162/REST_a_00646. URL https://doi.org/10.1162/REST_a_00646.

J. E. Arias, J. F. Rubio-Ramírez, and D. Waggoner. Inference based on structural vector autoregressions identified with sign and zero restrictions: Theory and applications. Econometrica, 86(2):685–720, 2018. URL https://EconPapers.repec.org/RePEc:wly:emetrp:v:86:y:2018:i:2:p:685-720.

G. Ascari and L. Fosso. The international dimension of trend inflation. Journal of International Economics, 2024.

G. Ascari, P. Bonomolo, M. Hoeberichts, and R. Trezzi. The euro area great inflation surge. Dnb analysis, De Nederlandsche Bank, Feb. 2023. URL https://www.dnb.nl/en/publications/research-publications/analysis/the-euro-area-great-inflation-surge/.

G. Ascari, D. Bonam, L. Mori, and A. Smadu. Fiscal Policy and Inflation in the Euro Area. Working Papers 820, DNB, Nov. 2024. URL https://ideas.repec.org/p/dnb/dnbwpp/820.html.

L. M. Ball, M. D. Leigh, and M. P. Mishra. Understanding U.S. Inflation During the COVID Era. IMF Working Papers 2022/208, International Monetary Fund, Oct. 2022. URL https://ideas.repec.org/p/imf/imfwpa/2022-208.html.

C. Baumeister, D. Korobilis, and T. K. Lee. Energy markets and global economic conditions. The Review of Economics and Statistics, 104(4):828–844, 07 2022. ISSN 0034-6535. doi: 10.1162/rest_a_00977. URL https://doi.org/10.1162/rest_a_00977.

M. Bańbura, E. Bobeica, and C. Martínez Hernández. What drives core inflation? The role of supply shocks. Working Paper Series 2875, European Central Bank, Nov. 2023. URL https://ideas.repec.org/p/ecb/ecbwps/20232875.html.

P. Beaudry, D. Galizia, and F. Portier. Putting the cycle back into business cycle analysis. American Economic Review, 110(1):1–47, January 2020. doi: 10.1257/aer.20190789. URL https://www.aeaweb.org/articles?id=10.1257/aer.20190789.

D. Bergholt, F. Canova, F. Furlanetto, N. Maffei-Faccioli, and P. Ulvedal. What Drives the Recent Surge in Inflation? The Historical Decomposition Roller Coaster. CEPR Discussion Paper 19005, CEPR, Apr. 2024a. URL https://cepr.org/publications/dp19005.

D. Bergholt, L. Fosso, and F. Furlanetto. Macroeconomic effects of the gender revolution. Working Paper Series 19, Norges Bank, 2024b. URL https://www.norges-bank.no/contentassets/16ce44f18bc84b229ef655e56434cba2/wp-19_2024.pdf?v=19122024110454.

F. Bianchi, S. C. Ludvigson, and S. Ma. Belief distortions and macroeconomic fluctuations. American Economic Review, 112(7):2269–2315, July 2022. doi: 10.1257/aer.20201713. URL https://www.aeaweb.org/articles?id=10.1257/aer.20201713.

F. Bianchi, C. Ilut, and H. Saijo. Diagnostic Business Cycles. The Review of Economic Studies, 91(1): 129–162, 02 2023a. ISSN 0034-6527. doi: 10.1093/restud/rdad024. URL https://doi.org/10.1093/restud/rdad024.

F. Bianchi, G. Nicolò, and D. Song. Inflation and Real Activity over the Business Cycle. NBER Working Papers 31075, National Bureau of Economic Research, Inc, Mar. 2023b. URL https://ideas.repec.org/p/nbr/nberwo/31075.html.

H. C. Bjørnland, L. Brubakk, and N. Maffei-Faccioli. Piecing the Puzzle: Real Exchange Rates and Long-Run Fundamentals. CAMA Working Papers 2025-10, Centre for Applied Macroeconomic Analysis, Crawford School of Public Policy, The Australian National University, Feb. 2025. URL https://ideas.repec.org/p/een/camaaa/2025-10.html.

O. Blanchard, E. Cerutti, and L. H. Summers. Inflation and Activity: Two Explorations and Their Monetary Policy Implications. Working Paper Series WP15-19, Peterson Institute for International Economics, Nov. 2015. URL https://ideas.repec.org/p/iie/wpaper/wp15-19.html.

