The Role of Gender in Employment Polarization*

Fabio Cerina† Alessio Moro‡ Michelle Petersen Rendall§

February 2018

Abstract

We document that U.S. employment polarization in the 1980-2008 period is largely generated by women. Female employment shares increase both at the bottom and at the top of the skill distribution, generating the typical U-shape polarization graph, while male employment shares decrease in a more 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 and broadly reproduces 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 most of the employment polarization generated by the model.

JEL Classification: E20, E21, J16.

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

*We thank David Autor, Vasco Carvalho, Luigi Guiso, Nezih Guner, Moshe Hazan, Esteban Jaimovich, Vahagn Jerbhasian, Omar Licandro, Rachel Ngai, Galo Nuño, Barbara Petrongolo, Chris Pissarides, Xavier Raurich, Christian Siegel, Kjetil Storesletten, Satoshi Tanaka, Marc Teigner, Montserrat Vilalta, Carlo Valdes, Fabrizio Zilibotti and seminar participants at LSE, UB, Bamberg, Cambridge, Copenhagen, Monash, New South Wales, Zurich, OECD, the 2017 Minnesota Workshop in Macroeconomic Theory, the CEPR Macroeconomics and Growth Programme (London Business School), the Barcelona GSE Summer Forum 2016, the II MadMac Conference in Growth and Development (Madrid), 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), the Second Marco Fanno Alumni Workshop (Milan) and the VII Workshop IIBEO (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.

§Monash University. E-mail: michelle.rendall@monash.edu.

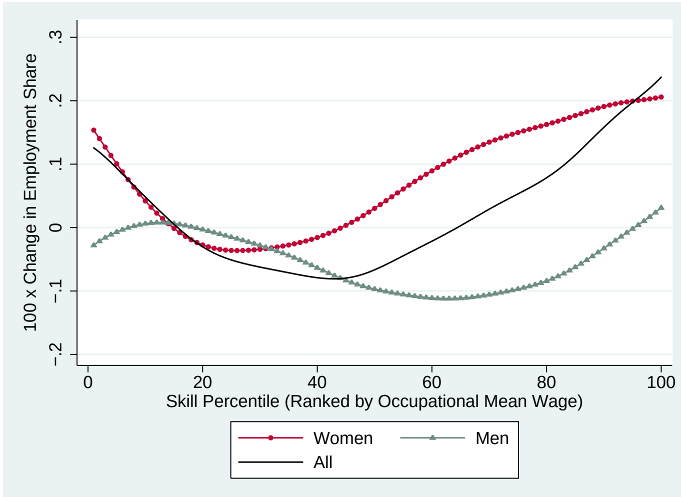

Figure 1: Changes in employment shares in the U.S. between 1980 and 2008 by skill percentile using a locally weighted smoothing regression. Data are from Census IPUMS 5 for 1980 and Census American Community Survey for 2008.

1 Introduction

Employment polarization in the U.S. has been extensively documented. The share of U.S. employment by skill rank in 1980 shows 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.[1] Less well known in the literature instead, is the behavior of polarization when distinguishing by gender, which we also report in Figure 1. As the red line suggests, aggregate employment polarization is largely driven by women. This group is 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. While changes in employment shares of men also display a U-shape, they are more homogeneous, ranging between -0.1 and 0 along the whole skill distribution except at the very top.[2] Figure 1 suggests that, while female labor force participation increases during the whole post-war period in the U.S., between 1980 and 2008 this increase occurs mainly at the top and at the bottom of the skill distribution, contributing to a large extent to the employment-polarization phenomenon. In this paper we use a calibrated macroeconomic model to show that the increase of employment shares of women at the extremes of the skill-distribution is attributable to a large extent to the emergence of skill-biased technological change after 1980.

[1]See Acemoglu and Autor (2011) and Autor and Dorn (2013) for instance.

[2]Up 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.

During the polarization era, commonly referred to as the 1980-2008 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.[3] This increase in the skill-premium is typically attributed to skill-biased technological change. In 1980, women work a high fraction of their time at home so the arrival of this type of technological change creates new incentives to increase participation in the labor market. These incentives are heterogeneous across individual, and foster employment rates especially at the top of the skill distribution. However, through general equilibrium effects, skill-biased technological change is also potentially able to increase employment rates at the bottom of the skill distribution. To see this, consider the following scenario. Skill-biased technological change 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 has the direct effect of increasing 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 main pool of workers in that market sector. With low wages in that sector, the individuals who are more likely to meet the increased demand for labor are low-skilled women working at home. 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 must be negative.[4]

The intution given above is consistent with changes in employment shares by gender and major occupational groups. Using IPUMS classification, we show that the pattern observed in Figure 1 is due to a larger increase of employment shares of women relative to men in two groups of occupation which display the bulk of their employment shares at the opposite extremes of the skill distribution: the increase at the top is driven by high-skilled managerial and professional specialty occupations while that at the bottom by low-skill services occupations. In middle-skill occupations employment shares of men and women are either similarly affected (technical, sales and administrative support occupations) or men experience a larger decline with respect to women (precision production, craft, and repair occupations, operators, fabricators and laborers).

[3] Heathcote, Storesletten, and Violante (2010) and Acemoglu and Autor (2011).

[4] Evidence from the literature on structural transformation appears to support this intuition. After being flat during the first part of the post-war period, the share of home production in total value added in the U.S. starts declining steadily between 1980 and the end of the 2000s (Moro, Moslehi, and Tanaka (2017)). This decline coincides with a rise in the market share of services substitutable to home production (typically low-skilled services - Autor and Dorn (2013) and Bridgman (2016)) and an acceleration in the share of modern market services (Moro, Moslehi, and Tanaka (2017)).

The pattern of employment shares displayed in Figure 1 is in stark contrast with the 1960-1980 period. For the latter we do not find evidence of employment polarization, and employment shares of women increase homogeneously along the skill distribution.[5] This observation is consistent with the idea that in the absence of skill-biased technological change, participation rates of women increase in a similar fashion for any skill level. Thus, if skill-biased technological change is a main driver of employment polarization in the model calibrated for the 1980-2008 period, its absence should produce changes in employment shares similar to those observed in the 1960-1980 period. In addition, as discussed in Acemoglu and Autor (2011) and Beaudry, Green, and Sand (2016), 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. 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, and the increase at the bottom being a feature of the 2000-2008 period. This time varying behavior of changes in employment shares along the skill distribution also poses a challenge to any theory of employement polarization. We thus assess the goodness of our theory by the ability of reproducing these "out of sample" observations both qualitatively and quantitatively.

Our model is designed to capture the potential general equilibrium effects that can occurr when women increase participation in the labor market in a period of sustained skill-biased technological change. We extend the canonical model discussed in 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. We 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. First, the literature identifies in structural change a main driver of the rise of female labor force participation and closing gender gaps in the labor markets.[6] Thus, while skill-biased technological change induces higher skilled women to participate more in the labor market, structural change fosters an increase in participation of all women. Second, the extent to which women participating more in the labor market substitute home production with market services depends on the relative price of the two. While a larger market wage increases the relative price of the home good, because it raises the opportunity cost of working at home, such relative

[5]Without considering agricultural occupations Bárány and Siegel (2018) report that employment polarization starts in the 1950s, but it is more pronounced in the 1980-2008 period. Here we include agriculture and find no polarization between 1960 and 1980.

[6]Ngai and Petrongolo (2017) and Rendall (2015).

price is also influenced by productivity growth in the market and at home. Thus, to assess the contribution of skill-biased technological change on gender employment polarization in a quantitative fashion, we need a setting that allows all these channels to operate.

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. 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. In this environment, as skill-biased technological change occurs, higher skilled individuals educate and increase employment shares at the top of the skill distribution. Higher wages in the market for these individuals imply less time devoted to home work and less home production which, other conditions equal, is replaced with purchases of substitutable services in the market.

Finally, 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. All sectors experience the same rate of skill-biased technological change and employ all types of workers, by gender and education level. However, the proportions of these groups are different across 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. 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 aggregate targets in the data, and evaluate its performance in replicating the main facts of employment-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, replicating fairly well employment polarization by gender, by marital status and by sector. In particular, it reproduces the higher employment polarization for women than for men. By running a counterfactual experiment in which we remove skill-biased technological change, employment polarization in the model is reduced substantially. Thus, our first contribution is to show that skill-biased technological change is a first order driver of employment polarization in our setting.

Given that a major difference between the pre- and the post-1980 is the emergence of skill-biased technological change in the more recent period, we then use the calibrated model to "forecast" changes in employment shares from 1960 to 1980, by feeding exogenous trends measured for that period. The model reproduces the flat behavior of changes in women's employment shares and the monotone behavior of men's employment shares. Our second contribution is then to show that the model can reconcile the evidence on changes in employment shares before and after 1980. To the best of our knowledge this is the first model that can account for such patterns both qualitatively and quantitatively.

We also question the ability of the model to reproduce employment polarization by decade during the 1980-2008 period. Using the benchmark calibration we feed the model with decade specific exogenous trends. We find that the model reproduces fairly well the tilting of polarization graphs observed across decades. From counterfactual experiments for each decade we are able to show that such tilting behaviour completely disappears once skill-biased technological change is removed. More precisely, in the first two decades skill-biased technological change affects mainly employment shares at the top of the skill-distribution, while in the last decade the effect is predominantly on the bottom of the distribution. From these exercises we conclude that the tilting is due to general equilibrium effects of skill-biased technological change that emerge over time. Thus, our third contribution is to show that skill-biased technological change can have a time-varying effect on the distribution of employment.

Finally, we note that this is the first paper that build a model that

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 have not been considered in previous literature, paying special attention on gender differences; 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: (1) the effect of female labor force participation on macroeconomic outcomes, (2) structural transformation, and (3) employment polarization. We connect these strands of the literature by showing that the process of employment 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. Ngai and Petrongolo (2017) and Rendall (2015), in a multi-sector model with home production, show how the process of *marketization,* occurring together with structural transformation, implies that women progressively abandon home production to work in the market.[7] Ngai and Petrongolo (2017) show that marketization and structural transformation explain together a fraction of the evolution of the gender gaps in wages and hours in the U.S. Our environment also builds on the insights in Buera and Kaboski (2012), who provide a theory predicting that the demand for skills in the labor market increases due to the rise of services that are skill-intensive, with a contemporaneous decline of home production. Buera, Kaboski, and Zhao (2013) evaluate quantitatively such theory by also introducing skill-biased technical change and gender, and find that both a higher demand for output which is skill-intensive and increasing female labor supply are key factors to explain the growth of services. In this paper we link skill-biased technological change to a gender and a sector dimension to study employment polarization.

