

The Role of Gender in Employment Polarization*

Fabio Cerina[†] Alessio Moro[‡] Michelle Petersen Rendall[§]

First version, February 2016 - This version, January 2017

Abstract

We document that employment polarization in the 1980-2008 period in the U.S. is largely generated by women. For the latter, employment shares increase both at the bottom and at the top of the skill distribution, generating the typical U-shape polarization graph, while for men employment shares decrease in a similar fashion along the whole skill distribution. We show that a canonical model of skill-biased technological change augmented with a gender dimension, an endogenous market/home labor choice and a multi-sector environment accounts well for gender and overall employment polarization. The model also accounts for the absence of employment polarization during the 1960-1980 period, which is due to the flat behavior of changes in women's employment shares along the skill distribution, and can reproduce the different evolution of employment shares across decades during the 1980-2008 period. The faster growth of skill-biased technological change since the 1980s accounts for a substantial part of the employment polarization generated by the model.

JEL Classification: E20, E21, J16.

Keywords: Job Polarization, Gender, Skill-biased technological change, Home Production.

*We thank Vasco Carvalho, Luigi Guiso, Nezh Guner, Moshe Hazan, Vahagn Jerbhasian, Rachel Ngai, Barbara Petrongolo, Xavier Raurich, Christian Siegel, Satoshi Tanaka, Marc Teigner, Montserrat Vilalta, Fabrizio Zilibotti and seminar participants at LSE, University of Barcelona, University of Cambridge, Monash University, University of New South Wales, University of Zurich, the CEPR Macroeconomics and Growth Programme Meeting (London Business School), the Barcelona GSE Summer Forum 2016, the V Workshop on Structural Transformation and Macroeconomic Dynamics (Kent), the 2016 RIDGE December Forums in Buenos Aires, Montevideo and Rio de Janeiro, the 3rd Workshop of the Australasian Macroeconomic Society (Brisbane), the SIE (Bocconi), the SAEe (Bilbao), and the VII Workshop on Institutions, Individual Behavior and Economic Outcomes (Alghero) for the useful comments. The usual disclaimers apply.

[†]University of Cagliari and CRENoS. E-mail: fcerina@unica.it.

[‡]University of Cagliari. E-mail: amoro@unica.it.

[§]University of Zurich. E-mail: michelle.rendall@econ.uzh.ch.

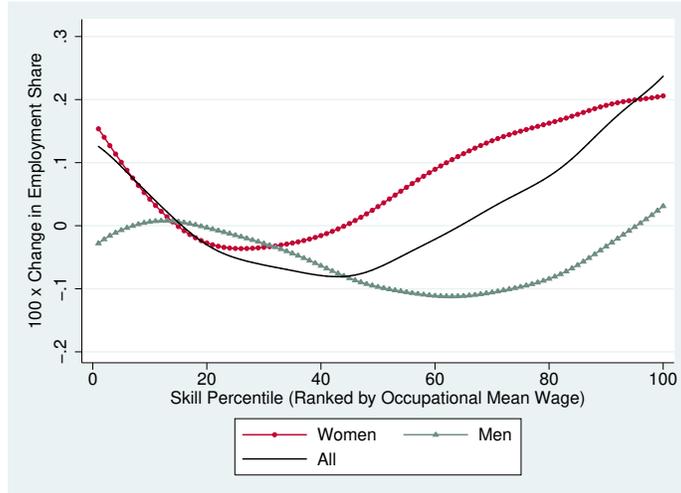


Figure 1: Changes in employment shares in the U.S. between 1980 and 2008 by skill percentile using a locally weighted smoothing regression. Data are from Census IPUMS 5 for 1980 and Census American Community Survey for 2008.

1 Introduction

Employment polarization in the U.S. has been extensively documented. [Autor and Dorn \(2013\)](#) show the change from 1980 to 2008 in the share of U.S. employment by skill rank and find an increase in employment shares both at the bottom and the top of the skill distribution, combined with a decline in the middle. This pattern, reported by the black continuous line in Figure 1, has become a well-known stylized fact. Less well known in the literature instead, is the behavior of job polarization when distinguishing by gender, which we also report in Figure 1. As the red line suggests, the overall phenomenon of job polarization is mainly driven by women. These individuals are responsible for the rise at the bottom and most of the rise at the top of the skill distribution, and for a small (relative to the aggregate economy) decline in the middle. Men instead, see their employment shares decline along the whole skill distribution except at the very top. The U-shape at the aggregate level emerges from the aggregation of these two groups.¹

In this paper we establish a number of facts on employment polarization that have not been reported in previous literature. First, starting from the evidence in Figure 1, and focusing on marital status, we document that the increase of female employment shares at the top of the distribution is mainly due to married women, while single women are responsible for the increase of employment shares at the bottom. We then show that changes of employment shares of single individuals are in general flatter than those of married along

¹Up to approximation due to the locally weighted smoothing regression, the black solid line is the vertical sum of the gender lines. See section 3 for the formula used to compute polarization by gender.

the skill distribution. This evidence raises the issue of why employment shares along the distribution evolve differently depending on the marital status. A second set of observations concerns employment polarization by broad sectors of economic activity. We document that changes in employment shares are U-shaped and positive along the skill distribution for the services sector, while they are flat for manufacturing. This observation suggests that the U-shape at the aggregate level is the result of the aggregation of a different pattern of the two sectors. Finally, we document that during the 1960-1980 period the change of employment shares of women along the skill distribution is flat, and there is no job-polarization. This fact requires an understanding of what triggered the different behavior of women after 1980, as men display no employment polarization before and after that date.

The facts discussed above suggest that to gain additional insights on the process of employment polarization, a theory that explains the observed demographic and sectoral differences in employment changes is needed. In this paper, we build such a theory by focusing on the following observations on the polarization era (commonly referred to as the 1980-2008 period). First, during this period the average growth of the skill premium of college graduates relative to workers with less than a college degree is substantial, while the same average growth is around zero between 1963 and 1980 ([Heathcote, Storesletten, and Violante \(2010\)](#) and [Acemoglu and Autor \(2011\)](#)). The increase in the skill premium incentivizes education during the 1980-2008 period, with an increase of a factor of 2.24 in the fraction of educated women, compared to a factor of 1.40 in the fraction of educated men. Second, after 1980 the share of home production (traditionally a female intensive sector) in total value added starts declining steadily until the end of the 2000s, while it has been flat during the rest of the post-war period ([Moro, Moslehi, and Tanaka \(forthcoming\)](#)). This decline coincides with a rise in the market share of services substitutable to home production, which are typically low-skilled services ([Autor and Dorn \(2013\)](#) and [Bridgman \(2016\)](#)) and an acceleration of modern market services with respect to the pre-1980 period ([Moro, Moslehi, and Tanaka \(forthcoming\)](#)). To see how these facts can be drivers of job polarization, consider the following theoretical environment. Increasing skill-biased technological change improves market opportunities for high-skilled individuals. This induces a high-skilled woman currently working at home to enter the labor market and obtain a high-skilled job. This event has three potential effects on employment shares. First, it increases employment shares at the top of the skill distribution. Second, as the agent abandons home production, she is likely to purchase substitutes for this in the market, typically represented by low-skilled services. By increasing the demand for low-skilled services the agent fosters an increase in employment shares of low-skilled individuals, who represent the bulk of employment in that market sector. Finally, as the change in employment shares at the top and the bottom of

the skill distribution is positive, the change of employment shares in the middle turns out to be negative. This example suggests that by writing a model that allows to focus on the effect of skill-biased technological change on the home/market working decision of particular demographic groups, it is possible to generate a U-shaped pattern of changes in employment shares along the skill distribution as the one observed in Figure 1 for females and for the aggregate economy.

With the above intuition in mind, we extend the canonical model of skill-biased technological change (Acemoglu and Autor (2011)) by introducing the three building blocks of our theory: i) a gender dimension; ii) an endogenous home/market labor supply; and iii) a multi-sector environment. Such a model allows us to obtain predictions about changes in employment shares of the two gender along the skill distribution as skill-biased technological change occurs. We stress here that our three building blocks are crucial for the model to reproduce the data pattern in Figure 1. In fact, it is well known that in general the canonical model is not able to make predictions on job polarization.² This is because, with only one good produced in the economy, there are no interesting goods-demand effects that emerge from changes in employment shares of skilled and unskilled individuals. In our setting instead, in addition to technological change, the demand for a particular type of good concurs to determine employment shares of the various types of workers employed in that sector.

We thus assume that there are three market sectors and a home sector. The three sectors are *modern services* (services without a home produced counterpart), *substitutable services* (services with a home produced counterpart) and *manufacturing*. We borrow this characterization of the structure of the economy from the literature of structural transformation for two reasons. In the first place, following the intuition given above, we require an environment in which there is a market sector that produces an output that is substitutable to home production, and a modern sector attracting most high-skilled individuals when skill-biased technological change occurs. Secondly, we need to model different employment opportunities in the market for men and women. Recent work suggests women have a comparative advantage in services relative to manufacturing.³ We thus assume that a female agent with the same characteristics of a male agent has a comparative advantage in both kinds of services (modern and substitutable) relative to manufacturing.⁴ This is the only difference between a man and a women in the model.

²Acemoglu and Autor (2011).

³See Ngai and Petrongolo (2013) and references therein.

⁴In addition, we also assume gender-biased technological change within each firm, following the specification in Heathcote, Storesletten, and Violante (2010). We do this because, although the mechanism working through comparative advantage makes the gender wage gap close endogenously in the model, its effect is not quantitatively strong enough, as discussed in Rendall (2010) and Ngai and Petrongolo (2013).

Agents in the economy are heterogeneous in that each agent is born with a triple of skills, one for each market sector. Each of these skills determines the amount of efficiency units per unit of time that the agent can supply in the corresponding market sector. Agents are also allowed to obtain education by paying a cost. If an agent becomes educated she increases her skill levels by a certain amount. Thus, for a unit of time supplied and the same skill level, the productivity and the wage received by the educated individual are larger than those of the uneducated. An agent, taking as given market wages, makes a contemporaneous decision on the sector in which to work and whether to obtain education or not. Finally, the household side of the model is closed by determining the marital status of each agent. We assume in the model that a fraction of agents is single and the rest is paired to an agent of the other gender to form a two-person household. The difference between the two types of agents is that a married couple maximizes the unique utility function of the household and each agent participates in home production.

Each market sector is given by a competitive representative firm that employs four types of labor: educated males, educated females, uneducated males and uneducated females. It is important to note here that all sectors can employ any type of worker, by gender, skill and education level. However, the proportions of these groups will be different in the three market sectors and calibrated to the data. The production function of each market sector is affected by three types of exogenous technological change: labor productivity growth, skill-biased technological change and gender-biased technological change. We also assume that the home production of each household is produced by a competitive firm. The only technological change at home is labor productivity growth.

We calibrate the model to two equilibria representing the years 1980 and 2008 to match a set of targets in the data, and evaluate its performance in replicating the main facts of job-polarization. The two equilibria differ in the following exogenous dimensions: i) the level of labor productivity of market sectors and the home sector; ii) the level of skill-biased technology; iii) the level of gender-biased technology; and iv) marriage rates. Given these differences, the model endogenously generates heterogeneous changes of employment shares along the skill distribution. Our first contribution is to show that the model replicates fairly well employment polarization by gender, by marital status and by sector.

We then use the model to study why polarization emerges in the 1980-2008 period.⁵ To answer this question we use the calibrated model to “forecast” changes in employment shares from 1960 to 1980, by feeding trends of exogenous factors for that period. The data

⁵Without considering agricultural occupations [Barany and Siegel \(2015\)](#) suggest that job-polarization starts in the 1950s, but it is more pronounced in the 1980-2008 period. Here we include agriculture and find not polarization between 1960 and 1980.

show that the 1960-1980 period is characterized by an increase in employment shares of women that is homogeneous along the skill distribution. Men instead display a monotone behavior of changes in employment shares, with those in the bottom part of the distribution displaying a negative change and those at the top a positive one. The resulting overall shape is a monotone function along the skill distribution. Thus, the data suggest that employment polarization is absent mainly because women do not display the same U-shape of the 1980-2008 period. When we feed the model with the exogenous trends for the 1960-1980 period, this accounts well for the flat behavior of changes in women's employment shares and reasonably well for the monotone behavior of men's employment shares. By running counterfactual experiments we find that the absence of job-polarization is largely due to the absence of skill-biased technological change during that period.