P. Bordalo, N. Gennaioli, Y. Ma, and A. Shleifer. Overreaction in macroeconomic expectations. American Economic Review, 110(9):2748–82, September 2020. doi: 10.1257/aer.20181219. URL https://www.aeaweb.org/articles?id=10.1257/aer.20181219.

M. D. Bordo and T. F. Helbling. International business cycle synchronization in historical perspective. Working Paper 16103, National Bureau of Economic Research, June 2010. URL http://www.nber.org/papers/w16103.

D. M. Byrne, J. G. Fernald, and M. B. Reinsdorf. Does the united states have a productivity slowdown or a measurement problem? Brookings Papers on Economic Activity, pages 109–157, 2016. ISSN 00072303, 15334465. URL http://www.jstor.org/stable/43869021.

F. Canova, D. Lopez-Salido, and C. Michelacci. The effects of technology shocks on hours and output: a robustness analysis. Journal of Applied Econometrics, 25(5):755–773, 2010. doi: https://doi.org/10.1002/jae.1090. URL https://onlinelibrary.wiley.com/doi/abs/10.1002/jae.1090.

F. Canova, D. Lopez-Salido, and C. Michelacci. The ins and outs of unemployment: An analysis conditional on technology shocks. The Economic Journal, 123(569):515–539, 2013. doi: https://doi.org/10.1111/j.1468-0297.2012.02548.x. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1468-0297.2012.02548.x.

V. Cerra, A. Fatás, and S. C. Saxena. Hysteresis and business cycles. Journal of Economic Literature, 61(1):181–225, March 2023. doi: 10.1257/jel.20211584. URL https://www.aeaweb.org/articles?id=10.1257/jel.20211584.

A. Cesa-Bianchi, R. Harrison, and R. Sajedi. Decomposing the drivers of Global R. Bank of England working papers 990, Bank of England, July 2022. URL https://ideas.repec.org/p/boe/boeewp/0990.html.

G. Cette, J. Fernald, and B. Mojon. The pre-Great Recession slowdown in productivity. European Economic Review, 88(C):3–20, 2016. doi: 10.1016/j.euroecorev.2016. URL https://ideas.repec.org/a/eee/eecrev/v88y2016icp3-20.html.

O. Coibion and Y. Gorodnichenko. What can survey forecasts tell us about information rigidities? Journal of Political Economy, 120(1):116–159, 2012. ISSN 00223808, 1537534X. URL http://www.jstor.org/stable/10.1086/665662.

O. Coibion and Y. Gorodnichenko. Information rigidity and the expectations formation process: A simple framework and new facts. The American Economic Review, 105(8):2644–2678, 2015. ISSN 00028282. URL http://www.jstor.org/stable/43821351.

R. K. Crump, S. Eusepi, M. Giannoni, and A. Sahin. A unified approach to measuring ustar. Brookings Papers on Economic Activity, pages 143–239, 2019.

M. Del Negro, F. Schorfheide, F. Smets, and R. Wouters. On the fit of new keynesian models. Journal of Business & Economic Statistics, 25(None):123–143, April 2007. doi: None. URL https://ideas.repec.org/a/bes/jnlbes/v25y2007p123-143.html.

M. Del Negro, D. Giannone, M. Giannoni, and A. Tambalotti. Safety, liquidity, and the natural rate of interest. Brookings Papers on Economic Activity, 48(1 (Spring)):235–316, 2017.

S. Delle Chiaie, L. Ferrara, and D. Giannone. Common factors of commodity prices. Journal of Applied Econometrics, 37(3):461–476, 2022. doi: https://doi.org/10.1002/jae.2887. URL https://onlinelibrary.wiley.com/doi/abs/10.1002/jae.2887.

J. Durbin and S. J. Koopman. A simple and efficient simulation smoother for state space time series analysis. Biometrika, 89(3):603–615, 2002. ISSN 00063444. URL http://www.jstor.org/stable/4140605.

S. Eickmeier and B. Hofmann. What drives inflation? Disentangling demand and supply factors. CEPR Discussion Paper 18378, CEPR, Aug. 2023. URL https://cepr.org/publications/dp18378.

J. Fernald and S. Ramnath. The acceleration in u.s. total productivity after 1995: the role of information technology. Economic Perspectives, 28(Q I):52–67, 2004. URL https://EconPapers.repec.org/RePEc:fip:fedhep:y:2004:i:qi:p:52-67:n:v.28no.1.

