# Essays in Empirical Macroeconomics

PhD Thesis by

**Luca Fosso**

UNIVERSITÀ
DI PAVIA

Supervisor: **Prof. Guido Ascari**
Dipartimento di Scienze Economiche e aziendali
Università degli Studi di Pavia
August 2022

*to Maria Luisa,*

*without whom I would have never started this journey*

# ACKNOWLEDGEMENTS

I am immensely indebted to Guido. Without his guidance, support, commitment and above all patience I would have never been able to go through these challenging years. Thanks to our innumerable conversations I had the chance to sharpen and greatly expand my understanding of economics. He is a great source of inspiration and it is a privilege being one of his students.

I am also deeply thankful to Francesco. During my period in Oslo, he was way more than a mentor. He is an inexorable source of great research ideas and there hasn't been a day with him that has not been enriching. I am truly proud of being one of his mentees, co-authors and friends.

I would like to express my gratitude to the Norges Bank family, who welcomed me as part of the unit from the very first day of my internship. A special mention goes to Drago, who was a fantastic co-mentor. He is one of the brightest economists I have ever met - a true rockstar - and I am grateful for having the opportunity to work with him. A big thank to Kasper, Knut Are, Nicolò, Eleonora, Mikkel, Ørjan, Pål, Akash, Tobi and all the other friends and colleagues. I will always keep fond memories of my period at Norges Bank. It would be unfair if I would not thank Giuse for his great company and for keeping Francesco and I cheerful with his amazing recipes.

Special thanks go to all the people I met at the ECB. Alessandro for his wise pieces of advice during lunch-breaks together. Peter for our endless exchange of ideas about research and economics. Massimo, Michele and Simona for believing in me. Also, thanks to all IPA people who made my traineeship a truly remarkable experience.

I would also like to express my gratitude to the Department of Economics of the University of Pavia, which always supported me in the pursuit of this PhD.

Dulcis in fundo, I will always be immensely indebted to my beloved family, whose unconditional love gave me the strength to overcome any challenges I encountered. I am fully aware of the incredible effort they put to see my dreams come true and I will always be grateful.

# SUMMARY

In chapter 1, we present a trend-cycle BVAR decomposition to explore the role of imported intermediate input prices in explaining the puzzling behavior of inflation in the US over the last decades. The main contribution is twofold: (i) propose a unified framework to jointly study the slow-moving and the business cycle dynamics of inflation and (ii) provide compelling empirical evidence in favour of an important role played by imported input prices in affecting not only the *cyclical* component, but also the *trend* component of inflation.[1]

Chapter 2 evaluates the macroeconomic effects of the secular reduction in gender inequality (measured in terms of female-to-male employment and wage ratios) in the US labor market. For this purpose, we design a SVAR that maps empirical trends into selected (macro and gender-specific) structural trends. Identification is achieved by informing the VAR with the analytical solutions derived from a neoclassical model with gender. We find that gender-specific labor demand shocks are the dominant drivers of gender wage gap and employment gap and play a key role in driving the trend component in employment and in GDP, thus being important contributors of growth in the US economy. Gender-specific labor supply shocks are minor in the baseline model but become important once the skill dimension is taken into account: a positive shock to the supply of female skilled workers is compensated by a negative shock to the supply of unskilled female workers.[2]

Chapter 3 revisits the nexus between EA fundamentals and oil price spikes. In a recursively identified BVAR augmented with Baumeister and Hamilton (2019) shocks, we build an empirical laboratory to test if the feedback of the EA macroeconomy to an oil price shock is dependent on the nature of the underlying shock itself. In contrast to the previous literature, which extensively documents the effects of oil supply shocks, we focus on an oil price surge triggered by (i) an increase in oil demand induced by concerns about future supply of oil and (ii) a favourable world economic activity shock. Our results show that, on the one hand, an oil-specific demand shock generates a *temporary*

[1]From chapter 1 the paper entitled "The Inflation Rate Disconnect Puzzle: On the International Component of Trend Inflation and the Flattening of the Phillips Curve" was derived, jointly written with Prof. Guido Ascari.

[2]Chapter 2 is the result of the visiting period at the Norges Bank as Research Intern. From this chapter, we derived the paper entitled "The Macroeconomic Effects of the Gender Revolution", jointly written with D. Bergholt and F. Furlanetto.

increase of inflation and a *persistent* drop of output, mainly due to a forceful recessionary component embedded in the shock. On the other hand, in response to a world economic activity shock, output *temporarily* improves, while inflation *persistently* increases through an amplifying transmission mechanism acting via firms costs.

# Contents

# List of Figures

# List of Tables

# Chapter 1

# The Inflation Rate Disconnect Puzzle

## 1.1 Introduction

For decades Central Banks have been working strenuously to secure the health of the entire economic system, by controlling inflation dynamics. After the second oil shock, the Volcker administration in the early '80s put incredible effort to bring surging inflation expectations to a halt and rebuild the credibility of the Fed (e.g., Goodfriend and King, 2005). After the Volcker disinflation, inflation started stabilising at lower levels and became also more persistent. Recently, the 2009 great financial crisis brought about unprecedented peaks of unemployment rates without triggering deep deflationary pressures, as markets and central bankers expected. Then, following the massive injection of liquidity into the system and the recovery of the economy - the unemployment rate at the end of 2019 was 3.5%, its lowest level in almost half a century - inflation was expected to catch up, but it did not. This puzzling inflation dynamics, led the literature to investigate explanations for both the "missing disinflation" (e.g., Coibion and Gorodnichenko, 2015a) during the financial crisis and the "missing inflation"(e.g., Heise et al., 2020) afterwards. Moreover, the fact that inflation has been stable and persistently below the 2% target in the last two decades calls for the identification of possible deflationary forces.

We contribute to the literature on inflation dynamics in two dimensions. First, we investigate the role of an international supply component of the dynamics of inflation. During the last 30 years the world lived through the fall of the Berlin wall, the rise of emerging markets economies, the entrance of China in the WTO. These events transformed the world we live in and determined a tremendous rise in globalization and international trade, thanks to reduced transportation costs, trade liberalization, the development in ICT and the integration of emerging economies in international production networks through Foreign Direct Investments (FDIs) and Global Value Chains (GVCs). An ever-increasingly integrated world economy has important implications for domestic price dynamics. We can distinguish two main channels: one from the demand side, related to final goods, and the other from the supply side, related to internationally traded intermediate inputs. On the demand side, a more open economy makes domestic markets more contestable. The harsher competition from abroad reduces the ability of domestic firms to adjust prices and keep profit margins constant (e.g. Heise et al., 2020; Jaravel and Sager, 2019; Bai and Stumpner, 2019). However, in this paper we will not focus on the effects of international competition pressures on domestic prices of *final* goods, but on a second effect, namely the international fragmentation of production and the rise in GVCs. As the world economy becomes more integrated, domestic firms find it convenient to delocalise and off-shore part of their production, leading to the fragmentation of national value chains and a globally interconnected production network, summarized into the concept of GVCs. As a consequence, firms extensively use imported *intermediate inputs* of production. More than half the world's trade in 2019 was accounted for by trade in intermediate products.[1] The main implication of this phenomenon is that firms' costs could become disconnected from domestic conditions, because they depend on imported intermediate goods produced abroad.[2] Our empirical model allows for this implication and asks the data how much this international component of costs is important in shaping the dynamics of inflation in U.S. data.

Second, we propose a unified approach to model inflation dynamics, based on a VAR in deviations from its stochastic trends. In studying inflation empirically it is important to decompose and to model jointly the trend and the cyclical components. For example, the Phillips curve is a relation between the cyclical components of unemployment (or output) and inflation, but to determine the cyclical component one needs to take a stand on the trend. Moreover, we want to allow the international supply factors described above to potentially affect both the trend and the cyclical components of inflation. Regarding the former, we decompose trend inflation into two distinct components: a domestic one due to inflation targeting anchoring long-term inflation expectations and a foreign one due to imported cost-push (supply) factor. Regarding the latter, the cyclical block of the empirical model relates the cyclical co-movements of inflation to domestic real variables, wage inflation and international import prices.

The trend-cycle analysis is motivated by our main research question, that is, what is the role of international supply factors in determining the dynamics of: (i) trend inflation, hence possibly explaining the recent deflationary pressure on trend inflation; (ii) cyclical inflation, hence possibly explaining its recent puzzling cyclical behavior. It is worth noting that both issues are of first-order importance from a monetary policy perspective.

[1] The rest is divided between primary, consumer and investment goods, see UNCTAD, 2020, Key statistics and Trends in International Trade, https://unctad.org/system/files/official-document/ditctab2020d4_en.pdf.

[2] The IMF (2005) shows that the declining trend in U.S. core inflation mainly reflected the effects of imported semiconductors between 1997 and 2005. When removing semiconductors, the trend is essentially flat.

An inflation level persistently below target threatens the long-run mandate of the central bank. An inflation dynamic that is insensible to the domestic conditions, due to a large portion of firms' costs being imported, becomes harder to control by the central bank.[3]

The main result of our analysis is that the international cost-push factor affects both the trend and cyclical behavior of inflation. First, the imported international cost-push factor exerts a persistent deflationary pressure on trend inflation over the whole sample period. In particular, despite the switch in the monetary regime successfully anchored long-run expectations around the explicit target of 2%, the international cost-push factor prevented trend inflation to stay on target throughout '90s and over the last decade. Importantly, we did not find any evidence of a change over time of the relative importance of the international component of trend inflation. Second, in the empirical analysis of the cyclical block of the model, we investigate whether the inflation gap has become increasingly exogenous to the domestic block of the model, by means of Uhlig (2003) identification scheme of the impulse response analysis. The results show a strong flattening of the Phillips curve over time. The relationship between domestic labor market and inflation during the 1960Q1-1984Q4 period is solid, while it substantially weakens in the period 1985Q1-2019Q4. From the results on the second subsample, two facets of the business cycle emerge: (i) a business cycle shock responsible for the main bulk of fluctuations in real variables, as in Angeletos et al. (2020), that push wages upwards but do not affect inflation; (ii) a propagation mechanism that generates a strong co-movement between domestic inflation and international intermediate input prices, but orthogonal to the domestic real variables. Overall, we find evidence of an "inflation rate disconnect puzzle", whereby inflation becomes disconnected from the domestic labor market due to a significant decline in the wage pass-through from domestic slack to U.S. inflation. Business cycle movements in U.S. inflation are, instead, increasingly characterized by fluctuations originating abroad, through international inputs linkages, leaving little room for domestic slack to move inflation. Third, our results, therefore, find support for the globalization of inflation hypothesis (GIH) - meaning that international factors have progressively over time replaced domestic factors as globalization increases - in the cyclical block but not in the trend block of the empirical model, possibly providing a rationale for some conflicting results in the literature.

The paper is organized as follows. Section 1.2 reviews the related literature. Section 1.3 presents the data and the methodology to extract the low and the high frequency components of the endogenous variables. Section 1.4 describes the analysis of trend inflation. 1.5 contains the main results about the dynamics of inflation at business cycle frequencies and about the flattening of the Phillips Curve. Section 1.6 concludes.

[3] In a speech the former FED governor Yellen (2006) stressed the importance of taking into account the implications of globalization on both the *level* and the *dynamics* of inflation. See also Carney (2019) and Greenspan (2005) for a discussion on the implications of globalization for domestic monetary policy.

## 1.2 Related Literature

The literature studying the inflation process is extensive. Thus, here we do not aim to survey it comprehensively, but to place our paper into the more relevant recent literature to highlight our contribution.

Our work is mostly related to the literature investigating the global determinants of inflation. While we focus on the U.S., this literature identifies a global common component in inflation dynamics across countries and studies its effects on national inflation rates. Borio and Filardo (2007) is among the earliest paper proposing a more globally-oriented rather than country-oriented view of inflation by providing supporting evidence in favor of the GIH. They estimate the empirical Phillips curve for a panel of OECD countries and find that including proxies for the global factors - i.e., oil prices, world output gap, China's output gap, etc. - substantially improves the explanatory power. In addition, they find evidence for a considerable increase in the pass-through of international factors into domestic inflation gap since the '90s, with a limited role for domestic measures of slack. While Borio and Filardo (2007) focus on the effect of the GIH on the slope of the Phillips curve, Ciccarelli and Mojon (2010) introduce the notion of global inflation, identifying through a factor-augmented model a common global factor that accounts for a strong co-movement among 22 OECD inflation rates both at low and business cycle frequencies. They show that this common global factor has driven the reduction in the level and persistence of national inflation rates over time (see also Mumtaz and Surico, 2012, for similar results in a time-varying VAR with stochastic volatility). Interestingly, Ciccarelli and Mojon (2010) report that inflation has been dominated by the common component since the '60s with no evidence of change over time, which seems to contradict the GIH. However, the results in Borio and Filardo (2007) and Ciccarelli and Mojon (2010) (and likewise other surveyed below) do not necessarily stand in contradiction with each other. The former is about the Phillips Curve that links the inflation gap to measure of domestic (and international) slack, and hence, a cyclical phenomenon, while the latter is about the level and persistence of inflation, hence, more related to the slow-moving component of inflation. In other words, these results tackle two important, but different, questions. The former question is about whether the sensitivity of the cyclical behavior of inflation to domestic slack has changed over time due to the increase in globalization, the latter one is about whether there is a common global component driving the level of inflation. Our trend-cycle decomposition nicely lends itself to reconcile these results since we find evidence of a structural change in the relationship between U.S. inflation gap and measures of domestic slack, while we find no evidence of a change in the relationship between the domestic and the international component of trend inflation in U.S. data.

Our approach is close to Eo et al. (2020), Kamber and Wong (2020) and Hasenzagl et al. (2020), who use trend-cycle decomposition to study both the long-run permanent component (trend) and the business cycle component (gap) of inflation dynamics. Kamber and Wong (2020) use a FAVAR model in which they distinguish a foreign block and a domestic block, assume block exogeneity identification restrictions and then perform a Beveridge-Nelson decomposition to distinguish trend and cycle. They find that global factors can have a sizeable influence on the inflation gap, while they play only a marginal role in driving trend inflation. However, in their setup, foreign shocks are mainly due to commodity prices, and, contrary to us they do not use intermediate good prices. Eo et al. (2020) use good and service sectors inflation rates to retrieve the aggregate headline inflation trend. They use service sector as proxy for the non-tradable component of inflation and show that it is dominant in explaining aggregate trend inflation. They conclude, therefore, that international factors have a limited effect. Stock and Watson (2019) show different results decomposing U.S. inflation into its sectoral components and applying a band-pass filter to study the cyclical correlation between sectoral inflation and real activity. They find evidence of a strong and stable correlation over time for those sectors that are "domestically" determined. While the relatively more "internationally" determined sectors are loosely linked to domestic real activity. In the section analyzing the cyclical component, Hasenzagl et al. (2020) explore the possibility for inflation gap being synchronized with the proxies of global demand. We instead focus on the supply side including a measure for imported intermediate inputs inflation in the international block of the model. Moreover, while Hasenzagl et al. (2020) uses a semi-structural approach, we employ a different methodology based on a BVAR with common trends as in Del Negro et al. (2017b). Furthermore, in contrast with Hasenzagl et al. (2020), we depart from the standard assumption of trend inflation being the common trend between actual inflation and long-run expectations (e.g., Mertens, 2016) and allow for a low frequency supply factor to contribute in shaping trend inflation dynamics. This assumption is motivated by the strand of literature on informational frictions (e.g., Coibion and Gorodnichenko, 2012, 2015b; Coibion et al., 2018; Mertens and Nason, 2020). The main takeaway from this literature is that surveys are subject to non-negligible forecast errors, implying that agents' inflation expectations sluggishly incorporate the new incoming information.

Recently, Carriero et al. (2019) extend the analysis in Ciccarelli and Mojon (2010) and Mumtaz and Surico (2012) to a FAVAR that allows commonality in both levels and volatilities, showing that a substantial share of inflation volatility across countries is attributed to the common global factor that drives also trend and persistence. Moreover, they document this common factor to be highly correlated with China's PPI, thus supporting the argument in favor of a China supply shock. The China shock relates to the rapidly increasing participation of China in international trade starting from the '90s, eventually culminated in the entrance of China in WTO in 2001. On the one hand, import competition from low-wage countries could exert downward pressure on U.S. prices. Gamber and Hung (2001) report that some U.S. sectoral prices are sensitive to prices of imports in the same sector. Auer and Fischer (2010) document similar effects both on sectoral prices and on equilibrium inflation. Recently, Heise et al. (2020) tested the role of international pressures as potential candidate of the missing inflation puzzle over the last two decades. According to the authors, the slow inflation pick-up is attributable to smaller wage pass-through to inflation, whose decline has been set in motion by increasing import competition. Similarly to Heise et al. (2020), Jaravel and Sager (2019) show that the fall of U.S. consumer prices is driven by the decline of domestic firms' mark-ups set in motion by the entrance of Chinese competitors in the domestic market. On the other hand, imports from low-wage countries could help maintain low pressure on firms' costs if a large share of imports are intermediate goods. Our analysis aims at capturing this mechanism, as a possible explanation of the missing deflation/inflation puzzle observed in the last decade.[4] Using multi-country industry-level data, Auer et al. (2019) document that international input-output linkages account for half of the global component of producer price inflation (PPI), by creating underlying cost shocks that are propagated internationally through the global input-output network. Auer et al. (2017) extend the analysis in Borio and Filardo (2007) showing that the relative sensitivity of domestic inflation to domestic and to global slack in a Phillips Curve estimation depends on new GVCs proxies. They interpret this as supporting evidence in favor of cross-border trade in intermediate inputs as transmission channel of international slack feeding into domestic inflation.

Forbes (2019) uses three different empirical frameworks - univariate trend-cycle decomposition à la Stock and Watson (2007), Phillips curve estimation and principal components - to investigate the role of globalization for the dynamics in U.S. inflation. The main finding is that global factors - more specifically, commodity and oil prices, exchange rate, world slack, and GVCs - are important to explain the cyclical component of CPI inflation, while less so to explain both the trend component of CPI and the dynamics of core and wage inflation. Moreover, Forbes (2019) document that the role of these global factors in affecting CPI inflation has increased over the last decade providing a possible explanation for the flattening of the Philips Curve.

Finally, the debate on the flattening of the Phillips curve and missing deflation/inflation puzzle goes beyond the role of international factors. Coibion and Gorodnichenko (2015a) show the important role played by inflation expectations - see Coibion et al. (2018) for a survey. McLeay and Tenreyro (2020) and Bergholt et al. (2020) argue that the Phillips curve is in good health and inflation is under full control of monetary tools. The reason why it has become hard to estimate its slope empirically should be attributed to an identification problem, induced by optimal monetary policy. Consistent with these results, Hazell et al. (2020) attribute the flatter Phillips curve to the monetary policy regime switch started with Volcker administration. Importantly, our finding about the

[4] See Bobeica and Jarociński (2019) for a discussion on the missing (dis)inflation phenomenon in the Euro Area.

flattening of the Phillips Curve and the disconnection of inflation dynamics from the domestic business cycle do not suffer of such identification problem because we do not explicitly estimate a Phillips Curve and because our correlations are conditional on a specific business cycle shock.[5]

A more agnostic approach has been, instead, applied by Del Negro et al. (2020). They conduct a horse race between three candidate hypotheses of the poor responsiveness of inflation, namely the mismeasurement of labor market variables, the flattening of the aggregate supply curve and the flattening of the aggregate demand curve, associated with the systematic aggressive response of monetary policy to demand shocks. The three hypotheses are scrutinized through the lens of a VAR and a medium-scale DSGE. Their findings mainly attribute the silent behavior of inflation to the flattening of the aggregate supply curve and, to a lesser extent, to more aggressive monetary policy.[6]

While our results confirm some findings in the literature above, our analysis adds the importance of considering intermediate inputs - imported industrial supplies and materials, and oil - to explain both the deflationary forces acting on the trend level of inflation and the changing sensitivity of PCE U.S. inflation to measures of domestic slack. Our trend-cycle decomposition is key to these results and provide complementary evidence to the literature and a rationale for reconciling some existing conflicting evidence.

## 1.3 Methodology and Data

Assume the vector of observables $\chi_t$ to be the sum of two unobserved states:

$$\underset{(n\times 1)}{\chi_t} = \underset{(n\times q)}{\Lambda}\underset{(q\times 1)}{\bar{\chi}_t} + \underset{(n\times 1)}{\tilde{\chi}_t} \tag{1.1}$$

$\bar{\chi}_t$ describes the slow-moving (trend) dynamics due to permanent shocks. $\tilde{\chi}_t$, instead, defines the transitory (cyclical) fluctuations in the data. $\Lambda$ is sparse and accommodates the presence of $q \leq n$ common trends. We leave the discussion about the assumptions on $\Lambda$ to section 1.4, as it plays a crucial role for our work. Keep in mind, for the moment, that its elements can be either calibrated, estimated or both. Similarly to Del Negro et al. (2017b) and Hasenzagl et al. (2020), trends and cycles are assumed to be stochastic and

[5] See Barnichon and Mesters (2020) for an enlightening discussion of the issue and a proposed solution based on the Phillips Curve estimation conditional on monetary policy shocks.

[6] The New Keynesian DSGE literature provides some explanations for the missing deflation/inflation puzzle. Christiano et al. (2015) and Del Negro et al. (2015) show that the inclusion of financial frictions in a standard medium-size DSGE model could help in replicating the small drop in inflation at the darkest hour of the 2009 recession. Gilchrist et al. (2017) show that liquidity constraint firms have an incentive to raise prices in response to adverse financial shocks to avoid the deterioration of their liquidity position, thus mitigating the response of inflation to output fluctuations. Lindé and Trabandt (2019) attributes the missing deflation puzzle to the arising non-linearities in price and wage settings when the system is hit by large shocks.

evolve according to a random walk and an reduced-form invertible VAR, respectively:

$$\bar{\chi}_t = c_t + \bar{\chi}_{t-1} + e_t \qquad e_t \sim N(0_q, \Sigma_e) \tag{1.2}$$

$$\Phi(L)\tilde{\chi}_t = \varepsilon_t \qquad \varepsilon_t \sim N(0_n, \Sigma_\varepsilon), \tag{1.3}$$

where $c_t$ is meant to capture a common growth rate in the trends of real output and investment and is assumed to follow a random walk without drift. $\Phi(L) = I - \Phi_1 L - ... - \Phi_p L^p$ and $\Phi_k$, for $k = 1, ..., p$, are the lag coefficients matrices of dimension $(n \times n)$. Our model is equivalent to a VAR in deviations from its steady states à la Villani (2009) with the notable difference that in this specification steady states are allowed to be time-varying. The vector embedding permanent and transitory innovations is assumed to be i.i.d. and distributed according to a multivariate Normal distribution. This implies, in turn, that permanent and transitory shocks are mutually uncorrelated. The initial conditions of the unobserved states are assumed to be distributed according to:

$$\bar{\chi}_0 \sim \mathcal{N}\Big(\underline{\chi_0}, V_0\Big) \tag{1.4}$$

$$\tilde{\chi}_0 \sim \mathcal{N}\Big(0, V(\Phi, \Sigma_\varepsilon)\Big), \tag{1.5}$$

where $\underline{\chi_0}$ is the pre-sample mean and $V_0$ is the $(q \text{ x } q)$ identity matrix. The initial condition of the cycles is a vector of zeros, as we assume that cycles symmetrically fluctuate around a zero mean. $V(\Phi, \Sigma_\varepsilon)$ is the unconditional variance of the initial conditions for cyclical components and it is always well defined as we impose stationarity on $\tilde{\chi}_t$.

Finally, the priors of the model's coefficients are distributed according to:

$$\Sigma_e \sim \mathcal{IW}(\kappa_e, (\kappa_e + n + 1)\underline{\Sigma_e}) \tag{1.6}$$

$$\Sigma_\varepsilon \sim \mathcal{IW}(\kappa_\varepsilon, (\kappa_\varepsilon + n + 1)\underline{\Sigma_\varepsilon}) \tag{1.7}$$

$$vec(\Phi)|\Sigma_\varepsilon \sim \mathcal{N}(vec(\underline{\Phi}), \Sigma_\varepsilon \otimes \underline{\Omega})\mathcal{I}(vec(\Phi)) \tag{1.8}$$

$\mathcal{IW}$ is the Inverse-Wishart distribution with $\kappa$ degrees of freedom and mode equal to $\underline{\Sigma}$. We place a rather tight prior on the covariance matrix of permanent innovations $\Sigma_e$ being diagonal, by setting the hyperparameter $\kappa_e = 100$. Notice, however, that this does not prevent the data to speak in favor of correlations among permanent innovations, if this is the case. Following Del Negro et al. (2017b), the prior on the diagonal elements of $\Sigma_e$ is conservative in limiting the amount of variance attributable to the trends. We do so by normalizing the variances by 200, which implies that the standard deviation of the expected change in the trend over five decades is only 1 percentage point.

Turning to the cyclical block, the priors for the lag coefficients are standard Minnesota priors with overall tightness hyperparameter equal to 0.2 with the exception of the own-lag hyperparameter which is set equal to zero instead of 1, as we are characterizing the stationary behavior of the data. $\mathcal{I}(vec(\Phi))$ is an indicator function that is equal to value of one, if all the roots of $\Phi(L)$ are outside the unit circle, equal to zero, if the VAR is explosive. Since data are quarterly, the number of lags $p$ in the stationary VAR is set equal to 4. Finally, the model is linear in its equations and we employ the Kalman Filter to extract the unobserved components in the model and implement simulation smoothing techniques to generate the posterior distribution (see Carter and Kohn, 1994). The model samples 10000 draws from the Gibbs algorithm and retains the last 5000 draws.

**Data.** The model is estimated using the following variables in the vector $\chi_t$ : unemployment rate, $u_t$; real output per capita, $y_t$, ; real investment per capita, $i_t$; PCE headline inflation $\pi_t$;[7] 10-year ahead PCE headline inflation expectations, $\pi_t^e$; wage inflation, $\pi_t^w$; imported intermediate inputs price inflation, $\pi_t^m$; oil price inflation, $\pi_t^o$.[8] Coherently with (1.1), a bar over a variable indicates its trend component, while a tilde its cyclical component. The sample spans from 1960Q1 to 2019Q4.

## 1.4 The Anatomy of Trend Inflation

We distinguish two determinants of the persistent component of inflation, i.e., trend inflation.[9] First, we assume trend inflation to be primarily determined by monetary policy and its ability to anchor inflation expectations. We capture this by modelling a common trend between inflation and 10-years expected inflation, as often assumed in the literature (see, e.g., Mertens and Nason, 2020; Nason and Smith, 2020). Second, we assume that trend inflation can also be potentially influenced by slow-moving "cost-push" factors, imported from abroad, that are not under full control of monetary policy. This international determinant of the "cost-push" factors is an important focus of this work. We want to allow for the possibility that low frequency movements in international factor prices affect the persistent component of inflation. Imported intermediate inputs are the most direct proxy of the effect of intensive international input-output linkages, as consequence of globally fragmented production chains. Intensive trade in intermediate inputs should shrink the portion of firms' cost originated domestically and possibly exerts downward pressures to trend inflation, generating persistent deviations from the monetary trend that could not be easily offset by monetary policy. Thus, we propose the following

[7] The reason for preferring headline inflation to core inflation owe to the fact that the medium-term objective of the FED is headline inflation. Thus, we are interested in estimating a measure of trend inflation that is consistent with the definition of inflation target by the FED.

[8] We use the series of industrial supplies and materials as proxy for intermediate inputs. The Appendix 1.7.1 summarizes the code for all variables.

[9] We use both terms - persistent component of inflation and trend inflation- interchangeably, as Koester et al. (2021).

linear decomposition:

$$\bar{\pi}_t = \underbrace{\bar{\pi}^*_t}_{\text{monetary}} + \underbrace{\lambda_\pi \bar{g}_t}_{\text{international}} \tag{1.9}$$

$$\bar{\pi}^m_t = \bar{\pi}^{m,id}_t + \lambda_{\pi^m} \bar{g}_t \tag{1.10}$$

$$\bar{\pi}^o_t = \bar{\pi}^{o,id}_t + \lambda_{\pi^o} \bar{g}_t \tag{1.11}$$

where $\bar{\pi}^*_t$ is the time-varying inflation target, thus we label it as the monetary component of trend inflation. As said above, this also coincides with the trend component in long-term inflation expectations, as usually assumed in this literature. $\bar{\pi}^{m,id}_t$ and $\bar{\pi}^{o,id}_t$ are the idiosyncratic slow-moving components of import prices and oil prices, respectively. $\bar{g}_t$ is the international trend that links the persistent dynamics in international factor prices to domestic inflation. Persistent wedges between $\bar{\pi}_t$ and its monetary trend $\pi^*_t$ potentially arise from $\bar{g}_t$. The prior for the loading parameters in the $\Lambda$ matrix in (1.1), i.e., $\lambda_\pi, \lambda_{\pi^m}$ and $\lambda_{\pi^o}$ is assumed to be Gaussian with mean 1 and standard deviation 0.5 - Appendix 1.7.2 shows the $\Lambda$ matrix.

A large literature, convincingly argued that inflation expectations deviate from the full-information benchmark and non-negligible forecast errors could be due to informational frictions (see Coibion et al., 2018, for a survey of this literature). Following the hypothesis that GVCs have shaped inflation dynamics (e.g., Borio and Filardo, 2007; Auer et al., 2017, 2019; Forbes, 2019), we allow for the possibility that globalization has produced not only short/medium run business cycle effects, but also persistent shifts in the long-run component of inflation. $\lambda_\pi$ determines the difference between the persistent components of inflation expectations and of inflation. We allow for this de-anchoring between long-term inflation expectations and the statistical measure of trend inflation to depend on the low frequency movements of international factor prices. With this specification, we allow the data to tell us whether international supply side cost-push factors are important in shaping trend inflation dynamics. Nevertheless, it is worth noting that, if $\lambda_\pi$ is equal to 0, we are back to the conventional case in which inflation is solely a monetary phenomenon in the long-run and, therefore, under full control of monetary policy.

Figure 1.1 plots the estimated trend according to (1.9). Our estimate of trend inflation captures the inertial and lagging trend of long-run expectations relative to trend inflation, both in the upturn in the '70s and in the slowdown during the Volcker disinflation. These results are in line with the empirical investigation in Mertens (2016) and consistent with the notion that inflation expectations became unanchored from trend inflation during the '70s and early '80s.[10] We do not provide here a structural explanation of why the trend in long-term inflation expectations adjusted only sluggishly and lagged behind

[10] This would not be possible in a specification defining trend inflation as the common long-run component between inflation rate and 10-years expectations. In fact, in such specification the estimated trend would essentially be adherent to observable data on 10-years expectations.

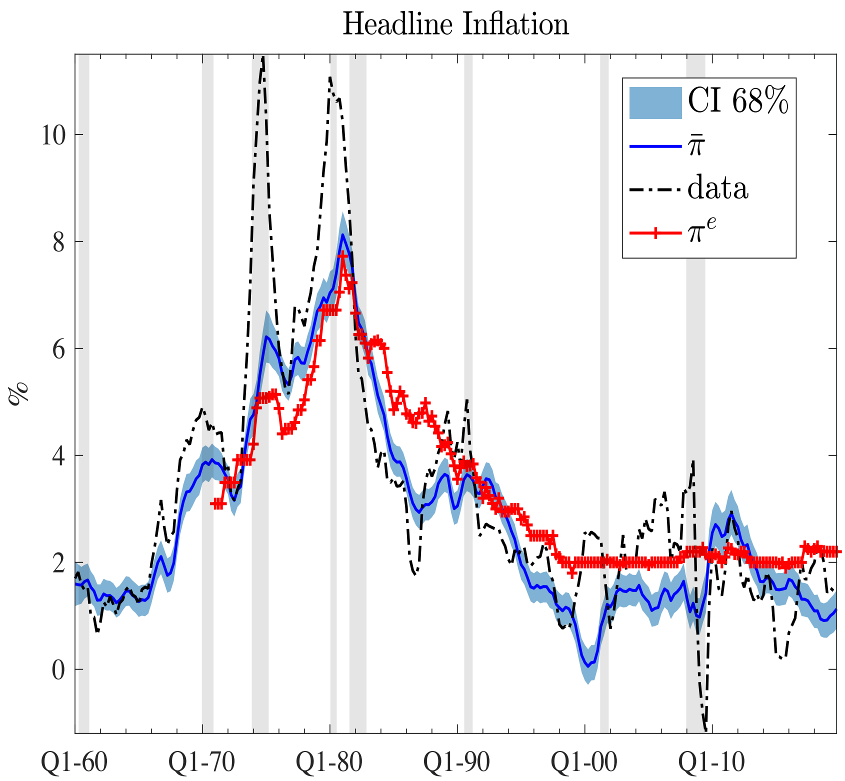


Figure 1.1: Trend inflation.

in incorporating into expectations the importance of these factors in influencing actual trend inflation dynamics. The empirical results of Coibion and Gorodnichenko (2012, 2015b), Mertens and Nason (2020), and Nason and Smith (2020), which are consistent with models of informational frictions in survey responses, provide a plausible possible explanation for the difference between the trend in inflation expectations and the one in actual inflation.

Regarding the cost-push factors, Figure 1.1 shows that their slow moving dynamics drive both the upturn of trend inflation in the pre-Volcker period and the sharp dip afterwards.

Figure 1.2 focuses on trend inflation dynamics over the last 30 years. Two main facts emerge. First, trend in inflation expectations is consistently above the trend in actual inflation and from the late '90s it is extremely well-anchored around the inflation target of 2%, showing the power of the inflation target regime in shaping long-term inflation expectations.[11] Second, the estimate of trend inflation falls below 2% from the mid '90s, showing a significant drop until early 2000s and then a steady recovery towards 2% until 2011, followed by a persistent drop thereafter. According to our estimates, the trend in the global cost-push component explains this discrepancy between the persistent component in inflation and the one in inflation expectations. Hence, from the '90s the

[11] Again these patterns are consistent with the results in Mertens (2016).

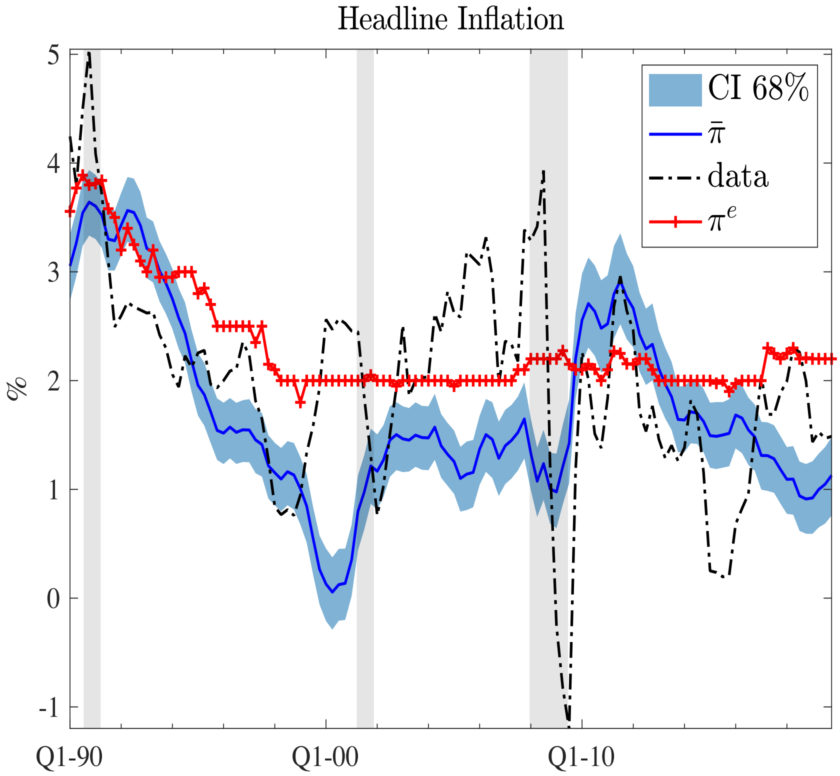


Figure 1.2: Trend inflation 1991Q1-2019Q4.

movements in trend inflation are dominated by movements in the cost-push component, which exerted a deflationary pressure on trend inflation, unanchoring the dynamics of trend inflation from the 2% target and from the trend in inflation expectations.

Our analysis provides a measurement of the relative importance of the global trend feeding into trend inflation. Figure 1.3 provides strong evidence in favor of a state of the world in which the slow-moving dynamics in international factor prices, i.e., $\bar{g}_t$, non-negligibly affect the low-frequency movements in inflation. The data could in principle reject that assumption by estimating a negligible role for this component of trend inflation, pushing $\lambda_\pi$ towards zero. The posterior distribution (blue) of $\lambda_\pi$, instead, moves decisively rightward and tightens, such that it does not include values smaller than 1. The posterior distributions for the other two loadings, $\lambda_\pi^m$ and $\lambda_\pi^o$, tightens too, but they remain centered in 1. Overall, the posterior distributions show a strong identification and refuse the possibility that $\lambda_\pi = 0$. Further corroborating evidence is provided by the Bayes factor in Table 1.1 that shows that the specification $\mathcal{M}_1$, allowing for the possibility of $\lambda_\pi \neq 0$ is strongly preferred to a benchmark "monetarist" model $\mathcal{M}_0$, which imposes $\lambda_\pi = 0$.[12]

Furthermore, we find no evidence of significant time variation in these coefficients when we estimate the model for the pre- and post-Volcker sample. The strong deflationary pressure in Figure 1.1 derives entirely from the international component, and starts

[12]See Jeffreys (1998) for the thresholds values.

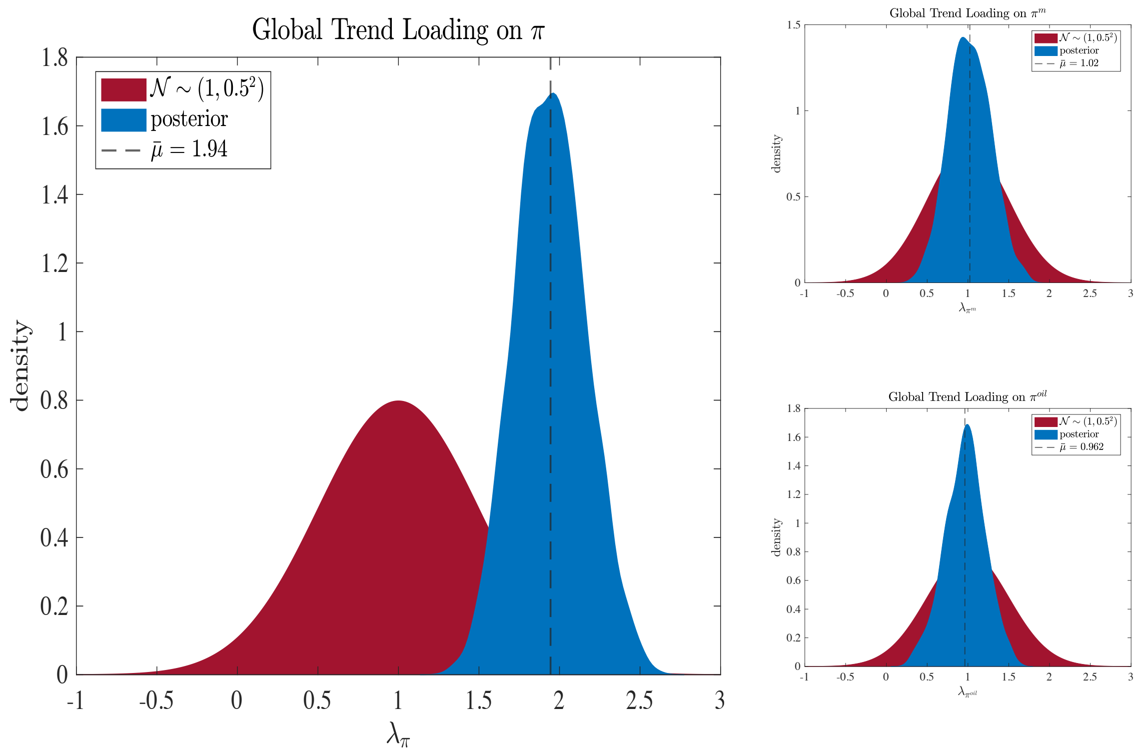


Figure 1.3: Prior and posterior distributions of the loadings

Table 1.1: Bayes Factor

| | Bayes Factor |
|---|---|
| $log(\mathcal{M}_0) - log(\mathcal{M}_1)$ | -6.70 |

already from mid-‘80s and particularly so from the mid-‘90s onward. This is roughly the period in which China embraced globalization becoming an increasingly important hub in international trade - the China shock. Branstetter and Lardy (2008) document the rapid growth of Chinese exports in the decade prior to WTO accession and the dramatic acceleration of liberalization of trade and foreign direct investment (FDI) in the mid-‘90s. Coase and Wang (2012) argue that important reforms in the beginning of the ‘90s enabled the emergence of a common national market, imposing the discipline of competition to economic agents and state-owned enterprises, and thus paving the way for the transformation of the Chinese economy.[13]

Overall, the results show that the decline of U.S. trend inflation from the ‘80s has been driven both by substantial improvements in monetary policy, leading to more anchored inflation expectations, and by a deflationary pressure of the international price of intermediate goods.