Finally, we use the model to analyze employment polarization by decade during the 1980-2008 period. [Acemoglu and Autor \(2011\)](#) show that the shape of the black line in [Figure 1](#) is the result of a different evolution of employment shares in the three decades 1980-1990, 1990-2000, and 2000-2008.⁶ In particular, the polarization graph displays a clockwise tilting over time, with the increase at the top of the skill distribution determined mainly in the 1980-2000 period, with the increase at the bottom being a feature of the 2000-2008 period. The model reproduces fairly well such tilting, due to a time-varying effect that (a constant growth of) skill-biased technological change has on the equilibrium of the model. In particular, the tilting is due to indirect effects of skill-biased technological change that emerge over time: consumption spillovers from skilled to unskilled individuals and q-complementarity in production between educated and uneducated workers. We thus conclude that, in contrast with the standard result of the canonical model, skill-biased technological change is a first order driver of employment polarization in our setting.

The remainder of the paper is as follows. Section 2 discusses the related literature; section 3 establishes some facts on employment polarization in the U.S. that has not been considered in previous literature; section 4 presents the model; section 5 discusses the calibration and section 6 provides the benchmark results; section 7 presents the comparison between the model and the data for the 1960-1980 period; section 8 analyzes the different behavior across decades during the 1980-2008 period. Finally, section 9 concludes.

2 Related Work

This paper links three fields of research that so far have intersected only marginally. The one on the effect of female labor force participation on macroeconomic outcomes, the one on

⁶See figure 10 in [Acemoglu and Autor \(2011\)](#).

structural transformation, and that on employment polarization. We connect these strands of the literature by showing that the process of job polarization can be accounted for by women entering the labor market in a multi-sector environment.

Heathcote, Storesletten, and Violante (2010) use a dynamic one-sector heterogeneous agents model with both skill-biased and gender-biased technological change to study the rise of wage inequality in the U.S. They find that women participating more in the labor market over time play a key role in shaping this process. Here we study the effect of increasing market hours of women on changes in employment shares along the skill distribution. To do this we introduce the production function used by Heathcote, Storesletten, and Violante (2010) for each market sector in our model.⁷ As discussed above, the multi-sector assumption is key to generate a demand for services that substitute for home production when women work more in the labor market. This mechanism is described in Rendall (2015) and Ngai and Petrongolo (2013) in a multi-sector model with home production. They describe how the process of *marketization*, occurring together with structural transformation, implies that women progressively abandon home production to work in the market.⁸ Ngai and Petrongolo (2013) show that marketization and structural transformation explain together a fraction of the evolution of the gender wage gaps of wages and hours in the U.S.

Recently, several contributions proved that the process of structural transformation affects several dimensions of the macroeconomy, including aggregate productivity (Duarte and Restuccia (2010) and Herrendorf and Valentinyi (2012)), growth (Moro (2015)), volatility (Carvalho and Gabaix (2013) and Moro (2012)), the amount of skill-biased technological change (Buera, Kaboski, and Rogerson (2015)) and, last but not least, employment levels (Rogerson (2008)). However, few works relate this process to job-polarization.⁹ Autor and Dorn (2013) provide an explanation of job-polarization based on a mechanism that has a flavor of structural transformation. They show how a two-sector environment with high-skilled workers, low-skilled workers and capital can generate employment polarization when there is technological change that reduces the price of capital over time. On another note, Barany and Siegel (2015) are the first to suggest that structural transformation can *per-se* be a main driver of employment polarization. By assuming a utility function in high-skilled services, low-skilled services and manufacturing with a low elasticity of substitution, productivity

⁷Thus, skill-biased technological change occurs also in the substitutable services sector.

⁸The relationship between home production and structural transformation has been extensively studied in the literature. See among others Rogerson (2008), Ngai and Pissarides (2008) and Rendall (2015). Consistent with a lower working time at home, Bridgman (2016) and Moro, Moselehi, and Tanaka (forthcoming) show that, when measuring home production at factor prices, the value added share of home in total value added (i.e. GDP plus home production) starts declining after 1980.

⁹In recent work Duernercker and Herrendorf (2016) study the relationship between structural change and the change in occupations composition in the U.S. but don't focus on employment polarization.

trends of these three sectors imply that the share of manufacturing shrinks with respect to the other two sectors. As by definition low-skilled services employ mostly workers at the bottom of the skill distribution, high-skilled services those at the top, and manufacturing the middle ones, the process of structural transformation generates job polarization in their environment. While we use a specification of preferences similar to [Barany and Siegel \(2015\)](#), and a Roy-type model, we depart in several dimensions from their framework. First, we allow for different labor inputs by gender and education, and skill-biased and gender-biased technological change in production. This implies that each of our market sectors employ all types of workers. Second, we construct polarization graphs from the model’s outcome, which allow us to make a close comparison with the data by skill level. Finally, and most importantly, we show that employment polarization is largely a female phenomenon.

Our modeling strategy is also related to the intuition discussed in [Manning \(2004\)](#) and [Mazzolari and Ragusa \(2013\)](#). The idea is that consumption “spillovers”, i.e., an increase in high-skill workers in the market, who have a high opportunity cost of working at home, increases the demand for services in the market that have a home counterpart.¹⁰ While [Manning \(2004\)](#) and [Mazzolari and Ragusa \(2013\)](#) take a local approach, by correlating an increase in high-skilled workers with the demand for low-skilled workers in the same geographical area, we show that this mechanism is quantitatively relevant in general equilibrium and it is mainly due to female agents increasing participation in the labor market. Closely related to this idea is also the work of [Hazan and Zoabi \(2015\)](#), who argue that the increase in income inequality over the last thirty years created a group of women who can afford services that are substitutable to home (in particular child care), and another one which supplies these services. They find that, opposite to the past, highly educated women increased their fertility rate during the 2000s, due to the reduction in the relative cost of child care in the market. Interestingly, this period coincides with a large increase of employment shares at the bottom of the distribution, which is captured by our model when we analyze polarization by decades.

Finally, note that the above considerations about gender and female labor force participation are only partially addressed in the polarization literature. [Acemoglu and Autor \(2011\)](#) provide evidence on *wage polarization* by gender, but not on employment polarization. In that chapter, they also provide a description of the canonical model with skill-biased technological change and show why it cannot address job-polarization issues. They suggest that a theory of job-polarization should consider a clear distinction between skills and tasks. There

¹⁰[Cortes and Tessada \(2011\)](#), instead, find that an increase in the supply of immigrants (typically producing services that are substitutable to home production) increases market hours of women at the top of the skill (i.e. wage) distribution.

are several advantages in doing this, in particular that of being able to study one of the main drivers of job-polarization, which is the process of *routinization*.¹¹ Here we take another approach and show that considering gender and home production, a model of skill-biased technological change can generate a pattern of employment polarization that is comparable with the one in the data. To the best of our knowledge this is the first general equilibrium model that can be used to produce polarization graphs that are comparable to the ones commonly used in the literature to analyze the data, like those in Figure 1.¹²

3 Facts on Employment Polarization

Figure 1 in the introduction is obtained by computing, for each percentile i , the formula

$$\frac{H_{i,2008}}{H_{2008}} - \frac{H_{i,1980}}{H_{1980}}, \quad (1)$$

where H_t is total hours worked in the economy in year t and $H_{i,t}$ is total hours worked in percentile i in year t . Consider now the following decomposition of this formula

$$\frac{H_{i,2008}}{H_{2008}} - \frac{H_{i,1980}}{H_{1980}} = \left(\frac{H_{i,2008}^f}{H_{2008}} - \frac{H_{i,1980}^f}{H_{1980}} \right) + \left(\frac{H_{i,2008}^m}{H_{2008}} - \frac{H_{i,1980}^m}{H_{1980}} \right) \quad (2)$$

where $H_{i,t}^f$ is total hours worked by women in percentile i in year t , and $H_{i,t}^m$ is the corresponding measure for men. The equality follows from the fact that total hours in percentile i in year t are given by female plus male hours, $H_{i,t} = H_{i,t}^f + H_{i,t}^m$. The first term on the right hand side of (2) gives the the red line in Figure 1 while the second term provides the green line in the same figure. In this section we use decompositions of equation (1) to establish six facts on job-polarization in the U.S.¹³ All facts, except for number 5, refer to the period 1980-2008. Note that we define employment-polarization as a situation in which employment shares increase at the bottom and at the top of the skill distribution, with a decline in the middle.¹⁴

Fact 1 (Gender): There is employment polarization for women and not for men (top-left panel of Figure 2).

¹¹See Autor, Levy, and Murnane (2003) and subsequent literature.

¹²In contemporaneous research, Cortes, Jaimovich, and Siu (2016) document that one third of the disappearance of routine occupations are due demographic changes in the U.S. In a similar vein, we show here that demographics is key for the whole process of employment polarization. In section 3 below we also discuss some evidence on gender and routine occupations.

¹³To do this, we use versions of (2) that consider different subgroups of the population.

¹⁴Thus, when the graph displays a U-shape, but employment shares at the bottom of the distribution are negative, there is no employment polarization.

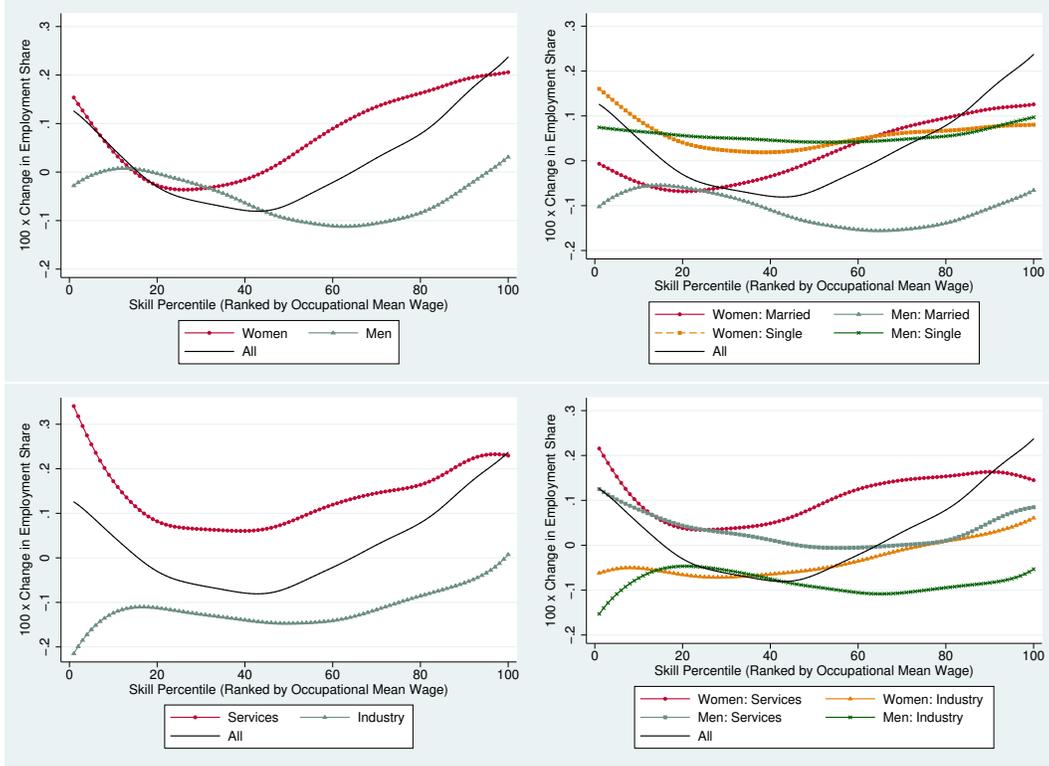


Figure 2: Job polarization by gender, marital status, and broad sector of economic activity.

Fact 2 (Marital Status): Married women contribute to the increase at the top of the distribution more than single women and men. Single women contribute to the increase at the bottom of the distribution more than married women and men. Single and married men display a flat behavior along the skill distribution, with the former displaying positive changes and the latter negative ones (top-right panel of Figure 2).

Fact 3 (Sectors): Changes in employment shares of occupations in services display a U-shaped behavior and are positive along the whole skill distribution. Changes in employment shares of occupations in manufacturing display a relatively flat (with respect to services) behavior, and are negative along the whole distribution (bottom-left panel of Figure 2).

Fact 4 (Gender and Sectors): In services, both women (to a larger extent) and men (to a lesser extent) display a U-shaped behavior of changes in employment shares. In manufacturing, changes in employment shares of men display a flat behavior along the distribution, while women increase their employment shares in that sector at the top of the distribution (bottom-right panel of Figure 2).

