MACROECONOMIC EFFECTS OF THE GENDER REVOLUTION*

DRAGO BERGHOLT†, LUCA FOSSO‡ AND FRANCESCO FURLANETTO§

December 2024

Abstract: U.S. labor market data exhibit a major, secular decline in the employment and wage gaps between males and females. In this paper, we identify the underlying, structural forces and quantify the spillover from this gender convergence to the broader macroeconomy. A novel time series model maps empirical trends in data into (aggregate and gender-specific) structural trends. Identification is achieved with restrictions derived from a neoclassical model with gender-specific labor. Empirically, we find that secular changes in female-specific labor productivity account for approximately one-third of economic growth in the postwar U.S. economy, in addition to most of the observed gender convergence.

Keywords: *Economic growth, gender inequality, labor market, productivity, VAR with common trends.*
JEL Classification: *C32, E10, E13, J2.*

*This work should not be reported as representing the views of the European Central Bank or Norges Bank. The views expressed are those of the authors and do not necessarily reflect those of the European Central Bank or Norges Bank. We thank for useful comments Stefania Albanesi, Jonas Arias, Philippe Andrade, Guido Ascari, Regis Barnichon, Martin Beraja, Paolo Bonomolo, Fabio Canova, Efrem Castelnuovo, Andrea Cerrato, Larry Christiano, Marco Del Negro, Martin Eichenbaum, Domenico Giannone, Luca Fornaro, Peter Karadi, Joseba Martinez, Leonardo Melosi, Karel Mertens, Silvia Miranda Agrippino, Evi Pappa, Giorgio Primiceri, Lukasz Rachel, Giuseppe Ragusa, Giovanni Ricco, Juan Rubio-Ramirez, Aysegul Sahin, Stefan Schiman, seminar participants at Northwestern University, Boston Fed, Boston College, Chicago Fed, Dallas Fed, Fed Board, University of Texas Austin, University of California Davis, Cleveland Fed, Washington University of Saint Louis, European Central Bank, Norges Bank, University of Aix-Marseille, University of Milano-Bicocca, 2021 Sailing the Macro Conference in Ventotene, 2021 conference in honor of Fabio Canova in Hydra, 2022 Padova Macro Talks, 2022 Barcelona Summer Forum, 2022 Dolomiti Macro Meetings in Selva Val Gardena, 2022 Forskermøtet in Stavanger, 2023 SNDE conference in Orlando, 2023 IAAE conference in Oslo, 2023 CEF conference in Nice, 2023 Dynare Conference in Malta, 2023 Fall Midwest Macro Meetings in Lubbock, 2024 Workshop in Empirical Macroeconomics King's College London, 2024 RCEA International Conference in London, 2024 T2M Conference in Amsterdam and 2024 SMN Conference in Madrid.

†Norges Bank. Email: drago.bergholt@norges-bank.no
‡European Central Bank. Email: luca.fosso@ecb.europa.eu
§Norges Bank. Email: francesco.furlanetto@norges-bank.no. Corresponding author.

1 INTRODUCTION

Women's increased labor market participation is arguably one of the fundamental changes observed in modern economies during the last century. Consider, for example, the U.S. labor market data in Figure 1: in the 1960s, the employment rate for females was less than half of the employment rate for males. But the female-to-male employment ratio increased steadily throughout the 1980s and 1990s, before converging to around 85 percent in recent decades. Gender differences in wages display a similar picture. Although women's hourly wages stayed relatively flat at 60 percent of men's wages until the mid 1970s (despite a substantial employment catch-up during that period), they have since outgrown male wages at about the same pace as the additional growth in female employment, resulting in a major wage convergence between the genders. In total, more than 60 percent of the female-to-male *employment gap*, and about half of the female-to-male *wage gap*, have disappeared in the last 5-6 decades. It is hardly a coincidence that Goldin (2006) refers to a "quiet revolution" when describing these trends.

The goal of our paper is twofold: first, we want to quantify the consequences of this gender revolution for the U.S. macroeconomy. In particular, we estimate the spillover effects on economic growth in terms of U.S. GDP, employment, and productivity. An important part of our motivation is the concern that talent may be significantly misallocated when only a minority of women participate in paid work, as advocated by Hsieh, Hurst, Jones, and Klenow (2019). Second, we aim to shed light on the structural drivers behind the gender convergence in employment and wages. At first glance, the observation that these two trends co-move may suggest that labor demand factors have been dominant, as stressed by Aguiar and Hurst (2007) and Fukui, Nakamura, and Steinsson (2023). However, the aggregate time series shown in Figure 1 are silent about any reallocation across different skill segments of the labor force, as well as reallocation across sectors. Our aim is to appropriately disentangle labor demand and labor supply factors once we control for the fact that the large increase in employment of female workers was concentrated in the market for high-skilled workers and in the service sector.

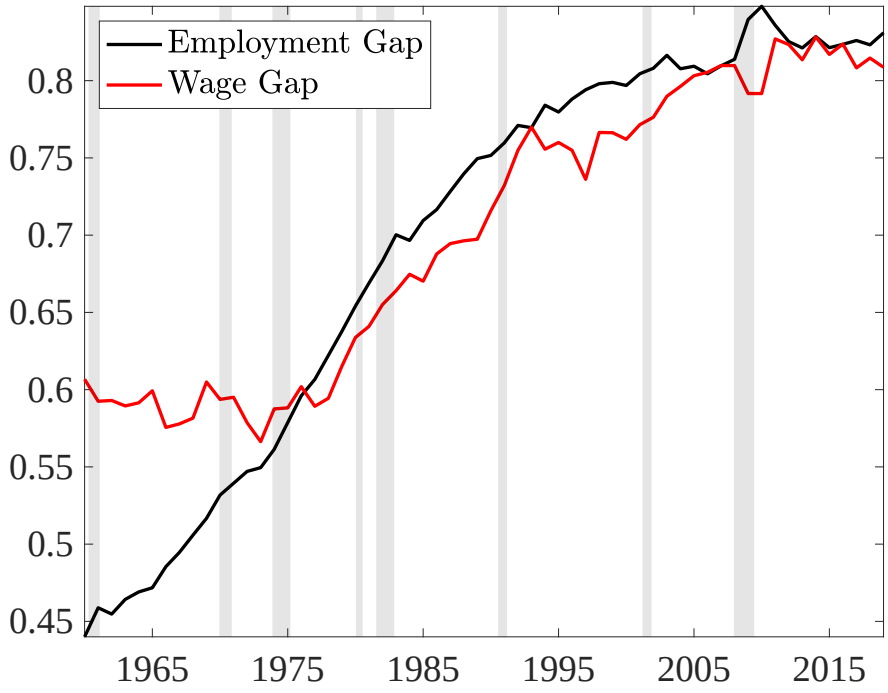
To identify the structural trends of interest in this paper, we develop a neoclassical model involving gender-specific labor which builds on Albanesi (2024) and Fukui et al. (2023).[1] The theoretical framework allows us to derive mutually exclusive identification restrictions on three gender-neutral macro trends and two gender-specific labor market trends. In turn, we impose these restrictions on a Structural Vector Autoregressive (SVAR) model fitted to relevant macro data, as well as to data on differences between females and males in wages and employment. The resulting econometric framework allows us to infer structural gender trends empirically, and to quantify their importance for the U.S. macroeconomy.

As our key focus is on slow-moving, structural drivers that persist way beyond common business cycle horizons, we use a *SVAR model with common trends*, as in Del Negro, Giannone, Giannoni, and Tambalotti (2017) and Crump, Eusepi, Giannoni, and Sahin (2019).[2] The model can be seen as a multivariate, unobserved components model, in

[1] Albanesi (2024) estimates a real business cycle model with gender-specific labor to account for jobless recoveries. Fukui et al. (2023) extend the model with home production and open economy features, showing that women's rising participation did not crowd out males' participation, thus, being an expansionary factor for the macroeconomy.

[2] More recent extensions of the same model to explain inflation dynamics include Ascari and Fosso (2024),

Figure 1: Gender differences in employment and wages



Notes: The employment (wage) gap is defined as the female-to-male ratio in employment (wage) rates. Based on data from The Current Population Survey, the United Census Bureau, the Bureau of Labor Statistics, and authors' calculations.

which the variables enter in levels, and transitory and permanent components in data are disentangled from each other. We seek to quantify the latter. As an example, let us consider GDP. Our model decomposes observed GDP dynamics into a cyclical and a permanent component. In turn, the permanent component—understood as the empirical or reduced-form GDP trend—is a function of underlying, structural drivers such as productivity and demographics (both of which can be gender-neutral or gender-specific). However, the mapping from these structural drivers to the empirical trend in GDP is unknown ex-ante. It is exactly this identification problem that we address with restrictions from economic theory. In particular, the theoretical framework presented here implies a log-linear mapping between trend GDP and (gender-neutral) technology shocks, (gender-neutral) automation shocks, (gender-neutral) labor supply shocks, as well as gender-specific labor demand and labor supply shocks. We disentangle these five *structural* forces based on their long-run impact on economic variables. This represents a key distinction from previous studies estimating VARs with stochastic trends. Del Negro et al. (2017) and Crump et al. (2019), for example, estimate selected common trends in data, but remain silent about the underlying *structural* sources. Our framework, instead, allows us to estimate the mapping from structural to empirical trends, and to instruct a Bayesian algorithm we deploy with prior information derived from theory.

We estimate a baseline model where standard macro data for the U.S. economy are linked to data on aggregate gender differences in wages and employment. The idea is to

Bianchi, Nicolò, and Song (2023), Hasenzagl, Pellegrino, Reichlin, and Ricco (2022) and Maffei-Faccioli (2024).

infer the structural drivers of trends in wage and employment differences between females and males, and to quantify their importance for the U.S. macroeconomy. In a second step, following Dolado, Motyovszki, and Pappa (2021) we use information on individuals' education and sector of employment contained in the Current Population Survey (CPS), and estimate extensions of the baseline model with data on gender differences in skills and sectors of employment. The results from these more granular extensions are then compared with those from the baseline. The idea is to separate "within-skill" and "within-sector" fixed effects from "within-gender" fixed effects. Thus, we can disentangle fundamental gender trends from trends in skills and sectoral composition, and gauge compositional effects that may bias our baseline results.

Our main empirical result concerns the macroeconomic spillover from gender-specific labor market trends in our data: these trends, which capture the observed convergence between females and males, are also essential for macroeconomic growth in the postwar U.S. economy. During the 70s, 80s and 90s for example, they account for 30-50% of the overall trend increase in GDP, and for 20-40% of the overall trend increase in productivity. Moreover, they are responsible for a sizable share of the slowdown in trend GDP growth during the last 20 years, coinciding with a leveling-off of the gender convergence compared with earlier decades. In total, gender-specific factors explain almost one-third of GDP growth in the postwar U.S. economy, and they prevented an overall employment decline during this period. Importantly, these results are obtained when we control for aggregate macro trends such as the evolution of total factor productivity.

In addition, structural factors that originate on the *demand side* of the labor market explain (i) most of the long-term gender convergence in employment and wages, (ii) the end of the gender convergence observed in the last 20-30 years, as well as (iii) almost all *net spillovers* from gender-specific labor market trends to the macroeconomy. From an econometric point of view, our empirical model suggests that the female-to-male employment ratio tends to increase permanently in periods when there is a permanent rise in the wages of females relative to the wages of males. This is consistent with a story about female-specific productivity growth, which shifts labor demand towards females. Supply-side explanations, by contrast, should have implied stagnant wage growth of females in periods with strong female employment growth. This is not what we typically see in our aggregate data.

Finally, while the net spillover from labor supply factors to the macroeconomy is limited, we document an important role for *skill-biased trends in females' labor supply*: when gender gap data by skill segments are considered, a supply-driven expansion of female labor in high-skill jobs emerges, and at the same time a contraction in the supply of female labor in low-skill jobs. In sum, these two forces have led to substantial reallocation of females from low-skill to high-skill jobs. Moreover, the compositional effects associated with these supply-driven labor flows across skill segments have likely contributed to the overall convergence in female and male wages. At the same time, they have largely counteracted each other in aggregate data on gender pay and employment gaps. This is partly why the net spillover effect from gender-specific labor supply factors to the macroeconomy are so limited in our baseline estimation. But trends in females' labor supply are still important for the macroeconomy.

Our paper speaks to a large literature studying the gender revolution. A useful distinction for our purposes is between papers discussing labor demand factors and labor supply

factors. Among the former, Galor and Weil (1996) emphasize technological factors that favored the demand for women in combination with an increase in the returns to intellectual skills (Beaudry and Lewis (2014), Rendall (2024)) and the rise of the service sector (Ngai and Petrongolo (2017), Buera, Kaboski, and Zhao (2019)), while Jones, Manuelli, and McGrattan (2015) and Hsieh et al. (2019) point to a reduction in gender discrimination and a reduction in barriers to schooling as important drivers of the convergence in wages. Among the latter, Albanesi and Olivetti (2016) and Goldin and Katz (2002) document the importance of advances in maternal health and contraception, Fernández, Fogli, and Olivetti (2004) emphasize cultural factors developed during World War II, Attanasio, Low, and Sánchez-Marcos (2008) point to the crucial role of availability and affordability of child care, while Greenwood, Seshadri, and Yorukoglu (2005) propose a model in which the emergence of home appliances favors females' 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 provides a clear link between gender trends and macroeconomic outcomes.

Within the large literature studying the gender convergence, our paper is closely related to two seminal papers quantifying the role of gender for macroeconomic dynamics in calibrated quantitative set-ups. Both Hsieh et al. (2019) and Heathcote, Storesletten, and Violante (2017) propose a decomposition of US macroeconomic growth in structural models with gender and find a major role for gender forces. Our results are in the same ballpark as theirs, but obtained using a simple time series model which is substantially less parameterized and estimated (rather than calibrated) on US quarterly data. Albanesi (2024) is also of particular interest. Using an estimated real business cycle model with a gender dimension, she finds that women's relative labor supply and productivity have been important for changes in the business cycle, and more generally, for the economic performance in the US.

The explicit use of theory to form priors for an estimated, empirical time series model relates our work to the well-known DSGE-VAR methodology proposed by Del Negro and Schorfheide (2004). As in their case, we are concerned about the potential misspecification induced by the tight cross-equation restrictions featured by a fully specified theoretical model. Therefore, we use theory only as a prior to inform the VAR with common trends. Differently from Del Negro and Schorfheide (2004), we focus on the variables' permanent components and not on the cycle. Moreover, we form priors about specific structural elasticities, rather than the full covariance structure implied by theory.

The rest of the paper is organized as follows: section 2 introduces the empirical trend-cycle model that we fit to U.S. labor market data (aggregate and gender-specific). Section 3 describes the theoretical framework that disciplines our empirical assessment of trends, while section 4 derives the exact, theory-robust identification restrictions. Section 5 documents the paper's main empirical results. Section 6 takes stock and puts our results in perspective. Section 7 relates the gender convergence in wages and employment to trends in the skills and the rise of services, while section 8 considers structural changes specific to male labor. Section 9 concludes.

2 A TIME SERIES MODEL WITH COMMON TRENDS

The model that we estimate is a multivariate time series model with unobserved components, see Watson (1986), Stock and Watson (1988, 2007), and Villani (2009). It decomposes a vector of observable data into *two unobservable, stochastic components*: the first component is characterized by cyclical but transitory fluctuations. The second component captures permanent changes, or trends, in data. One distinctive contribution of our paper is to map these trends into a vector of underlying, structural drivers. Importantly, the structural drivers may give rise to common trends in data, even if the structural drivers themselves are orthogonal to each other. To fix ideas, consider an $n \times 1$ vector of data Y_t, which is the sum of two unobserved states:

$$Y_t = \hat{Y}_t + \bar{Y}_t, \tag{1}$$

$\hat{Y}_t$ and $\bar{Y}_t$ represent the cycles and trends in our data, respectively. As such, equation (1) is a purely statistical (reduced-form) decomposition of the data. The focus of our paper will be on the *empirical trends* in $\bar{Y}_t$ and, more precisely, on the underlying causal drivers behind $\bar{Y}_t$. We suppose that $\bar{Y}_t$ can be decomposed into $q \leq n$ *structural trends*, collected in the $q \times 1$ vector X_t:

$$\bar{Y}_t = \mathcal{V} X_t \tag{2}$$

Here, $\mathcal{V}$ is a $n \times q$ matrix that maps the reduced-form trends into structural, economic factors. Similarly to Del Negro et al. (2017) and Crump et al. (2019), we treat the structural trends as separate random walk processes, potentially with a drift:

$$X_t = c + X_{t-1} + u_t, \qquad u_t \sim \mathcal{N}(0_q, \Sigma_u) \tag{3}$$

Throughout we assume that the covariance matrix Σ_u is diagonal. This is standard given that u_t is a vector of structural trend shocks. Nevertheless, the presence of non-zero, off-diagonal elements in the matrix $\mathcal{V}$ still implies common, stochastic trends across the individual components of $\bar{Y}_t$. A vital part of our analysis will be to identify $\mathcal{V}$, given that we are interested in the causal drivers of trends in data. Moreover, since our focus is on trends rather than on business cycle fluctuations, only a minimal set of restrictions is imposed on $\hat{Y}_t$. In particular, we model $\hat{Y}_t$ as a vector autoregressive (VAR) process in reduced-form:

$$\Phi(L)\hat{Y}_t = e_t, \qquad e_t \sim \mathcal{N}(0_n, \Sigma_e) \tag{4}$$

$\Phi(L) = I - \Phi_1 L - ... - \Phi_p L^p$ is an $n \times n$ matrix of lag coefficients. Σ_e is freely estimated without any restrictions on the off-diagonal elements. However, we assume that permanent and transitory shocks are mutually uncorrelated, i.e. that $cov(u_t, e_t) = 0$. In the parlance of Watson (1986), the model is an "independent trend-cycle decomposition" and trends do not affect the cycle by construction.[3]

Equations (1)-(4) constitute the model that we confront with data. It is the combination of relevant (aggregate and gender-specific) labor market data, together with a proper identification of $\mathcal{V}$, that allows us to infer gender-specific trends and quantify their importance for the U.S. macroeconomy. A first methodological contribution of this paper

[3] If anything, a violation of this assumption may bias estimates in the cyclical block given by (4). However, we view the assumption as rather innocuous, given that our sole interest is in secular trends and not in cyclical fluctuations.

is to derive identification restrictions on $\mathcal{V}$ from economic theory, and then to use these restrictions to *estimate* the elements in $\mathcal{V}$. This is a key difference from previous studies, which instead have calibrated $\mathcal{V}$. A second novel aspect is that the vector X_t is composed by *unobservable* variables. In previous studies, the permanent component of variables are related to *observable* variables: as an example, Del Negro et al. (2017) assume that inflation and long-run inflation expectations share the same permanent component. In that case, data on inflation expectations are used to obtain a better estimate of trend inflation. In our case, we are not interested in obtaining a refined estimate of the permanent components but rather to identify their unobservable structural drivers.