## 1.5 Inflation Dynamics and the Business Cycle

This section focuses on the stationary block of the model and analyses the business cycle behavior of inflation, filtered by the potential noise arising from lower frequencies. In the remainder of this section we present the results on: (i) the cyclical measures of real economic activity; (ii) the disconnection of inflation dynamics from real activity; (iii) the flattening of the slope of the Phillips curve; (iv) the (non-flattening) slope of the wage Phillips curve.[14]

### 1.5.1 A Common Business Cycle Index

A key feature of our trend-cycle decomposition is that we do not impose any structural assumptions on the relations between cyclical components. For instance, we do not relate output gap to unemployment gap by means of a semi-structural Okun’s law as often done in the literature. We leave the cycles unconstrained and let the data speak out for themselves. Notwithstanding, Figure 1.4 shows that the estimated cyclical components of real variables - unemployment and output - clearly share a common pattern, suggesting that business cycle fluctuations are originated from a common propagation mechanism. The model retrieves two business cycle measures that are directly comparable and highly correlated with the CBO estimates of unemployment gap and output gap - yellow dashed lines in Figure 1.4.

[13] After the famous “southern tour” of Dei Xiaoping in 1992, China implemented price reform in 1992, tax reform in 1994, and began to privatize state enterprises in the mid-1990s. See also, e.g., Storesletten and Zilibotti (2014) and Autor et al. (2016).

[14] Needless to say that the trend analysis of the previous Section and the cyclical analysis in this Section are tightly connected via our trend-cycle econometric specification: they depend one from the other.

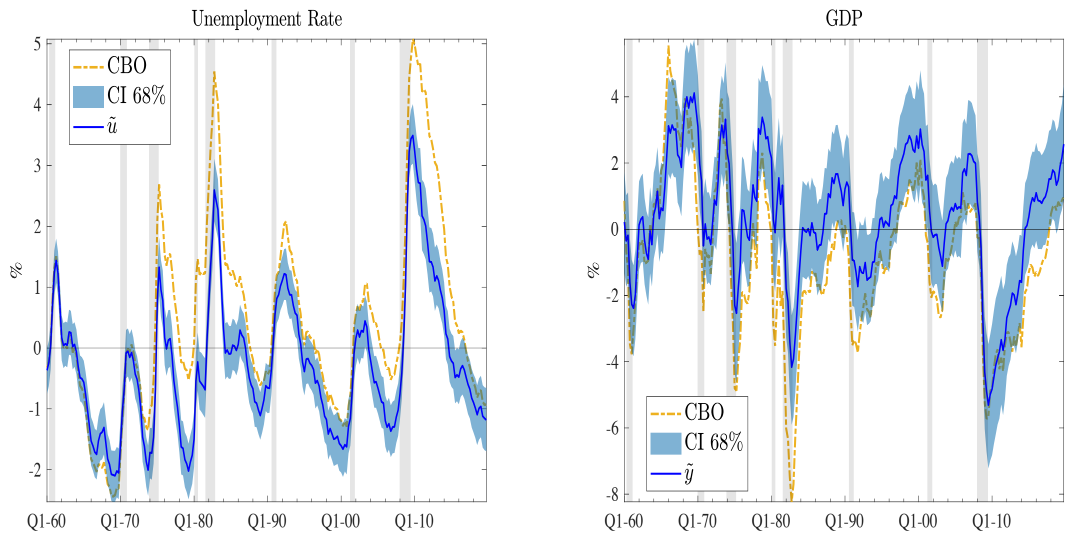


Figure 1.4: Business Cycle indexes

### 1.5.2 The Inflation Rate Disconnect Puzzle

To shed light on the relationship between the cyclical variables, we employ a methodology based on the analysis of impulse response functions presented in Uhlig (2003) (see also Giannone et al., 2019). This methodology extracts the convolution of reduced-form shocks which maximizes the forecast error variance (FEV) of one specific variable over some interval horizon - which is 6Q-32Q forecast horizon in our specific case.

Our main purpose, therefore, is to use this template to study the evolution of the relationship between inflation dynamics and the domestic real business cycle. More precisely, we are interested in extracting the convolution of shocks that is responsible for the cyclical fluctuations of unemployment, which we will refer to as the "unemployment shock" or "business cycle shock" interchangeably. To avoid potential misunderstandings, we are reluctant to assign any structural interpretation to this shock - though results show that the shock we retrieve resembles the *main business cycle shock* of Angeletos et al. (2020). In fact, our focus is specifically on understanding if: (i) this convolution of shocks explaining unemployment fluctuations transmits to domestic price inflation; (ii) these underlying forces moving unemployment can be considered to be the same as those generating inflation dynamics, implying in turn the existence of a common propagation mechanism that yields to a strong co-movement between unemployment and inflation; (iii) if there has been a break in transmission mechanism linking unemployment to domestic inflation dynamics, leading to an "inflation rate disconnect puzzle" over the last 30 years. With respect to (ii), analogously to unemployment, we extract an "inflation shock", defined as the major contributor of the business cycle inflation dynamics and compare the impulse responses with those of the "business cycle shock". For what concerns (iii), we proceed by re-estimating the model for the subsamples spanning 1960Q1-1984Q4 and 1985Q1-2019Q4 and then compare the results. However, it is important to stress that the estimation of the trends for the two sub-samples is very close to the one for the whole sample in the previous Section. By the same token, we would obtain similar result had we done the same analysis in this Section using the cyclical components coming from the trend estimated from the whole sample, as in the previous Section.

Finally, it is worth stressing out that our approach is consistent with theory - where short-run dynamics is usually studied with variables expressed in log-deviations from steady states - as we study the cyclical behavior of variables in deviations from their long-run equilibrium, thus, explicitly cleaning them out of their low frequency features.

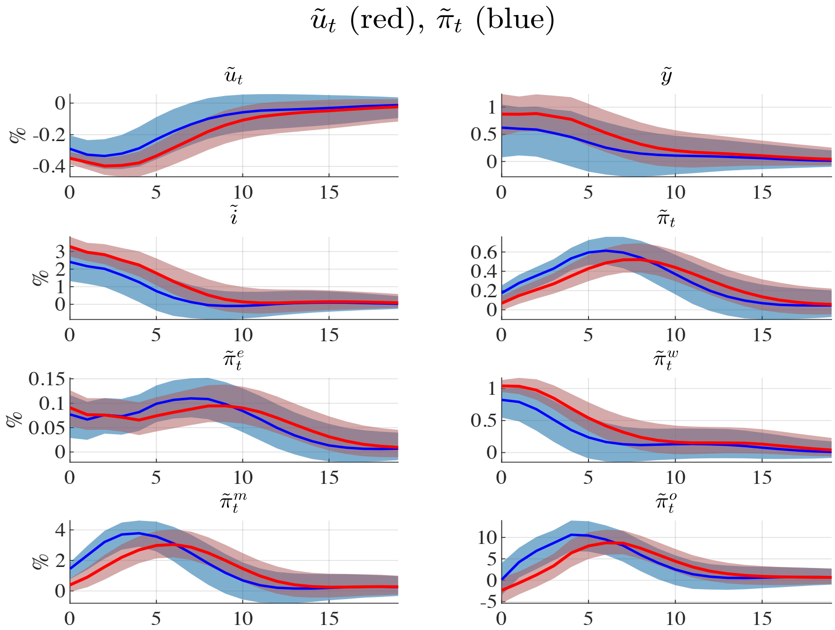


Figure 1.5: IRFs maximizing the FEVD of $\tilde{u}$ and $\tilde{\pi}$. Sample 1960Q1-1984Q4; median response and 68% uncertainty band.

Let us consider the 1960Q1-1984Q4 subsample in Figure 1.5 first. The red lines visualize the responses with respect to the "unemployment shock", and the blue lines the ones to the "inflation shock". The two set of impulse responses describe a common propagation mechanism that is not only responsible for the co-movement among real variables, but that also transmits to inflation and international prices. Similar to the results in Angeletos et al. (2020), the shock resembles a demand-type shock, as it generates an opposite co-movement between unemployment and inflation. Furthermore, it also draws a clear lead-lag relationship between domestic slack and inflation. Real variables strongly move together and tend to converge back to the steady state in approximately two years. For both the domestic real variables and wage inflation the peak of the shock realizes within one year. Instead, for inflation and inflation expectations the peak materializes after approximately two years before being entirely absorbed approximately 4 years after the shock hits the economy. In addition, the shock spillovers to imported prices of intermediate inputs and oil, whose responses are positive and significant for approximately 3 years. The "unemployment shock" generates wage inflationary pressures that encourage agents to revise the expectations upward and that pass-through inflation.

Table 1.2 displays how much the "unemployment shock" explains of the forecast error variance of the variables in our BVAR. It explains a large fraction of the volatility of real variables - 39% for the output gap and 67% for the investment gap.[15] Table 1.2 confirms the sound link between real slack and inflation, as the "business cycle shock" is responsible for approximately 84% and 60% of the business cycle fluctuations of $\tilde{\pi}^w$ and $\tilde{\pi}$, respectively. In addition, it also significantly contributes to inflation expectations as well as to the price dynamics of imported industrial supplies and materials and of oil. The third and fourth rows in Table 1.2 refer to the FEVD of "inflation shock". It explains, indeed, roughly 66% and 50% of the FEV of the unemployment gap and wage inflation gap, respectively. This first set of results shows that, in the pre-Great Moderation era, our econometric analysis points to well-functioning wage and price Phillips curves, linking the domestic labor market to inflation via wage inflationary pressures.

Figure 1.6 and 1.7 show the IRFs to the convolutions of shocks that maximize the FEV of imported industrial supplies and material and of oil prices, respectively. The dynamic responses to an imported intermediate inputs price shock are essentially indistinguishable from the ones arising from the "business cycle shock", implying that, according to our approach, the drivers of the business cycle behavior of imported intermediate input prices are indistinguishable from the ones of domestic variables. In addition, the shock maximizing the FEV of oil inflation gap displays large credibility bands, making any interpretation of this convolution of shocks difficult. The large credibility bands, indeed, span a wide range of dynamic responses that may be associated with both demand and supply shocks.

To sum up, these results lead us to conclude that in the sample 1960Q1-1984Q4 our approach retrieves a domestic "business cycle shock", associated with inflationary pressure through wages, and fluctuations in prices of imported intermediate inputs and oil.

Let us now move to the subsample spanning from 1985Q1 to 2019Q4. Figure 1.8 shows the responses to the "unemployment shock" (red) and the "inflation shock" (blue). Although wages soundly react to the "unemployment shock", this time the responses of inflation and inflation expectations are muted. By comparing the impulse responses of all variables with respect to the two shocks, we no longer find evidence of a common propagation mechanism driving both the real side and the nominal side of the economy. As a matter of fact, by just looking at Figure 1.8, we can clearly distinguish two facets of the business cycle. On the one hand, a "business cycle shock" drives the real domestic

[15] In section 1.7.7 of the Appendix, we show that, similarly, a shock that maximize the FEV of the output gap or the investment gap explains a substantial fraction of the volatility of the other two real variables, respectively.

Table 1.2: Forecast error variance decomposition.
68% uncertainty band in squared brackets.

| **1960Q1-1984Q4** | | | |
|---|---|---|---|
| $\tilde{u}_t$ shock | | | |
| $\tilde{u}_t$ | $\tilde{y}_t$ | $\tilde{i}_t$ | $\tilde{\pi}_t$ |
| 0.9071 | 0.3900 | 0.6743 | 0.5920 |
| [0.8806,0.9285] | [0.3081,0.4868] | [0.6246,0.7227] | [0.5119,0.6636] |
| $\tilde{\pi}_t^e$ | $\tilde{\pi}_t^w$ | $\tilde{\pi}_t^m$ | $\tilde{\pi}_t^o$ |
| 0.5668 | 0.8382 | 0.4731 | 0.3148 |
| [0.4835,0.6461] | [0.7996,0.8723] | [0.3987,0.5501] | [0.2536,0.3806] |
| $\tilde{\pi}_t$ shock | | | |
| $\tilde{u}_t$ | $\tilde{y}_t$ | $\tilde{i}_t$ | $\tilde{\pi}_t$ |
| 0.6556 | 0.2474 | 0.4043 | 0.7996 |
| [0.5529,0.7498] | [0.1705,0.3455] | [0.3164,0.5039] | [0.7564,0.8385] |
| $\tilde{\pi}_t^e$ | $\tilde{\pi}_t^w$ | $\tilde{\pi}_t^m$ | $\tilde{\pi}_t^o$ |
| 0.6654 | 0.5025 | 0.7093 | 0.4715 |
| [0.5972,0.7278] | [0.3883,0.6134] | [0.6524,0.7640] | [0.4175,0.5330] |
| **1985Q1-2019Q4** | | | |
| $\tilde{u}_t$ shock | | | |
| $\tilde{u}_t$ | $\tilde{y}_t$ | $\tilde{i}_t$ | $\tilde{\pi}_t$ |
| 0.9264 | 0.4428 | 0.8655 | 0.2797 |
| [0.9082,0.9422] | [0.3180,0.5570] | [0.8355,0.8914] | [0.2064,0.3603] |
| $\tilde{\pi}_t^e$ | $\tilde{\pi}_t^w$ | $\tilde{\pi}_t^m$ | $\tilde{\pi}_t^o$ |
| 0.1880 | 0.8502 | 0.0663 | 0.0593 |
| [0.1181,0.2860] | [0.8093,0.8858] | [0.0451,0.0929] | [0.0398,0.0849] |
| $\tilde{\pi}_t$ shock | | | |
| $\tilde{u}_t$ | $\tilde{y}_t$ | $\tilde{i}_t$ | $\tilde{\pi}_t$ |
| 0.1431 | 0.1079 | 0.1221 | 0.7172 |
| [0.0855,0.2451] | [0.0720,0.1638] | [0.0768,0.2057] | [0.6642,0.7701] |
| $\tilde{\pi}_t^e$ | $\tilde{\pi}_t^w$ | $\tilde{\pi}_t^m$ | $\tilde{\pi}_t^o$ |
| 0.1557 | 0.2336 | 0.59404 | 0.7485 |
| [0.1098,0.2154] | [0.16812,0.3329] | [0.5230,0.6511] | [0.6836,0.7965] |

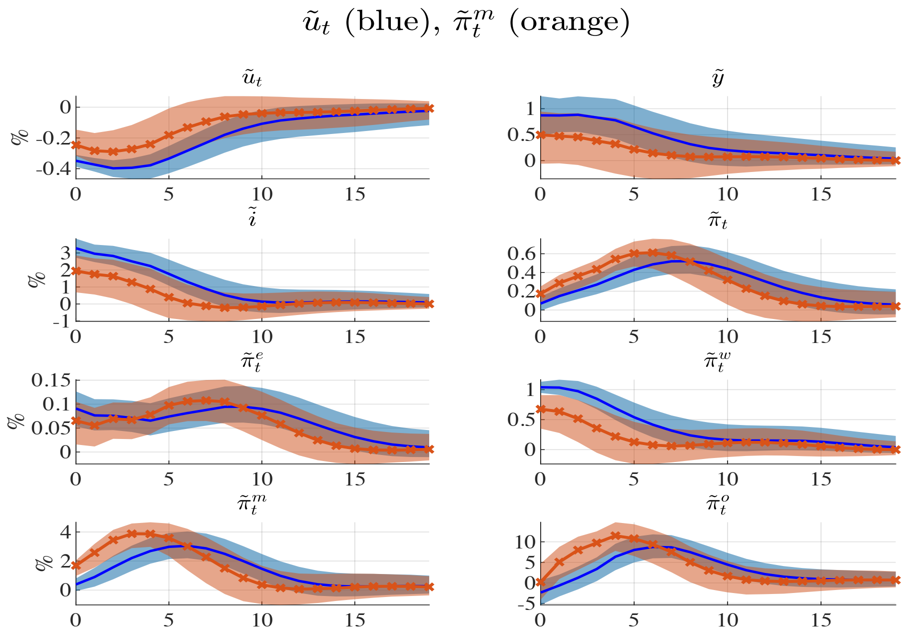


Figure 1.6: IRFs maximizing the FEVD of $\tilde{u}$ and $\tilde{\pi}^m$. Sample 1960Q1-1984Q4; median response and 68% uncertainty band.

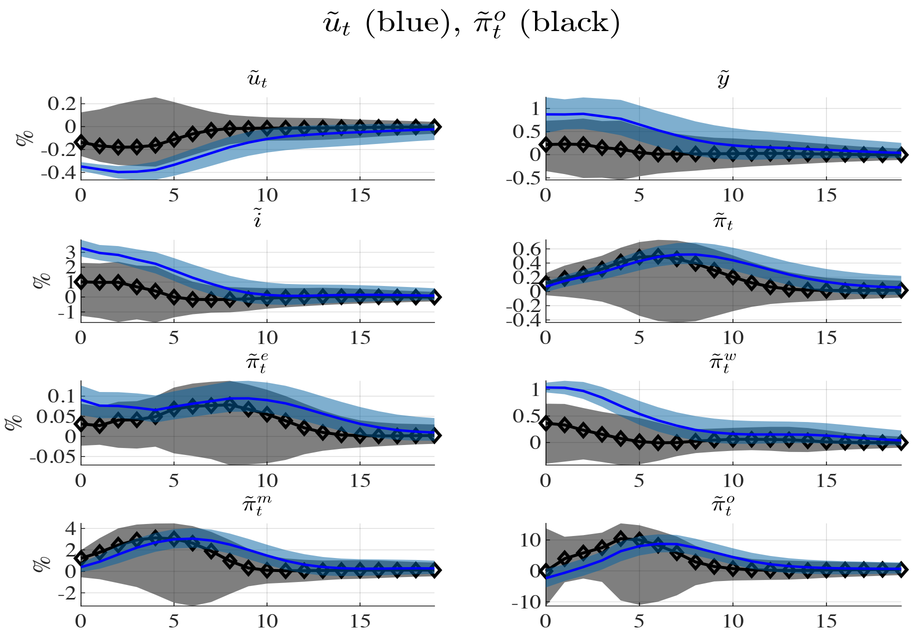


Figure 1.7: IRFs maximizing the FEVD of $\tilde{u}$ and $\tilde{\pi}^o$. Sample 1960Q1-1984Q4; median response and 68% uncertainty band.

variables and wage inflation, but no longer feeds into inflationary pressure - both actual inflation and inflation expectations are muted. In addition, it no longer transmits to international prices, as can be seen by the non-significant responses of $\tilde{\pi}^m$ and $\tilde{\pi}^o$. On the other hand, an"inflation shock" produces a strong positive co-movement with imported intermediate input prices - see the (blue) responses in the last row of Figure 1.8 -, but that affects only slightly the domestic real variables.

This *disconnection* is confirmed by Table 1.2. The "business cycle shock", indeed, characterizes approximately 93% of the total volatility of the unemployment gap, 87% of the investment gap, 44% of output gap. Furthermore, it accounts for 85% of wage inflation total volatility, therefore generating a strong wage inflationary pressures as in the pre-Great Moderation period. However, the portion of FEV of domestic inflation and inflation expectations explained by the shock has declined by more than a half. It seems that the main transmission channel, via wage pressure, no longer passes-through to inflation. Note that the muted response of inflation can not be due to an aggressive or improved monetary policy reaction, as suggested by McLeay and Tenreyro (2020) among others. As a matter of fact, our analysis features conditional correlations and therefore, the unemployment gap should have also remained closed conditional to the "unemployment shock", if an improved monetary policy was the reason behind the poor response of inflation.

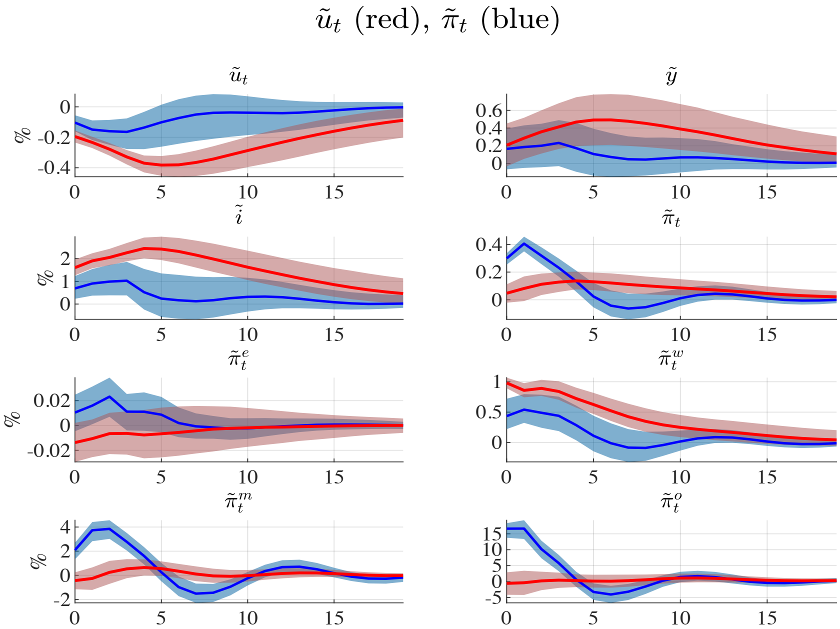


Figure 1.8: IRFs maximizing the FEVD of $\tilde{u}$ and $\tilde{\pi}$. Sample 1985Q1-2019Q4; median response and 68% uncertainty band.

Inflation, therefore, seems to have become increasingly decoupled from the domestic business cycle and the local labor market. And yet it moves: how? Figure 1.8 plots the impulse responses (blue lines) to the "inflation shock". The responses to this shock reveal a disconnection between the real side and the price side over the business cycle. First, the shock is essentially orthogonal to the real side of the economy because the impulse responses are no longer interchangeable to those of the "unemployment shock" - the credibility bands no longer overlap - and most importantly the real variables barely move and output not significantly so. Second, there is a strong positive - and precisely estimated

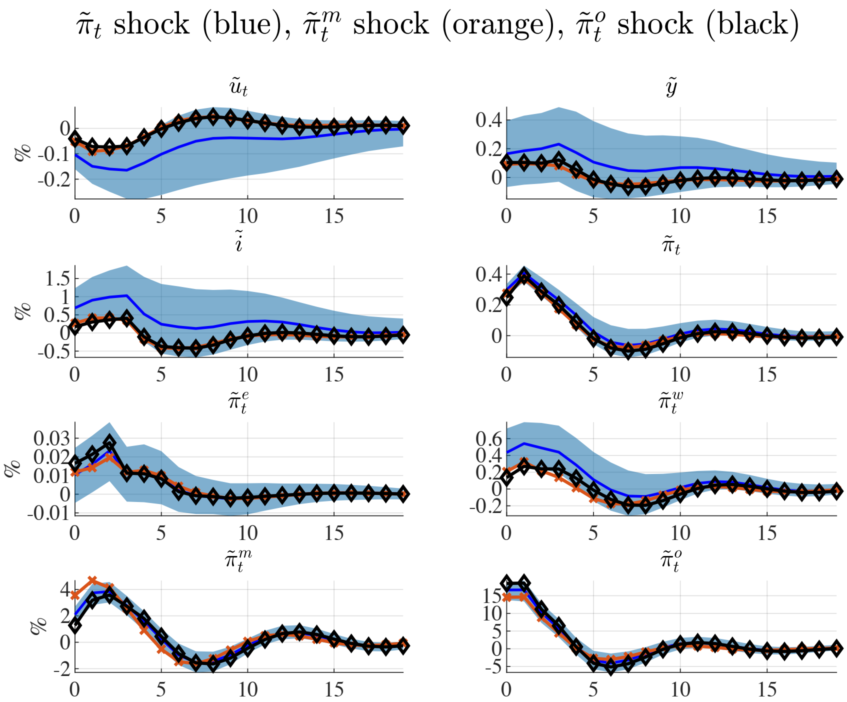


Figure 1.9: IRFs maximizing the FEVD of $\tilde{\pi}$ (blue –), $\tilde{\pi}^m$ (orange -×-) and $\tilde{\pi}^o$ (black -◇-). Sample 1985Q1-2019Q4; median response and 68% uncertainty band.

- co-movement between domestic inflation, imported intermediate inputs inflation and oil inflation. The results in Table 1.2 are even more striking: the shock generates a large part of the international components of prices inflation, i.e., 59% of the variance of imported intermediate goods inflation - $\tilde{\pi}_t^m$ - and 75% of the one of oil prices - $\tilde{\pi}_t^o$, however it generates a small fraction of the volatility of the wage inflation (23%) and an even smaller fraction of the volatility of inflation expectation (16%), unemployment (14%), output (11%) and investment (12%) gaps. These results suggests that, in this second subsample, the forces driving inflation dynamics are different from the ones driving the domestic cycle.

In order to gain additional evidence, we estimate the impulse responses generated from a shock maximizing the cyclical FEV of either imported intermediate inputs inflation or oil inflation. If the strong co-movement originates from a common *international* propagation channel, then the impulse responses should be interchangeable to the blue ones in Figure 1.8, that is, the underlying shocks should explain approximately the same amount of business cycle volatility of these three variables and should not co-move with any of the real domestic variables. The three sets of impulse responses are jointly plotted in Figure 1.9. Strikingly, they draw the same patterns and are almost perfectly correlated - not surprisingly, domestic variables react relatively more to the domestic inflation shock. The results, thus, seem to support our hypothesis of the existence of a "global shock" acting as a common propagation mechanism characterizing the cyclical behavior of all the three measures of prices. Both identifications explains around 60% of total fluctuations in domestic inflation and less than 10% of the real block of the model (see Appendix 1.7.7).

Few important comments are in order to avoid confusion about the interpretation our results for this second subsample. First, once again, it is important to stress that we are not trying to attach any structural interpretation to this shock, because the procedure by itself is not built to identify structural shocks. It uncovers a linear combination of reduced-form shocks, that we label "global shock", that drives most of the cyclical movements of domestic inflation and of inflation in imported industrial supplies and materials and in oil. By construction, we have nothing to say about what causes these correlations. As a matter of fact, they can arise both from surprises in global demand as well as global supply of internationally traded intermediate inputs. Second, our focus is on understanding if such convolution of shocks play a significant role in explaining the *disconnection* between U.S. inflation and domestic slack. In this respect, we think it is fair to say that this "global shock" seems exogenous to U.S. domestic cycle, because the IRFs imply no or minor movements in the U.S. real variables and they are quite precisely estimated (see Figure 1.8). Let's explore in details a possible counterargument. One might argue that some part of this convolution of reduced-form shocks could be due to an increase in U.S. (or global) demand, which raises the world demand for raw materials and intermediate inputs, pushing up international prices. If that would be the case, then we should observe a positive correlation also of U.S. real variables with the "global shock", while we don't. We might not observe this positive correlation because some other parts of the reduced-form convolution of shocks could be due to a typical supply shock to oil, for example, which raises the world prices for raw materials and intermediate inputs, but has a negative impact on U.S. real variables. The net effect of these two parts would then yield that U.S. real variables are uncorrelated with this "global shock". If this was the case, however, we would observe large confidence bands around the response of U.S. variables to the "global shock", because these variables sometime would go up - when due to U.S. demand - and sometime they would go down - when due to an oil supply shock. However, this is not the case: the IRFs are tight, which can be rationalized with the above argument only if these supposed supply and demand shocks are somewhat highly correlated, which is an obviously impossible assumption to defend. Moreover, such an argument would be difficult to square with the results in the bottom panel of Table 1.2.

To conclude, let us be crystal clear on the main takeaways from the analysis in this Section. The aim of the analysis is not to identify a "supply shock". First, the analysis suggests that in the second subsample the drivers of U.S. domestic business cycle do not explain much of the FEV of inflation and vice versa. Second, it suggests the existence of an international 'supply channel' for U.S. inflation, because in the second subsample dynamics of inflation share common drivers with the dynamics of inflation of international input prices - imported intermediate goods and oil. These drivers, however, seem not to affect the domestic U.S. real variables. These two facts naturally lend themselves to a natural explanation of the flattening of the Phillips Curve.

### 1.5.3 The Flattening of the Phillips Curve

Figure 1.10 displays the reduced-form slope of the Phillips curve obtained by projecting the model consistent inflation gap onto the space spanned by unemployment gap for pre- and Great Moderation time windows. When comparing the two periods, we find that the estimated reduced-form slope has significantly declined (in absolute value) at 68% credibility level from approximately 0.73 (median of the red density) down to 0.4 (see the blue density).

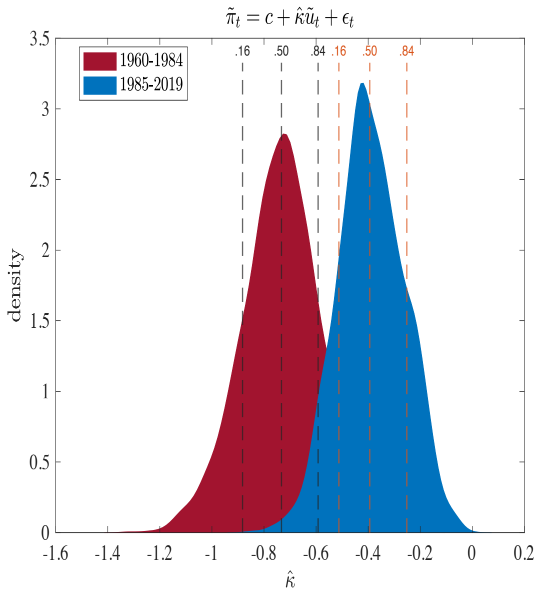


Figure 1.10: Distribution of the reduced-form slope of the Phillips Curve.

Additional support is provided from the two scatter plots in Figure 1.11, showing the relationship between the inflation gap and the unemployment gap in the two sub-periods. The scatter plot in the right-hand panel of Figure 1.11 exhibits a flatter relationship, in accordance with the evidence in the literature about the flattening of the Phillips Curve for the Great Moderation period. While we do not perform a structural estimation of a Phillips curve, the cyclical components estimated with Kalman filter naturally lend themselves to get an estimate of a reduced-form of Phillips curve relationship that should concern the cyclical components, given that our approach has already filtered out the low frequencies.[16] The results clearly shows evidence of a flattening of the Phillips curve.

[16]Think about the standard NK Phillips Curve, in which variables are usually expressed in log-deviation from trend.

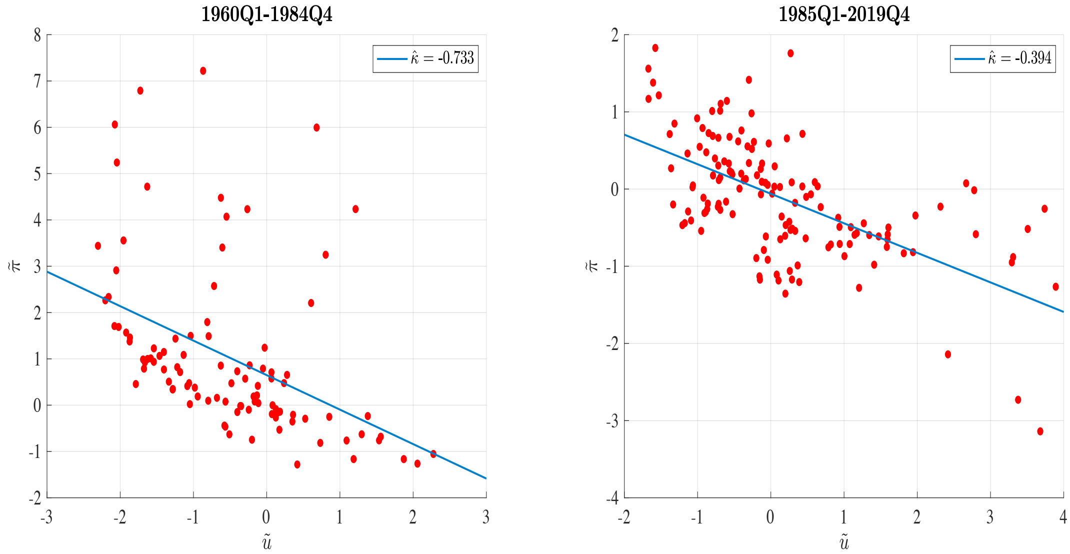


Figure 1.11: Scatter plot of the cyclical components of unemployment. Left panel: 1960Q1-1984Q4 ; right panel: 1985Q1-2019Q4.

### 1.5.4 The Wage Phillips Curve is Alive and Well

This Section shows that the labor market conditions are still a reliable barometer of business cycle that materializes in wage pressure also in the post-'90s sample. As a result, the flatness of the price Phillips Curve is not due to a flat wage Phillips Curve.

Recently, there has been a growing interest in exploring the hypothesis of a flatter wage Phillips curve relationship. Galí and Gambetti (2019), for instance, test this hypothesis by estimating a VAR with time-varying parameters, and provide evidence against the flattening of the wage Phillips curve. We estimate the shock that maximizes the cyclical variance of wages over business cycle frequencies. Had the relationship between wages and unemployment stayed strong and stable also in the second subsample, that is, had the slope of the wage Phillips Curve remained steep and constant, we should observe wage dynamics to share the same main sources of fluctuations of the unemployment gap in both subsamples. Figure 1.12 and 1.13 plot the responses to the shock that maximizes the FEVD of $\tilde{\pi}^w_t$ for the 1960Q1-1984Q4 and the 1985Q1-2019Q4 periods, respectively. On top of these responses, we also plot the impulse response functions to the shock that maximizes the FEVD of the unemployment gap, $\tilde{u}_t$, - from Figures 1.5 and 1.8, respectively. The plots visualize the presence of an unique common underlying propagation mechanism, suggesting a strong link between $\tilde{\pi}^w_t$ and $\tilde{u}_t$ in both subsamples. Hence, these results, together with the results reported in Table 1.2, point to a rejection of the flattening of the wage Phillips curve, providing additional evidence against the hypothesis of the change in the functioning of the labor market as being responsible for the flattening of the Phillips curve.

The evidence therefore indicates that wage inflation respond to the domestic business cycle in both subsamples, but it does not pass-through to inflation anymore in the sec-

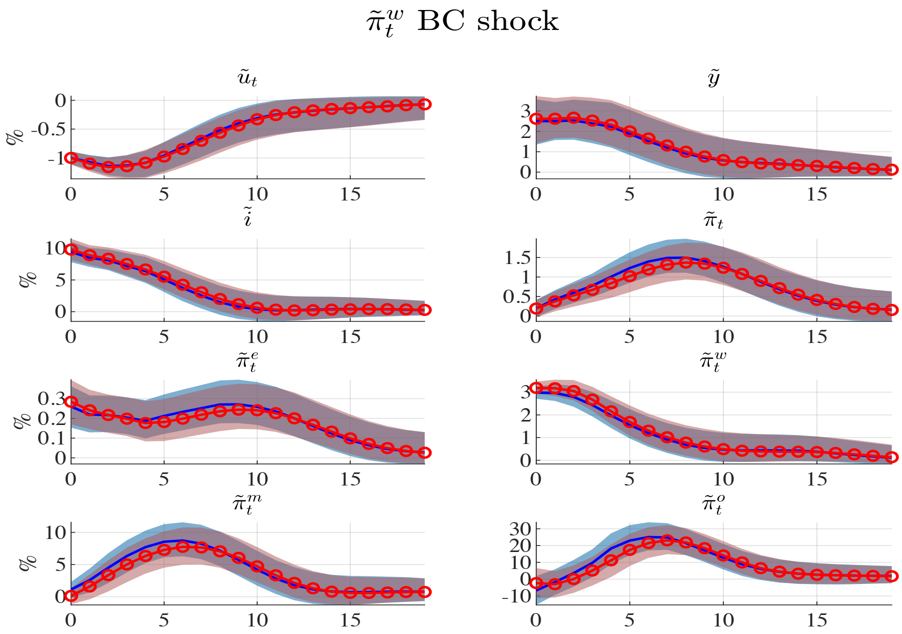


Figure 1.12: IRFs maximizing the FEVD of $\tilde{\pi}^w$ (red -○-) and $\tilde{u}$ (blue –). Sample 1960Q1-1984Q4; median response and 68% uncertainty band.

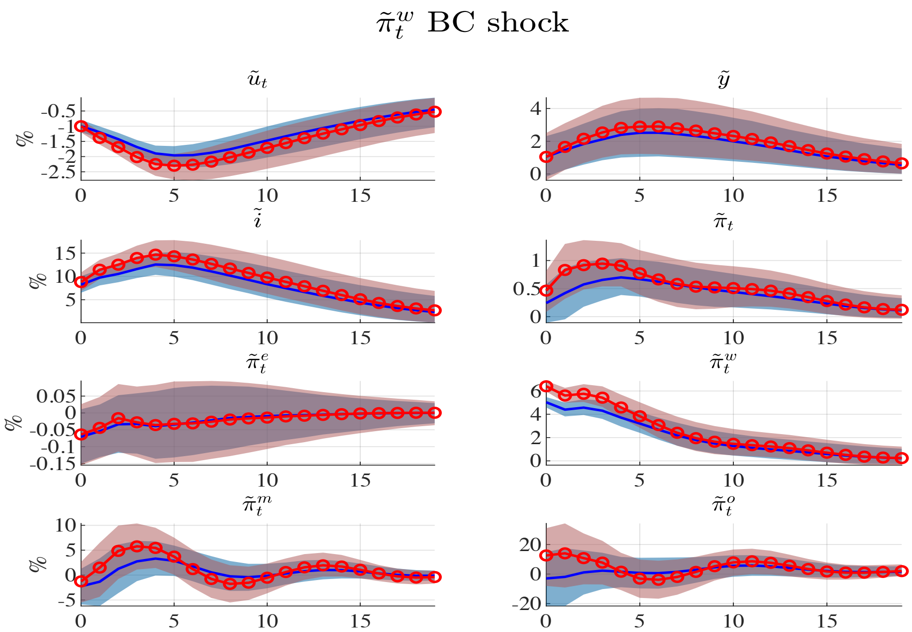


Figure 1.13: IRFs maximizing the FEVD of $\tilde{\pi}^w$ (red -○-) and $\tilde{u}$ (blue –). Sample 1985Q1-2019Q4; median response and 68% uncertainty band.

ond subsample.[17] Figure 1.14 confirms this. Borrowing from the literature on dynamic multipliers and, similarly to Conti et al. (2019) and Forbes (2019), we compute dynamic

[17]Other studies report similar results, see Coibion and Gorodnichenko (2015a), Galí and Gambetti (2019), Forbes (2019), Del Negro et al. (2020) and Heise et al. (2020). Similar results have been found on Euro area inflation by Conti et al. (2019), who claim that inflation has become inelastic to the wage pass-trough stemming from an aggregate demand shock.

pass-through effects[18] for both subsamples as the ratio between the cumulative IRs to the "business cycle shock" of price and wage inflation, which Figure 1.14 displays. On the one hand, in the first subsample, the wage response almost fully passes-through to inflation after 4 years, consistent with the timing of the dynamic response of inflation. On the other hand, in the second subsample, the pass-through is four times smaller over 5 years horizon. The fact that the increase in wage inflation no longer passes-through to price inflation is consistent with the observation that the "business cycle shock" is not associated with an upward revision of inflation expectations in the second subsample - see Figures 1.5 and 1.8. Indeed, there is evidence of a strong anchoring of long-term inflation expectations in this second subsample.