Recently, a number of 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, especially relevant for our work, employment levels (Rogerson (2008)). However, few works relate this process to employment-polarization.[8] Autor and Dorn (2013) provide an explanation of employment-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, Bárány and Siegel (2018) 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 trends of these three sectors imply that the

[7]The relationship between home production and structural transformation has been extensively studied in the literature. See, among others, Rogerson (2008) and Ngai and Pissarides (2008).

[8]In recent work Duernecker 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.

share of manufacturing shrinks with respect to the other two sectors. By their 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. Therefore, the process of structural transformation generates employment polarization in this environment. While we use a specification of preferences similar to Bárány and Siegel (2018), 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, also increases the demand for services in the market that have a home counterpart. We build on such intuition by adding the gender dimension and giving a primary role to skill-biased technological change, allowing for an explanation of the different pattern of employment polarization across gender. 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 increase 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.

In addition to consumption spillovers, in our model skill-biased technological change triggers another kind of of complementarity between top and bottom skills in our setting, that of production complementarity between educated and uneducated workers within the firm. A similar mechanism is proposed in Eeckhout, Pinheiro, and Schmidheiny (2014) to explain the thick tails of the skill distribution in large cities in the U.S., that is, the fact that these cities disproportionately attract both high- and low-skilled workers, while average skills are constant across city size.

We conclude this section relating to the well known *routinization hypothesis*.[9] This process, driven by the rise of information technology, implies that workers employed in occupations containing a large share of routinary tasks become redundant, because these tasks are

[9] See Autor, Levy, and Murnane (2003) and subsequent literature.

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. Acemoglu and Autor (2011) provide a description of the canonical model with skill-biased technological change and show that, absent any distinction between tasks and skills, it cannot reproduce the aggregate employment-polarization pattern. Building on that result, Autor and Dorn (2013) show that instead, the aggregate pattern of employment polarization can be generated in a model displaying routine-biased technological change, in which computer capital is complementary to skilled labor and substitutable to middle-skilled one.[10] Here we 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. As we discuss in the next section, our explanation is orthogonal to the routinization hypothesis so that both channels (skilled- and routine-biased technological change) can simultaneously be playing a role in driving employment polarization.[11] We stress here that our theory is able to account for at least two empirical patterns on which the routinization hypothesis is mostly silent. The first is the difference of employment polarization patterns across gender. The second is the tilting of the polarization graph in the three decades between 1980 and 2008.

3 Empirical evidence

3.1 Facts on Employment Polarization

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

$$\frac{H_{i,2008}}{H_{2008}} - \frac{H_{i,1980}}{H_{1980}}, \tag{1}$$

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

$$\frac{H_{i,2008}}{H_{2008}} - \frac{H_{i,1980}}{H_{1980}} = \left(\frac{H^f_{i,2008}}{H_{2008}} - \frac{H^f_{i,1980}}{H_{1980}}\right) + \left(\frac{H^m_{i,2008}}{H_{2008}} - \frac{H^m_{i,1980}}{H_{1980}}\right) \tag{2}$$

[10]See also Sevinç (2017), who discusses the ability of a model of skill-biased technological change to reproduce changes of employment shares by skill level when the latter is measured using, instead of the average wage of the occupation, the share of graduates working in that occupation. He suggests that this alternative ranking of skills produces a pattern of employment shares which is more in line with the predictions of a canonical skill-biased technological change model.

[11]In contemporaneus research, Cortes, Jaimovich, and Siu (2016) document that one third of the disappearance of routine occupations are due to demographic changes in the U.S. In a similar vein, we focus here on the role of specific demographics groups in accounting for the whole process of employment polarization.

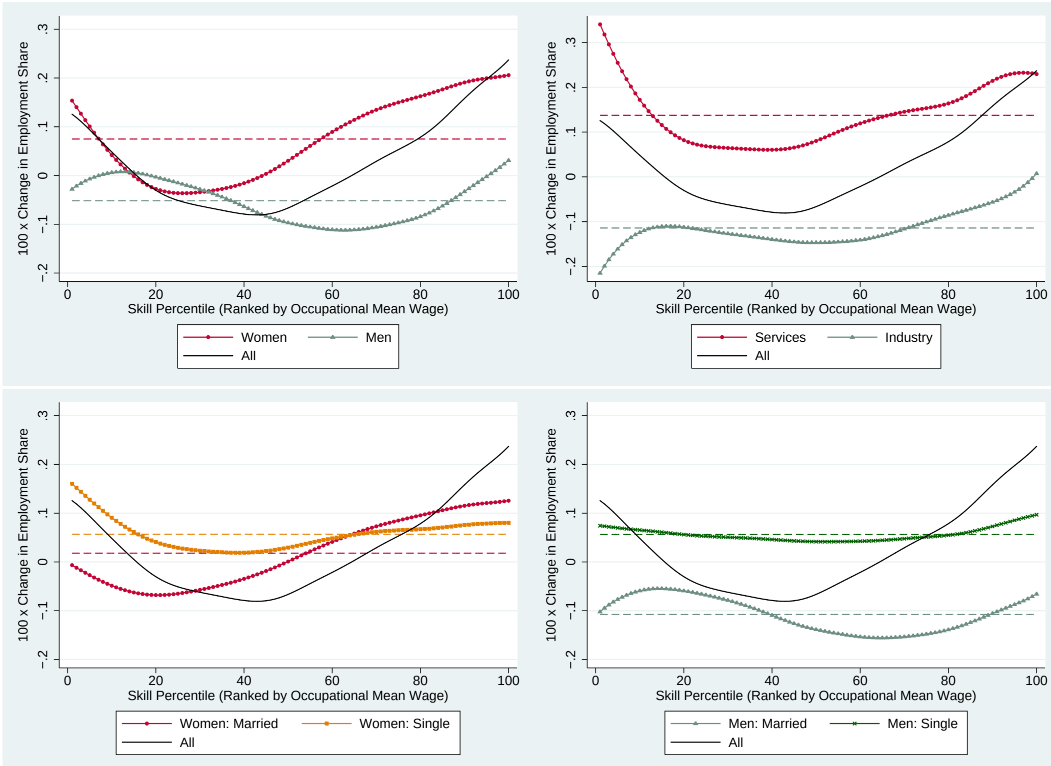

Figure 2: Employment polarization by gender, marital status, and broad sector of economic activity. For each color, the horizontal dashed line indicates the average change of employment shares for the corresponding group. In the top panels the vertical sum of the colored lines gives the black line. In each of the bottom panels the vertical sum of the colored lines of married and singles gives the correponding gender line of the top-left panel.

where $H^f_{i,t}$ is total hours worked by women in percentile i in year t, and $H^m_{i,t}$ 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, $H_{i,t} = H^f_{i,t} + H^m_{i,t}$. 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 five facts on employment-polarization in the U.S., with special attention to the gender dimension and factors that can affect the allocation of working hours by gender, like the marital status and the broad sector of economic activity (i.e. manufacturing and services).[12] All facts below, except for number 4, refer to the period 1980-2008. Also, in each polarization graph we report horizontal lines measuring the average change in employment shares of the corresponding group of interest.

Fact 1 (Gender): Employment polarization is larger for women than for men (top-left

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

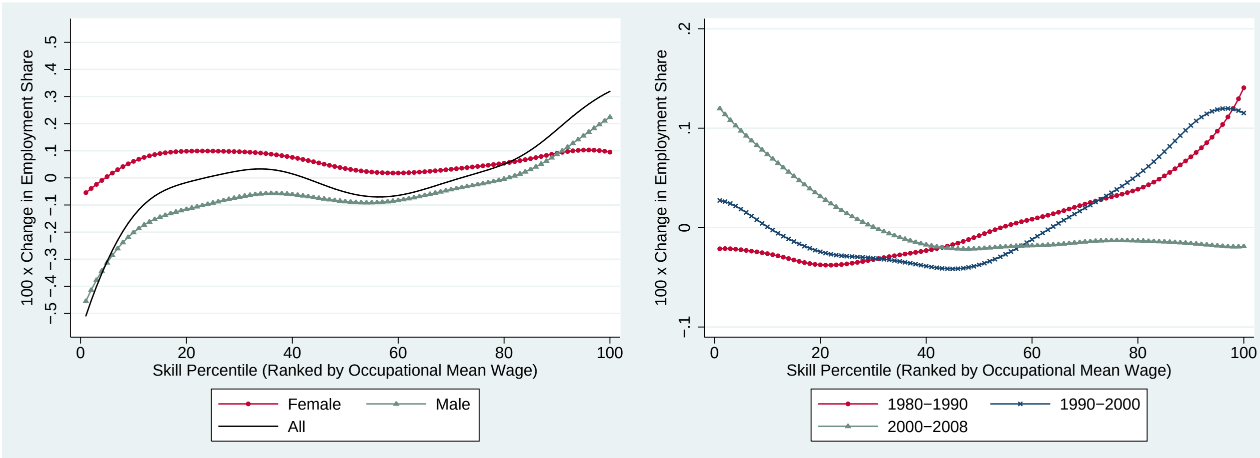

Figure 3: employment polarization in the 1960-1980 period (left) and in the three decades during the 1980-2008 period (right).

panel of Figure 2).[13]

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

Fact 3 (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 (bottom panels of Figure 2).

Fact 4 (Employment polarization before 1980): Employment polarization is absent in the 1960-1980 period. Changes in employment shares of women are homogeneous along the skill distribution, while those of men are increasing along the distribution for (almost) any percentile (left panel of Figure 3).

Fact 5 (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).

As discussed in the introduction, fact 1 suggests a key role of women in generating employment polarization. Fact 2 suggests that employment polarization is more evident in

[13] See the discussion below for a description of the measures we use to state that employment polarization is larger for women than for men.

the broad sector in which women display a comparative advantage (services) than in the sector that typically favors men work (manufacturing).[14] Fact 3 highlights that the marital status plays an important role in shaping the polarization curve for women, while for men it mainly 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. Fact 4 documents a striking difference of changes in employment shares for both gender between the pre- and post-1980. Finally, fact 5 suggests that even after 1980, there is not a consistent pattern of changes in employment shares along the skill distribution. A theory that aims at accounting for overall employment polarization should be able to account for this evidence. In particular it should potentially explain the different role of the various demographic groups and the time-varying pattern of employment polarization. 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 to 5.

3.2 Assessing Gender Differences in Employment Polarization

While there is not a universally accepted measure that allows to compare two sets of data and determine which one is more polarized, we discuss here some measures to assess gender differences in employment polarization and support the claim that changes in employment shares of women are more polarized then those of men.