J. Fernald, R. Inklaar, and D. Ruzic. The productivity slowdown in advanced economies: Common shocks or common trends? Review of Income and Wealth, n/a(n/a), 2023. doi: https://doi.org/10.1111/roiw.12690. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/roiw.12690.

J. G. Fernald. Trend breaks, long-run restrictions, and contractionary technology improvements. Journal of Monetary Economics, 54(8):2467–2485, November 2007. URL https://ideas.repec.org/a/eee/moneco/v54y2007i8p2467-2485.html.

J. G. Fernald. Productivity and Potential Output before, during, and after the Great Recession. NBER Macroeconomics Annual, 29(1):1–51, 2015. doi: 10.1086/680580. URL https://ideas.repec.org/a/ucp/macann/doi10.1086-680580.html.

J. Fisher, L. Melosi, and S. Rast. Long-Run Inflation Expectations. Mimeo, Jan. 2024.

J. Galí. Monetary policy, inflation, and the business cycle: An introduction to the new keynesian framework. In Princeton University Press. 2008. URL https://EconPapers.repec.org/RePEc:pup:chapts:8654-1.

D. Giannone and G. Primiceri. The drivers of post-pandemic inflation. Working Paper 32859, National Bureau of Economic Research, August 2024. URL http://www.nber.org/papers/w32859.

D. Giannone, M. Lenza, and G. Primiceri. Prior selection for vector autoregressions. The Review of Economics and Statistics, 97(2):436–451, 2015. URL https://EconPapers.repec.org/RePEc:tpr:restat:v:97:y:2015:i:2:p:436-451.

D. Giannone, M. Lenza, and G. E. Primiceri. Priors for the Long Run. Journal of the American Statistical Association, 114(526):565–580, April 2019. doi: 10.1080/01621459.2018.148. URL https://ideas.repec.org/a/taf/jnlasa/v114y2019i526p565-580.html.

M. Goodfriend and R. G. King. The incredible Volcker disinflation. Journal of Monetary Economics, 52(5):981–1015, July 2005. URL https://ideas.repec.org/a/eee/moneco/v52y2005i5p981-1015.html.

A. L. Grant and J. C. Chan. A bayesian model comparison for trend-cycle decompositions of output. Journal of Money, Credit and Banking, 49(2-3):525–552, 2017. doi: https://doi.org/10.1111/jmcb.12388. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/jmcb.12388.

I. Hajdini, A. Shapiro, A. Smith, and D. Vallenas. Inflation since the pandemic: Lessons and challenges. Finance and Economics Discussion Series, pages 1–1, 08 2025. doi: 10.17016/FEDS.2025.070.

J. D. Hamilton. Causes and consequences of the oil shock of 2007-08. Working Paper 15002, National Bureau of Economic Research, May 2009. URL http://www.nber.org/papers/w15002.

T. Hasenzagl, F. Pellegrino, L. Reichlin, and G. Ricco. A Model of the Fed's View on Inflation. The Review of Economics and Statistics, 104(4):686–704, 07 2022. ISSN 0034-6535. doi: 10.1162/rest_a_00974. URL https://doi.org/10.1162/rest_a_00974.

K. Holston, T. Laubach, and J. C. Williams. Measuring the natural rate of interest: International trends and determinants. Journal of International Economics, 108(S1):59–75, 2017. doi: 10.1016/j.jinteco.

2017.01. URL https://ideas.repec.org/a/eee/inecon/v108y2017is1ps59-s75.html.

L. Kilian. Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. American Economic Review, 99(3):1053–69, June 2009. doi: 10.1257/aer.99.3.1053. URL https://www.aeaweb.org/articles?id=10.1257/aer.99.3.1053.

T. A. Lubik, C. Matthes, and F. Verona. Assessing U.S. aggregate fluctuations across time and frequencies. Bank of Finland Research Discussion Papers 5/2019, Bank of Finland, 2019. URL https://ideas.repec.org/p/zbw/bofrdp/rdp2019_005.html.

N. Maffei-Faccioli. Identifying the sources of the slowdown in growth: Demand vs. supply. Working Paper 2021/9, Norges Bank, Sept. 2021. URL https://ideas.repec.org/p/bno/worpap/2021_9.html.

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

L. Mori. Fiscal shocks and the surge of inflation. University of Padova, available at https://sites. google.com/view/lorenzomori, 2024.