3 A STYLIZED MODEL WITH GENDER-SPECIFIC LABOR

In this section, we present a neoclassical model that allows us to derive theory-based identification assumptions and prior distributions for elements in the matrix $\mathcal{V}$. The theoretical model builds on the previous key contributions by Fukui et al. (2023) and Albanesi (2024). 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. Importantly, the representative household consists of *two different worker types*. Here we refer to these two types as females and males.

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 β is the time discount rate and Λ_t represents the shadow value of income. The firm's period profit is equal to

$$\Pi_t = Y_t - W_{f,t} L_{f,t} - W_{m,t} L_{m,t} - P_{I,t} I_t, \tag{5}$$

where 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 the firm's 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

$$Y_t = A_t L_t^{\alpha_t} K_{t-1}^{1-\alpha_t}, \tag{6}$$

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

$$L_t = \left[\alpha_l \left(A_{m,t} L_{m,t} \right)^{\frac{\gamma-1}{\gamma}} + (1 - \alpha_l) \left(A_{f,t} L_{f,t} \right)^{\frac{\gamma-1}{\gamma}} \right]^{\frac{\gamma}{\gamma-1}} \tag{7}$$

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. Importantly, a rise in $A_{f,t}$ may, for example, represent less discrimination of females at the workplace (Hsieh et al. (2019) and Jones et al. (2015)), or an expansion of women's rights (Doepke and Tertilt (2009)). A specification where discrimination is modeled as a tax on female labor is, as pointed out by Fukui et al. (2023), isomorphic to our female-specific

productivity shocks. However, a rise in $A_{f,t}$ may also capture technological innovations that increase the returns to social (Cortes, Jaimovich, and Siu (2023)) or intellectual skills (Rendall (2024)) in which females have a comparative advantage, or a secular increase in the service sector where females are disproportionately employed. The scope for such compositional effects is investigated in section 7.

We allow for a time-varying weight α_t on aggregate labor. One possible interpretation of a decline in α_t is that it follows from labor-displacing automation, see Acemoglu and Restrepo (2020) and Bergholt, Furlanetto, and Maffei-Faccioli (2022). But it may also follow from outsourcing or offshoring, or any other labor-saving, technical change.

Finally, physical capital dynamics are given by

$$K_t = (1-\delta) K_{t-1} + I_t. \tag{8}$$

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

$$P_{I,t} = \mathcal{Q}_t \tag{9}$$

$$\mathcal{Q}_t = \beta \mathbb{E}_t \frac{\Lambda_{t+1}}{\Lambda_t} \left[(1-\alpha_{t+1}) \frac{Y_{t+1}}{K_t} + \mathcal{Q}_{t+1} (1-\delta) \right] \tag{10}$$

$$W_{m,t} = \alpha_t \alpha_l \frac{Y_t}{L_t} \left(\frac{L_t}{L_{m,t}} \right)^{\frac{1}{\gamma}} A_{m,t}^{\frac{\gamma-1}{\gamma}} \tag{11}$$

$$W_{f,t} = \alpha_t (1-\alpha_l) \frac{Y_t}{L_t} \left(\frac{L_t}{L_{f,t}} \right)^{\frac{1}{\gamma}} A_{f,t}^{\frac{\gamma-1}{\gamma}} \tag{12}$$

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, gender-specific labor demand is increasing in aggregate activity, decreasing in the gender-specific wage rate, and, since $\gamma > 1$, increasing 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

$$\mathcal{U}_t = \frac{C_t^{1-\sigma}}{1-\sigma} \exp\left(-\Psi_t^{-1} \frac{(1-\sigma) \tilde{L}_t^{1+\varphi}}{1+\varphi} \right) \tag{13}$$

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

$$\tilde{L}_t = \left[\left(\frac{L_{m,t}}{\Psi_{m,t}} \right)^{\frac{1+\lambda}{\lambda}} + \left(\frac{L_{f,t}}{\Psi_{f,t}} \right)^{\frac{1+\lambda}{\lambda}} \right]^{\frac{\lambda}{1+\lambda}} \tag{14}$$

$\Psi_{m,t}$ and $\Psi_{f,t}$ are gender-specific labor supply shocks, $\lambda > 0$ governs the household's willingness to substitute female with male labor. Importantly, a rise in $\Psi_{f,t}$ captures

supply-side factors that effectively increase females' attachment to the labor market, such as advances in maternal health, contraception products, the emergence of home appliances, availability and the cost of child care, or cultural factors.

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:

$$\Lambda_t = C_t^{-\sigma} \exp\left(-\Psi_t^{-1} \frac{(1-\sigma)\,\tilde{L}_t^{1+\varphi}}{1+\varphi}\right) \tag{15}$$

$$\Lambda_t = \beta \mathbb{E}_t \Lambda_{t+1} (1 + r_t) \tag{16}$$

$$W_{m,t} = \Psi_t^{-1} C_t \tilde{L}_t^{\varphi - \frac{1}{\lambda}} L_{m,t}^{\frac{1}{\lambda}} \Psi_{m,t}^{-\frac{1+\lambda}{\lambda}} \tag{17}$$

$$W_{f,t} = \Psi_t^{-1} C_t \tilde{L}_t^{\varphi - \frac{1}{\lambda}} L_{f,t}^{\frac{1}{\lambda}} \Psi_{f,t}^{-\frac{1+\lambda}{\lambda}} \tag{18}$$

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 increasing in the aggregate and gender-specific labor supply shocks.

The optimality conditions above allow for quite rich interactions between the labor supply decisions of females and males. Consider, for example, an increase in female labor, which raises the household's overall labor dis-utility $\tilde{L}_t$. Ceteris paribus, it is not clear how male workers would tend to respond. On one side, the male may want to take out more vacation to compensate for the household's overall decline in leisure. However, as his own valuation of leisure generally depends on the spouse's labor allocations, a sufficiently strong degree of gender complementarity implies that the male benefits less from leisure if his spouse is busy working. In that case, even his own labor supply increases when the spouse works more. Importantly, λ governs the household's willingness to substitute work across genders in our framework. If the female works more, ceteris paribus this implies a reduction (increase) in the male's labor supply if and only if $\varphi > \lambda^{-1}$ ($< \lambda^{-1}$). Note that Fukui et al. (2023) set $\lambda = \frac{1}{\phi}$ in their baseline model. This knife-edge parameterization effectively makes the marginal dis-utility of female and male labor independent of how much the spouse is working. However, we want to be agnostic about the net effect when our empirical model is confronted with data.

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:

$$l_{f,t} = \left(\frac{1-\alpha_l}{\alpha_l}\right)^{\gamma} w_{f,t}^{-\gamma} a_{f,t}^{\gamma-1} \tag{19}$$

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 λ determines how responsive supply is to relative wage changes:

$$l_{f,t} = w_{f,t}^{\lambda} \psi_{f,t}^{1+\lambda} \tag{20}$$

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. From now on, we assume that shifts in $a_{f,t}$ and $\psi_{f,t}$ can be attributed to female-specific productivity and labor supply, implicitly keeping male-specific productivity and labor supply constant. However, this assumption is largely inconsequential, as shown in section 8 where we identify shocks for each gender directly.

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

$$w_{f,t} = \left(\frac{1-\alpha_l}{\alpha_l}\right)^{\frac{\gamma}{\gamma+\lambda}} a_{f,t}^{\frac{\gamma-1}{\gamma+\lambda}} \psi_{f,t}^{-\frac{1+\lambda}{\gamma+\lambda}} \tag{21}$$

$$l_{f,t} = \left(\frac{1-\alpha_l}{\alpha_l}\right)^{\frac{\gamma\lambda}{\gamma+\lambda}} a_{f,t}^{\frac{(\gamma-1)\lambda}{\gamma+\lambda}} \psi_{f,t}^{\frac{(1+\lambda)\gamma}{\gamma+\lambda}} \tag{22}$$

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. The there "Macro shocks" A_t, Ψ_t, and α_t play no role here.[4] 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.

4 FROM THEORY TO TREND IDENTIFICATION IN DATA

Next, we explain how the neoclassical theory just described is used in practice to discipline our empirical analysis, allowing us to identify the underlying, structural trends in data.

4.1 MONTE CARLO SIMULATIONS

The main purpose of the theoretical model is to infer a set of theory-consistent restrictions on $\mathcal{V}$. These restrictions should be (i) informative enough so that they ensure the identi-

[4]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 γ and λ, and remains even if we were to introduce business cycle frictions such as nominal price rigidities.

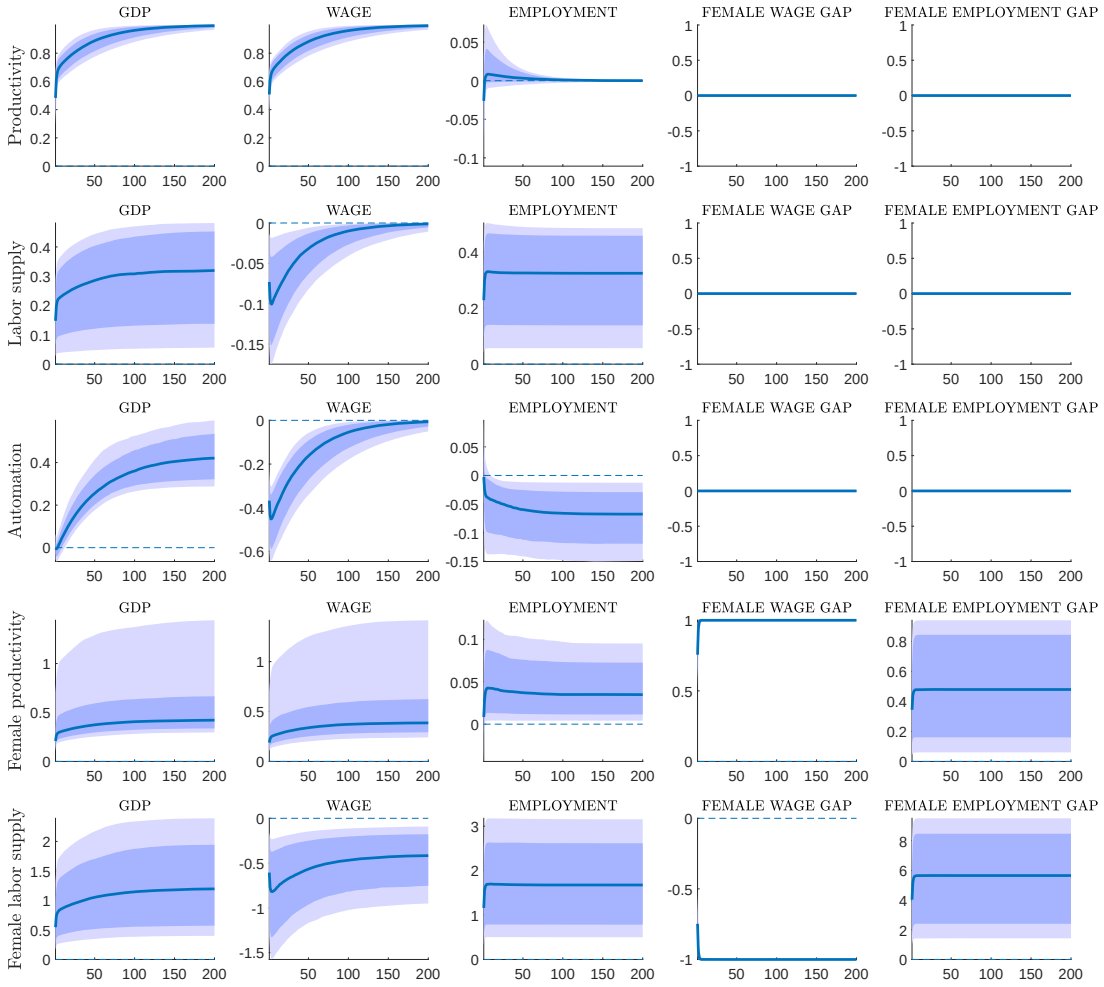
fication of all estimated elements in $\mathcal{V}$, and (ii) sufficiently agnostic so that our analysis remains robust to reasonable parametric perturbations of the underlying theory.

To this end, we conduct the following simulation exercise: first, we draw a parameter vector $\theta = [\sigma, \varphi, \gamma, \lambda, ...]'$ which includes all parameters of interest in the theoretical model, including state variables such as initial wage and employment gaps. In order to be as agnostic as possible, we draw each parameter independently from a uniform distribution specified further below. Second, conditional on θ, we solve the theoretical model and compute the impulse responses of macro and gender variables to each of the structural trend shocks. 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 Monte Carlo exercise leaves us, at each time horizon and conditional on each shock, with an entire distribution of structural outcomes for the endogenous variables of interest. The distribution visualizes 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 the existing literature. In particular, for the three "macro" parameters σ, φ and α 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 φ^{-1}, both from microeconomic and macroeconomic literature, are well within the bounds used here. Moreover, the bounds on α allow the model to cover a wide range of labor income shares, including all those observed in the postwar US economy. For the two parameters of key interest, namely those that reflect firms' and households' ability to substitute between female and male labor, we choose the following uniform distributions: $\gamma \sim U(1, 11)$ and $\lambda \sim U(0, 1)$. These distributions are broadly informed by Albanesi (2024), but with wide bounds that allow us to consider an array of different scenarios. 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 U.S. economy (see Figure 1).

Impulse responses from the simulation exercise are presented in Figure 2. We restrict attention to the five variables needed to separately identify our five structural trends. The first three rows document the responses to changes in gender-neutral macro trends, while the last two rows display responses to innovations in the two gender trends. Since the goal is to infer the permanent effects of structural change, we focus on the long end of the impulse responses. We start by commenting on the responses to permanent macro shocks. A few remarks are in place: first, consistent with the analytical solutions in equations (21)-(22), aggregate shocks have no effects whatsoever on the wage and employment gaps between females and males, and this is true at all horizons. Second, conditional on a permanent rise in aggregate productivity, GDP and aggregate wages rise one-for-one with the productivity increase in the long run, while aggregate, long-run employment is unaffected. Third, in response to a permanent rise in aggregate labor supply, we obtain long-run increases in GDP and aggregate employment of equal magnitude, but no long-run effects on aggregate wages. Fourth, a permanent rise in automation generates a long-run increase in GDP, a long-run decline in aggregate employment, but has zero long-run

Figure 2: Simulated impulse responses from the theoretical model



Notes: Impulse response functions from simulations of the theoretical model. Pointwise median, 90% and 68% bands based on 1,000 independent draws from the parameter distributions. The y-axes measure responses in percent, the x-axes represent time in quarters.

effects on aggregate wages. These dynamics are perfectly conventional and hold also in the standard neo-classical growth model without gender (Bergholt et al. (2022)).

Finally, the two gender-specific trend shocks in Figure 2 are normalized so that they each cause a unit change in the long-run wage gap between females and males. Conditional on a permanent rise in female-specific productivity, both wages and employment of females rise relative to that of males. Since this shock increases average labor productivity in the economy, all three macro variables rise as well in the long run. A permanent fall in females' labor dis-utility, in contrast, causes female wages to decline relative to the wage of males, while the female-to-male employment ratio permanently increases. Moreover, GDP and employment both rise in the long run while the aggregate wage rate falls.

4.2 REVISITING THE MAPPING TO EMPIRICAL TRENDS

We are now in a position to specify a baseline, theory-consistent mapping $\mathcal{V}$ from the structural trends X_t to trends in data $\bar{Y}_t$:

$$\underbrace{\begin{bmatrix} G\bar{D}P_t \\ \bar{W}_t \\ \bar{L}_t \\ \bar{w}_{f,t} \\ \bar{l}_{f,t} \end{bmatrix}}_{\bar{Y}_t} = \underbrace{\begin{bmatrix} 1 & 1 & 1 & \nu_{14} & \nu_{15} \\ 1 & 0 & 0 & \nu_{24} & \nu_{25} \\ 0 & 1 & \nu_{33} & \nu_{34} & \nu_{35} \\ 0 & 0 & 0 & 1 & -1 \\ 0 & 0 & 0 & \lambda & \gamma \end{bmatrix}}_{\mathcal{V}} \underbrace{\begin{bmatrix} A_t \\ \Psi_t \\ \alpha_t \\ a_{f,t} \\ \psi_{f,t} \end{bmatrix}}_{X_t} \tag{23}$$

The first column of $\mathcal{V}$ imposes restrictions on the stochastic technology trend. We normalize it to have a unit effect on GDP and wages, and zero long-run effects on the remaining variables. The second column of $\mathcal{V}$ governs the spillover from the aggregate labor supply trend, which is normalized to cause a unit increase in GDP and employment, while the third column in $\mathcal{V}$ imposes that automation permanently increases GDP and lowers employment. The latter restriction implies that $\nu_{33} < 0$. Consistent with Figure 2 we impose zero-restrictions on the remaining elements in $\mathcal{V}$ for both labor supply and automation. Note that most zero-restrictions are based on the assumption that macro shocks do not permanently affect gender gaps in wages and employment. This assumption, which greatly facilitates identification of $\mathcal{V}$, effectively rules out compositional effects of gender-neutral macro shocks. However, we inspect the possibility of compositional effects in section 7.

The fourth and fifth columns in $\mathcal{V}$ govern the long-run implications of permanent changes in female-specific labor productivity and labor supply, respectively. We normalize both of these trends so that they have a unitary effect on the wage gap between females and males. That is, rather than estimating $a_{f,t}$ and $\psi_{f,t}$, we identify $\tilde{a}_{f,t} \equiv a_{f,t}^{\frac{\gamma-1}{\gamma+\lambda}}$ and $\tilde{\psi}_{f,t} \equiv \psi_{f,t}^{\frac{1+\lambda}{\gamma+\lambda}}$. This is particularly convenient, as can be seen from the normalized, log-linear versions of (21)-(22):

$$\hat{w}_{f,t} = c_{w,f} + \hat{\tilde{a}}_{f,t} - \hat{\tilde{\psi}}_{f,t} \tag{24}$$

$$\hat{l}_{f,t} = c_{l,f} + \lambda\hat{\tilde{a}}_{f,t} + \gamma\hat{\tilde{\psi}}_{f,t} \tag{25}$$

A hat means that the variable is expressed in logarithms, with $c_{w,f}$ and $c_{l,f}$ being reduced-form constants. Importantly, the gender substitution elasticities λ and γ, two structural parameters of particular interest, enter directly as coefficients in (25). This means that $\lambda = \nu_{54}$ and $\gamma = \nu_{55}$ can be read directly from the estimated matrix $\mathcal{V}$.