First, the standard deviation of employment shares with respect to the gender-specific mean is 0.084 for women and 0.044 for men, implying that changes in female employment shares along the skill distribution display twice the variation than that of males. While a high standard deviation does not imply polarization, the top-left panel of Figure 2 shows that the positive change at the bottom and at the top percentiles relative to the gender-specific mean (the dashed line) is larger for women than for men. At the same time, the highest negative change (in absolute value) for women, around percentile 25, is larger than the highest negative change (in absolute value) for men, which is around percentile 60. Thus, the extreme points in Figure 2 suggest a more pronounced V shape, around the average change, for women than for men. In addition, we compute male and female changes in employment shares by percentiles *within each gender*, i.e. using as denominator in 1980 and 2008 the gender-specific amount of hours worked instead of the total one.[15] This provides a measure

[14] Ngai and Petrongolo (2017) and Rendall (2010).

[15] More precisely, in place of the expressions included in (2), we focus on the following two formulas

$$\left(\frac{H^f_{i,2008}}{H^f_{2008}} - \frac{H^f_{i,1980}}{H^f_{1980}}\right); \left(\frac{H^m_{i,2008}}{H^m_{2008}} - \frac{H^m_{i,1980}}{H^m_{1980}}\right),$$

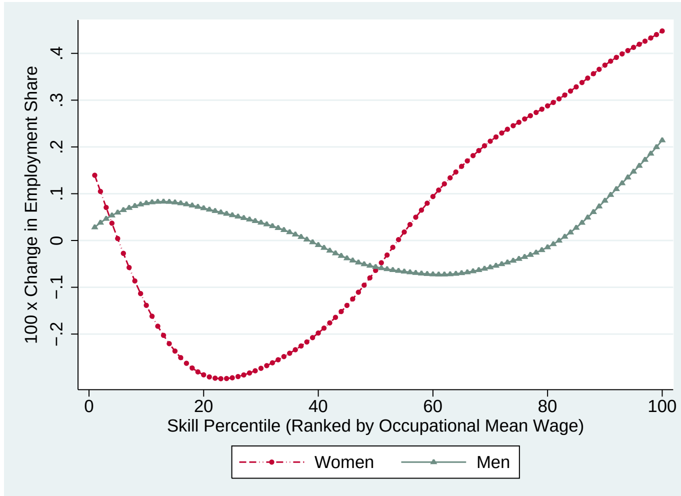

Figure 4: Changes in employment shares in the U.S. between 1980 and 2008 by skill percentile normalized for gender specific total working hours.

of the heterogeneity of changes in employment shares for each gender which is independent (in an accounting perspective) from the fact that women gain on average employment shares relative to men. The measure, plotted in Figure 4 for both gender, reinforces the view that employment is substantially more polarized for women than for men. In this case the standard-deviation over the gender-specific mean is 0.24 for women and 0.07 for men.

Next, we follow an approach similar to Goos and Manning (2007) and fit, for the aggregate economy, for women, and for men, the following quadratic relationship:

$$\triangle n_j = \beta_0 + \beta_1 j + \beta_2 j^2, \tag{3}$$

where $j = 1, ..100$ is the percentile number and $\triangle n_j$ is the change in log-employment at percentile j. Table *1* reports the results. Similar to Goos and Manning (2007), we find a positive coefficient of the quadratic term and a negative one on the linear term, suggesting the presence of polarization for the three groups. However, the coefficient of the quadratic term in the aggregate regression is driven mainly by women, who display a value of that coefficient almost four time larger than that of men. This implies a higher convexity of the quadratic relationship for women than for men. Also, the coefficient of the quadratic term is statistically significant for women and at the aggregate level, but not for men. The better fit of the quadratic relationship for women relative to men is also highlighed by the large value of the R-squared. In section 6 and 7 we perform similar regressions for the model, and use

where $H^f_{1980} = \sum_i H^f_{i,1980}$, $H^m_{1980} = \sum_i H^m_{i,1980}$, $H^f_{1980} = \sum_i H^f_{i,2008}$, and $H^f_{2008} = \sum_i H^f_{i,2008}$. Notice that, unlike in (2), the sum of these gender measures does not deliver the aggregate polarization curve.

Table 1: *Quadratic fit of the data*

	(1)	(2)	(3)
	All 1980-2008	Males 1980-2008	Females 1980-2008
Rank	-0.784	-0.777	-1.083
S.E.	(0.552)	(0.612)	(0.694)
t-statistic	1.419	1.270	1.561
Rank^2	0.978*	0.583	2.114**
S.E.	(0.530)	(0.587)	(0.665)
t-statistic	1.848	0.993	3.178
Observations	100	100	100
R-squared	0.055	0.026	0.330

Table 2: *Change in employment shares by groups of occupations and gender*

Occupational Group	Wage 1980	% Emp. share 1980			Change until 2008		
		All	Male	Fem.	All	Male	Fem.
Managerial, professional	2.99	23.96	15.78	8.18	12.01	2.90	9.11
Precision production, craft, repair	2.85	13.80	13.09	0.72	-3.56	-3.45	-0.11
Operators, fabricators, laborers	2.62	21.75	16.28	5.47	-8.79	-5.91	-2.88
Technical, sales, admin support	2.52	30.04	12.65	17.39	-1.91	-0.89	-1.02
Farming, forestry and fishing	2.49	0.14	0.13	0.01	0.04	0.01	0.03
Service	2.30	10.30	4.90	5.40	2.22	0.66	1.55
TOTAL	2.68	100.00	62.84	37.16	0.00	-6.68	6.68

the estimated coefficients on the quadratic term to provide a measure of the performance of the theory relative to the data.

Finally, we analyze the difference between the two gender by fousing on the dynamics of employment shares by broad groups of occupations. The zoom on occupations allows us to relate our argument to the mainstream explanation for the decline of employment shares in the middle of the skill distribution, the *routinization* hypothesis.[16] This process makes workers employed in occupations containing a large share of routinary tasks redundant, as the latter are taken up by computers. The evidence provided in the literature suggests that these types of occupations are in the middle of the skill distribution in 1980. Table 2 reports the dynamics of employment shares by gender and broad groups of occupations (ordered by their mean log hourly wage in 1980) as categorized by the IPUMS database.[17]

[16] Acemoglu and Autor (2011), Acemoglu and Autor (2012), Autor and Dorn (2013).

[17] Details on how occupations are grouped can be found at the web page https://usa.ipums.org/usa-action/variables/OCC1990#codes_section

The aggregate trend of employment polarization is confirmed: highest paid occupations (managerial, professional specialty occupations) are associated with the largest increase in the employment share (from 24% to 36%, a one-half increase). These occupations are also those where cognitive, "abstract", creative, problem-solving, socially interactive skills tend to be more required. On the other hand, even the group of lowest paid occupations (service occupations) increases its employment share by about one fourth (from 10.30% to 12.52%). This category includes jobs that involve assisting or caring for others: food service workers, security guards, janitors and gardeners, cleaners, home health aides, child care workers, hairdressers and beauticians, and recreation occupations. Service occupations are also those where manual tasks are more concentrated and hence cannot be easily automated. At the same time, the remaining groups of occupations in the middle of the wage distribution are all declining their employment shares between 1980 and 2008. This is particularly true for the two groups "precision production craft and repair occupations" (from 13.80% to 10.24%) and "operators, fabricators and laborers" (from 21.75% to 12.95%) where routine tasks are highly concentrated and therefore are highly substitutable with computers. Thus, as extensively reported in the literature, the data by broad occupation groups are consistent with the routinization hypothesis.[18] However, the evidence reported in Table 2 is also consistent with the mechanisms proposed in this paper, where the increase of labor participation by high-skilled women after 1980 due to skill biased technological change has a key role. The first element we emphasize is the remarkable differences between men and women in the dynamics of the employment shares among different occupational groups. Such differences can be only partly captured by a general "level" effect implying that women increase their total employment share by 6.68 percentage points (with a corresponding decrease for men), as the changes in the employment shares are highly asymmetric along the skill distribution. In particular, women more than double their employment share in occupations at the upper tail of the distribution (from 8.18% to 17.29%), while men share increases only by less than 20% (from 15.78% to 18.68%). This is consistent with the view according to which skill-biased technological change attracts to the market mainly (high-skilled) women who previously worked at home.[19] On the other extreme, the only other group of occupations where women increase their employment share is the service one: here female employment share grows by almost 30% (from 5.4% to almost 7%) while men only by 13% (from 4.90% to 5.56%). These service occupations are highly concentrated in sectors producing services which are highly substitutable to household production (especially child care workers, gardeners, cleaners,

[18]Our table can be compared to Table 1 in Autor and Dorn (2013). The main difference is that we aggregate occupations according to the criteria suggested by the IPUMS database itself, while they do it following their own aggregation criteria. Moreover, we add the gender dimension.

[19]In 1980, educated men worked 83% of their time while educated women only 49% (CENSUS).

home health aides). Some of these jobs (especially food service workers, security guards, janitors) also support the jobs of high-skilled workers and therefore they are complementary to highest paid occupations. Hence, this evidence appears consistent also with the hypothesis of extreme skill complementarity according to which the raise in labor participation by high-skilled women in the 80s generated an indirect increase in employment shares in occupations at the bottom of the distribution, both because of *consumption* spillovers (replacement of household services with market services) and *production* complementarities (more managerial jobs requires more cleaners, food workers, guards and so on).

The evidence produced in this section suggests that, in an accounting perspective women entering the labour market asymmetrically along the skill distribution have a key role in determining employment polarization in the aggregate during the period 1980-2008. Accordingly, in the next sections we investigate whether these patterns can be accounted for by skill-biased technological change in a general equilibrium setting. Disentangling the effect of our theory from that of routinization and assessing the respective contribution of the two in explaining job polarization remains an open issue that we leave 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.

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).[20] Hence, each agent is endowed with a triple of skills $a^i = [a^i_{ss}, a^i_g, a^i_{ms}]$, where $i = f, m$, and f stands for female and m for male. Thus, there exist two density functions of agents with characteristics $[a^i_{ss}, a^i_g, a^i_{ms}]$. 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 $f(a^i) = f\left(a^i_{ss}, a^i_g, a^i_{ms}\right)$.

Each agent is also endowed with one unit of time. She splits this between work at home

[20]The methodology to define substitutable services in the data is described in Section 5.