J. C. Morley, C. R. Nelson, and E. Zivot. Why are the beveridge-nelson and unobserved-components decompositions of gdp so different? The Review of Economics and Statistics, 85(2):235–243, 05 2003. ISSN 0034-6535. doi: 10.1162/003465303765299765. URL https://doi.org/10.1162/003465303765299765.

J. M. Nason and G. W. Smith. Measuring the slowly evolving trend in us inflation with professional forecasts. Journal of Applied Econometrics, 36(1):1–17, 2021. doi: https://doi.org/10.1002/jae.2784. URL https://onlinelibrary.wiley.com/doi/abs/10.1002/jae.2784.

P. Perron. The great crash, the oil price shock, and the unit root hypothesis. Econometrica, 57(6): 1361–1401, 1989. ISSN 00129682, 14680262. URL http://www.jstor.org/stable/1913712.

C. Pinilla-Torremocha. Persistent and transitory inflation in the euro area: Insights from global and domestic shocks. mimeo, 2025.

A. H. Shapiro. Decomposing Supply and Demand Driven Inflation. Journal of Money, Credit and Banking, 2024. Forthcoming.

C. A. Sims. Using a likelihood perspective to sharpen econometric discourse: Three examples. Journal of Econometrics, 95(2):443–462, April 2000. URL https://ideas.repec.org/a/eee/econom/v95y2000i2p443-462.html.

J. H. Stock and M. W. Watson. Testing for common trends. Journal of the American Statistical Association, 83(404):1097–1107, 1988. doi: 10.1080/01621459.1988.10478707. URL https://www.tandfonline.com/doi/abs/10.1080/01621459.1988.10478707.

J. H. Stock and M. W. Watson. Why Has U.S. Inflation Become Harder to Forecast? Journal of Money, Credit and Banking, 39(s1):3–33, February 2007. URL https://ideas.repec.org/a/mcb/jmoncb/v39y2007is1p3-33.html.

L. Summers. Demand side secular stagnation. American Economic Review, 105(5):60–65, 2015. URL https://EconPapers.repec.org/RePEc:aea:aecrev:v:105:y:2015:i:5:p:60-65.

M. Villani. Steady-state priors for vector autoregressions. Journal of Applied Econometrics, 24(4):630–650, 2009. ISSN 08837252, 10991255. URL http://www.jstor.org/stable/40206295.

M. W. Watson. Univariate detrending methods with stochastic trends. Journal of Monetary Economics, 18(1):49–75, July 1986. URL https://ideas.repec.org/a/eee/moneco/v18y1986i1p49-75.html.

M. Woodford. Interest and Prices: Foundations of a Theory of Monetary Policy. Princeton University Press, 2003. URL http://www.jstor.org/stable/j.ctv30pnvmf.

S. Zaman. A Unified Framework to Estimate Macroeconomic Stars. Working Papers 21-23R2, Federal Reserve Bank of Cleveland, Oct. 2021. URL https://ideas.repec.org/p/fip/fedcwq/93166.html.

Appendix

A The Empirical Model in Detail

This section exposes the state-space representation of the model together with the prior assumptions on the unknown parameters and the Gibbs sampling algorithm to estimate the latent states.

A.1 The State-space Representation

The state-space is expressed in terms of a measurement equation:

$$y_t = \Lambda\, \chi_t, \tag{A.1}$$

and a transition equation, which evolves according to:

$$\chi_t = \Phi\, \chi_{t-1} + \mathcal{R}\, \nu_t \tag{A.2}$$

where $\chi_t = [\underset{q^* \times 1}{y_t^*} \quad \underset{\tilde{q} \times 1}{\tilde{y}_t}]$ stacks the latent star and gap variables.