The restrictions imposed on each column in $\mathcal{V}$ are mutually exclusive, which is what we need to separately identify the five stochastic trends of interest. The aggregate technology trend, for example, is the only aggregate macro trend that makes the real wage co-move with GDP in the long-run and, consistent with the balanced growth path assumption, is the only one that has zero long-run effect on employment. Aggregate labor supply and automation can be disentangled in data because the former implies long-term co-movement GDP and employment, while the latter crowds out labor. Finally, female-specific labor demand and labor supply are separable from aggregate macro trends because they are the only drivers of long-run wage and employment gaps between females

Table 1: Prior distributions and posterior estimates

		Prior		Posterior		
		Density	Support	Mean	Mode	90% HPD
ν_{14}	$a_f \to G\bar{D}P$	Uniform	$[0, 1]$	0.91	0.98	(0.74, 0.99)
ν_{24}	$a_f \to \bar{W}$	Uniform	$[0, 1]$	0.60	0.67	(0.33, 0.81)
ν_{34}	$a_f \to \bar{L}$	Uniform	$[-0.5, 0.5]$	0.30	0.36	(0.09, 0.46)
ν_{15}	$\psi_f \to G\bar{D}P$	Uniform	$[0, 2]$	0.42	0.10	(0.02, 1.04)
ν_{25}	$\psi_f \to \bar{W}$	Uniform	$[-2, 0]$	-0.15	-0.02	(-0.48, -0.00)
ν_{35}	$\psi_f \to \bar{L}$	Uniform	$[0, 3]$	0.54	0.29	(0.04, 1.23)
$-\nu_{33}$	$\alpha \to \bar{L}$	$\Gamma(0.3, 0.15)$	$(0, \infty)$	0.39	0.40	(0.21, 0.60)
λ	$a_f \to \bar{l}_{f,t}$	$\Gamma(1, 0.5)$	$(0, \infty)$	1.53	1.63	(0.95, 2.18)
γ	$\psi_f \to \bar{l}_{f,t}$	$\Gamma(3, 1.5)$	$(0, \infty)$	3.26	3.17	(2.44, 4.20)

Notes: The posterior moments are generated from the last $10,000$ of $50,000$ draws generated from the RW Metropolis-Hastings algorithm. $\Gamma(\mu, \sigma^2)$ refers to the Gamma distribution with mean μ and variance σ^2.

and males. Moreover, they are uniquely identified because they imply opposite signs on the co-movement between gender gaps in wages and employment.[5]

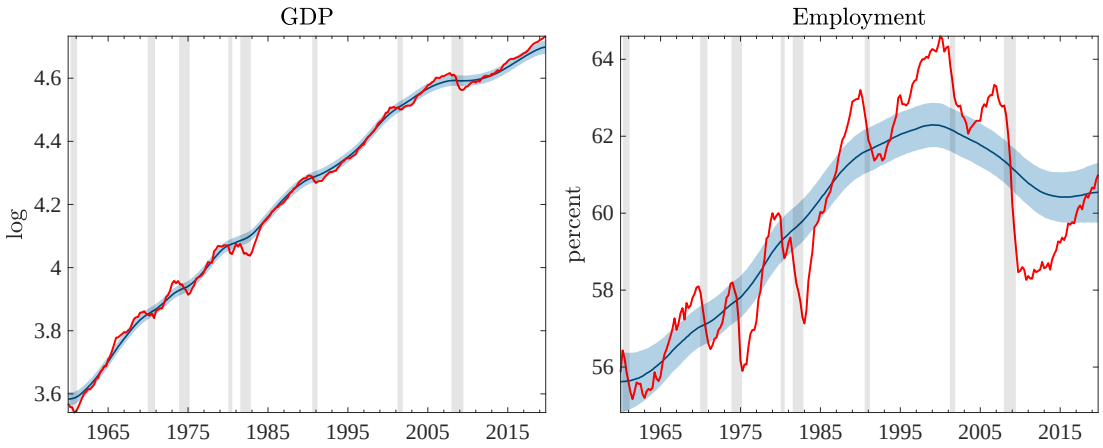
4.3 Priors

The last step is to specify prior shapes for the estimated parameters. Table 1 summarizes our choice of priors. We aim for an agnostic approach and use uniform priors for all elasticities governing the feedback from gender trends to the aggregate macroeconomy. The *support* of each uniform prior largely reflects the uncertainty bands computed during the Monte Carlo exercise, as shown in Figure 2. That is, consistent with theory, the female-specific productivity shock behaves qualitatively as a technology shock in the aggregate given our priors, while the female-specific labor supply shock behaves as a gender-neutral labor supply shock for aggregate variables. The prior for the aggregate employment response to automation has a Gamma distribution, reflecting that automation (a decline in α_t) crowds out employment. Note that we impose the prior on $-\nu_{33}$, since the Gamma distribution has a positive support.

The final two parameters that we estimate are γ and λ. The gender-specific labor demand elasticity γ has been quantified in a few existing studies. Weinberg (2000), for example, finds that γ is around 2.4 in the US, while Acemoglu, Autor, and Lyle (2004) report a slightly higher value of 3. Albanesi (2024) and Fukui et al. (2023) have considered values between 4 and 5. Thus, we choose a Gamma-prior for γ centered around 3 with most of the probability mass located between 1 and 5. The evidence on λ, which captures complementarity between males' and females' leisure time, is more scant. Ngai

[5]Note that the zero restrictions on employment in response to a technology shock and on aggregate wages in response to automation and labor supply shocks are not crucial for identification. In Appendix D.2, we show that results change only marginally if we relax these assumptions. In addition, a more conservative set of priors tilted against macro effects of gender shocks is evaluated in Appendix D.1.

Figure 3: Estimated empirical trend of real GDP and aggregate employment



Notes: observed data (red solid line), median trend estimate (blue solid line), and 68% coverage bands (blue shaded areas). The grey areas represent NBER recessions.

and Petrongolo (2017) use a value of 0.19 based on micro-evidence from Goux, Maurin, and Petrongolo (2014).[6] Thus, we choose a Gamma prior for λ with about 60% of the probability mass below one, and where the estimate from Goux et al. (2014) is covered by the 90% credible prior bands (even though much higher values are allowed as well during estimation). A defining feature of our priors is that firms can switch between female and male labor more easily than households ($\lambda < \gamma$) at the prior mode, a feature that seems highly reasonable.

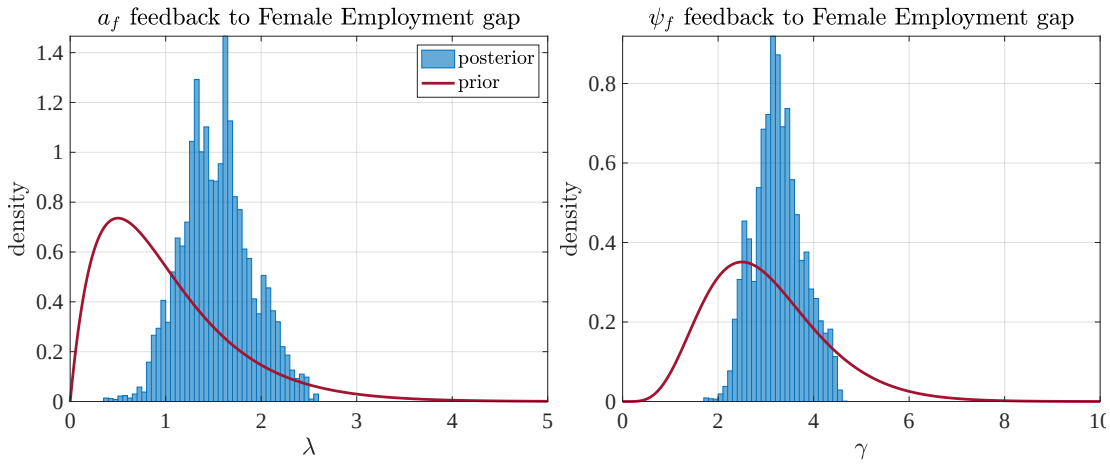
5 Empirical Results

In this section, we present the main empirical results based on the estimated time series model described in section 2. Given the theoretical restrictions derived in section 3 and section 4, the vector of observable variables Y_t includes: (i) real GDP, (ii) real aggregate wages, (iii) the aggregate employment-to-population ratio, (iv) the ratio of female-to-male wages, and finally (v) the ratio of female-to-male employment. All variables enter the system in log-levels. The model is estimated over the sample period 1960:Q1-2019:Q4, and we choose $p = 4$ lags in the system given that data are observed at a quarterly frequency. We use Bayesian methods to estimate the model. In particular, a Gibbs sampler is designed to generate $50,000$ draws, where the first 80% of the draws are discarded as a burn-in sample and the last 20% serve to generate posterior moments. The algorithm includes a Metropolis-Hastings step that draws from the posterior of the elasticities in $\mathcal{V}$. Details on the estimation steps, data sources and construction are laid out in Appendixes A and B, respectively.[7]

[6] Goux et al. (2014) exploit a workweek reduction policy in France to obtain an estimate that reflects pure cross-hour effects across partners and not income effects.

[7] All our data are published by the Bureau of Economic Analysis and the Bureau of Labor Statistics, and can be downloaded from the FRED website. In the baseline specification, data on employment and wages include both single and married individuals. In Appendix D.3, we restrict our attention to married couples only. Olsson (2024) and Albanesi and Prados (2022) model explicitly the heterogeneity in marital status.

Figure 4: Posterior distributions of coefficients λ and γ.



5.1 Permanent and transitory components

The first set of results is related to the decomposition of observable data into variable-specific, permanent and transitory components. Figure 3 reports the estimated permanent components of aggregate output and employment. Our model-implied trend estimates seem to be largely consistent with popular narratives for trends in US macro data: The GDP trend, for example, has displayed lower average growth in the last 20 years of data, and it leveled off completely during (and just after) the 2008 financial crisis. Moreover, the model assigns a large share of the observed employment decline in the last 15 years to permanent factors: it essentially concludes that the employment rate was back to its own trend by the end of our sample, despite being 3-4 percentage points lower than before the financial crisis. A secular decline in trend employment starting in the late 1990s is important for this result. Interestingly, our model-implied estimate of the output gap—defined as the difference between observed GDP and the inferred permanent counterpart (see Figure C.2 in the appendix)—exhibits a correlation coefficient of 0.88 with the output gap reported by the Congressional Budget Office (CBO). The latter constitutes a classic benchmark in the literature. Such a high correlation is neither obvious nor targeted, as our SVAR is not informed by data on the CBO estimates. We conclude that our model offers a reasonable description of trends and cycles in GDP in the postwar US economy, and that it can be used as a laboratory to investigate the structural drivers of trends in data.[8]

5.2 Estimated parameters in $\mathcal{V}$

Posteriors for the nine estimated parameters in $\mathcal{V}$ are summarized in Table 1, while Figure 4 plots the full posterior distributions of λ and γ. Despite using a prior with substantial mass below 1, we obtain a posterior density for λ that is centered around 1.5. This implies that a 1 percent increase in the wage of females relative to that of males, when caused by permanent productivity improvements among females, is associated with a 1.5 percent increase in the employment rate of females relative to that of males. Moreover, most of the posterior probability mass for λ is located between 1 and 2. The posterior mean value for

[8] An equivalent decomposition is provided in Appendix C for the remaining variables.

γ, in contrast, is 3.3. Thus, a permanent rise in females' labor supply, scaled to reduce the relative wage of females by 1 percent, increases the employment rate of females relative to males by more than 3 percent at the posterior mean. This number is close to, but somewhat higher than those obtained by Acemoglu et al. (2004) and Weinberg (2000) (they report values of 3 and 2.5, respectively). However, the posterior distribution for γ in Figure 4 covers well both of these estimates.

When it comes to the remaining coefficients in Table 1, the posterior distributions for ν_{14}, ν_{24} and ν_{34} are shifted further away from zero compared with the priors. These elasticities govern the sensitivity of aggregate output, wages and employment to a given change in females' productivity. The shift is particularly pronounced for output, where most of the posterior mass is located close to the upper bound of the uniform prior. The estimated feedback elasticities for female-specific labor supply, ν_{15}, ν_{25} and ν_{35}, reveal a different pattern. Here, they all move from the prior towards zero in absolute values, reflecting that a given change in females' labor supply has smaller effects on the macroeconomy in data than what our priors would indicate. However, to gauge the quantitative role of females-specific labor supply shocks, we would have to take into account the movement in $\psi_{f,t}$ as well. This is done below. Finally, the employment sensitivity to a given change in automation, ν_{33}, has a posterior centered around -0.4. Taking the normalization of this trend into account, a back-of-the-envelope calculation suggests that most of the labor productivity improvements arising from automation can be attributed to higher output rather than to job destruction.[9]

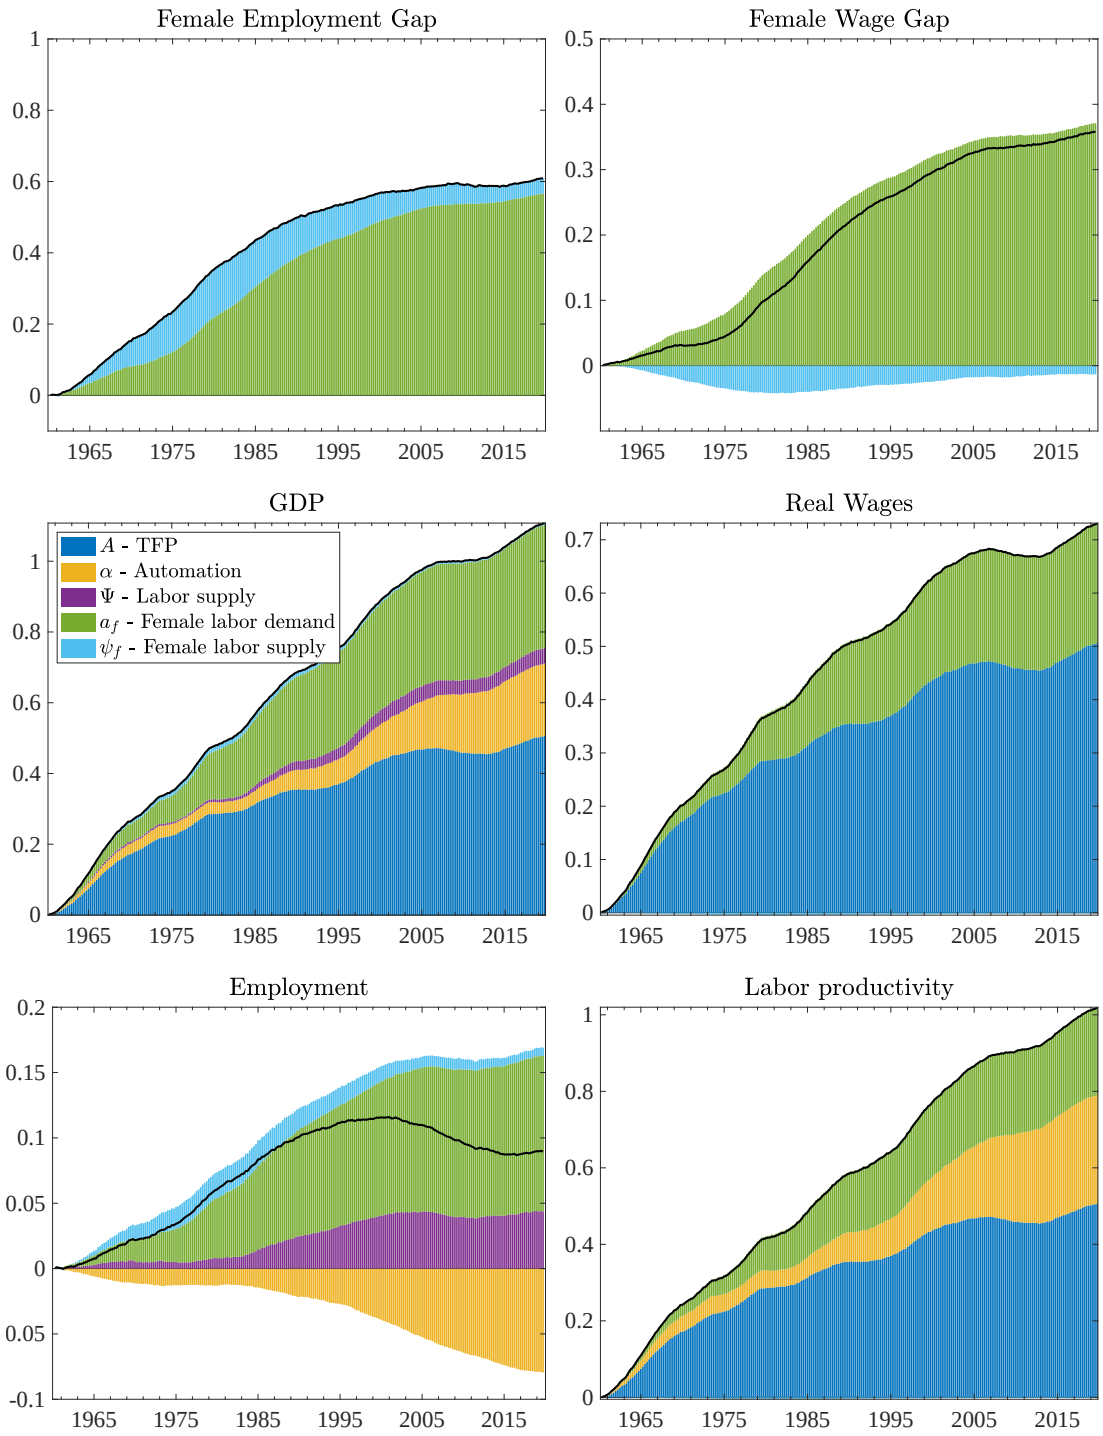
5.3 Trend Decompositions

Estimated contributions of the different structural factors to each empirical trend in our data are documented in Figure 5. Let us first consider the trends in female-to-male employment and wage ratios, which we decompose into female-specific labor demand (green) and female-specific labor supply (light blue). Recall that these are the only structural factors that affect gender gaps in wages and employment in our framework. We find it instructive to distinguish between three separate phases in our sample: in the first 15-20 years of data, both female-specific demand and female-specific supply contributed significantly to a secular increase in the employment rate of females relative to that of males. However, the prominent role of women's labor supply also kept their wage growth relatively modest, explaining why the wage gap between women and men remained somewhat stagnant during this period. Then, starting around 1980, the relative labor supply increase among female workers ceased to take place, causing women's wages to significantly outgrow the wage of males. The final phase in our data started in the late 1990s. Since then, the wage and employment growth among females have been much more modest and more in line with what we observe for men. While some convergence has taken place, the gender gaps have been much more stable in the last 20-25 years of our data. Importantly, our model largely attributes this observation to lower, female-specific productivity growth.