(l) and work in the market ($1-l$). Thus, a unit of time of agent of type i, depending on the sector it is employed, can correspond to: i) a^i_{ms} efficiency units of labor to production in sector ms; ii) a^i_g efficiency units of labor to production in sector g; iii) a^i_{ss} 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^i_j to $\left(a^i_j\right)^{1+\zeta}$.[21] As in Heathcote, Storesletten, and Violante (2010), we assume that agents draw the cost of education χ^i from a gender specific distribution such that $\log(\chi^i) \sim N\left(\mu^i_\chi, \left(\sigma^i_\chi\right)^2\right)$, $i = f, m$. By acquiring education, the agent upgrades her wage per unit of efficiency, $w^{i,e}_j$, from that of uneducated, $w^{i,0}_j$, to that of educated individuals, $w^{i,1}_j$, 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 $a^i = [a^i_{ss}, a^i_g, a^i_{ms}]$ and facing a vector of equilibrium market wages $w^{i,e} = [w^{i,e}_{ss}, w^{i,e}_g, w^{i,e}_{ms}]$ is then a pair $(e^*, j^*) = [e(a^i, w^{i,e}, i), j(a^i, w^{i,e}, i)] \in \{0, 1\} \times \{ss, g, ms\}$ such that

$$(e^*, j^*) = argmax_{(e,j)} \left[w^{i,e}_j \left(a^i_j\right)^{(1+e\zeta)} - e\chi^i\right] \tag{4}$$

Notice that conditional on $e^* = 0$, $(0, j^*) = argmax_{(0,j)} \left[w^{i,0}_j \left(a^i_j\right)\right]$, 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 $w^{i,0}_j a^i_j$. By contrast, conditional on $e^* = 1$, $(1, j^*) = argmax_{(1,j)} \left[w^{i,1}_j \left(a^i_j\right)^{(1+\zeta)} - \chi^i\right]$, 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 $argmax_{(0,j)} \left[w^{i,0}_j a^i_j\right] \neq argmax_{(1,j)} \left[w^{i,1}_j \left(a^i_j\right)^{(1+\zeta)} - \chi^i\right]$, 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

[21] This 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).

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 education level e and in which sector j to work to maximize her wage net of education costs, $w_j^{i,e}\left(a_j^i\right)^{(1+e\zeta)} - e\chi^i$. 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

$$W\left(a_{j^*}^i, w_{j^*}^{i,e^*}, e^*\right) = w_{j^*}^{i,e^*}\left(a_{j^*}^i\right)^{(1+e^*\zeta)} - e^*\chi^i \tag{5}$$

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

$$U^z = \left((\omega_{ms})^{1/\sigma}\left(\frac{c_{ms}^z}{\kappa^z}\right)^{\frac{\sigma-1}{\sigma}} + (\omega_g)^{1/\sigma}\left(\frac{c_g^z}{\kappa^z}\right)^{\frac{\sigma-1}{\sigma}} + (\omega_s)^{1/\sigma}\left(\frac{\tilde{c}_{ts}^z}{\kappa^z}\right)^{\frac{\sigma-1}{\sigma}}\right)^{\frac{\sigma}{\sigma-1}}, \tag{6}$$

$$\tilde{c}_{ts}^z = \left(\psi\left(c_{ss}^z\right)^{\frac{\gamma-1}{\gamma}} + (1-\psi)\left(c_h^z\right)^{\frac{\gamma-1}{\gamma}}\right)^{\frac{\gamma}{\gamma-1}} + \bar{c} \tag{7}$$

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.[22] Following the findings in Moro, Moslehi, and Tanaka (2017) 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

[22] So 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.

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

$$Y_h^z = A_h L^z, \tag{8}$$

where

$$\begin{aligned} L^c &= A_h \left[\varphi_h^c \left(l^f \right)^{\frac{\eta-1}{\eta}} + (1-\varphi_h^c) \left(l^m \right)^{\frac{\eta-1}{\eta}} \right]^{\frac{\eta}{\eta-1}}, &(9)\\ L^f &= A_h \left(\varphi_h^f \right)^{\frac{\eta}{\eta-1}} l^f, &(10)\\ L^m &= A_h \left(\varphi_h^m \right)^{\frac{\eta}{\eta-1}} l^m, &(11) \end{aligned}$$

The budget constraint changes across household types being

$$p_{ms} c_{ms}^z + p_g c_g^z + p_{ss} c_{ss}^z = E^z, \tag{12}$$

where

$$\begin{aligned} E^c &= W\left(a_{j^*}^i, w_{j^*}^{i,e^*}, e^*\right)(1-l^f) + W\left(a_{j^*}^i, w_{j^*}^{i,e^*}, e^*\right)(1-l^m), &(13)\\ E^f &= W\left(a_{j^*}^i, w_{j^*}^{i,e^*}, e^*\right)(1-l^f), &(14)\\ E^m &= W\left(a_{j^*}^i, w_{j^*}^{i,e^*}, e^*\right)(1-l^m). &(15) \end{aligned}$$

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 $W\left(a_{j^*}^f, w_{j^*}^{f,e^*}, e^*\right)$ irrespective of her husband's choice as the households maximizes total utility.[23]

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 (6), subject to the service aggregator (7), the budget constraint (12) and the home production technology constraint (8).

From first order conditions we obtain the relative time of work at home of spouses, which,

[23] A similar discussion can be made for a married men.

in an interior solution, is given by

$$\frac{l^f}{l^m} = \left(\frac{\varphi_h}{1-\varphi_h} \frac{W\left(a^m_{j^*}, w^{m,e^*}_{j^*}, e^*\right)}{W\left(a^f_{j^*}, w^{f,e^*}_{j^*}, e^*\right)} \right)^{\eta}. \tag{16}$$

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

$$p^c_h = \frac{1}{A_h} \left[\varphi^{\eta}_h \left[W\left(a^f_{j^*}, w^{f,e^*}_{j^*}, e^*\right) \right]^{1-\eta} + (1-\varphi_h)^{\eta} \left[W\left(a^m_{j^*}, w^{m,e^*}_{j^*}, e^*\right) \right]^{1-\eta} \right]^{\frac{1}{1-\eta}}. \tag{17}$$

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

$$p^i_h = \frac{W\left(a^i_{j^*}, w^i_{j^*}, e^*\right)}{A_h} \left(\varphi^i_h\right)^{-\frac{\eta}{\eta-1}}. \tag{18}$$

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 (17) and (18) 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 the market for the two types of households. Equations (17) and (18) also highlight that individuals with a higher wage also face a higher price of home production, which implies that substitutable services in the market are relatively cheaper for these individuals. Thus, as skill-biased technological change occurs in the model, educated indviduals experience an increase in the relative price of the home good, and react by increasing the demand for substitutable services in the market. In this way, the boost of productivity of high-skilled workers creates a feedback effect on employment levels of substitutable services, which is absent in models of skill-biased technological change, suggesting that the latter can have a non-monotonic effect on changes of employment shares along the skill distribution.

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

$$Y_j = A_j N_j, \tag{19}$$

where

$$N_j = \left[\phi_j \left(\varphi_j N_j^{f,1} + (1-\varphi_j) N_j^{m,1} \right)^{\frac{\eta_s - 1}{\eta_s}} + (1-\phi_j) \left(\varphi_j N_j^{f,0} + (1-\varphi_j) N_j^{m,0} \right)^{\frac{\eta_s - 1}{\eta_s}} \right]^{\frac{\eta_s}{\eta_s - 1}}, \tag{20}$$

and $N_j^{i,e}$ 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).

Note that we keep the assumption of gender-biased technological change within each firm, as in Heathcote, Storesletten, and Violante (2010), because, although the gender wage gap closes endogenously in the model, its effect is not quantitatively strong enough. We show in section 6 that gender-biased technological change has a homogeneous effect on employment shares along the skill distribution, so its absence would not change the predictions regarding employment polarization. However, without gender-biased technological change the model could not match the closing of aggregate gender wage gaps.

The representative firm operating in sector j maximizes profits

$$\pi_j = p_j Y_j - w_j^{f,1} N_j^{f,1} - w_j^{m,1} N_j^{m,1-} w_j^{f,0} N_j^{f,0} - w_j^{m,0} N_j^{m,0} \tag{21}$$

subject to (19) and (20).

First order conditions imply

$$\frac{\phi_j \left(\varphi_j N_j^{f,1} + (1-\varphi_j) N_j^{m,1} \right)^{-\frac{1}{\eta_s}}}{(1-\phi_j) \left(\varphi_j N_j^{f,0} + (1-\varphi_j) N_j^{m,0} \right)^{-\frac{1}{\eta_s}}} = \frac{w_j^{m,1}}{w_j^{m,0}} \tag{22}$$

$$\frac{\varphi_j}{1-\varphi_j} = \frac{w_j^{f,e}}{w_j^{m,e}} \tag{23}$$

Equation (22) shows that, other conditions equal, skill-biased technological change due to an increasing ϕ_j, raises the skill premium. Equation (23) 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 $\left\{w_{ss}^{f,1}, w_g^{f,1}, w_{ms}^{f,1}, w_{ss}^{m,1}, w_g^{m,1}, w_{ms}^{m,1}, w_{ss}^{f,0}, w_g^{f,0}, w_{ms}^{f,0}, w_{ss}^{m,0}, w_g^{m,0}, w_{ms}^{m,0}\right\}$, a set of choices for each agent (e^*, j^*) and a set of allocations for each household $\left\{c_{ss}^z, c_g^z, c_{ms}^z, l^{fz}, l^{mz}\right\}$ such that:

1. Given wages and prices, the choice (e^*, j^*) maximizes wages net of education costs for agent i by solving (4);
2. Given wages, prices, and (e^*, j^*) of each household member, the allocation $\left\{c_{ss}^z, c_g^z, c_{ms}^z, l^{fz}, l^{mz}\right\}$ maximizes utility (6) of the household subject to the budget constraint (12);
3. Given wages and prices, each representative firm in sectors ss, g, and ms maximizes profits (21);
4. Labor markets in sectors ss, g, and ms clear;
5. Goods markets in sectors ss, g, and ms clear.

5 Mapping between Model and Data

We calibrate the model to two equilibria to replicate a series of aggregate 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) the level of gender-biased technology; and iv) marriage rates. We divide this section in three parts. In subsection 5.1 we describe the targets and the calibration of the parameters. We devote subsection 5.2 to explain the mapping between the model and the data regarding sectors and demographic trends. Finally, in subsection 5.3 we describe the mapping between an occupation in the data and in the model and the methodology used to construct employment polarization graphs from the model's equilibrium.

5.1 Calibration

A number of parameters, $\{\sigma, \gamma, \eta, \eta_s\}$, are set from previous studies. 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 $a_j \in [\underline{a}_j, \overline{a}_j]$ and with men and women drawing from the same ability distribution by sector when born. The lower bound of ability in the substitutable service sector is $\underline{a}_{ss} = 0.5$. 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 $A_{h,1980} = 1$. 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 (2017) for the 1978-2010 period. We choose the value of 0.001 in our calibration. 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.

