Female Employment, Structural Transformation, Endogenous Sorting

Gonul Sengul*
Ozyegin University

December 6, 2025

Abstract

The significant increase in female employment and narrowing of the gender wage gap across developed economies is primarily attributed to structural transformation and the expansion of the service sector. However, this paper demonstrates that existing structural transformation models are quantitatively insufficient to account for the observed convergence without assuming exogenous changes in female comparative advantage. I extend this framework by incorporating idiosyncratic worker productivity and Roy-type self-selection, endogenizing the within-sector compositional changes. A benchmark simulation accounts for only half the observed employment rise and predicts a widening gender wage gap. Successful replication of the data dynamics requires a shift in female talent over time. Specifically, the model necessitates that female talent must have simultaneously increased in its mean and decreased in its dispersion (higher and more concentrated talent). This finding is robust to the era of calibration and is valid for a set of developed countries. The paper establishes that achieving gender convergence that is comparable with the data requires both structural transformation and better and more concentrated female talent.

JEL Code: J16, J21, J22, J24, J31.
Keywords: Female employment, structural transformation, service sector, labor productivity, worker heterogeneity, Roy model.

*Economics Department, Ozyegin University. Email: gonul@utexas.edu.

Introduction

The rise in female employment and the narrowing of the gender wage gap is one of the most significant labor market shifts across developed countries in recent decades. The leading explanation attributes this trend primarily to structural transformation, the reallocation of labor toward the service sector driven by uneven sectoral productivity growth (Olivetti and Petrongolo (2016)). Because women have comparative advantage in services, the expanding service sector disproportionately benefits female workers. Existing structural transformation models, such as by Ngai and Petrongolo (2017), provide a powerful quantitative framework for female employment based on this link. However, their full success in matching the magnitude of the rise in female employment and the convergence of gender gaps is contingent on exogenously specified changes in female comparative advantage. What drives the endogenous evolution of female labor intensity within sectors remains an open question. Capturing this dynamic is crucial, as it holds the key to fully understanding the rise of female employment and is essential for guiding effective policy aimed at closing the persistent gender gaps globally.

This paper extends the structural transformation framework to study the endogenous evolution of female labor intensity by introducing idiosyncratic worker productivity and Roy-type self-selection. This modeling approach is crucial, as it generates the required within-sector compositional changes endogenously through household labor supply responses. My quantitative analysis yields two key results regarding the sources of the rise in female employment and the convergence of the gender wage gap. First, I confirm that structural transformation alone is quantitatively insufficient. The model accounts for only half of the observed rise in female employment and predicts a widening gender wage gap. Second, I demonstrate that matching the observed data dynamics requires a change in female productivity relative to males. Specifically, the data implies that female talent must have simultaneously increased in its mean and decreased in its dispersion (i.e., concentrated talent) over the sample period. This concentrated talent is essential because it simultaneously raises average female productivity (increasing their intensity) and drives wage convergence, effects that are absent in the structural transformation mechanism alone.

The model features a representative household composed of a continuum of male and female members, making consumption and labor supply decisions to allocate time to goods, services, home production, or leisure. The structural transformation mechanism is captured by non-homothetic preferences and uneven sectoral productivity growth. The extension introduces idiosyncratic, sector-specific (for market services) productivity for each member, drawn from a gender-specific distribution. The household member's time allocation decision, and therefore the market sector into which they sort, is driven by their individual wages and sector-specific productivity draws (a Roy-type self-selection). Exogenous and uneven sector productivity growth leads to labor reallocation, which in turn endogenously changes the average productivity of workers remaining in or entering a market sector. This link is the mechanism through which the overall structural transformation alters the gender composition and labor intensity

within sectors.

I begin the quantitative analysis by calibrating the model to the baseline labor allocation and gender wage gap moments for the U.S. economy in 1983, utilizing data from the ILO and the American Time Use Survey (ATUS). I then establish a benchmark by simulating the model forward to 2018, imposing only the observed sectoral productivity growth rates while holding the gender productivity distributions constant over time. This benchmark scenario reveals three major divergences from the observed data trends. First, the model accounts for only half of the observed increase in female employment. Second, instead of narrowing, the model predicts a widening gender wage gap. Finally, the model incorrectly attributes almost all the change in female employment to between-sector forces, a composition starkly inconsistent with the data that points to strong within-sector growth.

The limitations of the benchmark model stem from endogenous labor sorting and the implied changes in average worker productivities. When the structural transformation expands the services sector, it draws in less productive workers, lowering the sector's average productivity. Since women start with lower employment shares, productivity decline due to selection is more pronounced for them, causing a strong dilution of average female productivity in the service sector. Lower average worker productivity weakens the growth of female employment and widens the gender wage gap. To fully match the data, I then re-calibrate the model allowing the female productivity distribution to change over time to simultaneously match the increase in female employment and the narrowing wage gaps. This exercise reveals that female talent must have simultaneously increased in its mean and decreased in its dispersion (concentrated talent) over the sample period. This shift is essential for narrowing gender wage gaps and driving within-sector growth of female intensity, thereby providing the full accounting of the observed labor market changes.

The main finding of this paper, that structural transformation alone is insufficient to account for the observed gender gap convergence and a shift toward higher and more concentrated female talent is also essential, is robust to the time-period or the country chosen. First, to check whether the result is driven by the sample period, I redo the calibration for the U.S. economy using the 1970?2007 period and the data moments reported by Ngai and Petrongolo (2017). Second, to test whether this finding is specific to the US economy or it has a cross-country generality, I calibrate the model to U.K, France, Italy, and Spain. Crucially, both exercises yield the same qualitative requirement: the model consistently demands a simultaneous increase in the mean and decrease in the dispersion of the female productivity distribution (compared to the male productivity distribution) to match the observed trends in female employment and wage convergence.

This paper primarily contributes to the structural transformation literature that explains the rise of female employment.[1] Olivetti and Petrongolo (2016) study the long

[1]There is also broader structural transformation literature that studies cross country differences in labor allocations. Rogerson (2008), Adamopoulos and Akyol (2009), and İmrohoroğlu et al. (2014) are examples.

run trends in female outcomes for 18 developed countries, and find that female labor force participation increased substantially, albeit unevenly across countries and over time. They use shift-share decomposition and show that there is contribution from both between industry and within-industry changes. Akbulut (2011) and Ngai and Petrongolo (2017) use a structural transformation model to show that the rise of service sector can explain s significant portion of the increase in female employment in the US. These papers assume that females have comparative advantage over males in service sector jobs. An expansion in service sector creates demand for female employment and hence draws females from out of labor force to the labor market. Rendall (2017a) provides a theory for the comparative advantage of women in services over goods sectors. In Akbulut (2011), only men work in goods sector and only women work in home production, and both inputs are perfect substitutes. In Ngai and Petrongolo (2017), each gender contributes to goods, services and home production, as well as leisure, and they are imperfect substitutes. Rendall (2017b) adds tax to a structural transformation model and studies the effect of tax regimes on the rise of female employment while comparing the experiences of the US and Europe.

The paper is also connects with the literature that explores various factors driving female labor supply, including technological change, fertility, childcare policy, and human capital accumulation. This literature provides compelling explanations for changes in women's market opportunity and time use. For instance, Greenwood et al. (2005) highlight how the availability of household appliances increased market work by reducing home production time. Other work, such as Fogli and Veldkamp (2011), investigates how social learning accelerates participation trends. My work complements these broader supply-side perspectives by focusing on the endogenous selection and changing quality of human capital that women bring to the market, which I show to be a necessary condition for gender convergence within the structural transformation framework.

Finally, I contribute to the literature that uses Roy-type selection and heterogeneous agents to study gender outcomes. With a similar modeling approach, Hsieh et al. (2019) uses Roy model to analyze the impact of misallocation on aggregate economic growth due to discrimination. Cerina et al. (2021) explore the role of female labor in job polarization, showing how skill-biased change and home production links shifted female demand. My work embeds the selection mechanism within the canonical structural transformation model, making the changing productivity distribution the central source of endogenous within-sector change. I also relate to work like Erosa et al. (2017) that examines the interaction of time allocation and occupational choice to explain gender wage gaps.

