

THE ROLE OF GENDER IN EMPLOYMENT POLARIZATION*BY FABIO CERINA,  ALESSIO MORO,  AND MICHELLE RENDALL¹ *University of Cagliari, Italy and CRENoS; University of Cagliari, Italy; Monash University, Australia and CEPR*

We document that U.S. employment polarization in the 1980–2017 period is largely generated by women. In addition, we provide evidence that the increase of employment shares at the bottom of the skill distribution is generated in market sectors producing services representing home production substitutes. We use a calibrated macroeconomic model to show that a rising skill premium induces high-skilled women to participate more in the labor market, reduce home working hours, and demand more home substitutes in the market, thus fostering a rise in employment shares at the bottom of the skill distribution. Counterfactual experiments suggest that without the large increase in the skill premium of high-skilled women, employment polarization would have been substantially reduced, and changes of employment shares at the bottom of the distribution would have been negative.

1. INTRODUCTION

The U.S. labor market experienced a sharp increase in the wage premium of highly educated agents since 1980 (Katz and Murphy, 1992; Autor et al., 2008). In that year, highly educated women still worked roughly half the time of highly educated men in the market (Ngai and Petrongolo, 2017), and were substantially more involved in home production activities with respect to their males counterparts (Aguiar and Hurst, 2007). In this context, the rising skill premium had nontrivial implications for the dynamics of hours worked in the market by gender and skill level that were observed in the decades after 1980 (Heathcote et al., 2010). In this article, we take a step further in this direction, and investigate how the differential evolution of hours worked by gender and skill level, triggered by the rising skill premium, might have contributed to shape the well-documented phenomenon of employment polarization in the United States.

Employment polarization refers to the contemporaneous increase of employment shares at the bottom and at the top of the U.S. skill distribution, with a consequent decline of shares in

*Manuscript received December 2020; revised May 2021.

¹ We thank the editor Dirk Krueger, several referees, David Autor, Vasco Carvalho, Luigi Guiso, Nezh Guner, Moshe Hazan, Esteban Jaimovich, Vahagn Jerbhasian, Omar Licandro, Rachel Ngai, Galo Nuño, Barbara Petrongolo, Chris Pissarides, Xavier Raurich, Andrew Rendall, Christian Siegel, Kjetil Storesletten, Satoshi Tanaka, Marc Teigner, Montserrat Vilalta, Carlo Valdes, Fabrizio Zilibotti and seminar participants at LSE, UB, Bamberg, Cambridge, Copenhagen, Monash, National University of Singapore, New South Wales, Zurich, OECD, 2017 Minnesota Workshop in Macroeconomic Theory, CEPR/Bank of Italy “Labour market participation: Forces at work and policy challenges” (Rome), CEPR Macroeconomics and Growth Programme (LBS), Barcelona GSE Summer Forum 2016, II MadMac Conference in Growth and Development (Madrid), V Workshop on Structural Transformation and Macroeconomic Dynamics (Kent), 2016 RIDGE December Forums in Buenos Aires, Montevideo and Rio de Janeiro, 3rd Workshop of the Australasian Macroeconomic Society (Brisbane), SIE (Bocconi), SAEe (Bilbao), Second Marco Fanno Alumni Workshop (Milan), and VII Workshop IIBEO (Alghero) for the useful comments. Moro thanks the Region of Sardinia for financial support, research grant F72F16002980002. Finally, we thank Sara Pau for excellent research assistance. The usual disclaimers apply. Please address correspondence to: Alessio Moro, Department of Economics and Business, University of Cagliari, Viale Sant’Ignazio 17, 09123 Cagliari, Italy. E-mail: amoro@unica.it.

the middle of the distribution. How then could an increase in the market premium for individuals located at the top of the skill distribution trigger such a nonmonotone aggregate pattern of employment shares? To show how, we consider the following scenario. A rise in the education premium 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. By increasing the demand for these *substitutable services*, the agent fosters an increase in employment shares of low-skilled individuals, who represent the main pool of workers in that market sector (Autor and Dorn, 2013; Cerina et al., 2019). The pull of demand for such type of labor creates incentives for low-skilled women working at home to also participate more in the market, thus increasing employment shares at the bottom of the skill distribution. 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. Although the same mechanism would in principle also apply to men, this category already displays high market participation rates at the beginning of the 1980s, so that the female response to more favorable market condition is likely to have been substantially stronger both at the top and at the bottom of the skill distribution.

Although employment polarization in the United States has been extensively documented, less well known in the literature is the behavior of employment shares along the skill distribution when distinguishing by gender. In this article, we provide evidence that employment polarization in the United States has been substantially stronger for women. This demographic group is responsible for most of the rise at the top and a substantial rise at the bottom of the skill distribution. In contrast, while changes in employment shares of men also display a mild U-shape, they are more homogeneous along the skill distribution. Motivated by the intuition discussed above, we complement this empirical finding with evidence on employment polarization by sectors of economic activity. More specifically, we split the economy in two broad sectors, the first producing services which are substitutes to home production and the second including the rest of the economy. This classification allows us to document that the increase in the employment shares at the bottom of the skill distribution is almost entirely driven by the first sector, whereas the upper tail is entirely generated by the second. Jointly taken, the evidence by gender and sector is consistent with the idea that the asymmetric entry of women in the labor market since the beginning of the 1980s is a key driver of employment polarization in the United States through the mechanism described before, in which high-skilled women entering the labor market and abandoning home production create a demand pull for low-skilled labor.

To investigate whether the mechanisms discussed above display quantitative relevance in generating employment polarization, we use a static calibrated macroeconomic model. The model is designed to capture the potential general equilibrium effects that can emerge when women increase participation in the labor market in a period of rising skill premium. It features exogenous technological change and an exogenously changing demographic structure. These factors allow us to ensure that the model matches the main *aggregate* features of the U.S. labor market, including the average skill premium and the average gender wage gap. In this environment, we solve for the endogenous equilibrium allocation of hours by agents with heterogeneous skills, who take as given the demographic structure and technological change at each point in time and so determine the *distributional* features of the labor market in the model. From the changes in hours worked of agents between two different points in time, we compute the changes in employment shares along the skill distribution and compare them to employment polarization patterns in the data.

Our theory rests on three building blocks: (i) a gender dimension; (ii) an endogenous home/market labor supply; and (iii) a multisector environment. We assume that there are two market sectors and a home sector. The two market sectors are *substitutable services* (services with a home produced counterpart) and *nonsubstitutable services* (the rest of the economy).

Each market sector is given by a competitive representative firm that employs four types of labor: highly educated males, highly educated females, less educated males, and less educated females. Although wages of each type of labor are endogenously determined in equilibrium, the forces driving the average changes in the skill premium and the gender wage gap are modeled, as in Heathcote et al. (2010), as exogenous shifts in the factor loads in production technologies. We also follow these authors in labeling these shifts *skill-biased technological change* (SBTC) and *gender-biased technological change* (GBTC).

The quantitative strategy is to evaluate the performance of the model in replicating the main facts of employment polarization when comparing two equilibria, meant to represent the U.S. economy in the years 1980 and 2017. The two equilibria differ exogenously in the demographic structure and the level of skill-biased technology and of gender-biased technology. All parameters in the model are pinned down by 1980 aggregate targets, except for the two types of technological change, which are disciplined by the increase in the aggregate skill premium and the decrease in the gender wage gap between the two years. Given these differences, the model endogenously generates heterogeneous changes of employment shares along the skill distribution between the two equilibria, which replicate fairly well employment polarization by gender and by sector as observed in the data. In particular, the model produces higher employment polarization for women than for men and suggests that, consistent with the data, the increase at the bottom of the distribution is mainly driven by substitutable services whereas the upward shift at the top of is entirely driven by the rest of the economy.

By running a counterfactual experiment in which we remove SBTC, the U-shape of employment polarization in the model is reduced substantially, and the change in employment shares at the bottom of the distribution turns from positive to negative. As SBTC affects both highly educated males and females, we also run a counterfactual in which we tax wages of highly educated women by 36% in the 2017 equilibrium, that is, roughly the increase in wages experienced by this group between 1980 and 2017 in the data. In this case, the increase of employment shares at the bottom and at the top of the skill distribution between the two equilibria is substantially reduced with respect to the benchmark case, confirming that the wage-induced higher participation of highly educated woman in the labor market is a key driver of employment polarization. In fact, the same tax to wages of less educated women, highly educated men, or less educated men, cannot produce a simultaneous decline of employment shares at the top and at the bottom of the skill distribution.

Finally, we use the calibrated model to “predict” changes in employment shares from 1960 to 1980. During this period, we find no employment polarization in the data and, as discussed above, there is no clear trend in the skill premium. To run the exercise, we feed exogenous growth rates in SBTC and GBTC previously measured in the literature for the 1960–80 period, the former being substantially lower than that of our benchmark calibration for 1980–2017, contrary to the latter whose value is similar to the 1980–2017 period.¹ The model broadly reproduces the flat behavior of changes in women’s employment shares and the monotone behavior of men’s employment shares. Thus, the theory can reconcile the evidence on changes in employment shares before and after 1980. Also, by counterfactually introducing a 36% subsidy to highly educated women in 1980 (in order to let them enjoy in the 1980 equilibrium the actual wage of 2017), the model predicts a substantial increase in employment shares both at the top and at the bottom of the skill distribution with respect to the 1960–80 benchmark. The main takeaway is that, during the 1960–80 period, without an increasing skill premium, there was no specific incentive for high-skilled women to increase participation in the labor market more than women of other skill levels. Without a higher participation (and wages) of high-skilled women, there were no pull of demand for substitutable services and no increase of employment shares at the bottom of the skill distribution.

¹ Values for the 1960–80 period are taken from Heathcote et al. (2010).

The remainder of the article is as follows: Section 2 discusses the related literature; Section 3 establishes some facts on employment polarization in the United States that have not been considered in previous literature, by focusing on a gender and sectoral dimension; Section 4 presents the model; Section 5 discusses the quantitative strategy; Section 6 provides the benchmark results; and Section 7 the counterfactual analysis. Finally, Section 8 concludes.

2. RELATED WORK

Our theory builds on the canonical model of SBTC (Katz and Murphy, 1992). While in its simple version the canonical model of SBTC cannot account for the polarization of employment observed in the United States (Acemoglu and Autor, 2011), we show that by extending it with gender, a multisector environment and home production activity, it can produce a U-shape of changes in employment shares mainly triggered by the effect of SBTC on a particular demographic group, that of high-skilled women. Our results challenge the conventional view suggesting that employment polarization cannot be addressed without an explicit distinction between skills and tasks in the model (Acemoglu and Autor, 2011, p. 1117). In the literature considering the differential effect that technological change has on occupations with different *task content*, the work which is most closely related to ours is Autor and Dorn (2013). They show that employment polarization can be qualitatively accounted for by *routinization*, a particular type of technological process which drives workers out of middle-skill (routinary) occupations and into low-skilled (nonroutine/manual) occupations. In Appendix B, we provide evidence that while employment shares of routine occupations decline in the middle, they also increase substantially at the bottom of the skill distribution, a finding that casts some doubts on the ability of the routinization hypothesis to account for changes of employment shares of low-skilled individuals. Our sectors classification (substitutable and nonsubstitutable services) and our theory suggest that employment changes at the bottom of the skill distribution are accounted mostly by the sector producing substitutable services. We find that within this sector, both types of occupations contribute to employment share growth of low-skilled workers: for instance, for the first percentile of the skill distribution, around 2/3 of employment growth is accounted for by nonroutine and 1/3 by routine occupations. This supports the view that it is the nature of production (substitutable services), instead of the task content of the occupations (routine or nonroutine), what matters for the change of employment shares at the bottom of the skill distribution.²

The period after 1980 in the United States is characterized by unprecedented changes in female labor markets, which suggest a potential role for women in generating employment polarization. Mulligan and Rubinstein (2008) show that females relative wages start increasing at the end of the 1970s and relate this to the increasing overall wage inequality, which induces more skilled women to invest in human capital, thus affecting the composition of the female workforce. Albanesi (2020) shows that female labor force participation increases until the end of the 1990s, then flattens out and declines after the Great Recession. She finds a key role of these pattern in accounting for both the cyclicity of aggregate hours since 1980 and in the jobless recoveries starting in the early 1990s. Ngai and Petrongolo (2017) and Rendall (2017) show how, in multisector models with home production, the process of “marketization” implies that women progressively abandon home production to work in the market, and especially in the services sector, as stressed by Lee and Wolpin (2006). Our environment also builds on the insights in Buera and Kaboski (2012), who provide a theory predicting that

² Tüzemen and Willis (2013) document that a large fraction of polarization occurs within broad sectors of economic activity rather than across sectors. Although our categorization of sectors does not coincide with Tüzemen and Willis (2013), their Chart 4 shows that the sectors which experience the largest increase in the fraction of low-skilled jobs is their “others” sector, which includes our definition of the substitutable services sector. In the other three sectors they consider (Construction, Manufacturing, and Health and Education) the change in the fraction of low-skilled jobs is either very small, or negative.