Fact 5 (Employment polarization before 1980): Employment polarization is absent in the 1960-1980 period. Changes in employment shares of women are homogeneous along

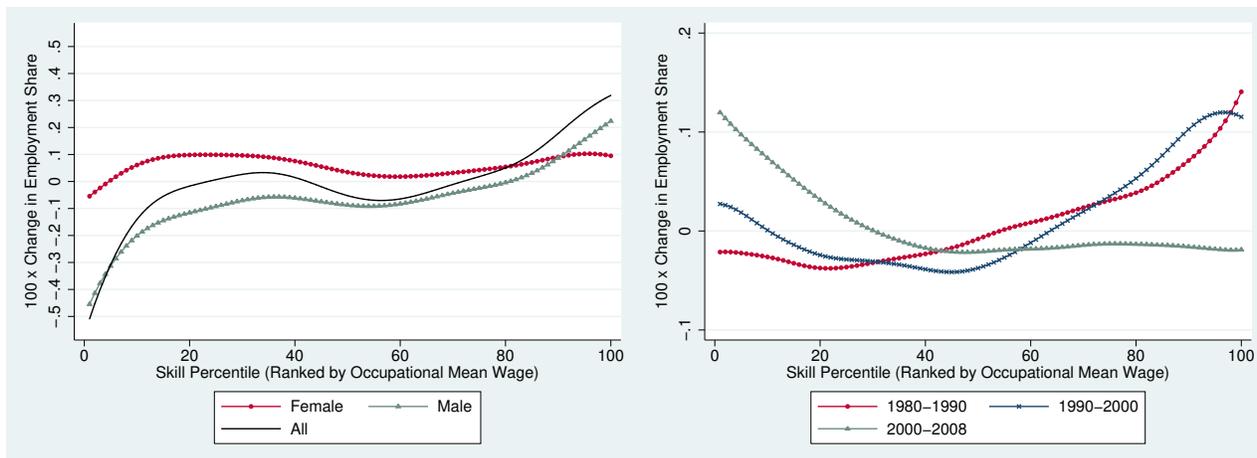


Figure 3: Job polarization in the 1960-1980 period (left) and in the three decades during the 1980-2008 period (right).

the skill distribution, while those of men are increasing along the distribution for (almost) any percentile (left panel of Figure 3).

Fact 6 (Employment polarization by decade): As documented in [Acemoglu and Autor \(2011\)](#), the change in employment shares is monotonically increasing in the 1980-1990, U-shaped in the 1990-2000 and monotonically decreasing in the 2000-2008 period (right panel of Figure 3).

Taken together, this evidence suggests a key role of women in generating employment polarization. This group increases employment shares in services, especially at the top and at the bottom of the distribution. Men instead see on average a decline of their employment shares, and the bulk of this decline occurs in the manufacturing sector. In addition, the marital status appears to play an important role, especially for women, while for men it only provides a level effect. Also, changes in employment shares of single women and single men are similar along the whole distribution, while those of married women and married men diverge substantially. Thus, a theory that aims at accounting for overall employment polarization should potentially explain the different role of the various demographic groups in shaping this phenomenon. In the following section we use a modified version of the canonical model of skill-biased technological change and show that this is broadly consistent with Facts 1-6.

Before moving to the model, we stress here that the mainstream explanation for the decline of employment shares in the middle of the skill distribution is *routinization*.¹⁵ This process, driven by the rise of information technology, makes workers performing jobs con-

¹⁵See, for instance, [Acemoglu and Autor \(2011\)](#), [Acemoglu and Autor \(2012\)](#), [Autor and Dorn \(2013\)](#).

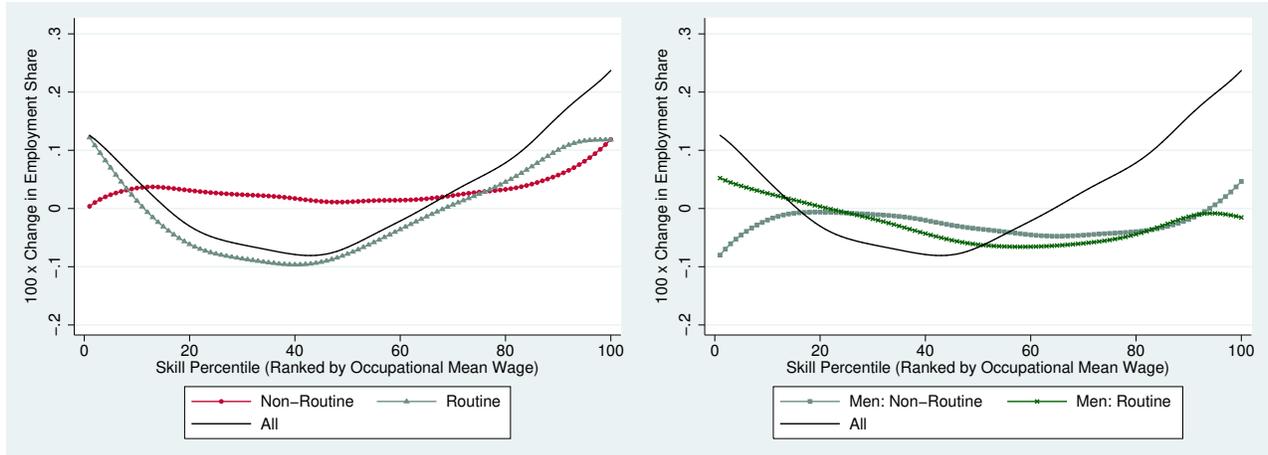


Figure 4: Job polarization in routine and non-routine occupations. Left: men routine and non-routine. Right: overall routine and non-routine.

taining a large share of routinary tasks redundant, as the latter are taken up by computers. The evidence provided in the literature suggests that these type of occupations were in the middle of the skill distribution in 1980. Thus, if routinization is the main driver of employment shares in the middle of the distribution, polarization graphs should be different for routinary and non-routinary occupations. We report these graphs in the left panel of Figure 4.¹⁶ The panel confirms the well know fact in the literature that the decline of employment shares in the middle of the distribution is due to routine occupations disappearing. However, the green line also suggests that even the increase at the bottom and at the top is due largely to routine occupations. Put it differently, the U-shape is entirely generated by routine occupations, with the non-routine occupations displaying changes in employment shares that are similar along the distribution, with a larger increase only after the 90th percentile.

The gender dimension allows us to address the routinization hypothesis from another angle. In fact, the difference between routinary and non-routinary occupations should be especially evident for men, the demographic group that loses the bulk of employment shares over time, as shown in the left panel of 1. Instead, the left panel of Figure 4 shows that polarization graphs for these two groups trace each other along the whole skill distribution and are remarkably different from the black aggregate line. Thus, the gender dimension suggests that in addition to the routine/non-routine dichotomy, other channels might be determining the U-shape at the aggregate level.

Finally, in Table 1, we perform an accounting exercise to measure the contribution to job-polarization of different pairs of categories that sum to the total population: female/male, industry/services and routine/non routine. In each exercise we fix the hours of one of the

¹⁶We use the definition of routine and non-routine occupations in Autor and Dorn (2013).

Table 1: *Job-Polarization Accounting*

% Explained by	<i>D1</i>	<i>D4</i>	<i>D5</i>	<i>D9</i>	<i>D10</i>
Women (Men 1980)	66	68	18	131	90
Men (Women 1980)	23	17	62	-8	30
Industry (Services 1980)	-166	144	144	-19	21
Services (Industry 1980)	169	14	1	92	79
Routine (Non-Routine 1980)	85	115	103	74	64
Non-Routine (Routine 1980)	8	-12	-1	25	36

two categories to their 1980 level. Thus, for instance, the first row of the table reports how much of the variation in employment shares by decile can be accounted for by changes in women hours, as men hours are held fixed to their 1980 level.¹⁷ We report the deciles that characterize job polarization, namely the first, the middle and the last deciles. By fixing men hours to 1980, the change in women hours can account for a large share of the change in deciles 1, 4, 9 and 10. Instead, by fixing women hours to 1980, men can account for a significant portion of job-polarization for decile 5. When repeating the exercise for industry/services, it appears how services account for deciles at the bottom and the top of the distribution and industry for those in the middle. Instead, when distinguishing routine and non-routine, the third and fourth lines show how routine occupations account for a large portion of employment polarization in each decile. Non-routine occupation instead, account for a small portion at all percentiles. While routinization represents an important driver of the disappearance of middle skill occupations, our aim in the next sections is to show that when considering the gender dimension, skill-biased technological change is a key determinant of Facts 1-6. We thus leave the introduction of routinization in this context for future research.

4 Model

The model economy consists of three market sectors, modern services, ms , substitutable (to home) services, ss , manufacturing goods g , and a home sector, h . The environment is static such that given the fundamentals at time t , the equilibrium of the model is uniquely determined in that period.

¹⁷By construction then, the sum of each pair of rows does not sum to 100.

4.1 Agents

There are two masses of agents in the economy, one of female agents and one of male agents. The female and the male population can be of different size. Both types of agents are heterogeneous such that each one has a skill level to work in services that are substitutable to home production (ss), a skill level to work in manufacturing (g for goods) and a skill level to work in modern services (ms).¹⁸ Hence, each agent is endowed with a triple of skills $a^i = [a_{ss}^i, a_g^i, a_{ms}^i]$, where $i = f, m$, and f stands for female and m for male. Thus, there exist two density functions of agents with characteristics $[a_{ss}^i, a_g^i, a_{ms}^i]$. Each characteristic is between a_{min} and a_{max} and an agent of type i is perfectly identified by a point in the support of the trivariate distribution $f(a^i) = f(a_{ss}^i, a_g^i, a_{ms}^i)$.

Each agent is also endowed with one unit of time. She splits this between work at home (l) and work in the market ($1 - l$). Thus, a unit of time of agent of type i can supply: i) a_{ms}^i efficiency units of labor to production in sector ms ; ii) a_g^i efficiency units of labor to production in sector g ; iii) a_{ss}^i efficiency unit of labor to production in sector ss ; and iv) 1 efficiency unit of labor to production in the home sector h .

4.2 Education and job decision

The education level and the sector where the agent works are jointly chosen. There are two different education levels $e = 0, 1$. When the agent chooses $e = 1$, she pays the fixed cost χ^i and increases her ability from a_j^i to $(a_j^i)^{1+\zeta}$.¹⁹ As in [Heathcote, Storesletten, and Violante \(2010\)](#), we assume that agents draw the disutility cost of education χ^i from a gender specific distribution such that $\log(\chi^i) \sim N(\mu_\chi^i, (\sigma_\chi^i)^2)$, $i = f, m$. By acquiring education, the agent upgrades her wage per unit of efficiency, $w_j^{i,e}$, from that of uneducated, $w_j^{i,0}$, to that of educated individuals, $w_j^{i,1}$, where $j = ss, g, ms$ is the sector where the agent decides to work. Since there are two education levels and three market sectors, the agent, depending on her skill vector, and taking as given the equilibrium (gender-specific) market wages per unit of efficiency in the three sectors and for each level of education, chooses the pair $(e, j) \in \{0, 1\} \times \{ss, g, ms\}$ in order to maximize her efficiency wage net of education costs.

The optimal choice by an agent of gender i , with ability $a^i = [a_{ss}^i, a_g^i, a_{ms}^i]$ and facing a vector of equilibrium market wages $w^i = [w_{ss}^i, w_g^i, w_{ms}^i]$ is then a pair $(e^*, j^*) =$

¹⁸The methodology to define substitutable services in the data is described in Section 4.

¹⁹This complementarity assumption is driven by the complementarity between skill levels and education attainment documented in recent work. See [Heckman, Stixrud, and Urzua \(2006\)](#) and [Findeisen and Sachs \(2015\)](#).

$[e(a^i, w^i, i), j(a^i, w^i, i)] \in \{0, 1\} \times \{ss, g, ms\}$ such that

$$(e^*, j^*) = \operatorname{argmax}_{(e,j)} \left[w_j^{i,e} (a_j^i)^{(1+e\zeta)} - e\chi^i \right] \quad (3)$$

Notice that conditional on $e^* = 0$, $(0, j^*) = \operatorname{argmax}_{(0,j)} [w_j^{i,0} (a_j^i)]$, so that the agent chooses to work in the sector j which, given her ability and the market wages per unit of efficiency, ensures the highest efficiency wage $w_j^{i,0} a_j^i$. By contrast, conditional on $e^* = 1$, $(1, j^*) = \operatorname{argmax}_{(1,j)} [w_j^{i,1} (a_j^i)^{(1+\zeta)} - \chi^i]$, the agent chooses to work in the sector which ensures the highest actual wage net of the education cost.