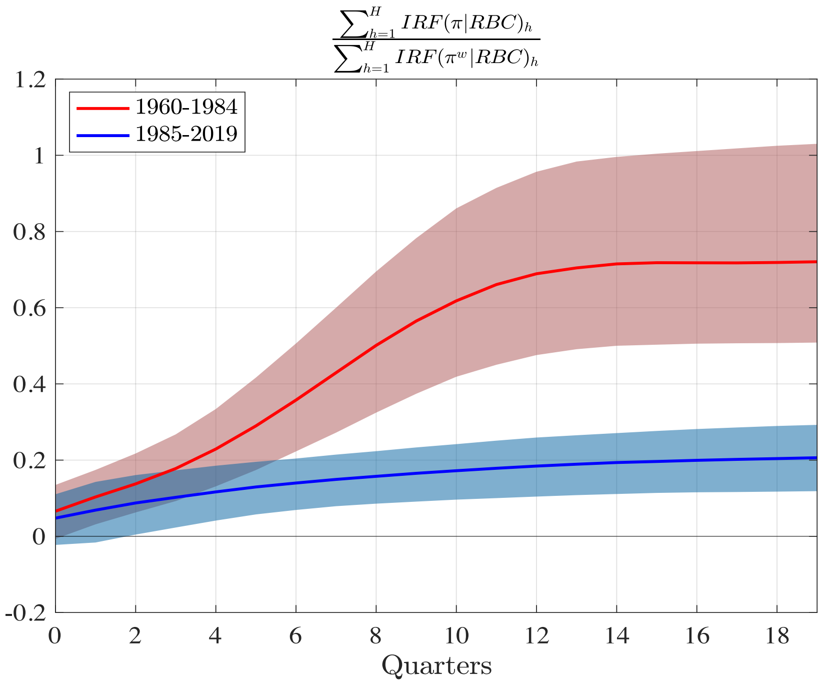


Figure 1.14: response of inflation to the "business cycle" shock over the response of wage inflation and 68% uncertainty band.

In Appendix 1.7.5, we provide additional evidence of this *disconnection* by inspecting the spectral densities of the estimated cycles. In the 1960Q1-1984Q4 period, both unemployment and inflation cyclical variance peaks after approximately 25 quarters, within the conventional 6-32 business cycle window. However, in the 1985Q1-2019Q4 period we observe a timing disconnection between the cycle of unemployment and inflation. On the one hand, similarly to what documented by Beaudry et al. (2020), the cycle of unemployment becomes more persistent peaking after 40 quarters and does not display other significant peaks within the usual business cycle windows. On the other hand, the cycle of inflation peaks within the 6-32 frequency band, and is characterized by multiple local peaks sharing the same timing of the cyclical peaks observed in the spectral densities of

[18]Here we slightly abuse of terminology, as these are dynamic conditional correlations.

international factor prices.

Therefore, this analysis provides additional evidence of the disconnection between inflation and the domestic business cycle. Heise et al. (2020) focus on a demand side channel, whereby the decrease of the wage pass-through to U.S. inflation is due to the increased imported competition that reduces domestic firms' ability to change prices of *final goods* in response to fluctuations in the domestic labor market. By looking, instead, at prices of *intermediate inputs*, we provide evidence that there is also an important supply side channel to the explanation, as international cost-push factors affect firms' costs and thus domestic inflation. Our results, therefore, should be seen as complementary to the ones in Heise et al. (2020).

Overall, the takeaways from our analysis can be summarized as follows. First, in the 1960Q1-1984Q4 sample, the link between inflation and domestic economic slack was strong, with a common propagation mechanism jointly driving inflation and the real domestic cycle. Second, according to our procedure, the drivers that generate most of the volatility in the real block and in the labor market variables in the 1985Q1-2019Q4 sample are no longer associated with inflationary pressures. The reason for this lies on the limited pass-through of wages to prices, despite the sound response of wage inflation. The lower responsiveness of inflation cannot thus be attributed to structural breaks in the labor markets, because wages are still tightly linked to the other measures of economic slack. Third, while disconnected from the domestic real and labor market variables, the cyclical behavior of inflation is subject to the same international driving forces of international prices of intermediate inputs - imported industrial supplies and materials - and oil. These results suggest that the share of cyclical inflation dynamics explained by domestic variables has decreased substantially, while the cost-push component generated internationally via international linkages has grown in importance. Despite our procedure does not allow us to provide a structural interpretation of this "global shock", it indicates that this shock is not endogenous to the domestic cycle. Had it been endogenous, we would have IRs similar to those obtained in the first subsample or large credibility bands.

## 1.6 Concluding Remarks

This paper proposes a novel approach to study the role of international factors contributing both to the decline of trend inflation and to the flatness of the slope of the Phillips curve. We implement a multivariate unobserved component analysis, which enables to explicitly isolate the frequencies of interest and then analyze both the trend and the cycle components.

In the analysis of the slow-moving drivers of inflation, we allow for the presence of common stochastic trends and propose an anatomy of trend inflation, decomposing it into two distinct low-frequency components. Our anatomy distinguishes between (domestic) monetary and international - imported intermediate goods - determinants of trend inflation. The results attribute the decline in trend inflation observed from the mid-‘80s to both the monetary policy regime switch toward explicit targeting and the dynamics of international prices of imported intermediate goods.

From the analysis on the cyclical block of the model several facts emerge. First, the Phillips Curve relationship between the unemployment gap and the inflation gap shows a strong flattening over time.

Second, the impulse response analysis in the 1985Q1-2019Q4 sample uncovers two main facets of business cycle. On the one hand, in accordance to the results in Angeletos et al. (2020), there is a common propagation mechanism among the cyclical components of the main real variables - i.e., a “business cycle shock” - which is non-inflationary. This is mainly attributable to the fact that the transmission channel from wage to inflation dynamics has broken down in this subsample - while was active in the 1960Q1-1984Q4 sample. The analysis on the wage Phillips curve relationship in Section 1.5.4 shows that the relationship between wage inflation and unemployment has remained relatively stable over time and, thus, the labor market continues to be a reliable barometer of business cycle pressure. The muted response of inflation is, thus, neither due to an aggressive monetary policy reaction, nor to a structural change in the local labor market dynamics. In fact, our analysis relies on conditional correlations and, therefore, the unemployment gap should have also remained closed had it been monetary policy the reason behind the muted response of inflation. On the other hand, the business cycle behavior of inflation in the 1985Q1-2019Q4 sample is mainly characterized by a shock originating abroad, which generates the main bulk of volatility in international prices of intermediate goods and is almost orthogonal to the domestic slack. Therefore, we conclude that, in the sample 1985Q1-2019Q4, domestic inflation disconnected from the local labor market and increasingly co-moved with the prices of imported intermediate inputs and oil, through international linkages.

Overall, our results suggest that the international component of inflation should not only be considered a business cycle phenomenon, but it could also leave long-run scars on the level of inflation. These results pose some new and crucial challenges for the ability of monetary policy to govern both the persistent component and cyclical behavior of inflation. Central banks should be aware that an external factor beyond their direct control could potentially undermine both the long-run mandate of the central bank, in terms of achieving the inflation target, and the stabilization of inflation dynamics around the target, because of the insensitivity of inflation to the domestic business cycle conditions.

## 1.7 Appendix

### 1.7.1 Data

All data are available in FRED website with the exception of long-run inflation expectations. The long-run PCE inflation expectations are obtained from the Survey of Professional Forecasters from 2007 onward, while for the period from 1970 to 2006, we use the survey-based long-run (5- to 10-years ahead) PCE inflation expectations series of the Federal Reserve Board's FRB/U.S. econometric model.[19] All nominal variables,

Table 1.3: Data codes

| DATA | CODE |
|---|---|
| Unemployment Rate | UNRATE |
| Real Gross Domestic Product per capita | A939RX0Q048SBEA |
| Gross Domestic Product | GDP |
| Gross Private Domestic Investment | GPDI |
| Personal Consumption Expenditures: Durable Goods | PCDG |
| Personal consumption expenditures (implicit price deflator) | DPCERD3Q086SBEA |
| 10-year ahead PCE expected inflation rate | |
| Compensation of Employees: Wages and Salary Accruals | WASCUR |
| Imports of goods: Industrial supplies and materials, except petroleum (chain-type price index) | B649RG3Q086SBEA |
| Spot Crude Oil Price: West Texas Intermediate (WTI) | WTISPLC |

namely PCE inflation, long-run expected inflation rate, wage inflation, imported intermediate input inflation, oil inflation are expressed in year-on-year percentage changes. Real investment per capita is the sum gross private domestic investment and durable goods multiplied by real GDP per capita and divided by nominal GDP. Finally, all variables but unemployment rate are in logs.

[19]The log-run inflation expectations series is the same used by Del Negro et al. (2017b) and is available at https://github.com/FRBNY-DSGE/rstarBrookings2017.

### 1.7.2 Model in Matrix Notation

$$\begin{bmatrix} u_t \\ y_t \\ i_t \\ \pi_t \\ \pi_t^e \\ \pi_t^w \\ \pi_t^m \\ \pi_t^o \end{bmatrix} = \begin{bmatrix} 0 \\ 1 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} c_t + \begin{bmatrix} 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 1 & \lambda_{4,2} & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & \lambda_{7,2} & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & \lambda_{8,2} & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \bar{\pi}_t^e \\ \bar{g}_t \\ \bar{u}_t \\ \bar{y}_t^{id} \\ \bar{i}_t^{id} \\ \bar{\pi}_t^{w,id} \\ \bar{\pi}_t^{m,id} \\ \bar{\pi}_t^{o,id} \end{bmatrix} + \begin{bmatrix} \tilde{u}_t \\ \tilde{y}_t \\ \tilde{i}_t \\ \tilde{\pi}_t \\ \tilde{\pi}_t^e \\ \tilde{\pi}_t^w \\ \tilde{\pi}_t^m \\ \tilde{\pi}_t^o \end{bmatrix}$$

### 1.7.3 Gibbs Sampler For The Estimation of The VAR

The model is estimated employing a Gibbs sampler, which is structured into two steps:

1. The algorithm draws from the joint distribution $\bar{y}_{0:T}, \tilde{y}_{-p+1:T}, \lambda | c, \phi, \Sigma_e, \Sigma_\varepsilon, y_{1:T}$, which is given by the product of the marginal posterior of $\lambda$ conditional on the other parameters $\lambda | c, \phi, \Sigma_e, \Sigma_\varepsilon, y_{1:T}$:

$$p(\lambda | c, \phi, \Sigma_e, \Sigma_\varepsilon, y_{1:T}) \propto L(y_{1:T} | \lambda, c, \phi, \Sigma_e, \Sigma_\varepsilon) p(\lambda)$$

   The posterior distribution of $\lambda$ is approximated implementing a Metropolis Hastings step within the Gibbs sampler. The posterior of the states $\bar{y}_{0:T}, \tilde{y}_{-p+1:T}$ conditional on $\lambda$ and the other parameters is estimated using Durbin and Koopman (2002b)'s simulation smoother to draw the latent states.

2. The second step involves the estimation of the two VARs. The posterior distribution of $\Sigma_e$ are given by:

$$\Sigma_e | \bar{y}_{0:T} \sim \mathcal{IW}(\underline{\Sigma}_e + \hat{S}_e, \kappa_e + T),$$

   where $\hat{S}_e$ is the sum of squared errors of the latent trends. The posterior distributions of the coefficients of the stationary VAR are given by:

$$\Sigma_\varepsilon | \tilde{y}_{0:T} \sim \mathcal{IW}(\underline{\Sigma}_\varepsilon + \hat{S}_\varepsilon, \kappa_\varepsilon + T)$$
$$p(\phi | \Sigma_\varepsilon, \tilde{y}_{0:T}) \sim \mathcal{N}(vec(\hat{\Phi}), \Sigma_\varepsilon(\tilde{X}\tilde{X}' + \underline{\Omega}^{-1})^{-1}),$$

   where $XX' = \sum_{t=1}^T \tilde{x}_t \tilde{x}_t'$, $\hat{S}_\varepsilon = (\tilde{X}\tilde{X}' + \underline{\Omega}^{-1})^{-1}(\tilde{X}\tilde{\boldsymbol{y}} + \underline{\Omega}^{-1}\underline{\Phi})$, $\hat{S}_\varepsilon = \boldsymbol{\varepsilon}\boldsymbol{\varepsilon}' + (\hat{\Phi} - \underline{\Phi})'\underline{\Omega}^{-1}(\hat{\Phi} - \underline{\Phi})$ and $\boldsymbol{\varepsilon} = \tilde{\boldsymbol{y}} - \hat{\Phi}'\tilde{X}$.

### 1.7.4 Identification Scheme

The impulse response analysis follows Uhlig (2003). The reduced-form VAR of cyclical components is given by:

$$\tilde{\chi}_t = C(L)\varepsilon_t \qquad \varepsilon_t \sim N(0_n, \Sigma_\varepsilon) \tag{1.12}$$

where $C(L) = \Phi(L)^{-1}$ and $\varepsilon_t = A\upsilon_t$ composite innovations. Let A be the impulse matrix obtained from some decomposition of the $\Sigma_\varepsilon$:

$$E[\varepsilon_t \varepsilon_t'] = \Sigma_\varepsilon = AE[\upsilon_t \upsilon_t']A' = AA'$$

Now, assume $\hat{A}$ being an alternative decomposition of $\Sigma_\varepsilon$. For sake of simplicity let $\hat{A}$ be the Cholesky triangular factor, such that:

$$\Sigma_\varepsilon = \hat{A}\hat{A}'$$

Then, there must exist an orthonormal matrix Q that enables to reconcile $\hat{A}$ with A:

$$A = \hat{A}Q \tag{1.13}$$

Now, the k-th step ahead forecast error is given by:

$$\epsilon_{t+k} = \sum_{i=0}^{k} \hat{B}_i Q \upsilon_{t+k-i} \tag{1.14}$$

where $\hat{B}(L) = C(L)\hat{A}$. The variance covariance matrix of the k-th step ahead forecast error is given by $\Sigma_\varepsilon(k) = \sum_{i=0}^{k} \hat{B}_i \hat{B}_i'$. It is possible to further decompose the variance so to get the contribution of the j-th shock:

$$\Sigma_\varepsilon(k, j) = \sum_{i=0}^{k} (\hat{B}_i q_j)(\hat{B}_i q_j)' \tag{1.15}$$

The goal is to find the impulse vector that maximizes the forecast error variance of the selected variable over a specific forecast horizon, say $[\underline{k}, \overline{k}]$ , as follows:

$$\begin{aligned} \sigma_\varepsilon^2(\underline{k}, \overline{k}; q_1) &= q_1' \left( \sum_{k=\underline{k}}^{\overline{k}} \sum_{i=0}^{k} \hat{B}_i' \hat{B}_i \right) q_1 \\ \sigma_\varepsilon^2(\underline{k}, \overline{k}; q_1) &= q_1' \mathcal{S} q_1 \end{aligned} \tag{1.16}$$

Recall that $q_1$ is a column of the orthonormal matrix Q, so it must be orthonormal itself. Finally, if we write the program in its Lagrangian form:

$$\mathcal{L}(q_1) = q_1' \mathcal{S} q_1 - \lambda[q_1' q_1 - 1]$$
$$\text{F.O.C.:}$$
$$\mathcal{S} q_1 = \lambda q_1.$$

The problem eventually boils down to an eigenvector-eigenvalue problem. Hence, the orthonormal vector $q_1$ is the eigenvector associated with the largest eigenvalue of the forecast error variance over a specific frequency interval.

### 1.7.5 Further Evidence From Spectral Analysis

**Refresh on Spectral Density Estimation**

Take zero-mean covariance stationary random variable $y_t$ and compute its sample autocovariance function:

$$\hat{\gamma}_y(j) = \frac{1}{T} \sum_{t=j+1}^{T} (y_t - \bar{y})(y_{t-j} - \bar{y}),$$

where $\bar{y} = \frac{1}{T}\sum_{t=1}^{T} y_t$ is the sample mean. The sample autocovariance is the object we will plug into the Fourier transform to express the atucovariance structure of $y_t$ as function of waves. In other words, we will make use of the Fourier transform to map the autocovariance structure from the time domain to the frequency domain. The theoretical spectrum is retrieved by means of the discrete Fourier transform:

$$f_y(\omega) = \frac{1}{2\pi} \sum_{j=-\infty}^{+\infty} \gamma_y(j) e^{-i\omega j}$$

$\omega = \frac{2\pi k}{T}$ is the frequency (i.e.: how quickly the process oscillates). The theoretical spectral density of $y_t$ is given by:

$$f_y(\omega) = \frac{1}{2\pi} \sum_{j=-\infty}^{+\infty} \gamma_y(j) \cos(\omega j)$$
$$f_y(k) = \frac{1}{2\pi} \sum_{j=-\infty}^{+\infty} \gamma_y(j) \cos\left(\frac{2\pi k}{T} j\right)$$

In practice, the estimation is far from being a trivial task. This is because data are finite. This implies that the empirical counterpart of the theoretical spectrum is a truncated version called periodogram:

$$\hat{f}_y(k) = \frac{1}{2\pi} \sum_{j=-(T-1)}^{(T-1)} \hat{\gamma}_y(j) \cos\left(\frac{2\pi k}{T} j\right)$$

Tough unbiased, the periodogram is an inconsistent estimator of the theoretical spectrum. To cope with the large variance associated with inconsistency, an auxiliary function is convoluted with the autocovariance so to smooth the variance. The auxiliary function is called window. Window functions are basically weighting functions and there are of different families. For the estimation of cycles' spectra in our paper we use Hamming window smoothing[20]. Once the window smoothing is applied to the periodogram, the estimator looks like:

$$\hat{f}_y(k) = \frac{1}{2\pi} \sum_{j=-(T-1)}^{(T-1)} w_k(j)\hat{\gamma}_y(j) \cos\left(\frac{2\pi k}{T} j\right)$$

#### Results From Spectral Analysis

Cycles are by construction zero-mean covariance stationary processes, thus they are suitable candidates for spectral analysis. We, therefore, perform the spectral analysis to further corroborate the above results. First, the analysis provides an additional robustness check of the loosening of the relationship between the unemployment gap and the inflation gap since the '90s. Second, if the cyclical behaviour of inflation is originated by a shared propagation mechanism with international prices, then the volatility peak should be tuned on the approximately the same frequencies. The same intuition applies for real variables, so our prior is that they are all tuned on the same frequencies, as well.

Figure 1.15 plots the spectral densities for unemployment and inflation gap over the two samples. The first row compares the spectral densities over the 1960Q1-1984Q4 sample. Unemployment and inflation gap are approximately synchronized on the same frequencies and the peak for both variables occurs at Q25. In contrast, when analysing the 1985Q1-2019Q4 sample (second row), the variance peak of unemployment gap occurs around Q40, suggesting a more persistent behaviour of unemployment in the Great Moderation era. This result is consistent with a recent work of Beaudry et al. (2020), who run a spectral analysis on labour market variables, finding that the variance peak of unemployment realizes after roughly 40Q. Moving to inflation gap, it seems hard to clearly detect a global peak. However, it seems that most of the fluctuations materializes at higher frequencies compared with unemployment gap, with a local peak around Q25.

Figure 1.17 visualizes the spectral density of imported intermediate inputs inflation for the two subsamples. Comparing them to the densities of domestic inflation clearly

[20] As robustness check, we also apply Bartlett triangular smoothing and the results do not change.

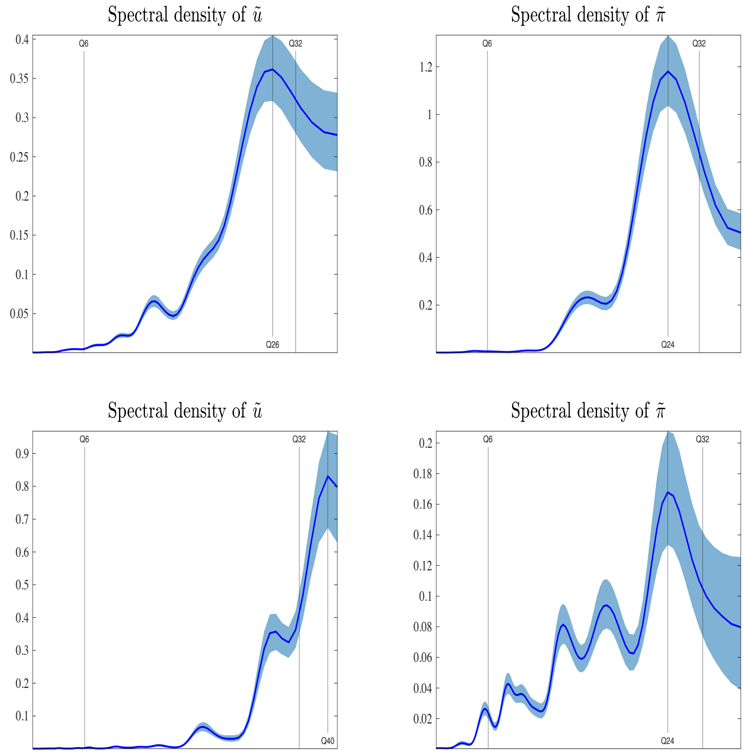


Figure 1.15: First row: sample 1960Q1-1984Q4; the second row: sample 1985Q1-2019Q4.

emerges that the two variables are synchronized on the same frequencies.

Finally, the spectral densities of real variables in the second subsample corroborates once more the hypothesis of a common propagation mechanism. Despite the large uncertainty bands, Figure 1.16 shows that all the gaps in the real variables exhibit the same timing of the realization of the variance peak as the unemployment gap, thus supporting the existence of the MBC as Angeletos et al. (2020). Wage inflation exhibits a peak after 36Q, providing additional evidence in support of an alive-and-well wage Phillips Curve.

To sum up, the main takeaways from the spectral analysis can be summarized as follows. Firstly, the unemployment gap and the inflation gap shared the frequency peak, when considering the pre-Great Moderation era. In contrast, since 1990s, their fluctuations peak at very distant frequencies, implying a disconnection between the two variables. Secondly, domestic inflation and imported inflation spectrum peaks are steadily synchronized throughout the two samples. Last but not least, all the gaps in real variables, and to some extent in wage inflation, exhibit their variance peak at the same frequencies of unemployment gap, supporting labor market variables as good barometers of business cycle pressure.

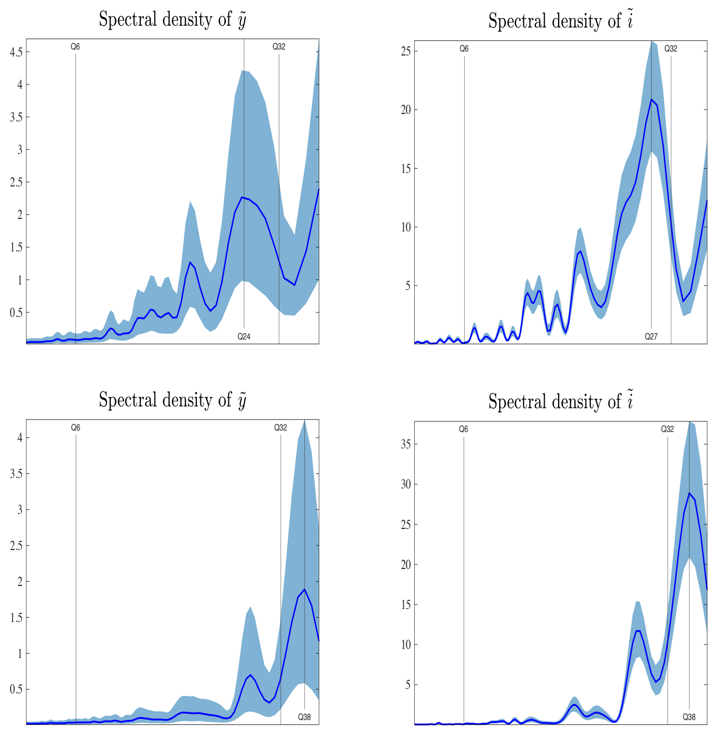


Figure 1.16: First row: sample 1960Q1-1984Q4; the second row: sample 1985Q1-2019Q4.

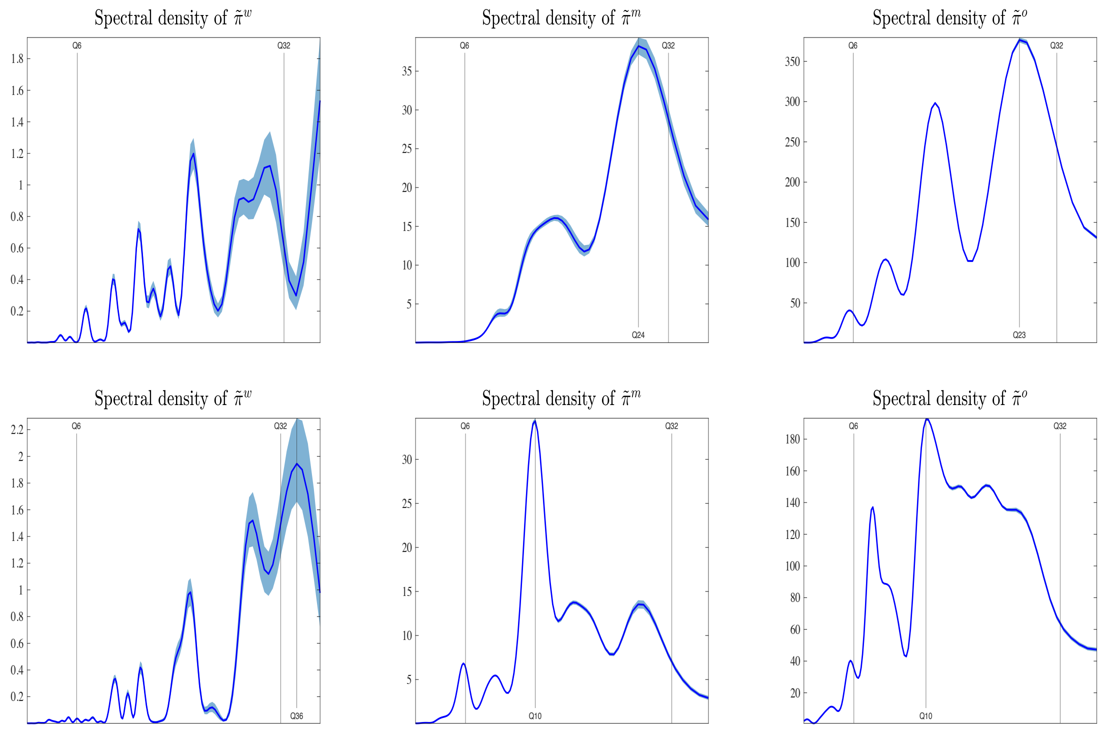


Figure 1.17: First row: sample 1960Q1-1984Q4; the second row: sample 1985Q1-2019Q4.

### 1.7.6 Additional Figures

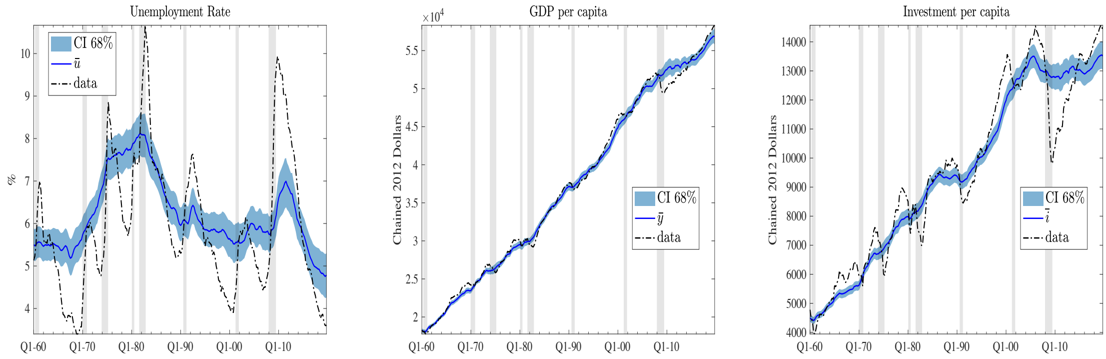


Figure 1.18: Trends of real variables.

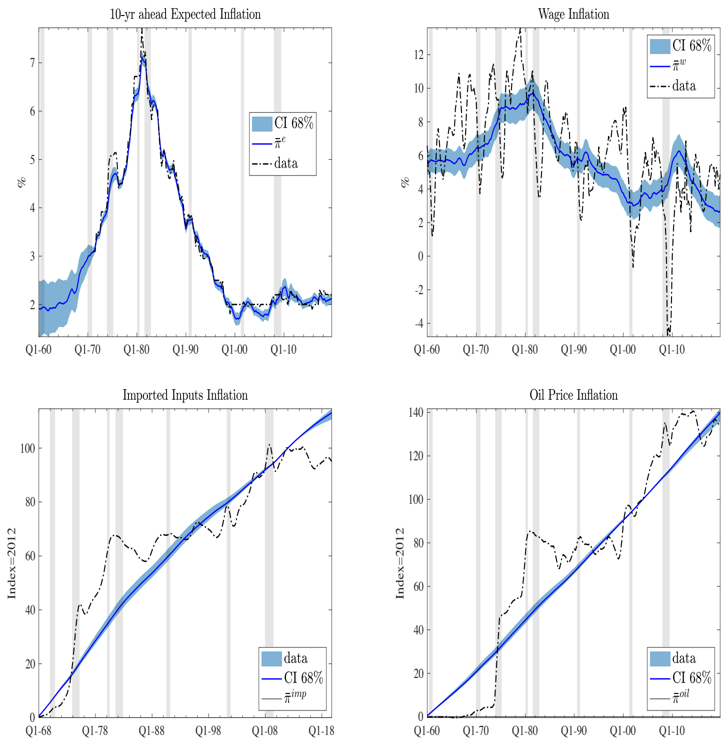


Figure 1.19: Trends of nominal variables.

### 1.7.7 Additional Tables

Table 1.4: Forecast error variance decomposition.
68% uncertainty band in squared brackets.

| **1985Q1-2019Q4** | | | |
|---|---|---|---|
| **$\tilde{\pi}_t^m$ shock** | | | |
| $\tilde{u}_t$ | $\tilde{y}_t$ | $\tilde{i}_t$ | $\tilde{\pi}_t$ |
| 0.0566 | 0.0679 | 0.0568 | 0.5498 |
| [0.0378,0.0843] | [0.0461,0.0972] | [0.0378,0.0840] | [0.4665,0.6248] |
| $\tilde{\pi}_t^e$ | $\tilde{\pi}_t^w$ | $\tilde{\pi}_t^m$ | $\tilde{\pi}_t^o$ |
| 0.1288 | 0.0981 | 0.7507 | 0.5611 |
| [0.0902,0.1815] | [0.0706,0.1319] | [0.7073,0.7894] | [0.4739,0.6427] |
| **$\tilde{\pi}_t^o$ shock** | | | |
| $\tilde{u}_t$ | $\tilde{y}_t$ | $\tilde{i}_t$ | $\tilde{\pi}_t$ |
| 0.0501 | 0.0738 | 0.0548 | 0.5963 |
| [0.0332,0.0756] | [0.0522,0.1034] | [0.0357,0.0794] | [0.5318,0.6572] |
| $\tilde{\pi}_t^e$ | $\tilde{\pi}_t^w$ | $\tilde{\pi}_t^m$ | $\tilde{\pi}_t^o$ |
| 0.1753 | 0.0994 | 0.5221 | 0.9198 |
| [0.1264,0.2374] | [0.0725,0.1310] | [0.4712,0.5693] | [0.8970,0.9394] |
| **$\tilde{\pi}_t^w$ shock** | | | |
| $\tilde{u}_t$ | $\tilde{y}_t$ | $\tilde{i}_t$ | $\tilde{\pi}_t$ |
| 0.8776 | 0.4148 | 0.8145 | 0.3127 |
| [0.8451,0.9055] | [0.2978,0.5195] | [0.7746,0.8539] | [0.2396,0.3845] |
| $\tilde{\pi}_t^e$ | $\tilde{\pi}_t^w$ | $\tilde{\pi}_t^m$ | $\tilde{\pi}_t^o$ |
| 0.1673 | 0.8984 | 0.0786 | 0.0672 |
| [0.1074,0.2550] | [0.8685,0.9222] | [0.0550,0.1087] | [0.0468,0.0951] |

Table 1.5: Forecast error variance decomposition.
68% uncertainty band in squared brackets.

| **1960Q1-1984Q4** | | | |
|---|---|---|---|
| **$\tilde{y}_t$ shock** | | | |
| $\tilde{u}_t$ | $\tilde{y}_t$ | $\tilde{i}_t$ | $\tilde{\pi}_t$ |
| 0.3444 | 0.6689 | 0.3156 | 0.1880 |
| [0.1946,0.5594] | [0.6330,0.7061] | [0.1838,0.4947] | [0.0989,0.3343] |
| $\tilde{\pi}^e_t$ | $\tilde{\pi}^w_t$ | $\tilde{\pi}^m_t$ | $\tilde{\pi}^o_t$ |
| 0.1862 | 0.3403 | 0.1556 | 0.1109 |
| [0.1003,0.3386] | [0.1914,0.5422] | [0.0822,0.2730] | [0.0613,0.1885] |
| **$\tilde{i}_t$ shock** | | | |
| $\tilde{u}_t$ | $\tilde{y}_t$ | $\tilde{i}_t$ | $\tilde{\pi}_t$ |
| 0.7867 | 0.3886 | 0.7657 | 0.4367 |
| [0.7369,0.8326] | [0.3116,0.4812] | [0.7299,0.7979] | [0.3466,0.5258] |
| $\tilde{\pi}^e_t$ | $\tilde{\pi}^w_t$ | $\tilde{\pi}^m_t$ | $\tilde{\pi}^o_t$ |
| 0.4305 | 0.7715 | 0.3585 | 0.2549 |
| [0.3462,0.5209] | [0.7194,0.8184] | [0.2823,0.4369] | [0.1979,0.3217] |
| **1985Q1-2019Q4** | | | |
| **$\tilde{y}_t$ shock** | | | |
| $\tilde{u}_t$ | $\tilde{y}_t$ | $\tilde{i}_t$ | $\tilde{\pi}_t$ |
| 0.4479 | 0.6315 | 0.4309 | 0.1073 |
| [0.0765,0.8416] | [0.5839,0.6901] | [0.0788,0.8080] | [0.0434,0.2233] |
| $\tilde{\pi}^e_t$ | $\tilde{\pi}^w_t$ | $\tilde{\pi}^m_t$ | $\tilde{\pi}^o_t$ |
| 0.0611 | 0.3925 | 0.0440 | 0.0409 |
| [0.0261,0.1414] | [0.0677,0.7328] | [0.0232,0.0721] | [0.0215,0.0700] |
| **$\tilde{i}_t$ shock** | | | |
| $\tilde{u}_t$ | $\tilde{y}_t$ | $\tilde{i}_t$ | $\tilde{\pi}_t$ |
| 0.9161 | 0.4465 | 0.8742 | 0.2661 |
| [0.8938,0.9348] | [0.3405,0.5583] | [0.8471,0.8976] | [0.1214,0.3479] |
| $\tilde{\pi}^e_t$ | $\tilde{\pi}^w_t$ | $\tilde{\pi}^m_t$ | $\tilde{\pi}^o_t$ |
| 0.1926 | 0.8382 | 0.0667 | 0.0609 |
| [0.1087,0.2920] | [0.7778,0.8748] | [0.0468,0.0926] | [0.0415,0.0869] |

# Chapter 2

# The Macroeconomic Effects of the Gender Revolution

## 2.1 Introduction

Women's increased participation to the labor market is one the most significant changes that modern economies have experienced during the last century. This prominent change manifested as a slow-moving trend, a "quiet revolution" as named by Goldin (2006). The employment rate for females was less a half of the employment rate for males in 1960 in the US, as shown in Figure 2.1. Since then, the female employment gap (i.e. the ratio of female employment over male employment) increased steadily until 70 per cent in the mid 1980s and converged more gradually to a value of 85 percent around 2020. At the same, a similar convergence happened in terms of wages. The gender wage gap, i.e. the ratio between women's wages and men's wages, stayed relatively flat until 1975, at a time when women's employment was already rising significantly, but grew steadily in pair with employment since then.

The goal of our paper is to investigate the macroeconomic consequences of the gender revolution. More specifically, we quantify its impact on economic growth in the US in terms of GDP, employment and various measures of productivity. In fact, one can think that talent was highly misallocated in an economy in which only 45 percent of females were working, as shown in Hsieh et al. (2019). In addition, our second question of interest is to investigate what factors lie behind the gender converge in employment and wages. The fact that these two trends co-move seems to suggest that labor demand factors must be dominant. However, it is undeniable that a large shock to the supply of skills, and thus a labor supply factor, has also affected the US economy. Our aim is to appropriately disentangle labor demand and labor supply factors once we take into account that a large increase in employment of female workers was concentrated in the market for high-skilled workers.

In order to answer our research questions, we use neoclassical macroeconomic theory in combination with a Structural Vector Autoregressive (SVAR) model, arguably the most standard tool for time series analysis, estimated on macroeconomic data and on gender-specific variables obtained by combining CPS data in the gender dimension with data from Dolado et al. (2021) to account for the skill dimension. Usually, SVAR models are mainly used to study cyclical fluctuations. In contrast, one key aspect of the question at hand is the focus on the macroeconomic effects of slow moving trends, thus implying that SVAR models need to be amended substantially. In practice, we use a SVAR model with common trends as in Del Negro et al. (2017a). The model can be seen as a multivariate unobserved component model in which the variables enter in levels and in which our interest is on the permanent component, and not on the cyclical component as it is often the case in the literature. Put differently, rather than focusing on structural shocks driving the cyclical component, we conduct structural analysis on the permanent component and decompose it into various structural components, both aggregate and gender-specific. As an example, let us focus on GDP. Our model decomposes GDP dynamics into a cyclical component and a permanent component which, in turn, is driven by (gender neutral) technology shocks, (gender neutral) labor supply shocks, gender-specific labor demand shocks and gender-specific labor supply shocks. Naturally, it is crucial to identify these structural drivers. Our key contribution is to derive identifying restrictions from a neoclassical model with gender which allows to disentangle the four driving forces based on their long-run impact on macroeconomic variables. While Del Negro et al. (2017a) calibrate the matrix linking permanent components with structural shocks, we estimate the matrix with Bayesian methods, using our neoclassical model with gender as a guideline to set the priors.