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 et al. (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. We build on these contributions by linking SBTC to a gender and sector dimension to study employment polarization.³

Our theory also builds on the intuition discussed in Manning (2004) and Mazzolari and Ragusa (2013). The idea is that of *consumption spillovers*, that is, the fact that an increase in high-skill workers in the market, with a high opportunity cost of working at home, also increases the demand for services in the market that have a home counterpart. Cerina et al. (2019) show that the degree of employment polarization in the United States increases with city size, as larger cities attracted both more high-skilled workers, and more low-skilled individuals, the latter attracted by consumption spillovers generated by the former. Also closely related to this idea, and with a gender dimension to it, is Hazan and Zoabi (2015), who argue that the increase in income inequality over the last 30 years in the United States created a group of women (the most skilled) who can afford services that are substitutable to home production (in particular child care), and another one which supplies these services. As a result, and in contrast with the evidence from earlier periods, which shows a monotonically declining pattern, fertility displays a U-shape by education level during the 2000s. Interestingly, this period coincides with the large increase of employment shares at the bottom of the distribution observed in U.S. data.

3. EMPIRICAL EVIDENCE

The first part of this section is dedicated to reporting the main facts on employment polarization in the United States that our theory is meant to account for. In the second part, we present microevidence on hours worked and consumption of different demographic categories, to provide additional empirical support for the mechanisms produced by our model.

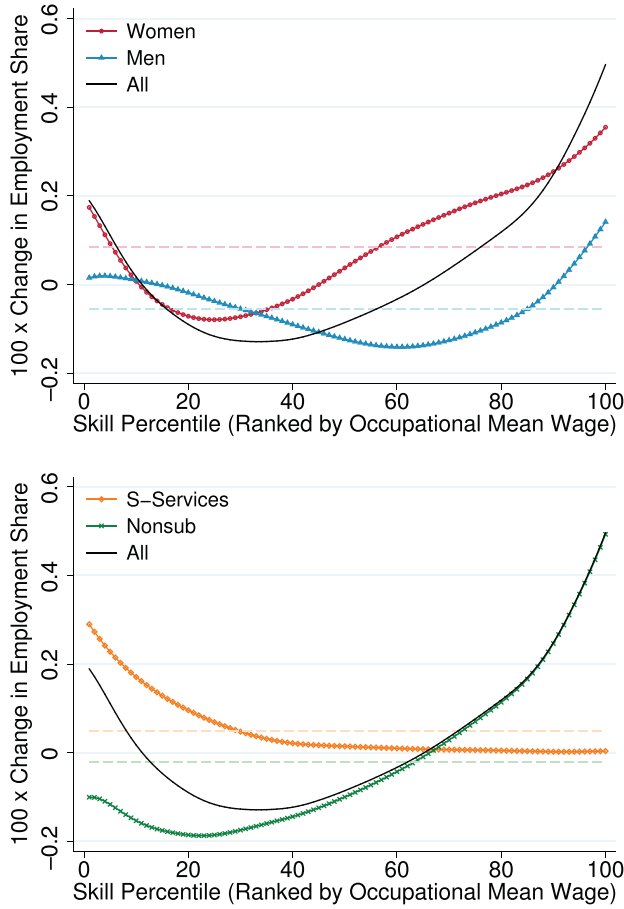
3.1. *Employment Polarization.* The reference period is 1980–2017. We first partition the economy into two sectors, *substitutable services* (services with a home produced counterpart) and *nonsubstitutable services* (the rest of the economy).⁴ Within each sector we compute, for the year 1980, the average wage of each occupation at the three-digit level. Then, we rank these occupations from the two sectors into a unique classification according to their average wage, and construct percentiles of employment by using hours worked of each occupation and Census weights. By keeping the same ranking of occupations for 2017, we construct employment polarization graphs by measuring the change in employment shares of each 1980 percentile and then applying a locally weighted smoothing regression.⁵

The resulting overall polarization for the U.S. economy is reported by the solid black line in the two panels of Figure 1, which displays the well-documented U-shaped pattern. In addition to the overall pattern, there are two partitions of the population we are interested in: by

³ Black and Spitz-Oener (2010) suggest that there is some evidence of employment polarization by gender in Germany and they refer to the WP version Black and Spitz-Oener (2007). However, in contrast with our findings for the United States there is no evidence of a U-shape of employment polarization for Germany, as there is no increase in employment shares at the bottom of the skill distribution for either women or men (figure 3 in Black and Spitz-Oener, 2007). This lack of polarization is in line with other studies for Germany (Rendall and Weiss, 2016).

⁴ To define substitutable services, we use the procedure in Moro et al. (2017). See Appendix A for details.

⁵ We adopt the same methodology as in Acemoglu and Autor (2011), only adapted to consider that we have two sectors. In the data, certain occupations are in both sectors (e.g., secretaries), whereas others are only in one (professors). In this case, in our method we have three occupations and wage rates in 1980 whereas Acemoglu and Autor (2011) computes an average wage for secretaries in the U.S. economy in 1980 and one for professors. So, instead of three occupations and wage rates in 1980, they have two. Besides that, the two methods are identical. Appendix C shows that the resulting differences between the two methods in the polarization graphs are very minor.



NOTES: The vertical sum of the colored lines gives the black line. Horizontal lines measure the average change in employment shares of the corresponding group of interest. Sources of data are the Census IPUMS for 1980 and Census American Community Survey for 2017. Occupation and earnings measures refer to prior year’s employment. The figure plots log changes in employment shares by 1980 occupational skill percentile rank using a locally weighted smoothing regression (bandwidth 0.8 with 100 observations). Skill percentiles are measured as the employment-weighted percentile rank of an occupation’s mean log wage in the Census IPUMS 1980 (5% extract). The mean log wage of each occupation is obtained from workers’ yearly working hours multiplied by Census weights. We use consistent occupation codes for 1980 and 2017 from Acemoglu and Autor (2011).

FIGURE 1

EMPLOYMENT POLARIZATION BY GENDER AND SECTOR FOR THE 1980–2017 PERIOD

gender, and by sector. Up to approximation due to smoothing, the aggregate polarization line computes, for each percentile i , the following difference:

$$(1) \quad \frac{H_{i,2017}}{H_{2017}} - \frac{H_{i,1980}}{H_{1980}},$$

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 . We then consider the following decomposition of (1):

$$(2) \quad \frac{H_{i,2017}}{H_{2017}} - \frac{H_{i,1980}}{H_{1980}} = \left(\frac{H_{i,2017}^f}{H_{2017}} - \frac{H_{i,1980}^f}{H_{1980}} \right) + \left(\frac{H_{i,2017}^m}{H_{2017}} - \frac{H_{i,1980}^m}{H_{1980}} \right),$$

where $H_{i,t}^f$ is total hours worked by women in percentile i in year t , and $H_{i,t}^m$ is the corresponding measure for men. The equality in Equation (2) follows from the fact that total hours in percentile i in year t are given by female plus male hours, $H_{i,t} = H_{i,t}^f + H_{i,t}^m$. The first term on the right-hand side gives the red dots line (females) in the top panel of Figure 1 whereas the second term provides the blue triangles line (males) in the same panel. We also apply Equation (2) to the two sectors instead of the two gender, by replacing $H_{i,t}^f$ and $H_{i,t}^m$ with $H_{i,t}^s$, total hours worked in the substitutable sector in percentile i in year t , and $H_{i,t}^n$, which is the corresponding measure for the nonsubstitutable sector. The two terms on the right-hand side of Equation (2) so constructed give the yellow diamonds (substitutables) and the green crosses lines (nonsubstitutables) in the second panel of Figure 1.⁶ We highlight the following empirical observations:

Empirical observation #1. Gender (first panel of Figure 1):

- Relative to men, women account for a larger fraction of the increase of employment shares at the top and at the bottom of the skill distribution.

Empirical observation(s) #2. Sectors (second panel of Figure 1):

- Nonsubstitutable services account for the increase of employment shares at the top of the skill distribution.
- Substitutable services account for the increase of employment shares at the bottom of the skill distribution.

Empirical observation #1 suggests a key role of women in generating employment polarization. This demographic group accounts for most of the rise at the top and for a rise at the bottom of the skill distribution which coincides with the aggregate black line. In contrast, while changes in employment shares of men also display a mild U-shape, they are more homogeneous along the skill distribution, and they show a change close to zero at the bottom of the skill distribution. Empirical observation #2 suggests that without the substitutable services sector there would be no employment polarization in the economy. Although some U-shape is retained by the green crosses line in the bottom panel of Figure 1, employment shares at the bottom of the distribution display negative changes, and those at the top positive ones. In fact, the positive change of employment shares at the bottom of the skill distribution emerges only in the substitutable services sector. This is consistent with the conjecture that the increase in market hours of high-skilled women and the corresponding reduction of home production during the period considered increases demand, and so employment, in the sector producing substitutable services. Observations #1 and #2 represent the main empirical facts that we require our theory to account for.⁷

3.2. *Microevidence.* In this article, we argue that the acceleration of the skill premium since the beginning of the 1980s is a key driver of employment polarization in the United States during the 1980–2017 period. The main conjecture is that the increase in the skill premium during this period induced higher labor market participation of high-skilled women, who formerly worked mostly at home (thereby increasing the employment shares in high-skilled occupations) and indirectly increased the demand for market substitutes of home production (thereby increasing the employment shares in low-skilled occupations). This mechanism should imply that skilled women, more than any other demographic category, responded to higher relative wages by supplying more market hours and demanding more substitutable services. We thus devote this section to discuss whether the microevidence on hours worked in

⁶ Note that, as for the overall polarization line, we use locally weighted smoothing regressions (bandwidth 0.8 with 100 observations) to plot the two terms on the right-hand side of Equation (2).

⁷ In addition to the evidence presented here, in Appendix C we discuss measures supporting our claim that changes in employment shares of women are more polarized than those of men and also provide additional evidence on the dynamics of employment shares by gender and broad groups of Census occupations.

TABLE 1
WEEKLY MARKET AND HOME HOURS BY DEMOGRAPHIC CATEGORY

	Women		Men	
Weekly Market H	H_{1980}	$H_{2017}-H_{1980}$	H_{1980}	$H_{2017}-H_{1980}$
No college	17.61	5.04	35.01	-3.20
College	22.86	8.16	40.45	-2.23
No college married	15.66	6.17	36.60	-0.88
College married	19.66	8.81	41.47	-1.48
	Women		Men	
Weekly Home H	H_{1975}	$H_{2017}-H_{1975}$	H_{1975}	$H_{2017}-H_{1975}$
No college	23.87	-6.61	6.42	3.69
College	21.94	-7.38	6.36	2.94
No college married	25.51	-6.33	6.28	4.89
College married	24.61	-8.28	6.46	3.62

NOTES: Data for working hours are from 1980 Census and 2017 ACS whereas data for home hours are drawn from the TUESA for 1975 and ATUS for 2017. Details of the data treatment are given in Appendix E.

the market and at home, and the consumption behavior across demographic categories is consistent with the mechanism we propose in this article. To do this, consistent with the model in Section 4 in which SBTC directly increases productivity of highly educated workers, we consider education as a proxy for skills.

In what follows, we use the 1980 Census and the 2017 American Community Survey (ACS) for market hours. For home hours, there are five time use surveys before the American Time Use Survey (ATUS), which starts in 2003, but none of them covers the year 1980.⁸ We use information from the 1975–76 Time Use in Economics and Social Accounts (TUESA) as our initial year of the home work sample, and the 2017 ATUS as the final year. An advantage of TUESA is that it collected diaries of multiple agents within the household, whereas all other surveys before the ATUS collected information for only one individual per household. This feature allows us to measure how female home hours worked change within the couple over the sample period. Finally, we use consumption expenditure data from the consumption expenditure survey (CEX) to study the change in the share of consumption devoted to substitutes of home production. Appendix E provides the details of the construction of the data for each data set we use in this section.

3.2.1. Hours worked in the market and home. The top panel of Table 1 shows that both women with and without a college degree increase hours worked in the market, but the former group displays an increase of 8.16 hours, compared to an increase of only 5.04 hours of the latter. In addition, married women increase hours worked in the market more than single women, both for the group with a college degree (8.81 hours of married women vs. 8.16 of the whole female population with a college degree) and for the group without (6.17 hours of married women vs. 5.04 of the whole female population without a college degree). Thus, the largest increase in market hours is displayed by highly educated married women. Changes in hours worked in the market by men are substantially smaller than for women, ranging from the decrease of 3.20 hours by the whole male population without a college degree, to a decrease of 0.88 hours of married men without a college degree. The bottom part of Table 1 shows the change in home hours worked for women and men. Married women with a college degree display a decline in home work of 8.28 hours, which suggests that this group is the one

⁸ See Aguiar et al. (2012) for a discussion of the five surveys.