The remaining 28 parameters: (1) ability and return to education $\{\overline{a}_{ss}, \underline{a}_{ms}, \overline{a}_{ms}, \underline{a}_g, \overline{a}_g, \zeta\}$, (2) productivity (market and home) $\{\{\varphi_{j,1980}, \phi_{j,1980}\}_{j=ms,g,ss}, \varphi_h, \varphi_h^f, \varphi_h^m\}$, (3) preferences and distribution of education cost $\{\omega_{ms}, \omega_g, \psi, \bar{c}, \mu_\chi^f, \mu_\chi^m, \sigma_\chi^f, \sigma_\chi^m\}$, 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.[24] Table 3 lists the parameter values used in the simulation and the standard errors obtained using a nonparametric bootstrap, by sampling individuals with replacement. While the calibration procedure matches all 28 parameters to 28 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, $\{\overline{a}_{ss}, \underline{a}_{ms}, \overline{a}_{ms}, \underline{a}_g, \overline{a}_g\}$ (5 targets): male modern services and 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 in 1980. Gender gaps in the market in 1980, $\{\varphi_{j,1980}\}_{j=ms,g,ss}$ (3 targets): aggregate gender wage gap, female to male industry hours gap, female substitutable services to modern services hours gap. Education determinants, $\{\zeta, \mu_\chi^f, \mu_\chi^m, \sigma_\chi^f, \sigma_\chi^m, \{\phi_{j,1980}\}_{j=ms,g,ss}\}$ (8 targets): the male and the

[24] Note that $\omega_{ss} = 1 - \omega_{ms} - \omega_g$.

Table 3: *Model Parameters*

Estimated	**Type**	**Value**	**S.E.**
$\{\underline{a}_{ss}, \overline{a}_{ss}\}$	Substitutable services ability	$\{0.50, 3.37\}$	$\{-, 0.0011\}$
$\{\underline{a}_{ms}, \overline{a}_{ms}\}$	Modern services ability	$\{1.05, 4.87\}$	$\{0.0014, 0.0008\}$
$\{\underline{a}_{g}, \overline{a}_{g}\}$	Manufacturing ability	$\{0.77, 4.40\}$	$\{0.0055, 0.0009\}$
ω_{ms}	Consumption market weight modern services	0.43	0.0002
ω_{g}	Consumption market weight manufacturing	0.33	0.0002
ψ	Substitutable services weight	0.25	0.0003
φ_h^c	Home female-labor weight	0.54	0.0006
φ_h^f	Single female home labor weight	0.41	0.0016
φ_h^m	Single male home labor weight	0.50	0.0022
$\varphi_{ms,1980}$	Female-labor weight in modern services	0.34	0.0002
$\varphi_{g,1980}$	Female-labor weight in manufacturing	0.31	0.0002
$\varphi_{ss,1980}$	Female labor weight in substitutable services	0.37	0.0002
ζ	Schooling factor	0.21	0.0010
μ_χ^f	Mean of the cost of education female	0.64	0.0063
μ_χ^m	Mean of the cost of education male	1.26	0.0199
σ_χ^f	Variance of the cost of schooling female	0.94	0.0048
σ_χ^m	Variance of the cost of schooling male	1.05	0.0229
$\phi_{ms,1980}$	Educated workers labor weight in modern services	0.34	0.0003
$\phi_{g,1980}$	Educated workers labor weight in manufacturing	0.32	0.0007
$\phi_{ss,1980}$	Educated workers labor weight in substitutable services	0.38	0.0008
$\bar{c}$	Non-homothetic consumption in traditional services	-0.09	0.0002
γ_{ms}	Annual growth in A_{ms}	0.004	0.0002
γ_{ss}	Annual growth in A_{ss}	0.017	0.0002
γ_{g}	Annual growth in A_{g}	0.034	0.0002
γ_{ϕ}	Skill-biased tech. change (annual growth rate in ϕ_j)	0.013	0.0001
γ_{φ}	Gender-biased tech. change (annual growth rate in φ_j)	0.005	0.0001

Predet.	**Type**	**Value**
σ	Substitutability between broad cons. 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
γ_h	Annual growth in A_h	0.001

Note: The first set of parameters is estimated (except $\underline{a}_{ss}$) while the second set is predetermined. Column S.E. displays, for estimated parameters, the standard errors obtained obtained through parametric bootstrap with 500 repetitions.

female college wage premium in 1980, the share of educated men and women 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 sectoral productivity growth rates $\{\gamma_j\}_{j=ms,g,ss}$ (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 aggregate gender wage gap between 1980 and 2008. All targets are computed using the 1980 Census and the 2008 American Community Survey unless noted. We emphasize here that parameters are calibrated using *aggregate* targets and the model endogenously produces heterogeneous changes of employment shares along the skill distribution.

5.2 Sectors and demographic trends

To define services sectors that are substitutable to home we use the procedure in Moro, Moslehi, and Tanaka (2017). 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 in U.S. data are summarized in table 5. To match such trends in the model, we create probability weights for each type of agent related to the different marriage rates (including the assortative mating patterns). More specifically, we create a matrix of the population of size 50,000x2 of males and females, respectively. We assume column one is made up of only men and column two of only women. The demographic structure is then constructed in three steps. First, male agents are created with random draws of abilities from the three uniform distributions by sector. Female agents (by row) draw from the same uniform distributions, but adjusted by a correlation coefficient between men and women by sector. That is, each row represents a *potential* household, which has

Table 4: *Model Targets*

Type	**Data**	**Model**
1980 - ability $\left(\{\underline{a}_j, \overline{a}_j\}_{j=ms,g,ss}\right)$		
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 $\left(\{\mu_\chi^i\}_{i=m,f}\right)$		
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 $\left(\psi, \varphi_h^c, \varphi_h^f, \varphi_h^m\right)$		
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 $\left(\sigma_\chi^m, \sigma_\chi^f\right)$		
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 5: *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

the correlation of abilities of married couples found in U.S. wage data by sector.[25] Second, each agent in a row chooses her/his education outcome independently of the other gender in the same row. Given the cost distribution of education, there is only imperfect sorting into college. Lastly, each row is given a probability of being either married or single such that the shares in table 5 are matched. Thus, while the population of agents is the same in each steady state, both aggregate marriage trends and assortative mating patterns are identical in the model and data.

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 aggregate marriage rate decreases and 4) assortative matching by education level increases.

5.3 Constructing Employment Polarization Graphs in Data and Model

To construct the data graphs we proceed as follows. Within each market sector in 1980 (modern services, manufacturing and substitutable services) we compute the average wage in 1980 of each occupation at the three digit level. Then, we rank these occupations from

[25]To compute the correlation between husband and wife wages, we compute female wages by sector correcting for selection bias using the Heckman correction, and then 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 substitutable services and 0.26 for modern services.

the three sectors into a unique classification according to their average wage and construct occupation percentiles using Census weights and hours worked of each occupation. By keeping the same ranking in 2008 we construct employment polarization graphs by measuring the change in employment share of each 1980 percentile and using a locally weighted smoothing regression.[26]

Note that we use the same method as in Acemoglu and Autor (2011) to construct the polarization graphs in the data, only adapted to consider that we have three sectors in the model. For instance, in the data, certain occupations are in all three sectors (e.g., secretaries), while others are only in one of the three (miners). In this case, 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. As Appendix A shows, the differences between the two methods are very minor, and mostly at the right tail of the overall distribution.

To our knowledge, this is the first paper that compares employment polarization graphs in the data with the outcome of a general equilibrium model. In the model we do not explicitely introduce occupations, thus one challenge is how to draw polarization graphs that are comparable with those in the data in a meaningful way. To do this we first note that to construct polarization graphs in the data, occupations are ranked using wages in 1980 as a proxy for skills. Thus, an interpretation of an occupation in the data is that of a group of workers (those employed in that occupation) with similar skills (who are paid a similar wage). We thus assume that the model counterpart of an *occupation* in the data is a *set of workers with similar skills* and working in a given sector (modern services, manufacturing or substitutable services).[27]

Stemming from the above interpretation, we generate polarization graphs in the model as follows. First, using the 1980 equilibrium, we rank agents working in a given sector according to their ability in that sector ($a^{1+e\zeta}$). We then create equally sized bins of workers with similar ability along the sector skill distribution. Thus, each bin is the model counterpart of an occupation in a given sector in the data. We distinguish by sector because, for instance, the ability level of a worker in manufacturing cannot directly compared with the ability level of a worker in modern services so we construct bins within each sector.[28] Next, we compute the average wage in each bin in the 1980 equilibrium and rank all bins from the three market

[26] Bandwidth 0.8 with 100 observations.

[27] In doing this, we depart from the common assumption in the literature on *routinization*, which considers an occupation as a set of tasks to be performed.

[28] This drives the same choice for the data, as discussed above. That is, in the data we consider "secretary" in manufacturing and "secretary" in modern services as two distinct occupations.

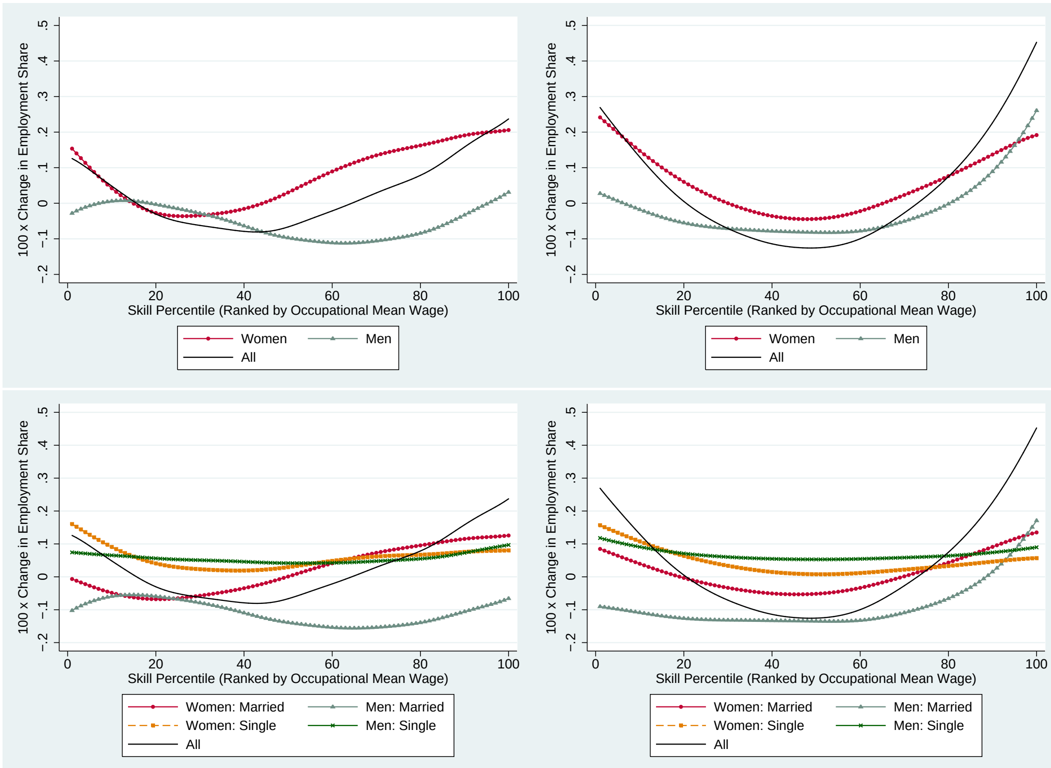

Figure 5: Job polarization in the data (left) and in the model (right). First row: gender; second row: marital status and gender.

sectors into a unique aggregate classification according to the wage. Finally, we construct occupation percentiles using hours worked by all workers in each bin in the 1980 equilibrium. This ranking is then kept for the 2008 equilibrium and the polarization graph is generated as in the data by measuring the change in employment share of each 1980 percentile and using the same locally weighted smoothing regression as for the data.