$\Lambda = [\underset{n \times q^*}{\Lambda^*} \quad \underset{n \times \tilde{q}p}{\tilde{\Lambda}}]$, $\tilde{\Lambda} = [\underset{n \times \tilde{q}}{\tilde{\Lambda}_0} \quad \underset{n \times \tilde{q}}{\tilde{\Lambda}_1} \quad \cdots \quad \underset{n \times \tilde{q}}{\tilde{\Lambda}_{p-1}}]$ and p are the lags of the autoregressive coefficients of the law of motion of $\tilde{y}_t$. The state matrices are:

$$\Phi = \begin{bmatrix} I & 0 \\ 0 & \tilde{\Phi} \end{bmatrix}, \qquad \mathcal{R} = \begin{bmatrix} I & 0 \\ 0 & \tilde{R} \end{bmatrix},$$

$$\tilde{\Phi} = \begin{bmatrix} \tilde{\Phi}_1 & \tilde{\Phi}_2 & \tilde{\Phi}_3 & \dots & \tilde{\Phi}_p \\ I & 0 & \dots & \dots & 0 \\ 0 & I & \dots & \ddots & 0 \\ \vdots & \dots & \ddots & \vdots & \\ 0 & \dots & \dots & I & 0 \end{bmatrix}, \qquad \tilde{R} = \begin{bmatrix} I \\ 0 \\ \vdots \\ 0 \\ 0 \end{bmatrix}$$

.

Finally, note that in the model specification presented in Section 2,

$$\tilde{\Lambda}_0 = \begin{bmatrix} 4 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 4 \end{bmatrix}, \qquad \tilde{\Lambda}_1 = \begin{bmatrix} -4 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & -4 \end{bmatrix}$$

where the first and the last entries in the main diagonal of $\tilde{\Lambda}_0$ and $\tilde{\Lambda}_1$ inform the state-space to estimate the deviations of real output and commodity prices from their log-level trends. This enables the model to estimate a measure of the output gap and of the commodity price cycle. The entries are normalized to 4, because real output and commodity prices enter the model in annualized terms.

A.2 Prior Assumptions

The initial conditions of the steady states are distributed according to $y_0^* \sim \mathcal{N}(\underline{y_0^*}, I_q)$. The prior means $\underline{y_0^*}$ are defined using averages of input data over the pre-sample period, which spans 1954Q1-1959Q4. The initial conditions of the cycles are distributed according to $\tilde{y}_0 \sim \mathcal{N}(0_n, I_n)$. This assumption implies that cycles fluctuate symmetrically around a zero mean. Finally, the priors for the remainder model's coefficients are distributed according to:

$$\Sigma^* \sim \mathcal{IW}(\kappa^*, (\kappa^* + n + 1)\underline{\Sigma}^*) \tag{A.3}$$

$$\tilde{\Sigma} \sim \mathcal{IW}(\tilde{\kappa}, (\tilde{\kappa} + n + 1)\underline{\tilde{\Sigma}}) \tag{A.4}$$

$$vec(\tilde{\Phi})|\tilde{\Sigma} \sim \mathcal{N}(vec(\underline{\tilde{\Phi}}), \tilde{\Sigma} \otimes \underline{\Omega})\mathcal{I}(\tilde{\Phi}), \tag{A.5}$$

where and $\mathcal{I}(\tilde{\Phi})$ is an indicator function that is equal to one, when the VAR of the cycle block is stationary, zero otherwise. $\mathcal{IW}$ is the Inverse-Wishart distribution with κ degrees of freedom and mode $\underline{\Sigma}$. The prior mode of trend innovations $\underline{\Sigma}^*$ is assumed to be diagonal. The priors on the main diagonal entries are conservative in limiting the amount of variance attributable to the trends. Similar to Del Negro et al. (2017), all prior volatilities are normalized by 60, which implies that the standard deviation of the expected change in the trends over six decades is only one percentage point. The degrees of freedom $\kappa^* = 100$ imply a rather tight prior around the mode. Moving to the cycle block, the priors for the lag coefficients are standard Minnesota with an overall tightness hyperparameter equal to 0.2, following Giannone et al. (2015), and the own-lag hyperparameters centered around zero, instead of one, since this is the stationary block. The prior mode of the transitory innovations $\underline{\tilde{\Sigma}}$ is assumed to be an identity matrix and rather uninformative, as the degrees of freedom are $\tilde{\kappa} = n + 2$.