The second and third rows of Figure 5 document how gender-specific labor market factors have affected aggregate trends in the post-war US macroeconomy. Female-specific

[9] A unit innovation to our normalized automation shock raises output by 1 percent and lowers employment 0.4 percent. Thus, labor productivity increases by 1.4 percent, and more than two-thirds of this is due to higher output.

Figure 5: A structural decomposition of the empirical trends



Notes: Empirical trends (black lines) vs. the contributions of individual structural trends (colored bars). Vertical axes represent log deviations from initial values, horizontal axes measures time in quarters. Pointwise median estimates are reported.

productivity growth (green), in particular, accounts for a sizable share of the overall increase in trend GDP and real wages. Moreover, female-specific productivity growth explains most of the postwar rise in employment rates, which would have fallen on average

Table 2: Growth factors in the US economy

A. Baseline account

	(1) GDP						(2) Labor productivity				
	Total	A	Ψ	α	a_f	ψ_f	Total	A	α	a_f	ψ_f
1960-1969	2.8	1.8	0.1	0.3	0.5	0.1	2.5	1.8	0.4	0.3	0.0
1970-1979	2.1	1.1	0.0	0.1	0.8	0.1	1.7	1.1	0.1	0.5	0.0
1980-1989	2.1	0.7	0.2	0.2	1.0	0.0	1.7	0.7	0.3	0.7	0.0
1990-1999	2.0	0.8	0.2	0.4	0.6	0.0	1.8	0.8	0.6	0.4	0.0
2000-2009	1.2	0.3	0.0	0.6	0.3	0.0	1.4	0.3	0.9	0.2	0.0
2010-2019	1.1	0.5	0.0	0.4	0.2	0.0	1.2	0.5	0.6	0.1	0.0

	(3) Aggregate employment					(4) Aggregate wages			
	Total	Ψ	α	a_f	ψ_f	Total	A	a_f	ψ_f
1960-1969	0.3	0.1	-0.1	0.2	0.1	2.1	1.8	0.3	0.0
1970-1979	0.4	0.0	0.0	0.3	0.1	1.6	1.1	0.5	0.0
1980-1989	0.5	0.2	-0.1	0.4	0.0	1.4	0.7	0.7	0.0
1990-1999	0.2	0.2	-0.2	0.2	0.0	1.2	0.8	0.4	0.0
2000-2009	-0.1	0.0	-0.2	0.1	0.0	0.5	0.3	0.2	0.0
2010-2019	-0.1	0.0	-0.2	0.1	0.0	0.6	0.5	0.1	0.0

B. Counterfactual: no gender trends

	(1) GDP						(2) Labor productivity				
	Total	A	Ψ	α	a_f	ψ_f	Total	A	α	a_f	ψ_f
1960-1969	2.7	2.1	0.3	0.3	–	–	2.5	2.1	0.4	–	–
1970-1979	2.1	1.7	0.3	0.1	–	–	1.8	1.7	0.1	–	–
1980-1989	2.1	1.4	0.4	0.3	–	–	1.8	1.4	0.4	–	–
1990-1999	2.0	1.2	0.3	0.5	–	–	1.8	1.2	0.6	–	–
2000-2009	1.0	0.4	0.0	0.6	–	–	1.3	0.4	0.9	–	–
2010-2019	1.1	0.6	0.1	0.4	–	–	1.2	0.6	0.6	–	–

	(3) Aggregate employment					(4) Aggregate wages			
	Total	Ψ	α	a_f	ψ_f	Total	A	a_f	ψ_f
1960-1969	0.2	0.3	-0.1	–	–	2.1	2.1	–	–
1970-1979	0.3	0.3	0.0	–	–	1.6	1.6	–	–
1980-1989	0.3	0.4	-0.1	–	–	1.4	1.4	–	–
1990-1999	0.1	0.3	-0.2	–	–	1.2	1.2	–	–
2000-2009	-0.2	0.0	-0.2	–	–	0.5	0.5	–	–
2010-2019	0.0	0.1	-0.1	–	–	0.6	0.6	–	–

Note: Structural decompositions of the average, annual trend growth rates of (1) real GDP, (2) labor productivity, (3) aggregate employment and (4) aggregate wages (by decade), into total factor productivity A, automation α, gender-neutral labor supply Ψ, female-specific productivity a_f, and female-specific labor supply ψ_f. Panel A: baseline model with gender-specific trend shocks. Panel B: restricted model with dogmatic zero-priors on the elasticities governing feedback from gender-specific shocks to the macroeconomy. All numbers are point-wise median estimates.

since 1960 in the absence of this secular trend. This suggests a rather muted crowding out of men when women enter the labor market, a point which we return to in section 8.

Female-specific productivity is also an important driver of aggregate labor productivity—here measured as output per worker—suggesting that women entering the labor market caused productivity gains beyond the mere scale effects of having a greater labor force. Finally, we note that while the trend in female-specific labor supply (light blue) has contributed to aggregate employment, it has played only a minimal role for GDP, real wages, and aggregate labor productivity. This finding suggests that the periods with extraordinary growth in females' labor supply did not coincide with extraordinary increases in trend GDP, causing the posterior for ν_{15} to shift towards zero (compared with their respective prior distributions).

When it comes to aggregate, gender-neutral factors, total factor productivity (blue) seems to be the quantitatively most important driver of trend GDP, wages and labor productivity in our data. Aggregate labor supply (purple) never plays an important role for GDP, but has stimulated total employment. Finally, labor-displacing automation (yellow) has not only contributed significantly to higher GDP and labor productivity over time, but also to lower employment. In particular, together with the slowdown of gender-specific trends, automation appears to be the main cause of the secular decline in aggregate employment rates over the recent decades. Automation also explains more or less the entire disconnect between real wages and labor productivity in our sample.[10]

Table 2 summarizes our account of economic growth in the postwar US economy. For each decade, we decompose the average, annual growth rates in the trend components of GDP and labor productivity into the estimated contributions of the five structural drivers. For now, we restrict attention to Panel A, which documents the growth accounting implied by the baseline model. Importantly, both GDP growth and labor productivity growth have declined substantially in our data, from about 2.8 and 2.5 percentage points per year in the 1960s, to 1.1-1.4 percentage points in the last 20 years. The decline has mostly taken place in two waves, between the 1960s and 1970s, and at the beginning of the 2000s. Also the growth rate of wages has fallen consistently in our sample, while employment started to stagnate in the 1990s. The slow-down in economic growth (as well as the timing of the two waves) is well-documented and has motivated a large literature on the possible causes, see, e.g. Syverson (2017) for a review.

For the sample as a whole, we find that most of the slow-down is attributed to total factor productivity, which in annual growth terms fell by more than 50% over the first two decades.[11] However, up until the 1980s, lower total factor productivity growth was substantially counteracted by a secular increase in female-specific productivity: its contribution to trend GDP growth went from 0.5 percentage points per year in the 1960s to 1 percentage point in the 1980s. For labor productivity, its contribution increased from 0.3 to 0.7 percentage points per year. Thus, female-specific labor productivity doubled its annual growth rate—and as a result—its contribution to trend growth, between the 1960s and the end of the Cold War.

This picture has changed fundamentally in the last 30 years of data: not only has the overall growth rate of gender-neutral macro trends continued to decline between the

[10]This result implies that the automation trend is the dominant driver of the labor share decline, as in Bergholt et al. (2022). This is visually confirmed in Figure C.3 in the appendix, where we compute and structurally decompose the model-implied trend in the US labor income share.

[11]The annual growth rate in total factor productivity fell from 1.8 percentage points in the 1960s to 0.7 percentage points in the 1980s.

1990s and the 2010s, but also two-thirds of female-specific labor productivity growth has disappeared during this period. The latter result is imperative for growth accounting in recent decades according to our model: between the 1990s and 2010s, the slow-down in female-specific labor productivity is responsible for about half of the overall decline in the growth rates of trend GDP and labor productivity.

Overall, we arrive at three main takeaways from the estimation of our empirical model: first, gender-specific labor market trends are quantitatively important for the US macroeconomy. For example, they account for almost one-third of the overall postwar increase in trend GDP, and more than one-fifth of the postwar increase in labor productivity. Second, about fifty percent of the slowdown in economic growth observed in the last 25 years is attributed to a slowdown in the employment convergence between females and males. Third, the catch-up of females' employment observed in the last 60 years is mainly, if not entirely, a consequence of labor demand factors. Notably, all these results are confirmed also when using a prior specification tilted against macro effects of gender shocks, as shown in Appendix D.1.

6 Taking stock

The important impact of gender-specific labor market trends for macroeconomic growth documented in this paper is consistent with the analyses by Hsieh et al. (2019) and Heathcote et al. (2017) in the context of calibrated models. Hsieh et al. (2019) develop a Roy model of occupational choice and conclude that declining obstacles to human capital accumulation, as well as reduced discrimination, may explain around half of GDP per-capita growth between 1960 and 2010. Heathcote, Storesletten, and Violante (2010) study a neoclassical growth model with incomplete markets and overlapping generations, and find that female-specific demand factors explain most of the increase in females' labor supply, which in turn drove half of the growth in earnings per capita between 1967 and 2002 (Heathcote et al., 2017). While the conclusions reached in these papers seem largely consistent with ours, we take a quite different approach by estimating an empirical time series model with substantially fewer cross-equation restrictions.

6.1 Why is female-specific productivity important?

To better understand why the data prefer such an important role for female-specific productivity growth, we find it informative to confront the implications of trend shocks in our estimated model with empirical patterns in data.

The first important piece of information is the observed, common rise of women's wages and employment relative to males. Recall that, by construction, none of the gender-neutral macro shocks can account for this empirical feature. Moreover, equations (24)-(25) demonstrate that the co-movement between wage and employment gaps naturally arises from gender-specific labor productivity $a_{f,t}$, as opposed to gender-specific labor supply $\psi_{f,t}$. Notably, the pre-1975 period represents an exception. Then the wage gap was rather stagnant compared with the employment gap, suggesting that female-specific labor supply must have played an important role as well. This is indeed captured in our trend estimates, as illustrated in Figure 5.

To appreciate why the estimation procedure also chooses female-specific productivity as a driver of macroeconomic variables such as GDP, it is important to understand how this trend departs from the gender-neutral alternatives. A first natural comparison is with total factor productivity, the largest contributor to macroeconomic growth according to our model. This is the only macroeconomic driver that can jointly capture the prominent upward trends of GDP, labor productivity, and aggregate wages in data. Automation cannot account for the trend in wages, while neither the wage trend nor the labor productivity trend can be accounted for by gender-neutral labor supply. However, total factor productivity implies a balanced growth path in our framework, where GDP, labor productivity and wages all *grow at the same rate*.[12] In US macro data, instead, GDP and labor productivity have tended to outgrow aggregate wages. The only other macroeconomic driver that can capture this phenomenon is automation. However, automation by itself has limited explanatory power because it implies (i) an extreme wage disconnect where wages do not grow at all, and (ii) an overall decline in aggregate employment. While automation seems likely to be important for GDP and labor productivity in periods with falling employment rates, both of these implications are on average at odds with the last 60 years of US labor market data.

Female-specific productivity speaks to all of these observations. Qualitatively, female-specific productivity behaves similarly to total factor productivity in the sense that both *GDP, labor productivity and wages rise* in response to the shock. But quantitatively we allow female productivity to have different effects across these three variables, in contrast to total factor productivity. The special case of balanced growth following changes in females' productivity, can be captured in our empirical framework with the parametric restrictions $\nu_{14} = \nu_{24}$ and $\nu_{34} = 0$. While these restrictions are satisfied at the *prior* mean in our baseline specification[13] (see Table 1), the posterior parameters are updated to explain trends in data. Table 1 summarizes how the model chooses to quantitatively match female-specific productivity with our macroeconomic time series: at the posterior mean, a shock to women's productivity—normalized to increase the relative wage of women by 1 percent—is associated with a 0.9 percent increase in GDP, a 0.6 percent increase in the aggregate real wage, and a 0.3 percent increase in the employment rate. Thus, higher female-specific productivity does not only account for the joint rise of these three variables over the full sample, but also for the slow-down in wage growth compared with GDP, as well as the increase in trend employment. Importantly, these effects are quantitatively relevant in our estimation because larger-than-normal growth in the wage of women tends to coincide with larger-than-normal growth in GDP, labor productivity, and aggregate wages, as well as a rise in aggregate employment. In the same vein, the flattening of the wage gap trend in the early 2000s coincides with lower income growth and a reversal of the aggregate employment rate.

Finally, we note that $\nu_{14} - \nu_{34} \approx \nu_{24}$ at the posterior mean, implying that female-specific productivity causes labor productivity and wages to respond similarly. In turn, this means that female-specific productivity has a negligible effect on the overall labor

[12] Critically, we consider deviations from balanced growth in Appendix D.2. The role of gender forces is even larger in that specification.

[13] Thus, at the prior mean, female-specific productivity shocks are separately identified from total factor productivity solely because the latter comes with zero-restrictions on the responses of gender gap variables.

income share. We conjecture that this is preferred by data because shifts in the labor income share, which are rather small and concentrated in time, tend to be orthogonal to the permanent changes in wage and employment gaps in our sample.

Next, we ask the question "how would the growth accounting change if we were to ignore the gender trends in US labor market data?" The answer to this question provides the last piece of information in understanding why female-specific productivity matters for economic growth. Thus, we now re-estimate the model but impose dogmatic zero-priors on the elasticities ν_{14}, ν_{15}, ν_{24}, ν_{25}, ν_{34}, and ν_{35}. This implies that the two gender-specific trends only affect female-to-male differences in wages and employment, while total factor productivity, gender-neutral labor supply and automation are forced to account for all trends in aggregate macro data.

Panel B in Table 2 reports the results from this exercise. Two important observations stand out: first, decade-specific trend growth rates in GDP and labor productivity are almost unaffected by the zero-restrictions imposed on the two gender-specific trends. Thus, it is the structural composition of different drivers, and not the level of economic growth, that is affected by these restrictions. Second, regarding the structural composition, the contribution attributed to gender trends in our baseline (see Panel A) is now mostly shifted to total factor productivity, while the contribution attributed to automation remains almost unchanged. As an example, for the full sample our baseline results imply that the two gender shocks on average account for just over 30% of total GDP growth, and just over 20% of labor productivity growth. The corresponding numbers for total factor productivity are 43% and 48%, respectively. In Panel B the contributions of total factor productivity have increased to 63% for GDP and 67% for labor productivity, while the gender trends attribute zero by assumption. The intuition for this finding is discussed above: compared with total factor productivity, female-specific productivity growth helps to understand not only the secular increase in females' wages and employment, but also the joint but uneven rise of GDP, wages and employment in US macro data.

6.2 JOBLESS RECOVERIES

Our analysis reveals that a halt in women's labor market catch-up has contributed significantly to the observed slowdown in US economic growth in recent decades. This has important implications for macroeconomic recovery after recessions. Albanesi (2024), Fukui et al. (2023) and Olsson (2024) find that lower trend growth, rather than changes in the cyclical properties of the economy, may explain the slow employment recoveries observed after recent recessions (the recovery may naturally appear slower in data if the underlying, unobservable employment trend shifts down). Our model disentangles a few drivers of the employment trend, as shown in Figure 5 and panel A.3 in Table 2. The average growth rate of trend employment declined by 0.3 percentage points between the 1980s and 1990s, and by an additional 0.3 percentage points between the 1990s and the 2000s. About half of this decline is explained by a slowdown in female-specific labor demand in our framework. This finding is at least qualitatively in line with Albanesi (2024), Fukui et al. (2023) and Olsson (2024). The remaining half is attributed, first, to accelerating automation in the 1990s, and second, to lower gender-neutral labor supply in the 2000s. These results seem consistent with Jaimovich and Siu (2020), who emphasize job

polarization as an explanation for jobless recoveries.[14] But they are also fully consistent with Pugsley and Şahin (2019), who document a link between the decline in entrepreneurship and aggregate employment dynamics. They find that changing worker demographics and a secular increase in import competition are the most promising explanations for the decline in entrepreneurship. Arguably, the labor supply shock and the automation shock are likely to capture these forces in our framework.

7 Trends in skills, the rise of services, and compositional effects

Given the simplicity of our baseline set-up, aggregate forces like skill-biased technological progress or the process of structural transformation from goods to services may be captured as gender-specific shocks as long as they affect the two genders differently (Cerina, Moro, and Rendall, 2021). In this section, we want to control for the role of skills and sectors and check whether genuine gender forces are still important for the macroeconomy and whether we confirm the dominant role for demand-side factors to explain the gender convergence.