Our main result is that the forces driving the gender convergence in employment and wages are important also for trend GDP growth in the US. They account for up to a half between 1960 and 2000 while their contribution slows down to 15 per cent during the last 20 years, in keeping with the flattening of the employment gap dynamics. In addition, gender specific forces explain also a substantial share of aggregate employment dynamics over the entire sample period. When it comes to the individual role of the two gender specific forces, our model, perhaps not surprisingly, attributes a dominant role to gender-specific labor demand factors. However, it is too premature to declare that gender-specific labor supply factors are irrelevant. In a second step, in fact, we re-estimate our model by using data on the gender employment gap and on the gender wage gap within skilled workers. In the context of this more disaggregate exercise, we confirm that gender shocks are important for US economic growth. In this case, however, the gender employment gap is driven by both a gender labor demand shock and a gender supply shock. Both shocks are quantitatively important. Thus, using data disaggregated by skill allow us to identify an important shock to the supply of skills that was almost irrelevant in the baseline model. Why does such a shock emerge? The gender employment gap within skilled worker has converges faster than the aggregate gender employment gap. At the same time, the gender wage gap within skilled workers has seen a slower convergence than its aggregate counterpart, a fact documented also by Taniguchi and Yamada (2020). Both features are consistent with an important role for a shock to the labor supply of skilled female workers. In contrast, when we focus only on unskilled workers, we document that the gender employment ratio has barely moved over the last 60 years in the US while the gender wage gap has converged significantly faster within unskilled workers than within skilled workers. When we estimate our model using data on gender employment gap and gender wage gap only for skilled workers we find that a negative shock to the supply of unskilled workers is needed to reconcile the absence of convergence in employment and the strong convergence in wages between males and females. All in all, it seems that our baseline finds no role for the gender specific labor supply shock because a positive shock to the supply of skilled female workers is compensated by a negative shock to the supply of unskilled female workers. We conclude that labor supply factors are in fact important once the skill dimension in taken into account.

Our paper contributes to two strands of the literature. First, we contribute to a large literature studying the gender revolution. A useful distinction for our purposes is between papers emphasizing labor demand factors from labor supply factors. Among the former, Galor and Weil (1996) and Buera and Kaboski (2012) emphasize technological factors that favored the demand for women in combination with an increase in the returns to intellectual skills and the rise of the service sector, while Hsieh et al. (2019) point to a reduction in gender discrimination as an important driver of the reduction in the gender wage gap. Among the latter, Albanesi and Olivetti (2016) and Goldin and Katz (2002) document the importance of advances in maternal health and contraception, Fernández et al. (2004) emphasize the importance of cultural factors developed during World War II, Attanasio et al. (2008) point to the crucial role of availability and affordability of child care, while Greenwood et al. (2005) propose a model in which the emergence of home appliances favors female's market production at the expense of home production. We contribute to this literature by proposing a horse race between labor supply and labor demand factors in the context of a macroeconomic time-series model. While less detailed in terms of the underlying transmission mechanisms, our analysis provide a clear link between gender trends and macroeconomic outcomes.

We contribute also to a recent literature emphasizing the role of gender for macroeconomic dynamics in quantitative set-ups. Heathcote et al. (2017) and Hsieh et al. (2019) propose a decomposition of US macroeconomic growth in structural models with gender. Albanesi (2019) estimates with Bayesian methods a real business cycle model with gender with a focus on the importance of gender trends to account for jobless recoveries, while Fukui et al. (forthcoming) propose a similar model with a focus on the fact that rising female participation has not crowded out male participation and has thus been an expansionary factor for the US economy. In a similar spirit to Heathcote et al. (2017) and Hsieh et al. (2019), we propose a decomposition of US trend economic growth but, differently from their approach, within a Structural Vector Autoregression framework, whose identification is motivated by theory.

Finally, the remainder of the paper is organised as follows. Section 2.2 presents the empirical model. The analytical solutions from the neoclassical model are derived in section 2.3. Sections 2.4 and 2.5 discuss the identification strategy. Section 2.6 is devoted to results. Concluding remarks are exposed in section 2.7.

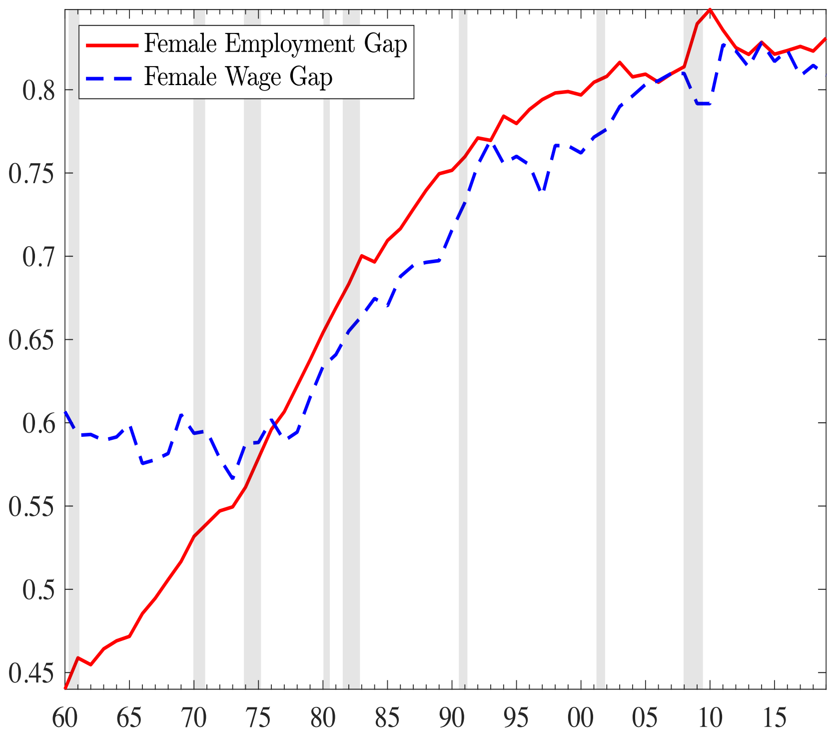


Figure 2.1: Gender gaps are defined as the ratio between female and male employment/wages.

## 2.2 A SVAR with Gender

Assume the $n \times 1$ vector of data $Y_t$ to be the sum of two unobserved states:

$$Y_t = \bar{Y}_t + \hat{Y}_t, \tag{2.1}$$

$\bar{Y}_t$ and $\hat{Y}_t$ describe the slow-moving (trend) and the cyclical fluctuations in the data, respectively. So far, this is a purely *statistical* decomposition. The main focus of our analysis will be on the trend block of this decomposition and, more precisely, on the underlying *structural* forces driving the empirical secular trends observed in the data.

Therefore, we assume the reduced-form (empirical) trends being a function of $q \leq n$ *structural* trends:

$$\bar{Y}_t = \mathcal{V} X_t \tag{2.2}$$

$X_t$ is the $q \times 1$ vector of structural trends and $\mathcal{V}$ is the $n \times q$ matrix that maps the reduced-form trends into structural ones. Importantly, $\mathcal{V}$ embeds the long-run identifying restrictions required to uniquely pin down the structural trends and reconcile $X_t$ with $\bar{Y}_t$. Similarly to Del Negro et al. (2017a) and Ascari and Fosso (2021), *structural* trends and cycles are assumed to be stochastic. Specifically, trends evolve according to a random walk:

$$X_t = c + X_{t-1} + u_t, \qquad u_t \sim \mathcal{N}(0_q, \Sigma_u) \tag{2.3}$$

Since the cycle block of the model is not the core focus of our investigation, we limit the amount of restrictions to imposing that the vector $\hat{Y}_t$ follows a *stationary* reduced-form VAR:

$$\Phi(L)\hat{Y}_t = e_t, \qquad e_t \sim \mathcal{N}(0_n, \Sigma_e) \tag{2.4}$$

$\Phi(L) = I - \Phi_1 L - ... - \Phi_p L^p$ is the $n \times n$ matrix of lag coefficients. Since we use quarterly data, the number of lags is chosen to be $p = 4$. Finally, the vector stacking trend and cycles innovations is assumed to be identically and independently distributed. This assumption implies that permanent and transitory shocks are mutually uncorrelated - i.e.: $cov(u_t, e_t) = 0$. This suggests, for instance, that by construction a technology trend shock does not affect the cycle, as instead conventional in the RBC literature. Though being a strong assumption, it should not be of great concern in our case, provided that we are solely interested in modelling the secular trends dynamics.[1] The initial conditions of the structural trends are distributed according to $X_0 \sim \mathcal{N}(\underline{X_0}, I_q)$. In principle we do not have information about $\underline{X_0}$. However, one can use the information on $\underline{\bar{Y}_0}$ - the initial conditions of the empirical trends[2] - as well as, on the prior coefficients in $\mathcal{V}$. Then, one can retrieve $\underline{X_0}$ by solving the system in eq. (2.2) - provided that the number of structural trends $q = n$, as it is the case in our model.[3] The initial conditions of the cycles are distributed according to $\hat{Y}_0 \sim \mathcal{N}(0_n, I_n)$. This assumption implies that cycles fluctuate symmetrically around a zero mean. Finally, the priors for the remainder model's

[1]Recently, Furlanetto et al. (2021) discuss the existence of hysteresis effects of recessions since the Great Moderation. They find evidence of recessions leaving long-run scars on economic capacity. In principle, our model can accommodate looser assumptions on the covariance structure of shocks, but this goes beyond the scope of this paper and we leave it to future research.

[2]Specifically, we set $\underline{\bar{Y}_0}$ equal to the average of the HP-filter trend growth rate from the pre-sample data.

[3]In section 2.8 of the appendix, we show how to derive priors for the structural trends by by solving the implied system in eq. (2.2).

coefficients are distributed according to:

$$\Sigma_u \sim \mathcal{IW}(\kappa_u, (\kappa_u + n + 1)\underline{\Sigma_u}) \quad (2.5)$$

$$\Sigma_e \sim \mathcal{IW}(\kappa_e, (\kappa_u + n + 1)\underline{\Sigma_e}) \quad (2.6)$$

$$\tilde{\Phi}|\Sigma_e \sim \mathcal{N}(\underline{\tilde{\Phi}}, \Sigma_e \otimes \underline{\Omega})\mathcal{I}(\tilde{\Phi}), \quad (2.7)$$

where $\tilde{\Phi} = vec(\Phi)$ 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 $\kappa$ degrees of freedom and mode $\underline{\Sigma}$. We assume the prior mode of trend shocks $\underline{\Sigma_u}$ to be diagonal and its non zero elements are retrieved in the same fashion as we discussed for the initial conditions of the structural trends $\underline{X_0}$.[4] We impose a rather tight prior around the prior mode, by setting the degrees of freedom $\kappa_u = 100$. Moving to the cycle block, the priors for the lag coefficients are standard Minnesota with overall tightness hyperparameter equal to 0.2 and the own-lag hyperparameters centered around zero, instead of one, since we are dealing with the stationary block. The prior mode of the transitory innovations $\underline{\Sigma_e}$ is assumed to be an identity matrix and rather uninformative, as the degrees of freedom are $\kappa_e = n + 2$.

**Data.** In the baseline specification of the model presented in 2.6.1, the vector of endogenous $Y_t$ includes: (i) real GDP; (ii) real aggregate wages; (iii) employment-to-population ratio; (iv) female-to-male employment ratio; (v) female-to-male wages ratio. All variables enter the model in log-levels.[5] Finally, the model samples 50000 draws from the Gibbs algorithm and retains the last 10000 draws. We use Kalman Filter to estimate the unobserved structural trends and implement Carter and Kohn (1994) simulation smoothing to generate the posterior distributions. The posterior estimates of coefficients in $\mathcal{V}$ are obtained by including a Metropolis-Hastings step within the Gibbs sampler.

**Baseline setup.** Consistent with eq. 2.2, the goal is to decompose the empirical trends into three aggregate macro drivers: (i) technology denoted by $A_t$; (ii) automation $\mathcal{M}_t$; (iii) labor supply $\Psi_t$. Technology and labor supply trends are standard, while the automation trend is meant to capture those slow-moving factors contributing to the secular decline of labor share - see e.g.: Autor and Salomons (2018), Acemoglu and Restrepo (2020), Bergholt et al. (2019). In addition, the long-run behavior of employment and wage gender ratios is mainly driven by two forces: (iv) gender-specific labor demand $a_{f,t}$; (v) gender-specific labor supply $\psi_{f,t}$. These are meant to capture any *positive* and *negative* co-movement between employment and wage gap, respectively. Thus, $X_t$ is a $5 \times 1$ vector, so the number of observables $n$ is equal to the number of structural trends $q$. Given the system of observables $Y_t$, the main challenge is to find a unique solution that

[4] We retrieve the standard deviations of the HP-filter trend growth rates from the pre-sample and solve the system of $n$ equations in $n$ unknowns implied by eq. (2.2) to derive the priors for the structural trends shocks volatilities. See section 2.8.4 of the Technical Appendix for the details.

[5] See the Appendix 2.8 for the description and source of the data.

allows us to reconcile the reduced-form objects in $\bar{Y}_t$ with the structural ones stacked in $X_t$. In fact, this is not a trivial task, as we need at least $\frac{n(n-1)}{2} = 10$ additional restrictions, according to Rothenberg (1971) condition. Furthermore, though extensive evidence documents the long-run impact of main macro trends - such as technology and labor supply -, there exists only scarce evidence so far about the relationship between gender-specific trends and the macroeconomy. Therefore, in order to derive reasonable identifying restrictions, we decide to let economic theory guide us. Consistent with a growing body of literature studying the role of gender-specific shock within structural models (e.g.: Albanesi (2019)), the next section presents a simple neoclassical model with gender which enables us to derive the theory-based identifying restrictions that will inform the identification matrix $\mathcal{V}$ in eq. 2.2 when estimating the SVAR.

## 2.3 A Neo-Classical Model with Gender

The model economy is populated by a unit mass of identical firms, and a unit mass of identical households who own equal shares in the firms. A representative firm chooses labor inputs and capital investments in order to maximize a properly discounted sum of expected lifetime profits, $\mathbb{E}_t \sum_{s=t}^{\infty} \beta^{s-t} \frac{\Lambda_s}{\Lambda_t} \Pi_s$. For each period $t$ we denote the rational expectations operator (conditional on the information currently available) by $\mathbb{E}_t$. $\beta \mathbb{E}_t \frac{\Lambda_s}{\Lambda_t}$ captures households' discounting of the future where $\beta$ is the time discount rate and $\Lambda_t$ represents the shadow value of income. The firm's period profit is equal to

$$\Pi_t = Y_t - W_{f,t} L_{f,t} - W_{m,t} L_{m,t} - P_{I,t} I_t. \tag{2.8}$$

$Y_t$ represents output, $W_{f,t}$ ($W_{m,t}$) represents the real wage rate specific to female (male) labor, and $L_{f,t}$ ($L_{m,t}$) is the quantity of female (male) labor used in production. $I_t$ represents gross investments in physical capital. The relative price of investments is given by $P_{I,t}$. The firm's maximization problem is subject to the production function

$$Y_t = A_t L_t^{\alpha_t} K_{t-1}^{1-\alpha_t}, \tag{2.9}$$

where $K_{t-1}$ stands for physical capital currently in place, and $L_t$ is an aggregation of male and female labor:

$$L_t = \left[ \alpha_l \left( A_{m,t} L_{m,t} \right)^{\frac{\gamma-1}{\gamma}} + (1-\alpha_l) \left( A_{f,t} L_{f,t} \right)^{\frac{\gamma-1}{\gamma}} \right]^{\frac{\gamma}{\gamma-1}} \tag{2.10}$$

$A_t$, $A_{m,t}$ and $A_{f,t}$ are aggregate and gender-specific productivity shocks, respectively, while $\gamma > 1$ governs the degree of substitution between genders when firms produce. Note that we allow for a time varying weight $\alpha_t$ on aggregate labor. One possible interpretation of a decline in $\alpha_t$ is that it follows from labor-displacing automation, see Acemoglu and

Restrepo (2020) and Bergholt et al. (2019). Finally, physical capital dynamics are given by

$$K_t = (1-\delta)\, K_{t-1} + I_t. \tag{2.11}$$

The representative firm's first order conditions with respect to investments, capital, and male and female labor, are summarized below:

$$P_{I,t} = \mathcal{Q}_t \tag{2.12}$$
$$\mathcal{Q}_t = \beta \mathbb{E}_t \frac{\Lambda_{t+1}}{\Lambda_t} \left[ (1-\alpha_{t+1}) \frac{Y_{t+1}}{K_t} + \mathcal{Q}_{t+1} (1-\delta) \right] \tag{2.13}$$
$$W_{m,t} = \alpha_t \alpha_l \frac{Y_t}{L_t} \left( \frac{L_t}{L_{m,t}} \right)^{\frac{1}{\gamma}} A_{m,t}^{\frac{\gamma-1}{\gamma}} \tag{2.14}$$
$$W_{f,t} = \alpha_t (1-\alpha_l) \frac{Y_t}{L_t} \left( \frac{L_t}{L_{f,t}} \right)^{\frac{1}{\gamma}} A_{f,t}^{\frac{\gamma-1}{\gamma}} \tag{2.15}$$

The first optimality condition states that firms invest until the price of investment is equal to $\mathcal{Q}_t$, the shadow value of one more unit of installed capital in the next period. The second optimality condition defines the shadow value of capital: it is the properly discounted sum of next period's marginal product of capital and the continuation value net of depreciation. The two last optimality conditions pin down optimal firm demand for male and female labor, respectively. Everything else equal, the gender-specific labor demand is increasing in aggregate activity, decreasing in the gender-specific wage rate, and, if $\gamma > 1$ ($< 1$), increasing (decreasing) in gender-specific productivity.

The representative household is populated by an equal number of male and female workers. In each period it chooses a plan for consumption and labor supply in order to maximize expected lifetime welfare $\mathbb{E}_t \sum_{s=t}^{\infty} \beta^{s-t} \mathcal{U}_s$, where

$$\mathcal{U}_t = \frac{C_t^{1-\sigma}}{1-\sigma} \exp\left( -\Psi_t^{-1} \frac{(1-\sigma)\, \tilde{L}_t^{1+\varphi}}{1+\varphi} \right) \tag{2.16}$$

represents the period utility function. Aggregate labor dis-utility $\tilde{L}_t$ is increasing in male and female labor:

$$\tilde{L}_t = \left[ \left( \frac{L_{m,t}}{\Psi_{m,t}} \right)^{\frac{1+\lambda}{\lambda}} + \left( \frac{L_{f,t}}{\Psi_{f,t}} \right)^{\frac{1+\lambda}{\lambda}} \right]^{\frac{\lambda}{1+\lambda}} \tag{2.17}$$

$\Psi_{m,t}$ and $\Psi_{f,t}$ are gender-specific labor supply shocks, $\lambda > 0$ governs the household's willingness to substitute female with male labor. The representative household's first order conditions with respect to consumption, bond savings, as well as supply of male and female labor respectively, are summarized below:

$$\Lambda_t = C_t^{-\sigma} \exp\left(-\Psi_t^{-1} \frac{(1-\sigma)\,\tilde{L}_t^{1+\varphi}}{1+\varphi}\right) \tag{2.18}$$

$$\Lambda_t = \beta \mathbb{E}_t \Lambda_{t+1} (1 + r_t) \tag{2.19}$$

$$W_{m,t} = \Psi_t^{-1} C_t \tilde{L}_t^{\varphi - \frac{1}{\lambda}} L_{m,t}^{\frac{1}{\lambda}} \Psi_{m,t}^{-\frac{1+\lambda}{\lambda}} \tag{2.20}$$

$$W_{f,t} = \Psi_t^{-1} C_t \tilde{L}_t^{\varphi - \frac{1}{\lambda}} L_{f,t}^{\frac{1}{\lambda}} \Psi_{f,t}^{-\frac{1+\lambda}{\lambda}} \tag{2.21}$$

The first optimality condition equates the shadow value of income with the marginal utility of consumption. The second optimality condition states the optimal, intertemporal consumption plan. The two last optimality conditions illustrate that, everything else equal, the optimal supply of gender-specific labor is increasing in the gender-specific wage rate, decreasing in aggregate consumption, and decreasing the aggregate and gender-specific labor dis-utility shocks. Finally, an increase in male labor for example, which in turn raises aggregate labor dis-utility $\tilde{L}_t$, implies a reduction (increase) in female labor supply if and only if $\varphi > \lambda^{-1}$ $(< \lambda^{-1})$. Importantly, $\lambda$ governs the household's willingness to substitute work across genders. A sufficiently low value of $\lambda$ implies complementarity of gender-specific labor dis-utility to such an extent that more time spent working for the male causes a decline in the joy of leisure for the female.

In order to characterize gender differences in the labor market, we find it instructive to focus on the female wage gap $w_{f,t} = \frac{W_{f,t}}{W_{m,t}}$, as well as the female employment gap $l_{f,t} = \frac{L_{f,t}}{L_{m,t}}$. This notation allows us to combine the firm's optimality conditions with respect to male and female labor in order to express relative labor demand, which is downward sloping in the $(w_{f,t}, l_{f,t})$-space:

$$l_{f,t} = \left(\frac{1-\alpha_l}{\alpha_l}\right)^{\gamma} w_{f,t}^{-\gamma} a_{f,t}^{\gamma-1} \tag{2.22}$$

The slope coefficient $-\gamma$ determines how responsive demand is to relative wage changes. It follows naturally that shifts in $l_{f,t}$ not associated with changes in $w_{f,t}$ are driven by the "ratio shock" $a_{f,t} = \frac{A_{f,t}}{A_{m,t}}$, which we interpret as a relative demand shifter. In a similar way, we can combine the household's optimality conditions with respect to male and female labor in order to express relative labor supply, which is sloping upwards in the $(w_{f,t}, l_{f,t})$-space. The slope coefficient $\lambda$ determines how responsive supply is to relative wage changes:

$$l_{f,t} = w_{f,t}^{\lambda} \psi_{f,t}^{1+\lambda} \tag{2.23}$$

We interpret the "ratio shock" $\psi_{f,t} = \frac{\Psi_{f,t}}{\Psi_{m,t}}$ as a supply shifter which effectively soaks up all the variation in relative labor supply not associated with movements in the wage gap.

Combining the two previous equations, one arrives at the following analytical solutions for the wage and employment gaps between females and males:

$$w_{f,t} = \left(\frac{1-\alpha_l}{\alpha_l}\right)^{\frac{\gamma}{\gamma+\lambda}} a_{f,t}^{\frac{\gamma-1}{\gamma+\lambda}} \psi_{f,t}^{-\frac{1+\lambda}{\gamma+\lambda}} \tag{2.24}$$

$$l_{f,t} = \left(\frac{1-\alpha_l}{\alpha_l}\right)^{\frac{\gamma\lambda}{\gamma+\lambda}} a_{f,t}^{\frac{(\gamma-1)\lambda}{\gamma+\lambda}} \psi_{f,t}^{\frac{(1+\lambda)\gamma}{\gamma+\lambda}} \tag{2.25}$$

Importantly, the female-specific demand shock $a_{f,t}$ implies co-movement between wage and employment gaps across genders, while the female-specific supply shock $\psi_{f,t}$ implies negative co-movement. Moreover, while macroeconomic shocks may drive the *absolute level* of female wages and employment, only the two "ratio shocks" $a_{f,t}$ and $\psi_{f,t}$ can affect $w_{f,t}$ and $l_{f,t}$, i.e. the *relative* wage and employment levels of female workers. "Macro shocks" such as $A_t$ and $\Psi_t$ play no role here.[6] A corollary statement is that gender-specific labor market variables and their aggregate counterparts display proportional responses to macroeconomic shocks (e.g. $L_{f,t} \propto L_t$). Importantly, these model implications form a key part of our identification scheme in the empirical section, allowing us to disentangle the different structural drivers of $w_{f,t}$ and $l_{f,t}$ in data.

Finally, we also note that the model presented here nests as special cases some recent theoretical contributions in the literature. Fukui et al. (forthcoming), for example, implicitly assume $\gamma = \infty$, i.e. perfect gender substitutability within the firm. By construction this causes the wage gap to be driven solely by gender biased demand shocks, as can be seen from the analytical solution for $w_{f,t}$ given above. Albanesi (2019), in contrast, implicitly assumes that $\lambda = \frac{1}{\phi}$. This knife-edge parametrization effectively makes the supply of female and male labor independent of how much the spouse is working, as emphasized earlier. While our estimation procedure allows for these special cases, a potentially important contribution of this paper is to quantify the degree of gender complementarity –both on the firm and household side– in the labor market. To this end we offer a novel identification approach which allows us to estimate directly the values of gender substitution elasticities $\gamma$ and $\lambda$ using state-of-the-art Bayesian techniques on macroeconomic data.

### 2.3.1 Simulation results and our implied identification scheme

As stressed earlier, we use the model presented above as a laboratory to infer a set of theory-consistent identification assumptions which are imposed on $\mathcal{V}$. The aim is to arrive at a set of restrictions on $\mathcal{V}$ which are robust to all reasonable parametrizations of the theoretical model. To this end we conduct the following simulation exercise: first we

[6] This is true not only in the long run but also within the business cycle. The irrelevance of aggregate macro shocks for gender gaps is a consequence of the constancy of gender substitution elasticities $\gamma$ and $\lambda$, and remains even if we were to introduce business cycle frictions such as nominal price rigidities.

draw a parameter vector $\theta = [\sigma, \varphi, \gamma, \lambda, ...]'$ which includes all parameters of interest in the model. In order to be as agnostic as possible, we draw each parameter independently from a uniform distribution specified further below. Second, conditional on $\theta$, we solve the model numerically and compute the impulse responses of macro and gender variables to each of the structural trend shocks in the model. For our purpose the main interest lies in the long end of the impulse responses, i.e. the long run effects of various trend shocks. For that reason we compute the perfect foresight solution of the model. We repeat the exercise $1,000$ times and save all impulse responses. This leaves us, at each time horizon and conditional on each shock, an entire distribution of structural outcomes for the endogenous variables of interest. The distribution visualises variation in outcomes due to parameter uncertainty.

Regarding the uniform distributions for the structural parameters, we impose bounds that are wide enough so that they span the set of values proposed in existing literature. In particular, for the three "macro" parameters $\sigma$, $\varphi$ and $\alpha$ we choose $\sigma \sim U(1,5)$, $\varphi \sim U(0,4)$ and $\alpha \sim U(0.5, 0.7)$ respectively. Note that common values for the aggregate Frisch elasticity $\varphi^{-1}$, both from microeconomic and macroeconomic research, are well within the bounds used here. Moreover, the bounds on $\alpha$ allow the model to cover a wide range of labor income shares, including all those observed in the postwar US economy. There is substantially less external information available about the key gender parameters of interest, $\gamma$ and $\lambda$. For these we choose the following uniform distributions: $\gamma \sim U(1,10)$ and $\lambda \sim U(0,1)$. Regarding $\gamma$, we note that the median value is $\gamma \approx 5$, which is similar to the value proposed by Albanesi (2019). Fukui et al. (forthcoming) in contrast assume $\gamma = \infty$ in their baseline, but also analyze the case with $\gamma = 5$ as a robustness test. Finally, note that $\varphi = \frac{1}{\lambda}$ for median values of $\varphi$ and $\lambda$. This is exactly the special case considered by both Albanesi (2019) and Fukui et al. (forthcoming). However, our chosen bounds allow us to investigate very different scenarios as well, including those with substantial complementarity (as well as substitutability) across genders when the household decides on labor supply. When solving the model, the initial wage and employment gaps may matter for the long run outcomes of structural shocks. Therefore, we choose to draw initial wage and employment gaps from $w_{f,0} \sim U(0.56, 0.85)$ and $l_{f,0} \sim U(0.44, 0.85)$ respectively. These bounds are chosen so that they cover both the highest and the lowest wage and employment gaps observed in the postwar US economy.

The impulse responses resulting from the simulation exercise are presented in figures 2.2 and 2.3. Recall that the main goal is to disentangle the different structural drivers of the main macro aggregates, as well as, of wage and employment gaps in the long-run. Therefore, the focus will be on the long-run co-movement across of the variables.

Let us start with the impulse responses to aggregate shocks in figure 2.2. First, consistent with the analytical solutions in eq. 2.25, aggregate shocks have no effects whatsoever on both wage and employment gaps at any horizons. Second, conditional on

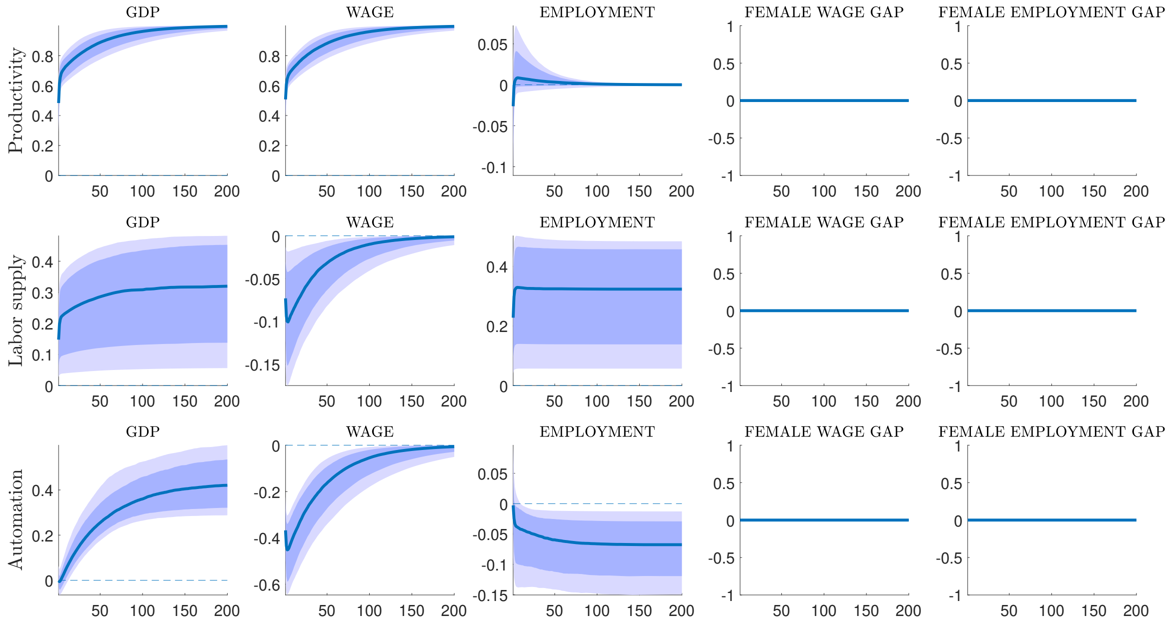


Figure 2.2: Impulse Responses to aggregate macro shocks.

an aggregate productivity shock, GDP and aggregate wages display the same response both in the sign and the magnitude, while aggregate employment is not affected in the long-run. Then, in response to an aggregate labor supply shock, the responses show a positive long-run co-movement between GDP and aggregate employment, but no effects on aggregate wages. Finally, an automation shock generates a negative co-movement between GDP and aggregate employment in the long-run, but no effects on aggregate wages.

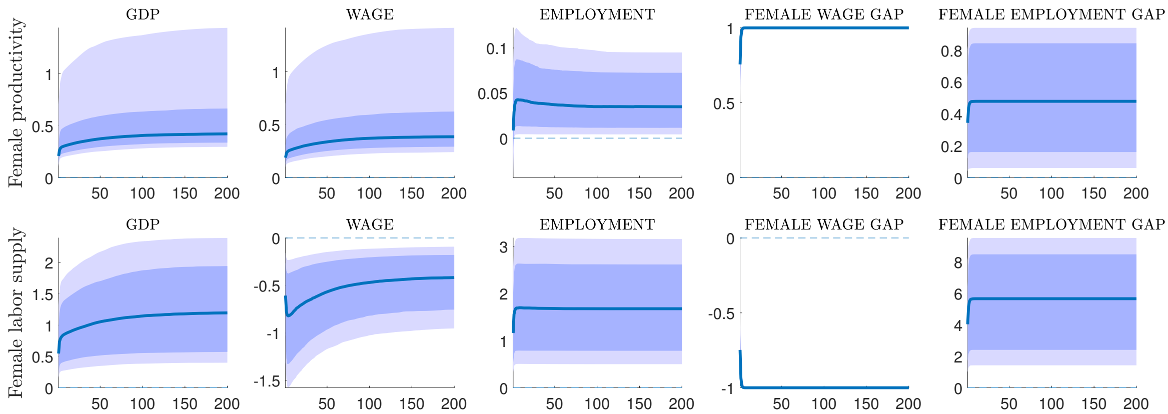


Figure 2.3: Impulse responses to gender-specific shocks.

Let us now move to the gender-specific shocks in figure 2.3. Both shocks are normalised to produce a unit effect on the wage gap at any horizons.[7] Conditional on a gender-specific productivity shock, both the wage gap and the employment gap display a *positive*

[7]This normalisation reveals particularly useful when estimating the empirical model, as it will enable us to interpret the long-run effects of these two shocks on employment gap in terms of the elasticities $\gamma$ and $\lambda$.

response in the long-run. In addition, all aggregate macro variables also move in tandem. Conversely, in response to a gender-specific labor supply shock, wage gap and employment gap now exhibit a *negative* long-run co-movement. In addition, GDP and employment are affected positively, while aggregate wages display a negative response at any horizon.

Overall, these results provide us with a set of theory-consistent identifying restrictions that will inform matrix $\mathcal{V}$ in the empirical model. Specifically, these restrictions enable us to identify: (i) three aggregate trends which characterise the long-run behavior of GDP, employment and wages, but have no effects whatsoever on gender differentials; (ii) a gender-specific labor demand trend that is responsible for the positive co-movement between wage gap and employment gap; (iii) a gender-specific labor supply yielding to an opposite co-movement between wage gap and employment gap.

## 2.4 Theory-Based Identifying Restrictions

The analytical solutions in section 2.3 provide us with a set of identifying restrictions which will be employed in the VAR to derive a unique mapping between the *empirical* and the *structural* trends. Equation 2.26 describes the most important relation in the manuscript as it draws a direct link between theory and empiric and it embeds all the information needed to unfold the reduced-form trends into the underlying structural drivers. On the left hand-side the empirical trends are stacked into the vector $\bar{Y}_t$, while on the right-hand side the unobserved structural trends are enclosed in the vector $X_t$. Finally, matrix $\mathcal{V}$ summarises all the identifying restrictions needed to reconcile $X_t$ with $\bar{Y}_t$.[8]

$$\underbrace{\begin{bmatrix} G\bar{D}P_t \\ \bar{W}_t \\ \bar{E}_t \\ \bar{E}_{f-m,t} \\ \bar{W}_{f-m,t} \end{bmatrix}}_{\bar{Y}_t} = \underbrace{\begin{bmatrix} 1 & 1 & \nu_{13} & \nu_{14} & \nu_{15} \\ 1 & 0 & 0 & \nu_{24} & \nu_{25} \\ 0 & \nu_{32} & \nu_{33} & \nu_{34} & \nu_{35} \\ 0 & 0 & 0 & \lambda & \gamma \\ 0 & 0 & 0 & 1 & -1 \end{bmatrix}}_{\mathcal{V}} \underbrace{\begin{bmatrix} A_t \\ \mathcal{M}_t \\ \Psi_t \\ a_{f,t} \\ \psi_{f,t} \end{bmatrix}}_{X_t} \tag{2.26}$$

The entries of $\mathcal{V}$ are interpreted as the feedbacks of the structural trends to the empirical ones. For instance, the first column embeds the feedbacks of the technology trend $A_t$ to each empirical trend in $\bar{Y}_t$, the second column contains the feedbacks of the automation trend $\mathcal{M}_t$, and so on. One clear advantage of eq. 2.26 is that it reconciles reduced-form and structural trends by means of a *linear* relationship, which can be easily accommodated by any standard linear state-space representation of the data. Therefore, the "only" challenge is to find an appropriate parametrization of $\mathcal{V}$, such that structural trends are uniquely identified. At this point, there are two options: (i) trust the model solutions dogmatically and calibrate $\mathcal{V}$ accordingly; (ii) allowing for some uncertainty around the

[8] In section 2.2, $\mathcal{V}$ will be the matrix of loadings of a linear VAR with common trends.

implied model solutions and employ them as guidance in the formation of prior beliefs. Despite calibration might sound appealing, as it reduces model uncertainty substantially, it would make the empirical exercise completely useless. The reason for this lies upon the fact that the model should not be viewed as an exhaustive representation of the economy. In fact, the main advantage is its relatively stylised formulation, which allows us to come up with some closed-form solutions. In a Bayesian spirit, the model solutions should be deemed as a useful starting point in the formulation of the prior restrictions for the relevant parameters in $\mathcal{V}$. Consistently, the theoretical model is meant to be complementary - not exclusive nor exhaustive - to the structural VAR in the next sections.

Concretely, this is how we proceed. Whenever the results from the theoretical model are associated with well-documented facts in the literature, the restrictions will incorporate them with a bell-shaped type of prior restriction. Alternatively, for those elements in $\mathcal{V}$ for which existing literature does not provide any assistance, the restrictions will be informed both in the *sign* and in the *magnitude* by the simulated impulse responses in figures 2.2 and 2.3. Hence, whenever we are unable to form sensible prior restrictions, the impulse responses will serve as guideline. This will be the case for the feedbacks to macro of gender-specific trends, for instance. Summarising, identification is accomplished by imposing *support* restrictions on these elasticities. This has the twofold benefit of jointly: (i) *bounding* the magnitude to plausible values, by assigning little or no probability mass to unsuitable draws; (ii) determining the *sign* of the long-run co-movement in the data. A similar, though not identical, reasoning has also been employed in the oil market literature to correct one undesirable property of agnostic sign restrictions in VARs. Specifically, Kilian and Murphy (2012a) and Baumeister and Hamilton (2019) showed that sign restrictions alone are not enough to identify the structural drivers of the global oil market. This is due to the fact that sign restrictions deliver *set identification* by definition. If, on the one hand, this has the advantage of being a more agnostic approach, on the other hand, it has the potential drawback of retaining implausible draws in terms of magnitude. In the oil market example, this resulted in unrealistically large short-run elasticity of oil demand with respect to an adverse shock in the global supply of oil. To correct this, the authors complemented the signs with boundary restrictions on the impact elasticity of the oil demand curve. Though sharing some similarities, our specific case is different in two ways: (i) our restrictions are meant to discipline the long-run behavior of the data; (ii) most importantly, we move away from set identification, as the estimation of all the long-run elasticities in $\mathcal{V}$ allows us to exactly pin down the magnitude effect of all structural trends, thereby achieving *point identification.* This is also the very precise reason why we do not enter into the business of estimating a VAR in first differences and identify the trend shocks by means of signs and zero long-run restrictions. This approach belongs to the universe of set identification schemes, as well, and, contrary to our identification strategy, does not allow to neither directly pin down the long-run elasticities of trend shocks nor explicitly quantify the relative importance of the underlying structural drivers for secular macro trends.