The remainder of the paper is structured as follows: Section 1 introduces the general equilibrium model and formalizes the Roy-type selection mechanism. Section2 describes the data, estimation, and calibration strategy. Section 3 presents the main quantitative results, establishing the limitations of the structural transformation benchmark and demonstrating the necessity of the shift toward higher and more concentrated female talent. Section 4 discusses the robustness and generality of the main findings, and

Section 5 concludes.

1 The Model

The model has an infinitely lived representative household that makes consumption and labor supply decisions, two market sectors producing goods and services, and a home sector producing home services. Producing sectors have exogenous sectoral productivity growth, the model is otherwise static and hence time subscripts are omited.

The household is composed of a continuum of female ($g = f$) and male ($g = m$) members, each endowed with one unit of time. The representative household maximizes utility, defined over composite consumption good (C) and total leisure time(L):

$$u(C, L) = [(1 - \alpha_l)C^{\epsilon_l} + \alpha_l L^{\epsilon_l}]^{\frac{1}{\epsilon_l}},$$

where α_l is the share parameter, and ϵ_l governs the elasticity of substitution between consumption and leisure. ϵ_l is negative so that consumption and leisure are gross complements. An increase in sectoral productivity increases both consumption and leisure. The composite consumption good, C, is defined as:

$$C = [\alpha_g G^{\epsilon} + (1 - \alpha_g)S_c^{\epsilon}]^{\frac{1}{\epsilon}}.$$

G is the market-produced goods and S_c is combined market and home services. α_g is the share parameter of goods, and ϵ governs the degree of substitution between goods and composite services. Goods and services are complements as well ($\epsilon < 0$). S_c is comprised of services produced in the market (S) and at home (H), aggregated through the following form:

$$S_c = [\alpha_s S^{\eta} + (1 - \alpha_s)H^{\eta}]^{\frac{1}{\eta}},$$

where α_s is the share parameter of market services and η governs the degree of substitution between home and market services. Home services and market services are assumed to be substitutes ($\eta > 0$).

Uneven exogenous sectoral productivity growths generate structural transformation and marketization through different income elasticities of items in the consumption bundle. If there is a higher productivity growth in goods sector, higher consumption of G will increase the demand for composite service good S_c as they are complements. An increase in S will raise the demand for labor in market service sector. This reallocation of labor from goods to services is structural transformation. If productivity in market services grows faster than the productivity in home services, as we assume that home and market service products are substitutes, labor will relocate from home services to market services, which is known as marketization.

Lastly, leisure is aggregated as

$$L = \left[\phi_l L_f^{\mu_l} + (1 - \phi_l)L_m^{\mu_l}\right]^{\frac{1}{\mu_l}}, \tag{1}$$

where ϕ_l is share parameter while μ_l governs the elasticity of substitution. Female and male leisure are complements. Hence, as the households demands more leisure, it does so from both genders.

The household's budget constrain is:

$$(1+\tau_c)(P_gG + P_sS) \leq (1-\tau_w)(\bar{W}_f + \bar{W}_m) + T,$$

where P_j is the price of good $j \in \{g, s\}$, and $\bar{W}_x$ is the total income the household of gender x earns. τ_c and τ_w are consumption and labor income taxes, respectively, that the government collects and distributes as a lump-sum transfer (T).

Each household member has a unit measure of time that they use to produce either in the market or at home, or produce leisure. Members of each gender have the same productivity in home services and leisure production. A worker of gender x can produce z_j amount of output in sector $j \in \{g, s\}$. Worker productivity in market sectors (z_g, z_s) comes from the distribution $\Gamma_x(z_g, z_s)$ for each gender x. I assume that worker productivity is independent across sectors and comes from a Frechet distribution (a la Hsieh et al. (2019)).

Idiosyncratic productivity in each sector has the following CDF:

$$Pr(Z \leq z) = e^{-(\frac{z-m}{s})^{-\theta}},$$

where s is the scale parameter, θ is shape parameter, and the support is such that $x > m$. I assumes that $m = 0$ for both genders. The mean of the distribution increases with s_x and its variances decreases with θ_x. The joint cumulative distribution for gender x becomes:

$$G_x(z_g, z_s) = Pr_x(Z_g \leq z_g, Z_s \leq z_s) = e^{-\sum_j (z_j/s_x)^{-\theta_x}}.$$

A household member chooses to work in market sector j_x^* if:

$$j_x^* = argmax\{w_{xg}z_g,\ w_{xs}z_s\},$$

where w_{xj} is the efficiency wage of gender x in sector j and $w_{xj}z_j$ is the earnings of a worker of gender x with productivity z_j in sector j.

The worker decides to work in market sector j_x^*, rather than working at home or producing leisure if

$$j_x^* = argmax\{w_{xj_x^*}z_{j_x^*},\ \tilde{w}_{xh},\ \tilde{w}_{xl}\},$$

where $\tilde{w}_{xh}$ and $\tilde{w}_{xl}$ are the shadow prices of producing home services and leisure, respectively.

The household allocates non-participant members to home services and leisure production. Leisure is produced via equation 1, and home services production is:

$$H = A_hN_h = A_h\left[\phi_h E_{fh}^{\mu} + (1-\phi_h)E_{mh}^{\mu}\right]^{\frac{1}{\mu}},$$

where A_h is sectoral productivity that grows at an exogenous rate r_h, μ governs the elasticity of substitution between male and female labor inputs, and ϕ_h is the female weight in production.

Producers in market sectors hire labor and produce through the following CES production form:

$$Y_j = A_j N_j = A_j \left[\phi_j Y_{fj}^{\mu} + (1-\phi_j) Y_{mj}^{\mu}\right]^{\frac{1}{\mu}}, \quad j \in \{g, s\},$$

These firms operate in perfectly competitive markets to maximize their profits:

$$\max P_j Y_j - w_{fj} Y_{fj} - w_{mj} Y_{mj},$$

where A_j is sectoral productivity that grows at an exogenous rate r_j, μ is the parameter governing the elasticity of substitution between male and female labor inputs, and ϕ_j is the female weight in sector j. Production by gender x is defined as

$$Y_{xj} = \int_{\Omega_{xj}} z_j d\Gamma_x(z_g, z_s),$$

where Ω_{xj} is the set of members of gender x that choose working in sector j over working in other sectors or producing leisure.

1.1 Equilibrium

Given exogenous labor productivity growth rates r_g, r_s, r_h, idiosyncratic productivity distribution Γ_x for gender $x \in \{f, m\}$, prices P_s, P_g and efficiency unit wages w_{fg}, w_{fs}, w_{mg} and w_{ms},

- The household chooses the amount of goods and market services to purchase, and allocates members across market work, home services production and leisure to maximize its utility, subject to the budget constraint.
- Firms hire workers to maximize their profits.
- Prices clear the markets.