7.1 Quantifying the role of skills

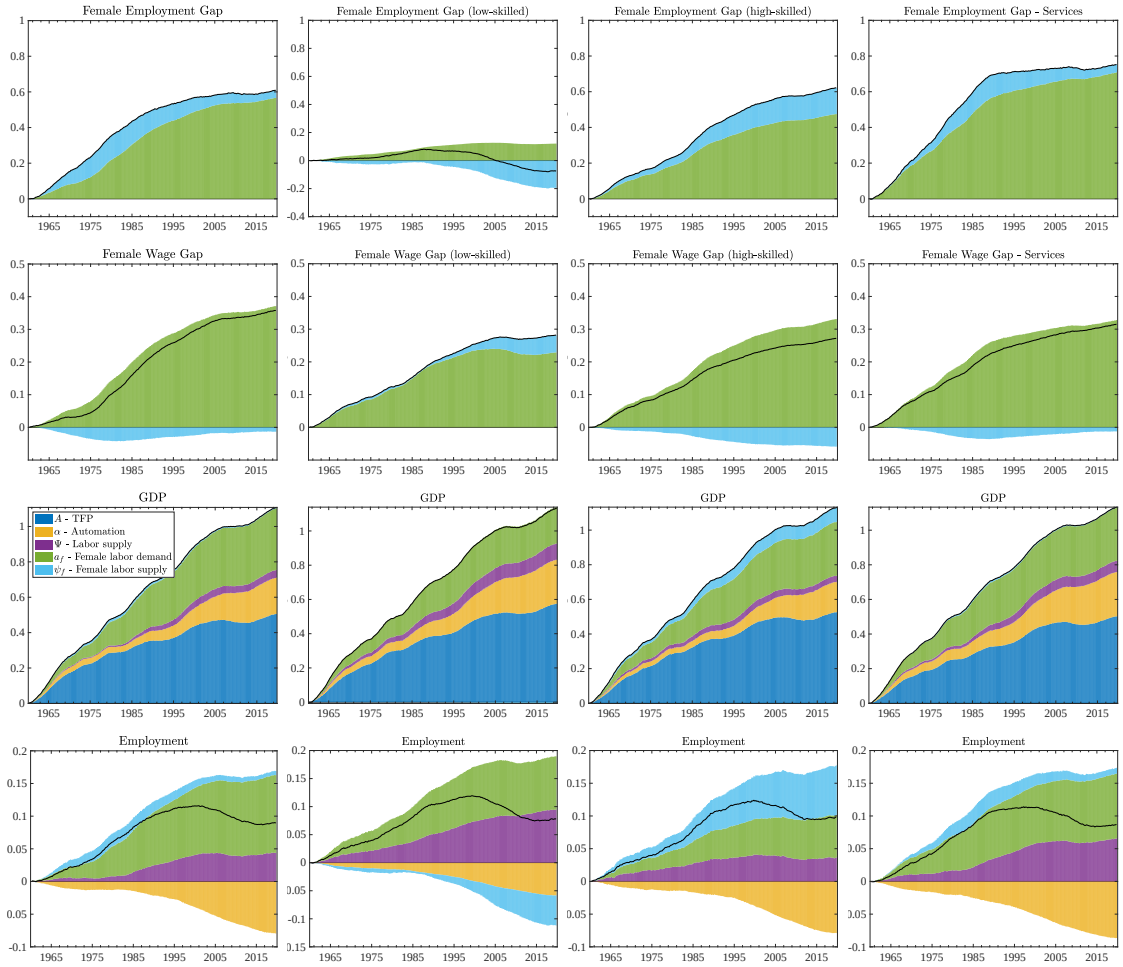
To understand how trends in skills may have affected gender gaps in the US labor market, we use data on individuals' wages and employment, conditional both on gender and the level of education level from the Current Population Survey (CPS).[15] Building on Dolado et al. (2021), we partition workers' skills by their level of education: individual high-skilled workers are characterized as those who have at least some college experience, while low-skilled workers do not have any college experience. Appendix B explains how we fetch quarterly, seasonally adjusted time series from the CPS dataset, and also plots the data series that we construct and use.

We re-estimate the empirical model using data on wage and employment gaps (between female and male workers) *within* each skill segment: compared with the baseline specification in section 5, we replace the aggregate wage and employment gap variables with their counterparts for high- and low-skill workers, respectively. The macroeconomic variables—aggregate GDP, wages and employment—are kept as in the baseline model. Moreover, the identification scheme is as before, but now with restrictions on skill-biased gender gaps in the last two columns and rows of $\mathcal{V}$. For example, an increase in high-skilled females' productivity (relative to high-skilled males' productivity) permanently raises both the wage and employment rate of high-skilled females, compared with that of high-skilled males. An increase in the relative supply of high-skilled females, in contrast,

[14] Autor, Levy, and Murnane (2003) and a large subsequent literature demonstrates that job polarization is primarily due to progress in technologies that substitute for labor in performing routine tasks. Similar effects can be induced by outsourcing and offshoring. All these forces are in principle captured by our automation trend, which can be broadly interpreted as representing general, labor-saving technological progress.

[15] The CPS is a monthly survey published by the United Census Bureau and the Bureau of Labor Statistics that provides an extensive body of data on households' demographic and labor force characteristics as well as employment, earnings and hours of work, starting from 1979.

Figure 6: *Structural* drivers of the empirical trends – accounting for skills and services



Notes: The colored bars display the point-wise median evolution of the empirical trends attributable to each structural trend. Colors legend: TFP A (blue); automation α (yellow); neutral labor supply Ψ (purple); female-specific labor demand a_f (green); female-specific labor supply ψ_f (light-blue). From left to right: first column, all workers (baseline); second column, low-skilled workers; third column, high-skilled workers; fourth column, workers employed in service industries.

raises the female share of high-skilled employment but lowers their wages compared with high-skilled males. This exercise allows us to quantify the role of shocks to the gender trends *within* each skill segment, thus, factoring out the fixed skill effects on aggregate gender gap trends.

Estimated empirical and structural trends for wage and employment gaps by skills are shown in Figure 6. We also include the baseline decompositions of aggregate gender gap variables in the left column to facilitate comparison. Consider, first, the estimated empirical trends: interestingly, while the ratio of female-to-male employment among low-skilled workers has trended slightly down on average, the relative employment of females among high-skilled workers displays a clear upward trend in our data. And this trend is quantitatively very similar to its aggregate counterpart. Moreover, female wages have outgrown male wages in both low-skilled and high-skilled jobs.

A couple of remarks are in place: first, these results speak against the idea that fundamental, gender-neutral trends in skills can account for the secular rise of female labor.

If such gender-neutral trends had been the main force, we would have expected secular reallocation towards high-skilled jobs for both females and males, with limited gender convergence in wages and employment *within* each skill segment. Instead, the wage and employment trends of females have evolved quite differently than those of males even after controlling for skills, especially in the market for high-skilled jobs.

Second, compared to the aggregate wage gap trend, we find that the catch-up of female wages has been less pronounced *both* in low-skilled and high-skilled jobs. This observation suggests that compositional effects may play an important role. Women in particular have moved into skill-intensive occupations at such a pace that the overall female share in high-skill jobs has increased significantly. Since high-skill jobs generally tend to come with a skill premium on the wage rate, this contributes to a secular increase in females' wages over time, above and beyond that accounted for by gender trends within the skill segments.

The *structural* decompositions reported in Figure 6 shed further light on empirical trends in our data: the market for low-skilled labor, in particular, is characterized by a contraction of females' labor supply relative to that of males over time. This contraction has been especially relevant since the 2000s, which is when the decline of females' relative employment in low-skilled jobs started to take place. In fact, our results imply that the ratio of low-skilled female-to-male employment would have increased, rather than decreased, in the absence of a falling supply of low-skilled female labor. Then the gender differences in low-skilled employment would have been determined solely by the upward trend in female-specific demand, which has been present in both low- and high-skilled labor markets. The high-skill labor market, in contrast, displays a more prominent role for female-specific demand. Moreover, both demand and supply forces have contributed to a relative increase in females' employment in high-skill jobs. But the role of female-specific supply has been quantitatively more important for high skills than in the aggregate where we abstract from skills altogether. This is likely to capture a large increase in educational attainment among women in our data. In fact, the average education level has been higher for women than for men since 1993, and the gap has continued to grow since then (see Albanesi and Şahin (2018) for more details).

Figure 6 also reports the estimated empirical and structural trends in aggregate GDP and employment when we substitute data on aggregate wage and employment gaps with their counterparts for low- and high-skilled labor, respectively. The empirical macro trends are similar across specifications. However, the decomposition into structural drivers changes significantly: when the model is informed by data on low-skilled wage and employment gaps, it suggests that female-specific supply of low-skilled labor has been contractionary for aggregate employment on average. Disciplined by data on high-skill wages and employment, instead, the model concludes that female-specific supply of high-skilled labor has been responsible for nearly half of the total increase in the post-war, aggregate employment rate. Interestingly, the model with low-skill labor assigns the bulk of the fall in aggregate employment during the 2000s to an accelerating contraction in low-skilled labor supply of females. The model with high-skill labor, instead, suggests that aggregate employment fell because the supply-driven increase of high-skilled females seen in earlier decades leveled off, while at the same time automation started to take place at a larger scale. In sum, our partitioning of the data by gender and skills uncovers important reallocation patterns: increased supply of high-skilled females during the last three decades

has coincided and largely been counteracted by declining supply of low-skilled females. Thus, reallocation trends in females' labor supply have been quantitatively important for U.S. labor markets, despite their apparent disconnect with macroeconomic aggregates when we abstract from skills as in our baseline model. In that sense, this extension provides an important caveat to our baseline model: labor demand is still the main driver of the gender convergence, but gender-specific labor supply trends are more important once skills are taken into account.

7.2 THE SECULAR RISE OF THE SERVICE SECTOR

The U.S. economy has gone through a major, sectoral transformation in the last decades. Approximately 60% of the total workforce was employed in the service sector in the 1980s according to our CPS data.[16] By the end of 2019, this share has increased beyond 75%. At the same time, the employment share in manufacturing declined from 30% in the 1980s to 10% in 2019. Ngai and Petrongolo (2017) document that the rise of services has paved the way for female labor. The idea is that service production requires more *cognitive* skills in which females have a comparative advantage, as opposed to the need for *brawn* skills in the production of manufactured goods.[17]

One open question is whether the process of structural transformation from goods to services can account in and of itself for gender trends or whether within-sector gender specific forces are needed to match the data. To answer this question, we replace the aggregate gender gaps in employment and wages with their counterparts in the service sector in our model.

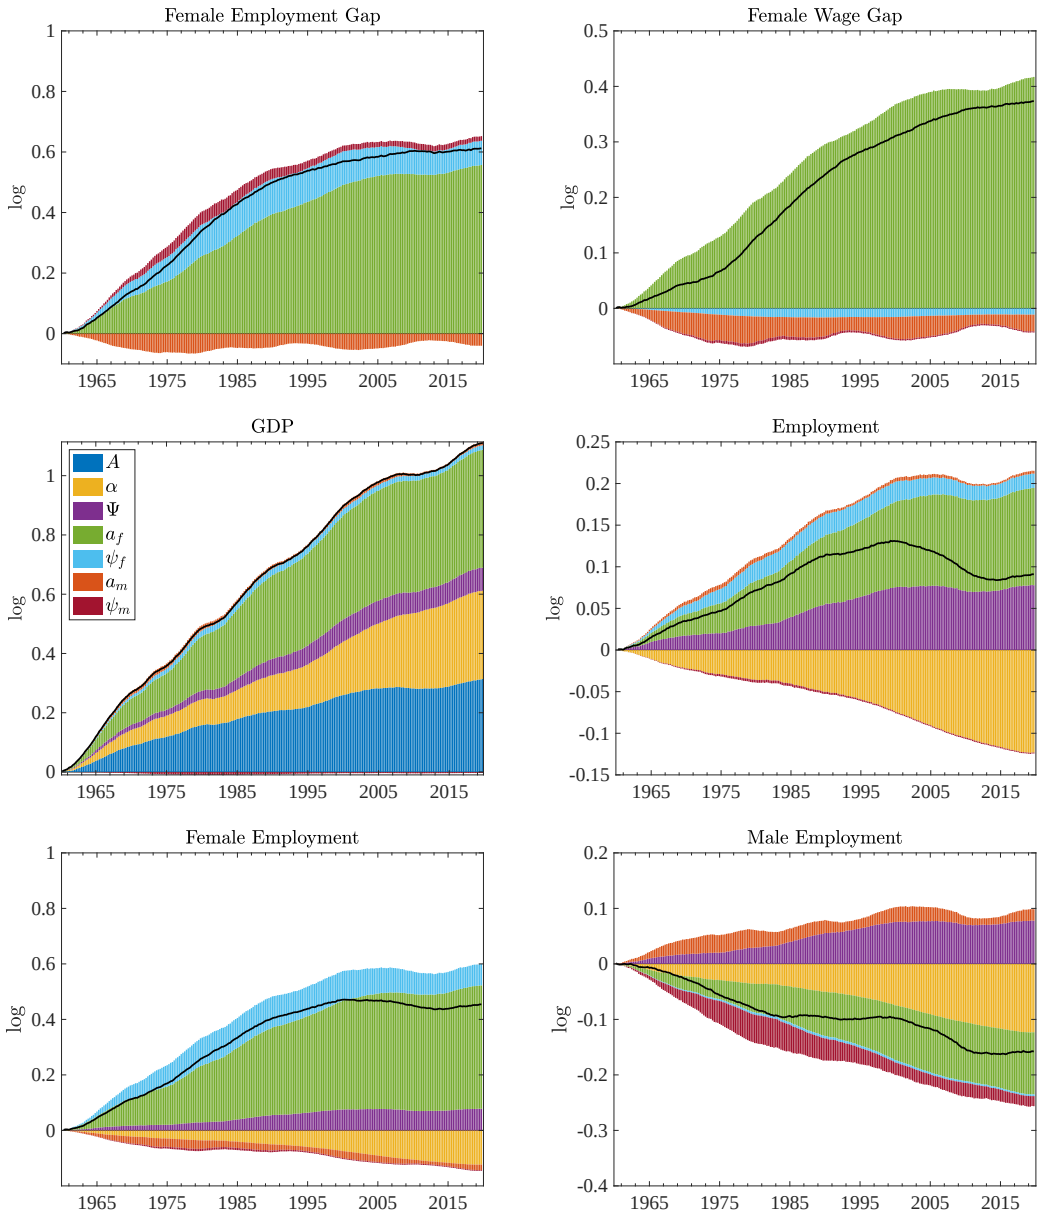
Figure 6 summarizes the results. The subplots in the first two rows compare estimated gender gap trends in the total economy (first column from the left) with those in the service sector alone (last column). We find evidence of a strong gender convergence within the service sector and this gender convergence is mainly driven by labor demand factors. In addition, the effects on GDP and employment are comparable the baseline model, as shown in last two rows in Figure 6. This means that the rise of the service sector is not per se the main reason for why gender matters for growth: within-sector labor demand forces like gender-specific technological progress (Heathcote et al. (2010) and Heathcote et al. (2017)) or reduction in discrimination (Jones et al. (2015)) play an important role to account for the observed trends in gender ratios. Notably, within-sector labor demand forces are crucial also in Ngai and Petrongolo (2017) to match the data on gender ratios. We recover this result in the context of a flexible estimated (rather than calibrated) model.

The last two rows in Figure 6 (fourth column) document how gender-specific trends within the service sector have affected the macroeconomy. Overall, the main takeaways from section 5 remain, with female-specific labor demand being the *only* significant gen-

[16] The CPS survey contains information on the industries in which households supply labor. Therefore, we use information on more than 35 private industries and classify them into aggregate *manufacturing* and *service* sectors. Appendix B provides an overview on the industry-specific data we use, including the classification into manufacturing and services.

[17] Kuhn, Manovskii, and Qiu (2024) explore the opposite direction of causality and find an important impact of the rise of female employment on the process of structural transformation in a model with gender complementarity in sectoral production. Buera and Kaboski (2012) and Buera et al. (2019) study the role of skill-biased technological change to explain the rise of the service sector in models with and without a family structure.

Figure 7: Adding male-specific trends to the system



Notes: The colored bars display the point-wise median evolution of the empirical trends attributable to each structural trend. Colors legend: TFP A (blue); automation α (yellow); neutral labor supply Ψ (purple); female-specific labor demand a_f (green); female-specific labor supply ψ_f (light-blue); male-specific labor demand a_m (red); male-specific labor supply ψ_m (bordeaux).

der trend for the long-run increase in real GDP.

8 MALE-SPECIFIC LABOR MARKET TRENDS

Our maintained assumption so far has been that the labor market convergence between females and males has been driven by a secular change in females' productivity and labor supply, while male-specific productivity and labor supply have been kept fixed.

Now we relax the above-mentioned assumption and allow for the presence of male-specific trend shocks $A_{m,t}$ and $\Psi_{m,t}$, in addition to the female-specific shocks $A_{f,t}$ and $\Psi_{f,t}$. Contractionary, male-specific shocks may, for example, capture the impact of video gaming and recreational computing on the labor supply of young men, as discussed in Aguiar, Bils, Charles, and Hurst (2021). The presence of male-specific trends implies that the wage and employment gaps $w_{f,t}$ and $l_{f,t}$ are driven by four gender trends in total, and we need to impose additional identification restrictions on the system. To this end, we augment the baseline empirical model with two more observables—the male employment rate and male wages—both observed in log-levels. The joint restrictions that we impose on female-to-male employment and wage gaps, as well as on the levels of males' employment and wages, allow us to identify all of the gender-specific trends in the system. Details about the identification strategy and the theory-consistent priors are discussed in Appendix E where we present also impulse responses to male-specific shocks in the theoretical model which is once again our reference to set the priors in the empirical model.

Results are presented in Figure 7. The first row plots the decomposition of trend employment and wage gaps between females and males. Both gaps are driven almost exclusively by female-specific shocks. Thus, ignoring male-specific shocks was largely inconsequential. The second row in Figure 7 reveals that female-specific shocks also remains an important driver of GDP and aggregate employment. The two male trends, in contrast, are never important at the aggregate level.

Finally, the third row plots the decomposition of trends in the levels of female and male employment. A few comments are in place: first, the disappearance of aggregate employment growth observed in the last 20 years is entirely driven by female labor which stopped growing around 2000. The employment rate of males, instead, has declined every decade since 1960 up until the financial crisis, and has since then been relatively flat. Second, our account of trends in the female employment rate is similar to that of aggregate employment, albeit with a relatively larger role for female-specific productivity and labor supply (as opposed to the gender-neutral shocks). This is not surprising, given that these shocks have a smaller weight in aggregate employment. Third, the specification with males can shed light on whether secular increases in female employment crowd out male labor. The crowding-out elasticity is a key statistic in Fukui et al. (2023). On average over the sample, an increase in female employment of 1 percentage point, when driven by female-specific demand, leads to a decline in male employment of around 0.30 percentage points according to our estimates. Such a muted crowding out elasticity implies relatively large effects on aggregate employment and economic activity when women enter the labor market. The corresponding crowding out elasticity conditional on female-specific labor supply is even smaller, around 0.1. Part of the difference could be that the gender-specific productivity shocks imply stronger income effects on spouses' labor supply because they have a greater impact on the family's total income. This would be in line with our theoretical model, as shown in Figure E.1 in the Appendix. By comparison, Fukui et al. (2023) who do not make an explicit distinction between demand and supply-driven forces in their empirical section, report a value of 0.18. However, they refer to "relative" crowding out elasticities across regions with different exposure to gender trends, making a comparison less straightforward.