Note also that it can be that $\operatorname{argmax}_{(0,j)} [w_j^{i,0} a_j^i] \neq \operatorname{argmax}_{(1,j)} [w_j^{i,1} (a_j^i)^{(1+\zeta)} - \chi^i]$, so that the sector which ensures the maximum wage with education investment might be different from the sector which ensures the maximum wage without education. Put it differently, we allow for an interaction between human capital investment and structural change: on the one hand, investing in human capital might be convenient only if the agent switches to another sector; on the other hand, switching to another sector might be profitable only conditional on human capital investment.

4.3 Consumption and time allocation decisions

Before choosing the consumption and time allocations, each agent chooses the combination of education level e and in which sector j to work to maximize her wage net of education costs, $w_j^{i,e} (a_j^i)^{(1+e\zeta)} - e\chi^i$. This implies that this wage is taken as given in the maximization problem involving consumption and labor. We define the maximum efficiency wage net of education for an agent of type i as follows

$$W(a_{j^*}^i, w_{j^*}^{i,e^*}, e^*) = w_{j^*}^{i,e^*} (a_{j^*}^i)^{(1+e^*\zeta)} - e^*\chi^i \quad (4)$$

being e^* and j^* the level of education and the sector of work optimally chosen by an agent of type i .

Regarding the consumption and time allocation there are three kinds of decision units (i.e. households) in the model, $z = c, f, m : 1$) a household c , which is formed by a couple of a female and a male individual; 2) a single female f ; 3) a single male m . The utility function of a decision unit $z = c, f, m$ is

$$U^z = \left((\omega_{ms})^{1/\sigma} \left(\frac{C_{ms}^z}{\kappa^z} \right)^{\frac{\sigma-1}{\sigma}} + (\omega_g)^{1/\sigma} \left(\frac{C_g^z}{\kappa^z} \right)^{\frac{\sigma-1}{\sigma}} + (\omega_s)^{1/\sigma} \left(\frac{\tilde{C}_{ts}^z}{\kappa^z} \right)^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}, \quad (5)$$

$$\tilde{c}_{ts}^z = \left(\psi (c_{ss}^z)^{\frac{\gamma-1}{\gamma}} + (1 - \psi) (c_h^z)^{\frac{\gamma-1}{\gamma}} \right)^{\frac{\gamma}{\gamma-1}} + \bar{c} \quad (6)$$

where c_{ms}^z is consumption of *modern services*, c_g^z is consumption of manufacturing, \tilde{c}_{ts}^z represents *traditional services (ts)*, which is an aggregator of c_{ss}^z , consumption of substitutable services and c_h^z , which is consumption of home services. The parameter κ^z represents economies of scale for the couple.²⁰ Following the findings in [Moro, Moslehi, and Tanaka \(forthcoming\)](#) we assume that the income elasticity of traditional services is different from that of modern services, and introduce the negative non-homothetic term \bar{c} .

The first three types of consumption are purchased in the market, while home services are produced within the household. Each agent is endowed with 1 unit of time and each household devotes a fraction of this time to home production and the remaining time to market work. In the case of the couples, $z = c$, both male and female labor is used to produce home services. This is not so when the decision unit is a single women ($z = f$, no male labor is available) or when it is a single man ($z = m$, no female labor is available). For each type of household, home services are produced according to the following technology

$$Y_h^z = A_h L^z, \quad (7)$$

where

$$L^c = A_h \left[\varphi_h^c (l^f)^{\frac{\eta-1}{\eta}} + (1 - \varphi_h^c) (l^m)^{\frac{\eta-1}{\eta}} \right]^{\frac{\eta}{\eta-1}}, \quad (8)$$

$$L^f = A_h \left(\varphi_h^f \right)^{\frac{\eta}{\eta-1}} l^f, \quad (9)$$

$$L^m = A_h \left(\varphi_h^m \right)^{\frac{\eta}{\eta-1}} l^m, \quad (10)$$

The budget constraint changes across household types being

$$p_{ms} c_{ms}^z + p_g c_g^z + p_{ss} c_{ss}^z = E^z, \quad (11)$$

where

$$E^c = W \left(a_{j^*}^i, w_{j^*}^{i,e^*}, e^* \right) (1 - l^f) + W \left(a_{j^*}^i, w_{j^*}^{i,e^*}, e^* \right) (1 - l^m), \quad (12)$$

$$E^f = W \left(a_{j^*}^i, w_{j^*}^{i,e^*}, e^* \right) (1 - l^f), \quad (13)$$

$$E^m = W \left(a_{j^*}^i, w_{j^*}^{i,e^*}, e^* \right) (1 - l^m). \quad (14)$$

²⁰So in the calibrated model we will have $\kappa^f = \kappa^m = 1$ and $\kappa^c = 1.5$ from the scale equivalence computed by the OECD.

We highlight that when $z = c$ (when the decision unit is a couple) every female agent always works in the market in the sector with the highest $W(a_{j^*}^f, w_{j^*}^{f,e^*}, e^*)$ irrespective of her husband's choice as the households maximizes total utility.²¹

Each decision unit $z = c, f, m$ chooses the amount of consumption of each good c_j and the time devoted to home production by men and women l^m and l^f in order to maximize utility (5), subject to the service aggregator (6), the budget constraint (11) and the home production technology constraint (7).

From first order conditions we obtain the relative time of work at home of spouses, which, in an interior solution, is given by

$$\frac{l^f}{l^m} = \left(\frac{\varphi_h}{1 - \varphi_h} \frac{W(a_{j^*}^m, w_{j^*}^{m,e^*}, e^*)}{W(a_{j^*}^f, w_{j^*}^{f,e^*}, e^*)} \right)^\eta. \quad (15)$$

Thus, the time of work at home of a female agent increases with the wage and the ability of the male in the market (which can be boosted by education) and declines with the wage and the ability of herself in the market.

From utility maximization we can derive an implicit price for home services, which is the key dimension in which singles and married are different. For married, this is given by

$$p_h^c = \frac{1}{A_h} \left[\varphi_h^\eta \left[W(a_{j^*}^f, w_{j^*}^{f,e^*}, e^*) \right]^{1-\eta} + (1 - \varphi_h)^\eta \left[W(a_{j^*}^m, w_{j^*}^{m,e^*}, e^*) \right]^{1-\eta} \right]^{\frac{1}{1-\eta}}. \quad (16)$$

The price of home services is household specific, which is due to the fact that, the higher the efficiency wage of a member of the household, the higher the opportunity cost of working at home rather than in the market. Thus, the model predicts that households with higher abilities tend to work more in the market and less at home, compared with households with lower abilities.

The home price for a single individual is

$$p_h^i = \frac{W(a_{j^*}^i, w_{j^*}^{i,e^*}, e^*)}{A_h} (\varphi_h^i)^{-\frac{\eta}{\eta-1}}. \quad (17)$$

This implicit price is increasing in ability so that a single agent with higher ability works more in the market and less at home, compared with a single agent with lower abilities. By comparing (16) and (17) it is also possible to see that changes in market conditions (i.e. wages) have a different effect on the price of home production of married and singles, which translates, *ceteris paribus*, into a different decisions on how much to work at home and in

²¹A similar discussion can be made for a married men.

the market for the two types of households.

4.4 Firms and sectors

There is a representative firm in each market sector $j=ms, g, ss$. Each representative firm has the following production function

$$Y_j = A_j N_j, \quad (18)$$

where

$$N_j = \left[\phi_j \left(\varphi_j N_j^{f,1} + (1 - \varphi_j) N_j^{m,1} \right)^{\frac{\eta_s - 1}{\eta_s}} + (1 - \phi_j) \left(\varphi_j N_j^{f,0} + (1 - \varphi_j) N_j^{m,0} \right)^{\frac{\eta_s - 1}{\eta_s}} \right]^{\frac{\eta_s}{\eta_s - 1}}, \quad (19)$$

where $N_j^{i,e}$ is the aggregator of labor efficiency units of agents of gender $i = m, f$ and education level $e = 0, 1$ in sector j . Our production function follows [Heathcote, Storesletten, and Violante \(2010\)](#) in displaying 1) perfect substitutability across gender; 2) gender-biased technology (through the parameter φ_j) 3) imperfect substitutability across education levels ($\eta_s > 1$ being the elasticity of substitution between educated and non-educated workers); and 4) skilled-biased technology (through the parameter ϕ_j).

The representative firm operating in sector j maximizes profits

$$\pi_j = p_j Y_j - w_j^{f,1} N_j^{f,1} - w_j^{m,1} N_j^{m,1} - w_j^{f,0} N_j^{f,0} - w_j^{m,0} N_j^{m,0} \quad (20)$$

subject to (18) and (19).

First order conditions imply

$$\frac{\phi_j \left(\varphi_j N_j^{f,1} + (1 - \varphi_j) N_j^{m,1} \right)^{-\frac{1}{\eta_s}}}{(1 - \phi_j) \left(\varphi_j N_j^{f,0} + (1 - \varphi_j) N_j^{m,0} \right)^{-\frac{1}{\eta_s}}} = \frac{w_j^{m,1}}{w_j^{m,0}} \quad (21)$$

$$\frac{\varphi_j}{1 - \varphi_j} = \frac{w_j^{f,e}}{w_j^{m,e}} \quad (22)$$

Equation (21) shows that, other conditions equal, skill-biased technological change due to a time varying ϕ_j , raises the skill premium. Equation (22) shows that gender-bias technological change, in the form of growing φ_j , directly affects the wage ratio between males and females. Note, however, that the initial value of φ_j can be different across sectors, so that the aggregate gender wage gap is determined endogenously and changes over time, even without gender-

biased technological change.

4.5 Definition of equilibrium

The equilibrium is defined as a set of prices $\{p_{ss}, p_g, p_{ms}\}$, a set of wages per unit of efficiency $\{w_{ss}^{f,1}, w_g^{f,1}, w_{ms}^{f,1}, w_{ss}^{m,1}, w_g^{m,1}, w_{ms}^{m,1}, w_{ss}^{f,0}, w_g^{f,0}, w_{ms}^{f,0}, w_{ss}^{m,0}, w_g^{m,0}, w_{ms}^{m,0}\}$, a set of choices for each agent $(e^*, j^*) = [e(a^i, w^i, i), j(a^i, w^i, i)]$ and a set of allocations for each household $\{c_{ss}^z, c_g^z, c_{ms}^z, l^{fz}, l^{mz}\}$ such that:

1. Given wages and prices, the choice (e^*, j^*) maximizes wages net of education costs for agent i by solving (3);
2. Given wages, prices, and (e^*, j^*) of each household member, the allocation $\{c_{ss}^z, c_g^z, c_{ms}^z, l^{fz}, l^{mz}\}$ maximizes utility (5) of the household subject to the budget constraint (11);
3. Given wages and prices, each representative firm in sectors ss , g , and ms maximizes profits (20);
4. Labor markets in sectors ss , g , and ms clear;
5. Goods markets in sectors ss , g , and ms clear.

5 Calibration

We calibrate the model to two equilibria to replicate a series of targets of the U.S. economy in the years 1980 and 2008. We allow for the following exogenous differences between the two equilibria: i) the level of labor productivity of market sectors and the home sector; ii) the level of skill-biased technology; iii) level of gender-biased technology; and iv) the marriage market.

A number of parameters, $\{\sigma, \gamma, \eta, \eta_s\}$, are set from previous study. Following [Ngai and Pissarides \(2008\)](#) we set $\sigma = 0.3$ and $\gamma = 2.3$. η is estimated in [Knowles \(2013\)](#) to 3, while the elasticity of substitution between educated and uneducated workers is taken from [Heathcote, Storesletten, and Violante \(2010\)](#) and set to $\eta_s = 1.43$.

Ability is assumed to be uniformly distributed, with $a_j \in [\underline{a}_j, \bar{a}_j]$ and with men and women drawing from the same ability distribution by sector when born. Spouses' abilities are correlated with correlation coefficient ρ_j . These correlations are computed using data on U.S. wages. To compute the correlation between husband and wife wages, we first compute female wages by sector correcting for selection bias using the Heckman correction, and second correlate wages of husbands and wives that work in the same sector. The correlation,

averaging from 1978 to 2010, is 0.32 for manufacturing, 0.25 for low-skilled services and 0.26 for high skilled services. These values provide our correlation of skills measure.