Moving forward, to better understand the intuition behind our identification strategy, let us now walk the reader through the rows and columns of $\mathcal{V}$ in eq. 2.26. The first column refers to the effects of the technology trend $A_t$. Provided that output and aggregate wages are affected in the same way by $A_t$, the first two elements are normalised to unity. In this context, normalisation has the advantage of favouring a parsimonious implementation of the identifying restrictions, thereby reducing the space of the estimated coefficients and limiting model uncertainty. Consistent with the assumptions of the theoretical model, the zero effect on employment derives from balance growth path.[9] The second column defines the impact of the automation trend $\mathcal{M}_t$. By normalising the effect on GDP, it is possible to solely estimate the feedback to aggregate employment $\nu_{32}$ and interpret it relative to the effect on GDP. Consistent with the literature studying the macroeconomic implications of automated tasks performed by robots in place of human workforce - see e.g.: Acemoglu and Restrepo (2020), Autor and Salomons (2018) and Bergholt et al. (2019) -, an automation shock yields a negative co-movement between output and aggregate employment in the long-run. Hence, provided that the effect on GDP is normalised to 1, the support of the feedback to employment is negative. Labor supply $\Psi_t$ - third column - yields the same long-run effects on GDP and aggregate employment, therefore we assume $\nu_{13} = \nu_{33}$. In addition, since it is well-known from neo-classical theory that the elasticity of GDP and aggregate employment with respect to labor supply is equal to $-\frac{1}{1+\phi}$, where $\phi$ is the inverse Frisch elasticity of the households labor disutility in the theoretical model, then we assume $\nu_{13} = \nu_{33} = -\frac{1}{1+\phi}$.

The last two columns of $\mathcal{V}$ capture the long-run effects of gender-specific labor market trends, namely the gender-specific labor demand $a_f$ and labor supply $\psi_f$. The former captures any *positive* co-movement between employment gap and wage gap, while the latter is responsible for any *negative* co-movement between the two gender ratios. Besides, the two gender-specific trends are normalised to have unit effect on the female-to-male wage gap.[10] This implies that the feedbacks to macro are expressed in relative terms to the effects on the wage gap. Furthermore, this normalisation reveals to be particularly convenient not only for the same reason mentioned above, but also for two additional reasons. First, scaling to unity the elasticities of the female-to-male wage gap, allows us to express the elasticities of female-to-male employment gap in terms of the parameters $\lambda$ and $\gamma$, which govern the degree of substitutability across gender in the household and firm sectors in the theoretical model, respectively. This not only enables us to estimate two structural parameters of the model directly in the data, but it also has the merit

[9] One can, in principle, relax this assumption and estimate the long-run effects of technology on employment, consistent with the recent facts documented by Boppart and Krusell (2020), among others.
[10] The gender-specific shocks are normalised to have unit long-run effect on the wage gap also in the theoretical model.

of greatly simplifying the task of choosing the prior distribution for the long-run effects of $a_f$ and $\psi_f$ on female-to-male employment gap, as there exists an empirical literature studying the degree of substitution between females and males in the production sector, as well, as within the household - see e.g. Olivetti and Petrongolo (2014) for a discussion. Second, by fixing the elasticity of female-to-male wage gap to unity, we also avoid to run into weak identification due the additional uncertainty arising from the estimation of the volatilities of structural trends shocks in the VAR. More precisely, if we were not normalising to unity, outcomes of the empirical model from relatively *large* shock volatility and *small* elasticity would have been observationally equivalent to outcomes from relatively *small* shock volatility and *large* elasticity.

The design of appropriate prior restrictions turns out to be a delicate task when it comes to the feedbacks of gender-specific shocks to aggregate macro trends - i.e.: the $3 \times 2$ submatrix of coefficients in the top-right corner. This is due to the fact that these coefficients are non-linear functions of the deep parameters of the model and there is no existing literature guiding us. For this reason, we decide to limit our confidence about the prior beliefs inherited from the theoretical model by distributing the priors *uniformly* along the entire support of the admissible long-run impulse responses implied by the model and let the data speak. To see this, take for example the response of GDP to a gender-specific labor supply shock as displayed in the bottom-left panel of figure 2.3. In the long-run, the set of admissible responses spans approximately from 0.5 to 2, with the median point equal to 1. Thus, according to the model, a gender-specific labor supply always generates a statistically different from 0 effect on GDP in the long-run. However, we prefer to be more agnostic and incorporate in our prior the possibility that gender-specific labor supply might have no effect whatsoever on GDP in the long-run and, in addition, assign the same probability as the median effect of 1. Therefore, in this case the prior restriction for the feedback of the gender-specific labor supply to GDP trend will be a uniform over the interval [0,2]. Yet, the set of admissible long-run responses of aggregate employment to a gender-specific labor supply ranges between 0.5 to 3 (see figure 2.3). Despite the theoretical model predicts once again significant aggregate effects of gender-specific shocks, we opt to be more agnostic and allow the prior to inherit the possibility that the gender-specific labor supply might have *no* long-run effects on aggregate employment. Hence, we assume the long-run effect to be uniformly distributed along the interval [0,3]. We proceed accordingly also for the other feedbacks to macro.

Finally, the $2 \times 3$ block of zeros in the bottom-left corner nails down the long-run *neutrality* of aggregate macro trends - i.e.: technology, labor supply, automation - on gender differentials. There is no reason to believe that a shock to fundamentals would favour one gender at the detriment of the other, at least in the long-run. Composition effects over the different phases of the business cycle might in principle temporarily exacerbate or relax gender gaps in the labor market, but these effects should eventually fade away in the long-run.

## 2.5 Priors on $\mathcal{V}$

Matrix $\mathcal{V}$ can be conveniently expressed as $\mathcal{V}(\nu)$ function of $\nu$, the vector free parameters. Precisely, $\nu = [\nu_{13}\ \nu_{14}\ \nu_{15}\ \nu_{24}\ \nu_{24}\ \nu_{32}\ \nu_{33}\ \nu_{34}\ \nu_{35}\ \lambda\ \gamma]'$ is the $10 \times 1$ vector stacking the long-run elasticities onto which we place the identifying restrictions discussed in section 2.4. Therefore, the prior of $\nu$ is denoted by $p(\nu)$ and is the product of independent Gamma, Gaussian or Uniform distributions for each element in the vector $\nu$.

Let us start first with the priors of the labor supply feedbacks to GDP and employment. We know from theory that the long-run effect of an increase of labor force participation is equal to $-\frac{1}{1+\phi}$, where $\phi$ is the inverse Frisch elasticity. Consistent with the empirical literature on the Frisch elasticity, the prior of $\phi$ is centered around 4. Notice, however, that the prior is not formed directly on the parameter $\phi$, but on the implied long-run elasticity $\frac{1}{1+\phi}$. Therefore, a prior mean $\underline{\phi} = 4$ implies $\underline{\nu_{13}} = \underline{\nu_{33}} = \frac{1}{1+\phi} = 0.2$. The little disagreement on the value of Frisch elasticity is inherited in the shape of the prior density, which belongs to the Gamma family and assigns very little probability mass on values above 0.4, implying tiny chances on values of $\phi < 1.5$. Moving to the effect of a permanent shock favouring automated labor $\mathcal{M}_t$, this has been normalised to have unit effect on GDP, implying that the effect on aggregate employment - defined by $\nu_{32}$ - is expressed relative to the effect on GDP. Thus, consistent with the impulse responses from the theoretical model and the state-of-the-art literature on the macroeconomic consequences of automation, we assume a Gamma prior density narrowly centered around a mode of 0.25 and virtually no probability of observing an impact below 0.1 and above 0.75 on employment, relative to the long-run effect on GDP. Also, notice that the Gamma distribution is defined only on a positive support, therefore the negative effect on employment is obtained by simply placing a minus in front of each draw from the prior density for each Gibbs sampling iteration.

Recall from the discussion in section 2.4 that there is *no* established agreement on the macroeconomic implications of gender-specific shocks. As a matter of fact, the goal of this paper is precisely to bridge this gap in the literature. Therefore, an appropriate formulation of the priors of the gender-specific labor demand and supply feedbacks to GDP, wages and employment should remain agnostic on their relative importance, but at the same time inherit some features of the theoretical model which help to grease the wheels of the identification procedure. Consistently, all the gender-specific prior feedbacks to macro are assumed to be uniform over a range of possible values, as predicted by the simulated impulse responses in the theoretical model. To see this, let us start with the gender-specific labor demand $a_f$ feedback to macro. Figure 2.3 suggests that output and wages are positively affected in the same way in the long-run. Besides, according to the shaded areas, the distribution of responses is right skewed - i.e.: asymmetric - and always statistically different from zero. Nonetheless, we prefer to remain more agnostic and assume the priors to be symmetrically and uniformly distributed between zero and one and let, therefore, the data speak in favour of no effects whatsoever on output and wages if this is, indeed, the case. The simulated response of aggregate employment is positive and mildly statistically significant in the long-run. Consistent with the approach so far, we place a loose prior centered at zero and ranging over the interval [-0.5,0.5]. Let us now move to the gender-specific labor supply $\psi_f$ feedbacks to macro. In this case, the impulse responses in figure 2.3 displays a negative co-movement between output and wages at all horizons. Hence, the prior densities will span over [0,2] and [-2,0] intervals for GDP and wages, respectively. Regarding the long-run impact on employment, a uniform prior centered at 1.5 and ranging between 0 and 3 is assumed.

The remaining free parameters to restrict are $\lambda$ and $\gamma$. In matrix $\mathcal{V}$, these parameters represent the feedbacks of gender-specific labor demand and labor supply to the female employment gap, respectively. In contrast to the priors for the feedbacks to macro, we depart from the uniform assumption and rely on bell-shaped type of priors. The reason lies upon the fact that prior knowledge from empirical literature on $\lambda$ and $\gamma$ - elasticities of substitution in the household and production sectors - allows us to narrow our prior beliefs. Weinberg (2000), for example, estimates the elasticity $\gamma$ to be equal to 2.4 in the US, while Acemoglu et al. (2004) report a slightly higher value of 3.[11] Consistently, the density prior is assumed to be Gamma centered around 3 - assigning little probability on draws above five and below one. The choice of downplaying values below one is also coherent with the theoretical assumption and discussion in Section 2.3 about $\lambda < \gamma$ - i.e.: it is harder to substitute female with male labor within the household than in the firm sector.[12]

The macro literature studying the elasticity of labour supply is much less abundant. To the best of our knowledge, there is only one study trying to shed light on the aggregate labor supply elasticity across gender, that is, Keane and Rogerson (2012) who show that relatively small (micro) labor supply elasticities can be reconciled with aggregate elasticities ranging between 1 and 2. In a review of the micro literature, in fact, Blundell and Macurdy (1999) conclude that $\lambda$ must falls in a range of values below one. Accordingly, the prior distribution is assumed to be a right-skewed Exponential and rather conserva-

[11] Johnson and Keane (2013) obtain estimates of $\gamma$ spanning between 1.85 and 2.2 in a dynamic general equilibrium model estimated to fit US data. Furthermore, cross-country empirical evidence exhibits very similar magnitudes to those reported for the US. Hamermesh (1993), for instance, reports values of approximately 2.3 and 2 for Australia and UK, respectively.

[12] In a robustness check we aim at testing the assumption of $\gamma = \infty$ as in Fukui et al. (forthcoming). To this end, we replace the Gamma prior with a $N(20, 5)$, with the idea of letting $\gamma$ to navigate across a wide interval of values. The discussion on the implications of this robustness exercise is left to the section 2.6 on the results.

tive, as it condenses most of the probability mass on values smaller than one.[13] This not only allows us to assume $\lambda < \gamma$, but also nest the specific case of Albanesi (2019), in which $\lambda = \frac{1}{\phi}$, provided that our prior mean of $\phi$ is equal to 4. Nonetheless, data is allowed to speak in favour of values above 1, provided that the support of the Exponential is $[0, \infty)$. Finally, all prior densities are summarised in table 2.1.

| Mnemonic | Feedback | Density |
|---|---|---|
| $\nu_{13}$ | $\Psi \rightarrow G\bar{D}P$ | $\Gamma(6, .03)$ |
| $\nu_{14}$ | $a_f \rightarrow G\bar{D}P$ | $\mathcal{U}(0, 1)$ |
| $\nu_{15}$ | $\psi_f \rightarrow G\bar{D}P$ | $\mathcal{U}(0, 2)$ |
| $\nu_{24}$ | $a_f \rightarrow \bar{W}$ | $\mathcal{U}(0, 1)$ |
| $\nu_{25}$ | $\psi_f \rightarrow \bar{W}$ | $\mathcal{U}(-2, 0)$ |
| $\nu_{32}$ | $\mathcal{M} \rightarrow \bar{E}$ | $\Gamma(6, .05)$ |
| $\nu_{34}$ | $a_f \rightarrow \bar{E}$ | $\mathcal{U}(-.5, .5)$ |
| $\nu_{35}$ | $\psi_f \rightarrow \bar{E}$ | $\mathcal{U}(0, 3)$ |
| $\lambda$ | $a_f \rightarrow \bar{E}_{f-m,t}$ | $Exp(2)$ |
| $\gamma$ | $\psi_f \rightarrow \bar{E}_{f-m,t}$ | $\Gamma(6, .5)$ |

Table 2.1: The second column defines what each coefficient stands for. E.g.: $\Psi \rightarrow G\bar{D}P$ refers to the effect of labor supply trend on GDP trend.

## 2.6 Results

The presentation of results is divided into two subsections. Subsection 2.6.1 shows the results from the baseline specification. Then, in order to account for the potential role that education (skills) might have played in driving the secular trends, we disaggregate female employment and wage gaps by skills and estimate two alternative specifications. Accordingly, in subsection 2.6.2, we present the results from two alternative models in which employment and wage gaps are expressed in terms of differentials between (i) *skilled* - i.e.: highly educated - workers; (ii) *unskilled* workers

### 2.6.1 Baseline Model

The baseline specification is based on the same information set as described in the paragraph on data in Section 2.2. Therefore, the vector of observables includes three aggregate macro variables: output, wages and employment; and the two gender gaps -

[13]In order to use consistent priors for both $\lambda$ and $\gamma$, we perform a robustness check and replace the Exponential density prior for $\lambda$ with a Gamma distribution. Since the Gamma distribution is defined only on the support $(0, \infty)$, while the Exponential is supported over the set $[0, \infty)$, the robustness check excludes draws being exactly equal to zero. This different parametrization does not affect the results.

i.e.: female-to-male employment and wage ratios.

**Coefficients in $\mathcal{V}$.** Figure 2.4 displays the posterior distributions of the coefficients defining the effects of the aggregate labor supply and automation trends on GDP and employment. The left-hand panel shows the posterior of aggregate supply feedback to GDP and employment with a posterior median of approximately 0.18. Recall that this number refers to the coefficients $\nu_{13} = \nu_{33} = \frac{1}{1+\phi}$. This implies that the posterior median estimate of the Inverse Frisch elasticity $\phi$ is approximately equal to 4.5. The posterior distribution of the coefficient of automation on employment is close to its prior, though the mass concentrates even more around a median of 0.35.

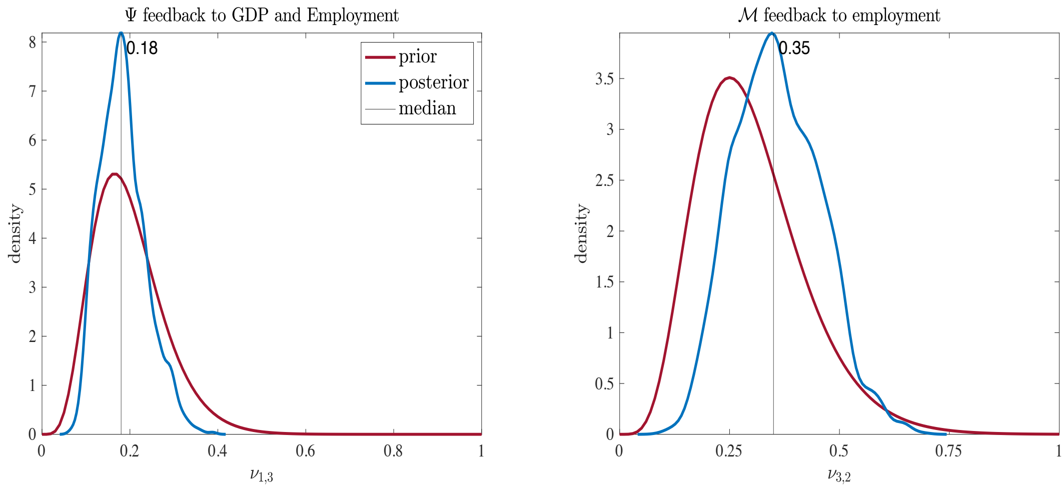


Figure 2.4: Prior distributions (red) and posterior distributions (blue) of the aggregate labor supply and automation feedbacks to GDP and employment.

Figure 2.5 displays the posterior distributions of the gender-specific labor demand and labor supply feedbacks to the aggregate macro variables. Recall from the discussion in section 2.5 that the priors (in red) of all the gender-specific shocks' feedbacks to macro aggregate variables span *uniformly* over the support defined in accordance to the set of admissible (shaded areas of) impulse responses in the theoretical model, as displayed in figure 2.3.

Let us now comment figure 2.5 column by column, so to compare the implication of the two gender-specific shocks on each aggregate macro variable. First, the left-hand side column refers to the long-run elasticities of GDP trend. On the one hand, the posterior estimate of the elasticity to the gender-specific labor demand $a_f$ - top left panel - concentrates most of the mass towards values ranging between 0.75 and 1. On the other hand, the elasticity to the gender-specific labor supply $\psi_f$ displays a posterior distribution that shifts toward values close to zero with most of the mass concentrated around values below 0.5.

Second, the long-run elasticity of aggregate wages to gender-specific labor demand - top-middle panel - displays a posterior median of approximately 0.6 and it excludes zero-effects on wages at any significance level. Conversely, the posterior density in the bottom-middle panel implies essentially no long-run effects of gender-specific labor supply on aggregate wages.

Finally, the feedbacks to aggregate employment are showed in the third column. The posterior of the elasticity with respect to the gender-specific labor demand shifts to the right and does not include zero value any significance level. Most of the mass is condensed in the interval 0.25 to 0.5, as a matter of fact. The posterior density of the gender-specific labor supply gravitates to the interval between 0 and 1, with a median close to 0.5.

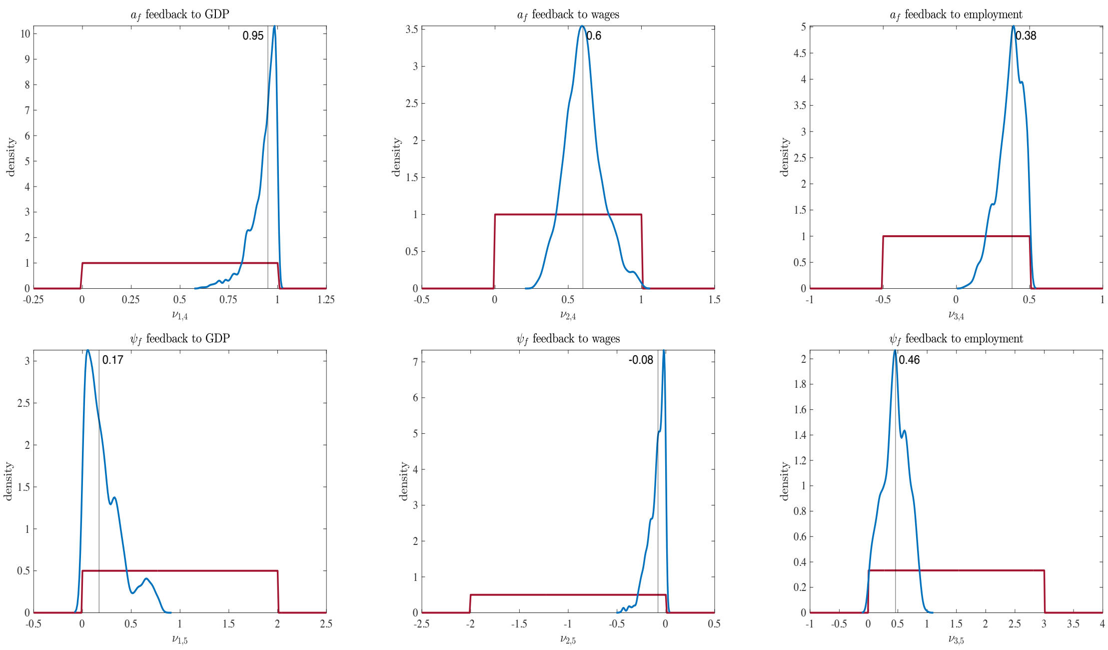


Figure 2.5: Prior distributions (red) and posterior distributions (blue) of the gender-specific labor demand and supply - $a_f$ and $\psi_f$, respectively - feedbacks to macro variables.

We conclude the discussion on the coefficients of matrix $\mathcal{V}$ with figure 2.6, which visualizes the posterior distributions of the elasticities $\lambda$ and $\gamma$, the feedbacks of gender-specific labor demand and supply shocks to female employment gap. Recall that these two parameters have a direct link with theory, as they describe the degree of female-male substitutability in the household and firm sectors of our theoretical model. In addition, it is important to stress once more that, differently from the other feedbacks to macro, the prior densities of $\lambda$ and $\gamma$ are inherited by the literature studying the labor demand and supply elasticities between women and men. This is indeed the reason why the priors (in red) are not distributed uniformly in the probability space. Let us now, discuss the posterior estimates in figure 2.6. The first result is displayed in the left-hand side panel. Despite a rather conservative on $\lambda < 1$, the data clearly speak in favour of the Keane and Rogerson (2012) argument, according to which labor supply elasticity in the aggregate should fall within the interval of values ranging between 1 and 2. That is, our posterior density exhibits a median around 1.5, with the probability mass mostly concentrated within the [1,2] interval. Thus, differently from the micro literature on the labor supply elasticity, our model shows that data assigns almost no probability to $\lambda < 1$, in line with the facts documented by Keane and Rogerson (2012). Moving to $\gamma$, the prior density is rather loose as it tries to inherit the prior knowledge from the macro literature on the gender-specific labor demand elasticity of substitution. Consistent with this empirical literature, the posterior density narrows around values between 2 and 3, with a median estimate (2.4) essentially identical to the one estimated by Weinberg (2000), but also not significantly different from Acemoglu et al. (2004), who estimated $\gamma$ to be close to 3.

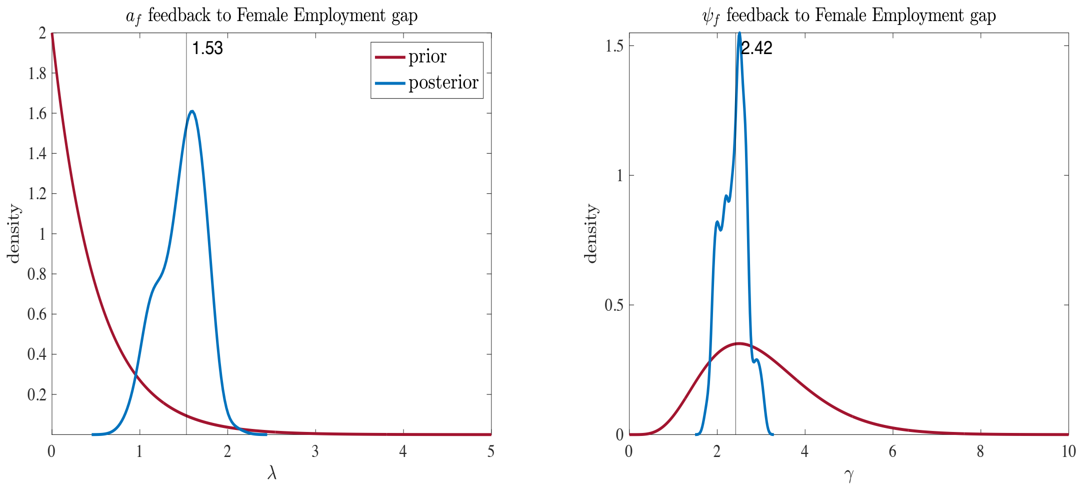


Figure 2.6: Prior distributions (red) and posterior distributions (blue) of the coefficients $\lambda$ and $\gamma$.

Overall, from the estimation of the long-run elasticities embedded in $\mathcal{V}$ some interesting facts emerge. First, data speak in favour of a *non-negligible* long-run feedback of gender-specific shocks to aggregate macro variables. Second, on the one hand, the posterior densities of the labor supply feedbacks tend to concentrate around values close to 0 for both GDP, wages and - to a lesser extent - employment. On the other hand, the posterior estimates of the labor demand feedbacks shift away from values close to the 0-neighborhood. Accordingly, we expect gender-specific labor demand to play a relatively more important role than labor supply. Third, the posterior estimate of $\lambda$ is consistent with the facts document by Keane and Rogerson (2012), implying - differently from the micro literature - values of $\lambda$ above 1 in the aggregate. As concerns $\gamma$, the posterior estimate is perfectly comparable with the results documented in the literature, as it ranges between 2 and 3.

**Trends Decomposition**. Let us now discuss the contribution of the different structural trends to each empirical trend in the model. Figure 2.7 decomposes the empirical trends of female employment gap and wage gap into gender-specific labor demand (green) and supply (light blue). These are the only shocks affecting gender differentials in the long-run. This is a restriction coming from the theoretical model and summarised in eq.

2.26 in section 2.4. Three facts worth noting. First, in the period 1960-1980, both positive labor demand and supply forces are in place. This explains the steep convergence of employment gap, while the wage gap tends to stagnate over this period. Second, Starting from roughly 1980 until mid '90s, the dominant structural force driving the gender convergence in both employment and wage gap is labor demand. Third, since mid '90s, the gender convergence essentially stops and both employment and wage gap exhibit a clear plateau lasting until today, implying a slow-down in the demand and supply forces favouring the gender convergence in the labor market.

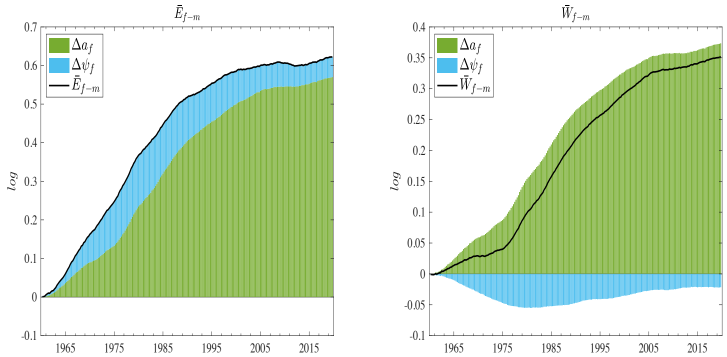


Figure 2.7: Cumulative contribution of each *structural* trend to the *empirical* trend.

The aim of this paper is also to quantify the feedback effect of these gender-specific secular trends to the macroeconomy. Figure 2.8 presents an anatomy of the structural secular drivers for the US main macro aggregates. The results can be summarised by five main facts. First, the neutral technology trend $A$ (blue area) confirms to be the main source of long-run co-movement between GDP and Wages. Second, perhaps unexpectedly, aggregate labor supply $\Psi$ (purple) does never play a significant role neither for GDP nor for employment. Third, the trend in labor-displacing automation $\mathcal{M}$ (yellow) contributes significantly to GDP and employment, especially starting from approximately 1985. In particular, together with the slow-down of gender-specific trends, it is the main responsible of the negative trend experienced by aggregate employment from late '90s. Four, consistent with the posterior estimates of the long-run elasticities in $\mathcal{V}$, gender-specific trends are important for the macroeconomy. This can be clearly seen by the large green area driving GDP, wage and employment trends upward. Five, the gender-specific trends (green and light blue in right-hand side panel) are the sole driver of the positive trend in the US aggregate employment in the post-war period. Had it not been for the gender-specific trends, aggregate employment in the US would have stagnated.

One fine feature of our empirical framework is that it accommodates for the possibility of performing a growth accounting exercise. This enables us to deliver a more comprehensive discussion on the macroeconomic effects of the gender-specific shocks, as we complement the results in figure 2.7 and 2.8 with the growth accounting tables 2.2 and 2.3. Table 2.2 indeed reports the key result from the baseline specification. In the period

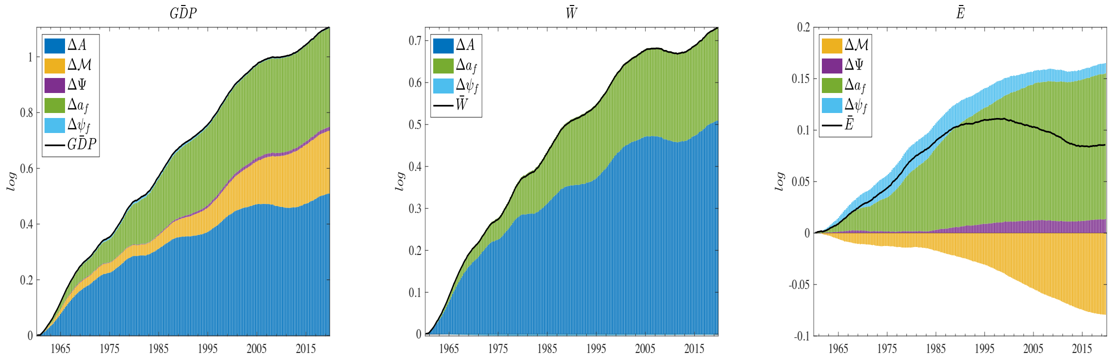


Figure 2.8: Cumulative contribution of each *structural* trend to the *empirical* trend.

1970-1990, gender-specific shocks are responsible on average of approximately 1pp of the annualised GDP trend growth rate, which is about 50% of total growth. Furthermore, the significant contribution of gender-specific shocks is entirely attributable to labor demand factors, with labor supply playing no role. Yet, the 1pp loss in trend growth in the last twenty years is mainly due to the break in the gender convergence. As a matter of fact, the contribution of aggregate trends remained stable around 1pp over the last 50 years, while gender-specific trends plateaued precisely when GDP trend growth started its decline.

The second key results comes from table 2.3 decomposing the drivers of employment trend growth rate over time. In this case, at least for the first 20 years of the sample - i.e.: 1960-1979 -, both gender-specific labor demand and supply positively contribute to employment. Starting from the '80s, however, labor supply fades away and the main contribution comes from labor demand. Aggregate shocks play almost no role until early '90s, when the automation trend picks-up and together with the break in the gender convergence drag employment trend growth rate downward.

| $\Delta G\bar{D}P_t$ | | | | |
|---|---|---|---|---|
| | Aggregate shocks | $a_f$ | $\psi_f$ | avg. ann. growth rate |
| 1960-1969 | 2.09 | 0.56 | 0.05 | 2.7 |
| 1970-1979 | 1.23 | 0.88 | 0.04 | 2.15 |
| 1980-1989 | 1.00 | 1.06 | -0.01 | 2.05 |
| 1990-1999 | 1.38 | 0.60 | -0.02 | 1.96 |
| 2000-2009 | 0.91 | 0.29 | -0.02 | 1.18 |
| 2010-2019 | 0.94 | 0.15 | -0.01 | 1.08 |

Table 2.2: Relative contribution of aggregate and gender-specific trend shocks to the annualised trend growth rate of GDP. Average of each corresponding decade.

| $\Delta \bar{E}_t$ | | | | |
|---|---|---|---|---|
| | Aggregate shocks | $a_f$ | $\psi_f$ | avg. ann. growth rate |
| 1960-1969 | -0.08 | 0.22 | 0.14 | 0.28 |
| 1970-1979 | -0.04 | 0.35 | 0.12 | 0.43 |
| 1980-1989 | -0.05 | 0.42 | -0.03 | 0.34 |
| 1990-1999 | -0.13 | 0.24 | -0.05 | 0.06 |
| 2000-2009 | -0.22 | 0.12 | -0.05 | -0.15 |
| 2010-2019 | -0.13 | 0.06 | -0.02 | -0.09 |

Table 2.3: Relative contribution of aggregate and gender-specific trend shocks to the annualised trend growth rate of employment. Average of each corresponding decade.

Overall, there are three main takeaways from the baseline specification. First, gender-specific trends are important for the macroeconomy. Second, according to our model, the slow-down in US economic growth in the last 25 years can be mainly attributed to the stagnating gender-convergence in the labor market. Last but not least, the Gender Revolution is mainly - if not entirely - a matter of labor demand factors.

However, in the light of the massive increase in female labor force participation over the last 50 years, this last statement is a hard bite to digest and arguing that gender-specific labor supply shocks do not play any significant role would probably be a hasty conclusion. In the interest of parsimony our baseline model might have potentially omitted some relevant features in the data, that once taken into account may change the overall picture. Data over the second half of the last century, for instance, not only document a massive entrance of women in the labor force, but especially an ever-increasingly number of highly educated (skilled) women entering the labor market. Now, suppose that the increase in women labor force participation is asymmetric, in the sense, that it is mainly induced by women who accessed higher levels of education. Then, this "asymmetric" shock may trigger a composition effect eventually leading to an increase in the average wage of women, as the share of highly educated workers in the women labor force has increased, ceteris paribus. Since the baseline specification does not control for the *skill dimension*, it would commingle this gender-specific *skilled* labor supply shock with labor demand factors, leading us to the wrong conclusion that labor demand factors are the *only* drivers of the Gender Revolution. In the following subsection, as a robustness check, we investigate the implications of our empirical model when controlling for the *skill dimension* of the labor force. Concretely, based on the CPS dataset used by Dolado et al. (2021), we estimate two alternative specifications to the baseline, namely (i) a model with gender differential for *skilled* workers and (ii) a model with gender differentials for *unskilled* workers. Then, we set them side by side with our baseline specification and draw the conclusions.

### 2.6.2 The Role of Skills in the Gender Revolution

With respect to the baseline model, these alternative specifications differ only in the two gender variables embedded in the vector of endogenous $Y_t$. More precisely, the *aggregate* female employment and wage gaps are now replaced by the female employment and wage gaps of (i) *skilled* and (ii) *unskilled* workers. All the prior assumptions and the estimation steps remain the same as the baseline. In particular, we find that the uniform priors for the feedbacks of the gender-specific shocks to macro variables are reasonable also in the context of these alternative specifications. In addition, regarding the prior densities for $\lambda$ and $\gamma$, there exists very scarce evidence documenting how these elasticities change when focusing on skilled/unskilled workers only. To the best of our knowledge, Acemoglu et al. (2004) is the only study documenting some heterogeneity between skilled and unskilled workers about the labor demand elasticity $\gamma$, though not about labor supply elasticity. According to the authors, the elasticity of substitution between skilled men and women should be higher - ranging between 4 and 10 - than in the case of unskilled men and women - spanning between 2.5 and 4 -. It is important to stress, nonetheless, that this study is based on data referring to the period 1940-60, which makes it unsuitable for the formulation of our priors, given that the degree of substitutability back then was likely to be very different from subsequent decades. In their multi-sector equilibrium model, Olivetti and Petrongolo (2014), instead, propose different calibrations for the labor demand elasticity within skilled workers - i.e.: $\gamma_S = 3.5$; $\gamma_S = 5$ - and unskilled workers - i.e.: $\gamma_U = 1.5$; $\gamma_U = 2.5$ -, while assuming no heterogeneity in the labor supply elasticity $\lambda$. Provided that the calibrations proposed by Olivetti and Petrongolo (2014) fall under the bell of the prior densities for both parameters, we finally decide to not change them.

**Skilled Workers.** First, compared to the aggregate, focusing only on skilled workers allows to detect a *positive* gender-specific labor supply shock, which significantly contributes to the steep convergence in employment levels (right panel) since early '80s and at the same time responsible for the slow convergence in wages (left panel).

Second, figure 2.10 rejects the conclusion in favor of the dominant role of gender-specific labor demand over supply trend shock in the skilled sector. As a matter of fact, both gender-specific shocks are relevant for the aggregate macro variables. In particular, the gender-specific labor supply shock contribution to GDP (light blue) is no longer negligible, as it was the case in the baseline specification. The contribution of the gender-specific labor supply shock is even relatively more important than labor demand when looking at the contribution to aggregate employment (right panel). Notice, in addition, that the light blue area has absorbed part of the contribution explained by the gender-

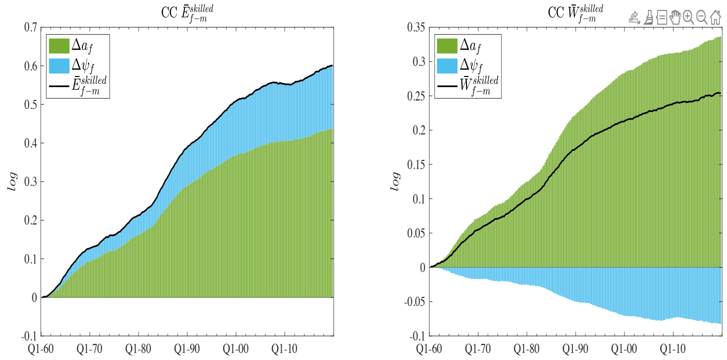


Figure 2.9: Contribution of each *structural* trend to the *empirical* trend. **Skilled workers**

specific labor demand (green), while the aggregate shocks - i.e.: labor supply (purple) and automation (yellow) - explains essentially the same as in the baseline results. This is a very important result as it confirms our concerns about the possibility of missing some features in the data that at the aggregate level were not possible to capture and that require to zoom in to avoid rash conclusions.

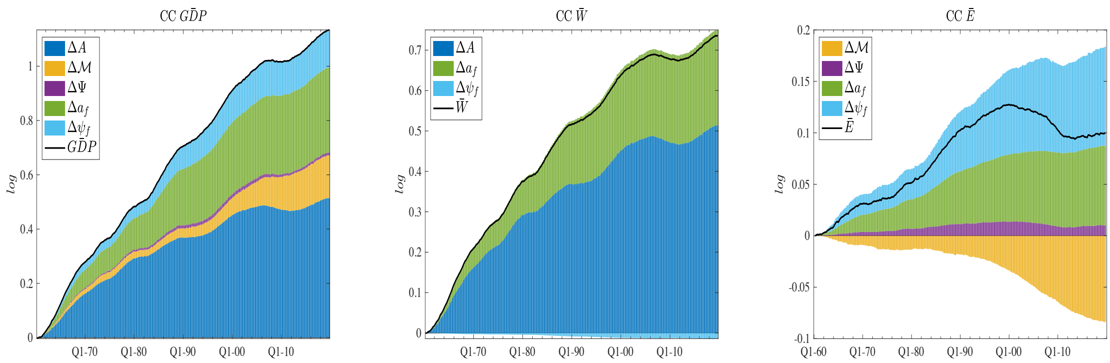


Figure 2.10: Contribution of each *structural* trend to the *empirical* trend. **Skilled workers**

**Unskilled Workers.** Let us now discuss the results from the specification controlling for *unskilled* workforce. the left panel of figure 2.11 shows as the gender convergence in unskilled employment has been much more weaker than for skilled labor. By comparing the magnitudes on the y-axis, the convergence peak of unskilled employment - that stopped in 1990 - is comparable to the employment convergence reached by female skilled workers already in the '70s. In addition, starting from 1990 a strong negate gender-specific labor supply shock dragged *unskilled* employment gap down until the end of the sample to levels of gender inequality similar to those observable in the 60s. Yet, this forceful negative labor supply shocks produced a steeper convergence in the wage rates in the final part of the sample, when the gender-specific labor demand substantially slowed down.