TABLE 2
HOME HOURS WITHIN THE COUPLE

Weekly Home H	H_{1975}	Wife $H_{2017}-H_{1975}$	Husband $H_{2017}-H_{1975}$	Household $H_{2017}-H_{1975}$
No college married women	25.51	-6.33	4.74	-1.60
College married women	24.61	-8.28	4.34	-3.94

NOTES: Data are drawn from TUESA for 1975 and ATUS and CPS for 2017. Details of the data treatment are given in Appendix E.

experiencing the largest change in both market and home hours.⁹ All men categories instead display an increase in hours worked at home.¹⁰

3.2.2. Home hours within the couple. To put the difference between highly educated and less educated married women in Table 1 into perspective, we also investigate the change in hours worked at home within the couple. For 1975, the TUESA provides hours worked at home and the education level of the wife and the husband within the household. Thus, we can measure the average hours worked of husbands married to women of a certain education level. For 2017, we match the ATUS with the CPS to identify married male individuals by the education of the woman and, as for 1975 in TUESA, we compute the average hours worked of husbands married to women of a certain education level. More precisely, for each year we split the group of married women into highly educated and less educated, and for each group we measure the average hours worked by husbands of women in that group. Results are reported in Table 2. Less educated married women decrease home work by 6.33 hours between 1975 and 2017. Their representative husband increase home hours by 4.74, so that the representative household in which the women is without a college degree decreases home work by 1.60 hours. Instead, highly educated married women decrease home work by 8.28 hours between 1975 and 2017, whereas their representative husband increases hours worked at home by 4.34. This implies a decline of home work for the household of 3.94 hours, almost double that of households in which a women without a college degree is present. This evidence suggests that it is female home hours within the couple that fall, and that they fall most strongly at the top of the skill distribution, confirming that married highly educated women are the demographic group displaying the largest increase in market hours and the largest decline in home hours.

3.2.3. Expenditure data. As long as home productivity is stagnant, as documented in Bridgman (2016), and Moro et al. (2017) for the period 1980-2010 in the U.S., the decline in home hours translates into a decline in home production, with the lower production of home services potentially replaced by purchases of substitutable service in the market. According to the theory proposed, we should then observe different consumption patterns for substitutable services by different demographic categories: a larger increase in market hours and a larger decline in home hours should be associated with a larger increase in consumption of substitutable services. Evidence from the CEX data support this idea. Using harmonized CEX data from 1980 to 2003, we find that the increase in the share of consumption devoted to substitutes of home production increases more for married women with a college degree than for

⁹ Aguiar and Hurst (2007) document that less educated women and men experience greater increases in leisure with respect to their more educated counterparts. Except for married men, the evidence in Table 1 is consistent with this finding, as the total change in total hours worked (i.e., market plus home) is larger for highly educated agents than for less educated ones, which indicates a larger increase of leisure for the less educated individuals.

¹⁰ Note that the mechanism proposed in this article implies that hours at home for women of lower skill levels also decline, because of the increased demand for low-skilled services generated by larger market hours and wages, and fewer hours worked at home by highly educated women. Thus, the theory is also consistent with a reduction of home hours of low-skilled women (due to increasing demand for substitutable services) similar to the reduction of home hours of high-skilled women (due to SBTC).

TABLE 3
CONSUMPTION EXPENDITURE SHARES OF HOME PRODUCTION SUBSTITUTES

	Share Increase 1980–2003
No college single	–1.9%
College single	0%
No college married	1.1%
College married	1.8%

NOTES: Data are from the Consumption Expenditure Survey (CEX). Details of the data treatment are given in Appendix E.

other female categories.¹¹ Table 3 reports that the increases in the share of consumption devoted to substitutes of home production are: 1.8 percentage points for married women with a college degree; 1.1 percentage points for married women without a college degree; 0 for single women with a college degree; and –1.9 percentage points for single women without a college degree.

4. MODEL

The environment is static such that, given the fundamentals at time t , the equilibrium of the model is uniquely determined in that period. We are interested in two types of endogenous allocations by agents and in their role in generating employment polarization: (i) the allocation of time between market and home work; and (ii) the allocation of consumption between market and home. In the data, these choices depend on the specific demographic characteristics of agents, as documented in Subsection 3.2. For this reason, we posit an exogenous demographic structure in which each agent is identified by gender, marital status, and education level. In the quantitative part, we calibrate this structure for each year of interest, and study the endogenous allocation of time and consumption in this context.

4.1. Agents and Demographics. There are two different masses of agents in the economy, female and male agents, which can be of different size. We denote gender by the index $i = f, m$. Agents can work in one of two market sectors. The first produces services that are substitutes to home production. We label this as sector s . The second market sector produces the rest of goods and services in the economy, and we label it sector n (for nonsubstitutable). Agents of both gender are heterogeneous, such that each agent has a skill level to work in sector s and a skill level to work in sector n . Hence, an agent of gender i is identified by a couple of skill levels $\{a_s^i, a_n^i\}$, with $j = s, n$.

Agents display one of two possible education levels, which we indicate by e . An agent with $e = 1$ is highly educated, whereas an agent with $e = 0$ is less educated. High education magnifies the agent's ability such that her a_j^i becomes $(a_j^i)^{1+\zeta}$, with $\zeta > 0$. Thus, we define the *effective efficiency units* as \widehat{a}_j^e , where

$$\widehat{a}_j^e = \begin{cases} a_j^i, & \text{if } e = 0 \\ (a_j^i)^{1+\zeta} & \text{if } e = 1 \end{cases}$$

for $j = s, n$. Thus, high education magnifies the skill of the agent in both sectors.¹²

¹¹ We limit the analysis to the 1980–2003 period because of major changes to the CEX since 2003 (available at https://data.nber.org/data/ces_cbo.html) and the lack of properly harmonized data beyond that year (Goldenberg and Ryan, 2009).

¹² Along the text, we refer to ability, skill level, and efficiency units as synonyms. Thus, we also refer to effective ability, effective skill level, and effective efficiency units as such.

Finally, each agent is assigned a marital status $q = c, u$ (c for couples and u for unpaired) such that she can be in a single household or in a married household, the latter being formed by a female and a male agent. We thus have four categories according to gender and marital status: (i) a women in a married couple household fc ; (ii) a man in a married couple household mc ; (iii) a women in a single female household fu ; and (iv) a man in a single male household mu .

4.2. *Firms and Market Sectors.* There is a representative firm in each market sector $j=n, s$. Each firm has the production function

$$(3) \quad Y_j = A_j N_j,$$

where A_j is a scale parameter and

$$(4) \quad N_j = \left[\theta_j \left(\varphi_j N_j^{f1} + (1 - \varphi_j) N_j^{m1} \right)^{\frac{\eta_s - 1}{\eta_s}} + (1 - \theta_j) \left(\varphi_j N_j^{f0} + (1 - \varphi_j) N_j^{m0} \right)^{\frac{\eta_s - 1}{\eta_s}} \right]^{\frac{\eta_s}{\eta_s - 1}},$$

and N_j^{ie} is total labor efficiency units of agents of gender $i = m, f$ and education level $e = 0, 1$ in sector j . The production function follows Katz and Murphy (1992) in that different demographic groups are regarded as distinct labor inputs and Heathcote et al. (2010) in displaying (i) perfect substitutability across gender of equal education level; (ii) GBTC (through time variation in φ_j); (iii) imperfect substitutability across education levels ($\eta_s > 1$ being the elasticity of substitution between highly educated and less educated workers); and (iv) SBTC (through time variation in θ_j). We also follow Heathcote et al. (2010) in interpreting GBTC broadly as a glass ceiling against women that is weakened over time due to changes in social norms and antidiscrimination legislation, something that Goldin (2006) defines as the “quiet revolution.”

For the representative firm in sector j , an efficiency unit of gender i with education e costs w_j^{ie} . The firm maximizes profits given by

$$(5) \quad \pi_j = p_j Y_j - w_j^{f1} N_j^{f1} - w_j^{m1} N_j^{m1} - w_j^{f0} N_j^{f0} - w_j^{m0} N_j^{m0}$$

subject to (3) and (4).

First-order conditions imply

$$(6) \quad \frac{\theta_j \left(\varphi_j N_j^{f1} + (1 - \varphi_j) N_j^{m1} \right)^{-\frac{1}{\eta_s}}}{(1 - \theta_j) \left(\varphi_j N_j^{f0} + (1 - \varphi_j) N_j^{m0} \right)^{-\frac{1}{\eta_s}}} = \frac{w_j^{m1}}{w_j^{m0}},$$

$$(7) \quad \frac{\varphi_j}{1 - \varphi_j} = \frac{w_j^{fe}}{w_j^{me}}.$$

Equation (6) shows that, other conditions equal, SBTC, due to an increasing θ_j , raises the skill premium. Equation (7) shows that GBTC, in the form of growing φ_j , directly affects the wage ratio between males and females.¹³

¹³ 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 GBTC.

4.3. *Work and Consumption Allocations.* Given the exogenous demographic structure, agents choose endogenously: (a) in which sector to work and (b) their consumption and time allocation. Each agent is endowed with one unit of time, which she splits between work at home (l) and work in one of the two market sectors ($1 - l$).¹⁴ A unit of time of an agent of gender i , depending on her skills, the sector of employment and her education, can correspond to: (i) \widehat{a}_n^{ie} efficiency units of labor to production in sector n ; (ii) \widehat{a}_s^{ie} efficiency units of labor to production in sector s ; and (iii) 1 efficiency unit of labor to production in the home sector, which we label h .

At the level of sector choice the decision unit is the agent. The total wage an agent receives in the market depends on her skills, gender, education, and the sector she decides to work in. An agent of gender i with education e working in sector j receives a wage *per unit of efficiency* equal to w_j^{ie} . Thus, the agent, conditional on working a positive amount of time in the market ($1 - l > 0$), depending on her skill vector and her education, and taking as given the equilibrium (gender-specific) market wages per unit of efficiency in the two sectors, chooses to work in the market sector $j \in \{s, n\}$ that maximizes her efficiency wage $w_j^{ie} \widehat{a}_j^{ie}$.

At the level of consumption and time allocation, there are three kinds of *decision units* (i.e., households) $z = c, fu, mu$: households of type c in which a female and a male individual maximize a unitary utility function for the couple; households of type fu in which a single female maximizes her utility; and households of type mu in which a single male maximizes his utility.¹⁵ The utility function of a decision unit z is

$$(8) \quad U^z = \left((\omega)^{1/\sigma} (c_n^z)^{\frac{\sigma-1}{\sigma}} + (1 - \omega)^{1/\sigma} (c_t^z)^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}},$$

$$(9) \quad c_t^z = \left(\psi (c_s^z)^{\frac{\gamma-1}{\gamma}} + (1 - \psi) (c_h^z)^{\frac{\gamma-1}{\gamma}} \right)^{\frac{\gamma}{\gamma-1}},$$

where σ is the elasticity of substitution among broad consumption categories, γ is the elasticity of substitution between home production and substitutable services and ω and ψ are weights between zero and one. In addition, c_n^z is consumption of nonsubstitutable, c_t^z represents an aggregator which we label as *traditional services* (t), in turn composed of c_s^z , which is consumption of substitutable services and c_h^z , which is consumption of home services.

The first two types of consumption are purchased in the market, whereas home services are produced within the household. In the case of couples, both male and female labor is used to produce home services. This is not the case when the decision unit is a single woman ($z = fu$, no male labor is available) or when it is a single man ($z = mu$, no female labor is available). For each type of household, home services are produced according to the following technology:

$$(10) \quad Y_h^z = A_h L^z,$$

where

$$(11) \quad L^c = \left[\varphi_h^c (l^{fc})^{\frac{\eta-1}{\eta}} + (1 - \varphi_h^c) (l^{mc})^{\frac{\eta-1}{\eta}} \right]^{\frac{\eta}{\eta-1}},$$

¹⁴ In equilibrium, it might be the case that a married individual chooses not to work in the market, that is, $l = 1$ and $1 - l = 0$.

¹⁵ For singles, an agent coincides with a decision unit. For this reason, and without ambiguity, we use fu and mu , which we already used at the end of Subsection 4.1 to identify single agents of the two gender, with the index z of decision units.

$$(12) \quad L^{fu} = \xi_h^{fu} l^{fu},$$

$$(13) \quad L^{mu} = \xi_h^{mu} l^{mu},$$

and η is the elasticity of substitution between male and female labor in a married household, A_h , ξ_h^{fu} , and ξ_h^{mu} are positive scale parameters, $0 < \varphi_h^c < 1$ is a weight, l^{fc} and l^{mc} are female and male labor devoted to home production in a married household, respectively, l^{fu} is labor devoted to home production in a single female household, and l^{mu} is labor devoted to home production in a single male household. The budget constraint of a household of type z is

$$(14) \quad p_n c_n^z + p_s c_s^z = E^z,$$

where

$$(15) \quad E^c = w_j^{fe} \widehat{a}_j^{fe} (1 - l^{fc}) + w_j^{me} \widehat{a}_j^{me} (1 - l^{mc}),$$

$$(16) \quad E^{fu} = w_j^{fe} \widehat{a}_j^{fe} (1 - l^{fu}),$$

$$(17) \quad E^{mu} = w_j^{me} \widehat{a}_j^{me} (1 - l^{mu}).$$

Each household chooses the amount of consumption and the time devoted to home production in order to maximize utility (8), subject to the traditional services aggregator (9), the budget constraint (14), and the home production technology constraint (10).