6 Results

Figure 5 presents the comparison between polarization in the data and the respective polarization graphs generated by the model.[29] 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

[29] To compare the model with the data in terms of aggregate results, in Appendix C we report a table with the aggregate change in employment shares by gender, sector and education level and another table with changes in the employment shares by deciles for women, men and the overall population.

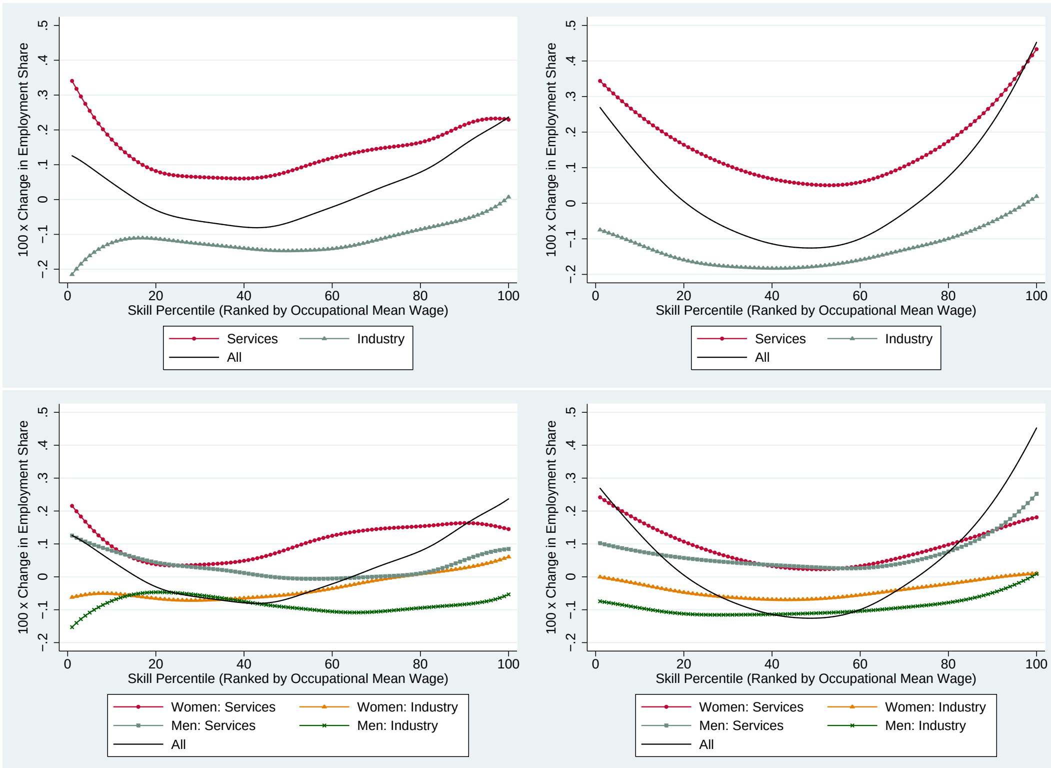

Figure 6: Job polarization in the data (left) and in the model (right). First row: sectors; second row: sectors and gender.

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

In table 6 we report the quadratic fit of equation (3) for the model and compare it with the data. The estimated coefficients suggest that the model produced too much convexity with respect to the data. This is true for the three groups All, Males and Female. However, we are interested here in assessing the model's ability to reproduce *relative* polarization of the two gender. By taking the ratio of the coefficient of one gender to the one for the aggregate, we obtain for males a value of 0.60 (0.583/0.978) in the data and of 0.63 (1.318/2.102) in the model. For females, the corresponding figures are 2.16 (2.114/0.978) in the data and of 1.85 (3.883/2.102) in the model. Thus, according to this metric, the model reproduces fairly well the higher polarization observed in the data for women relative to men.

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 distri-

Table 6: *Quadratic fit of the data and the model*

	Data			Model		
	All (A) 1980-2008	Males (M) 1980-2008	Females (F) 1980-2008	All (A) 1980-2008	Males (M) 1980-2008	Females (F) 1980-2008
Rankˆ2	0.978*	0.583	2.114**	2.102***	1.318***	3.883***
S.E.	(0.530)	(0.587)	(0.665)	(0.265)	(0.263)	(0.273)

bution. 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 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 flatter 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 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 in-

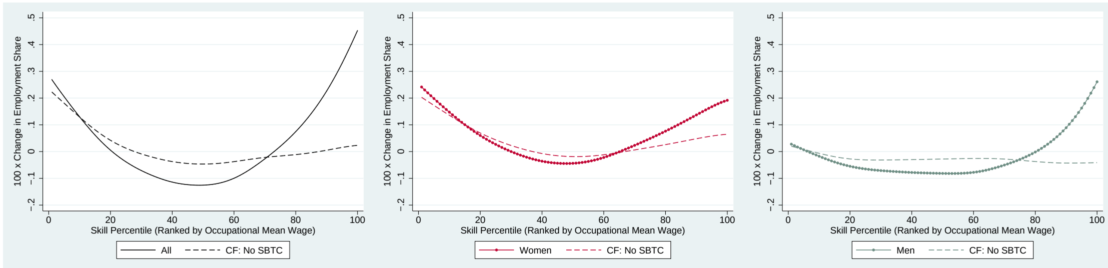

Figure 7: Counterfactual: No skill-biased technological change

crease 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.[30]

By estimating equation (3) for this counterfactual exercise, we find that the value of the coefficient of the quadratic term for aggregate employment polarization becomes 0.785, which implies a ratio with respect to the benchmark case of 0.37. This number confirms that skill-biased technological change has a preminent role in explaining aggregate employment polarization. According to this metric, by closing this channel, aggregate polarization is reduced by 63%.

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. This is confirmed when estimating equation (3) for this counterfactual. We obtain values for the coefficient of the Rankˆ2 term of 1.951 for the aggregate corresponding to a ratio with respect to the benchmark of 0.93, which confirms the fact that gender biased technological change has a negligible impact on aggregate employment polarization.

7 Predicting the Pre-Polarization Era

The results 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

[30]We 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.

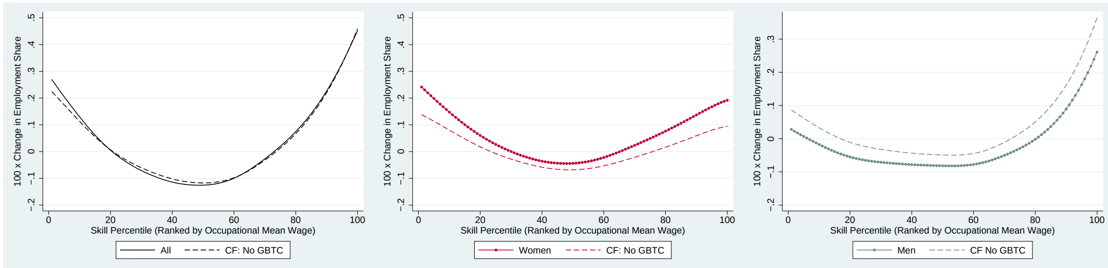

Figure 8: Counterfactual: No gender-biased technological change

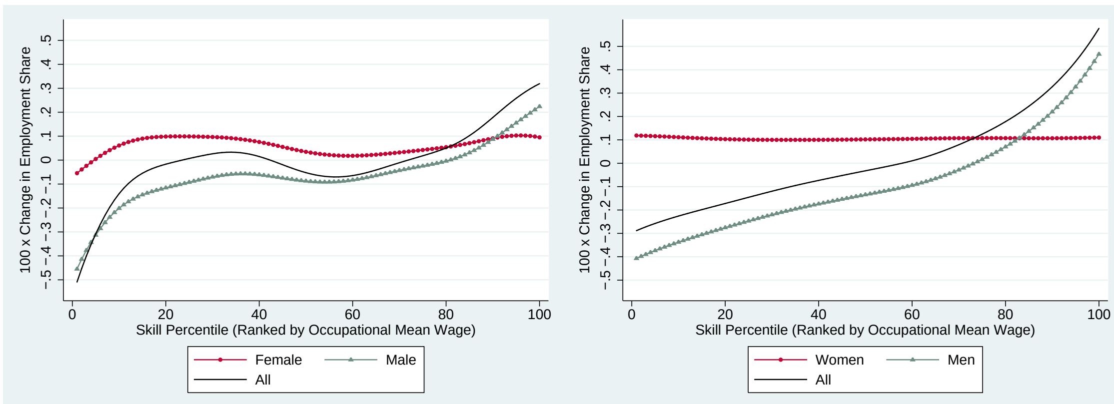

Figure 9: Job polarization in the data (left) and in the model (right) during the period 1960-1980.

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). Thus we set an average growth rate for skill-biased technological change of -0.0066 and an average growth rate for gender-biased technological change of 0.0064 during the 1960-1980 period in our experiment.[31] For home labor productivity we follow Bridgman (2016), who measures an average growth rate of 2.5% for the pre-1980 period. We assume that labor productivity in

[31] We thank Kjetil Storesletten for providing the numbers. Heathcote, Storesletten, and Violante (2010) compute the implied skill-biased and gender-biased technological change for the period 1966-2005. We use their numbers for the 1966-1980 period to compute an average growth rate that we apply to the 1960-1980 period. To be consistent with our benchmark calibration we also need to scale the numbers in Heathcote, Storesletten, and Violante (2010) by an appropriate factor. For convenience, we describe how to compute this factor in the next section. Note that measures of technological change in Heathcote, Storesletten, and Violante (2010) are appropriate in our setting because we employ the same production function. 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.

the three market sectors displays the same growth rate as in the 1980-2008 period.[32] Finally, we also match the demographic trends from 1960 to 1980.[33]

Figure 9 presents the comparison between employment polarization in the data and the corresponding pattern generated by the model for the period 1960-1980.[34] As the black 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.