Let us begin with firms' profit maximization conditions, which yield the efficiency wages:

$$P_j \frac{\partial Y_j}{\partial Y_{fj}} = w_{fj}, \qquad P_j \frac{\partial Y_j}{\partial Y_{mj}} = w_{mj} \quad j \in \{g, s\}.$$

This implies that the marginal rate of technical substitution in efficiency units across genders should be equal to the efficiency wage ratio:

$$\frac{\partial Y_j / \partial Y_{fj}}{\partial Y_j / \partial Y_{mj}} = \frac{\phi_j Y_{fj}^{\mu-1}}{(1-\phi_j) Y_{mj}^{\mu-1}} = \frac{w_{fj}}{w_{mj}},$$

Let $\bar{y}_{xj}$ be the average productivity of workers of gender x that chose to work in sector j. Then, $Y_{xj} = E_{xj} \bar{y}_{xj}$, and the equation above becomes

$$\frac{E_{fj}}{E_{mj}} = \left(\frac{\phi_j}{1-\phi_j}\right)^{\frac{-1}{\mu-1}} \left(\frac{w_{fj}}{w_{mj}}\right)^{\frac{1}{\mu-1}} \frac{\bar{y}_{mj}}{\bar{y}_{fj}}. \tag{2}$$

The optimal ratio of female to male employment depends on the weight parameters in the production function ϕ_j, the gender efficiency wage gap in the sector, $\frac{w_{fj}}{w_{mj}}$, and the average gender productivity gap of workers in that sector $\frac{\bar{y}_{mj}}{\bar{y}_{fj}}$. Note that the female share in a sector increases with ϕ_j and the ratio of average productivity values of females to that of males. Moreover, the average worker productivity in a sector declines as more workers are hired. Hence, with structural transformation average worker productivity will increase in goods sector (as it sheds labor) and decreases in market services. Relative decline or rise in the ratio of worker productivity of females and males in a sector will determine the gender composition.

The household, given prices and wages, chooses amount of G and S to purchase from the market to maximize its utility. To do so, the household will want to equate the marginal rate of substitution across goods and market services to their relative prices. The household will also allocate male and female labor to the production of goods, market services, home home services as well as to leisure, to maximize its utility.

A worker of gender x chooses one market sector over another if that sector generates higher income than the other one. Let p_{xj} denote the probability that an individual of gender x chooses sector $j \in \{g, s\}$. Suppose goods sector is chosen. Then we have:

$$p_{xg} = Pr[\frac{\omega_{xg}}{\omega_{xs}} z_g > z_s],$$

$$p_{xg} = \frac{1}{\bar{\omega}_{xg}} \Big(1 - e^{-(\frac{\tilde{\omega}_x}{\omega_{xg}})^{-\theta_x} \bar{\omega}_{xg}}\Big),$$

where $\bar{\omega}_{xg} = [1 + (\frac{\omega_{xg}}{\omega_{xs}})^{-\theta_x}]$. Note that a change in the shape parameter affects the probability of a worker choosing one sector over the other directly, but not the change in the scale parameter.

Expected productivity of a worker in sector g, conditional on worker choosing that sector, is y_{xg}:

$$y_{xg} = E[z_g | \ \frac{\omega_{xg}}{\omega_{xs}} z_g > z_s]$$

$$y_{xg} = s_x \bar{\omega_{xj}}^{\frac{1}{\theta}-1} \frac{\int^{\frac{\tilde{\omega}_x}{\omega_{xj}}} u^{\frac{-1}{\theta_x}} e^{-u} du}{Pr(z^{*j} > \frac{\tilde{\omega}_x}{\omega_{xj}})},$$

where $\int^{\frac{\tilde{\omega}_x}{\omega_{xg}}} u^{-1/\theta_x} e^{-u} du$ is incomplete gamma function. Note that scale parameter s_x affects the worker productivity directly.

The household will allocate non-participant members to home services production and leisure so that marginal returns are the same for each gender.

$$\frac{\partial U}{\partial H} \frac{\partial H}{\partial E_{xh}} = \frac{\partial U}{\partial L} \frac{\partial L}{\partial E_{xl}}, \quad x \in \{f, m\}. \tag{3}$$

A household member with idiosyncratic productivity (z_g, z_s) participating in the market gains $w_{j_x^*} z_{j_x^*}$ amount of income, where j_x^* is the chosen market sector. As this income will relax the budget constraint, its utility benefit is $\lambda w_{j_x^*} z_{j_x^*}$, where λ is the shadow price of the budget constraint. Hence, the worker chooses market work if

$$\frac{\partial U}{\partial H} \frac{\partial H}{\partial E_{xh}} \leq \lambda w_{j^*} z_{j^*}. \tag{4}$$

Note that in equilibrium we have the same shadow wage for each gender for home work and leisure. Household member works at sector $j*_x$ if the income from that sector brings a higher benefit than working at home or producing leisure.

In equilibrium, all economic agents' decisions are consistent, and markets clear. The household maximizes its utility by allocating the time of its male and female members across production of goods, market services, home services, and leisure. These workers, who are heterogeneous in their productivity in market sectors, self-select into the sector with the highest return. Firms, in turn, hire labor to maximize profits, which results in efficiency wages for each gender in each sector. Market wages influence which workers choose to enter each sector, and this endogenous sorting determines the average productivity of the workforce, which in turn, affects the gender composition within each sector and influences average wages. Worker heterogeneity creates a dynamic outcome of structural transformation and gender-specific labor reallocation. This endogenous process allows the model to capture changes in gender composition within sectors, which is a central focus of this paper.

2 Quantitative Analysis

I now turn to the quantitative exercise. I calibrate the model to the US economy between 1983 and 2018. I determine the elasticity parameters based on established literature values to match the number of calibrated parameters to the data moments. I start with elasticity of substitution between goods and composite services in consumption, which is $1 - 1/\epsilon_l$. Boppart and Ngai (2021) calibrates the elasticity to be 0.1 in a balanced growth model that studies the evolution of market and leisure time trends in the USA. I set the ϵ_l parameter to -9 so that the elasticity is 0.1. Like consumption and leisure, goods and composite services are also complements. ϵ is the parameter that governs the elasticity of substitution between these products and there is a wide range of elasticity values used in the literature. For instance, Rogerson (2008) uses an elasticity of 0.44 while Ngai and Petrongolo (2017) use an elasticity of 0.002. I set this elasticity to be 0.05. Changing this parameter affects calibrated value of α_g while leaving the main results qualitatively the same. η governs the elasticity of substitution between home and market services. I set the elasticity to 2, as in Ngai and Petrongolo (2017). μ is the parameter that governs the elasticity of substitution between male and female labor in production. Ngai and Petrongolo (2017) sets this value to 0.559, and Acemoglu et al. (2004) uses 0.68. I use the value Ngai and Petrongolo (2017) uses. I

also get μ_l value from Ngai and Petrongolo (2017), which is -4.26. Table 1 presents the parameter values.

Table 1: Model Parameters

Parameter	Value	Source	Parameter	Value	Source
ϵ	-19	-	ϕ_g	0.233	Calibrated
η	0.5	Ngai and Petrongolo (2017)	ϕ_s	0.366	Calibrated
μ	0.5595	Ngai and Petrongolo (2017)	ϕ_h	0.391	Calibrated
μ_l	-4.26	Ngai and Petrongolo (2017)	ϕ_l	0.388	Calibrated
ϵ_1	-9	Boppart and Ngai (2021)	A_g^b	1	Normalization
α_l	0.041	Calibrated	A_s^b	1	Normalization
α_g	0.076	Calibrated	A_h^b	2	Normalization
α_s	0.531	Calibrated	A_l^b	2	Normalization
θ_m	3	Normalization	A_g^e	2.05	Calibrated
s_m	2	Normalization	A_s^e	1.61	Calibrated
θ_f	3.69	Calibrated	A_h^e	2.32	Calibrated
s_f	2	Normalization	A_l^e	2.80	Calibrated

Notes: Subscripts $\{l, h, g, s\}$ stand for leisure, home, goods, and services, respectively. Superscripts $\{b, e\}$ stand for beginning year and end year, respectively. All parameters in the top left are preference parameters while all the parameters on the right side of the table are production parameters. θ and s are shape and scale parameters of Frechet distributions for females (f) and males(m).