Initial productivities by sector, including the home sector, are normalized to one, $A_{j,1980} = 1$ and $A_{h,1980} = 1$. The lower bound of ability in the substitutable service sector is $\underline{a}_{ss} = 0.5$. For education, we follow Heathcote, Storesletten, and Violante (2010): individuals draw a disutility cost from the distribution $\log(\chi^i) \sim N\left(\mu_\chi^i, (\sigma_\chi^i)^2\right)$, where $i = f, m$. As in Heathcote, Storesletten, and Violante (2010), mean education costs for both men and women are matched to the share of educated men and women in 1980, respectively. The remaining 26 parameters: (1) ability and return to education $\{\bar{a}_{ss}, \underline{a}_{ms}, \bar{a}_{ms}, \underline{a}_g, \bar{a}_g, \zeta\}$, (2) productivity (market and home) $\{\{\varphi_j, \phi_j\}_{j=ms,g,ss}, \varphi_h, \varphi_h^f, \varphi_h^m\}$, (3) preferences $\{\omega_{ms}, \omega_g, \psi, \bar{c}, \sigma_\chi^f, \sigma_\chi^m\}$, and (4) time trends for sector productivity, gender-biased and skill-biased technological change $\{\{\gamma_j\}_{j=ms,g,ss}, \gamma_\varphi, \gamma_\phi\}$ are calibrated to match a number of moments.²²

In addition, OECD economies of scales assume the first adult in consumption accounts for 1.0, but the second adult accounts for a factor of 0.5 in a multi-person household. Therefore, it is assumed that $\kappa = 1.5$ for married households and $\kappa = 1$ for single households. Table 2 lists the parameter values used in the simulation.

While the calibration procedure matches all 26 parameters to 26 moments concurrently, by minimizing the distance between data targets and model moments, some targets are more informative for certain parameters than others. Below we outline the general strategy.

Ability parameters, $\{\bar{a}_{ss}, \underline{a}_{ms}, \bar{a}_{ms}, \underline{a}_g, \bar{a}_g\}$ (5 targets): male modern services (industry) to substitutable services wage premiums and the standard deviation of log male wages of full-time full-year workers from the CPS in 1980 in the three market sectors. Relative weights in consumption, $\{\omega_{ms}, \omega_g\}$ (2 targets): share of hours in modern services and industry in 1980. Home production $\{\varphi_h^f, \varphi_h^m, \varphi_h^c, \psi\}$ (4 targets): married male market hours, single male market hours, married female market hours and single female market hours. Gender gaps in the market in 1980, $\{\varphi_j\}_{j=ms,g,ss}$ (3 targets): aggregate gender wage gap, female to male industry hours gap, female substitutable services to modern services hours gap. Education determinants, $\{\zeta, \sigma_\chi^f, \sigma_\chi^m, \{\phi_{1980}\}_{j=ms,g,ss}\}$ (6 targets): the male and the female college wage premium in 1980, the relative hours of uneducated (LTC=less than college) to educated in manufacturing and substitutable services and the fraction of educated women and educated men in 2008. Non-homotheticity in consumption $\{\bar{c}\}$ and sector growth rates $\{\gamma_j\}_{j=hs,ls,g}$ (4 targets): the changes over time of hours in industry, hours in modern services, industry to substitutable services wage, and modern services to substitutable services wage. Skill-biased technological change, $\{\gamma_\phi\}$ (1 target): the growth in the male college wage premium between 1980 and 2008. Gender-biased technological change, $\{\gamma_\varphi\}$ (1 target): the growth in the

²²Note that $\omega_{ss} = 1 - \omega_{ms} - \omega_g$.

Table 2: *Model Parameters*

Parameter	Type	Value
$\{\underline{a}_{ss}, \bar{a}_{ss}\}$	Substitutable services ability	{0.50, 3.37}
$\{\underline{a}_{ms}, \bar{a}_{ms}\}$	Modern services ability	{1.05, 4.87}
$\{\underline{a}_g, \bar{a}_g\}$	Manufacturing ability	{0.77, 4.40}
σ	Substitutability between broad consumption categories	0.3
γ	Substitutability between home and market services	2.3
η	Gender substitutability at home (married only)	3
η_s	Substitutability educated/uneducated in production	1.43
ω_{ms}	Consumption market weight modern services	0.43
ω_g	Consumption market weight manufacturing	0.33
ψ	Substitutable services weight	0.25
φ_h^c	Home female-labor weight	0.54
φ_h^f	Single female home labor weight	0.41
φ_h^m	Single male home labor weight	0.50
φ_{ms}	Female-labor weight in modern services	0.34
φ_g	Female-labor weight in manufacturing	0.31
φ_{ss}	Female labor weight in substitutable services	0.36
ζ	Schooling factor	0.21
μ_χ^f	Mean of the cost of schooling female	0.64
μ_χ^m	Mean of the cost of schooling male	1.26
σ_χ^f	Variance of the cost of schooling female	0.94
σ_χ^m	Variance of the cost of schooling male	1.05
$\phi_{ms,1980}$	Educated workers labor weight in modern services	0.34
$\phi_{g,1980}$	Educated workers labor weight in manufacturing	0.32
$\phi_{ss,1980}$	Educated workers labor weight in substitutable services	0.38
\bar{c}	Non-homothetic consumption in low skilled service aggregator	-0.09
γ_h	Annual growth in A_h	0.001
γ_{hs}	Annual growth in A_{hs}	0.004
γ_{ls}	Annual growth in A_{ls}	0.017
γ_g	Annual growth in A_g	0.034
γ_ϕ	Skill-biased technological change (annual growth rate in ϕ_j)	0.013
γ_φ	Gender-biased technological change (annual growth rate in φ_j)	0.005

aggregate gender wage gap between 1980 and 2008. All targets are computed using the 1980 Census and the 2008 American Community Survey unless noted. Home labor productivity growth γ_h is measured to be 0.1 percent in [Bridgman \(2016\)](#) and -0.4 percent in [Moro, Moslehi, and Tanaka \(forthcoming\)](#) for the 1978-2010 period. We choose the value of 0.001 in our calibration.

To define services sectors that are substitutable to home we use the procedure in [Moro, Moslehi, and Tanaka \(forthcoming\)](#). First, from time use surveys we select home activities that are considered home production. We follow [Bridgman, Dugan, Lal, Osborne, and Villones \(2012\)](#) and [Landefeld and McCulla \(2000\)](#) and define seven broad categories: “*cooking*”, “*house work*”, “*odd jobs*”, “*gardening*”, “*shopping*”, “*child care*”, and “*travel*”, where the last one is intended as travel related to the other six categories. We then use the 1990 CENSUS classification (3 digits) to select industries producing an output that is “close” in nature to the output produced by the seven home activities. Selected industries are: Bus service and urban transit; Taxicab service; Retail bakeries; Eating and drinking places; Liquor stores; Private households; Laundry, cleaning, and garment services; Beauty shops; Barber shops; Dressmaking shops; Miscellaneous personal services; Nursing and personal care facilities; Child day care services; Family child care homes; Residential care facilities, without nursing.

Changes in the demographic structures are summarized in table 4. It is worth emphasizing the following differences between the two equilibria: 1) the share of educated individuals grows for both men and women but relatively faster for women; 2) among the latter, the share of educated individuals increases faster for married rather than for single women; 3) the marriage rate decreases.

Finally, note that this is the first paper that compares polarization graphs in the data with the outcome of a general equilibrium model. Thus, one challenge is how to draw polarization graphs in the model that are comparable with those in the data. We proceed as follows. First, within each market sector in the model we create equally sized bins of workers with similar ability along the sector skill distribution using the 1980 equilibrium. We do this because, for instance, the ability level of a worker in manufacturing cannot directly compared with the ability level of a worker in high-skilled services. Next, we compute the average wage in each bin. Then, we rank all bins from the three market sectors into a unique classification by using the average wage in each bin in the 1980 equilibrium. This ranking is then kept for the 2008 equilibrium to construct polarization graphs. We apply the same method to the data. That is, within each market sector in 1980 we create bins of workers with similar wage, and compute the average wage in each bin. Then we rank bins from the three sectors into a unique classification using the average wage in each bin. Then, by keeping the same

Table 3: *Model Targets*

Type	Data	Model
1980 - ability ($\{\underline{a}_j, \bar{a}_j\}_{j=ms,g,ss}$)		
Male industry to substitutable services wage	1.33	1.41
Male modern services to substitutable services services wage	1.42	1.48
Standard deviation of industry log male wages	0.27	0.31
Standard deviation of substitutable services log male wages	0.28	0.28
Standard deviation of modern services log male wages	0.29	0.34
1980 - education cost ($\{\mu_\chi^i\}_{i=m,f}$)		
Fraction of educated men in 1980	0.16	0.16
Fraction of educated women in 1980	0.13	0.13
1980 - consumption ($\{\omega_j\}_{j=ms,g,ss}$)		
Share of hours in industry	0.35	0.35
Share of hours in modern services	0.59	0.57
1980 - home production ($\psi, \varphi_h^c, \varphi_h^f, \varphi_h^m$)		
Married male market hours	0.78	0.95
Single male market hours	0.61	0.51
Married female market hours	0.34	0.36
Single female market hours	0.49	0.48
1980 - Gender weights in the market ($\{\varphi_j\}_{j=ms,g,ss}$)		
Aggregate Gender Wage Gap	0.59	0.46
Female to male industry hours gap	0.32	0.32
Female subst. serv. to modern serv. hours gap	0.17	0.14
1980 - education ability returns ($\zeta, \{\phi_{j,1980}\}_{j=ms,g,ss}$)		
Female college wage premium	1.57	1.62
Male college wage premium	1.54	1.65
Share of LTC Hours in manufacturing	0.88	0.84
Share LTC Hours in substitutable services	0.92	0.79
Variance of education cost ($\sigma_\chi^m, \sigma_\chi^f$)		
Fraction of educated men in 2008	0.28	0.27
Fraction of educated women in 2008	0.27	0.27
1980-2008 - non-homotheticity and productivity ($\bar{c}, \{\gamma_j\}_{j=ms,g,ss}$)		
Hours in industry (change over time)	0.67	0.72
Hours in modern services (change over time)	1.24	1.28
Industry to substitutable services wage (change over time)	0.99	0.94
Modern serv. to substitutable serv. wage (change over time)	1.19	1.10
1980-2008 - skill-biased and gender-biased technological change ($\{\gamma_j\}_{j=\phi,\varphi}$)		
Gender wage gap (change over time)	1.25	1.28
Relative college wages (change over time)	1.28	1.33

Table 4: *Demographic Changes*

	1980	2008
Singles		
Male	0.23	0.31
Female	0.26	0.30
Share Educated		
Single Men	0.16	0.19
Single Women	0.13	0.23
Married Men	0.20	0.34
Married Women	0.13	0.32
Couple Types		
Educated Couples	0.09	0.22
Educated Husband Only	0.11	0.12
Educated Wife Only	0.04	0.10
Uneducated Couples	0.76	0.56

ranking in 2008 we construct employment polarization graphs.

Note that in the data, certain occupations are in all three sectors (e.g., secretaries), but others are likely just in one of the three (miners). So in our method we have four occupations and wage rates in 1980. Instead, the original method in [Acemoglu and Autor \(2011\)](#) computes an average wage for secretaries in the U.S. economy in 1980 and one for miners. So instead of four occupations and wage rates in 1980, they have two. Besides that the two methods are identical, that is, we rank occupations by their wages in 1980 from 0 to 100. As [Appendix A](#) shows, the differences between the two methods in the data are very minor, and mostly at the right tail of the overall distribution. The reason for this difference is that in our methodology some occupation groups are more homogeneous.