A.3 Estimation of the State Space with Gibbs Sampling

Consider the unobserved states of the above model in the following stacked formulation:

$$\begin{bmatrix} y_t^* \\ \tilde{y}_t \end{bmatrix} = \begin{bmatrix} I & 0 \\ 0 & \tilde{\Phi} \end{bmatrix} \begin{bmatrix} y_{t-1}^* \\ \tilde{y}_{t-1} \end{bmatrix} + \begin{bmatrix} I & 0 \\ 0 & I \end{bmatrix} \begin{bmatrix} u_t^* \\ \tilde{u}_t \end{bmatrix} \tag{A.6}$$

and the covariance matrix of the model is given by Σ:

$$\Sigma = \begin{bmatrix} \Sigma^* & 0 \\ 0 & \tilde{\Sigma} \end{bmatrix} \tag{A.7}$$

Then, the model samples 150000 draws and throws the first 25000 draws from a Gibbs algorithm, according to the following steps:

1. Draw from the joint distribution $y_{0:T}^*, \tilde{y}_{-p+1:T} \mid \tilde{\Phi}, \Sigma^*, \tilde{\Sigma}, Y_{1:T}$ implementing Durbin and Koopman (2002) simulation smoothing algorithm.

2. Draw from the joint distribution $\tilde{\Phi}, \Sigma^*, \tilde{\Sigma} \mid y_{0:T}^*, \tilde{y}_{-p+1:T}, y_{1:T}$. The estimation of the remaining parameters is relatively straightforward, provided that the unobserved states follow rather standard vector autoregressive laws of motion.

   (a) **Star Block.** the posterior distribution of Σ^* is given by:

   $$p(\Sigma^* \mid y_{0:T}^*) = \mathcal{IW}(\underline{\Sigma}^* + \underbrace{\sum_{t=1}^{T}(y_t^* - y_{t-1}^*)(y_t^* - y_{t-1}^*)'}_{S^*}, \kappa^* + T)$$

   The posterior distribution of $y_{0:T}^*$ is obtained from a standard Normal.

   (b) **Gap Block.** The posterior distributions of the lag coefficients in $\tilde{\Phi}$ and the covariance matrix $\tilde{\Sigma}$ of the stationary VAR are standard:

   $$p(\tilde{\Sigma} \mid \tilde{y}_{0:T}) = \mathcal{IW}(\underline{\tilde{\Sigma}} + \tilde{S}, \tilde{\kappa} + T)$$

$$p(\tilde{\Phi} \mid \tilde{\Sigma}, \tilde{y}_{0:T}) = N\left(vec(\tilde{\Psi}), \tilde{\Sigma} \otimes \left(\sum_{t=1}^{T} \tilde{Z}_t \tilde{Z}_t' + \underline{\Omega}^{-1}\right)^{-1}\right)$$

where $\tilde{Z}_t = (\tilde{y}_{t-1}', \ldots, \tilde{y}_{t-p}')$,

$\tilde{\Psi} = \left(\sum_{t=1}^{T} \tilde{Z}_t \tilde{Z}_t' + \underline{\Omega}^{-1}\right)^{-1} \left(\sum_{t=1}^{T} \tilde{Z}_t \tilde{y}_t' + \underline{\Omega}^{-1} \underline{\tilde{\Phi}}\right)$,

$\tilde{S} = \sum_{t=1}^{T} \tilde{u}_t \tilde{u}_t' + (\tilde{\Psi} - \underline{\Phi})' \underline{\Omega}^{-1} (\tilde{\Psi} - \underline{\Phi})$.