To find data counterparts of time allocated to market, home work and leisure, I start with the ILO data on hours worked by gender and sector. The data is available for US starting in 1983. Using population, employment and average weekly hours worked data, I find aggregate hours allocated to goods and services. I compute the remaining time allocation between household work and leisure using American Time Use survey results. Details of these computations are discussed in the appendix. I calibrate $\alpha_l, \alpha_g, \alpha_s$ (the share parameters in the utility) to match the male aggregate hours in leisure, goods and service sectors in the beginning year. The remaining time is allocated, by definition, to home production. I calibrate female weights (comparative advantage parameters) in production of goods, market services, and home production (ϕ_g, ϕ_s, ϕ_h) to match beginning year female aggregate hours in these sectors. Hence, remaining female time is allocated to leisure. I then compute the leisure aggregation share parameter ϕ_l so that these leisure allocations are consistent with the equilibrium outcomes. Using these moments for 1983, I find that females have the highest comparative advantage in home production. Then, they are more productive in leisure, which is followed by services and then goods sectors (Table 1). Hence, in the model women have comparative advantage in services over goods sector, and in home production over these sectors.

Sectoral productivity values (A_j) for market work are normalized to 1 for the beginning year. I normalize beginning year sectoral productivity values for home and leisure

to 2.[2] Sectoral productivity values (A_j) grow at an exogenous rate. For the growth in sectoral productivity in goods and services, I use labor productivity growth rates in the data as calibration targets. In the model, sectoralal productivity (A_j) and labor productivity ($\bar{a_j}$) are not the same, but they are tightly linked:

$$\bar{a_j} = \frac{Y_j}{E_{fj} + E_{mj}} = \frac{A_j \left[\phi_j Y_{fj}^{\mu} + (1-\phi_j) Y_{mj}^{\mu}\right]^{\frac{1}{\mu}}}{E_{fj} + E_{mj}}.$$

I calibrate end year A_j values for goods and market services so that the model implied labor productivity growth rates in these sectors match those observed in the data. For the home sector, I assume labor productivity growth rate is 0.45 percent, as Ngai and Petrongolo (2017) reports. I also assume leisure productivity grows over time and I target male time allocation to market sectors in the end year to pin down the growth in leisure productivity. Productivity growth rates are reported and discussed in Section 2.1.

I assume that the worker productivity distribution is a two parameter Frechet distribution that differs by gender.[3] I normalize the male distribution shape parameter (θ_m) to 3 and scale parameter (s_m) to 2. I normalize female scale parameter (s_f) to 2, and calibrate female shape parameter (θ_f) for the beginning year to match gender wage gap (from ILO data) in 1983. I assume distribution parameters stay constant over time in this benchmark calibration. The effects of distribution parameter changes are discussed in Section 3. Lastly, all tax data is taken from McDaniel (2014), and the values are reported in the Table **??**.

2.1 Results

As the there is productivity heterogeneity across workers, average labor productivity ($\bar{a}_j$) and sectoral productivity (A_j) are different in the model. I calibrate the end year sectoral productivity (A_j) values such that the model implied labor productivity ($\bar{a}_j$) matches that of the data. Table 2 presents the resulting implied sectoral productivity (A_j) growth rates.

Consistent with structural transformation, we observe that the goods sector grows at a higher rate than the market services, which, in turn, grows at a higher rate than home production sector. Sectoral productivity growth is lower than the labor productivity growth in goods sector, because declining employment means the average remaining employee is more productive, demanding a lower sector productivity increase (A_j) to match the labor productivity growth observed in the data. The sectoral productivity growth rate in market services is lower than its corresponding labor productivity growth rate for the same reason. I calibrate the leisure productivity growth rate to match the male employment change in market (home and leisure allocation) hours as no change in

[2] As there is additional value added from worker productivity, we need a higher initial A values for home and leisure for computational reasons.

[3] Hsieh et al. (2019) use a Frechet distribution as well.

leisure productivity would require market hours to decline and leisure hours to increase. Moreover, the increase in leisure productivity is higher than that of home production to match the observed non-market work time change for females.

Table 2: Sectoral Productivity Growth Rates (%)

	Industry	Services	Home	Leisure
Sectoral Productivity Growth - Benchmark Model				
	2.08	1.37	0.42	0.96
Labor Productivity Growth - Data				
Data	2.48	1.10	0.45	

Notes: Sectoral productivity growth is the growth rate implied by the model calibration for each sector. Labor Productivity growth is the growth rate of the labor productivity, which is computed from the data for industry and services. Home sector labor productivity growth rate is taken from Ngai and Petrongolo (2017).

Table 3 presents data and the model implied employment allocations for both the initial year (1983) and the end year (2018). The model, by construction, matches all the beginning year values as they were targets used in calibration. For the end year (2018), I targeted male market work, which pins down leisure productivity value for that year. All other moments are equilibrium outcomes. Examining the results, I find that the model allocates more male labor to the services sector and consequently less to the goods sector, compared to the data. Similarly, the model predicts more time spent on leisure and less time on home production for males. For females, the total market work in the model is less than that observed int the data. This discrepancy stems from a higher decline in goods employment and a slower increase in services employment. Similar to the male allocation, the model predicts more female leisure and less female home hours compared to the data.

Table 3: Labor Allocations - Data vs Model

	Female				Male			
	Data		Model		Data		Model	
	1983	2018	1983	2018	1983	2018	1983	2018
g	3.97	2.84	3.97	2.57	14.76	11.87	14.76	9.62
s	18.15	24.52	18.15	22.37	21.97	24.20	21.97	26.44
h	37.52	35.11	37.52	32.30	24.27	24.50	24.27	22.37
l	40.35	37.53	40.35	42.76	39.00	39.43	39.00	41.56

Notes: Each row is the sector share of total hours for a gender. g is goods sector, s is service sector, h is home production, l is leisure. 1983 values were targeted to calibrate parameter values. Male non-work time (leisure plus home production) is targeted moments for 2018. .

To better asses the distance between model and the data moments, Table 4 shows

how several moments changed over time in the data and in the model. When we compare the model results in column three to the data counterparts (column two), we see a stark difference in the gender wage gap: it increases in the model, while it decreases in the data. This divergence is driven by the endogenous labor sorting and resulting compositional effects. As female employment in the declining goods sector falls, the remaining higher productive workers drive up the average female wage in that sector. Although the same effect occurs with average male wages in the goods sector, the lower female labor share in the goods sector makes the average female wages increase more in that sector. The same mechanism generates the opposite change in the market services sector. As more women enter the expanding services sector, the average female productivity and, consequently, the average female wage declines. Since this relative decline in average productivity is more pronounced for females (given their lower initial employment share in services), the average wage declines more for females in the services sector. Consequently, the overall wage gap increases. Furthermore, as the model allocates less women to the market in 2018 compared to data, it generates around half of the observed increase in female employment. As was targeted, the male employment change is exactly matched. Since the model allocates more men to services and less to goods sectors, the service share of total employment is higher in the model.

Table 4: Other Moments - Data vs Model

	Data	Benchmark	Counterfactual	
			No Structural Transformation	No Marketization
Changes over time (in %)				
Female-to-male wage rate	9.51	-6.43	-9.09	6.87
Female market work	23.66	12.72	22.23	-20.96
Male market work	-1.81	-1.81	7.66	-19.67
Services share of employment	12.68	17.37	8.43	11.18
Shift-share Analysis for Female Employment				
Within Factors	3.16	0.40	1.61	-2.19
Between Factors	2.38	2.89	1.41	1.81
Total change	5.54	3.29	3.02	-0.38

Notes: Change in male market work is a calibration moment, hence matched. "Between Factors" is the shift-share decomposition result for between-sector factors while "Within Factors" is the shift-share decomposition result for within-sector factors. Reported values for shift-share analysis are female shares in total employment.

Table 4 also reports the result of a shift-share decomposition to analyze how much of the increase in female share of total employment is due allocation across sectors (between factors) versus allocation within a sector (within factors). The last three rows of the Table 4 demonstrate the stark difference between the data and the baseline model: In the data, within-sector factors (changes in gender composition within sectors)

generated 3.2 percentage points of the 5.6 percentage points rise in the female share of total employment. The remaining 2.4 percentage points were due to between-sector factors. In the baseline model, 0.4 of the 3.3 percentage points rise in the female share of employment came from within-sector factors, and the remaining 2.9 percentage points were the contribution of between-sector factors. The model's sizable shortcoming in generating within-sector allocation is the primary reason the benchmark produces a significantly smaller overall increase in female employment compared to the data.