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 6. 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.

The counterfactual exercises in section 6 suggest that a key role in shaping employment polarization in the 1980-2008 period is skill-biased technological change. To study whether this factor per-se can explain the absence of employment polarization before 1980, we run a counterfactual for the 1960-1980 period where all trends *but* skill-biased technological change are those of the 1980-2008 period. Put it differently, we assume that all technological change in the model evolves at a constant rate (the average one implied by our calibration for the 1980-2008 period), except skill-biased technological change, which accelerates between the pre- and post-1980 periods.[35] Figure 11 shows that this unique difference makes changes in employment shares flatter for women, men and, consequently, for the overall economy. Not surprisingly, the effect of the counterfactual is similar to that in Figure 7, in which skill-biased technological change is set to zero. The exercise thus suggests that the different growth rate of skill-biased technological change between the pre- and the post-1980 period is

[32] Note 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.

[33] See the discussion for the benchmark case in subsection 5.2.

[34] Computing employment polarization for the 1960-1980 period requires dealing with occupations that are not present in both years. In Appendix B we discuss different methodologies to address this issue. The main results in this section are maintained regardless of the methodology.

[35] Strictly speaking, it switches from negative to positive growth.

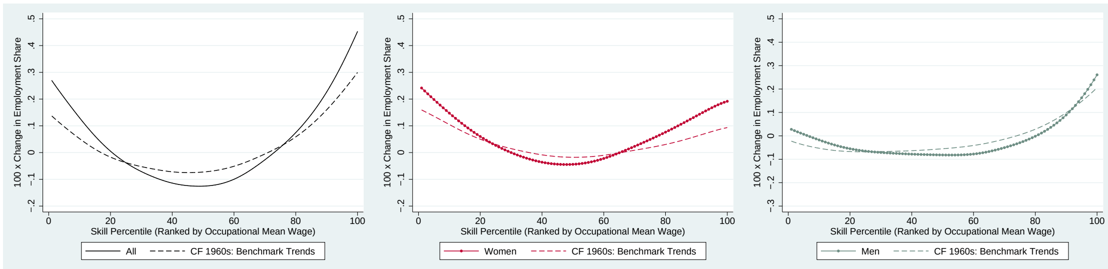

Figure 10: Job-polarization in the model during the period 1960-1980 when using the 1980-2008 exogenous trends.

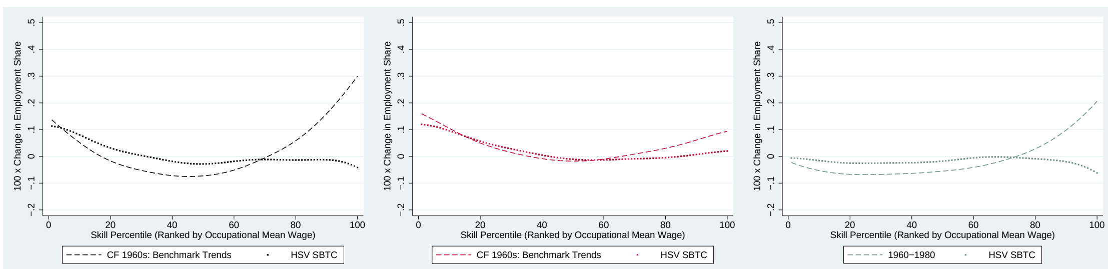

Figure 11: Job-polarization in the model during the period 1960-1980 when using the 1980-2008 exogenous trends. Counterfactual with 1960-1980 SBTC.

a first order determinant of the occurrence of employment polarization in the latest period. The remaining difference between Figure 9 and Figure 11 is given by the combined effect of gender-biased technological change and home productivity.

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 5 in subsection 3.1 and we report it for convenience also 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 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 in reproducing the observed changes across decades. To do this, we feed the model with decade specific measures of skill-biased and gender-biased technological change. As in the previous section, to compute these measures we use the yearly time-series of skill-biased and gender-biased technological change

Table 7: *Decade-specific exogenous trends*: SBTC and GBTC

	SBTC	GBTC
1980-1990	0.015	0.010
1990-2000	0.014	-0.001
2000-2008	0.008	0.007

between 1980 and 2008 derived in Heathcote, Storesletten, and Violante (2010). First, we use these numbers to compute the average growth for each decade (1980-90, 1990-2000, 2000-2008). Next, to be consistent with the total growth over the 1980-2008 period implied by our benchmark calibration (γ_ϕ and γ_φ in table 3) for each type of technological change, we multiply each decade specific average by a scaling factor. This is given by the ratio between total growth over the 1980-2008 period of each type of technological change in our calibration and that in Heathcote, Storesletten, and Violante (2010).[36] The values we obtain are summarized in table 7. Finally, to perform our quantitative exercise we adjust the average cost of education by decade to match the fraction of educated females and males in each decade.[37]

The right panel of Figure 12 reports the behavior of the model.[38] As in the data, the model can reproduce a tilt in the three lines, with no increase at the bottom in the 1980s, an increase at the bottom and at the top in the 1990s and a large increase at the bottom in the 2000s. However, the model produces an increase at the very top of the distribution in the 2000s, something that is absent in the data.

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.

[36] In this way, we tie our hands by preserving the relative growth across decades as measured by Heathcote, Storesletten, and Violante (2010) while at the same time we retain the total growth over the whole period implied by our benchmark calibration. Note that for consistency we applied the same scaling factor also in section 7 for the pre-1980 period, but the effect of the scaling factor on the results for that period is negligible.

[37] To compute skill-biased and gender-biased technological change, Heathcote, Storesletten, and Violante (2010) also allow the mean cost of education to vary over time. Thus, strictly speaking, their technology measures should be used only together with a time-varying cost of education, as we do here. In section 7 we do not change the cost of education in 1960 because the model performs well in replicating the fraction of educated: the shares of educated males and females in the model in 1960 are 0.0984 and 0.0600 compared with 0.1016 and 0.0597 in the data.

[38] We also report the comparison between model and data of the two gender across decades in Appendix D.

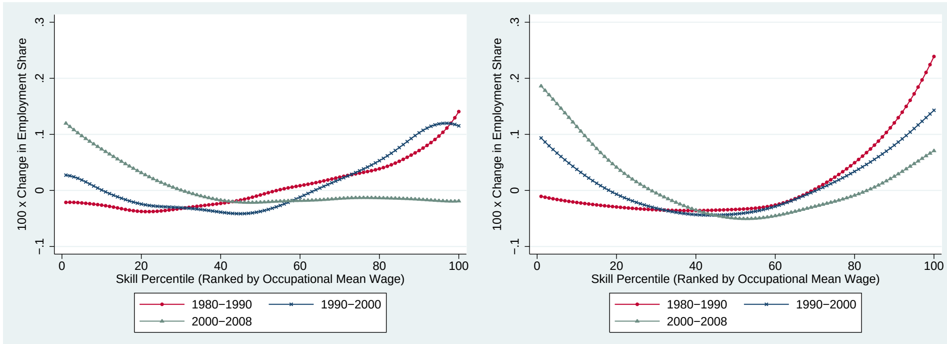

Figure 12: Job-Polarization by decade, 1980-2008. Data (Left) and Model (Right).

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 together with that of educated individuals (q-complementarity in production between educated and uneducated workers). In the model, the direct effect dominates in the first and second decade, while the indirect effects dominate in the last decade. The behavior of the model for the 2000s is consistent with the evidence discussed in Jorgenson, Ho, Stiroh, et al. (2005), p. 13, who find that the contribution to output growth of college-labor in the U.S. is substantially more important than that of non-college educated labor during the period 1977-2000, but that the sustained growth of the U.S. economy of the late 1990s allowed a large number of workers with low skills to obtain a job.

To study the direct and the indirect effects of skill-biased technological change we report, in Figure 13, counterfactual exercises for the three decades in which we set it to zero. Loosely speaking, removing skill-biased technological change should affect more the higher part of the distribution when the direct effects are quantitatively more important, and the lower part of the distribution if the indirect effects dominate. From the first row of Figure 13 it is clear that the tilting behavior across decades completely disappears at the aggregate level, once we remove skill-biased technological change. Also, the counterfactual confirms that skill-biased technological change has a time varying effect on changes in employment shares across the distribution. In the 1980-1990 period, by removing skill-biased technological change the increase of men shares at the top of the distribution completely disappears, while the effect on women at the top is similar but less substantial. 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