6 Results

[Figure 5](#) presents the comparison between polarization in the data and the respective polarization graphs generated by the model. The top-right panel of [Figure 5](#) shows that model performs well in replicating the main features of the data, in particular the standard pattern of overall polarization. Employment shares increase both at the bottom and the top of the skill distribution, while they decline in the middle of the distribution. The model also generates similar patterns with respect to the data when decomposed by gender. Women generate an increase in employment shares at the bottom and the top of the skill distribution, with a decrease in the middle. The behavior of men is also broadly consistent with the data, with a decrease of employment shares along most of the skill distribution, except for an increase

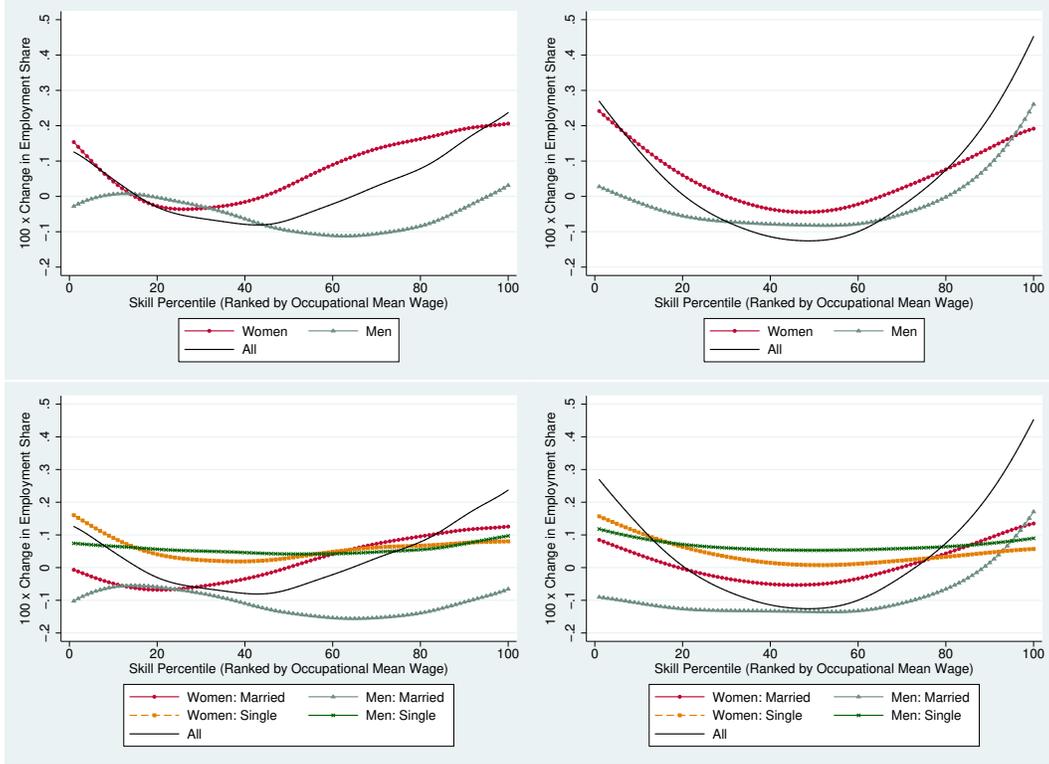


Figure 5: Job polarization in the data (left) and in the model (right). First row: gender; second row: marital status and gender.

at the top. However, such increase is too pronounced with respect to the data.²³

The second row of Figure 5 compares the performance of the model conditional on marital status. Similarly to the data, singles display a flatter behavior across the skill distribution with respect to married, and increase their employment shares along the whole skill distribution. This is due to the fact that couples can reallocate working hours within the family, while single individuals cannot. As in the data, married women are also largely responsible for the increase at the top of the distribution, while single women contribute to a large extent to the increase at the bottom. The intuition for this pattern can be found in the different fraction of married and single women that acquire education between the two equilibria. In the first equilibrium, the share of educated individuals single and married women is equal at 0.13. In contrast, in the second equilibrium, the share of educated individuals among married women increases by a factor of 2.53, while that of single women only by a factor

²³By running OLS regressions in which the dependent variable is the change in the employment share of percentile i in the data, and the independent variable is the corresponding change in the employment share of percentile i in the model, we find that the estimated coefficient is positive and significant, suggesting a good predicting power of the model. The coefficient of the regression is 0.56 when considering the overall population, 0.56 when considering only women, and 0.28 when considering only men.

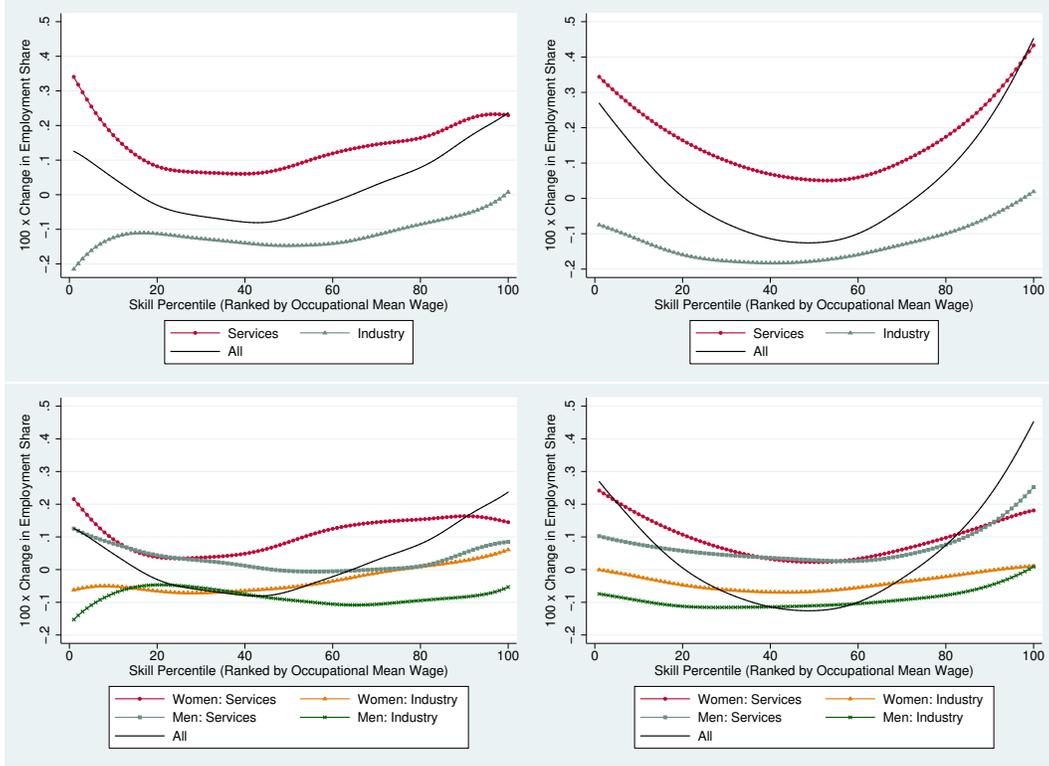


Figure 6: Job polarization in the data (left) and in the model (right). First row: sectors; second row: sectors and gender.

of 1.80. Hence, the former are more likely to satisfy the increase of educated labor demand while the latter are more likely to absorb the demand of uneducated labor.

We also report polarization across sectors in the four panels of Figure 6. The outcome of the model is again similar to the data. The first row of Figure 6 shows how the model reproduces job-polarization in services and the flat behavior of manufacturing (except at the top of the distribution) observed in the data. The second row of Figure 6 suggests that the model does well even when decomposing sectoral polarization by gender. In particular, it replicates the upward twists for women in services at the top and at the bottom of the distribution and the relative homogeneity and “flatness” of the negative change in men hours in manufacturing along the whole skill distribution. The latter behavior of men, when coupled with the strong female polarization, is key in explaining the downward twist in the middle of the distribution of the overall economy. In fact, this result suggests that the decline at the bottom of the overall distribution is the result of services occupation increasing at the middle less than in the rest of the distribution, and manufacturing occupations declining similarly along the whole distribution.

We conclude this section by running counterfactual exercises that help assessing the role

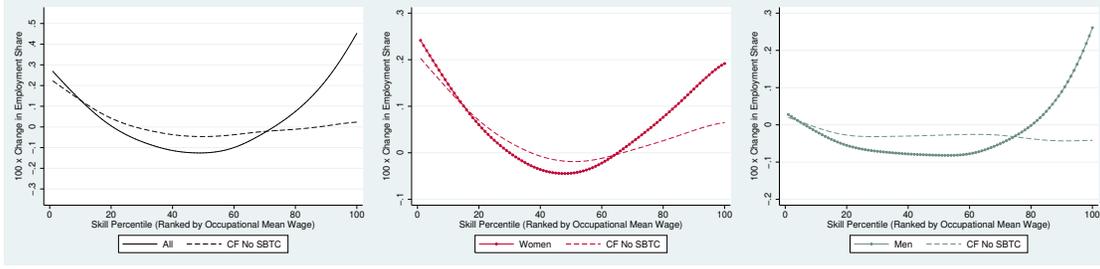


Figure 7: Counterfactual: No skill-biased technological change

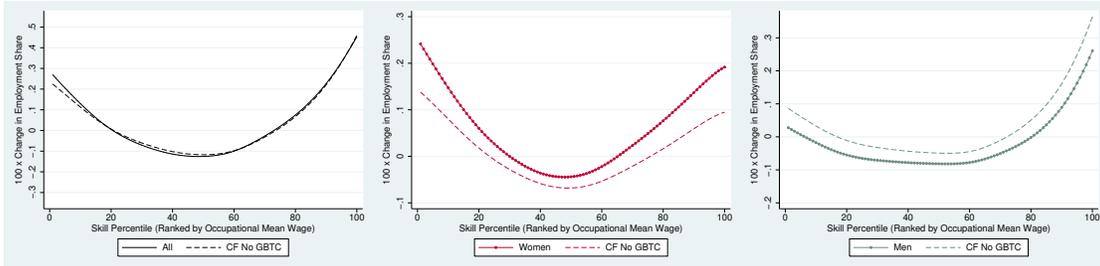


Figure 8: Counterfactual: No gender-biased technological change

of exogenous factors in shaping the results. In Figure 7 we set skill-biased technological change to zero. As the left-panel shows, removing this type of technological change makes employment polarization disappear. Changes in employment shares are roughly flat from percentile forty to the top of the skill distribution. For men, the effect is to make changes along the distribution entirely homogeneous (and roughly zero). For women, both the increase at the top and at the bottom of the distribution are smaller. Note that this can be interpreted as the existence of consumption spillovers from wealthy high-skilled women who increase the amount of time worked in the market and, as a consequence, demand services that are substitutable to home production, thus fostering the demand for low-skilled women.²⁴

Figure 8 displays the counterfactual in which gender-biased technological change is set to zero. The main effect is to shift up the graph for men and down the graph for women. The effect on overall polarization is negligible. This suggests that the gender wage gap channel increases market hours of women in a homogeneous way along the distribution and does not have a first order effect of the shape of employment polarization. It is, instead, quantitatively relevant for determining the position of the curves of the two gender.

²⁴We will return on the existence of consumption spillovers due to skill-biased technological change in section 8, when we analyze each one of the three decades in the 1980-2008 period.

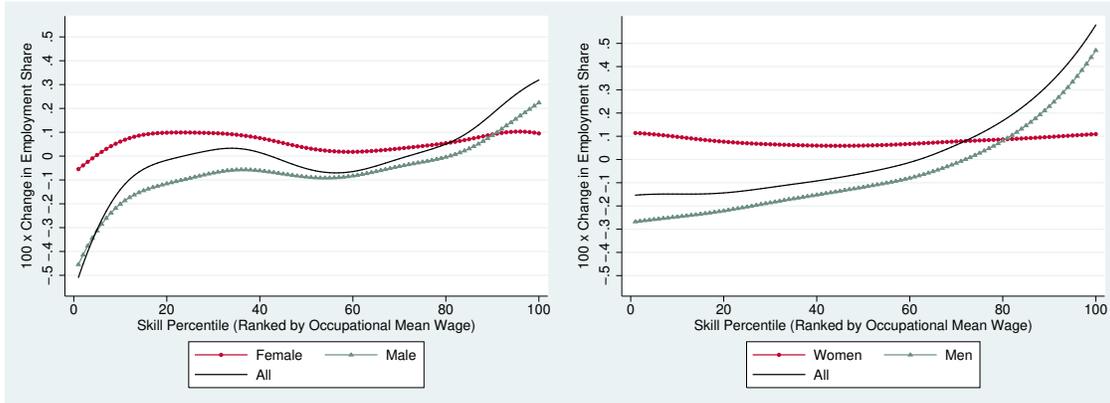


Figure 9: Job polarization in the data (left) and in the model (right) during the period 1960-1980.

7 Predicting the Pre-Polarization Era

The results in the calibrated model presented in Section 6 are driven by the exogenous factors evolving between 1980 and 2008. A natural out of sample test of the model is to study the behavior of employment shares when the trends in exogenous factors are those of the 1960-1980 period. If the calibrated model performs well outside the calibration period, then we can argue that the exogenous factors in our model are the key drivers of changes in employment shares since 1960.

As skill-biased and gender-biased technological change cannot be directly measured from the data, we rely on the results in [Heathcote, Storesletten, and Violante \(2010\)](#), who find, for the 1967-1980 period, an average growth rate of around zero for both types of technological change.²⁵ Thus we set zero growth for the 1960-1980 period in our experiment. For home labor productivity we follow [Bridgman \(2016\)](#), who measures an average growth rate of 2.5% for the pre-1980 period. Finally, we assume that labor productivity in the three market sectors displays the same growth rate as in the 1980-2008 period.²⁶

Figure 9 presents the comparison between employment polarization in the data and the corresponding pattern generated by the model for the period 1960-1980.²⁷ As the black

²⁵See their figure 3. Measures in [Heathcote, Storesletten, and Violante \(2010\)](#) are appropriate in our setting because we employ the same production functions. Although they have a unique production function at the aggregate level, while we have various sectors, skill-biased and gender-biased technological change are common across sectors in our model.