To evaluate the contributions of structural transformation and marketization to the rise of female employment, the last two columns of Table 4 report counterfactual exercise results. The fourth column of the table shows results of a counterfactual scenario in which structural transformation is absent (growth in goods sector productivity is as high as the growth in services sector productivity). In this scenario, there is significantly less reallocation of labor to services. We observe a rise in total male market employment, driven by the moderate increase in services employment and relatively less shedding of labor in the goods sector compared to the benchmark. Female employment increases significantly for the same reasons: There is an increase in services sector employment and limited shedding of labor in goods sector compared to the benchmark. As the shift-share analysis documents, there is a more pronounced contribution from between sector forces while there is a smaller increase in female share of total employment. The last column isolates the effect of marketization. In this counterfactual, home sector productivity increases as much as the market services sector productivity. This strong productivity growth in the non-market sector leads to a large decline in overall market employment. The employment decline in goods sector is much more pronounced than the decline in services. Almost all of the decline is absorbed by leisure, and the share of females in total employment decreases.

The model significantly overpredicts the percentage rise of services share of employment (by around 40 percent) and underpredicts the percentage rise of female employment (by around 50 percent). It also underpredicts the closing gender wage gap (predicts an increase) and underpredicts the contribution of within-sector factors to the rise of female share in total employment. Compositional changes due to the heterogeneity of workers amplify some of the failures and mitigate some other. With exogenous sectoral productivity changes, as workers leave goods sector, the average productivity of remaining employees increases, and this increase is larger for males as more males leave the sector. For services, average labor productivity declines with increase in employment and the decline is more pronounced for females as there are more females entering (starting with a lower productivity level). The increase in the average labor productivity in goods sector and the decreases in the average labor productivity in services amplifies the expansion of services sector. Moreover, since the marginal female employee in goods sector is more likely to be more productive (as there are fewer females working in that sector), the shedding of female labor would be slower than that of males. Similarly, as the marginal female employee in services is expected to be more productive in the beginning year (as there are fewer female employees), the service sector would pull relatively more women, contributing further to the rise of fe-

male employment. As discussed above, these compositional changes also generate the rise of gender wage gap. These findings lead to the following question: How much the female productivity distribution should change for the model to match the data in these aspects? The next section answers that question.

3 Role of Changing Productivity Distribution

The benchmark simulation in Section 2 demonstrated that structural transformation alone is quantitatively insufficient, shortcoming to account for the observed rise in female employment and predicting a widening gender wage gap. How much change in the female distribution is required to generate the rise in female employment as well as the decline in gender wage gap observed in the data? To answer this, I extend the calibration exercise from the previous section to also calibrate the end year scale and shape parameters of the female productivity distribution (s_f and θ_f, respectively) to match the change in female market employment and the gender wage gap to the data counterparts. The model requires the female distribution to have both a higher scale and a higher shape parameter compared to their beginning year values. A higher scale parameter increases the mean of the distribution and a higher shape parameter reduces dispersion.[4] Parameter values are reported in Table 5. Figure **??** displays the idiosyncratic worker productivity distributions across genders and time. Female distribution in 1983 has the same mean (by assumption) and a higher shape parameter (lower dispersion) compared to males. The model implies that over time female productivity concentrated more around a higher mean.

Table 5: Model Parameters

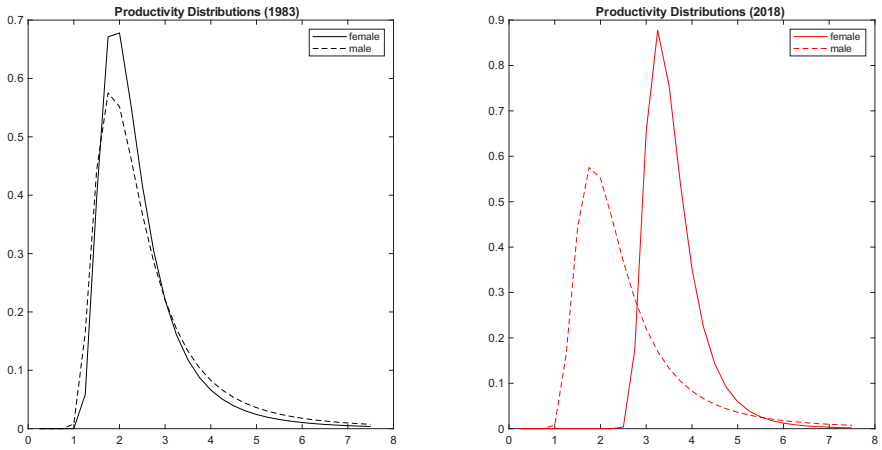
Beginning Year				End Year			
Parameter	Value	Parameter	Value	Parameter	Value	Parameter	Value
θ_m	3	A_g	1	θ_m	3	A_g	2.05
s_m	2	A_s	1	s_m	2	A_s	1.61
θ_f	3.69	A_h	2	θ_f	7.83	A_h	3.22
s_f	2	A_l	2	s_f	3.31	A_l	2.80

Notes: Subscripts $\{g, s, h, l\}$ stand for goods, market services, home, and leisure, respectively. A_j are sectoral productivity values. θ and s are shape and scale parameters of Frechet distributions for females (f) and males(m). Beginning year is 1983, end year is 2018.

This exercise also requires recalibrating the sectoral productivity values (A_j) so that the model's labor productivity growth rates remain consistent with the data. Table 5

[4]Both of the parameters affect these moments, but for the calibrated parameter values increase in mean due to scale parameter dominates and decrease in dispersion due to shape parameter dominates. The calibrated end year parameters together imply a decrease in the variance of the distribution. Also, a higher shape parameter implies a relatively thinner right tale.

Figure 1: Female Employment Elasticity Across Countries.



Notes: Calibrated female and male distributions of worker productivity.

displays the distribution and sectoral productivity values for the end year as a result of the new calibration. As Table 6 clarifies, with distribution change, we have a stronger structural transformation (the gap between goods sector and market services productivity growth rates is higher) but a weaker marketization (the gap between market services and home sector productivity growth rates is lower). This outcome is primarily driven by declining sectoral productivity in market services. With a much more productive female labor force (due to s_f and θ_f change), the model requires a much lower growth rate in A_s to generate the observed labor productivity changes in the market service sector.

Table 6: Sectoral Productivity Growth Rates (%)

	Industry	Services	Home	Leisure
Benchmark Model				
	2.08	1.37	0.42	0.96
Model with Required Distributional Change				
	2.08	0.88	0.38	1.43
Labor Productivity Growth				
Data	2.48	1.10	0.45	

Notes: Sectoral productivity growth is the growth rate implied by the model calibration for each sector. Labor Productivity growth is the growth rate of the labor productivity, which is computed from the data for industry and services. Home sector labor productivity growth rate is taken from Ngai and Petrongolo (2017).

Table 7 reports the employment allocations of the model when female productivity distribution becomes more concentrated around a higher mean over time. The model

delivers a stronger increase in female employment in this scenario, because demand for female labor relative to male labor increases with a higher female market productivity. Results in this section yield a larger increase in female employment compared to the benchmark calibration with constant productivity distributions, despite a weaker marketization. Moreover, while the female allocation between market sectors matches the data well, the model allocates less male labor to goods and more male labor to market services. As a result of this allocation, the share of market services in total market hours in the model remains higher than in the data.