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

$$(18) \quad \frac{l^{fc}}{l^{mc}} = \left(\frac{\varphi_h^c}{1 - \varphi_h^c} \frac{w_j^{me} \widehat{a}_j^{me}}{w_j^{fe} \widehat{a}_j^{fe}} \right)^\eta.$$

Thus, the time at home of the female agent increases with the wage and the effective efficiency units of the male in the market and declines with the wage and the effective efficiency units of herself in the market.

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

$$(19) \quad p_h^c(w_j^{fe} \widehat{a}_j^{fe}, w_j^{me} \widehat{a}_j^{me}) = \frac{1}{A_h} \left[(\varphi_h^c)^\eta [w_j^{fe} \widehat{a}_j^{fe}]^{1-\eta} + (1 - \varphi_h^c)^\eta [w_j^{me} \widehat{a}_j^{me}]^{1-\eta} \right]^{\frac{1}{1-\eta}}.$$

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 instead of 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 implicit home price for a single individual, for $i = fu, mu$, is

$$(20) \quad p_h^i(w_j^{ie}\widehat{a}_j^{ie}) = \frac{w_j^{ie}\widehat{a}_j^{ie}}{A_h} (\varphi_h^{iu})^{-\frac{\eta}{\eta-1}}.$$

This price is increasing in the efficiency wage of the agent, 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 (19) and (20), it also emerges 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 (19) and (20) 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 the skill-premium rises, highly educated individuals experience an increase in the price of the home good relative to the less educated, and react by increasing the demand for substitutable services in the market.

4.4. Definition of Equilibrium. The equilibrium is defined as a set of output prices $\{p_s, p_n\}$, a set of wages per unit of efficiency $\{w_s^{f1}, w_n^{f1}, w_s^{m1}, w_n^{m1}, w_s^{f0}, w_n^{f0}, w_s^{m0}, w_n^{m0}\}$, a sector choice (j) for each agent working a positive amount of time in the market, and a set of allocations for

- married households $\{c_s^c(\widehat{a}_j^{fe}, \widehat{a}_j^{me}), c_n^c(\widehat{a}_j^{fe}, \widehat{a}_j^{me}), l^{fc}(\widehat{a}_j^{fe}, \widehat{a}_j^{me}), l^{mc}(\widehat{a}_j^{fe}, \widehat{a}_j^{me})\}$,
- single female households $\{c_s^{fu}(\widehat{a}_j^{fe}), c_n^{fu}(\widehat{a}_j^{fe}), l^{fu}(\widehat{a}_j^{fe})\}$,
- single male households $\{c_s^{mu}(\widehat{a}_j^{me}), c_n^{mu}(\widehat{a}_j^{me}), l^{mu}(\widehat{a}_j^{me})\}$,

such that:

1. Given wages and prices, the choice of working sector j maximizes wage per efficiency unit for each agent;
2. Given wages, prices, and working sector j of each household member, the consumption and labor allocations maximize utility (8) of each household subject to the budget constraint (14) and the home production technology constraint (10);
3. Given wages and prices, each representative firm in sectors s and n maximizes profits (5);
4. Labor markets in sectors s and n clear, that is, in each sector total supply of efficiency units of each type equal total demand of those units;
5. Services and goods markets in sectors s and n clear.

5. QUANTITATIVE ANALYSIS

We design the quantitative exercise to compare two equilibria of the model, meant to represent the U.S. economy in 1980 and in 2017. The exogenous factors that differ between the two equilibria are:

- (i) demographics;
- (ii) the level of skill-biased technology; and
- (iii) the level of gender-biased technology.

Given the exogenous factors, agents endogenously decide consumption and work allocation in each equilibrium and, in doing so, they determine the distribution of employment along the skill distribution. We then use these distributions to construct employment polarization graphs for the model, and compare them with their data counterparts.

We divide this section in three parts. In Subsection 5.1, we briefly describe how we match the demographic structure. We devote Subsection 5.2 to explain the calibration strategy.

TABLE 4
DEMOGRAPHIC CATEGORIES

	1980	2017
Females		
Less educated single	0.135	0.128
Highly educated single	0.034	0.108
Less educated married (with less ed. spouse)	0.268	0.164
Highly educated married (with highly ed. spouse)	0.021	0.038
Highly educated married (with less ed. spouse)	0.009	0.020
Less educated married (with highly ed. spouse)	0.043	0.036
Males		
Less educated single	0.122	0.178
Highly educated single	0.028	0.070
Less educated married (with less ed. spouse)	0.268	0.164
Highly educated married (with highly ed. spouse)	0.021	0.038
Less educated married (with highly ed. spouse)	0.009	0.020
Highly educated married (with less ed. spouse)	0.043	0.036
Total	1.000	1.000

Finally, in Subsection 5.3 we describe the methodology used to construct employment polarization graphs from the model's equilibrium.

5.1. Demographics. Changes in the U.S. demographic structure that are relevant to our analysis are summarized in Table 4, where a highly educated agent is one with a college degree. In each year, there are 12 possible demographic groups to which an agent can belong. In the model, we assume a constant population of agents, which we randomly divide into 12 groups in such a way that the demographic structure is perfectly matched in each year. In Section 1 of the Online Appendix, we detail the formal procedure used to match the demographic structure in each year.

5.2. Parameters Calibration. For a group of parameters, specifically the elasticities $\{\sigma, \gamma, \eta, \eta_s\}$, we use values from previous studies. Estimates of the elasticity of substitution among broad consumption categories σ are typically between 0 and 0.3, as discussed in Ngai and Pissarides (2008). However, most studies consider a classification of broad categories into agriculture, manufacturing and services, or into goods and services. We are not aware of estimates of the elasticity of substitution between substitutable services and the rest of the economy. For this reason, we set σ to the value of the elasticity of substitution of the entire service sector with respect to total GDP estimated in Falvey and Gemmell (1996), who find a value of 0.3.¹⁶

Regarding the elasticity of substitution between home production and substitutable services γ , Rogerson (2007), Ngai and Pissarides (2011), and Moro et al. (2017) discuss that when considering a narrow set of market services, displaying high substitutability with home production, that parameter should be set higher than estimates in the literature, typically obtained considering substitutability between home production and total consumption, and in the range 1.49–2.30.¹⁷ Rogerson (2007) suggests a value of 5. Olivetti (2006) calibrates a value of 4 for that elasticity when considering substitutability between child care at home and in the market in the United States. Ragan (2013), instead, uses a substantially higher value of 6.66. We choose a conservative value $\gamma = 4$, as in Olivetti (2006) and provide a sensitivity analysis around this value in Section 3 of the Online Appendix.

Regarding the elasticity of substitution between highly educated and less educated workers η_s , Katz and Murphy (1992) estimate an elasticity of substitution between college and high

¹⁶ In Section 3 of the Online Appendix, we perform a sensitivity analysis on σ .

¹⁷ McGrattan et al. (1997), Chang and Schorfheide (2003), Rupert et al. (1995), and Aguiar and Hurst (2007).

TABLE 5
MODEL PARAMETERS

Estimated	Type	Value	S.E.
$\{\underline{a}_s, \bar{a}_s\}$	Substitutable services ability	{1,3.45}	{-,0.0154}
$\{\underline{a}_n, \bar{a}_n\}$	Nonsubstitutable services ability	{2.36,3.55}	{0.0075,0.0074}
ω	Consumption market weight nonsubs. services	0.88	0.0003
ψ	Substitutable services weight	0.03	0.0002
φ_h^c	Married female home labor weight	0.42	0.0004
ξ_h^{fu}	Single female home scale parameter	0.015	0.0001
ξ_h^{mu}	Single male home scale parameter	0.109	0.0004
$\varphi_{s,1980}$	Female labor weight in substitutable services	0.46	0.0003
$\varphi_{n,1980}$	Female labor weight in nonsubstitutable services	0.37	0.0003
ζ	Schooling factor	0.35	0.0051
$\theta_{s,1980}$	Highly educated workers labor weight subs. services	0.22	0.0007
$\theta_{n,1980}$	Highly educated workers labor weight nonsubs. services	0.37	0.0004
γ_θ	Skill-biased tech. change (yearly growth rate in θ_j)	0.012	0.00002
γ_φ	Gender-biased tech. change (yearly growth rate in φ_j)	0.006	0.0009
Predeter.	Type		Value
σ	Substitutability between broad cons. categories		0.3
γ	Substitutability between home and market services		4.0
η	Gender substitutability at home (married only)		3.0
η_s	Substitutability highly educated/less educated in production		1.41

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

school labor of 1.41 for the period 1963–87. Heathcote et al. (2010) follow them in setting a value for this elasticity of 1.43 for their aggregate production function. As we share the same functional form for the production function in each sector, and we assume a common rate of SBTC and GBTC at the sectoral level, we also set $\eta_s = 1.41$.¹⁸ The elasticity between the two types of labor in the home production function of the couple η is estimated in Knowles (2013) to 3 and we adopt this value.

Productivities by sector, including the home sector, are normalized to one, $A_j = 1$ and $A_h = 1$, and fixed over time.¹⁹ Sectoral ability is uniformly distributed between \underline{a}_j and \bar{a}_j , with $j = s, n$. Men and women draw ability by sector from two independent uniform distributions.²⁰ While the correlation of abilities across sectors is zero for an agent, the correlation across men and women in married couples is set to 0.26, which allows the model to match the wage correlation, adjusted for selection, of couples in the U.S. economy. We also normalize the lower bound of ability in substitutable services to $\underline{a}_s = 1$.

The remaining 15 parameters: (i) ability and return to education $\{\bar{a}_s, \underline{a}_n, \bar{a}_s, \zeta\}$, (ii) agents productivity (market and home) $\{\{\varphi_{j,1980}, \theta_{j,1980}\}_{j=s,n}, \varphi_h^c, \varphi_h^{fu}, \varphi_h^{mu}\}$, (iii) preferences $\{\omega, \psi\}$, and (iv) time evolution for GBTC and SBTC $\{\gamma_\varphi, \gamma_\theta\}$ are calibrated to match a number of moments. More specifically, we use a method of simulated moments (MSM).²¹ Table 5 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

¹⁸ That is, we assume that sectors experience the same technological change at each point in time, so it appears reasonable to calibrate such common technological change as if occurring in an aggregate production function.

¹⁹ As in Katz and Murphy (1992), we are interested in relative wage changes that emerge as a function of the relative supply of labor (through a market/home decision by the agents) and relative factor demand shifts (generated by SBTC and GBTC). For this reason, we follow them in abstracting from changes in absolute wages which might arise from unbiased technological change (i.e., TFP growth).

²⁰ Note that even if ability is uniformly distributed within each market sector, the model produces a dispersed wage distribution which is comparable to that in the data. We also performed a calibration in which we consider lognormal distributions for ability. We find that the results are very close to the ones obtained with uniform distributions. See Section 3 of the Online Appendix for details.

²¹ McFadden (1989).

matches all 15 parameters to 15 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. All parameters are matched to 1980 targets, except the time change of SBTC and GBTC which are matched to observed changes between 1980 and 2017 in the skill premium and the gender wage gap.²² Targets are computed using the 1980 Census and the 2017 American Community Survey. The relevant parameters and targets for 1980 are:

- Range of ability, $\{\bar{a}_s, \underline{a}_n, \bar{a}_n\}$ (three targets): Using male wage information, to avoid issues of women's changing selection in the data between the two years (Mulligan and Rubinstein, 2008), we match the wage premium in substitutable to nonsubstitutable services and the standard deviation of log male wages of full-time full-year workers from 1980 in the two sectors. The premium should be informative on the difference in average ability across the two sectors, whereas the standard deviations should be informative on the standard deviation of ability in both sectors. This allows us to match three targets with three moments as we assume $\underline{a}_s = 1$.
- Relative weight in consumption, $\{\omega\}$ (one target): we match the share of hours worked in substitutable services in 1980.
- Home production $\{\xi_h^{fu}, \xi_h^{mu}, \varphi_h^c, \psi\}$ (four targets): as home hours are defined as the remainder of total hours less market hours, we calibrate the four parameters to 1980 married male market hours, single male market hours, married female market hours, and single female market hours.²³
- Gender weights in 1980, $\{\varphi_{j,1980}\}_{j=s,n}$ (two targets): to determine the difference between men and women in terms of production weights, we use the aggregate gender wage gap and the female to male hours ratio in the nonsubstitutable sector in 1980.
- Education determinants, $\{\zeta, \{\theta_{j,1980}\}_{j=s,n}\}$ (three targets): the share of hours worked by less educated (i.e., less than college) in substitutable services, $\frac{N_s^{m0} + N_s^{f0}}{N_s^{m0} + N_s^{f0} + N_s^{m1} + N_s^{f1}}$, and the female and male college wage premium in 1980. This matches three targets to three moments, as hours of highly educated individuals, N_j^{g1} , are function of ζ , whereas the college wage premia are functions of $\{\theta_{j,1980}\}_{j=s,n}$ in Equation (6).