Figure 2.12 shows the trend decomposition of aggregate variables. Similarly to the results in the baseline model, the (*unskilled*) gender-specific labor demand trend dominates over gender specific labor supply for all the three macro variables. Only labor demand factors matter for the secular dynamics of GDP and wages and no effects whatsoever arising from labor supply factors. In contrast with the baseline model, in which labor supply factors played some role in driving the upward trend of aggregate employment, the *unskilled* gender-specific labor supply contributes to the slowdown of the employment trend (right panel) over the entire sample and plays a non-negligible role especially from the '90s.

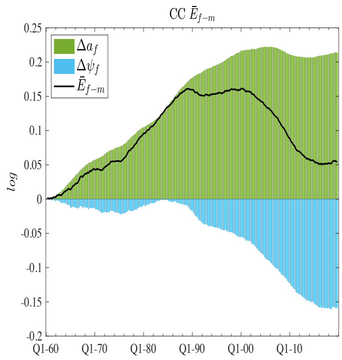


Figure 2.11: Contribution of each *structural* trend to the *empirical* trend. **Unskilled workers**

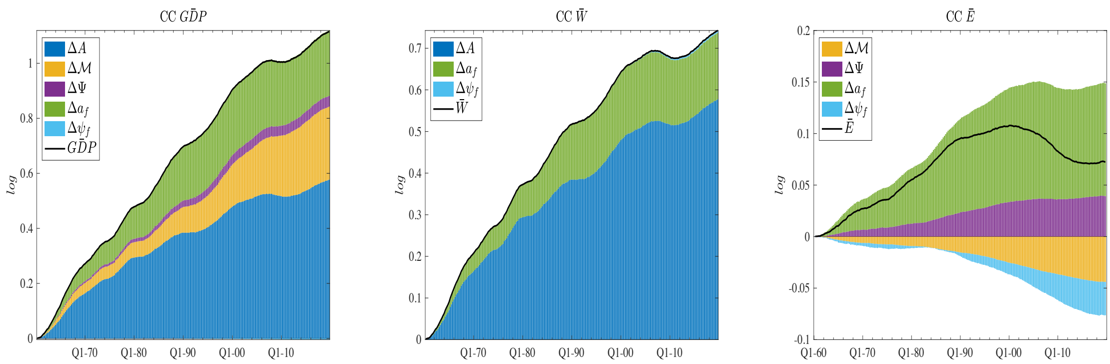


Figure 2.12: Contribution of each *structural* trend to the *empirical* trend. **Unskilled workers**

The key result from this section is that the Gender Revolution is more than simply labor demand factors. As a matter of fact, the results from the alternative specifications show that, once controlling for the skill dimension of the workforce, gender-specific labor supply factors do play a non-negligible role, both in driving gender differentials and macroeconomic outcome. The reason why labor supply factors disappear in the aggregate, paving the way to the hegemony of labor demand factors, is twofold: (i) there exists a *positive* labor supply shock to *skilled* workers that turns out to be commingled with the labor demand shock in the baseline model, leading us to hold only labor demand accountable for the macroeconomic effects of the Gender Revolution; (ii) a forceful *negative* labor supply shock to *unskilled* workers partially off-sets the effects of the labor supply shock to *skilled* workers, partially canceling out the contribution of labor supply factors to the gender differentials in the aggregate.

## 2.7 Conclusions

This paper investigates the macroeconomic implications of the Gender Revolution. It aims at quantifying its impact on post-war economic growth in the US in terms of GDP, employment and wages. In addition, this study aims to shed light on the factors behind the secular convergence of employment and wages between female and male. To this end, we estimate a SVAR with common trends à la Del Negro et al. (2017a) and propose a decomposition of the empirical (reduced-form) trends into selected unobserved structural trends, that is motivated by economic theory.

Our first contribution is methodological. We propose an explicit direct mapping between the unit roots characterising the steady state of a Neo-classical model and the (reduced-form) slow-moving trends observed in the data. We show that by means an appropriate parametrization of the matrix of long-run elasticities, denoted by $\mathcal{V}$, it is possible to exploit the information inherited from theory and uniquely pin down the underlying structural drivers that characterise the secular dynamics of selected variables in the data. This paper focuses on a specific application that aims to quantify the feedback of the Gender Revolution to the macroeconomy. Nonetheless, our identification procedure can be extended to the study of any other secular trends observed in the data.

Second, our empirical model documents the importance of gender-specific structural forces not only for the reduction of gender inequality (gender convergence) in the labor market, but also for economic growth. In particular, we show that gender-specific slow-moving trends account for up to 50% of the GDP trend growth rate over the period 1960-1990. Furthermore, according to our model, the flattening of the gender convergence started in '90s is accountable for the marked slow-down observed in trend growth over the last 25 years.

The third contribution of the paper is to quantify the individual role of the gender-specific labor demand and supply in driving the secular convergence of the gender differentials. From the estimation of our baseline model, labor demand shocks completely overshadow labor supply. Hence, according to the baseline results, one would conclude that labor supply plays no role whatsoever in explaining the gender convergence. However, this conclusion reveals to be hasty once controlling for the skill dimension of the labor market. Exploiting more disaggregated data we are able, in fact, to rationalise the main results of the baseline model. More specifically, the macroeconomic implications of the Gender Revolution are not solely due to labor demand, but also to (*skilled*) labor supply. The baseline model fails to account for this, because the *skilled* labor supply trend is commingled with the labor demand. Yet, the presence of two opposite labor supply shocks to skilled and unskilled workers cancels out in the aggregate eventually downplaying the role of labor supply in driving the secular convergence of gender differentials in the US labor market.

## 2.8 Appendix

### 2.8.1 Additional Figures & Tables

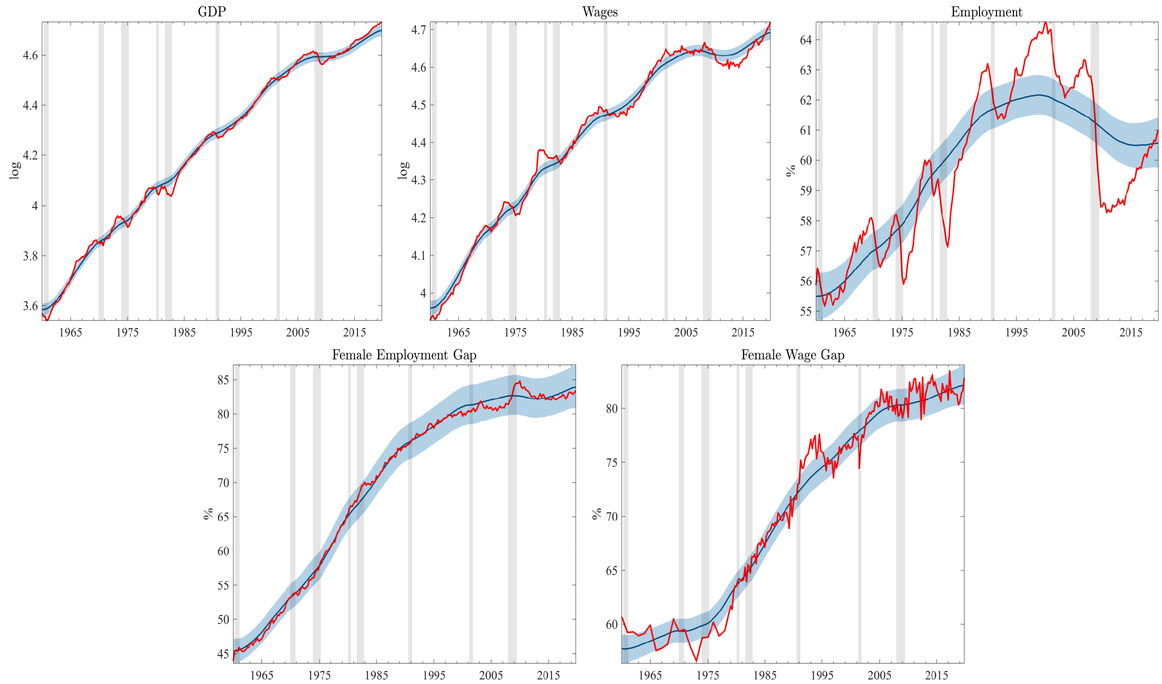


Figure 2.13: **Empirical trends baseline model**. Data in red, Median trend estimate in blue. 68% uncertainty band shaded area.

| $\Delta \bar{E}_{f,t}$ | | | |
|---|---|---|---|
| | $a_f$ | $\psi_f$ | avg. ann. growth rate |
| 1960-1969 | 0.90 | 0.71 | 1.61 |
| 1970-1979 | 1.42 | 0.63 | 2.05 |
| 1980-1989 | 1.70 | -0.17 | 1.53 |
| 1990-1999 | 0.97 | -0.29 | 0.68 |
| 2000-2009 | 0.47 | -0.25 | 0.22 |
| 2010-2019 | 0.25 | -0.08 | 0.17 |

| $\Delta \bar{W}_{f,t}$ | | | |
|---|---|---|---|
| | $a_f$ | $\psi_f$ | avg. ann. growth rate |
| 1960-1969 | 0.59 | -0.29 | 0.30 |
| 1970-1979 | 0.92 | -0.26 | 0.66 |
| 1980-1989 | 1.12 | 0.07 | 1.19 |
| 1990-1999 | 0.64 | 0.12 | 0.76 |
| 2000-2009 | 0.31 | 0.10 | 0.41 |
| 2010-2019 | 0.16 | 0.04 | 0.20 |

Table 2.4: **Baseline model.** Relative contribution of gender-specific shocks to the annualised trend growth rate of employment gap. Average of each corresponding decade.

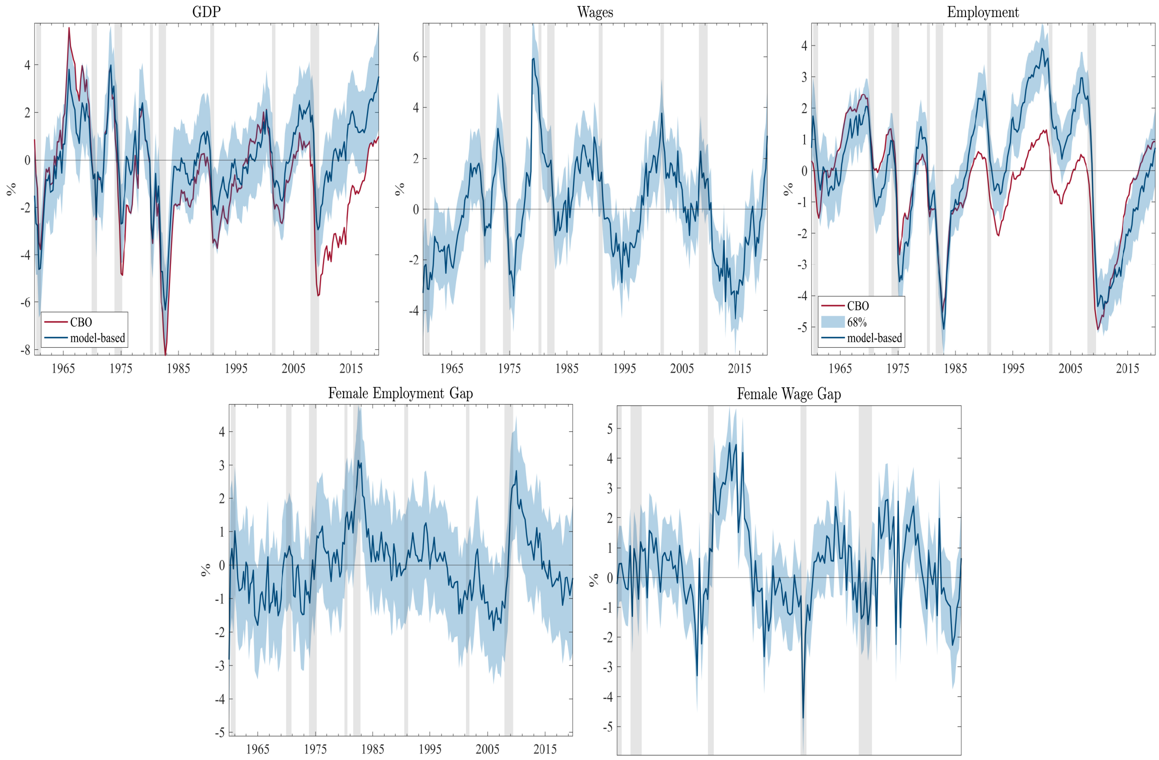


Figure 2.14: **Empirical Cycles baseline model**. Data in red, Median trend estimate in blue. 68% uncertainty band shaded area.

### 2.8.2 Data

Data available on FRED website is listed in the table below along with their identification code.

| DATA | CODE |
|---|---|
| Real Gross Domestic Product per capita | A939RX0Q048SBEA |
| Non-farm business sector: real compensation per hour | COMPRNFB |
| Employment level, thousands of persons | CE16OV |
| Population level, thousands of persons | CNP16OV |
| Employment-to-Population ratio | EMRATIO |
| Women Employment-to-Population ratio | LNS12300002 |
| Men Employment-to-Population ratio | LNS12300001 |
| Women nominal weekly earnings | LES1252882700Q |
| Men nominal weekly earnings | LES1252881800Q |

Table 2.5: Data codes

**Transformations.** Data only available at monthly frequency (e.g.: employment, population, etc.) are transformed into quarterly by taking the three-month average of each corresponding quarter. Real aggregate wages per capita is retrieved from the following product $COMPRNFB \times \frac{CE16OV}{CNP16OV}$. Regarding the gender employment and wage gaps, these are computed as the ratios of Women employment-to-Male employment and Women wages-to-Male wages. Before computing the ratios, the gender-specific wage rates are transformed into hourly wage rates, dividing them by 52 (number of working weeks in a year). Also, unfortunately gender-specific wage rates are only available from 1979Q1. To missing observations, the period spanning 1960-1978 is filled by the earnings data available from the Annual Social and Economic Supplements, Current Population Survey, U.S. Census Bureau. These data are at annual frequency, thus we decide to get the intra-annual observations by using standard interpolation techniques.

**Gender-specific Data by Sectors and Skills.** Following Dolado et al. (2021), we download data on hourly wages and employment by skills and sector from the publicly available Current Population Survey (CPS) produced by the United Census Bureau and the Bureau of Labor Statistics downloadable at `http://data.nber.org/morg/annual/`. Then, we edit the STATA code used by Dolado et al. (2021) to merge the CPS survey into a unified dataset and back out the gender employment and wage gaps by skills and sectors. We refer to the online Appendix of Dolado et al. (2021) for a detailed discussion on how the CPS dataset is merged. For the purpose of our analysis, we are interested in retrieving the employment level and hourly wage rate of women and men by skills and sectors. Regarding the skills dimension, we limit to split individuals into those with at least *some college* experience (skilled) and those who do not have any college education at all (unskilled). Hence, we are able to retrieve the employment level and hourly wages of women and men skilled and unskilled workers.

### 2.8.3 Estimation of the State Space Model with Gibbs Sampling

Consider the unobserved states of the model in section 2.2 in the following stacked formulation:

$$\begin{bmatrix} \mathcal{V}X_t \\ \hat{Y}_t \end{bmatrix} = \begin{bmatrix} \mathcal{V}c \\ 0 \end{bmatrix} + \begin{bmatrix} I & 0 \\ 0 & A \end{bmatrix} \begin{bmatrix} \mathcal{V}X_{t-1} \\ \hat{Y}_{t-1} \end{bmatrix} + \begin{bmatrix} I & 0 \\ 0 & I \end{bmatrix} \begin{bmatrix} \mathcal{V}u_t \\ e_t \end{bmatrix} \tag{2.27}$$

and the Covariance matrix of the model is given by $\Sigma$:

$$\Sigma = \begin{bmatrix} \mathcal{V}'\Sigma_u\mathcal{V} & 0 \\ 0 & \Sigma_e \end{bmatrix} \tag{2.28}$$

Then, the unobserved states of the model are estimated numerically according to a Gibbs sampler algorithm, structured in the following steps:

1. Draw from the joint distribution $X_{0:T}, \hat{Y}_{-p+1:T}, \nu \mid c, A, \Sigma_u, \Sigma_e, Y_{1:T}$, which is given by the product of the marginal posterior of $\nu$ - vector of free parameters in $\mathcal{V}$ - conditional on the other parameters $\nu \mid c, A, \Sigma_u, \Sigma_e, Y_{1:T}$ and the distribution of the unobserved states conditional on $\nu$ and the other parameters $X_{0:T}, \hat{Y}_{-p+1:T} \mid$

$\nu, c, A, \Sigma_u, \Sigma_e, Y_{1:T}$.

(a) $p(\nu \mid c, A, \Sigma_u, \Sigma_e, Y_{1:T}) \propto \mathcal{L}(Y_{1:T} \mid \nu, c, A, \Sigma_u, \Sigma_e)p(\nu)$,
where $\mathcal{L}(Y_{1:T} \mid \nu, c, A, \Sigma_u, \Sigma_e)$ is the likelihood of the data obtained from the Kalman filter applied to the state space of the model. The posterior of $\nu$ does not have known solution, therefore we approximate it by introducing a Metropolis Hastings step.

(b) Draws from p$(X_{0:T}, \hat{Y}_{-p+1:T} \mid \nu, c, A, \Sigma_u, \Sigma_e, Y_{1:T})$ are obtained implementing Durbin and Koopman (2002a) simulation smoothing algorithm.

2. Draw from the joint distribution $A, c, \Sigma_u, \Sigma_e \mid X_{0:T}, \hat{Y}_{-p+1:T}, Y_{1:T}$. The estimation of the remaining parameters is relatively straightforward, provided that the unobserved states follow rather standard vector autoregressive laws of motion.

(a) **Trend Block.** the posterior distribution of $\Sigma_u$ is given by:

$$p(\Sigma_u \mid X_{0:T}) = \mathcal{IW}(\underline{\Sigma_u} + \underbrace{\sum_{t=1}^{T}(X_t - X_{t-1})(X_t - X_{t-1})'}_{S_u}, \kappa_u + T)$$

The posterior distribution of the vector of drifts $c$ conditional on the $\Sigma_u$ and $X_{0:T}$ is obtained from a standard Normal.

(b) **Cycle Block.** The posterior distributions of the lag coefficients in $A$ and the covariance matrix $\Sigma_e$ of the stationary VAR are standard:

$$p(\Sigma_e \mid \hat{Y}_{0:T}) = \mathcal{IW}(\underline{\Sigma_e} + S_e, \kappa_e + T)$$
$$p(A \mid \Sigma_e, \hat{Y}_{0:T}) = N\left(vec(\mathcal{A}), \Sigma_e \otimes \left(\sum_{t=1}^{T} \hat{Z}_t \hat{Z}_t' + \underline{\Omega}^{-1}\right)^{-1}\right)$$

where $\hat{Z}_t = (\hat{Y}_{t-1}', \ldots, \hat{Y}_{t-p}')$, $\mathcal{A} = \left(\sum_{t=1}^{T} \hat{Z}_t \hat{Z}_t' + \underline{\Omega}^{-1}\right)^{-1} \left(\sum_{t=1}^{T} \hat{Z}_t \hat{Y}_t' + \underline{\Omega}^{-1}\underline{A}\right)$,
$S_e = \sum_{t=1}^{T} e_t e_t' + (\mathcal{A} - \underline{A})'\underline{\Omega}^{-1}(\mathcal{A} - \underline{A})$

The algorithm produces 50000 draws and discards the first 40000.

### 2.8.4 Priors for the volatilities of the structural trend shocks $u_t$

One non-trivial task for the econometrician is to come up with reasonable priors for the elements of $\sigma_{\mathbf{u}}^{\mathbf{2}} = [\sigma_A^2\ \sigma_{\mathcal{M}}^2\ \sigma_{\Psi}^2\ \sigma_{a_f}^2\ \sigma_{\psi_f}^2]'$, the vector stacking the shocks' volatilities of the *structural* trends in $X_t$. The reason is due to the fact that the structural trends are *unobservable* in the first place. Nonetheless, it is still possible to form fairly non-judgmental priors on these structural volatilities by combining two pieces of information we already possess, namely: (i) the data and (ii) the theory-based prior beliefs on the free parameters in $\mathcal{V}(\nu)$. To see this, recall that empirical and structural trends are linked by the linear relationship $\bar{Y}_t = \mathcal{V}X_t$ and that $X_t = c + X_{t-1} + u_t$. Without loss of generality, one can express the empirical trends in their growth rates, as follows:

$$\Delta\bar{Y}_t = \mathcal{V}(c + u_t)$$

This equation implies that the covariance matrix of the empirical trends in growth rates is denoted by $\Sigma_{\Delta\bar{Y}} = \mathcal{V}'\Sigma_u\mathcal{V}$. Then, provided that covariance matrix of the structural trends shocks $\Sigma_u$ is diagonal, the following linear relations apply:

$$\begin{aligned}
\sigma^2_{G\bar{D}P} &= \sigma^2_A + \sigma^2_{\mathcal{M}} + \nu^2_{13}\sigma^2_{\Psi} + \nu^2_{14}\sigma^2_{a_f} + \nu^2_{15}\sigma^2_{\psi_f} & \sigma^2_{\bar{W}} &= \sigma^2_A + \nu^2_{24}\sigma^2_{a_f} + \nu^2_{25}\sigma^2_{\psi_f} \\
\sigma^2_{\bar{E}} &= \nu^2_{32}\sigma^2_{\mathcal{M}} + \nu^2_{33}\sigma^2_{\Psi} + \nu^2_{34}\sigma^2_{a_f} + \nu^2_{35}\sigma^2_{\psi_f} & \sigma^2_{\bar{E}_{f-m}} &= \nu^2_{44}\sigma^2_{a_f} + \nu^2_{45}\sigma^2_{\psi_f} \\
\sigma^2_{\bar{W}_{f-m}} &= \sigma^2_{a_f} + \sigma^2_{\psi_f}
\end{aligned}$$

On the left hand side of each equation there are the volatilities of the empirical trends in growth rates, while on the right hand side there are the coefficients of $\mathcal{V}$ and volatilities of the structural shocks. The empirical volatilities are available in the data and the parameters $\nu_{ii}$ are simply the values around which the prior density of the long-run elasticities is centered. The only *unknowns* are the structural volatilities, in fact. It turns out that is straightforward to retrieve the structural volatilities in $\boldsymbol{\sigma^2_u}$, as they are the unknowns of a linear system of 5 equation in 5 unknowns and, therefore, there always exists a unique solution to the system.

Consistently, this is how we proceed in practice. First, back out the empirical volatilities from the HP-filter trend growth rates of the endogenous variables using a pre-sample training. Second, plug the empirical volatilities and the prior means of the parameters in $\mathcal{V}$. Finally, solve the system for the unknown volatilities and use them to center the prior density of the structural volatilities.

Notice that the very same reasoning applies when forming priors for the initial conditions and the drifts of the structural trends. Accordingly, the initial conditions $X_0$ should be centered around $X_0 = \mathcal{V}\bar{Y}_0$, with $\bar{Y}_0$ being the *last period*'s empirical trend *in levels* (last period in the training sample). As for the drifts, the constants $c$ should be centered around $c = \mathcal{V}\mathbb{E}(\Delta\bar{Y}_t)$, with $\mathbb{E}(\Delta\bar{Y}_t)$ being the average of the empirical trends *in growth rates* (in the training sample).

# Chapter 3

# Output, Inflation and Oil Prices in the Euro Area

## 3.1 Introduction

In April 2021 the monthly year-on-year percentage change of Brent crude oil price skyrocketed to 252%.[1] Oil price kept registering record highs throughout the 2021. Yet, the rising geopolitical tensions on the eastern Ukrainian border contributed to fan the flame by increasing market uncertainty about the future price of oil and natural gas. Against this background, the US and Euro Area inflation rates woke up after a long period of apparent hibernation. The US consumer price inflation was already beyond 5% in June and above 7% in December.[2] The HICP inflation in the Euro Area initially lagged behind, but then reached its all time high (above 4.8%) since the inception of the Euro in November and kept scaling up to 7.4% in March 2022.[3] The massive energy price swings together with the raise of inflation have brought about a revival of the discussion in policy and academic circles about the consequences of global oil price surges for domestic price stability and economic growth. The core question of the debate can be summarised with the formula: is inflation going to be *temporary* or *persistent*?

A popular argument in favour of inflation being transitory relies on the role played by global supply bottlenecks, being the result of temporary demand-supply mismatches arising from the rapid rebound of demand together with the slow re-openings from worldwide lock-downs, eventually intended to ease as supply gains capacity back. In a recent speech

[1] U.S. Energy Information Administration, Crude Oil Prices: Brent - Europe [DCOILBRENTEU], retrieved from FRED, Federal Reserve Bank of St. Louis; https://fred.stlouisfed.org/series/DCOILBRENTEU.

[2] U.S. Bureau of Labor Statistics, Consumer Price Index for All Urban Consumers: All Items in U.S. City Average [CPIAUCSL], retrieved from FRED, Federal Reserve Bank of St. Louis; https://fred.stlouisfed.org/series/CPIAUCSL.

[3] Eurostat, Harmonized Index of Consumer Prices: All Items for Euro area (19 countries) [CP0000EZ19M086NEST], retrieved from FRED, Federal Reserve Bank of St. Louis; https://fred.stlouisfed.org/series/CP0000EZ19M086NEST.

Phillip Lane[4] expressed this view by linking the surge of energy prices to supply bottlenecks, claiming it to be inflationary in the near-term, but not in medium-term. This lies upon the fact that the Euro Area is a net energy importer, so the spike in energy prices should potentially act as an adverse terms of trade shock eroding households disposable income and firms profits. This implies that, as the negative wealth effects materialise, the short-run inflationary burden will start relaxing, exerting downward pressure on the path of underlying inflation. This argument found favourable consensus in the light of the absence, at least until December 2021, of any second-round effects on negotiated wages.

However, there are reasons to believe that inflation persistence has been underestimated by policy makers and other factors pushing energy prices may leave their medium-term footprints on inflation. For instance, it is likely that wage growth in the Euro Area has remained subdued because of the job retention schemes which smoothed job losses, implying that in the wake of the economic recovery fewer workers transitioned away from unemployment and engaged wage negotiation.[5] Nonetheless, to the extent that high energy prices persist, workers might eventually claim an adjustment of their salaries generating additional inflationary pressure in the medium-term. Furthermore, in a recent study Celasun et al. (2022) show that Euro Area producer prices would have run at historical highs even in the absence of the pressures exerted by supply bottlenecks, implying that sound aggregate demand backed by the joint action of fiscal and monetary stance may have induced persistent effects.

Against this background, the aim of this paper is to contribute to the current debate by shedding light onto which structural drivers of energy prices, and in our specific case oil prices, are informative to EA monetary policy medium-term objectives. For this purpose, the empirical literature on the global oil market comes to help. In a recent paper, Baumeister and Hamilton (2019) propose a state-of-the-art identification scheme to retrieve three main drivers of global oil prices, namely: (i) oil-specific supply; (ii) oil-specific demand; (iii) world economic activity. In particular, we are interested in the last two structural forces. The reason is due to the fact that, despite the literature on global oil markets has stressed the importance of these two shocks in explaining the last 30 years fluctuations in oil global oil prices,[6] little attention has been posed in analysing their implications domestic growth and inflation. As a matter of fact, most of the studies investigating the macroeconomic consequences of oil price spikes focus on oil-specific supply shocks (e.g.: Baumeister et al. (2010), Peersman and Robays (2009),

[4]Welcome address by Philip R. Lane, Member of the Executive Board of the ECB, at the ECB Conference on Money Markets, 8 November 2021.

[5]In a recent speech Isabel Schnabel acknowledged the underestimation of inflation persistence by policy-makers and stressed the potential persistent effects wage growth revival. Speech by Isabel Schnabel, Member of the Executive Board of the ECB, at a workshop organised by the European House – Ambrosetti on "The Agenda for Europe: Macroeconomic and Structural Policy Challenges", Cernobbio, 2 April 2022.

[6]For a discussion see Kilian (2008), Kilian (2009), Kilian and Murphy (2012b)

among other).

The empirical exercise is framed within a recursive Bayesian VAR for the Euro Area augmented with the Baumeister and Hamilton (2019) structural shocks. This simple setting accommodates for the possibility to easily expand the information set and explore the transmission mechanisms whereby global shocks propagate to the Euro Area. This has the advantage, not only to test if these shocks can be held accountable for *persistently* higher inflation, but also to unveil the channels whereby inflation and other fundamentals are affected. Furthermore, it should be highlighted that the estimation window is 1999Q1-2019Q4, excluding therefore the pandemic. This is motivated by the fact that: (i) the ultimate goal of this exercise is to not to explain the relative contribution of the aforementioned global shocks to the recent events, but to uncover their effects and transmission channels; (ii) excluding the pandemic allows us to test if - with this information set at hands - one could form reasonable priors about the current inflation and output paths, contributing to the ongoing debate; (iii) from a purely practical viewpoint, expanding the sample to include the pandemic would have substantially complicated the estimation of the model, given the massive deviations between trough and peak.

The results of the paper can be summarised as follows. On the one hand, an oil-specific demand shocks induces *temporary* inflationary pressures and a *persistent* worsening of economic growth. This is the result of: (i) the direct income effect from the immediate raise of energy price, which erodes households disposable income; (ii) the uncertainty about future energy prices, that discourages demand for durable goods. On the other hand, an oil price increase induced by a global economic activity shock leads to a *temporary* improvement in the fundamentals and *persistent* inflationary pressures. The persistent effect on domestic inflation is set in motion by an increase in the domestic cost of production which are eventually discharged into final consumers prices, potentially triggering wage-price spirals.

Finally, the remainder of the paper is organised as follows. Section 3.2 outlines the methodology. Section 3.3 presents the results from the baseline specification of the model. The transmission mechanisms are investigate through sections 3.4 to 3.5. Section 3.6 provides additional support to the interpretation of results. Section 3.7 concludes.

## 3.2 A Bayesian VAR for the Euro Area

Consider the following BVAR representation of the Euro Area economy:

$$A_0^{-1}y_t = c + B_1 y_{t-1} + ... + B_p y_{t-p} + \varepsilon_t, \tag{3.1}$$

where $\varepsilon_t$ is the vector of zero-mean normally distributed innovations with covariance $\Sigma = A_0 A_0'$. The aim of this empirical exercise is to uncover the propagation effects of external (global) shocks on the Euro Area. It is, therefore, convenient to partition the vector $y_t$ into two blocks $[y_{1t}\ y_{2t}]'$. The first block $y_{1t}$ captures any movements in commodity prices and world demand. Hence, $y_{1t}$ defines the block of *predetermined* variables, which are not affected - at least on impact - by unexpected changes in Euro Area main aggregates. The block $y_{2t}$ describes the dynamics of the Euro Area macroeconomy. In the baseline specification this second block includes real GDP, unemployment rate, HICP and short-term interest rate.

The literature on global oil market comes to help in identifying the sources of fluctuations in commodity prices and world economic activity in the first place. Specifically, Baumeister and Hamilton (2019) combine signs together with restrictions on the short-run elasticities of the global oil market supply and demand curves to retrieve three structural shocks, namely: an oil-specific supply, a world economic activity and an oil-specific demand shock from a BVAR for the global oil market.[7] The work of Baumeister and Hamilton (2019) considerably simplifies the tasks demanded to the BVAR, as it allows one to plug the structural shock series as first in $y_{1t}$ block, postulate $A_0$ to be lower triangular and estimate the model recursively. This approach builds on the results of Plagborg-Møller and Wolf (2021), who prove under some mild conditions that the estimation of the structural impulse responses based upon instrumental variable *Local Projections* can be carried out by simply ordering the instrument first in a *recursive VAR*, even in cases of non-invertibility.

Therefore, we proceed by ordering first the structural shock series and augment the block of *predetermined* variables with the global price of Brent crude oil. This approach brings about a number of advantages. First, it is parsimonious, because it does not require to go through the cumbersome estimation of a VAR for the global oil market augmented with macroeconomic aggregates. As a matter of fact, by including the shock in the system, one can study the implications for the Euro Area economy in a small or medium-scale VAR, reducing the risk of incurring in the curse of dimensionality. Second, the structural shock series comes from a state-of-the-art identification procedure. Sign restrictions allow to pin down the co-movement between the variables and to correctly disentangle the different structural sources of fluctuations, as extensively documented by Peersman and Robays (2009), Baumeister et al. (2009), Kilian and Murphy (2012b), among others. In addition, the restrictions on the matrix of contemporaneous effects narrow the set of credible draws to only those in line with theory - we refer to Kilian and Murphy (2012b) and Baumeister and Hamilton (2019) for a detailed discussion.

However, this approach does not come without costs. Given its relatively straightfor-

[7]More specifically, an adverse oil-specific supply shock yields to *negative* co-movement between world oil production and oil prices and a *negative* effect on world economic activity on impact. An oil-specific demand shock yields to a *positive* co-movement between oil production and prices, while negatively affect economic activity. Finally, the world economic activity shock generates a *positive* co-movement between oil production, prices and economic activity.

ward recursive implementation, it allows to study the feedback to the macroeconomy of only one shock at time. Therefore, this would not be an appropriate apparatus if one is interested in designing a comprehensive framework to quantify the relative importance of each global shock. Nonetheless, this approach has the merit of: (i) understanding which sources of global fluctuations might be *looked-through* and those which instead require the prompt intervention of monetary policy; (ii) investigate the channels whereby macroeconomic aggregates are affected.

The model is estimated for the period 1999Q1-2019Q4 with variables entering in log-levels. The shock series are directly downloaded from Christiane Baumeister's website, the global price of Brent crude oil is available on FRED, while Euro Area macro aggregates are retrieved from Eurostat. I assume the VAR to have $p = 4$ lags. Given that all macro variables and the price of Brent are in levels, I employ standard Minnesota priors with overall tightness parameter equal to 0.2, as suggested by Giannone et al. (2015). In addition, differently from the remainder variables in $y_t$, whose own lag prior is centered around 1, the prior for the shock series is centered around 0, provided that describes a process without memory.

## 3.3 Baseline Results

In this section the results from the benchmark specification are reported. Recall that the focus will be on the effects of oil-specific demand and global economic activity shocks. The different potential transmission mechanisms will be investigated in the next sections.

Before presenting the baseline results, it is worth recalling the economic interpretation of the oil-specific demand and global economic activity shocks. The former is meant to capture any oil price increase driven by higher precautionary demand associated with market concerns about future availability of oil. The latter refers to any improvement in the world economic outlook which spurs demand for oil and other industrial commodities.

Figure 3.1 displays the impulse responses of the Euro Area main macroeconomic aggregates to an oil-specific demand (red) and a global economic activity (green) shock. Both shocks are normalised to raise the price of crude oil Brent by 10 percentage point after one quarter. Let us begin with the responses to an oil-specific demand shock. Output takes approximately 1 year to become statistically significant and permanently drop by 0.4pp after 5 years. Unemployment rapidly raises in the three years following the shock and plateaus at a higher level (+0.25pp). Moving to the nominal side of the economy, headline inflation and nominal rate positively co-move. Inflation peaks after 1 year and starts converging back to its trend thereafter. The systematic response of monetary policy implies a rate increase by 10 basis points in the first year after the shock, then gradually relaxing, as inflation normalises. Overall, an increase in global oil prices induced by concerns about future oil supply yields to output-inflation trade-off,

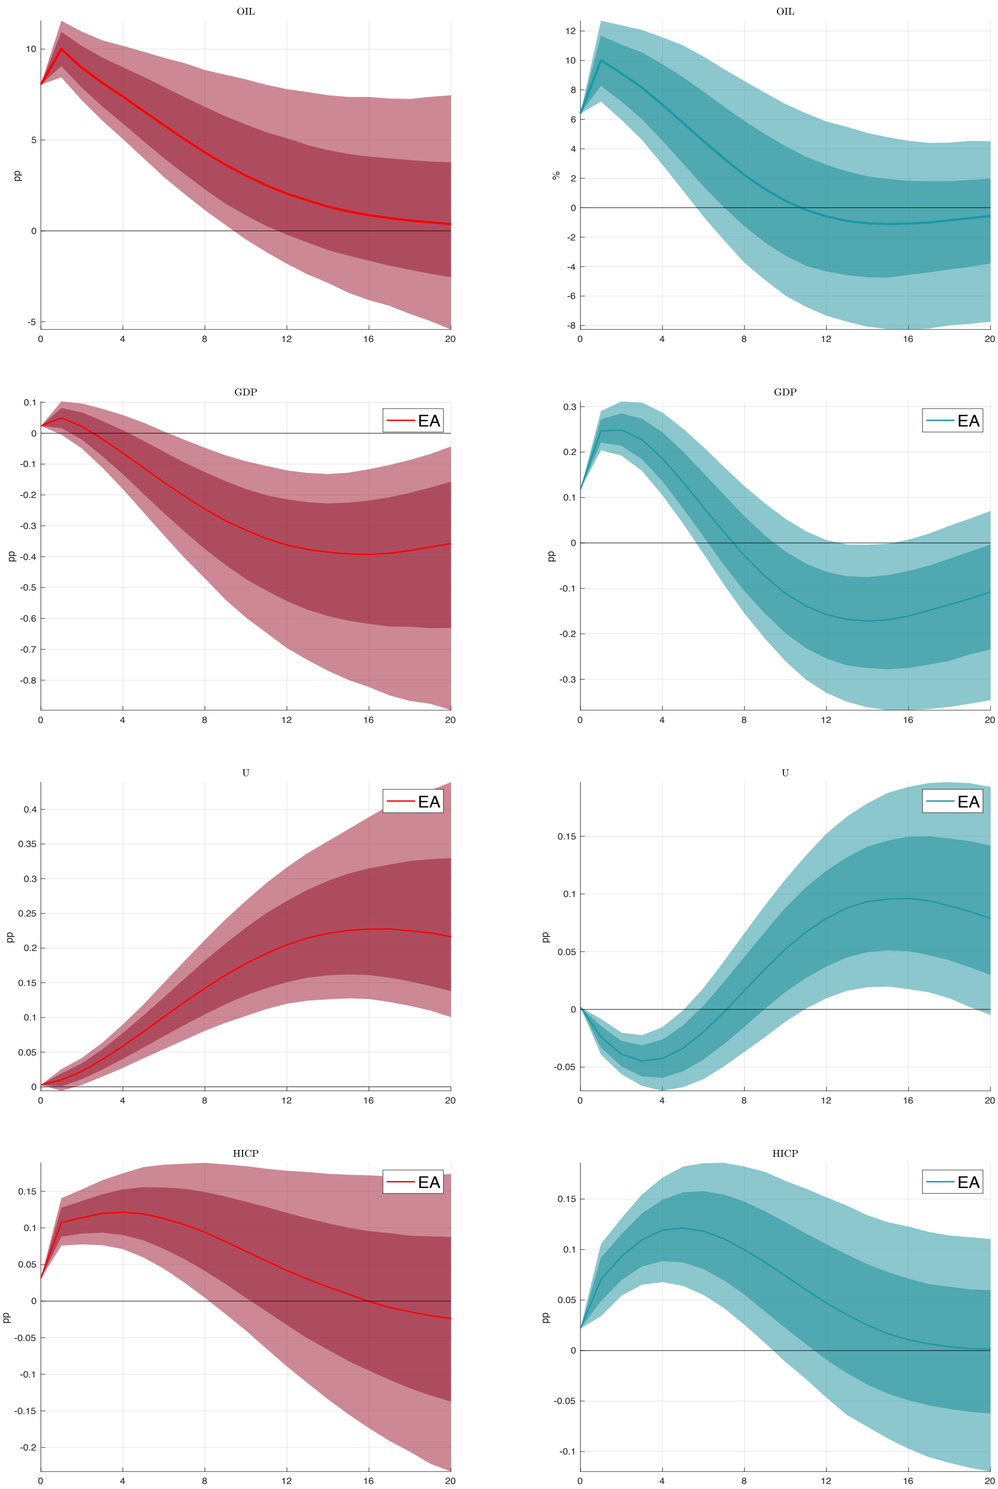


with the Central Bank systematically reacting to higher inflation.[8]

The second column of figure 3.1 shows the impulse responses to a favourable world economic activity shock. In contrast to the response to the oil-specific demand shock, output increases on impact and stays higher for approximately two years, before slowing

[8]One word of caution on terminology, one should not be mislead by the label oil-specific "demand" shock, because it feeds into the economy as an adverse aggregate "supply" shock, in the sense that generates an opposite co-movement between output and prices.