Table 7: Labor Allocations - Data vs Model

	Female				Male			
	Data		Model		Data		Model	
	1983	2018	1983	2018	1983	2018	1983	2018
g	3.97	2.84	3.97	2.38	14.76	11.87	14.76	10.46
s	18.15	24.52	18.15	24.98	21.97	24.20	21.97	25.61
h	37.52	35.11	37.52	33.79	24.27	24.50	24.27	25.86
l	40.35	37.53	40.35	38.85	39.00	39.43	39.00	38.07

Notes: Each row is the sector share of total hours for a gender. g is goods sector, s is service sector, h is home production, l is leisure. 1983 values were targeted to calibrate parameter values. Male non-work time (leisure plus home production) is targeted moments for 2018.

Table 8: Other Moments - Data vs Model

	Data	Benchmark	Counterfactual		
			No Change in Distribution	No Structural Transformation	No Marketization
Changes over time (in %)					
Female-to-male wage rate	9.51	9.51	-4.67	1.58	23.63
Female market work	23.66	23.66	16.18	45.10	-2.51
Male market work	-1.81	-1.81	2.21	15.86	-11.25
Services share (employment)	12.68	16.98	19.22	1.71	13.20
Shift-share Analysis for Female Employment					
Within Factors	3.16	2.37	-0.14	5.09	-0.17
Between Factors	2.38	3.18	3.19	0.31	2.40
Total change	5.54	5.54	3.05	5.41	2.23

Notes: All but service share changes reported on the top part of the table are calibration moments, and hence are matched. "Between Factors" is the shift-share decomposition result for between-sector factors while "Within Factors" is the shift-share decomposition result for within-sector factors. Reported values for shift-share analysis are female shares in total employment.

Table 8 shows various moments for the benchmark calibration (with female productivity distribution change) and the counterfactual analysis results as well as the shift-share decomposition results. Comparing bottom panels of Table **??** and Table

8, shift-share decomposition clearly demonstrates the strong within sector change that comes with changing female productivity distribution. We can also see through a counterfactual with no change in female productivity distribution (comparing column three to column 2 in Table 8), if the female productivity distribution had not changed, we would observed an increase in gender wage gap and a muted increase in female employment, consistent with the findings in benchmark calibration with constant productivity distributions. If goods sector productivity had not grown faster than the market services (no structural transformation), both female and male market work would have risen significantly with a small narrowing of gender wage gap. Hence, structural transformation (comparing column five to column three in Table 8) reduces market work. This is because with a higher income, household also increases leisure. If the home production productivity was as high as the market services sector (no marketization), the female and male employment would decline and gender wage gap would close significantly.

4 Robustness to Country and Period Selection

Is this finding, that talent shift is necessary complement to structural transformation to understand the recent shifts in gender gap outcomes, specific to the USA? To investigate the generality of these findings, I conduct the following calibration exercise for France, Spain, Italy, and the UK: I assume these countries have the same preference parameters as the USA, thereby ensuring that the results are not driven by differences in preferences. I normalize male productivity distribution for the beginning year to have $s_m = 2$ and $\theta_m = 3$, same values as for the US. Moreover, female productivity distribution has the same scale parameter (s_f) as the male productivity distribution. Country specific parameters left to be calibrated are the beginning year values for female comparative advantage parameters ϕ_j, female productivity distribution's shape parameter (θ_f), and the sectoral productivity levels A_j. Parameters to be calibrated in the end year are the sectoral productivity levels A_j, and the female productivity distribution parameters (s_f and θ_f).

I use beginning year hours allocations for males and females, beginning year gender wage gap, and the beginning year GDP of the country relative to the US GDP.[5] Data details are reported in Section A.1. For each country, there are eight data moments to target and eight parameters must be calibrated for the beginning year. The joint calibration of θ_f, A_j and ϕ_j values are not unique, hence I calibrate these parameters to values in neighborhoods so that the female share parameter in home production, ϕ_h for each country is be the same as in the US. Results are qualitatively robust to the choice of ϕ_h level.

For the end year, I take exogenous labor productivity growth rate for market sectors from data to pin down the market sectors' productivity values (A_g, A_s). As I do not have home sector productivity growth data, I target end year home and leisure allocation for

[5]I calculate sectoral output for all countries, and use end year prices for the US to get country GDP.

males to find end year home and leisure productivity levels (A_h and A_l, respectively). Lastly, I target female market share and gender wage gap in the end year to calibrate the end year female productivity distribution parameters.

Table 9 shows the labor productivity growth rates in the data and the model implied sectoral productivity growth rates. Comparing the difference between productivity growth rates in goods and services when assuming constant female distribution, we observe that there is structural transformation in all countries, and it is the strongest in France. Similarly, there is marketization in all countries, and it is the strongest in Spain and Italy. When we let the female productivity distribution change, all countries experience a stronger structural transformation and a stronger marketization.

Table 9: Sectoral Productivity Growth Rates (%)

Sector	USA	UK	Spain	Italy	France
Data growth rates					
industry	2.48	2.13	2.14	1.56	3.17
services	1.10	1.39	0.46	0.04	1.15
Model growth rates, constant distribution					
industry	2.08	1.97	1.63	1.54	1.96
services	1.37	1.61	0.92	0.96	0.52
home	0.42	0.95	-0.88	-1.37	0.11
leisure	0.96	1.57	0.67	0.73	0.50
Model growth rates, distribution change					
industry	2.08	1.87	1.37	1.23	2.01
services	0.88	1.02	0.17	0.32	0.53
home	0.38	0.27	-1.83	-2.14	0.04
leisure	1.43	1.69	0.82	0.70	0.86

Notes: Sectoral productivity growth is the growth rate implied by the model calibration for each sector. Labor Productivity growth is the growth rate of the labor productivity, which is computed from the data for industry and services. Home sector labor productivity growth rate is taken from Ngai and Petrongolo (2017).

Table 10 reports comparative advantage values for females (shares in CES production functions) for all countries. Consistent with the literature and the US baseline, females in all countries have a comparative advantage in market services over goods sectors, and a comparative advantage in home production over market services (France is an exception, where these share parameters are equal). By construction, these countries were constrained to have the same value for home production (ϕ_h). This calibration strategy was employed because the joint calibration of all comparative advantage values and the female productivity distribution's shape parameter does not yield unique values. To overcome this challenge and better compare countries, I choose the set of parameters such that the female comparative advantage value for home production (ϕ_h)

in all countries is the same as the value for the US. Italy is an exception because of a technical issue that is most likely resulting from forcing Italy to have the same preferences as the USA.[6] The restriction for the same ϕ_h affects the beginning and end year female productivity distribution shape parameter values while the fact that that parameter has to increase over time to match the data remains.

Table 10: Female Comparative Advantage Parameters

	USA	UK	Spain	Italy	France
industry	0.23	0.18	0.14	0.25	0.29
services	0.37	0.31	0.22	0.32	0.39
home	0.39	0.39	0.39	0.43	0.39
leisure	0.39	0.41	0.29	0.14	0.41

Comparative advantage value for home (ϕ_h) for all countries is set the USA value. Italy is the exception as a lower ϕ_h requires a very high shape parameter for females in the end year to match the data.

Table 12 reports the change over time in some moments in the data in the top panel, and in the model with a constant female productivity distribution (baseline calibration) in the middle panel. Consistent with the US baseline calibration results, the model underpredicts female employment growth and predicts that the gender wage gap to moves in the opposite direction across all countries. We observe differential experiences among these countries in their growth of services sector employment share, as well as in the relative contributions of within-sector and between-sector factors to the overall rise of the female share of employment.