Only two additional parameters, the growth rate of GBTC and SBTC, are necessary to compute the 2017 equilibrium, as the rest of parameters are shared between the two equilibria. As can be seen from Equations (6) and (7), the growth rate of SBTC and GBTC can be identified by the change in gender gaps and education premia over time. As such, we use the following targets:

- SBTC, $\{\gamma_\theta\}$ (one target): the growth rate in the aggregate college wage premium between 1980 and 2017.
- GBTC, $\{\gamma_\varphi\}$ (1 target): the growth rate in the aggregate gender wage gap between 1980 and 2017.

²² We also performed the calibration with a two-step procedure. First, we calibrate all parameters except the rates of technological change using only targets for the 1980. Then, we calibrate the rate of change of SBTC and GBTC taking as given the parameters obtained in the first step. This procedure provides virtually the same parameter estimates, confirming that in the one step procedure in the text, the 1980 targets pin down structural (i.e., time-invariant) parameters, whereas the changes between 1980 and 2017 in the skill premium and the gender wage gap pin down technological change. Note, however, that the procedure calibrating all parameters at the same time greatly reduces the computational burden when computing standard errors.

²³ In the model, home hours and market hours sum to one, so analyzing market hours is the equivalent of analyzing both. The hours data we present is adjusted to remove leisure hours. From Aguiar and Hurst, 2007 (who kindly shared their data with us), we know that all work (market and home) hours add up to an average of 52 hours per week irrespective of gender or year. Thus, hours shares refer to a total amount of yearly hours equal to 2,496, resulting from assuming individuals work an average of 48 weeks times 52 hours per week each year.

TABLE 6
TARGETS

Type	Data	Model
<i>1980 - ability</i> ($\{\underline{a}_j, \bar{a}_j\}_{j=s,n}$)		
Male college nonsubstitutable to substitutable services wage	1.32	1.32
Standard deviation of nonsubstitutable log male wages	0.30	0.26
Standard deviation of substitutable services log male wages	0.28	0.27
<i>1980 - consumption</i> (ω)		
Share of total hours worked in substitutable services	0.065	0.064
<i>1980 - home production</i> ($\psi, \varphi_h^c, \xi_h^{fu}, \xi_h^{mu}$)		
Fraction of hours worked in the market by married male	0.78	0.72
Fraction of hours worked in the market by single male	0.61	0.60
Fraction of hours worked in the market by married female	0.34	0.36
Fraction of hours worked in the market by single female	0.49	0.52
1980 - gender weights in the market ($\{\varphi_j\}_{j=s,n}$)		
Aggregate gender wage gap	0.61	0.58
Female to male nonsubstitutable hours gap	0.54	0.52
<i>1980 - education ability returns</i> ($\zeta, \{\theta_{j,1980}\}_{j=s,n}$)		
Female college wage premium	1.56	1.82
Male to female college wage premium	1.0	1.01
Share of hours worked by less educated in substitutable services	0.92	0.91
<i>1980–2017 - SBTC and GBTC</i> ($\{\gamma_j\}_{j=\theta,\varphi}$)		
Relative female to male wage (total growth rate over time)	0.34	0.33
Relative college wages (total growth rate over time)	0.19	0.19

Table 6 reports the data targets and the corresponding values produced by the calibrated model.

5.3. Constructing Polarization Graphs in the Model. To construct employment polarization graphs that are comparable to the data, we first define what an occupation is in our environment. In the data, an occupation is identified by the Census as a set of tasks (i.e., *most important activities or duties*) to be performed by the workers employed in that occupation.²⁴ The boundaries of this set are then arbitrary, as the concept of “most important activities and duties” is subjective, and in principle one could define one different occupation for each worker in the economy. In the same fashion, an occupation in the model can be defined in different ways. Our setting belongs to the class of SBTC models, which by definition do not distinguish between skills and tasks. Thus, one interpretation of our setting is that each worker, with her specific skill set, is performing a different task in the economy. Following this interpretation, and similarly to the data, we need a rule to create sets of tasks that identify occupations in the model. Our choice is to create *intervals of ability*, to represent the model’s counterpart of an occupation in the data.²⁵

Technically, we generate polarization graphs in the model as follows: First, we create equally sized bins of effective ability on the supports $[\underline{a}_j, \hat{a}_j]$, $j = n, s$. Thus, each bin is the

²⁴ More specifically, the Census reports: “Occupation data describe the kind of work the person does on the job. These data are derived from responses to write-in questions that are autocoded and clerically coded by Census Bureau staff, using the Census Occupation Code List developed for Census Bureau household surveys.” <https://www.census.gov/topics/employment/industry-occupation/about/occupation.html>. In addition: “These questions ask: “What was this person’s main occupation?” and “Describe this person’s most important activities or duties.” <https://www2.census.gov/programs-surveys/demo/guidance/industry-occupation/overview2019.pdf>.

²⁵ More precisely, we create *intervals of effective ability*. In adopting this definition, we are implicitly assuming that workers in the same occupation have similar skills. Although it is unlikely that this property holds in the data, we also note that, as occupations are ranked by mean wage, the heterogeneity in the skills of workers within an occupation does not play a role in constructing employment polarization graphs. In addition, to support our choice we show in Appendix D.1 that our definition of an occupation in the model produces an average skill level in each occupational percentile that is comparable to observable skill measures in the data.

model counterpart of an occupation in a given sector in the data.²⁶ We distinguish by sector because, for instance, the ability level of a worker in the substitutable service sector cannot be directly compared with the ability level of a worker in the nonsubstitutable sector.²⁷ Once occupations in the model are defined, the procedure is the same as for the data. We compute the average (efficiency) wage in each occupation in the 1980 equilibrium and then rank all occupations from the two market sectors into a unique aggregate classification according to their average wage. Then, we construct occupation percentiles using hours worked by workers in each bin in the 1980 equilibrium. This ranking of occupations is kept unchanged for the 2017 equilibrium and the polarization graph is generated as in the data by measuring the change in employment share of each 1980 percentile. In Section 2 of the Online Appendix we provide details of the procedure.

6. RESULTS

In this section, we report the main predictions of the calibrated model. In Subsection 6.1 we discuss the ability of the model in replicating a number of aggregate untargeted moments. In Subsection 6.2, we compare the changes of employment shares along the skill distribution generated by the model with those in the data, and assess the capability of the model in generating employment polarization by gender and sector.

6.1. Untargeted Moments. In the two equilibria, the demographic structure (marital status and education shares) is matched by construction. Furthermore, the calibration matches the rise in the aggregate college wage premium as well as the closing aggregate gender wage gap.²⁸ No other changes over time, either in hours or wages are targeted in the calibration. Nonetheless, the model does reasonably well in matching changes in hours worked between the two years. The model reproduces most of the rise in the fraction of market work by females (a 0.14 increase in the data vs. 0.13 in the model) and half of the fall in the fraction of male hours (-0.04 in the data vs. -0.02 in the model).²⁹ Breaking down the change by marital status, the model reproduces all the rise observed for married females (0.17 in both data and model) whereas it slightly overestimates the fall in married male hours (-0.01 in the data vs. -0.02 in the model).

Regarding the fraction of hours worked by women and men in the two sectors and by education level, the model also produces changes comparable to those in the data. In particular, in the data the larger increase in the share of hours worked by females in the nonsubstitutable sector is given by highly educated women (0.13 vs. 0.05 of less educated women) whereas the opposite occurs in the substitutable sector, where less educated women display the largest increase (0.05 vs. 0.02 of highly educated women). This pattern is consistent with the conjecture that the increase in hours worked at the top of the distribution is mainly driven by high-skilled women in the nonsubstitutable sector, whereas the increase at the bottom by low-skilled women in the substitutable sector. The model reproduces this pattern with a 0.07 change of highly educated women versus a -0.02 of less educated women in the nonsubstitutable sector, and a 0.12 of less educated women versus a 0.02 of highly educated women in the substitutable sector.


Regarding wages, highly educated women is the group that experiences the largest increase, both in the model and in the data (0.43 in the data vs. 0.40 in the model in the nonsubstitutable sector and 0.33 in the data vs. 0.43 in the model in the substitutable sector). Highly

²⁶ Then, an occupation is a convex set defined over the ability support of one sector.

²⁷ This drives the same choice for the data, as discussed in note 6 in Section 3. That is, in the data we consider a “secretary” in the substitutable sector and a “secretary” in the rest of the economy as two distinct occupations.

²⁸ See Table 6.

²⁹ All numbers in this subsection refer to the change in the fraction of hours worked in the market by the relevant demographic group. For instance, a 0.14 increase for women results from the difference between the average fraction of hours worked in the market by this group in 2017 (0.53) and the corresponding fraction in 1980 (0.39).



NOTES: First panel: aggregate; second panel: females; third panel: males.

FIGURE 2

EMPLOYMENT POLARIZATION IN THE DATA (LINE) AND IN THE BENCHMARK MODEL (BARS)

educated men, in comparison, show a 0.24 increase in the data and a 0.26 in the model in the nonsubstitutable sector and a 0.12 increase in the data and a 0.27 in the model in the substitutable sector. This is consistent with the idea that the faster increase in the skill-premium for highly educated women triggered a large increase of hours worked by this category. In Appendix E, we report all changes in hours worked and in wages for the eight types of labor in production in the data and in the model.³⁰


6.2. Employment Polarization in the Model. In accordance with the facts on employment polarization presented in Section 3, we evaluate the ability of the model in reproducing changes of employment shares along two dimensions: gender and sectors. In Figure 2, we report changes in employment shares by quintiles of occupation in 1980 for the model, and superimpose the data smoothed line.³¹

Similarly to the data, the model produces a U-shape for the overall economy with an increase in the share of employment of the first and the fifth quintiles, and a decline of employment shares of the middle quintiles. The model produces a more pronounced increase at the top of the distribution, and a similar one at the bottom, when compared to the data. As in the data, women display a well-defined U-shape of changes in employment shares. The increase at the top of the distribution is quantitatively larger, whereas the increase at the bottom appears in line with the data. Also similarly to the data, men produce a mild U-shape of changes of employment shares across quintiles. The positive change at the top quintile is similar to the data, whereas the middle quintiles are all negative, as in the data. Whereas in the data the first quintile for men is slightly positive, in the model it is negative. However, it is less negative than the middle three quintiles.

Figure 3 reports the results by sector. In the empirical analysis, we emphasized the remarkable differences in the pattern of changes in employment shares along the skill distribution across sectors. Although the nonsubstitutable sector displays a mild U-shape with positive changes in the employment shares only after the 70th percentile and negative changes otherwise, the substitutable sector exhibits an opposite behavior, with positive changes at the

³⁰ That is, by gender, sector, and education level.

³¹ In assessing the ability of the model in reproducing employment polarization, we are mainly interested in three sets of changes of employments shares, those at the bottom, in the middle and at the top of the skill distribution. The convention in the empirical literature is to use smoothed lines, so we also apply this representation to the data in Section 3 and retain it here. However, a property of the model is that changes in employment shares display a substantially smoother pattern along the skill distribution when compared to raw data, so there is little purpose of using smoothed graphs. In addition, employment shares in the data display an increase at the bottom only below the 20th percentile. For these reasons, and for the sake of an easy comparison with the data, we focus on quintiles for the model. Bar graphs then prove to be particularly useful when presenting the counterfactual exercises in the next section, as they allow to make a quintile by quintile comparison between the benchmark results of the model and each counterfactual.



NOTES: Left panel: substitutable services; right panel: nonsubstitutable services.

FIGURE 3

EMPLOYMENT POLARIZATION IN THE DATA (LINE) AND IN THE BENCHMARK MODEL (BARS)

bottom of the distribution and almost zero changes everywhere else. Once again, the benchmark calibration replicates well the above empirical regularities both qualitatively and quantitatively, in that changes in the employment shares for the nonsubstitutable sector are positive only for the fifth quintile and negative otherwise, whereas the increase in the employment shares for the substitutable sector are concentrated in the first quintile, being negligible in the other quintiles. Overall, the model calibrated to aggregate targets replicates well the change in employment shares along the skill distribution both by gender, with women displaying substantially stronger employment polarization than men and by sectors, with the substitutable sector alone driving the upward shift in the bottom tail, and the rest of the economy driving the upward shift in the upper tail.³²

7. COUNTERFACTUALS


In Section 6, we show how the benchmark calibration of the model can account for the main facts of employment polarization by gender and sectors. We devote this section to present three sets of counterfactual exercises that allow us to disentangle the sources of employment polarization in the model: the first focuses on the role of technology, the second on the skill premium by gender and education, whereas the third presents an out-of-sample analysis for the period 1960–80.

7.1. The Role of Technology. This section investigates what the role of technology is in generating the employment polarization patterns by gender and sectors delivered by the benchmark calibration. To do so, we perform three counterfactual experiments in which we set, first one at the time, and then together, the two types of technological change to zero. The benchmark calibration delivers an average yearly growth of SBTC of 1.2% and of GBTC of 0.6%.³³ Figure 4 shows the counterfactual in which SBTC is set to zero.