Overall, we conclude that i) male shocks play a minor role, ii) our model estimates a

rather small degree of crowding out, consistently with the large macro effects of gender shocks and iii) estimates of the degree of crowding out are shock-specific, a point that to the best of our knowledge is novel.

9 Conclusion

In this paper, we investigate and quantify the implications of gender-specific labor market trends for the U.S. macroeconomy. Using a SVAR model with common trends, we document 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 labor market trends account for up to 50% of trend growth in GDP over the period 1960-1990. Furthermore, the flattening of the gender convergence which started in the 1990s is key for the marked slowdown in trend growth observed over the last 25 years.

Importantly, we document that gender differences matter for the macroeconomy using a pure "macro" approach: an empirical time series model is disciplined by neoclassical theory and estimated on selected macroeconomic variables. In that sense, our "let the data speak" approach is complementary yet very different from the more heavily parametrized structural models, such as those put forward by Hsieh et al. (2019), Heathcote et al. (2010) and Heathcote et al. (2017). Interestingly, and somewhat unexpectedly, our flexible setup reaches very similar conclusions on quite granular *quantitative* results, like the implications of gender convergence for growth and productivity.

Regarding future growth prospects, one possible concern is that the muted growth since the 2000s represents a new normal unless the labor market participation among women starts accelerating again. Is this likely to happen? On one side, gender differences in the US labor market are still sizable, and experiences from other countries (see Albanesi, Olivetti, and Petrongolo (2023) for an international comparison) suggest that ample pockets of growth may still be available if the right institutional features are put in place.[18] However, further growth could also prove more difficult than in earlier decades given that females' employment rates are much higher now, and given the much smaller gap between females and males compared with the past. After all, policies cannot improve economic performance unboundedly, as the labor force participation rate has a natural upward bound of 100%. The extent to which labor market participation among women could start to rise again ultimately depends on why it stopped in mid-2000s, and the jury is still out on this question.[19]

[18] In addition, one can imagine that other long-lasting sources of growth can be exploited by addressing other forms of misallocation. For example, wage and employment gaps between native and migrants are still far from closed and there is substantial evidence of skill downgrading of migrants (Dustmann, Frattini, and Preston (2013)).

[19] A slowdown in gender convergence can be associated with cultural factors (Fogli and Veldkamp, 2011; Fernández, 2013), lack of family-friendly policies (Blau and Kahn, 2013), and increased income inequality inducing negative income effects on women married to high-earning husbands (Albanesi and Prados, 2022). Goldin (2014) and Goldin (2021) argue that the gender wage gap would be reduced further if firms did not disproportionately reward workers for long hours and duties difficult to plan in advance. Erosa, Fuster, Kambourov, and Rogerson (2022) validate this view in a Roy model with occupation-specific non-convex earnings functions.

While we believe that the application to the gender convergence is particularly interesting, our methodology can be applied to an array of questions concerning other secular trends as well. Examples include demographics, climate change, sectoral trends, immigration, as well as linkages between growth and inequality. In addition, our framework can be used to study gender differences at business cycle frequencies (cf. Albanesi (2024) and Albanesi and Şahin (2018)). We plan to investigate some of these topics in future research.

REFERENCES

Acemoglu, D., D. H. Autor, and D. Lyle (2004). Women, war, and wages: The effect of female labor supply on the wage structure at midcentury. *Journal of Political Economy 112*(3), 497–551.

Acemoglu, D. and P. Restrepo (2020). Robots and jobs: Evidence from US labor markets. *Journal of Political Economy 128*(6), 2188–2244.

Aguiar, M., M. Bils, K. K. Charles, and E. Hurst (2021). Leisure luxuries and the labor supply of young men. *Journal of Political Economy 129*(2), 337–382.

Aguiar, M. and E. Hurst (2007). Measuring trends in leisure: The allocation of time over five decades. *Quarterly Journal of Economics 122*(3), 969–1006.

Albanesi, S. (2024). Changing business cycles: The role of women's employment. *American Economic Journal: Macroeconomics, forthcoming*.

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

Albanesi, S., C. Olivetti, and B. Petrongolo (2023). Families, labor markets, and policy. In *Handbook of the Economics of the Family*, Volume 1, pp. 255–326. Elsevier.

Albanesi, S. and M. J. Prados (2022). Slowing women's labor force participation: The role of income inequality. NBER Working Paper 29675.

Albanesi, S. and A. Şahin (2018). The gender unemployment gap. *Review of Economic Dynamics 30*, 47–67.

Ascari, G. and L. Fosso (2024). The international dimension of trend inflation. *Journal of International Economics 148*.

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

Autor, D. H., F. Levy, and R. J. Murnane (2003). The skill content of recent technological change: An empirical exploration. *Quarterly Journal of Economics 118*(4), 1279–1333.

Beaudry, P. and E. Lewis (2014). Do male-female wage differentials reflect differences in the return to skill? Cross-city evidence from 1980–2000. *American Economic Journal: Applied Economics 6*(2), 178–194.

Bergholt, D., F. Furlanetto, and N. Maffei-Faccioli (2022). The decline of the labor share: New empirical evidence. *American Economic Journal: Macroeconomics 14*(3), 163–198.

Bianchi, F., G. Nicolò, and D. Song (2023). Inflation and real activity over the business cycle. NBER Working Paper 31075.

Blau, F. D. and L. M. Kahn (2013). Female labor supply: Why is the United States falling behind? *American Economic Review 103*(3), 251–256.

Boppart, T. and P. Krusell (2020). Labor supply in the past, present, and future: a balanced-growth perspective. *Journal of Political Economy 128*(1), 118–157.

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

Buera, F. J., J. P. Kaboski, and M. Q. Zhao (2019). The rise of services: The role of skills, scale, and female labor supply. *Journal of Human Capital 13*(2), 157–187.

Cerina, F., A. Moro, and M. Rendall (2021). The role of gender in employment polarization. *International Economic Review 62*(4), 1655–1691.

Cortes, G. M., N. Jaimovich, and H. E. Siu (2023). The growing importance of social tasks in high-paying occupations: Implications for sorting. *Journal of Human Resources 58*(5), 1429–1451.

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

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

Del Negro, M. and F. Schorfheide (2004). Priors from general equilibrium models for VARs. *International Economic Review 45*(2), 643–673.

Doepke, M. and M. Tertilt (2009). Women's liberation: What's in it for men? *Quarterly Journal of Economics 124*(4), 1541–1591.

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

Durbin, J. and S. J. Koopman (2002). A simple and efficient simulation smoother for state space time series analysis. *Biometrika 89*(3), 603–615.

Dustmann, C., T. Frattini, and I. P. Preston (2013). The effect of immigration along the distribution of wages. *Review of Economic Studies 80*(1), 145–173.

Erosa, A., L. Fuster, G. Kambourov, and R. Rogerson (2022). Hours, occupations, and gender differences in labor market outcomes. *American Economic Journal: Macroeconomics 14*(3), 543–590.

Fernández, R. (2013). Cultural change as learning: The evolution of female labor force participation over a century. *American Economic Review 103*(1), 472–500.

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

Fogli, A. and L. Veldkamp (2011). Nature or nurture? Learning and the geography of female labor force participation. *Econometrica 79*(4), 1103–1138.

Fukui, M., E. Nakamura, and J. Steinsson (2023). Women, wealth effects, and slow recoveries. *American Economic Journal: Macroeconomics 15*(1), 269–313.

Galor, O. and D. N. Weil (1996). The gender gap, fertility, and growth. *American Economic Review 86*(3), 374–387.

Giannone, D., M. Lenza, and G. E. Primiceri (2015). Prior Selection for Vector Autoregressions. *Review of Economics and Statistics 97*(2), 436–451.

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

Goldin, C. (2014). A grand gender convergence: Its last chapter. *American Economic Review 104*(4), 1091–1119.

Goldin, C. (2021). *Career and family: Women's century-long journey toward equity*. Princeton University Press.

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

Goux, D., E. Maurin, and B. Petrongolo (2014). Worktime regulations and spousal labor supply. *American Economic Review 104*(1), 252–276.

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

Hasenzagl, T., F. Pellegrino, L. Reichlin, and G. Ricco (2022). A model of the Fed's view on inflation. *Review of Economics and Statistics 104*(4), 686–704.

Heathcote, J., K. Storesletten, and G. L. Violante (2010). The macroeconomic implications of rising wage inequality in the United States. *Journal of Political Economy 118*(4), 681–722.

Heathcote, J., K. Storesletten, and G. L. Violante (2017). 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.

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

Jaimovich, N. and H. E. Siu (2020). Job polarization and jobless recoveries. *Review of Economics and Statistics 102*(1), 129–147.

Jones, L. E., R. E. Manuelli, and E. R. McGrattan (2015). Why are married women working so much? *Journal of Demographic Economics 81*(1), 75–114.

Kuhn, M., I. Manovskii, and X. Qiu (2024). Female employment and structural transfor-

mation. CEPR Discussion Paper 19231.

Maffei-Faccioli, N. (2024). Identifying the sources of the slowdown in growth: Demand vs. supply. *Journal of Applied Econometrics, forthcoming*.

Ngai, L. R. and B. Petrongolo (2017). Gender gaps and the rise of the service economy. *American Economic Journal: Macroeconomics 9*(4), 1–44.

Olsson, J. (2024). Singles, couples, and their labor supply: Long-run trends and short-run fluctuations. *American Economic Journal: Macroeconomics, forthcoming*.

Pugsley, B. W. and A. Şahin (2019). Grown-up business cycles. *Review of Financial Studies 32*(3), 1102–1147.

Rendall, M. (2024). Brain versus brawn: The realization of women's comparative advantage. CEPR Discussion Paper 18825.

Stock, J. H. and M. W. Watson (1988). Testing for common trends. *Journal of the American Statistical Association 83*(404), 1097–1107.

Stock, J. H. and M. W. Watson (2007). Why has US inflation become harder to forecast? *Journal of Money, Credit and Banking 39*(s1), 3–33.

Syverson, C. (2017). Challenges to mismeasurement explanations for the US productivity slowdown. *Journal of Economic Perspectives 31*(2), 165–186.

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

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

Weinberg, B. A. (2000). Computer use and the demand for female workers. *Industrial and Labor Relations Review 53*(2), 290–308.

APPENDIX

A A BVAR WITH COMMON TRENDS

This section discusses the prior assumptions we formulate on the free parameters of the model presented in section 2 – including the assumptions on the prior volatilities of the structural trends' shocks – and the instructions to estimate the model with a Gibbs sampling algorithm.

A.1 PRIOR ASSUMPTIONS

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, we can use the information on $\bar{Y}_0$ – the initial conditions of the empirical trends[20] – as well as on the prior coefficients in $\mathcal{V}$. Then, one can retrieve $\underline{X_0}$ by solving the system in eq. (2), provided that the number of structural trends $q = n$, as it is the case in our model. 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 coefficients are distributed according to:

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

$$\tilde{\Phi}|\Sigma_e \sim \mathcal{N}(\underline{\tilde{\Phi}}, \Sigma_e \otimes \underline{\Omega})\mathcal{I}(\tilde{\Phi}) \tag{A.2}$$

$$\Sigma_u \sim \mathcal{IW}(\kappa_u, (\kappa_u + n + 1)\underline{\Sigma_u}) \tag{A.3}$$

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. The prior on the lag coefficients is standard Minnesota with mean zero and overall tightness hyperparameter equal to 0.2 (Giannone, Lenza, and Primiceri (2015)). $\mathcal{IW}$ is the Inverse-Wishart distribution with κ degrees of freedom and mode $\underline{\Sigma}$. The prior on the transitory innovations is rather loose, with degrees of freedom $\kappa_e = n + 2$ to ensure the existence of the mean and the prior mode $\underline{\Sigma_e} = I$.

Next, the prior on the trends' shocks $\underline{\Sigma_u}$ is distributed according to an Inverse-Wishart, as well. The prior is rather tight, as we set the degrees of freedom $\kappa_u = 100$. Finally, the prior mode $\underline{\Sigma_u}$ is assumed to be diagonal. One non-trivial task is to come up with reasonable priors for the elements of $\sigma^2_{\mathbf{u}} = [\sigma^2_A \;\; \sigma^2_\Psi \;\; \sigma^2_\alpha \;\; \sigma^2_{a_f} \;\; \sigma^2_{\psi_f}]'$, the vector stacking the shocks' volatilities of the *structural* trends in X_t. The reason is because the structural trends are *unobservable* in the first place. However, 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 fix ideas, 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 the covariance matrix Σ_u is diagonal, the

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

following linear relations apply:

$$\sigma^2_{G\bar{D}P} = \sigma^2_A + \sigma^2_\Psi + \sigma^2_\alpha + \nu^2_{14}\sigma^2_{a_f} + \nu^2_{15}\sigma^2_{\psi_f} \qquad \sigma^2_{\bar{W}} = \sigma^2_A + \nu^2_{24}\sigma^2_{a_f} + \nu^2_{25}\sigma^2_{\psi_f}$$
$$\sigma^2_{\bar{E}} = \sigma^2_\Psi + \nu^2_{33}\sigma^2_\alpha + \nu^2_{34}\sigma^2_{a_f} + \nu^2_{35}\sigma^2_{\psi_f} \qquad \sigma^2_{\bar{E}_{f-m}} = \nu^2_{44}\sigma^2_{a_f} + \nu^2_{45}\sigma^2_{\psi_f}$$
$$\sigma^2_{\bar{W}_{f-m}} = \sigma^2_{a_f} + \sigma^2_{\psi_f}$$

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 ν_{ij} are simply the values around which the prior density of the long-run elasticities is centered. The only *unknowns* are the structural volatilities. It turns out that is straightforward to retrieve the structural volatilities in $\sigma^2_{\mathbf{u}}$, as they are the unknowns of a linear system of 5 equations 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 pre-sample training. Second, plug the empirical volatilities and the prior means of the parameters in $\mathcal{V}$. Solve the system for the unknown volatilities and use them to center the prior density of the structural volatilities.

Finally, 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).

A.2 Estimation of the State Space with Gibbs Sampling

Consider the unobserved states of the model in section 2 in the following stacked formulation:

$$\begin{bmatrix} \mathcal{V}X_t \\ \hat{Y}_t \end{bmatrix} = \begin{bmatrix} \mathcal{V}c \\ 0 \end{bmatrix} + \begin{bmatrix} I & 0 \\ 0 & A \end{bmatrix} \begin{bmatrix} \mathcal{V}X_{t-1} \\ \hat{Y}_{t-1} \end{bmatrix} + \begin{bmatrix} I & 0 \\ 0 & I \end{bmatrix} \begin{bmatrix} \mathcal{V}u_t \\ e_t \end{bmatrix} \tag{A.4}$$

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

$$\Sigma = \begin{bmatrix} \mathcal{V}'\Sigma_u\mathcal{V} & 0 \\ 0 & \Sigma_e \end{bmatrix} \tag{A.5}$$

Then, the model samples 50000 draws and retains the last 10000 draws from a Gibbs algorithm, according to 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 ν - 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 ν 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 ν is estimated by introducing a Metropolis-Hastings step.

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

2. Draw from the joint distribution $A, \Sigma_u, \Sigma_e \mid X_{0:T}, \hat{Y}_{-p+1:T}, Y_{1:T}$. 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 Σ_u is given by:

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

(b) CYCLE BLOCK. The posterior distributions of the lag coefficients in A and the covariance matrix Σ_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 $\hat{Z}_t = (\hat{Y}_{t-1}', \ldots, \hat{Y}_{t-p}')$,
$\mathcal{A} = \left(\sum_{t=1}^{T} \hat{Z}_t \hat{Z}_t' + \underline{\Omega}^{-1}\right)^{-1}\left(\sum_{t=1}^{T} \hat{Z}_t \hat{Y}_t' + \underline{\Omega}^{-1}\underline{A}\right)$,
$S_e = \sum_{t=1}^{T} e_t e_t' + (\mathcal{A} - \underline{A})'\underline{\Omega}^{-1}(\mathcal{A} - \underline{A})$

B DATA

The data used for the baseline model are available on the FRED website and listed in table B.1 below along with their identification code.

Table B.1: US data definitions and identification codes

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

Source: Federal Reserve Bank of St. Louis.

TRANSFORMATIONS. Data 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 are retrieved from the following product $COMPRNFB \times \frac{CE16OV}{CNP16OV}$. The female-to-male employment gap is defined as the ratio of the females (males) employment levels, scaled by their respective populations. Analogously, the female-to-male wage gap is defined as the ratio of the females over males hourly wage rates. Note that the females (males) wage rates are transformed into hourly wage rates dividing them by the usual number of working weeks in a year. Moreover, since gender-specific wage rates are only available from 1979Q1, the missing observations for the period spanning from 1960Q1 to 1978Q4 are filled with the earnings data available from the Annual Social and Economic Supplements, Current Population Survey published by the U.S. Census Bureau. These data are at annual frequency. We retrieve intra-annual observations using standard interpolation techniques.

GENDER-SPECIFIC DATA BY SECTORS AND SKILLS. We retrieve gender-specific data on hourly wages and employment by skills and sector by merging the monthly Current Population Survey (CPS) produced by the United Census Bureau and published by the Bureau of Labor Statistics. These data are publicly available at `http://data.nber.org/morg/annual/`. Concretely, the CPS survey is merged into a unified dataset that spans from 1979Q1 to 2019Q4 by editing the STATA do-file used by Dolado et al. (2021). The reader can refer to the online Appendix of Dolado et al. (2021) for a detailed exposition on how the CPS dataset is merged.

We are interested in retrieving the employment level and hourly wage rate of women and men by skills and sectors. As regards the skill dimension, workers are classified into two different types according to their education level. High-skilled workers are defined as

those individuals with at least *some college* experience. The survey respondents *without any college* experience are assigned to low-skilled worker-type. Moving to the sector dimension, the CPS includes information on more than 35 private industries in which respondents supply labor. The industries are then imputed to two aggregate sectors: (i) good-producing; (ii) services. We collect data on the service-providing industries listed in table B.2. Overall, the service sector is representative of almost 80% of total employment. The remainder share accrues to the good-producing sector and public employment.