I also re-calibrate the end year values of the model, targeting the change in female employment share and the gender wage gap, to study how the female productivity distribution should have changed for the model to deliver the outcomes observed in the data. The bottom panel of Table 12 reports the change over time in moments with a changing female productivity distribution. Consistent with the findings in the previous section, a change in the female productivity distribution not only allows the model to align with the data in evolution of gender gaps, but it also delivers a more aligned distribution of this change between within-sector versus between-sector factors while also bringing the growth rate of services share of employment closer to the data. Moreover, the results consistently show that for all countries, the female productivity distribution must have a higher mean and a smaller dispersion compared to their beginning year

[6]Enforcing US ϕ_h value in calibration generates difficulty to calibrate end year parameter values. As ϕ_h goes down, shape parameter of the end year female productivity distribution that is needed to match the data increases. To match the ϕ_h of US, Italy has to have a larger initial shape parameter, which requires even a much larger end year shape parameter than the current calibration. This exercise suggests that preferences in Italy are likely different from the preferences in the US.

values, in order for the model to match the observed female employment increase and the closing of the gender wage gap (Table 8). This analysis confirms the generality of the main finding: a shift toward higher and more concentrated female talent is a necessary complement to structural transformation across these developed economies.[7] Table 11 presents female productivity distribution parameter values for 1983 and 2018 for all countries while Figure 2 visualizes these distributions.

Table 11: Female Distribution Parameter Values

	USA	UK	Spain	Italy	France
1983					
shape	3.69	2.90	2.05	3.69	6.87
scale	2.00	2.00	2.00	2.00	2.00
2018					
shape	7.83	6.26	5.03	12.41	7.65
scale	3.31	4.76	9.28	4.35	3.04

Calibration results for female productivity distribution, which is a Frechet distribution with scale (s) and shape (θ) parameter. Beginning year scale parameter is normalized to 2 for all countries, all other parameters are calibration outcomes.

To check whether the results are driven by the time period, I also calibrate the US economy along with France, Spain, Italy, and the UK to 2003-2015 period, following European time use data availability, and get the same qualitative results (reported in the appendix). With this relatively short sample period, observed change in female employment (and overall employment allocations) is also small. Hence, the percent changes reported in Table A.4 are somewhat larger. Regardless of these differences in magnitude, I arrive at the same conclusion: the necessity of the shift toward higher and more concentrated female talent remains. Additionally, I confirm the consistency of the findings using a longer historical period, calibrating the US economy to the moments between 1960 and 2007 reported by Ngai and Petrongolo (2017), which also yields the same qualitative results.

Overall, the calibration for the UK, Spain, Italy and France, as well as calibrations for different time periods confirm the main finding of this paper: Over time more concentrated female productivity over a higher mean is the factor, along with structural transformation, that explains the rise of the female employment along with declining gender wage gap.

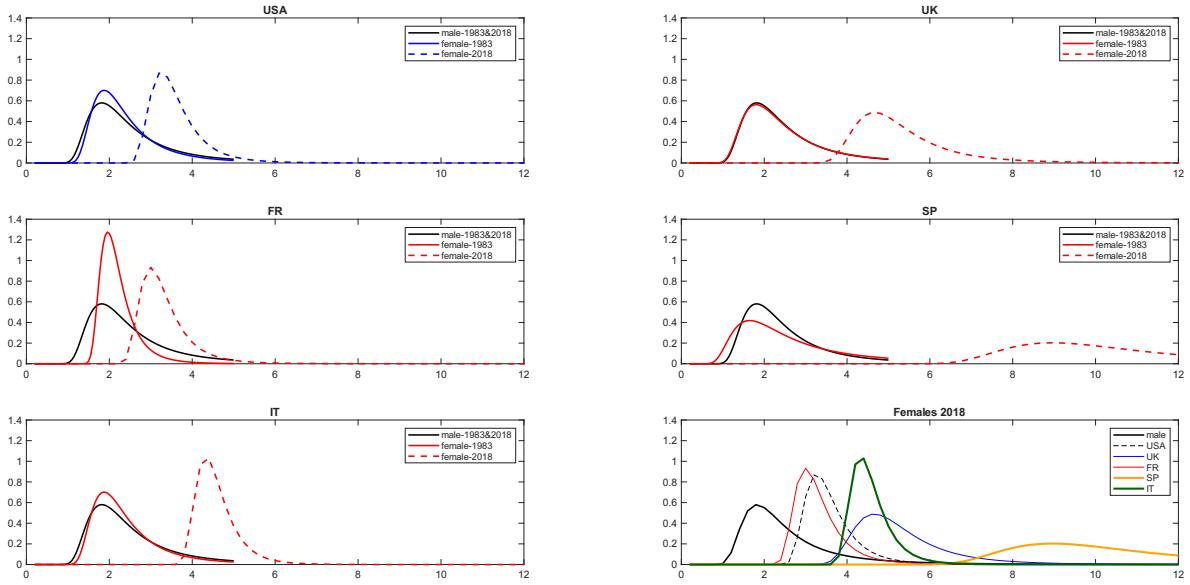
[7]Italy's data requires a more drastic change in the female productivity distribution which could be because the restricting assumption of the same preferences across countries could bind this country more.

Table 12: Other Moments as Changes Over Time - Data vs Model

	USA	UK	Spain	Italy	France
Data					
Female-to-male wage rate (%)	9.51	12.79	25.62	7.51	-1.26
Services share (employment, %)	12.68	28.99	56.66	71.43	33.35
Female market work(%)	23.66	45.07	90.67	58.77	11.58
Male market work (%)	-1.81	0.27	-1.61	8.45	-21.58
Shift-share Analysis for Female Employment					
Within Factors	3.16	3.70	8.48	3.15	3.64
Between Factors	2.38	4.84	6.24	5.24	4.76
Total Change	5.54	8.54	14.72	8.39	8.41
Model with constant distribution					
Female-to-male wage rate (%)	-6.43	-3.67	-4.89	-16.04	22.10
Services share (employment, %)	17.37	12.58	45.18	57.07	29.09
Female market work(%)	12.72	9.38	15.13	39.03	-29.40
Male market work (%)	-1.81	0.27	-1.61	8.45	-21.58
Shift-share Analysis for Female Employment					
Within Factors	0.40	0.12	-0.19	2.58	-5.76
Between Factors	2.89	1.81	3.39	2.76	3.40
Total Change	3.29	1.93	3.20	5.34	-2.36
Model with distribution change					
Female-to-male wage rate (%)	9.51	12.79	25.62	7.51	-1.26
Services share (employment, %)	16.98	16.48	54.09	63.29	31.91
Female market work(%)	23.66	45.07	90.67	58.77	11.58
Male market work (%)	-1.81	0.27	-1.61	8.45	-21.58
Shift-share Analysis for Female Employment					
Within Factors	2.37	5.78	9.31	4.73	4.34
Between Factors	3.18	2.75	5.42	3.65	4.06
Total Change	5.54	8.54	14.72	8.39	8.41

Notes: All but service share changes reported on the top part of the table are calibration moments, and hence are matched. "Between Factors" is the shift-share decomposition result for between-sector factors while "Within Factors" is the shift-share decomposition result for within-sector factors. Reported values for shift-share analysis are female shares in total employment.

Figure 2: Productivity Distributions



Notes: Calibrated female and male distributions of worker productivity.

5 Conclusion

Female employment increased significantly across developed economies over the last decades. While structural transformation, the reallocation of labor to the service sector, is a central reason for this rise, shift-share decompositions show substantial contribution from within-sector factors (changes in gender composition). This paper develops a structural transformation model augmented with worker heterogeneity and a Roy-type self-selection mechanism to endogenously account for these within-sector dynamics. The paper finds that when driven solely by observed sectoral productivity growth, the model underpredicts the rise in female market work and incorrectly predicts a widening of the gender wage gap.

The model implies that the simultaneous surge in female market work and the narrowing gender wage gap requires a shift in the supply of female talent as well as the structural transformation. To match the data, the female productivity distribution must evolve toward a higher mean and lower dispersion (more concentrated human capital). This necessity, which is robust across different time periods in the US and holds qualitatively for the UK, Spain, Italy, and France, confirms that labor supply dynamics were essential and complementary to the pull factors (demand shifts). This finding suggest that policies aimed at structural demand shifts must be coupled with simultaneous investment in female human capital accumulation to improve female labor market experience.

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.

Adamopoulos, T. and A. Akyol (2009). Relative underperformance Alla Turca. *Review of Economic Dynamics 12*(4), 697–717.