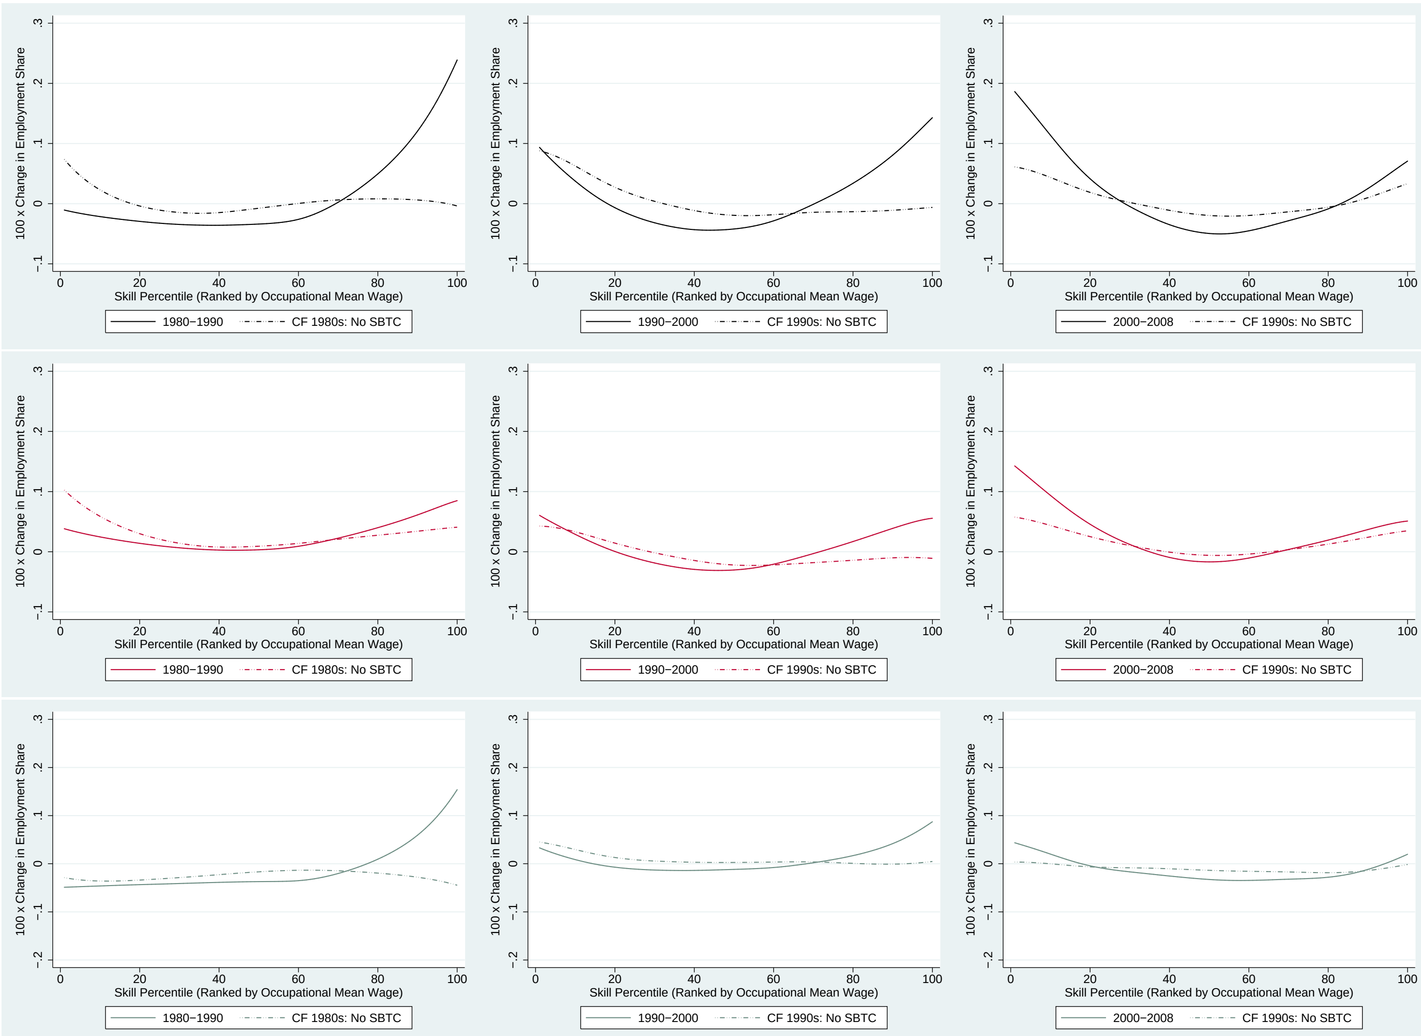

Figure 13: Job-Polarization by decade, 1980-2008. Counterfactual of skill-biased technological change.

removing skill-biased technological change reduces substantially the increase of employment shares at the top of the skill-distribution both for women and for men. Finally, the absence of skill-biased technological change during the 2000s removes the large increase at the bottom of the distribution observed in the data and generated in the model by women. This suggests that in this decade the indirect effects are the ones with quantitative relevance. These results rationalize the emergence of consumption spillovers during the late polarization era documented in Hazan and Zoabi (2015) for the U.S.

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 motivates the study of 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 construct 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.

Bárány, Z. L., and C. Siegel (2018): "Job polarization and structural change," *American Economic Journal: Macroeconomics*, 10(1), 57–89.

Beaudry, P., D. A. Green, and B. M. Sand (2016): "The Great Reversal in the Demand for Skill and Cognitive Tasks," *Journal of Labor Economics*, 34(S1), S199–S247.

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., and J. P. Kaboski (2012): "The Rise of the Service Economy," *The American Economic Review*, 102(6), 2540–69.

Buera, F. J., J. P. Kaboski, and R. Rogerson (2015): "Skill biased structural change," Discussion paper, National Bureau of Economic Research.

Buera, F. J., J. P. Kaboski, and M. Q. Zhao (2013): "The rise of services: the role of skills, scale, and female labor supply," 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.

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*.

EECKHOUT, J., R. PINHEIRO, AND K. SCHMIDHEINY (2014): "Spatial sorting," *Journal of Political Economy*, 122(3), 554–620.

FINDEISEN, S., AND D. SACHS (2015): "Designing efficient college and tax policies," .

GOOS, M., AND A. MANNING (2007): "Lousy and lovely jobs: The rising polarization of work in Britain," *The review of economics and statistics*, 89(1), 118–133.

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.

JORGENSON, D. W., M. S. HO, K. J. STIROH, ET AL. (2005): "Productivity, Volume 3: Information Technology and the American Growth Resurgence," *MIT Press Books*, 3.

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 (2017): "Does Home Production Drive Structural Transformation?," *American Economic Journal: Macroeconomics*, 9(3), 116–46.

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

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.

Sevinç, O. (2017): "Skill-biased technical change and Labor market polarization: the role of skill heterogeneity within occupations," .

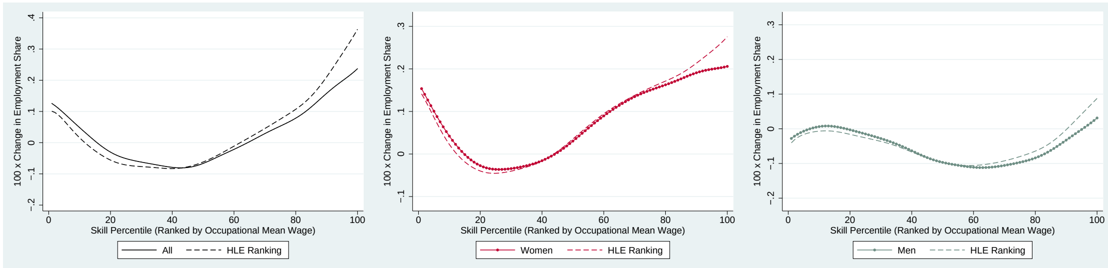

Figure 14: Data: Ranking Method. The dashed line is from Acemoglu and Autor (2011). The other line in each panel is our methodology as described in section 4.

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 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.[39] 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

[39]See subection 5.3.

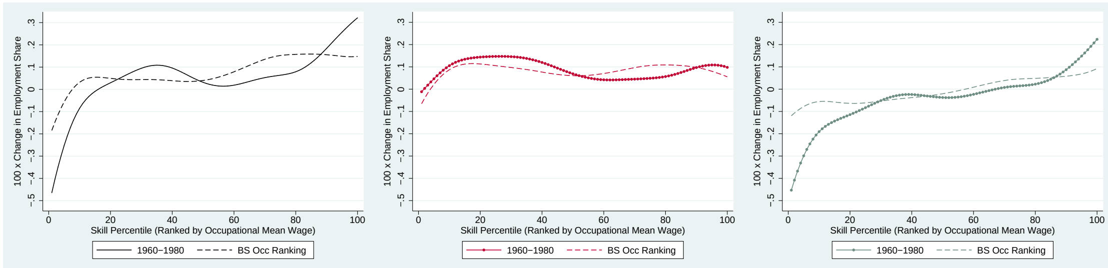

Figure 15: Different methodologies for the 1960-1980 period. The continuous line reports employment polarization as computed in Figure 9. The dashed line is constructed following the methodology in Bárány and Siegel (2018).

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 Bárány and Siegel (2018).[40] 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.[41] 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 Additional Tables

Table 8 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 between 1980 and 2008. Table 9 compares the change in hours worked by gender, education and sectors in the data and the benchmark model.

D Gender Behavior by Decade 1980-2008

Here we report the behavior of employment shares of women and men for the three decades 1980-2008 in the data and in the model. The comparison is reported in Figure 16. The tilting

[40] We thank Zsofia Barany and Christian Siegel for sharing the occupational classification codes.

[41] Note also that Bárány and Siegel (2018) drop occupations in agriculture while here we use them.

Table 8: *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

Table 9: *Percentage Change in hours worked across categories 1980-2008*

Category	% Change	
Women	Data	Model
Uned - ss services	0.02	0.04
Uned - manufacturing	-0.03	-0.05
Uned - ms services	0.11	0.07
Ed - ss services	0.02	0.03
Ed - manufacturing	0.01	-0.04
Ed - ms services	0.11	0.06
Men	Data	Model
Uned - ss services	0.02	0.01
Uned - manufacturing	-0.07	-0.14
Uned - ms services	0.02	0.02
Ed - ss services	0.01	0.02
Ed - manufacturing	-0.05	-0.11
Ed - ms services	0.04	0.05

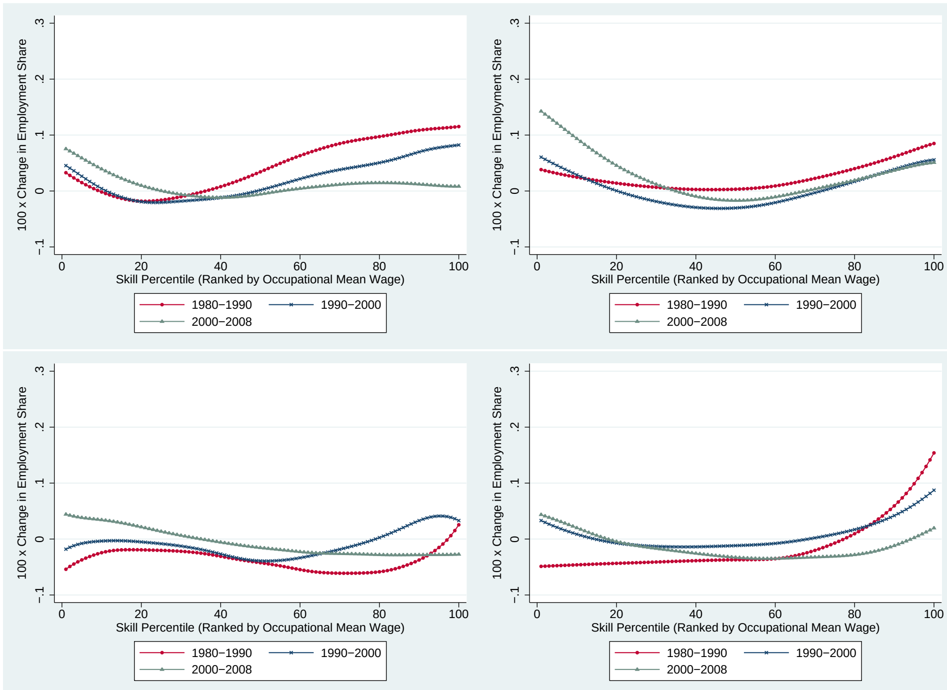

Figure 16: Job-Polarization by decade and gender,1980-2008. Data (Left) and Model (Right). First row: females; second row: males.

behavior across decades is apparent both for women and for men, although to a different extent. Thus, the model reproduces the tilting behavior across decades for the two gender.