B “Dogmatic” deterministic trend model

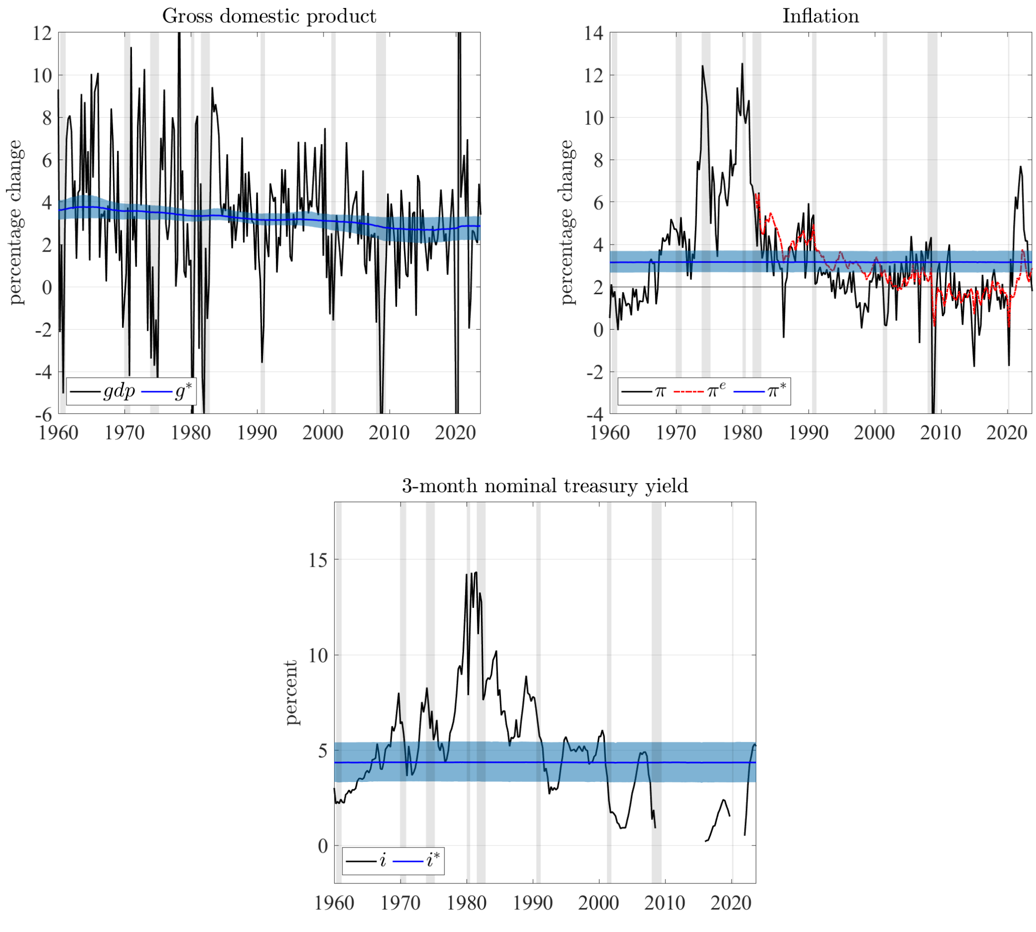

Figure B.1: Data (black dashed), median estimate (blue solid), 68% coverage bands (blue shaded areas), NBER recession dates (grey shaded areas).

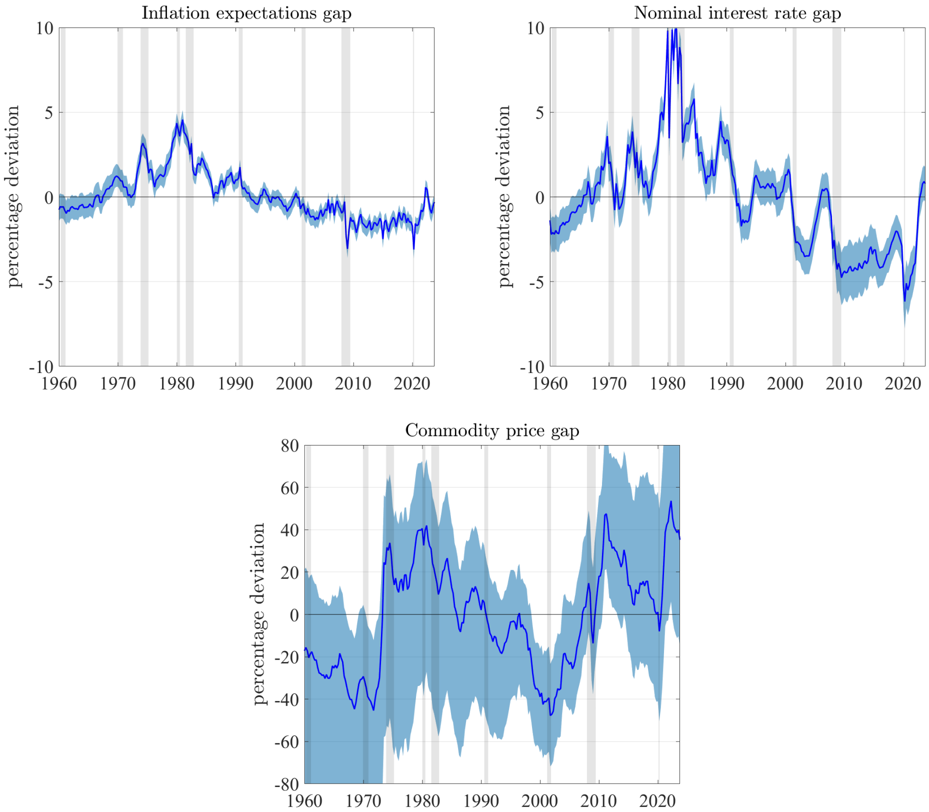

Figure B.2: Median estimate (blue solid), 68% coverage bands (blue shaded areas), NBER recession dates (grey shaded areas).