²⁶Note that in section 4 the growth rates of labor productivity are calibrated together with the rest of parameters. This is because, due to the presence of gender-biased and skill-biased technological change we cannot measure labor productivity with a growth accounting exercise. For this reason we assume that labor productivity growth is constant over time when projecting the model to 1960.

²⁷Computing employment polarization for the 1960-1980 period requires to deal with occupations that are not present in both years. In Appendix B we discuss different methodologies to address this issue. The main

line in the left panel shows, overall polarization is not present, and changes in employment shares are negative below the seventieth percentile and positive above. This trend is driven by men, who display a monotone behavior that is similar to the overall pattern. Women instead, display changes in employment shares which are similar along the whole distribution, ranging between 0 and 0.1.

To understand the role of exogenous factors in shaping the difference between the 1960-1980 and the 1980-2008 period, we now use the model to predict the change of employment shares that would have occurred between 1960 and 1980 if exogenous trends had been *those of the 1980-2008 period*. Results are reported in Figure 10. The dashed line in the left panel shows that the model produces employment-polarization, although this is less pronounced than in the benchmark case of section 5. The change in overall polarization is due to women, who display a reduction of employment shares both at the top and at the bottom of the skill-distribution, while employment shares of men display a change similar to the benchmark case.

To assess which channels are driving the different behavior of women in the 1960-1980 period in the two exercises, we run some additional counterfactual exercises. Figure 11 reports the change that occurs in the red dashed line in Figure 10 when we set, sequentially: i) skill-biased technological change to zero; ii) gender-biased technological change to zero and iii) home labor productivity growth at 2.5% per year. The left panel shows that by removing skill-biased technological change polarization is reduced to a large extent. Removing gender-biased technical change moves the line down, as in the benchmark case. Finally, home productivity moves the line up and removes entirely employment polarization for women. Thus, the counterfactuals suggest that the absence of skill-biased and gender-biased technological change in the 1960-1980 period are the key factors driving the absence of employment polarization. As in the results of the previous section, the former is the main driver of the shape of employment changes along the skill distribution, while the latter is the main determinant of the position of the curve.

8 Predicting decades in the 1980-2008 period

[Acemoglu and Autor \(2011\)](#) show that the shape of overall polarization between 1980 and 2008 results from the aggregation of a different behavior of changes in employment shares in the three decades. This is Fact 6 in section 3 and we report it for convenience in the left panel of Figure 12 . While some convexity in the shape is present for the three lines, there is a clockwise tilting behavior across decades. During the 1980s the change in employment

results in this section are maintained regardless of the methodology.

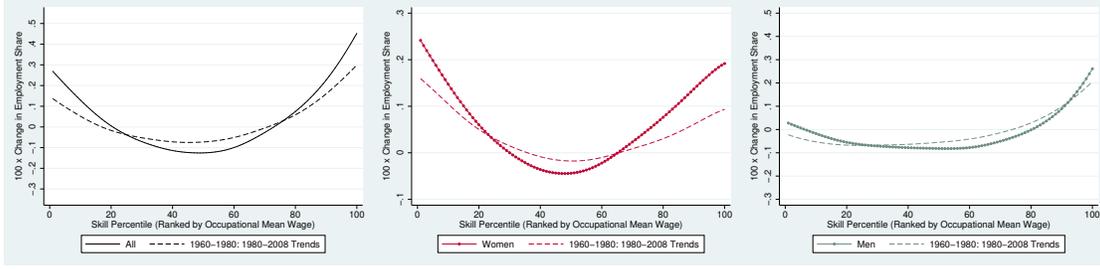


Figure 10: Job-polarization in the model during the period 1960-1980 when using the 1980-2008 exogenous trends.

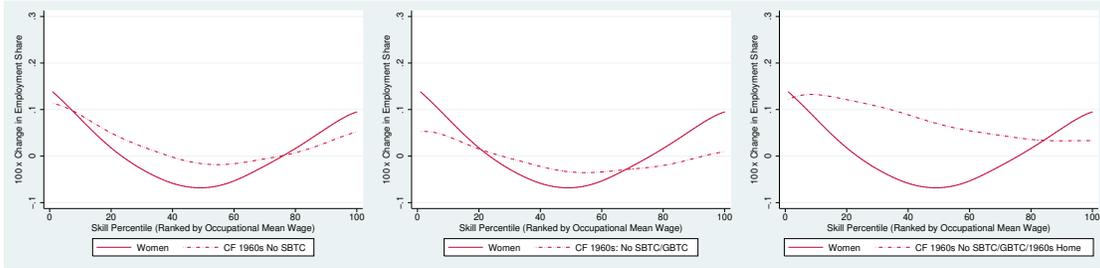


Figure 11: Job-polarization in the model during the period 1960-1980 when using the 1980-2008 exogenous trends. Counterfactuals for Women.

shares is increasing along the skill distribution. During the 1990s the graph displays a U-shape, while during the 2000s the large change in employment shares occurs at the bottom of the distribution.

In this section we test the performance of the model to reproduce the observed changes across decades. To do this, we feed the model with decade specific measures of skill-biased technological change and gender-biased technological change. The right panel of Figure 12 reports the behavior of the model. As in the data, the model can reproduce a tilt in the three lines, with (almost) no increase at the bottom in the 1980s, an increase at the bottom and at the top in the 1990s and an increase only at the bottom in the 2000s.

The rationale behind the good performance of the model is the changing effect over time of skill-biased technological change on employment shares, especially of women. To understand this time-varying impact, note first that skill-biased technological change has one direct effect and two indirect effects on employment shares. The former is the typical effect of skill-biased technological change which implies: an increase in the wage of educated individuals, in the number of educated individuals and in the amount of hours of the high-skilled in production. The latter effects are (i) a consumption spillover from the skilled (who work less at home) to the unskilled individuals due to a rise in the demand for substitutable market services and (ii) an increase in the labor demand of uneducated individuals when

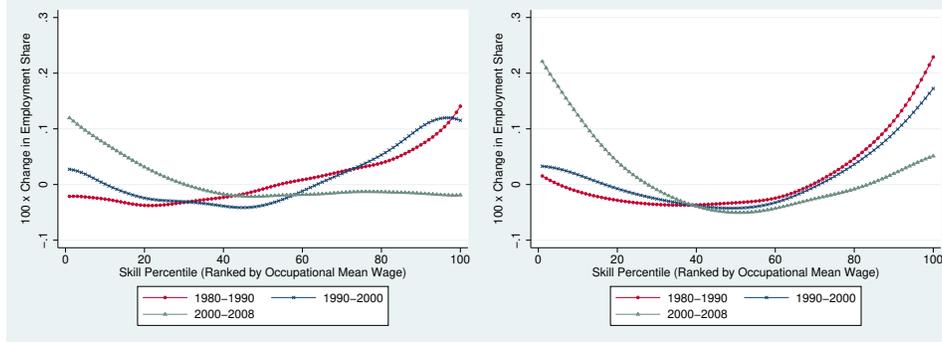


Figure 12: Job-Polarization by decade, 1980-2008. Data (Left) and Model (Right).

the hours of educated individuals increase in production (q-complementarity in production between educated and uneducated workers). In the model, the direct effect dominates in the first decade, while the indirect effects dominate in the last decade. In the middle decade the three effects are relevant.

To provide a narrative of the time-varying effect of skill-biased technological change we report, in Figure 13, counterfactual exercises for the three decades. In the 1980-1990 period skill-biased technological change mainly affects men, as displayed in the left panels of Figure 13. By removing skill-biased technological change the graph for men becomes flat, while the effect on women is less substantial, with the U-shape maintained also in the counterfactual. Thus, in the first decade, the direct effect appears as the one quantitatively relevant. In the 1990s skill-biased technological change becomes the main driver of employment polarization also for women. Average hours in the market for women do not increase during this period, but women move extensively along the skill distribution. The middle panel of the second row of Figure 13 shows that removing skill-biased technological change reduces substantially the increase of employment shares at the top of the skill-distribution while increasing the bottom. For men the effect is similar to the previous decade. Thus, during the 1990s both the direct and the indirect effect appear as quantitatively relevant. Finally, removing skill-biased technological change during the 2000s makes the change in employment shares of women (and men) entirely flat, removing the large increase at the bottom of the distribution observed in the data and in the benchmark case for the decades. This suggests that in this decade only the indirect effects have quantitative relevance. These results rationalize the emergence of consumption spillovers during the polarization era found in [Mazzolari and Ragusa \(2013\)](#) and [Hazan and Zoabi \(2015\)](#).

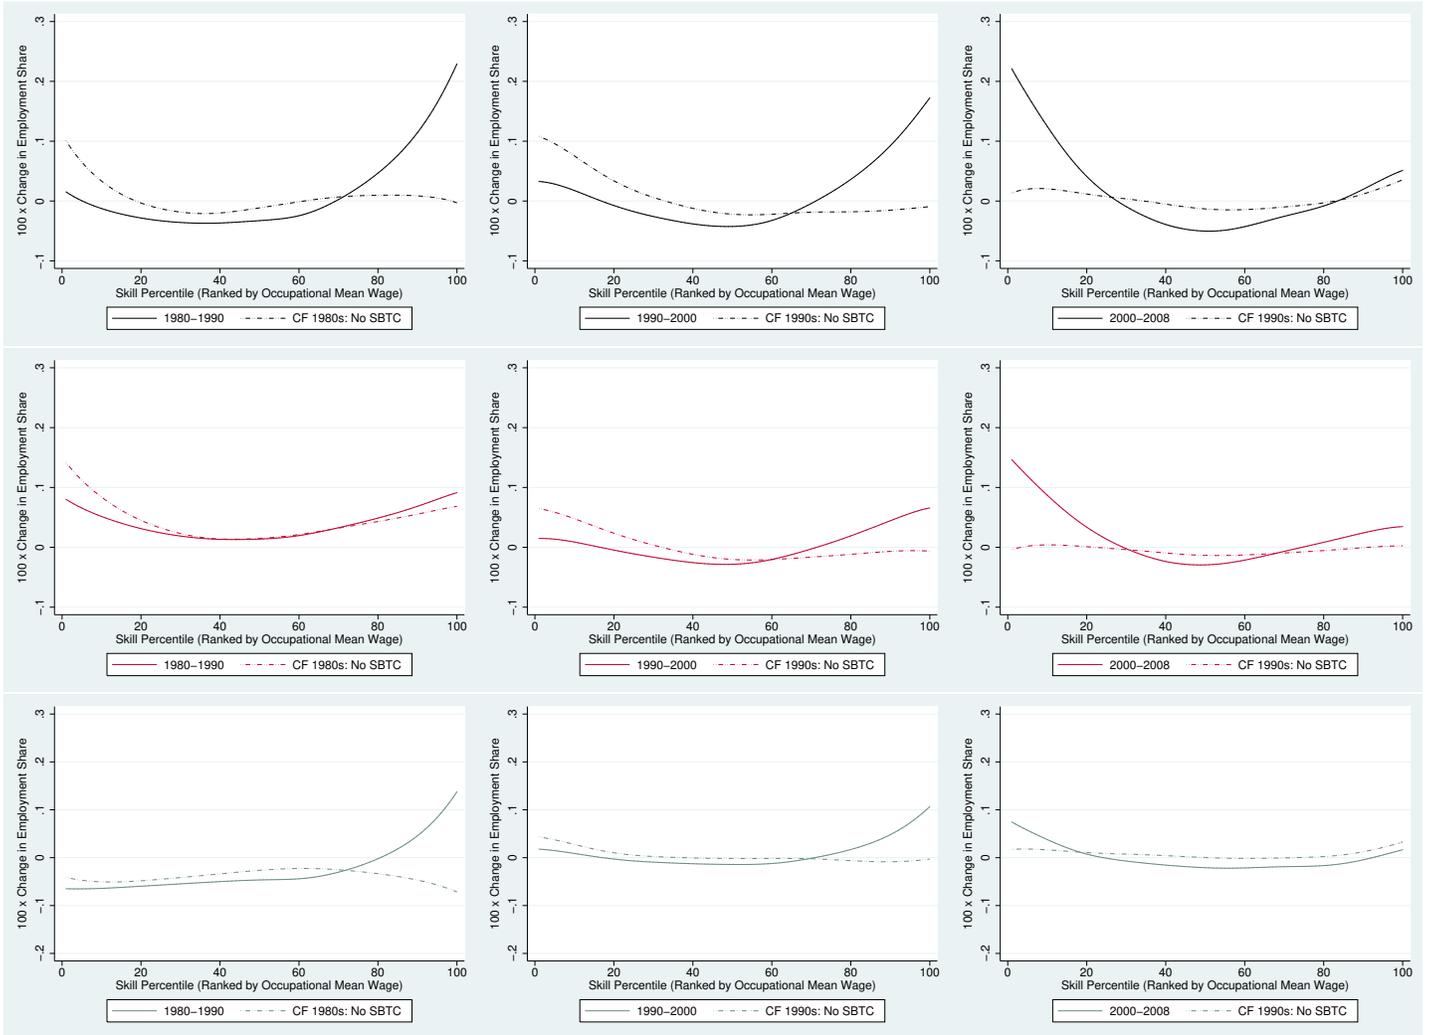


Figure 13: Job-Polarization by decade, 1980-2008. Counterfactual of skill-biased technological change.