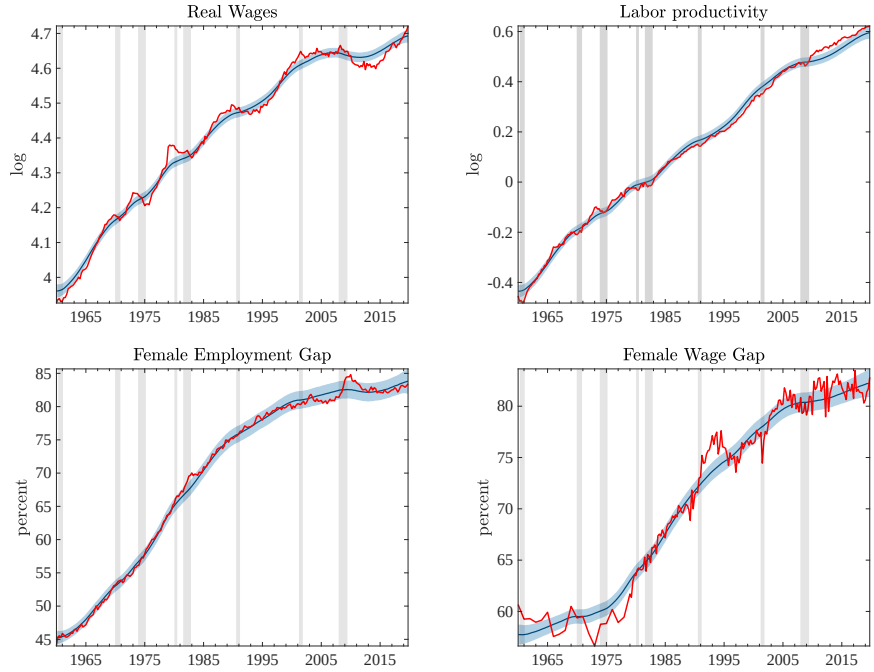
Table B.2: List of private service-providing industries

Wholesale trade	Retail trade
Transportation and warehousing	Publishing industries
Motion picture and sound recording	Broadcasting
Internet publishing and broadcasting	Telecommunications
Internet providers and data processing	Other information services
Finance, Insurance, Real Estate	Rental and leasing services
Professional and Technical services	Management of Co. and Ent.
Administrative and support services	Waste management and remediation
Educational services	Hospitals, Health care services
Social assistance	Food services and drinking places
Private households	

Source: Current Population Survey (CPS), United Census Bureau and Bureau of Labor Statistics.

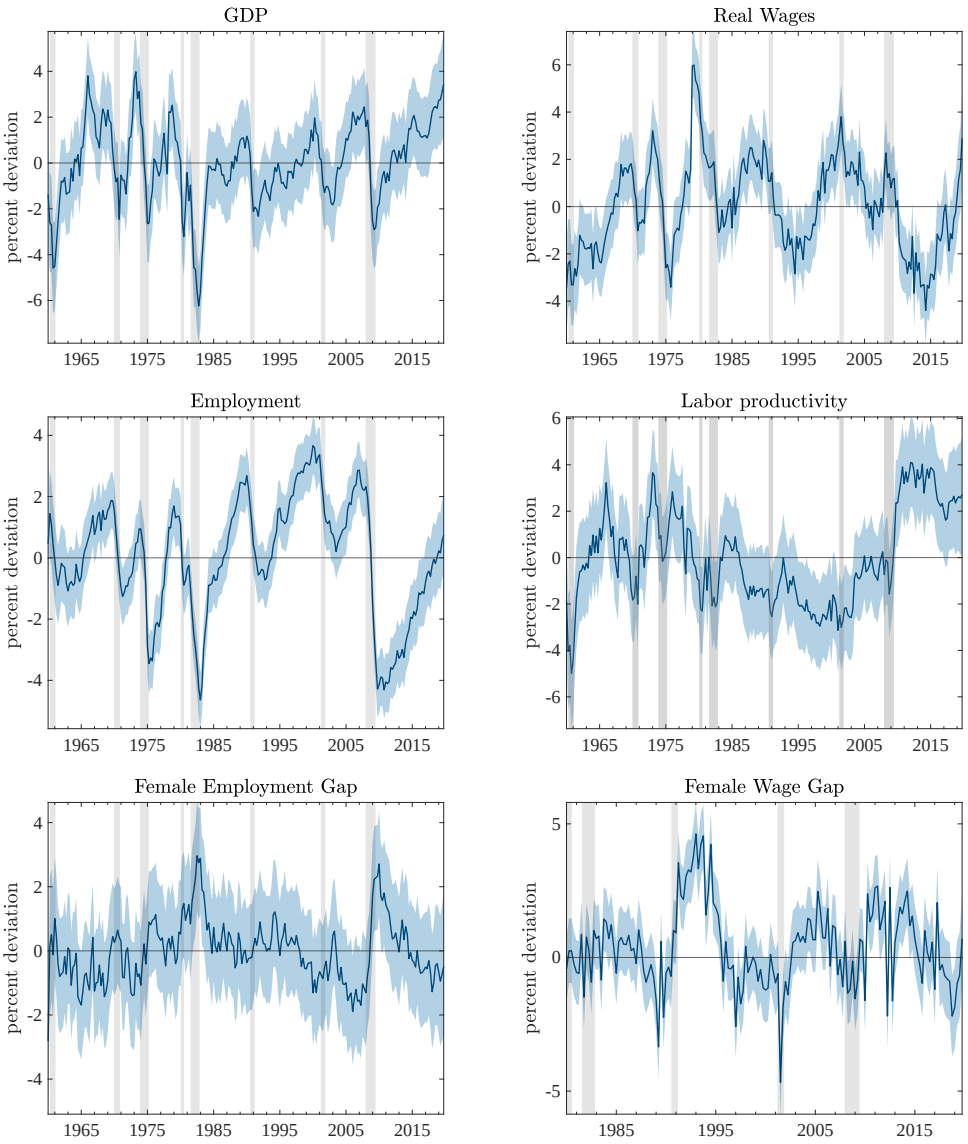
C BASELINE – ADDITIONAL FIGURES AND TABLES

Figure C.1: Empirical trends – baseline



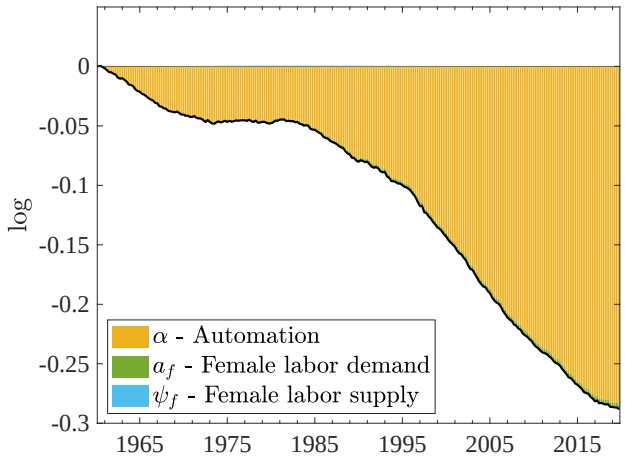
Notes: Observed data (red solid lines), median point estimates (blue solid lines). 68% coverage bands (blue shaded area).

Figure C.2: Empirical cycles – baseline



Notes: Median point estimates (blue solid lines). 68% coverage bands (blue shaded area).

Figure C.3: *Structural* drivers of the labor share – baseline model



Notes: The colored bars display the point-wise median evolution of the empirical trend of the labor share attributable to each structural trend.

Figure C.3 decomposes the secular decline in U.S. labor income share into neutral and gender-specific structural trends, based on our baseline model. The overwhelming contribution of the automation (yellow area) confirms previous findings from Acemoglu and Restrepo (2020) and Bergholt et al. (2022). At the same time, it also adds one important insight: gender-specific structural forces played *no role* in driving the secular decline of U.S. post-war labor share.

Table C.1: Drivers of trend growth rates of female employment and wage gaps – baseline

	$\Delta\bar{E}_{f,t}$			$\Delta\bar{W}_{f,t}$		
	a_f	ψ_f	Total	a_f	ψ_f	Total
1960-1969	0.83	0.72	1.55	0.55	-0.22	0.33
1970-1979	1.34	0.63	1.97	0.88	-0.19	0.69
1980-1989	1.70	-0.23	1.47	1.11	0.07	1.18
1990-1999	1.01	-0.31	0.70	0.66	0.10	0.76
2000-2009	0.51	-0.25	0.26	0.33	0.08	0.41
2010-2019	0.30	-0.13	0.17	0.19	0.04	0.23

Notes: Median point-wise contribution of the female-specific structural trends to the trend growth rates of female employment and wage gaps over time. Average of each corresponding decade.

D ROBUSTNESS EXERCISES

This section presents a series of robustness exercises that illustrate how our main empirical results are affected by various perturbations to the baseline setup. First, we tilt the priors governing feedback from gender trends to the macroeconomy towards zero. Second, we check for robustness with respect to the various zero-restrictions imposed on aggregate wages and employment. Finally, we interpret the household in our theoretical model more literally by restricting attention to gender data on married individuals, so that gender complementarity within the household in the model has a more meaningful counterpart in data.

D.1 STACKING THE CARDS AGAINST MACROECONOMIC FEEDBACK EFFECTS

Our baseline model features ample spillover from female-specific trend shocks to the macroeconomy. But how informative are data about this spillover? Here we address that question by re-estimating the baseline model with priors that are tilted towards minor spillover from gender shocks to the macroeconomy.

These new priors, as well as resulting posterior distributions, are summarized in Table D.1. Instead of the uniform priors used in the baseline, we choose an exponential prior distribution for ν_{14}, ν_{15}, ν_{24} and ν_{25}. These parameters govern the feedback from gender trends to aggregate GDP and wages, respectively. The exponential prior is tightly centered close to zero, with both the mean and the variance being equal to 0.25. For the feedback from gender shocks to aggregate employment we use a normally distributed prior with mean and variance equal to 0 and 0.1, respectively. This ensures that we are agnostic about the sign of the response of aggregate employment to gender shocks, while at the same time penalizing heavily responses that are quantitatively large in absolute value.

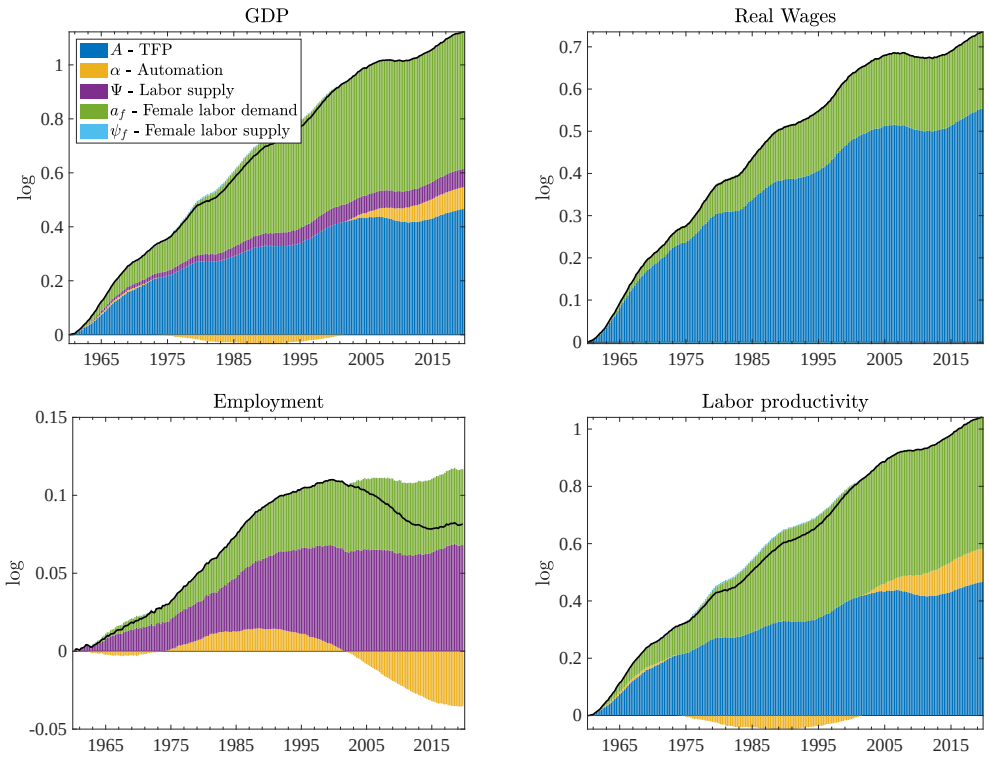
Table D.1: Prior distributions and posterior estimates

		Prior		Posterior		
		Density	Support	Mean	Mode	90% HPD
ν_{14}	$a_f \to G\bar{D}P$	$Exp(0.25)$	$[0, \infty)$	1.40	1.30	(0.85, 1.97)
ν_{24}	$a_f \to \bar{W}$	$Exp(0.25)$	$[0, \infty)$	0.74	0.77	(0.32, 1.17)
ν_{34}	$a_f \to \bar{E}$	$N(0, 0.1)$	$\mathbb{R}$	0.14	0.14	(0.00, 0.27)
ν_{15}	$\psi_f \to G\bar{D}P$	$Exp(0.25)$	$[0, \infty)$	0.16	0.06	(0.01, 0.46)
ν_{25}	$\psi_f \to \bar{W}$	$Exp(0.25)$	$[0, \infty)$	-0.11	-0.05	(-0.33, -0.00)
ν_{35}	$\psi_f \to \bar{E}$	$N(0, 0.1)$	$\mathbb{R}$	0.01	0.01	(-0.16, 0.18)
-ν_{33}	$\alpha \to \bar{E}$	$\Gamma(.3, 0.15)$	$(0, \infty)$	0.44	0.45	(0.23, 0.65)
λ	$a_f \to \bar{E}_{f-m,t}$	$\Gamma(1, 0.5)$	$(0, \infty)$	1.75	1.85	(1.18, 2.34)
γ	$\psi_f \to \bar{E}_{f-m,t}$	$\Gamma(3, 1.5)$	$(0, \infty)$	3.15	3.14	(1.87, 5.05)

Notes: The posterior moments are generated from the last 10,000 of 50,000 draws generated from the RW Metropolis-Hastings algorithm. $Exp(\mu)$ refers to the Exponential distribution with mean μ. $\Gamma(\mu, \sigma^2)$ refers to the Gamma distribution with mean μ and variance σ^2.

As documented in Table D.1, the elasticities governing feedback from female-specific productivity shift substantially away from zero even when we put low prior weight on such an outcome. The posteriors for ν_{14} and ν_{24}, which govern the feedback from female-specific productivity to GDP and aggregate wages, are centered around 1.4 and 0.75, respectively. Given that the estimates of λ and γ are relatively similar to those in the baseline, suggesting that also the in-sample estimates of realized productivity growth for female labor is rather similar across specifications, it follows that the high posterior values for ν_{14} and ν_{24} do indeed reflect a major role for female-specific productivity (instead of simply capturing lower estimates of female-specific productivity growth). The remaining parameters that govern macroeconomic feedback move less, especially those that determine the feedback from female-specific labor supply. Figure D.1 demonstrates that female-specific productivity remains important for the US macroeconomy even in a setting with very conservative priors against this outcome. Notably, the effect on labor productivity is even stronger in this specification since employment is slightly less affected by females' productivity compared with the baseline.

Figure D.1: Structural decomposition with priors stacked against macroeconomic feedback



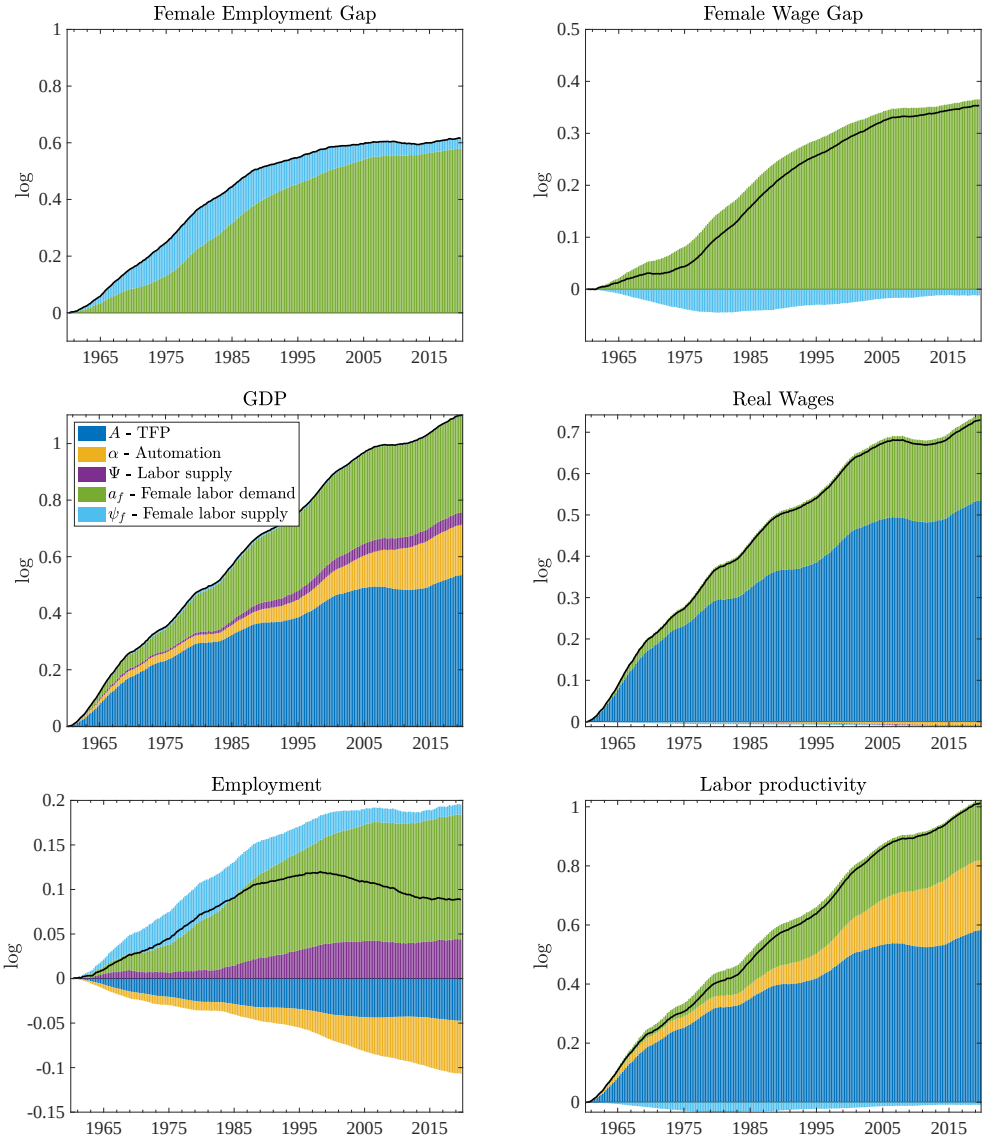
Notes: The colored bars display the point-wise median evolution of the empirical trends attributable to each structural trend.