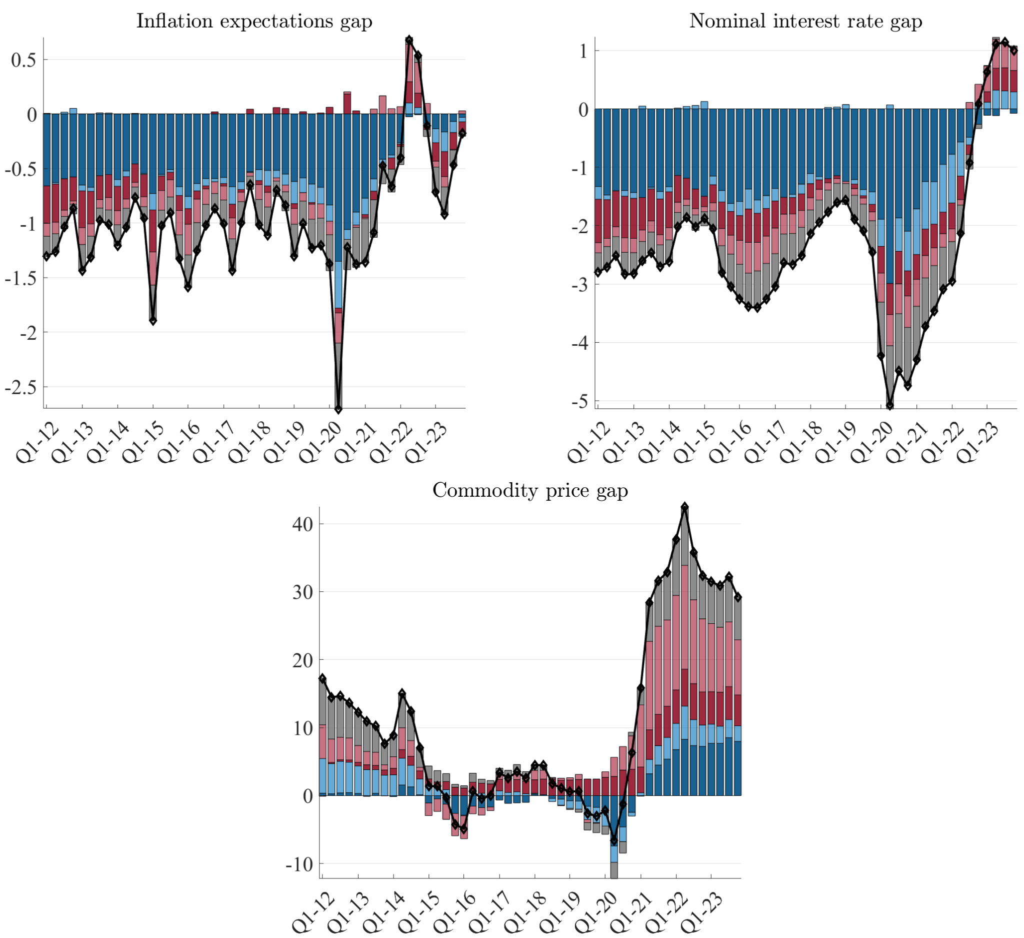

Figure B.3: Historical decomposition, percentage deviation from the time-varying trends, median point estimate. Top panel: inflation expectations gap ($\tilde{\pi}_{e,t}$). Middle panel: nominal interest rate gap ($\tilde{i}_t$). Bottom panel: commodities price gap ($\tilde{\pi}_{c,t}$). Note: grey bars refer to the residual shock.

Acknowledgements

I thank Guido Ascari, Fabio Canova, Efrem Castelnuovo, Michael Ehrmann, Björn Fischer, Francesco Furlanetto, Michele Lenza, Bernd Schnatz for the useful comments.

Luca Fosso
European Central Bank, Frankfurt am Main, Germany; email: luca.fosso@ecb.europa.eu

Postal address 60640 Frankfurt am Main, Germany
Telephone +49 69 1344 0
Website www.ecb.europa.eu

PDF ISBN 978-92-899-7483-7 ISSN 1725-2806 doi: 10.2866/3420463 QB-01-25-232-EN-N