The effect of removing SBTC is to reduce overall employment polarization substantially, as both the top and bottom quintiles decline. At the bottom quintile, the effect is strong enough to make the change in employment shares turn from positive in the benchmark to negative in the counterfactual. The gender panels in Figure 4 show that the main effect of removing SBTC is on women, as changes in employment shares are close to zero in the first four quintiles, and substantially reduced in the top quintile. For men, there is some redistribution of

³² In Appendix F, we show that the model also replicates fairly well wage polarization for both gender and for the aggregate economy.

³³ See Table 5.



NOTES: Left column: aggregate; middle column: females; right column: males.

FIGURE 4

COUNTERFACTUAL EXPERIMENTS, NO SBTC



NOTES: Left column: substitutable services; right column: nonsubstitutable services.


FIGURE 5

COUNTERFACTUAL EXPERIMENT, NO SBTC

employment shares across the skill distribution, which does not change the overall pattern. The counterfactual supports the idea that the increase in the skill premium since the 1980s (in the model driven by faster SBTC) attracted highly educated women in the labor market, thereby directly increasing the employment shares at the top of the skill distribution and indirectly increasing the employment shares at the bottom of the skill distribution via consumption spillovers enjoyed by workers in the substitutable sector. Although the same mechanism might also be triggered by highly educated men, these agents already display a high participation rate in the market in 1980 (Ngai and Petrongolo, 2017; Albanesi, 2020), so the quantitative effect of SBTC is substantially smaller than for highly educated women.

The above intuition is further supported by Figure 5, which shows the effects on sectors of removing SBTC. This exercise suggests that while SBTC directly shifts employment shares from the fifth to the other quintiles of the skill distribution in the nonsubstitutable sector, it also has a substantial impact on the substitutable sector by almost removing the positive change in the first quintile predicted by the benchmark calibration. We interpret this result as supportive of the existence of consumption spillovers from the upper to the bottom part of the skill distribution, and triggered by the rise in the skill premium.


Figure 6 shows the effect of removing GBTC. By observing the gender panels it appears that removing GBTC has mainly a level effect on both women and men, moving employment shares downward along the whole distribution for the former, and upward for the latter. For women, removing GBTC also slightly reduces employment polarization. The effect on overall polarization is to decrease the bottom quintile and slightly increase the top one due to a compositional effect between men and women. This exercise suggests that, while GBTC allows the model to match the closing gender wage gap, when compared to SBTC it plays a



NOTES: Left column: aggregate; middle column: females; right column: males.

FIGURE 6

COUNTERFACTUAL EXPERIMENTS, NO GBTC



NOTES: Left column: aggregate; middle column: females; right column: males.

FIGURE 7

COUNTERFACTUAL EXPERIMENTS, NO SBTC AND NO GBTC

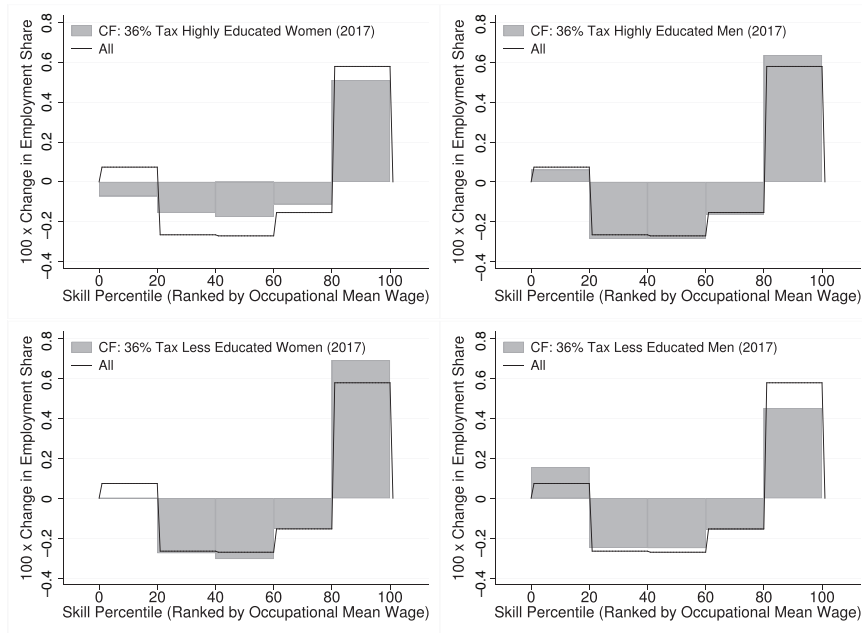
negligible role in generating heterogeneous changes of employment shares along the skill distribution, and cannot, by itself, generate overall employment polarization.³⁴

Finally, to investigate potential complementarities between the two types of technological change, we report the effect of removing both SBTC and GBTC. Figure 7 shows the results. Without the two types of technological change, the first four quintiles display a decline of similar magnitude for the overall economy, whereas the top quintile displays no change with respect to the benchmark. However, such unchanged pattern of the top quintile is the result of a substantial decline of employment shares at the top for women, and a corresponding increase for men. That is, although not visible in the aggregate, gender-specific graphs show that the decline of women’s market hours at the right tail are still correlated with a decline of women’s market hours at the left tail.

Figure 7 also allows us to assess the effect of demographic changes in the model. Absent technological change, the observed change in shares in the figure is due to the different demographic structure between the 1980 and the 2017 equilibrium. The larger fraction of highly educated agents in 2017 implies a substantial increase of employment shares at the top. However, Figure 7 also clarifies that the change in the demographic structure alone cannot generate employment polarization. Without SBTC magnifying wages at the top of the distribution, there is no increase of employment shares at the bottom.³⁵

³⁴ We introduce GBTC to allow the model to match the actual quantitative increase in female hours worked, which is an essential quantitative property the model should display, to be suited to studying employment polarization.

³⁵ Note that our SBTC counterfactual does not reduce the incentive to obtain education as would occur in a model where education is endogenous. Our approach here is similar to that of Buera et al. (Review of Economic Studies (forthcoming)) and it is motivated by the fact that we are interested in the impact of SBTC on time allocation instead of in education incentives. In doing so, however, we disregard the potential incentives that SBTC might have had on education rates. Thus, observed changes in education rates, which we take as exogenous, could be also interpreted as a response of agents to SBTC.



NOTES: Effects on overall employment polarization of a tax of 36% on wages of: highly educated females (top left), highly educated males (top right), less educated females (bottom left), and less educated males (bottom right).

FIGURE 8

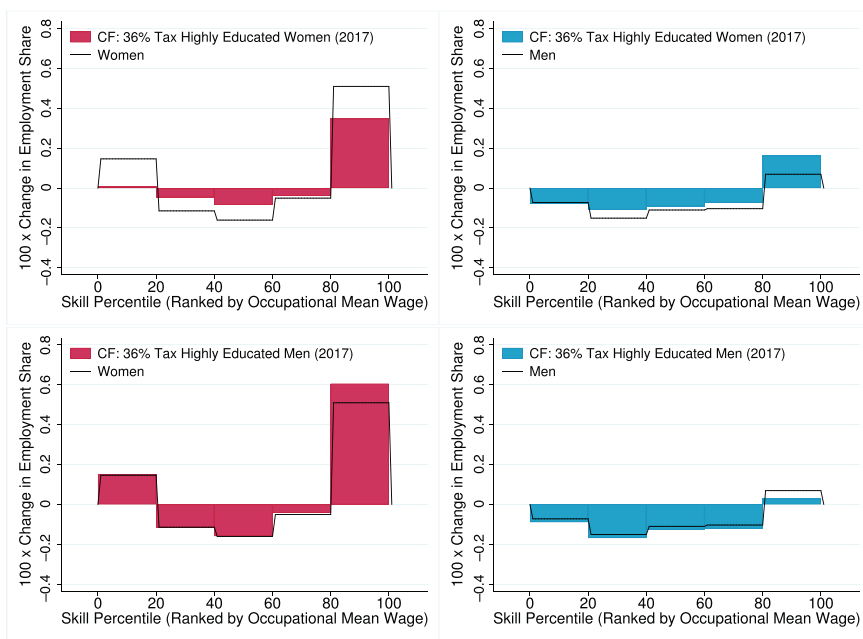
COUNTERFACTUAL EXPERIMENT

The counterfactual exercises in this section support the role of the rise of the skill premium after 1980 in generating the observed patterns of employment polarization by gender and sectors: by inducing a substantial increase in labor market participation of females at the top of the skill distribution, the rising skill premium also triggers a demand for services that are substitutable to home production, thus fostering labor demand for low-skilled (mostly female) workers in the sector producing these services.

7.2. Skill Premium by Gender and Education. The counterfactuals on technological change suggest that SBTC is a key factor in generating employment polarization. By fostering productivity of highly educated agents, it increases their wage in the market and simultaneously their opportunity cost of working at home. This fact, in turn, increases the implicit price of home services for these agents, and induces them to purchase substitutable services in the market. As SBTC affects highly educated agents of both gender, we are also interested in isolating the effect of the increase in the skill premium of highly educated women from that of highly educated men. To do this, we run a series of counterfactuals in which, in the 2017 equilibrium, we impose a tax on the wage of a specific group of agents.³⁶ In the first, we impose a tax on the wage of highly educated women. As wages of highly educated women are on average 36% higher in 2017 than in 1980, we set the tax to 36% of the wage. We then run other three counterfactuals in which we apply, in turn, a 36% tax on wages of highly educated males, less educated women, and less educated men. Results are reported in Figure 8.

The top-left panel of Figure 8 shows that the tax on educated women is sufficient to make employment polarization disappear, by dampening the positive change of employment shares

³⁶ Technically, we introduce the tax in the budget constraint of the relevant group, assuming that the tax rebates are wasted by the government, and solve for the 2017 equilibrium with all parameters as in the benchmark calibration. Then, we compute employment polarization by considering such a counterfactual economy for the 2017.



NOTES: Tax of 36% on highly educated females (top row) and highly educated males (bottom row). Effects on women left column and on men (right column).


FIGURE 9

COUNTERFACTUAL EXPERIMENT

in the first and fifth quintile and the negative change in the middle three quintiles. Note also that, as in Figure 4, the resulting pattern in this case is very similar to that of the nonsubstitutable sector in the data and in the benchmark model, as reported in the right panel of Figure 3. This suggests that, without the actual high wages of highly educated women in 2017, the model behaves very similarly to the whole economy without the substitutable sector. The remaining panels of Figure 8 report the counterfactuals in which the tax is applied to highly educated males (top-right panel), less educated women (bottom-left panel) and less educated men (top-right panel). None of these other counterfactual can produce a contemporaneous decline in the first and the fifth quintile, confirming that the key variable to generate employment polarization is the rising wage of highly educated women.

Figure 9 also shows the effect on the two gender of introducing taxation on highly educated categories. The top-row of Figure 9 shows the effect on women and men of a 36% tax on highly educated women. These panels confirm that the role of women in generating employment polarization is key, as both the top and the bottom quintile decline for this gender when taxing highly educated women. The effect on men is to increase all quintiles apart from the bottom one. The bottom row of Figure 9 shows the effect of taxing highly educated men by 36% on the two gender. Most of the effect is concentrated in the top quintile, with an increase in the shares of women that compensate the decline in the employment shares of men in that quintile.


Finally, Figure 10 shows the effect on sectors employment shares in the counterfactual in which we impose a 36% tax on highly educated women. The main effect is that of reducing the increase of shares both in the top quintile of nonsubstitutable services, and in the bottom quintile of substitutable services, again suggesting that the reduction in purchasing power of highly educated women alone reduces dramatically employment polarization because of a lack of consumption spillovers affecting the bottom of the skill distribution.



NOTES: Tax of 36% on highly educated females: substitutable services (left column); nonsubstitutable services (right column).

FIGURE 10

COUNTERFACTUAL EXPERIMENT



NOTES: The vertical sum of the colored lines gives the black line.


FIGURE 11

EMPLOYMENT POLARIZATION BY GENDER FOR THE 1960–80 PERIOD

7.3. *Prepolarization Era.* Our main focus in this article is the period after 1980, during which employment polarization is broadly considered to have appeared (Acemoglu and Autor, 2011; Acemoglu and Autor, 2012). Thus, our theory is grounded on changes in the structure of the U.S. economy that emerged after 1980. To provide additional support to the theory, we devote this section to assess the ability of the model in reproducing the different patterns of employment shares observed between the pre- and post-1980. The evidence for the period 1960–80 is reported in Figure 11.³⁷ During the pre-1980, changes of employment shares of men display a monotonically increasing pattern along the skill distribution, with large differences in magnitude across percentiles. Changes of employment shares of women, instead, are more homogeneous and mostly positive along most of the distribution. Thus, there is a striking difference of changes in employment shares for both gender between the pre- and post-1980. In particular, no curve displays a U-shaped pattern in the 1960–80 period.³⁸

³⁷ Computing employment polarization for the 1960–80 period requires dealing with occupations that are not present in both years. The occupations classification we adopt here is the one developed by Bárány and Siegel (2018), which provides a consistent set of occupations for the period 1960–80.

³⁸ Bárány and Siegel (2018) find some evidence of a U-shaped pattern of employment shares since 1950 in the United States but they also show that changes of employment shares at the bottom of the skill distribution are neg-



NOTES: Left panel: aggregate; middle panel: females; and right panel: males.