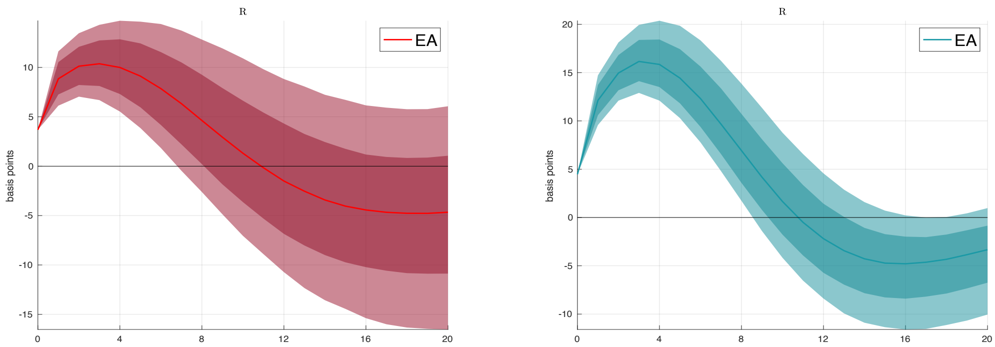


Figure 3.1: Impulse responses to an oil-specific demand shock (**red**) and global economic activity shock (**green**). Shaded areas are 68% and 90% credibility intervals, respectively.

down and particularly so after 3-4 years. The response of unemployment is sluggish; it timidly declines for one year before rebounding and reaching the peak effect after 4 years. Moving to the response of inflation, figure 3.1 displays an hump-shaped response peaking after 1.5 year. Compared to the oil demand shock, the raise in the short rate is approximately 5bps higher and more persistent.

Jointly assessing the two shocks, it is worth highlighting a couple of facts. First, given its recessionary effects, the oil-specific demand immediately generates a trade-off between inflation and output stabilisation which is not observed - at least in the first 2 years - in the case of a shock to world economic activity, implying in turn a more aggressive systematic response of the short rate. Second, The recessionary effects of the oil demand shock materialise immediately with a rapid surge in unemployment. Conversely, the world economic activity shock, improves fundamentals, but only temporarily - unemployment raises and significantly at 90% credibility after 3 years -. In addition, the 0.1pp peak of unemployment is reached after 4 years, while in response to the oil demand shock, unemployment is 0.1pp higher already after 1.5 year and steadily grows thereafter. Thus, the sluggish response of unemployment to the economic activity shock might be due to raising costs of production, while in response to the oil demand shock the rapid and persistent recessionary effects might be due to adverse income effects on the side of consumers. So far, this is purely speculative, but, if anything, these results call for a deeper analysis of the transmission mechanism, which is, indeed, what will be done in the next sections.

## 3.4 Effects on GDP Components

In this section, we will dissect GDP into its main components. The aim is to unveil the potential different channels thereby affecting output. The literature lists a number of different channels - falling under the umbrella of demand - which can ultimately drive the response of GDP. For instance, an extensive strand of literature has focused on the consequences of higher energy prices on consumption and savings decisions of agents - see e.g. Bernanke (1983), Bernanke (2006), Hamilton (2005), among others. This strand of literature poses particular emphasis on three main effects, namely: (i) *income effect*; (ii) *uncertainty effect*; (iii) *precautionary savings*. Regarding the income effect, Edelstein and Kilian (2009) - for instance - provide empirical evidence of a reduction of households' *disposable income* in response to raising energy prices, thereby contributing to a drop in overall consumption in the United States.[9] The uncertainty effect, instead, is meant to capture any postponements of potentially *irreversible* investments, such as expenditure in durables, associated with the uncertain future path of energy prices.[10] Finally, an increase in precautionary savings may occur if households believe that higher energy prices may translate into a higher probability of losing employment in the future.[11] The overall effect on output also depends on the response of the relative price of domestic goods in the international market. The inflationary pressure of global shocks may induce, for instance, a *depreciation* of the exchange rate vis-à-vis the main trading partners encouraging foreign demand for domestic products, thereby dampening the potential recessionary effects on output in the medium-run.

Before delving into the aforementioned demand channels, let us discuss the response of each component of GDP first. The first column of figure 3.2 visualises the responses to the oil-specific demand shock. Consumption drops persistently and keeps declining for three years before stabilising at -0.25pp, thereafter. Investment strongly co-moves with aggregate consumption, but it takes one year to see the adverse effects materialising. Notice also that investment decreases even more than the other components as it remains 1 percentage point lower after five years. Government spending lays dormant for the first three years and then permanently declines. This result seems to support the poor coordination between the ECB and member states' national fiscal authorities in the pre-pandemic sample period, as reported by Reichlin et al. (2021). Moving to real exports, it is rather interesting to observe an increase, though short-lived. Conventional theory suggests that following a global shock which raising domestic inflation, the domestic currency should depreciate in the short-run and, provided that aggregate demand is inelastic, real value of exports should decrease while quantities remain fixed. The chart, however, exhibits a short-run increase in real exports. This suggests that - at least in the short-term - euro appreciates vis-à-vis the US and the other main trading partners.[12]

[9] It should be stressed that the ultimate effect primarily depends on the price elasticity of households' energy demand. The more inelastic, the larger the effect on consumption decisions. Yet, even assuming that the demand for energy is rather inelastic, the consequences for aggregate real consumption remain bounded by the share of energy in the overall consumption basket.

[10] Bernanke (1983) was the first documenting this mechanism leading to a dip in investment and consumption of durable goods.

[11] The weakening of labor market conditions may reinforce labor market frictions. For example, it may induce reallocation of workforce away from more energy-intensive sectors, as shown by Hamilton (1988) in a simple neoclassical model.

[12] The impulse responses for the euro/dollar exchange rate and the BIS trade-weighted exchange rate

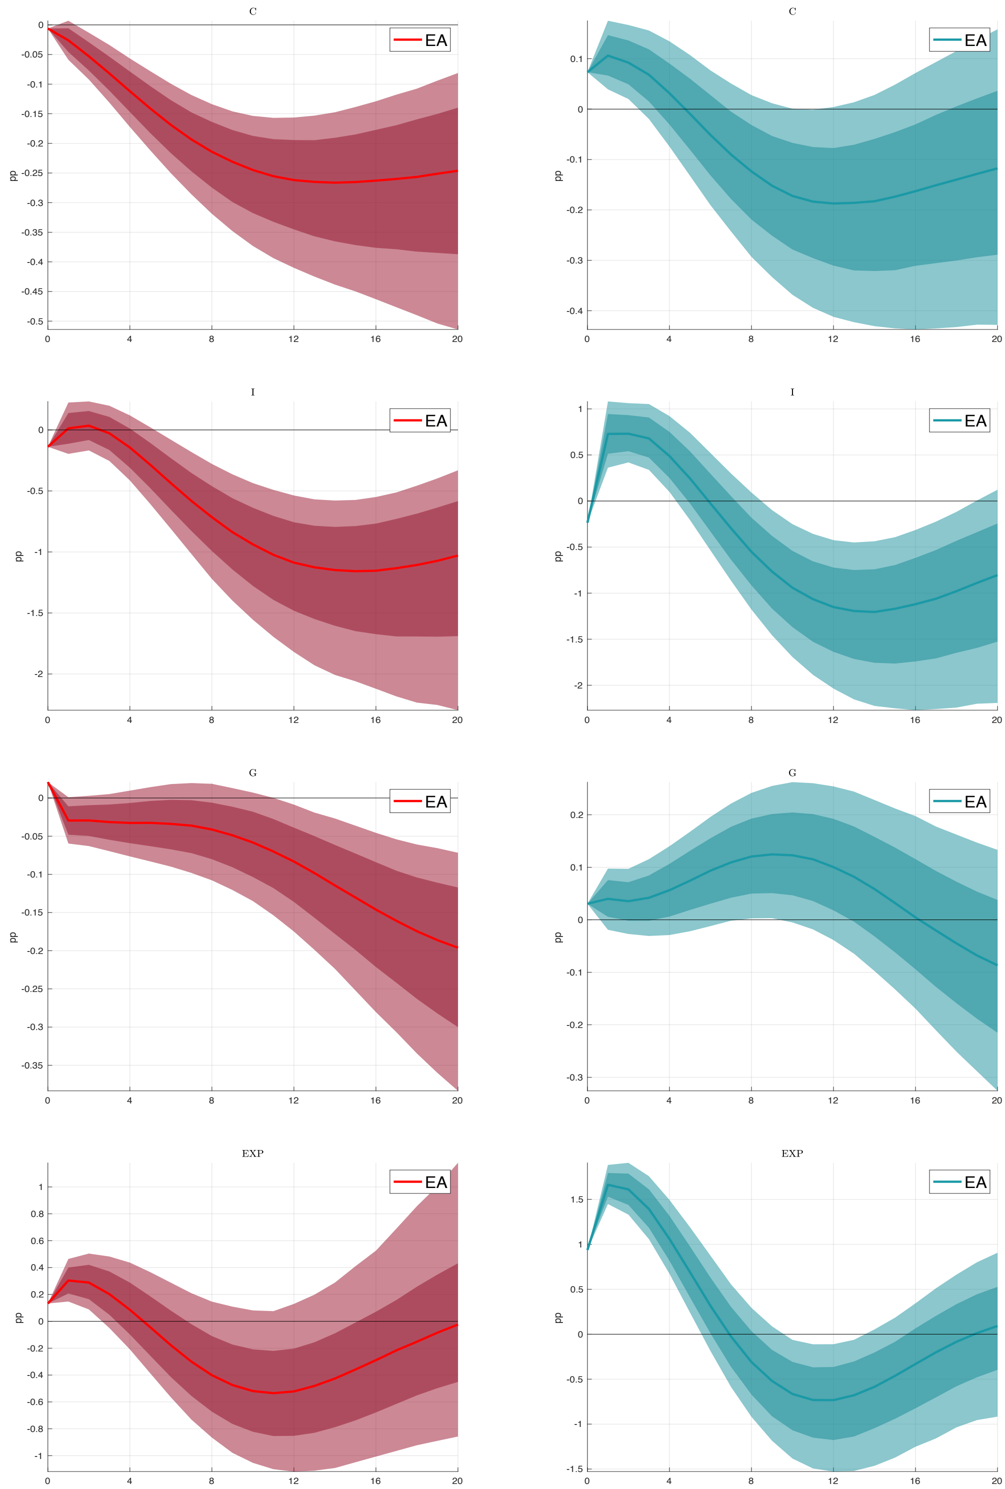


Figure 3.2: GDP components

Despite being an interesting result, the overall picture supports a substantial role played by consumption and investment in driving the response of output.

Moving to the second column of figure 3.2, in response to the economic activity

are available upon request. Both measure show a short-lived appreciation of the euro, which fades away within one year.

shock, both consumption and investment exhibit a temporary increase for one year before decreasing and persistently so, thereafter. In contrast to the responses to an oil demand shock, there is much more uncertainty around the medium-term fall in consumption, despite the median point estimates being rather close to each other. Comparing the red and green responses of investment, both display a large fall of approximately 1pp in the 3-4 years time window after the shock. In addition, the response to the economic activity shock seems to be less persistent and, as a matter of fact, converges back to zero more rapidly. government spending poorly reacts and the response is surrounded by large uncertainty, though it displays opposite sign compared to the response to the oil demand shock. Finally, the favourable economic activity shock induces a sound response of real exports, which increase by 1 percentage point on impact and reach the peak in the subsequent quarter (+1.5pp). The effect lasts for approximately two years before slowing down and finally converging back to its long-run trend. The strong transitory increase in real exports shows that the increase in domestic output is mainly driven by sound foreign demand.

### 3.4.1 The Demand Channel

The results from the previous section highlight two main facts. On the one hand, an oil-specific demand shock produces strong recessionary effects on private consumption and investment. On the other hand, in response of a favourable (world) economic activity shock, a temporary improvement in the world outlook spurs domestic demand components. However, once this initial positive effect vanishes, higher prices seem to slowdown output and raise unemployment. Therefore, the aim of this section is to shed light on the potential effects of the demand channels listed above in: (i) triggering a recession, in response of an oil-specific demand shock; (ii) contributing to the slowdown after 3-5 year, in response to the economic activity shock.

***Income effect.*** Figure 3.3 shows the impulse responses of some components of HICP, representing a substantial portion of households bills, namely energy, food (including alcohol and tobacco) and transport.[13] Let us start with the first row. Both shocks make energy bills heavier with a peak increase of approximately 1.5 percentage points within the first year and the response to the oil-demand shock (red) being more persistent. Moving to food-related items, it is interesting to observe that, on the one hand, the response to a shift in oil demand is muted and never statistically nor economically significant. On the other hand, in response to the economic activity shock, food tend to co-move with energy prices, although the peak effect is reached only with one year lag. The heterogeneous response of food prices echoes the argument proposed by Peersman (2022),

[13] In the period 1999-2019, transport share and food-related items account for 30% and 20% on average, respectively. While the weight of energy components is approximately 10%. Source: Eurostat.

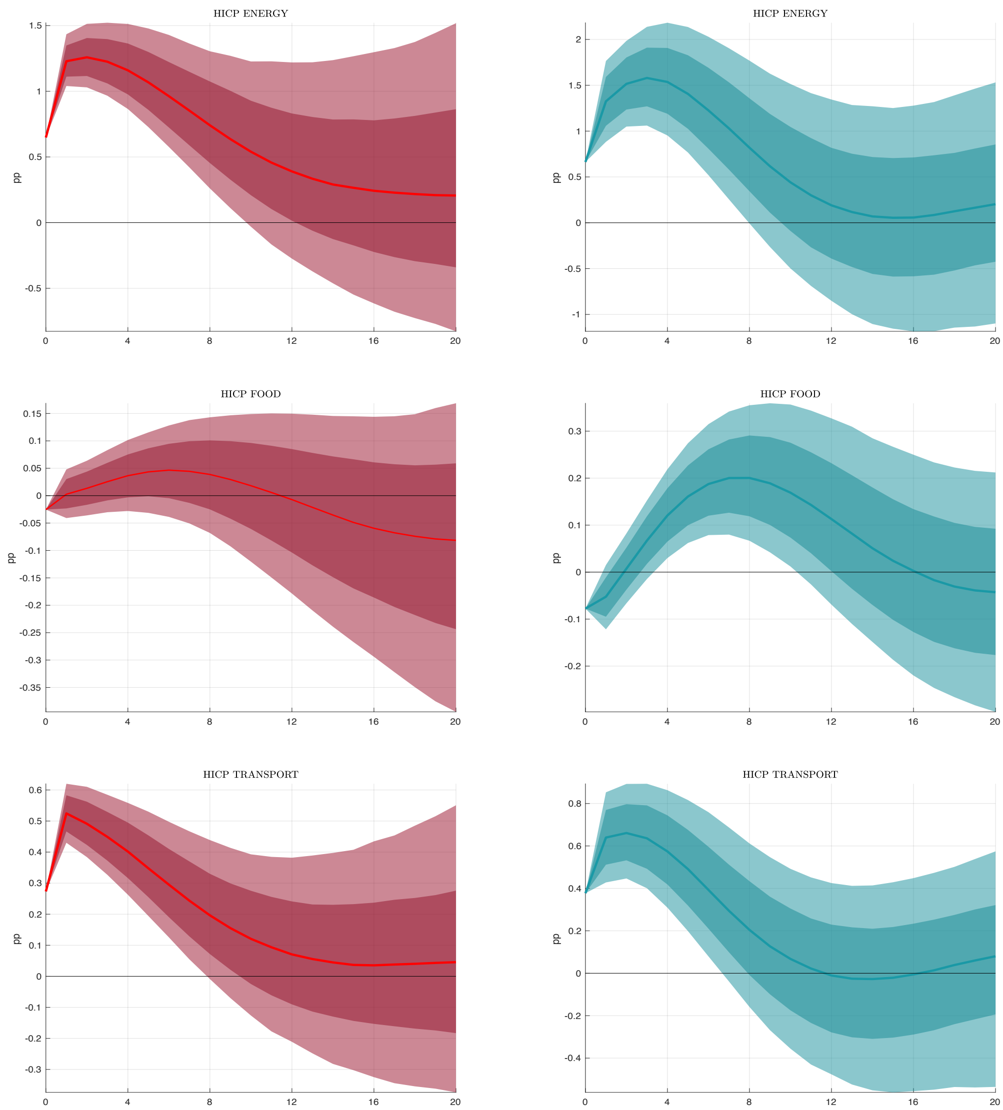


Figure 3.3: Income Effect

according to which the co-movement between food and other commodities, such as crude oil and more broadly industrial commodities, should not be deemed univocal. On the one hand, they tend to move in tandem in response to fluctuations in world economic activity, which raises or depresses demand for inputs according to the different phases of the global business cycle. On the other hand - as also documented by Blomberg and Harris (1995) -, differently from industrial commodities, food is subject to major supply disruptions, which are exogenous to the phases of the global business cycle and, in fact, related to climate and natural factors. Consistent with these views, our results show that - during the period 1999-2019 - food prices in the Euro Area are essentially disconnected from oil price shifts driven by precautionary demand for crude oil. As a matter of fact, food prices do react, though with some lag, to world economic activity shocks, in line with the conventional intuition that food prices are sensitive to movements in the global business cycle. Finally, conditional on either shocks, the response of transport consumer prices soundly increase for approximately two years, in line with the high energy-intensive feature of the transport sector.

Overall, by inspecting different components of HICP, which account on average for more than 50% of consumers' bills, we provide evidence of an *income effect*, which reduces the net disposable income of households. In particular, the results show that, in response to an increase in oil prices due to a demand shock, it is mainly energy and transport items weighting down the bill, while food-related items appear to be not affected. Shifts in economic activity, instead, move all three components in tandem, though food prices react with one year lag.

***Uncertainty effect.*** Figure 3.4 plots the response of consumption of durable goods. Here we are interested in testing the existence of an *uncertainty* transmission channel that discourages agents to committing to *irreversible* investments, such as the purchase of durable goods. The left-hand side chart shows a quite persistent dip in consumption for durables, on top of the *income effect*. This is consistent with the definition of the oil-specific demand shock itself. That is, concerns about the future path of oil supplies deters agents to invest in durables. The evidence of an *uncertainty effect* following an economic activity shock is less evident. There is some temporary decrease in durable goods in the 2-3 year time window after the shock, but the surrounding credibility sets are large.

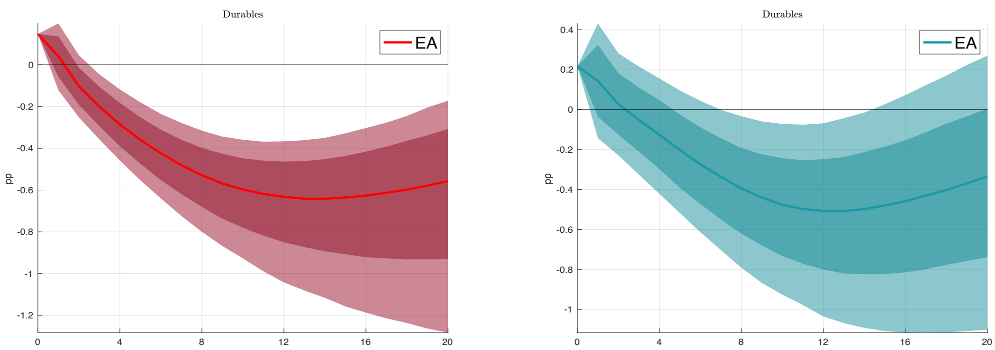


Figure 3.4: Uncertainty Effect

An important school of thought argues that the motor-vehicle industry is one of the major channels whereby energy price shifts propagate throughout the economy. Bresnahan and Ramey (1993), for instance, document that an unexpected increase in gasoline prices in the US unleashes a change in the composition of demand in the automotive sector which rewards more fuel-efficient car manufacturers at the expenses of less fuel-efficient ones. In particular, this has the consequences of promoting import of small-size vehicles from Europe and Asia, with a sharp drop in the demand of domestically produced large-size cars. Along the same lines, a more recent study by Ramey and Vine (2011) confirms previous findings and support the intuition of the motor-vehicle industry as one of the major channels of propagation in response to an energy price surge.

We are, therefore, interested in testing the role of the automotive sector as one of the main conductors of the *uncertainty channel* in the Euro Area, as well. Unfortunately, data on motor-vehicles consumption are not available. However, the Eurostat database on consumer sentiment provides us with a survey conducted by the European Commission in which interviewers ask households if they have the intention to buy a new car in the next 12 months. The idea is to capture whether households feel confident enough about fundamentals to commit to *irreversible* investments. In addition to the consumer sentiment on car purchases, figure 3.5 also plots the impulse response of the survey asking the willingness to buy a new house over the next 12 months.[14] None of the two consumer surveys reacts significantly. Regarding the poor response of the consumer sentiment on new car purchases, this may be due to the fact that, differently from the US, there is no reallocation of demand towards more fuel-efficient means of transport. Unfortunately, the data at disposal do not allow us to further investigate this hypothesis. Regarding the intention to buy a new house, Since early 2000s most of Euro Area countries have gone through steady increase of house prices, which might indeed suggest some excess of confidence regarding the real estate market. However, the results find inconclusive evidence of a change in agents' intentions of purchasing a new car or house in the upcoming future in the Euro Area, conditional on the two global shocks. This excludes the automotive and real estate sectors as main propagators of the *uncertainty* channel.

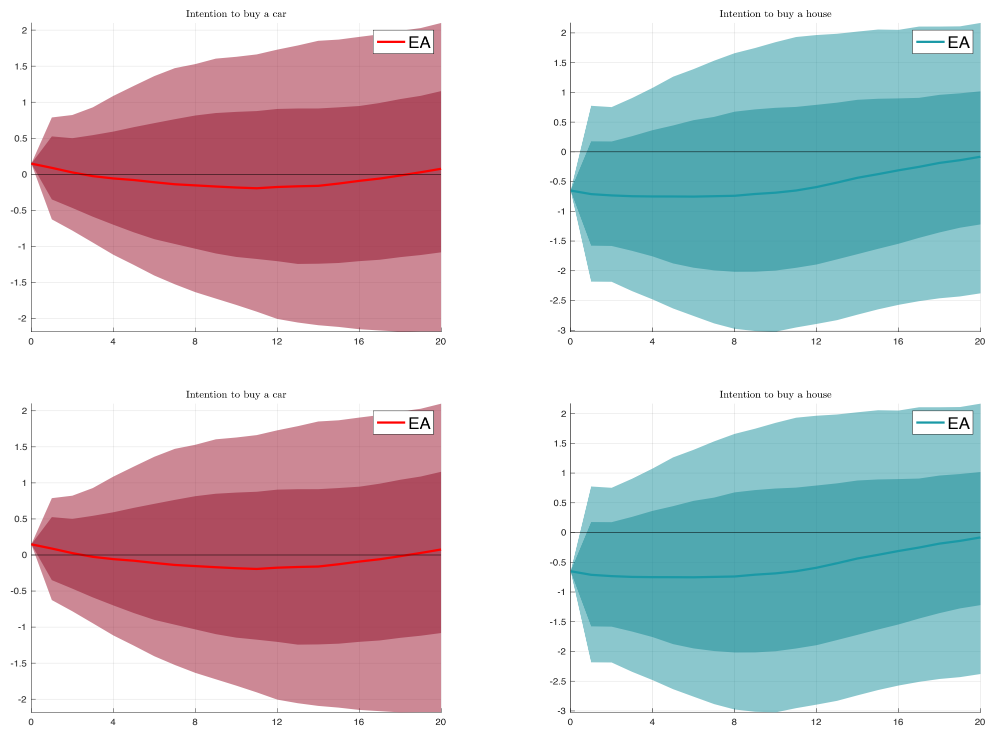


Figure 3.5: The automotive Sector and the uncertainty effect

[14]The questions asked to consumers in the survey are: (i) "How likely are you to buy a car over the next 12 months"; (ii) "Are you planning to buy or build a home over the next 12 months?".

***Precautionary Savings effect.*** In his seminal paper, Bernanke (1983) argues that there is an additional channel whereby a surge in energy prices might generate recessionary effects, that is, precautionary savings. The intuition is the following: if, in the wake of higher commodity prices, households believe that firms will eventually re-optimise their input mix, then they will revise the probability of losing their occupation upward. This, in turn, should encourage households to raise current savings in order to ensure themselves against the prospective of future labor income losses. Unfortunately, it is not a trivial task to retrieve a measure of savings rates. However, the survey on Consumers Sentiment released by the European Commission comes to help once again. More precisely, the survey contains answers to questions about the sentiment of households on savings and unemployment over the next twelve months, which can be used as proxies for our purpose.[15] Figure 3.6 exhibits the responses of consumers sentiment. In response to an oil demand shock, consumers expect an erosion of their savings in the following months and a worsening of the labor market conditions for at least three years after the shock. This result suggests that Euro Area consumers are indeed concerned about the future labor market outlook, as predicted by Bernanke (1983) for the US consumers, but - in the light of a worsening of the labor market - they expect to save less. Although data provide evidence in favour of a revision of households expectations about future labor market outcomes downward, we do not find evidence of precautionary savings. If anything, the response of households shows that they will actually draw from savings to smooth consumption over time. Moving to the response of households expectations following an economic activity shock, there is no evidence of a change in households savings expectations. The response is, indeed, not statistically significant. More action can be observed, instead, when looking at the response of households perception about future labor market conditions. Despite the transient and limited effect on current unemployment rate - see figure 3.1 -, households become more confident about the labor market for the first year after the shock, consistent with the initial improvement in output, consumption and investment. However, in the face of the inflationary effects, they start revising their beliefs as the economic conditions slowdown and current unemployment starts rising.

Overall, this section explores an array of different transmission mechanisms whereby global shocks rising commodity prices slowdown domestic economic growth and potentially trigger recessions. The following facts emerge.

First, from the analysis of GDP components, in response of an oil-specific demand shock, the Euro Area immediately enters into a recession, which is mainly driven by a persistent fall in aggregate consumption and investment. This is caused by: (i) the *in-*

[15] The questions asked to consumers in the survey are: (i) "Over the next 12 months, how likely is it that you save any money?"; (ii) "How do you expect the number of people unemployed in this country to change over the next 12 months?".

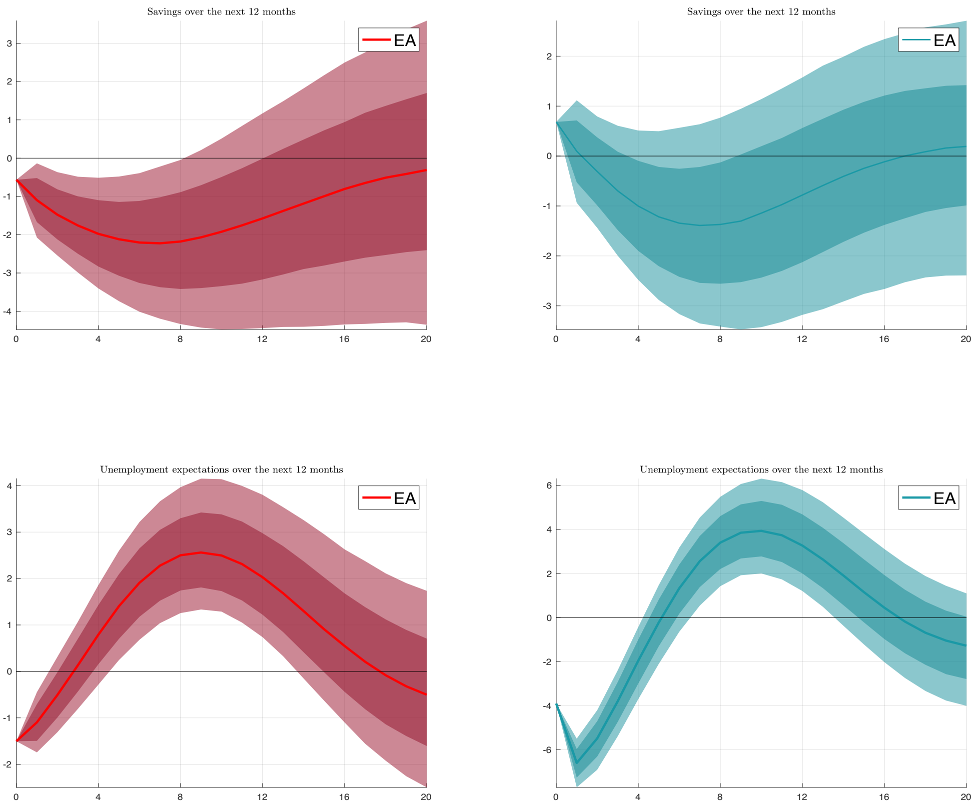


Figure 3.6: Precautionary Savings

*come* channel, which erodes disposable income, due to heavier energy-related bills;[16] (ii) the *uncertainty* channel, which makes households reluctant to commit to *irreversible* investments and substantially reduce their consumption of durable goods. Finally, there is no clear evidence of a *precautionary savings* channel, though - consistent with this argument - households expect a deterioration of the labor market conditions in the following 2-3 years after the shock.

Second, the world economic activity shock leads to a temporary improvement of fundamentals, which reverse after approximately one-two years, unveiling a slowdown in economic growth that is mainly driven by lower investment and higher unemployment. In this case, however, the fall of output and consumption is less pronounced, especially in comparison to the oil demand shock. Results show milder evidence of the presence of the demand channels under scrutiny. Similarly to the oil demand shock, there is compelling evidence in favour of the *income* channel. Conversely, the *uncertainty* channel hypothesis is much less clear than in response to an oil demand shock. Finally, there is no evidence whatsoever of a *precautionary savings* channel.

[16] i.e.: energy and transport components of HICP.

## 3.5 Cost Channel & Second-Round Effects

In this section, we explore the potential indirect effects of energy price increases. Despite the energy component of the Euro Area is rather small, the large effects on the macroeconomy suggest that there must be an endogenous transmission mechanism which passes from higher energy prices all the way through output (and its components) and inflation. This additional source of propagation may materialise in terms of *cost channel* effects on the *supply side* of the economy. The intuition is that an increase in energy prices raises the costs of domestic firms and, depending on their market power, the cost pressure can be partially or entirely transferred to final consumers, potentially leading to *persistently* higher inflation. Hence, testing the existence of a cost channel is particularly informative for monetary policy, because persistently higher prices for consumers may dis-anchor inflation expectations, pushing them to demand a re-negotiation of wages - the so-called *second-round effects* - triggering, in turn, *wage-price spirals*. Identifying such propagation mechanisms in advance is extremely relevant for monetary policy, as it may avoid interventions that require a high sacrifice in terms of employment to cool down inflation.

The most common measure to gain *prima facie* evidence on persistent effects on prices is to check the conditional response of core inflation, namely the measure of price changes excluding the most volatile components (energy and food). Core inflation embeds the evolution of prices of goods and services and it reflects the medium-run trajectory of inflation. Hence, if the inflationary pressure is there to stay, then it will transmit from the (energy) short-run component all the way to the (core) medium-run component. This implies that, conditional to a persistent inflationary shock, core and energy inflation move together, with energy leading core. Conversely, if inflationary pressures do reflect only idiosyncratic movements in the highly volatile energy components, then core inflation will *not* move in tandem with energy inflation.

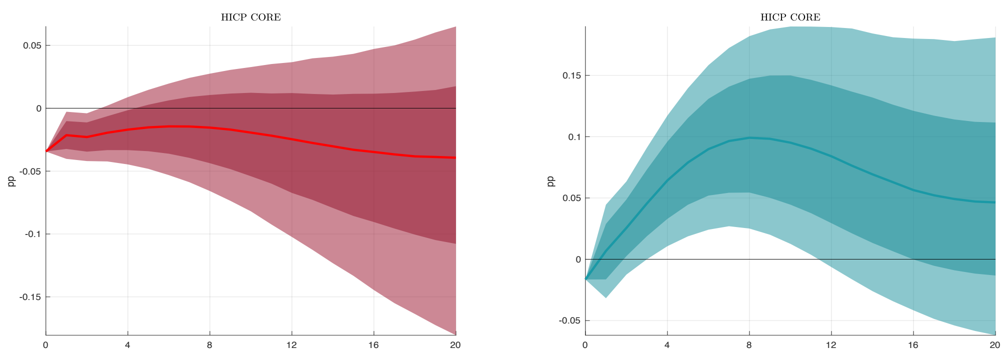


Figure 3.7: Core Inflation

Figure 3.7 plots the response of core inflation conditional on the oil demand shock (red) and the economic activity shock (green). There is no evidence of a significant response of core inflation conditional to an oil demand shock. As a matter of fact, most of the credibility set and the median are negative, implying, if anything, a negative conditional correlation between core inflation and energy inflation (see figure 3.3). This means that core inflation, intended as the measure embedding the medium-term trajectory of price dynamics, is making room to the temporary increase in the more volatile components, by incorporating agents anchored expectations. One possible, though not unique, explanation for this is given by Rostagno et al. (2019), who attribute the negative conditional correlation between core and energy inflation - in particular during the pre 2009 crisis period - to the well-functioning *self-stabilising* expectations mechanism implied by the 2% asymmetric target of the ECB. The intuition is the following: if the ceiling is *credible*, then agents expect that any deviations above it will trigger a prompt and aggressive response of monetary policy. This *self-stabilising* mechanism should, in turn, make the economy more resilient to inflationary shocks. However, given our empirical framework, it is worth inspecting more deeply the cost channel effects to understand if instead core inflation is dormant because of the genuine temporary features of the oil demand shock. In contrast to the negligible response to oil demand, the response of core inflation to the economic activity shock is significant and take two years to reach the peak. In this case, core and energy inflation are positively correlated, with energy leading core. Overall, figure 3.7 suggests that a shift in oil prices induced by world economic activity is a likely candidate of *persistent* inflationary pressures.

At this point, it is interesting to see if the *cost channel* hypothesis is confirmed by the data. There is an extensive literature that tries to shed light on how energy price spikes pass-through to the firms cost structure and translate into persistently higher inflation while depressing economic performance (i.e.: *stagflation*). Rotemberg and Woodford (1996), for instance, try to rationalise this in a neoclassical model with imperfect competition and separability between energy inputs and the capital-labor mix in the production function. Under these assumptions, they show that higher energy prices push firms to re-optimise the cost structure, as firms apply the mark-up on all inputs, therefore becoming eventually a source of large macroeconomic fluctuations. In a more recent paper, Barsky and Kilian (2004) question the conventional wisdom according to which energy shocks - in particular oil shocks - unambiguously lead to *stagflation* and, therefore, to permanent inflationary effects. Their argument focus on the conditional correlations between two different definition of inflation, namely the CPI and the GDP deflator. The main difference between the two lies upon the fact that the GDP deflator measures the price of the *domestic value-added*, therefore it excludes import prices and, so, energy prices by definition. Conversely, the CPI index includes the prices final consumers directly pay for a given good or service, including the imported ones. It is a *gross* measure, indeed. Given these differences, it is not clear a priori if GDP deflator should increase in the wake of higher energy prices.[17] The sign of the co-movement between GDP deflator and CPI should, indeed, suggest if inflation is likely to be a concern in the medium-term or not. More specifically, if they positively co-move, it means that energy price pressure does not remain bounded in the energy component of firms costs, but it passes-through to other inputs, thereby generating persistent effects on inflation. In contrast, if GDP deflator does not move in tandem with CPI, this means that the energy price pressure remains bounded in the share of energy-related components of firms costs. Therefore, the threaten of an upside risk for the medium-term inflation outlook should be deemed low and not subject to any monetary policy reactions.

The responses of GDP and import price deflator are displayed in figure 3.8. Conditional to an oil-specific demand shock, the response is negative and persistent. In this case, not only GDP deflator does not co-move with CPI, but it actually exhibits an opposite sign response. The reason for this lies on the definition of GDP deflator itself. Intuitively, in the absence of a *cost channel* on the firms side, if there is evidence of a direct *income channel* - eroding the disposable income on the household side -, and/or an *uncertainty channel* effect - discouraging both firms and households to commit to irreversible investments -, then not only the GDP deflator will not be positively correlated with CPI, but it will actually be negatively correlated, as in this specific case, reflecting the recessionary consequences of the oil demand shock. Moving to the response to an economic activity shock (top-right panel), the response of GDP deflator is positive and statistically significant for almost three years after the shock, implying a positive conditional correlation with HICP in figure 3.1. This, in turn, suggests that, in the presence of a upward shift of energy prices induced by a world economic activity shock, there exist a *cost channel* on the side of firms which passes-through the entire cost structure of domestic firms.

Overall, these results contribute in building additional evidence on: (i) a concerning persistent inflationary pressure of economic activity shocks; (ii) a limited effect on the medium-term inflation outlook, but a strong and immediate recessionary effect of oil-specific demand shocks.

Consistent with the aim of studying the indirect *cost channel* transmission mechanism, we proceed by inspecting the impulse responses of the domestic producer price. The advantage is twofold: (i) it measures the price received by firms for *domestically produced* goods, (ii) which are dispatched to the *domestic* market. Narrowing to the basket of domestically produced goods, allows us to test the existence of an indirect effect of energy prices on domestic firms costs. In addition, further narrowing to those goods that will be only delivered to the domestic market, enables us to avoid any potential blurring effects arising from the price received by domestic producers for the portion of production

[17] These considerations fall short if one considers net energy exporting countries. That is, taks the case an oil net exporting country. An increase in oil prices will actually boost domestic value-added, eliminating any ambiguity.