D.2 Revisiting the long-run restrictions on employment and wages

In this exercise, we relax the balanced-growth assumption which implies that employment is invariant to changes in TFP. We set a prior on the employment that allows for negative wealth effects on labor supply but does not rule out positive effects. We use a Normal

distribution centered around -0.1, a value broadly based on Boppart and Krusell (2020). In addition, we also relax the zero long-run effects of automation and labor supply on aggregate wages. In particular, we restrict the effects to be negative with most of the prior density concentrated around zero. As shown in Figure D.2, the data favor a small but non negligible negative effect of TFP shocks on employment. This implies that gender shocks play an even larger role in driving employment up, especially in the first part of our sample. The key role of gender shocks for GDP and productivity are confirmed also in this specification. Finally, we do not find evidence of a non-negligible effect of either automation or labor supply on aggregate wages.

Figure D.2: *Structural* drivers of empirical trends – relaxing balanced growth path

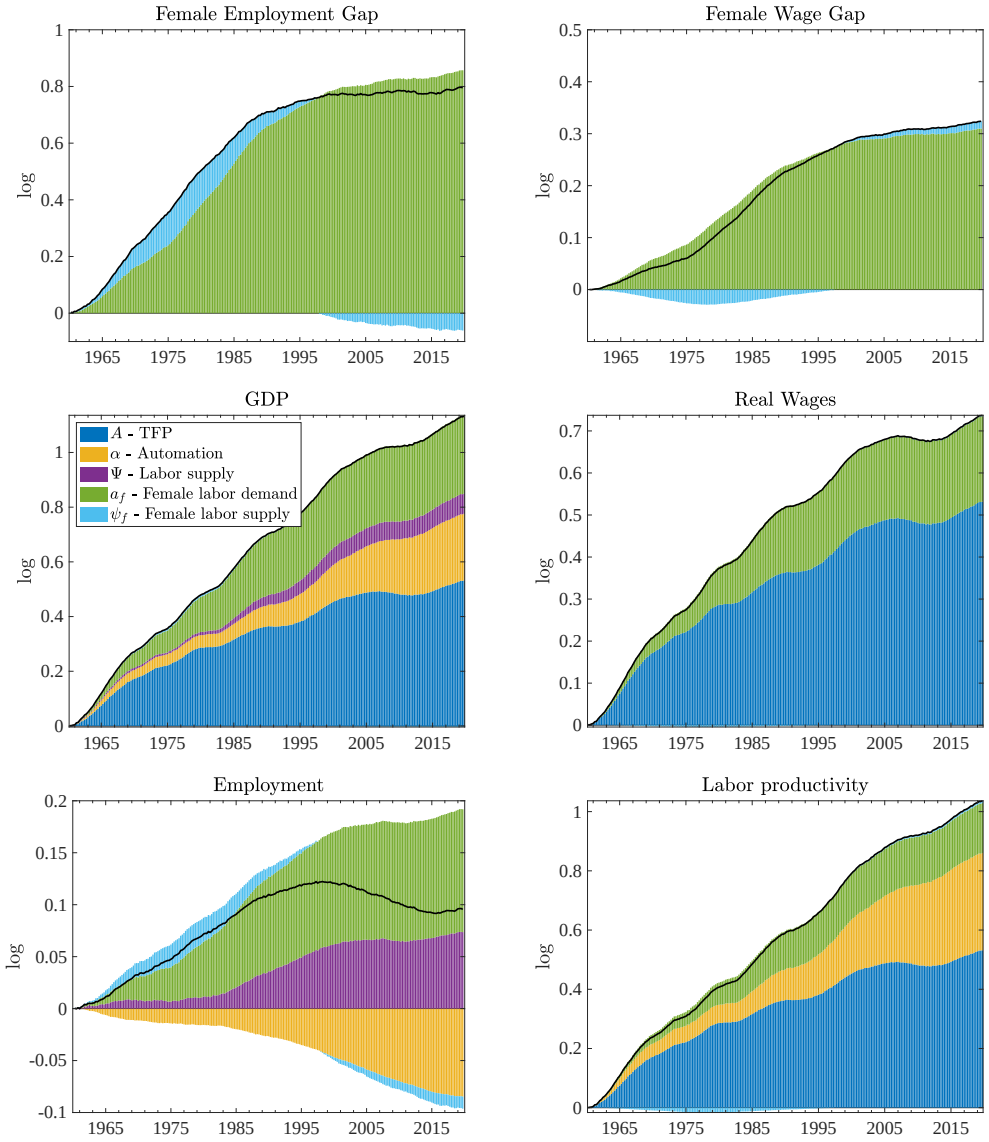


Notes: The colored bars display the point-wise median evolution of the empirical trends attributable to each structural trend.

D.3 WHAT IF WE CONSIDER ONLY MARRIED INDIVIDUALS?

In this exercise, we construct employment and wage gender gaps based only on married individuals to better reflect the counterpart in the theoretical model where decisions are taken at the household level. As shown in Figure D.3, all the main results are confirmed although the role of gender-specific labor supply shocks is further reduced in this exercise.

Figure D.3: *Structural* drivers of empirical trends – model with data on married individuals



Notes: The colored bars display the point-wise median evolution of the empirical trends attributable to each structural trend.

E A MODEL WITH MALE-SPECIFIC SHOCKS

This section outlines the prior assumptions of the model specification that jointly identifies female-specific and male-specific structural trends, as presented in section 8. We augment the baseline specification with male-specific data on employment and wage rate levels. This enables to identify a male-specific labor demand trend and a male-specific labor supply trend in addition to female-specific trends. $\mathcal{V}$ is modified accordingly:

$$\underbrace{\begin{bmatrix} G\bar{D}P_t \\ \bar{W}_t \\ \bar{E}_t \\ \bar{W}_{f-m,t} \\ \bar{E}_{f-m,t} \\ \bar{W}_{m,t} \\ \bar{E}_{m,t} \end{bmatrix}}_{\bar{Y}_t} = \underbrace{\begin{bmatrix} 1 & 1 & 1 & \nu_{14} & \nu_{15} & \nu_{16} & \nu_{17} \\ 1 & 0 & 0 & \nu_{24} & \nu_{25} & \nu_{26} & \nu_{27} \\ 0 & 1 & \nu_{33} & \nu_{34} & \nu_{35} & \nu_{36} & \nu_{37} \\ 0 & 0 & 0 & -1 & 1 & 1 & -1 \\ 0 & 0 & 0 & \gamma & \lambda & -\gamma & -\lambda \\ 1 & 0 & 0 & \nu_{64} & \nu_{65} & \nu_{66} & \nu_{67} \\ 0 & 1 & \nu_{73} & \nu_{74} & \nu_{75} & \nu_{76} & \nu_{77} \end{bmatrix}}_{\mathcal{V}} \underbrace{\begin{bmatrix} A_t \\ \Psi_t \\ \alpha_t \\ \psi_{f,t} \\ a_{f,t} \\ \psi_{m,t} \\ a_{m,t} \end{bmatrix}}_{X_t} \tag{E.1}$$

As in the baseline, the first three columns define the long-run effects of aggregate macro trends. Restrictions on GDP, wages, employment and the gender gaps are identical to the baseline. In addition, we assume that the long-run effect of technology, automation and labor supply on the level of males employment and wages is identical to their aggregate counterparts. This implies, for example, that the long-run feedback of automation to aggregate employment and males employment is identical – i.e., $\nu_{33} = \nu_{73}$. Together with the zero long-run restrictions on the gender differentials, such assumption preserves the long-run gender *neutrality* of macro trends.

The remainder columns identify the four gender-specific trends. A few remarks are in place. First, similarly to the baseline, female(male)-specific labor demand is separable from female(male)-specific labour supply because the former implies the same co-movement between gender gaps, while the latter implies a negative sign on the co-movement between gender gaps. Consistent with the results from the theoretical model in Figure E.1, female-specific and male-specific shocks are assumed to have non-negative long-run effect on GDP and are mutually exclusive via the opposite sign effect on the employment gap. The uniform priors for the remainder gender-specific feedback to macro are formulated using the impulse responses in Figure E.1 as a reference point. Furthermore, both the female-specific and the male-specific shocks are normalized to have unit long-run effects on the wage gap, so that the feedbacks to the employment gap can be interpreted in terms of γ and λ, as in the baseline. Furthermore, we assume that female-specific and male-specific shocks have symmetric effects on both the gender gaps and macro aggregates. This implies that the elasticities of macro aggregate with respect to male-specific shocks span the same uniform boundaries of the macro feedbacks to female-specific shocks. In this way, we remain agnostic about the *relative* strength of female- and male-specific shocks. Finally, we also estimate the effects of female-specific shocks to males employment and wages – i.e., ν_{64}, ν_{65}, ν_{74}, ν_{75}. As discussed in section 8, this is particularly useful because we can make inference on the crowding out effect, conditional on either a female-specific shocks. These effects are captured by ν_{74} and ν_{75}. The prior on both these elasticities is rather loose: it uniformly spans the probability set [-1,0]. This allows the likelihood to visit both regions with large and small crowding out effects. The

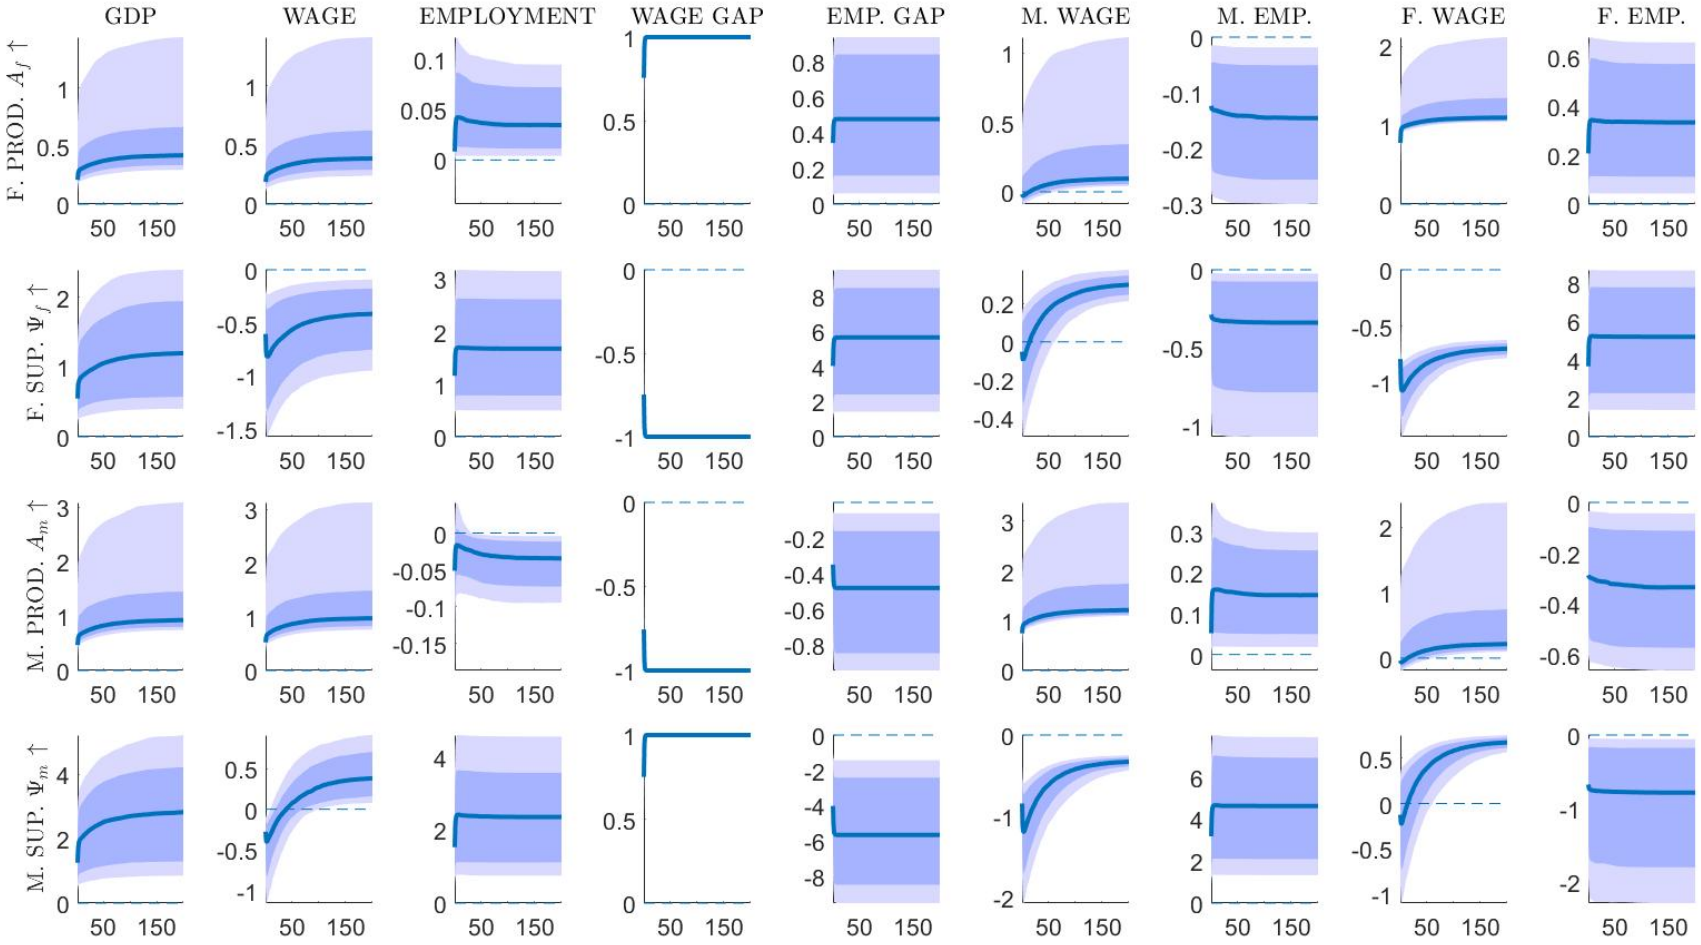
prior and posterior estimates are summarized in table E.1.

Table E.1: Prior distributions and posterior estimates

		Prior		Posterior		
		Density	Support	Mean	Mode	90% HPD
ν_{14}	$\psi_f \to G\bar{D}P$	Uniform	$[0, 2]$	1.61	1.39	(1.23, 1.92)
ν_{24}	$\psi_f \to \bar{W}$	Uniform	$[-2, 0]$	-0.24	-0.30	(-0.52, -0.03)
ν_{34}	$\psi_f \to \bar{E}$	Uniform	$[0, 3]$	2.27	1.94	(1.89, 2.69)
ν_{64}	$\psi_f \to \bar{W}_m$	Uniform	$[0, 1]$	1.60	1.69	(1.07, 2.03)
ν_{74}	$\psi_f \to \bar{E}_m$	Uniform	$[-1, 0]$	2.73	2.07	(1.98, -3.53)
ν_{15}	$a_f \to G\bar{D}P$	Uniform	$[0, 1]$	0.94	0.99	(0.80, 0.99)
ν_{25}	$a_f \to \bar{W}$	Uniform	$[0, 1]$	0.91	0.99	(0.77, 0.99)
ν_{35}	$a_f \to \bar{E}$	Uniform	$[-0.5, 0.5]$	0.26	0.31	(-0.04, 0.46)
ν_{65}	$a_f \to \bar{W}_m$	Uniform	$[0, 1]$	-0.01	0.00	(-0.21, 0.13)
ν_{75}	$a_f \to \bar{E}_m$	Uniform	$[-1, 0]$	-0.33	0.41	(-0.63, -0.10)
ν_{16}	$\psi_m \to G\bar{D}P$	Uniform	$[0, 2]$	0.62	0.57	(0.13, 1.13)
ν_{26}	$\psi_m \to \bar{W}$	Uniform	$[-2, 0]$	-1.51	-1.52	(-1.97, -0.72)
ν_{36}	$\psi_m \to \bar{E}$	Uniform	$[0, 3]$	0.81	0.91	(0.30, 1.25)
ν_{66}	$\psi_m \to \bar{W}_m$	Uniform	$[-1, 0]$	-0.40	-0.20	(-0.89, -0.06)
ν_{76}	$\psi_m \to \bar{E}_m$	Uniform	$[0, 10]$	6.10	5.80	(5.30, 7.29)
ν_{17}	$a_m \to G\bar{D}P$	Uniform	$[0, 1]$	0.63	0.98	(0.09, 0.95)
ν_{27}	$a_m \to \bar{W}$	Uniform	$[0, 1]$	0.72	0.82	(0.26, 0.94)
ν_{37}	$a_m \to \bar{E}$	Uniform	$[-0.5, 0.5]$	0.15	0.20	(-0.29, 0.44)
ν_{67}	$a_m \to \bar{W}_m$	Uniform	$[0, 2]$	1.68	1.95	(1.18, 1.96)
ν_{67}	$a_m \to \bar{E}_m$	Uniform	$[0, 1]$	0.54	0.68	(0.07, 0.92)
$-\nu_{33}$	$\alpha \to \bar{E}$	$\Gamma(0.3, .15)$	$(0, \infty)$	0.39	0.41	(0.20, 0.54)
λ	$a_{i=\{f,m\}} \to \bar{E}_{f-m,t}$	$\Gamma(1, .5)$	$(0, \infty)$	1.31	1.57	(0.81, 1.61)
γ	$\psi_{i=\{f,m\}} \to \bar{E}_{f-m,t}$	$\Gamma(3, 1.5)$	$(0, \infty)$	6.28	6.7	(5.48, 6.81)

Notes: The posterior moments are generated from the last $10,000$ of $50,000$ draws generated from the RW Metropolis-Hastings algorithm. $\Gamma(\mu, \sigma^2)$ refers to the Gamma prior density with mean μ and variance σ^2.

Figure E.1: Additional impulse responses: male vs female shocks



Notes: Impulse response functions from simulations of the theoretical model. Pointwise median, 90% and 68% bands based on 1,000 independent draws from the parameter distributions. The y-axes measure responses in percent, the x-axes represent time in quarters.