9 Conclusions

In this paper we study the role of gender in generating the phenomenon labeled employment polarization. We document that the emergence of employment polarization since the 1980 is largely a female phenomenon due to women increasing market hours of work asymmetrically along the skill distribution. This observation motivated us to study the optimal response of different demographic groups when skill-biased technological change occurs and home production is an option for the agents. To do this, we constructed a multi-sector general equilibrium model with an education and occupational choice. The model shows that by taking into account the endogenous response of heterogeneous individuals to technological changes, it is possible to account for overall, gender and marital status specific, and sectoral

job-polarization facts. In addition, the model helps to rationalize the absence of employment polarization before 1980 and the changing behavior of employment shares in the various decades during the polarization era.

The model suggests that there are two main drivers for the gender differences in job-polarization patterns. First, a general increase in working opportunities for women, homogeneous along the skill distribution (due to the rise of the service economy and to gender-biased technological change). This driver accounts for the opposite sign of employment changes between the two gender along the whole skill distribution: negative for men, positive for women. Second, an increase in working opportunities for educated workers (due to skilled-biased technological change). This driver has a key role in generating the U-shape in the change of overall and female employment shares along the skill distribution. By fostering an increase in the working time of skilled women (mainly married) it accounts for most of the upward twist at the top of the skill distribution. Also, by favoring a reduction in home production, it leads to an increase in the labor demand for substitutable market services, thereby accounting for most of the downward twist at the bottom of the skill distribution. Our results suggest that any policy aimed at affecting the overall pattern of employment polarization should consider the effect on the various demographic groups that are contributing to shape this phenomenon.

References

- ACEMOGLU, D., AND D. AUTOR (2011): “Chapter 12 - Skills, Tasks and Technologies: Implications for Employment and Earnings,” vol. 4, Part B of *Handbook of Labor Economics*, pp. 1043 – 1171. Elsevier.
- (2012): “What does human capital do? A review of Goldin and Katz’s The race between education and technology,” *Journal of Economic Literature*, 50(2), 426–463.
- AUTOR, D. H., AND D. DORN (2013): “The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market,” *American Economic Review*, 103.
- 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.
- BARANY, Z., AND C. SIEGEL (2015): “Job Polarization and Structural Change,” Discussion paper.
- BRIDGMAN, B. (2016): “Home productivity,” *Journal of Economic Dynamics and Control*, 71, 60 – 76.
- BRIDGMAN, B., A. DUGAN, M. LAL, M. OSBORNE, AND S. VILLONES (2012): “Accounting for household production in the national accounts, 1965–2010,” *Survey of Current Business*, 92(5), 23–36.
- BUERA, F. J., J. P. KABOSKI, AND R. ROGERSON (2015): “Skill biased structural change,” Discussion paper, National Bureau of Economic Research.
- CARVALHO, V., AND X. GABAIX (2013): “The great diversification and its undoing,” *The American Economic Review*, 103(5), 1697–1727.
- CORTES, G. M., N. JAIMOVICH, AND H. E. SIU (2016): “Disappearing Routine Jobs: Who, How, and Why?,” Working Paper 22918, National Bureau of Economic Research.
- CORTES, P., AND J. TESSADA (2011): “Low-skilled immigration and the labor supply of highly skilled women,” *American Economic Journal: Applied Economics*, 3(3), 88–123.
- DUARTE, M., AND D. RESTUCCIA (2010): “The Role of the Structural Transformation in Aggregate Productivity,” *The Quarterly Journal of Economics*, 125(1), 129–173.
- DUERNECKER, G., AND B. HERRENDORF (2016): “Structural Transformation of Occupation Employment,” *Working Paper*.

- FINDEISEN, S., AND D. SACHS (2015): “Designing efficient college and tax policies,” .
- HAZAN, M., AND H. ZOABI (2015): “Do highly educated women choose smaller families?,” *The Economic Journal*, 125(587), 1191–1226.
- 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.
- HECKMAN, J. J., J. STIXRUD, AND S. URZUA (2006): “The effects of cognitive and noncognitive abilities on labor market outcomes and social behavior,” Discussion paper, National Bureau of Economic Research.
- HERRENDORF, B., AND A. VALENTINYI (2012): “Which Sectors Make Poor Countries So Unproductive?,” *Journal of the European Economic Association*, 10(2), 323–341.
- KNOWLES, J. A. (2013): “Why are Married Men Working So Much? An Aggregate Analysis of Intra-Household Bargaining and Labour Supply,” *Review of Economic Studies*, 80(3), 1055–1085.
- LANDEFELD, J. S., AND S. H. MCCULLA (2000): “Accounting for Nonmarket Household Production within a National Accounts Framework,” *Review of Income and Wealth*, 46(3), 289–307.
- MANNING, A. (2004): “We Can Work It Out: The Impact of Technological Change on the Demand for Low-Skill Workers,” *Scottish Journal of Political Economy*, 51(5), 581–608.
- MAZZOLARI, F., AND G. RAGUSA (2013): “Spillovers from High-Skill Consumption to Low-Skill Labor Markets,” *The Review of Economics and Statistics*, 95.
- MORO, A. (2012): “The structural transformation between manufacturing and services and the decline in the US GDP volatility,” *Review of Economic Dynamics*, 15(3), 402 – 415.
- (2015): “Structural Change, Growth, and Volatility,” *American Economic Journal: Macroeconomics*, 7(3), 259–94.
- MORO, A., S. MOSLEHI, AND S. TANAKA (forthcoming): “Does home production drive structural transformation?,” *American Economic Journal: Macroeconomics*.
- NGAI, L. R., AND B. PETRONGOLO (2013): “Gender Gaps and the Rise of the Service Economy,” *Working Paper*, Centre for Economic Performance, LSE.

- NGAI, L. R., AND C. A. PISSARIDES (2008): “Trends in Hours and Economic Growth,” *Review of Economic Dynamics*, 11(2), 239–256.
- RENDALL, M. (2010): “Brain versus brawn: the realization of women’s comparative advantage,” *Available at SSRN 1635251*.
- (2015): “Female market work, tax regimes, and the rise of the service sector,” (492).
- ROGERSON, R. (2008): “Structural Transformation and the Deterioration of European Labor Market Outcomes,” *Journal of Political Economy*, 116(2), 235–259.

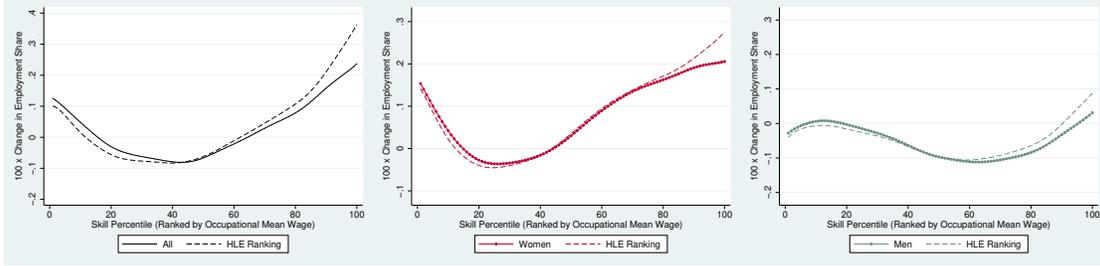


Figure 14: Data: Ranking Method. The dashed line is from [Acemoglu and Autor \(2011\)](#). The other line in each panel is our methodology as described in section 4.

Appendix

A Computing Job-Polarization

As outlined in the text we follow the methodology of [Acemoglu and Autor \(2011\)](#) in creating polarization graphs. For the benchmark graphs we use the 1980 Census of Populations (5% sample of the US) and the 2008 American Community Survey (ACS) (1% sample of the U.S). In sections 6 and 7 we also use the 1960 (1% sample of the US), 1990 and 2000 (5% sample of the US) Census of Populations. For detail on the data selection process and treatment see the Appendix A in [Autor and Dorn \(2013\)](#). The only difference here is the ranking methodology of occupations in 1980, since we not only compute average wages by occupation, but instead compute average wages by a combined measure of the three sectors and occupation Census classifications.²⁸ In Figure 14 we report the polarization graphs generated with the methodology in [Acemoglu and Autor \(2011\)](#) and ours. The resulting difference between the two ranking methods generates minor deviations.

B Treatment of the Data for the 1960-1980 period

To compute employment polarization in the 1960-1980 period we retain the same sample and data correction procedure as [Acemoglu and Autor \(2011\)](#) and [Autor and Dorn \(2013\)](#) from the polarization era for the 1960-1980 period. That is, we use all occupations that exist both in 1960 and 1980. However, as there are fewer Census occupations represented in 1960, to avoid losing a large share of the working population in 1980, we compute changes by decade (1960-1970 and 1970-1980) and then add the decades for the overall employment effect. More specifically, using the 1980 occupational ranking and applying the same procedure as

²⁸See Section 4.

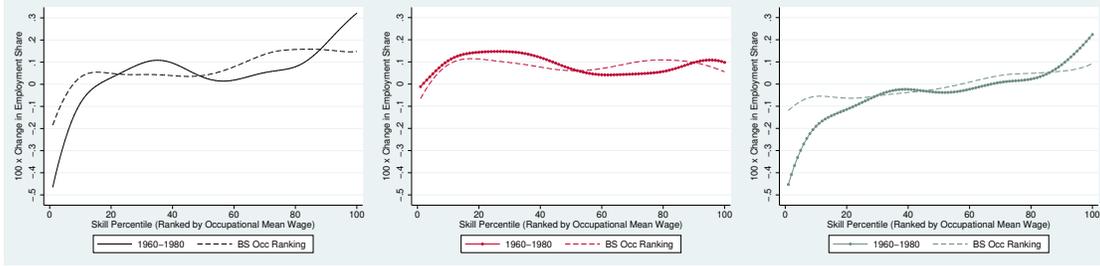


Figure 15: Different methodologies for the 1960-1980 period. The continuous line reports employment polarization as computed in Figure 9. The dashed line is constructed following the methodology in Barany and Siegel (2015).

in Acemoglu and Autor (2011) would require dropping 21.5 percent of the work force in 1980 and 6.6 percent in 1960. Instead, using a decade by decade approach drops 1 percent of the workforce in 1960, 13 percent in 1970 and 9.3 percent in 1980. Given that this methodology still does result in dropping a share of the workforce we also use the methodology in Barany and Siegel (2015).²⁹ This consists in creating a consistent occupational grouping from 1960 to 1980 to avoid dropping any of the work force. Note that with this alternative occupational classification each occupation is more heterogeneous than the original measure.³⁰ Also the new specification results in women showing no U-shape (employment-polarization) in the 1960-1980 period. In contrast, men’s changes in employment shares during the 1960-1980 period are sensitive to the methodology used. However, in general the 1960-1980 period consistently shows no polarization for the aggregate population.

C Table of polarization graphs

Table 5 summarizes the employment changes in the data and the benchmark model by skill decile. Each cell reports the average employment change (in percent) within a decile from 1980 to 2008.

²⁹We thank Zsofia Barany and Christian Siegel for sharing the occupational classification codes.

³⁰Note also that Barany and Siegel (2015) drop occupations in agriculture while here we use them.

Table 5: *Aggregate Results: Deciles*

	Decile									
	1	2	3	4	5	6	7	8	9	10
Data										
All	8.8	0.3	-5.1	-7.2	-7.7	-4.2	0.6	5.6	12	20.1
Women	9.6	-0.2	-3.5	-2.6	0.7	6.3	11.6	15.1	17.8	20.0
Men	-0.7	0.4	-1.6	-4.6	-8.4	-10.6	-11.0	-9.5	-5.8	0.2
Model										
All	19.8	5.7	-4	-9.7	-12.3	-11.5	-6.3	2.6	15.3	34.5
Women	19.4	9.7	2.5	-2.2	-4.3	-3.5	0.2	5.2	10.9	16.9
Men	0.5	-4	-6.5	-7.5	-8	-8.1	-6.5	-2.6	4.3	17.5