Akbulut, R. (2011). Sectoral changes and the increase in women's labor force participation. *Macroeconomic Dynamics 15*(02), 240–264.

Boppart, T. and L. R. Ngai (2021). Rising inequality and trends in leisure. *Journal of Economic Growth 26*(2), 153–185.

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

Erosa, A., L. Fuster, G. Kambourov, and R. Rogerson (2017). Hours, occupations, and gender differences in labor market outcomes.

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

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

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.

İmrohoroğlu, A., S. İmrohoroğlu, and M. Üngör (2014). Agricultural productivity and growth in Turkey. *Macroeconomic Dynamics 18*(5), 998–1017.

McDaniel, C. (2014). Average tax rates on consumption, investment, labor and capital in the OECD 1950-2003.

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

Olivetti, C. and B. Petrongolo (2016). The evolution of gender gaps in industrialized economies. *Annual Review of Economics* (Forthcoming).

Rendall, M. (2017a). Brain versus brawn: the realization of women's comparative advantage. *University of Zurich, Institute for Empirical Research in Economics, Working Paper* (491).

Rendall, M. (2017b). Female market work, tax regimes, and the rise of the service sector. *Review of Economic Dynamics*.

Rogerson, R. (2008). Structural transformation and the deterioration of European labor market outcomes. *Journal of Political Economy 116*(2), 235–259.

A APPENDIX

A.1 Data

I use average weekly hours worked, employment, and working age population data from the ILO database to compute aggregate hours worked by females and males in goods and market services sectors. Goods sector is agriculture, mining, manufacturing, utilities and construction, while market services is all the remaining sectors. Aggregate weekly hours are employment to population ratios multiplied by the average weekly hours worked for each sector-gender pair. Assuming 13 hours per day is available (excluding average time spent on sleep, personal care, etc.), I compute aggregate hours as a share of total weekly time available. To compute the sectoral labor productivity growth rates, I combine Gronningen 10 sector database (discontinued after 2013), OECD value added database (available starting from 1995 for many countries), and ILO aggregate hours worked data.

To compute time allocations in the data, I use US ATUS for 2003 on, and 2000 and 2010 European time use survey data for the other countries. I use time use survey data to compute what fraction of non-market time is allocated to home production and how that share changes over time. Using the growth rate of the home production share of non-market time I extrapolate the 1983 and 2018 values. Using ILO time allocation to market sectors and remaining non market time, I find overall allocation of time to goods, market services, home, and leisure for both genders. I use this method to get overall time trends (increase in leisure time, decline in both market work and home production time) in the data that are consistent with findings in the literature.

I get gender wage gap data from the ILO, with the exception of Spain. I get consumption and income tax data from the McDaniel (2014). I use Penn World Table data to compute relative per capita GDP values for 1983 and 2018.

A.2 Supplementary Tables

Table A.1: Average Consumption and Income Tax Rates

	USA	UK	Spain	Italy	France
1983					
consumption	0.10	0.12	0.08	0.12	0.07
income	0.08	0.15	0.12	0.15	0.22
2018					
consumption	0.11	0.13	0.09	0.14	0.12
income	0.08	0.17	0.17	0.21	0.24

Notes: Table documents average consumption and income tax rates, and it is taken from McDaniel (2014).

Table A.2: Female Comparative Advantage Parameters

	USA	UK	Spain	Italy	France
ind	0.25	0.20	0.17	0.29	0.28
srv	0.42	0.39	0.34	0.41	0.45
home	0.43	0.43	0.43	0.43	0.45
leisure	0.33	0.25	0.08	0.06	0.13

Table reports calibration results for female comparative advantage values for 2003-2015 period. Comparative advantage value for home (ϕ_h) for all countries is set the USA value. France is the exception as a lower ϕ_h requires a very high shape parameter for females in the end year to match the data.

Table A.3: Female Distribution Parameter Values

	USA	UK	Spain	Italy	France
1983					
shape	3.44	3.24	2.93	6.87	4.46
scale	2.00	2.00	2.00	2.00	2.00
2018					
shape	4.20	4.68	4.44	23.45	6.74
scale	2.33	2.49	4.03	2.81	2.42

Table reports calibration results for female productivity distribution parameter values for 2003-2015 period. Female productivity distribution is a Frechet distribution with scale (s) and shape (θ) parameter. Beginning year scale parameter is normalized to 2 for all countries, all other parameters are calibration outcomes.

Table A.4: Other Moments as Changes Over Time - Data vs Model

	USA	UK	Spain	Italy	France
Data					
Female-to-male wage rate (%)	5.89	1.83	10.27	4.61	0.60
Services share (employment, %)	3.56	14.72	34.01	24.89	15.96
Female market work (%)	-1.04	28.62	36.45	22.96	21.44
Male market work (%)	-2.98	15.28	1.44	8.51	9.52
Within Factors	-0.31	-0.51	1.22	-0.30	-0.58
Between Factors	0.79	3.11	5.74	3.15	3.11
Total Change	0.48	2.60	6.96	2.85	2.52
Model with constant distribution					
Female-to-male wage rate (%)	-1.90	-11.32	8.39	-15.15	-11.00
Services share (employment, %)	1.39	5.78	5.23	3.71	5.64
Female market work (%)	-0.44	33.26	-3.54	28.39	25.46
Male market work (%)	-2.98	15.28	1.44	8.51	9.52
Within Factors	0.33	2.20	-2.00	3.42	2.21
Between Factors	0.30	1.26	0.87	0.44	1.11
Total Change	0.63	3.46	-1.13	3.86	3.32
model with distribution change (%)					
Female-to-male wage rate (%)	5.89	1.83	10.27	4.61	0.60
Services share (employment, %)	0.89	6.44	12.00	6.48	6.59
Female market work (%)	-1.04	28.62	36.45	22.96	21.44
Male market work (%)	-2.98	15.28	1.44	8.50	9.52
Within Factors	0.28	1.15	4.77	2.05	1.18
Between Factors	0.20	1.46	2.20	0.80	1.34
Total Change	0.48	2.60	6.96	2.85	2.52

Notes: Table reports calibration outcomes for 2003-2015 period. All but service share changes reported on the top part of the table are calibration moments, and hence are matched. "Between Factors" is the shift-share decomposition result for between-sector factors while "Within Factors" is the shift-share decomposition result for within-sector factors. Reported values for shift-share analysis are female shares in total employment.

Table A.5: USA Moments as Changes over Time

			Changing Distribution			
	Data	Bm(a)	Cd	Cf - no dt	Cf - no st	Cf - no mt
Female-to-male wage rate (%)	30.65	-2.46	30.65	11.88	30.61	26.97
Female market work (%)	31.82	-6.21	31.82	-8.24	155.83	43.65
Male market work (%)	-14.89	-14.89	-14.89	-5.94	-3.14	-11.07
Services share of employment (%)	21.95	21.32	21.25	22.73	-10.19	23.88
Shift-share Analysis						
Within Factors	6.72	-0.73	6.65	-3.51	23.82	7.23
Between Factors	3.43	2.88	3.49	2.97	-0.42	3.95
Total change	10.14	2.15	10.14	-0.54	23.40	11.17

Notes: Table reports calibration outcomes for 1960-2007, using moments reported in Ngai and Petrongolo (2017). "Bm(a)" is the benchmark calibration that assumes female productivity distribution is constant over time. Male market work is a target moment. "Cd" is the calibration results with changing female productivity distribution. All but service share changes for this column are calibration moments, and hence are matched. The remaining three columns are different counterfactual exercises: "Cf - no dt" is counterfactual exercise with constant female productivity distribution, "Cf - no st" is the counterfactual exercise with no structural transformation, "Cf - no mt" is the counterfactual exercise with no marketization. "Between Factors" is the shift-share decomposition result for between-sector factors while "Within Factors" is the shift-share decomposition result for within-sector factors. Reported values for shift-share analysis are female shares in total employment.