FIGURE 12

EMPLOYMENT POLARIZATION IN THE DATA FOR THE 1960–80 PERIOD (LINE), IN THE MODEL FOR THE 1960–80 (FULL BARS), AND IN THE MODEL FOR THE BENCHMARK CALIBRATION FOR THE 1980–2017 PERIOD (EMPTY BARS)


Given the difference in the behavior of employment shares between the pre- and post-1980, a natural out-of-sample test of the theory is to study the behavior of the model when we feed it with the changes in exogenous factors measured for the 1960–80 period. For the 1980–2017 benchmark, SBTC and GBTC are calibrated together with the rest of parameters, by taking as given the evolution in the demographic structure. For the 1960–80, we want to keep the same parameterization of the model except for technological change, so we tie our hands on SBTC and GBTC measurement by relying on the results in Heathcote et al. (2010). Thus, we set an average growth rate for SBTC of -0.0066 (as opposed to a calibrated annual growth of 0.012 for the period 1980–2017) and an average growth rate for GBTC of 0.0064 during the 1960–80 period (as opposed to a calibrated value of 0.0056 for the period 1980–2017) in our experiment.³⁹ Finally, we match the demographic structure in 1960 and 1980.

Figure 12 presents the comparison between changes in employment shares in the data for 1960–80 and the corresponding pattern generated by the model, both for the 1960–80 and the 1980–2017 periods. Although the model cannot account for the large drop of employment shares in the first quintile of the 1960–80 period, it does relatively well in the out-of-sample exercise, by capturing the main trends by gender and so at the aggregate level. In particular, it replicates the increasing trend for men (with positive change in employment shares only after the 80th percentile) and the positive and roughly homogeneous change in the employment shares along the skill distribution for women.

Given the key role of highly educated women in generating employment polarization for the 1980–2017 period, we also perform a counterfactual experiment in which we impose a 36% subsidy to highly educated women in the 1980 equilibrium and recompute employment polarization for the 1960–80 period. The rationale for this counterfactual is the following. During the 1980–2017 period, wages of highly educated women increase by 36%, and taxing these wages by 36% in the 2017 equilibrium is sufficient to make employment polarization

ative for the 1950–80 period (see their Figure 1). In Cerina et al. (2020), we show that a U-shaped pattern mildly emerges for the 1960–80 period only when agriculture occupations are dropped from the sample, as in Bárány and Siegel (2018). In this article, instead, we follow the original methodology in Acemoglu and Autor (2011), who include agricultural occupations in the sample. When these occupations are included, no U-shape is appreciated, as reported in Figure 11. This reflects the substantial disappearance of agricultural employment during this period in the United States, and the fact that the employment share of agricultural occupations is 7.34% in 1960 and 2.88% in 1980. Thus, dropping agricultural occupations for the period 1960–80 would imply removing a large fraction of the labor force, concentrated mainly at the bottom of the skill distribution in 1980. See also Comin et al. (2020) for a recent paper analyzing employment polarization and including agricultural occupations in the sample.

³⁹ We thank Kjetil Storesletten for providing the numbers. Heathcote et al. (2010) compute the implied SBTC and GBTC for the period 1966–2005. We use their numbers for the 1966–80 period to compute an average growth rate that we apply to the 1960–80 period. Note that measures of technological change in Heathcote et al. (2010) are appropriate in our setting because we employ the same production function. Although they have a unique production function at the aggregate level, whereas we have two sectors, SBTC and GBTC are common across sectors in our model.



NOTES: First panel: aggregate; second panel: females; and third panel: males.

FIGURE 13

EMPLOYMENT POLARIZATION IN THE 1960–80 BENCHMARK (LINE) AND IN THE COUNTERFACTUAL IN WHICH THERE IS A 36% SUBSIDY TO HIGHLY EDUCATED WOMEN IN THE 1980 EQUILIBRIUM (BARS)

disappear.⁴⁰ In the 1960–80 period, the skill premium does not display a clear trend, being similar at the beginning and at the end of the period, especially for females (Autor et al., 2008, figure 5; Acemoglu and Autor, 2011, figure 4). Thus, it is reasonable to ask whether by artificially generating an increase in wages of highly educated women similar to the one observed in the 1980–2017 period is sufficient to generate employment polarization in the 1960–80 period. Results are reported in Figure 13, which compares the behavior of the model for the 1960–80 period reported in Figure 12 with this new counterfactual.

The increase in the wage for highly educated women is sufficient to generate employment polarization for women which, in turn, induces an increase of employment shares both at the top and at the bottom of the skill distribution at the aggregate level.⁴¹ This is also due to the fact that the effect on men is smaller than for women and mainly located at the very top of the skill distribution. Thus, the main message of this exercise, which is in line with the main predictions of the theory, is that female employment shares are flatter in 1960–80 with respect to 1980–2017 because in the earlier period the wage of high-skilled women did not increase as fast as in the later period, thereby preventing the emergence of the mechanism we highlight in this article for the post-1980 period.

8. CONCLUSIONS

We investigate the role of gender in generating employment polarization in the United States. Our main empirical findings suggest that women display a larger degree of employment polarization during the 1980–2017 period. We relate this fact with evidence by sectors, showing that the rise of employment shares at the bottom of the skill distribution is mainly driven by sectors producing services that are substitute for home produced services. We propose an extension of the theory of SBTC, in which the rising skill premium after 1980 triggers a large increase in market participation of high-skilled women. This increase, in turn, reduces home production and creates a demand for market services produced by workers at the bottom end of the skill distribution. The theory accounts reasonably well for overall, gender, and sectoral employment polarization, and allows a set of counterfactuals to be performed. These allow to show that indeed in the model the increase of wages of high-skilled women is a main determinant of employment polarization. Our results suggest that any policy aimed at affecting the overall pattern of employment polarization should consider the effects on the specific demographic groups that are contributing to this phenomenon.

⁴⁰ See Figure 8, top-left panel.

⁴¹ Strictly speaking, the negative change in the bottom quintile becomes smaller in absolute value.


APPENDIX A: SECTORS CLASSIFICATION

To define services sectors that are substitutable to home production, we use the procedure in Moro et al. (2017). First, from time use surveys we select home activities that are considered home production. We follow Bridgman et al. (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 (three 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; and residential care facilities, without nursing.

APPENDIX B: ROUTINE INTENSITY AND SECTORS

In this appendix, we show evidence on how the increase of employment shares at the bottom of the skill distribution is driven largely by routine occupations. This is reported in the left panel of Figure B.1. To construct the panel, we divide occupations into routine and nonroutine according to the occupations classification in Autor and Dorn (2013).⁴² Consistent with the routinization hypothesis, employment shares of routine occupations decline in the middle, and their change is mildly positive at the top of the distribution. However, routine occupations also drive the increase of employment shares at the bottom, with a contribution which is larger than that of nonroutine occupations. This evidence appears at odds with the idea that shares at the bottom of the distribution of employment increase because workers abandoning middle-skill routine occupations take low-skill manual occupations (Autor and Dorn, 2013).

The middle and right panels of Figure B.1 provide evidence by sectors as defined in this ar-



NOTES: Left panel: whole economy; middle panel: substitutable services; and right panel: nonsubstitutable services.

FIGURE B.1

EMPLOYMENT POLARIZATION BY ROUTINE AND NONROUTINE OCCUPATIONS

tile. The middle panel shows that in the substitutable sector, the increase at the bottom of the skill distribution is due to both routine and nonroutine occupations. Considering the bottom percentiles, nonroutine occupations contribute roughly 2/3 and routine roughly 1/3 of the increase. This suggests that it is the nature of the sector the occupation is in (substitutable ser-

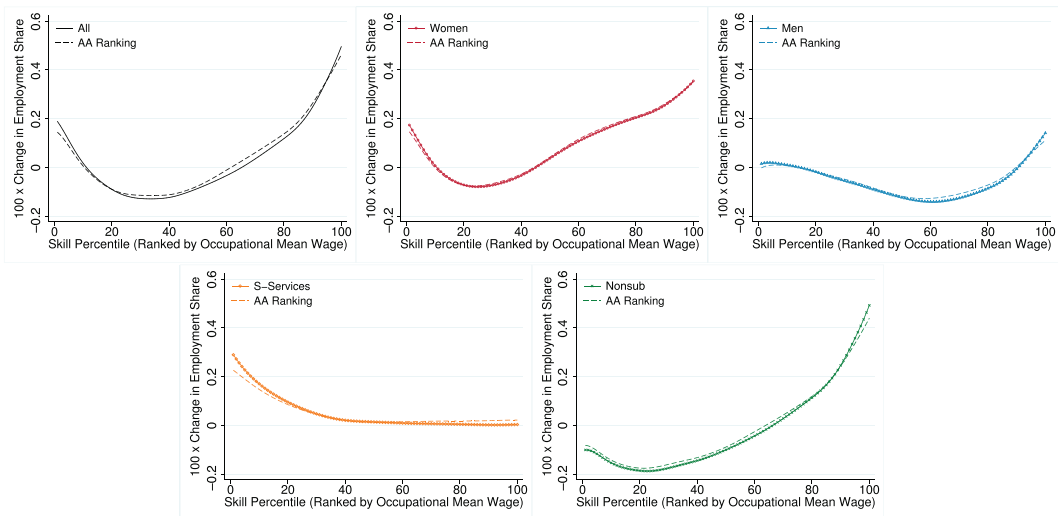
⁴² Autor and Dorn (2013, p. 1571), write “we first use the routine-task index (RTI) to identify the set of occupations that are in the top employment-weighted third of routine task-intensity in 1980. We refer to these as routine-intensive occupations.” Here, we follow the same procedure using their RTI by occupation and splitting the sample into routine and nonroutine occupations.

VICES), instead of the task content of the occupation, what matters most for the increase of employment shares at the bottom of the distribution.

The right panel of Figure A.1 reports evidence for the nonsubstitutable sector. In this sector, employment shares at the bottom do not grow, as predicted by our theory (see Figure 3). When splitting the sector by routine content, we observe that at the bottom of the distribution, employment shares of routine occupations mildly grow whereas those of nonroutine clearly decline. Again, this fact seems challenging for the routinization hypothesis, which suggests that routine occupations should experience a decline in their employment shares at the bottom of the distribution in favor of manual nonroutine occupations.

APPENDIX C: COMPUTING EMPLOYMENT POLARIZATION

As outlined in the text, we follow the methodology of Acemoglu and Autor (2011), who provide data and codes for creating employment polarization graphs.⁴³ We use their codes to construct employment polarization adapting them as needed for all the different splits we report in the article (i.e., gender and sector). For the benchmark graphs, we use the 1980 Census of Populations (5% sample of the United States) and the 2017 ACS (1% sample of the United States). In Section 7, we also use the 1960 (1% sample of the United States) Census of Populations. For details on the data selection process and treatment, see Acemoglu and Autor (2011). 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 two sectors and occupation Census classifications.⁴⁴ In Figure C.1, we



NOTES: AA ranking refers to the ranking used by Acemoglu and Autor (2011). Top row: aggregate, females and males; bottom row: substitutable services and nonsubstitutable services.

FIGURE C.1

EMPLOYMENT POLARIZATION COMPARING OCCUPATIONAL RANKINGS

compare the employment polarization graphs generated with the methodology in Acemoglu and Autor (2011) and ours. The two ranking methods generate minor deviations.

APPENDIX D: ADDITIONAL EVIDENCE

⁴³ <https://economics.mit.edu/faculty/acemoglu/data/handchapter>.

⁴⁴ See Subsection 5.3.

D.1. *Assessing gender differences in employment polarization.* As no universally accepted measure allows the comparison of relative polarization degrees between sets of data, 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 than those of men. First, the standard deviation of changes in employment shares by 1980 percentile (red dots line and blue triangles line in Figure 1) computed with respect to the gender-specific mean (i.e., the dashed lines in Figure 1) is 0.13 for women and 0.07 for men, implying that changes in female employment shares along the skill distribution display twice the variation than that of males. Although a higher standard deviation does not imply more polarization, the top panel of Figure 1 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, and relative to the gender-specific mean) 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 1 suggest a more pronounced U-shape, around the average change, for women than for men.

We also compute male and female changes in employment shares by percentiles *within each gender*, that is, using gender-specific total hours worked instead of aggregate total hours as the denominator in 1980 and 2017.⁴⁵ This provides a measure of the heterogeneity of changes in employment shares for each gender which is independent (in an accounting perspective) from the fact that, on average, female employment shares increase relative to male employment shares. The measure, plotted in Figure D.1 for both genders, reinforces the view that employment is substantially more polarized for women than for men. In this case, the standard deviation computed using the gender-specific mean is 0.34 for women and 0.12 for men.