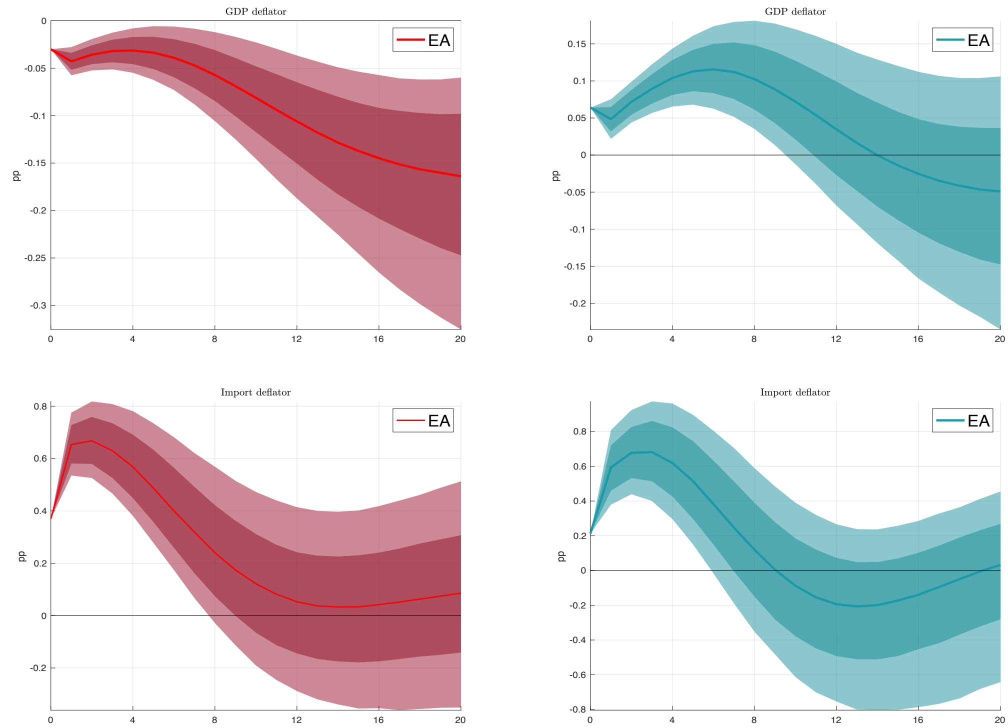


Figure 3.8: Cost channel: GDP and import deflators

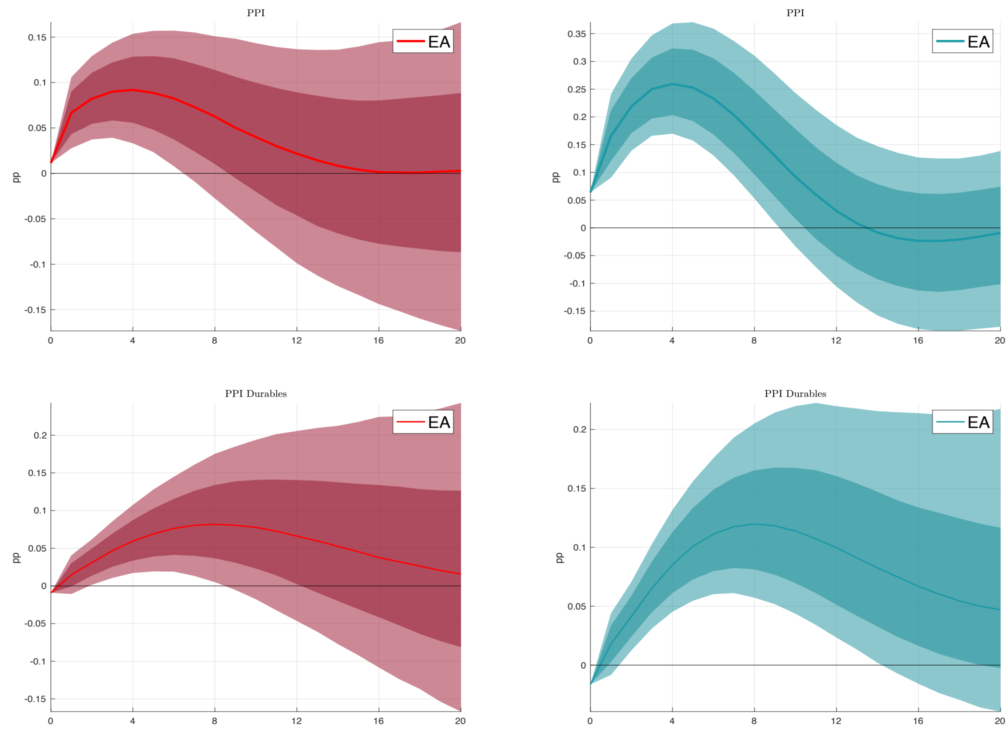


Figure 3.9: Cost channel: producer price index

allocated abroad.[18] Yet, given the availability of more disaggregated producer price data,

[18] In this case, the idea is that the law of one price does not hold and firms can discriminate on prices according to the destination of their output.

it is possible to check which sectors are able to transfer higher costs to final consumers. Figure 3.9 shows the response of overall PPI and PPI of durable goods. The top panels exhibit the response of overall PPI. Though sharing the same sign, the response to the economic activity shock is more pronounced and persistent. The response to the oil demand is instead limited and short-lived. This highlights again the strong recessionary component of the oil-specific demand shock, which prevents firms to discharge higher costs to final consumers. It is more interesting, however, analysing the different response of producer prices of durable goods across the two shocks. Let us start with the response to the oil demand shock in the bottom-left chart. The PPI of durables poorly react to the oil-specific demand shock, implying that the higher energy costs are not transferred to final consumers. This is consistent with a highly elastic demand for durable goods conditional to oil-specific demand shocks. The drastic fall in consumption of durables in figure 3.4 suggests, indeed, that firms producing durable goods will have hard times passing higher energy costs to final consumers. Conversely, in response to the economic activity shock (bottom-right panel), PPI of durables mimics the hump-shape response of core inflation and also shares the timing of the peak (i.e.: two years). The previous results provided compelling evidence of an indirect *cost channel* effects stemming from economic activity shocks and that the demand for durables is much less elastic than the case of an oil-specific demand shock hitting the economy. All together, this implies that firms producing durable goods are able to transfer higher costs of production to final consumers without facing a dramatic drop in demand, thereby contributing to the build-up of core inflation.

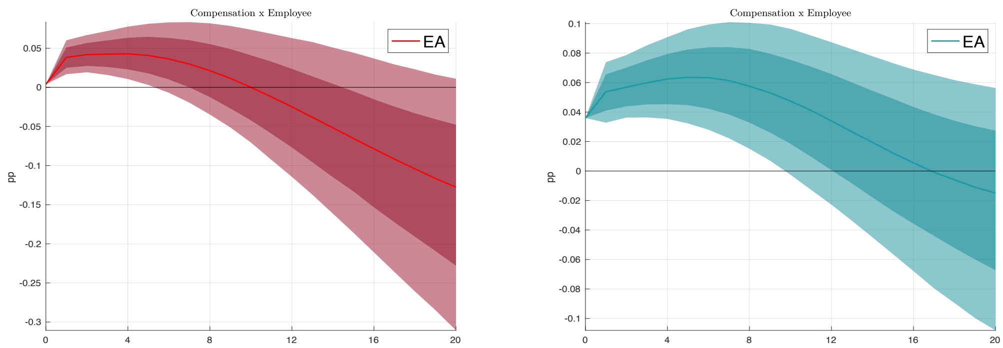


Figure 3.10: Second-round effects: compensation per employee

Finally, a comprehensive analysis of the potentially persistent effects on inflation also requires to look through the lens of the labor market. Wage dynamics, for instance, represent a good metric of how firmly anchored are expectations, in the first place. If, for example, following an inflationary perturbation, workers ask for a re-negotiation of their salaries, this should act as a warning bell for the policy-maker. By asking for an adjustment in their compensations, workers are signalling their concerns about inflation eroding their purchasing power. In other words, inflation expectations are *dis-anchoring*.

Then, conditional on the relative bargaining power of workers in the wage negotiation process, this may unleash a vicious spiral between wages and prices, thereby amplifying the inflationary pressures over time at the detriment of the medium-term inflation outlook.

The top panels of figure 3.10 display the response of compensation per employee to the oil demand and the economic activity shock, respectively. the oil demand shock has zero impact effect and only a very transient effect thereafter, which lasts less than a year and is limited in its magnitude. In response to an economic activity shock, instead, compensation per employee raises immediately and significantly so for approximately three years, though the magnitude effect is yet relatively limited. Overall, by looking at compensation per employee, we do not find overwhelming evidence in favour of strong adjustment in wages. Though being statistically significant and rather persistent, the response to the economic activity shock is limited in terms of economic magnitude.[19] It should be stressed, however, that this specific framework analyses one-time shocks, while in reality the economy is usually hit by sequence of shocks. This, in turn, would strengthen the evidence in favour of wage pressures.

To conclude, in this section we test the existence of a supply side amplification mechanism potentially leading to persistent effects on inflation, threatening the medium-term target of the Central Bank. The results can be summarised as follows.
First, there is no evidence of an indirect *cost channel* feeding into firms costs, when the economy is hit by an oil demand shock. As a matter of fact, core inflation does not react significantly and, if anything, it suggests that price pressure will ease in the medium-term. In addition, the conditional correlation of GDP deflator and energy HICP is negative, suggesting that the price pressure remains bounded to the energy component of firms costs and does not propagate to the other inputs. Yet, the response of PPI is short-lived and limited, corroborating even further the strong recessionary consequences for households and the high elasticity of demand for durables conditional on oil-specific demand shocks. Finally, there is no evidence of wage inflationary pressures as result of a workers adjusting their inflation expectations upward. Second, we find compelling evidence of persistent inflationary pressures in response to a shift in global energy prices induced by an improvement of the world outlook. GDP deflator and energy HICP exhibit positive conditional correlation, documenting the existence of an indirect cost mechanism channeling higher energy prices into the whole cost structure of domestic firms. Also, core inflation and PPI for durables strongly move in tandem, suggesting that firms producing durables are able to discharge the cost burden into final consumers via selling prices. Finally, though not overwhelming, there is supporting evidence of potential wage-price spirals unleashed by workers revising expectations upward, materialised in wage inflationary pressures.

[19]See Conti and Nobili (2019) for a recent discussion on potential sources of a weak wage-price nexus in the Euro Area over the last twenty years.

## 3.6 101 Economic Intuition

This section summarises and rationalises the previous results in a simple 101 macroeconomic AS-AD space. The idea is to convey a simple reasoning of why monetary policy should: (i) *look through* oil-specific demand shocks; (ii) be concerned of a raise in oil and energy prices induced by global economic activity shocks.

Let us start with the case of an oil-specific demand shock hitting the world economy. The left panel in figure 3.11 shows that an increase in oil prices induced by precautionary demand produces an outward shift of the demand curve *OD* in the global market for oil. As observed in the section 3.3, the raise in oil prices generates a fall in real output and an increase in consumer prices, leading, in turn, to a leftward shift of the aggregate supply curve in the right-hand side panel of figure 3.11 - moving from equilibrium 0 to 1. Furthermore, section 3.2 also provides compelling evidence of a strong recessionary propagation mechanism via (i) a *direct income channel* on households disposable income and (ii) an *uncertainty channel* - discouraging in particular demand for durable goods -, which act as an aggregate demand shifter cooling down the inflationary pressures. This automatic "stabilisation mechanism" makes the acceleration of prices to be *temporary*, but with the side effect of a *permanent* fall in output and income, as we move to intersection 2. Thus, from the point of view of a policy-maker whose ultimate goal is to stabilise inflation, she should not be much "concerned" about an oil demand shock.

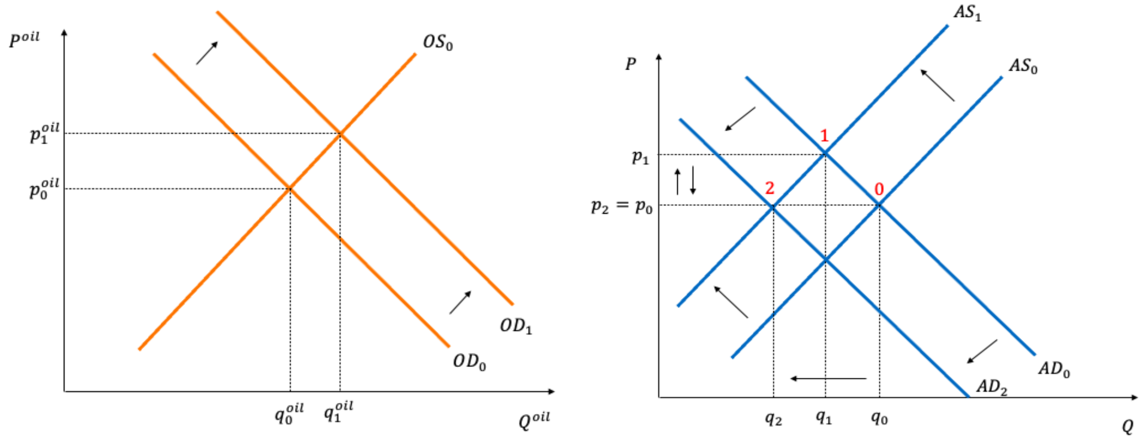


Figure 3.11: AS-AD: oil-specific demand shock

Different considerations arise in the case of an oil price surge associated with a world economic activity shock. Let us be backed again by the AS-AD diagram in figure 3.12. An increase in oil prices driven by the overheating of the world economic activity shifts the oil demand OD to the right once again. This time, however, the raise in oil prices translates into a temporary increase in output and inflation, as shown in section 3.3, leading in turn to a rightward shift of the AD curve along the AS, moving from point 0 to 1. The improvement in fundamentals, however, is short-lived. This is due to a *cost* *channel* propagation mechanism, which act on the supple side of the economy and yields a leftward shift of the aggregate supply curve, ending up in equilibrium 2. In this case, the world economic activity shock generates only a *temporary* improvement in output and income, but a *permanent* effect on prices, potentially dis-anchoring inflation expectations and unleashing wage-price spirals. It is, therefore, a call for duty for monetary policy to bring inflation back onto sustainable boundaries by moving the AD curve to the left. Clearly, this will push the economy into a recession, but it is important to stress that the amplitude of the subsequent recession crucially depends on the policy-maker ability to promptly react. Therefore, the earlier the intervention, the smaller the *sacrifice* required to the economy in terms of employment to bring inflation to a halt.

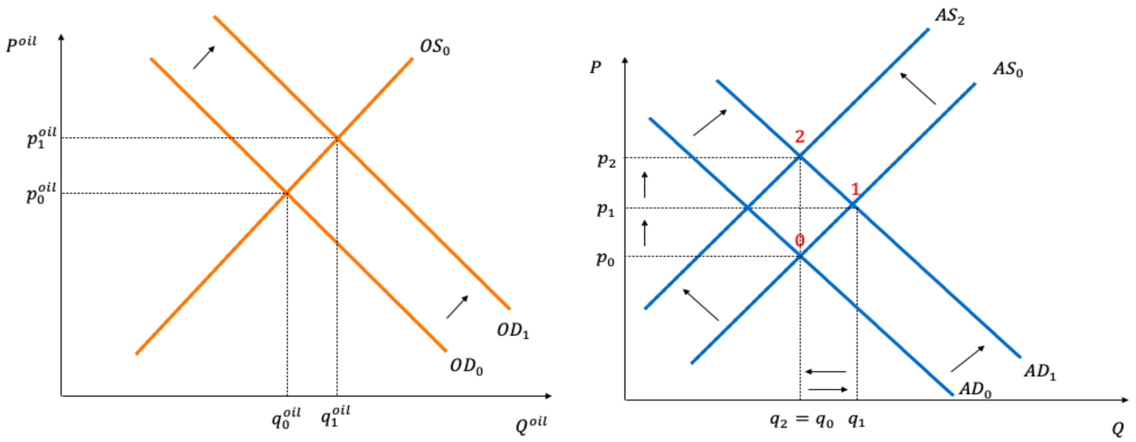


Figure 3.12: AS-AD: global economic activity shock

## 3.7 Conclusions

The latest energy price swings in the wake of the post-pandemic economic recovery and the rising geopolitical uncertainty about the Russian-Ukrainian conflict have brought about a revival of the discussion in policy and academic circles about the consequences of global oil price surges for the inflation outlook and the performance of domestic economies. In this paper we study the implications for the Euro Area and contribute to the discussion in the following dimensions: (i) designing a simple empirical laboratory yet based on the state-of-the-art identification schemes to study the implications of an increase in global oil prices; (ii) highlighting the importance of the underlying sources of oil price fluctuations to unveil the consequences for economic growth and inflation; (iii) inspecting the transmission mechanisms whereby the global shocks propagate throughout the economy. For this purpose, we estimate a recursive VAR augmented with the structural (global) shock series identified by Baumeister and Hamilton (2019). This relatively simple approach provides us with a tractable and ready-to-use tool to analyse the different implications and transmission mechanisms of these two shocks. Consistent with the ongoing policy discussion, the rationale is, therefore, to understand to what extent monetary policy should care about shifts in oil prices due to (i) *precautionary demand* induced by concerns about future supply of oil - i.e.: oil demand shock - and (ii) an overheating of the global business cycle. The results are summarised as follows.

First, oil price swings driven by oil-specific demand are responsible for *temporary* inflationary pressures and the immediate and *persistent* worsening of economic growth. This is due to the fact that oil-specific demand shocks embed a strong recessionary component, which materialises with the erosion of households disposable income - due to heavier energy bills - and a sharp drop in consumption for durable goods - related to higher uncertainty about future energy prices.

Second, a rise in oil prices associated with (global) economic activity shocks produces a *temporary* improvement of fundamentals and *persistent* inflationary pressures, eventually leading to a slow-down in domestic economic activity. The persistent effect on domestic inflation is set in motion by a forceful *cost channel* transmission mechanism affecting domestic firms' costs, which are eventually discharged into final consumer prices, potentially de-anchoring agents' expectations and triggering wage-price spirals.

Overall, the results suggest that not all oil price spikes are alike. On the one hand, the recessionary effect of the oil-specific demand shock acts as an "automatic machanism" that dampens inflationary pressures. On the other hand, the *cost channel* effect, following a global economic activity shock, transmits to domestic firms' costs and amplifies the inflationary impulse, endangering the medium-term inflation outlook. Consistent with our results, it is crucial - from the perspective of the monetary authority - to promptly identify which source of fluctuations is behind energy prices swings in order to avoid harsh unemployment-inflation trade-offs to converge back to the target.

# Bibliography


D. Acemoglu and P. Restrepo. Robots and Jobs: Evidence from US Labor Markets. Journal of Political Economy, 128(6):2188–2244, 2020. doi: 10.1086/705716. URL https://ideas.repec.org/a/ucp/jpolec/doi10.1086-705716.html.

D. Acemoglu, D. H. Autor, and D. Lyle. Women, War, and Wages: The Effect of Female Labor Supply on the Wage Structure at Midcentury. Journal of Political Economy, 112(3):497–551, June 2004. doi: 10.1086/383100. URL https://ideas.repec.org/a/ucp/jpolec/v112y2004i3p497-551.html.

S. Albanesi. Changing business cycles: The role of women's employment. Technical report, National Bureau of Economic Research, 2019.

S. Albanesi and C. Olivetti. Gender roles and medical progress. Journal of Political Economy, 124(3):650–695, 2016.

G.-M. Angeletos, F. Collard, and H. Dellas. Business-cycle anatomy. American Economic Review, 110(10):3030–70, October 2020.

G. Ascari and L. Fosso. The inflation rate disconnect puzzle: On the international component of trend inflation and the flattening of the phillips curve. Working paper no. 733, De Nederlandsche Bank, Dec. 2021.

O. Attanasio, H. Low, and V. Sánchez-Marcos. Explaining changes in female labor supply in a life-cycle model. American Economic Review, 98(4):1517–52, 2008.

R. Auer and A. M. Fischer. The effect of low-wage import competition on u.s. inflationary pressure. Journal of Monetary Economics, 57(4):491–503, 2010. ISSN 0304-3932. doi: https://doi.org/10.1016/j.jmoneco.2010.02.007. URL https://www.sciencedirect.com/science/article/pii/S030439321000019X.

R. Auer, C. Borio, and A. J. Filardo. The Globalisation of Inflation: the Growing Importance of Global Value Chains. Globalization institute working papers, Federal Reserve Bank of Dallas, Jan. 2017.

R. A. Auer, A. A. Levchenko, and P. Sauré. International inflation spillovers through input linkages. The Review of Economics and Statistics, 101(3):507–521, 2019.

D. Autor and A. Salomons. Is Automation Labor-Displacing? Productivity Growth, Employment, and the Labor Share. NBER Working Papers 24871, National Bureau of Economic Research, Inc, July 2018. URL https://ideas.repec.org/p/nbr/nberwo/24871.html.

D. H. Autor, D. Dorn, and G. H. Hanson. The china shock: Learning from labor-market adjustment to large changes in trade. Annual Review of Economics, 8(1):205–240, 2016. doi: 10.1146/annurev-economics-080315-015041. URL https://doi.org/10.1146/annurev-economics-080315-015041.

L. Bai and S. Stumpner. Estimating us consumer gains from chinese imports. American Economic Review: Insights, 1(2):209–24, September 2019. doi: 10.1257/aeri.20180358. URL https://www.aeaweb.org/articles?id=10.1257/aeri.20180358.

R. Barnichon and G. Mesters. Identifying Modern Macro Equations with Old Shocks*. The Quarterly Journal of Economics, 135(4):2255–2298, 06 2020. ISSN 0033-5533. doi: 10.1093/qje/qjaa022. URL https://doi.org/10.1093/qje/qjaa022.

R. B. Barsky and L. Kilian. Oil and the macroeconomy since the 1970s. The Journal of Economic Perspectives, 18(4):115–134, 2004. ISSN 08953309. URL http://www.jstor.org/stable/3216795.

C. Baumeister and J. D. Hamilton. Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks. American Economic Review, 109(5):1873–1910, May 2019. doi: 10.1257/aer.20151569. URL https://www.aeaweb.org/articles?id=10.1257/aer.20151569.

C. Baumeister, G. Peersman, I. Van Robays, and . The economic consequences of oil shocks: Differences across countries and time. Working papers of faculty of economics and business administration, ghent university, belgium, Ghent University, Faculty of Economics and Business Administration, 2009. URL https://EconPapers.repec.org/RePEc:rug:rugwps:09/630.

C. Baumeister, G. Peersman, and I. Van Robays. The economic consequences of oil shocks: Differences across countries and time. In R. Fry, C. Jones, and C. Kent, editors, Inflation in an Era of Relative Price Shocks. Reserve Bank of Australia, 2010. URL https://EconPapers.repec.org/RePEc:rba:rbaacv:acv2009-06.

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. Furlanetto, and N. M. Faccioli. The decline of the labor share: new empirical evidence. Working Paper 2019/18, Norges Bank, Oct. 2019. URL https://ideas.repec.org/p/bno/worpap/2019_18.html.

D. Bergholt, F. Furlanetto, and E. Vaccaro-Grange. The death and resurrection of the us price phillips curve. Manuscript, 2020.

B. S. Bernanke. Irreversibility, uncertainty, and cyclical investment. The Quarterly Journal of Economics, 98(1):85–106, 1983. ISSN 00335533, 15314650. URL http://www.jstor.org/stable/1885568.

B. S. Bernanke. Energy and the economy. Speech 221, Board of Governors of the Federal Reserve System (U.S.), 2006. URL https://ideas.repec.org/p/fip/fedgsq/221.html.

B. Blomberg and E. Harris. The commodity–consumer price connection: Fact or fable? Economic Policy Review, 2:21–38, 01 1995.

R. Blundell and T. Macurdy. Labor supply: A review of alternative approaches. In O. Ashenfelter and D. Card, editors, Handbook of Labor Economics, volume 3 of Handbook of Labor Economics, chapter 27, pages 1559–1695. Elsevier, 1999. URL https://ideas.repec.org/h/eee/labchp/3-27.html.

E. Bobeica and M. Jarociński. Missing disinflation and missing inflation: A var perspective. International Journal of Central Banking, 15(1):199–232, 2019. URL https://EconPapers.repec.org/RePEc:ijc:ijcjou:y:2019:q:1:a:5.

T. Boppart and P. Krusell. Labor supply in the past, present, and future: A balanced-growth perspective. Journal of Political Economy, 128(1):118–157, 2020. doi: 10.1086/704071. URL https://doi.org/10.1086/704071.

C. E. Borio and A. J. Filardo. Globalisation and inflation: New cross-country evidence on the global determinants of domestic inflation. 2007.

L. Branstetter and N. R. Lardy. China's Embrace of Globalization, page 633–682. Cambridge University Press, 2008. doi: 10.1017/CBO9780511754234.017.

T. Bresnahan and V. Ramey. Segment shifts and capacity utilization in the u.s. automobile industry. American Economic Review, 83(2):213–18, 1993. URL https://EconPapers.repec.org/RePEc:aea:aecrev:v:83:y:1993:i:2:p:213-18.

F. J. Buera and J. P. Kaboski. The rise of the service economy. American Economic Review, 102(6):2540–69, 2012.

M. Carney. The growing challenges for monetary policy in the current international monetary and financial system, jackson hole symposium 2019, 23 august 2019. Speech, Bank of England, 2019.

A. Carriero, F. Corsello, and M. Marcellino. The Global Component of Inflation Volatility. CEPR Discussion Papers 13470, C.E.P.R. Discussion Papers, Jan. 2019. URL https://ideas.repec.org/p/cpr/ceprdp/13470.html.

C. K. Carter and R. Kohn. On gibbs sampling for state space models. Biometrika, 81 (3):541–553, 1994.

O. Celasun, A. Mineshima, N.-J. Hansen, J. Zhou, and M. Spector. Supply bottlenecks: Where, why, how much, and what next? IMF Working Papers 2022/031, International Monetary Fund, 2022. URL https://EconPapers.repec.org/RePEc:imf:imfwpa:2022/031.

L. J. Christiano, M. S. Eichenbaum, and M. Trabandt. Understanding the great recession. American Economic Journal: Macroeconomics, 7(1):110–67, 2015.

M. Ciccarelli and B. Mojon. Global inflation. The Review of Economics and Statistics, 92(3):524–535, 2010.

R. H. Coase and N. Wang. Palgrave Macmillan, Basingstoke, 2012. ISBN 978-1-137-01936-3.

O. Coibion and Y. Gorodnichenko. What can survey forecasts tell us about information rigidities? Journal of Political Economy, 120(1):116–159, 2012.

O. Coibion and Y. Gorodnichenko. Is the phillips curve alive and well after all? inflation expectations and the missing disinflation. American Economic Journal: Macroeconomics, 7(1):197–232, January 2015a. doi: 10.1257/mac.20130306. URL https://www.aeaweb.org/articles?id=10.1257/mac.20130306.

O. Coibion and Y. Gorodnichenko. Information rigidity and the expectations formation process: A simple framework and new facts. American Economic Review, 105(8): 2644–78, 2015b.

O. Coibion, Y. Gorodnichenko, and R. Kamdar. The formation of expectations, inflation, and the phillips curve. Journal of Economic Literature, 56(4):1447–91, December 2018.

A. Conti and A. Nobili. Wages and prices in the euro area: exploring the nexus. Questioni di Economia e Finanza (Occasional Papers) 518, Bank of Italy, Economic Research and International Relations Area, 2019. URL https://EconPapers.repec.org/RePEc:bdi:opques:qef_518_19.

A. M. Conti, A. Nobili, et al. Wages and prices in the euro area: exploring the nexus. Technical report, Bank of Italy, Economic Research and International Relations Area, 2019.

M. Del Negro, M. P. Giannoni, and F. Schorfheide. Inflation in the great recession and new keynesian models. American Economic Journal: Macroeconomics, 7(1):168–96, January 2015. doi: 10.1257/mac.20140097. URL https://www.aeaweb.org/articles?id=10.1257/mac.20140097.

M. Del Negro, D. Giannone, M. Giannoni, and T. Andrea. Safety, liquidity, and the natural rate of interest. Staff Reports 812, Federal Reserve Bank of New York, May 2017a. URL https://ideas.repec.org/p/fip/fednsr/812.html.

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

M. Del Negro, M. Lenza, G. E. Primiceri, and A. Tambalotti. What's up with the phillips curve? Brookings Paper on Economic Activity, Spring:301–373, 2020.

J. J. Dolado, G. Motyovszki, and E. Pappa. Monetary policy and inequality under labor market frictions and capital-skill complementarity. American economic journal: macroeconomics, 13(2):292–332, 2021.

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

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

P. Edelstein and L. Kilian. How sensitive are consumer expenditures to retail energy prices? Journal of Monetary Economics, 56(6):766–779, September 2009. URL https://ideas.repec.org/a/eee/moneco/v56y2009i6p766-779.html.

Y. Eo, L. Uzeda, and B. Wong. Understanding Trend Inflation Through the Lens of the Goods and Services Sectors. Staff Working Papers 20-45, Bank of Canada, Nov. 2020.

R. Fernández, A. Fogli, and C. Olivetti. Mothers and sons: Preference formation and female labor force dynamics. The Quarterly Journal of Economics, 119(4):1249–1299, 2004.

K. J. Forbes. Inflation dynamics: Dead, dormant, or determined abroad? Brookings Papers on Economic Activity, Fall:257–319, 2019.

M. Fukui, E. Nakamura, and J. Steinsson. Women, wealth effects, and slow recoveries. American Economic Journal: Macroeconomics, forthcoming.

F. Furlanetto, A. Lepetit, Ørjan Robstad, J. Rubio-Ramírez, and P. Ulvedal. Estimating Hysteresis Effects. Working Papers 2021-11, FEDEA, June 2021. URL https://ideas.repec.org/p/fda/fdaddt/2021-11.html.

O. Galor and D. N. Weil. The gender gap, fertility, and growth. The American Economic Review, pages 374–387, 1996.

J. Galí and L. Gambetti. Has the u.s. wage phillips curve flattened? a semi-structural exploration. Working Paper 25476, National Bureau of Economic Research, January 2019.

E. Gamber and J. Hung. Has the rise in globalization reduced u.s. inflation in the 1990s? Economic Inquiry, 39(1):58–73, 2001. doi: https://doi.org/10.1111/j.1465-7295.2001.tb00050.x. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1465-7295.2001.tb00050.x.

D. Giannone, M. Lenza, and G. E. Primiceri. Prior Selection for Vector Autoregressions. The Review of Economics and Statistics, 97(2):436–451, 05 2015. ISSN 0034-6535. doi: 10.1162/REST_a_00483. URL https://doi.org/10.1162/REST_a_00483.

D. Giannone, M. Lenza, and L. Reichlin. Money, Credit, Monetary Policy, and the Business Cycle in the Euro Area: What Has Changed Since the Crisis? International Journal of Central Banking, 15(5):137–173, December 2019. URL https://ideas.repec.org/a/ijc/ijcjou/y2019q5a4.html.

S. Gilchrist, R. Schoenle, J. Sim, and E. Zakrajšek. Inflation dynamics during the financial crisis. American Economic Review, 107(3):785–823, March 2017.

C. Goldin. The quiet revolution that transformed women's employment, education, and family. American economic review, 96(2):1–21, 2006.

C. Goldin and L. F. Katz. The power of the pill: Oral contraceptives and women's career and marriage decisions. Journal of political Economy, 110(4):730–770, 2002.

M. Goodfriend and R. G. King. The incredible volcker disinflation. Journal of Monetary Economics, 52(5):981–1015, 2005.

A. Greenspan. Globalization: a speech at the council on foreign relations, new york, new york, march 10, 2005. Speech 76, Board of Governors of the Federal Reserve System (U.S.), 2005. URL https://EconPapers.repec.org/RePEc:fip:fedgsq:76.

J. Greenwood, A. Seshadri, and M. Yorukoglu. Engines of liberation. The Review of Economic Studies, 72(1):109–133, 2005.

D. S. Hamermesh. Labor demand. Princeton University Press, Princeton, New Jersey, 1993.

J. Hamilton. A neoclassical model of unemployment and the business cycle. Journal of Political Economy, 96(3):593–617, 1988. URL https://EconPapers.repec.org/RePEc:ucp:jpolec:v:96:y:1988:i:3:p:593-617.

J. Hamilton. Oil and the macroeconomy. The New Palgrave Dictionary of Economics, 09 2005.

T. Hasenzagl, F. Pellegrino, L. Reichlin, and G. Ricco. A model of the fed's view on inflation. Review of Economics and Statistics, 2020. doi: https://doi.org/10.1162/rest_a_00974. forthcoming.

J. Hazell, J. Herreño, E. Nakamura, and J. Steinsson. The slope of the phillips curve: Evidence from u.s. states. Working Paper 28005, National Bureau of Economic Research, October 2020.

J. Heathcote, K. Storesletten, and G. L. Violante. The macroeconomics of the quiet revolution: Understanding the implications of the rise in women's participation for economic growth and inequality. Research in Economics, 71(3):521–539, 2017.

S. Heise, F. Karahan, and A. Şahin. The Missing Inflation Puzzle: The Role of the Wage-Price Pass-Through. Journal of Money, Credit and Banking, 2020. forthcoming.

C.-T. Hsieh, E. Hurst, C. I. Jones, and P. J. Klenow. The allocation of talent and us economic growth. Econometrica, 87(5):1439–1474, 2019.

IMF. World Economic Outlook, April 2006: Globalization and Inflation. International Monetary Fund, USA, 2005. ISBN 9781589065499. doi: 10.5089/9781589065499.081. URL https://www.elibrary.imf.org/view/books/081/07993-9781589065499-en/07993-9781589065499-en-book.xml.

X. Jaravel and E. Sager. What are the price effects of trade? Evidence from the US and implications for quantitative trade models. CEP Discussion Papers dp1642, Centre for Economic Performance, LSE, Aug. 2019. URL https://ideas.repec.org/p/cep/cepdps/dp1642.html.

H. Jeffreys. The theory of probability. OUP Oxford, 1998.

M. Johnson and M. Keane. A dynamic equilibrium model of the us wage structure, 1968-1996. Journal of Labor Economics, 31(1):1 – 49, 2013. URL https://EconPapers.repec.org/RePEc:ucp:jlabec:doi:10.1086/666698.

G. Kamber and B. Wong. Global factors and trend inflation. Journal of International Economics, 122(C), 2020.

M. Keane and R. Rogerson. Micro and macro labor supply elasticities: A reassessment of conventional wisdom. Journal of Economic Literature, 50(2):464–476, 2012. ISSN 00220515. URL http://www.jstor.org/stable/23270026.

L. Kilian. The economic effects of energy price shocks. Journal of Economic Literature, 46 (4):871–909, 2008. ISSN 00220515. URL http://www.jstor.org/stable/27647084.

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.

L. Kilian and D. Murphy. Why agnostic sign restrictions are not enough: Understanding the dynamics of oil market var models. Journal of the European Economic Association, 10(5):1166–1188, 2012a. URL https://EconPapers.repec.org/RePEc:bla:jeurec:v:10:y:2012:i:5:p:1166-1188.

L. Kilian and D. P. Murphy. Why agnostic sign restrictions are not enough: Understanding the dynamics of oil market var models. Journal of the European Economic Association, 10(5):1166–1188, 2012b. ISSN 15424766, 15424774. URL http://www.jstor.org/stable/23251216.

G. Koester, E. Lis, C. Nickel, C. Osbat, and F. Smets. Understanding low inflation in the euro area from 2013 to 2019: cyclical and structural drivers. ECB Strategy Review Occasional Paper Series No. 280, 2021.

J. Lindé and M. Trabandt. Resolving the missing deflation puzzle. 2019.

M. McLeay and S. Tenreyro. Optimal inflation and the identification of the phillips curve. NBER Macroeconomics Annual, 34(1):199–255, 2020.

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

E. Mertens and J. M. Nason. Inflation and professional forecast dynamics: An evaluation of stickiness, persistence, and volatility. Quantitative Economics, 11(4):1485–1520, 2020.

H. Mumtaz and P. Surico. Evolving international inflation dynamics: World and country-specific factors. Journal of the European Economic Association, 10(4):716–734, 2012. ISSN 15424766, 15424774. URL http://www.jstor.org/stable/23251097.

J. M. Nason and G. W. Smith. Measuring the slowly evolving trend in us inflation with professional forecasts. Journal of Applied Econometrics, 2020. doi: https://doi.org/10.1002/jae.2784. Forthcoming.

C. Olivetti and B. Petrongolo. Gender gaps across countries and skills: Demand, supply and the industry structure. Review of Economic Dynamics, 17(4):842–859, October 2014. doi: 10.1016/j.red.2014.03.001. URL https://ideas.repec.org/a/red/issued/12-174.html.

G. Peersman. International food commodity prices and missing (dis)inflation in the euro area. The Review of Economics and Statistics, 104(1):85–100, 2022. URL https://EconPapers.repec.org/RePEc:tpr:restat:v:104:y:2022:i:1:p:85-100.

G. Peersman and I. V. Robays. Oil and the euro area economy. Economic Policy, 24 (60):603–651, 2009. ISSN 02664658, 14680327. URL http://www.jstor.org/stable/40272533.

M. Plagborg-Møller and C. K. Wolf. Local projections and vars estimate the same impulse responses. Econometrica, 89(2):955–980, 2021. URL https://doi.org/10.3982/ECTA17813.

V. Ramey and D. J. Vine. Oil, automobiles, and the us economy: How much have things really changed? In NBER Macroeconomics Annual 2010, volume 25, pages 333–367. National Bureau of Economic Research, Inc, 2011. URL https://EconPapers.repec.org/RePEc:nbr:nberch:12036.

L. Reichlin, G. Ricco, and M. Tarbé. Monetary-fiscal crosswinds in the european monetary union. BIS Working Papers 940, Bank for International Settlements, 2021. URL https://EconPapers.repec.org/RePEc:bis:biswps:940.

M. Rostagno, C. Altavilla, G. Carboni, W. Lemke, R. Motto, A. Saint Guilhem, and J. Yiangou. A tale of two decades: the ecb's monetary policy at 20. Working Paper Series 2346, European Central Bank, 2019. URL https://EconPapers.repec.org/RePEc:ecb:ecbwps:20192346.

J. J. Rotemberg and M. Woodford. Imperfect competition and the effects of energy price increases on economic activity. Journal of Money, Credit and Banking, 28(4):549–577, 1996. ISSN 00222879, 15384616. URL http://www.jstor.org/stable/2078071.

T. J. Rothenberg. Identification in parametric models. Econometrica, 39(3):577–91, 1971. URL https://EconPapers.repec.org/RePEc:ecm:emetrp:v:39:y:1971:i:3:p:577-91.

J. H. Stock and M. W. Watson. Why has us inflation become harder to forecast? Journal of Money, Credit and banking, 39:3–33, 2007.

J. H. Stock and M. W. Watson. Slack and Cyclically Sensitive Inflation. NBER Working Papers 25987, National Bureau of Economic Research, Inc, June 2019. URL https://ideas.repec.org/p/nbr/nberwo/25987.html.

K. Storesletten and F. Zilibotti. China's great convergence and beyond. Annual Review of Economics, 6(1):333–362, 2014. doi: 10.1146/annurev-economics-080213-041050. URL https://doi.org/10.1146/annurev-economics-080213-041050.

H. Taniguchi and K. Yamada. The race between technological progress and female advancement: Changes in gender and skill premia in oecd countries. arXiv preprint arXiv:2005.12600, 2020.

H. Uhlig. What drives gnp? Unpublished manuscript, Euro Area Business Cycle Network, 2003.

M. Villani. Steady-state priors for vector autoregressions. Journal of Applied Econometrics, 24(4):630–650, 2009.

B. A. Weinberg. Computer use and the demand for female workers. Industrial and Labor Relations Review, 53(2):290–308, 2000. ISSN 00197939, 2162271X. URL http://www.jstor.org/stable/2696077.

J. Yellen. Monetary policy in a global environment. FRBSF Economic Letter, (jun2):13, 2006. URL https://EconPapers.repec.org/RePEc:fip:fedfel:y:2006:i:jun2:n:2006-12-13.