TABLE D.1
CHANGE IN EMPLOYMENT SHARES BY GROUPS OF OCCUPATIONS AND GENDER

Occupational Group	Wage 1980	% Emp. Share 1980			Change until 2017		
		All	Male	Fem.	All	Male	Fem.
Managerial, professional	3.04	23.47	15.49	7.98	14.63	3.60	11.02
Precision production, craft, repair	2.87	14.91	13.97	0.93	-5.13	-5.18	0.05
Operators, fabricators, laborers	2.65	19.82	14.72	5.10	-8.25	-5.18	-3.07
Technical, sales, admin support	2.65	28.40	11.57	16.83	-3.21	-0.93	-2.29
Service	2.33	10.52	4.94	5.57	3.39	1.64	1.73
Farming, forestry and fishing	2.00	2.88	2.59	0.29	-1.42	-1.37	-0.04
Total	2.68	100.00	63.30	36.70	0.00	-7.43	7.43

D.2. *Occupations.* Table D.1 reports the dynamics of employment shares by gender and Census broad groups of occupations (ordered by their mean log hourly wage in 1980).⁴⁶ The aggregate trend of employment polarization is confirmed also by looking at broad occupation categories: highest paid occupations (managerial, professional specialty occupations) are associated with the largest increase in the employment share (from 23.47% to 38.10%, almost a 2/3 increase). On the other hand, even the employment share of the lowest paid occupations

⁴⁵ More precisely, in place of the expressions included in (2), we focus on the following two formulas:

$$\left(\frac{H_{i,2017}^f}{H_{2017}^f} - \frac{H_{i,1980}^f}{H_{1980}^f} \right); \left(\frac{H_{i,2017}^m}{H_{2017}^m} - \frac{H_{i,1980}^m}{H_{1980}^m} \right),$$

where $H_{1980}^f = \sum_i H_{i,1980}^f$, $H_{1980}^m = \sum_i H_{i,1980}^m$, $H_{2017}^f = \sum_i H_{i,2017}^f$, and $H_{2017}^m = \sum_i H_{i,2017}^m$. Notice that, unlike in (2), the sum of these gender measures does not deliver the aggregate polarization curve.

⁴⁶ Details on how occupations are grouped can be found at the web page https://usa.ipums.org/usa-action/variables/OCC1990#codes_section. Our table can be compared to table 1 in Autor and Dorn (2013). The main difference is that we aggregate occupations according to Census classification and add the gender dimension.




FIGURE D.1

CHANGES IN EMPLOYMENT SHARES IN THE UNITED STATES BETWEEN 1980 AND 2017 BY SKILL PERCENTILE NORMALIZED FOR GENDER-SPECIFIC TOTAL WORKING HOURS

(services) increases from 10.52% to 13.91%. 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. At the same time, the remaining middle-wage occupation group's employment shares decline between 1980 and 2017. This is particularly true for the two groups "precision production craft and repair occupations" (from 14.91% to 9.78%) and "operators, fabricators, and laborers" (from 19.82% to 11.57%).

The evidence reported in Table D.1 is consistent with the mechanisms proposed in this article, where the increase of labor participation by high-skilled women after 1980 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 7.43 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 7.98% to 19.00%), whereas the male share increases by less than 20% (from 15.49% to 19.09%). This is consistent with a rising skill-premium affecting primarily (high-skilled) women.⁴⁷ On the other extreme, the only other group of occupations where women increase their employment share is the service one: here the female employment share grows by almost 31% (from 5.57% to 7.30%), a growth rate which is comparable to that of males (from 4.94% to 6.58%). These occupations are highly concentrated in sectors producing services which are highly substitutable to household production (especially child care workers, gardeners, cleaners, and 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 the highest paid occupations. Hence, the evidence by broad groups of occupations is also consistent with the idea that rising high-skilled female labor force participation in the 1980s generated an indirect increase in employment shares in occupations at the bottom of the distribution.

APPENDIX E: MICRO DATA TREATMENT

⁴⁷ In 1980, highly educated men worked 83% of their time whereas highly educated women only 49% (Census).

E.1. *Weekly market hours.* Data are drawn from IPUMS - USA. We use data for 1980 (Census) and 2017 (American Community Survey). We use a repeated cross-sectional survey data to perform an OLS regression of market hours on education, marriage, and state fixed effects, exploiting an interaction term between education and marriage. In details, the adopted empirical model is the following:

$$\text{MarketHours}_i = \beta_0 + \beta_1 \text{Educated}_i + \beta_2 \text{Married}_i + \beta_3 (\text{Educated} * \text{Married})_i + \theta_s + \varepsilon_i,$$

where variable MarketHours_i is the yearly labor hours worked by the individual i in each year. It comes from the product of the variable UHRSWORK, namely, the usual hours worked per week, and the variable WKSWORK1, which indicates the weeks worked last year. The variable *Educated* assumes value 1 if the individual i has at least a college degree (at least 10 in the variable EDUC, corresponding to four years of college), and 0 otherwise. *Married* takes value 1 if the individual i records code 1 in the variable MARST, meaning “Married, spouse present,” 0 otherwise. θ_s are State dummies. Estimations are weighted with the variable PERWT, which indicates how many people in the U.S. population are represented by a given person in the sample. We run estimations separately for both years, 1980 and 2017. The final sample includes all individuals not currently in school aged 25–65.

E.2. *Weekly home hours.* In computing the mean weekly home hours, we follow a similar procedure using the data from TUESA and ATUS. Specifically, we use the same data set of Aguiar and Hurst (2007) for 1975 and to recover similar variables for 2017 we follow the same procedure as Gelber and Mitchell (2012) for ATUS. ATUS data are expressed in number of minutes respondents reported for the activities they engaged in over a 24-hour period. We convert minutes into hours per week dividing the total amount of minutes by 60 and multiplying by 7 days. Following Gelber and Mitchell (2012), home hours are the sum of time spent in preparing food, drinks, presentation, and clean up, housework like indoor cleaning and laundry, ironing, mending, clothes upkeep home repairs and car care, outdoor chores, yard work, raking leaves, plumbing, fixing furniture or appliances, gardening, animal care, and plant care. *Educated* takes value 1 if grade is greater than 15 where grade measures the highest grade “completed” by respondent and it is top-coded at 17. For the year 2017, *Educated* takes value 1 if the variable *peeduca* is greater than 42 since 43 means bachelor’s degrees, 44 Master’s Degrees, and so on. *Married* equals one if respondent is married, and zero otherwise. The empirical model is the same as for market hours with the only exception of the inclusion of State dummies. The information of respondent State of residence is not available in ATUS data. The adopted empirical model is the following:

$$\text{HomeHours}_i = \beta_0 + \beta_1 \text{Educated}_i + \beta_2 \text{Married}_i + \beta_3 (\text{Educated} * \text{Married})_i + \varepsilon_i.$$

The final sample includes all individuals not currently in school nor retired aged 25–65.

E.3. *Consumer expenditure.* To measure the share of consumption on home services substitutes, we compute expenditure shares from the CEX as the summation of “Food On-Premise,” “Alcohol On-Premise,” “Clothing Services,” “Barbershops, Beauty Parlors, Health Clubs,” “Domestic Service, Other Household Operation,” “Nursing Homes,” and “Nursery, Elementary, and Secondary Education,” relative to total consumption expenditure. As before, highly educated takes a value of 1 if years of schooling are at least 16. The adopted empirical model is the following:

$$\text{ExpenditureShareHome}_i = \beta_0 + \beta_1 \text{Educated}_i + \beta_2 \text{Married}_i + \beta_3 (\text{Educated} * \text{Married})_i + \varepsilon_i.$$

The final sample includes all individuals aged 25–65.




FIGURE F.1

SKILL CONTENT OF OCCUPATIONS: MODEL VERSUS DIFFERENT DATA MEASURES

TABLE F.1
MODEL AND DATA OCCUPATIONAL CORRELATIONS

		Model Ability
Data	Abstract tasks	0.74
	Years of education	0.73

APPENDIX F: SKILL MEASURES BY OCCUPATIONS

In constructing employment polarization graphs in the model, we assume that the model counterpart of an occupation in the data is an interval of effective ability. In adopting this definition, we are implicitly assuming that workers in the same occupation have similar skills. In the data, instead, an occupation might include agents with substantially different skill level. Although our assumption rule out such dispersion of ability within an occupation, it is also true that in the data occupations are ranked according to mean wage, such that the dispersion of ability within an occupation is also ruled out. To support our choice, however, we devote this appendix to investigate whether our definition of an occupation in the model produces a skill content in each occupation percentile that is comparable to observable measures of ability in the data.

To do this, we compare, for each occupational 1980 percentile, the average effective ability in the model’s equilibrium (normalized between zero and one) to: (i) years of education (normalized between zero and one) in each 1980 percentile in the data and (ii) an index (between zero and one) of how abstract intensive occupations in each percentile are.⁴⁸ Results are reported in Figure F.1. The two measures are commonly regarded as skill proxies (not involving wages) and both are highly positively correlated with the rank of the associated skill percentile (which, as in the model, involves wages). In addition, Table F.1 shows the correlation coefficient between the indices in the model and in the data across percentiles. Both measures in the data are highly correlated with the measure of effective ability in the model, suggesting that our theoretical definition of an occupation produces a skill content by occupational percentiles in 1980 that is comparable to the two observable measures of skills in the data.

APPENDIX G: UNTARGETED MOMENTS


Table G.1 reports changes in hours worked and in wages for the eight types of inputs in production in the data and in the model. We refer the reader to Subsection 6.1 for a discussion of the table.

⁴⁸ The abstract measure is from Autor and Dorn (2013).

TABLE G.1
UNTARGETED MOMENTS

Sector and education	Women				Men			
	<i>n</i> , 1	<i>n</i> , 0	<i>s</i> , 1	<i>s</i> , 0	<i>n</i> , 1	<i>n</i> , 0	<i>s</i> , 1	<i>s</i> , 0
Data hours	0.13	0.05	0.02	0.05	-0.03	-0.08	0.01	0.03
Model hours	0.07	-0.02	0.02	0.12	-0.09	-0.02	0.00	0.02
Data wages	0.43	0.23	0.33	0.23	0.24	0.00	0.12	-0.09
Model wages	0.40	0.12	0.43	0.21	0.26	0.00	0.27	0.08

NOTES: Change in the fraction of hours worked and wages by gender, sector (*n*, *s*), and education (0, 1) in the data and in the model. A highly educated agent in the data is one with at least a college degree. All numbers show the difference between the value in 2017 and in 1980. Numbers for wages refer to the difference between the two years of the relevant log wage relative to the log wage of less educated men in the nonsubstitutable sector.



NOTES: Left panel: aggregate; middle panel: females; and right panel: males.

FIGURE H.1

WAGE POLARIZATION IN THE DATA (LINE) AND IN THE BENCHMARK MODEL (BARS)

APPENDIX H: WAGE POLARIZATION

The phenomenon of employment polarization is sometimes studied jointly with the wage behavior along the skill distribution. For instance, Autor and Dorn (2013) show that during the 1980–2005 period there is both employment and wage polarization. However, in Cerina et al. (2020), we show that wage polarization can emerge or not in periods of employment polarization. In this appendix, we compare the behavior of wages along the skill distribution in the data and in the model for our reference period 1980–2017.

For both years, and for both data and model, we compute the average log hourly wage for each 1980 occupational percentile. We then compute the change between the two periods for each 1980 percentile, normalized by the change in the consumption price index (CPI).⁴⁹ Note, however, that by construction the model cannot match the growth in the average level of real wages between 1980 and 2017.⁵⁰ Consequently, to compare model and data we rescale figures such that aggregate real (i.e., relative to CPI) wage growth is zero for the 50th percentile both in the model and data.⁵¹ This normalization allows us to compare the change in real wages along the skill distribution regardless of the fact that the model does not display an average increase in real wages that is comparable to the data.

Figure H.1 reports the result for the data and for the model. For the 1980–2017 period, we find that wage polarization coexists with employment polarization. This is reported by the black lines in the three panels of Figure H.1. There is a well-defined wage polarization pattern

⁴⁹ Hourly wages are defined as yearly wage and salary income divided by the product of weeks worked times usual weekly hours. We use the historical CPI index computed by the BLS for the data while we use a Fisher price index to construct the CPI in the model.

⁵⁰ See note 20 on this point.

⁵¹ More specifically, if real wage growth for the 50th percentile of the aggregate wage polarization graph is *x*% in the data and *y*% in the model, we subtract *x*% from all wage growth percentiles in the data (for aggregate, female and males graphs) and *y*% from all wage growth percentiles in the model (for aggregate, female and males graphs).

for both females and males, which results in aggregate wage polarization. Wage increases are stronger for women than for men both at the top and the bottom of the distribution, whereas the decline in the middle is stronger for men. The model broadly reproduced these patterns, although the change of female wages is quantitatively stronger in the data than in the model.

SUPPORTING INFORMATION

Additional supporting information may be found online in the Supporting Information section at the end of the article.

Table 1 - An example of the demographic matrix

Figure 1 - Employment polarization in the data (line) and in the model (bars).

Figure 2 - Employment polarization in the data (line) and in the model (bars).

Figure 3 - Employment polarization in the data (line) and in the model (bars).

Figure 4 - Employment polarization in the data (line) and in the model (bars).

Figure 5 - Employment polarization in the data (line) and in the model (bars).

Figure 6 - Employment polarization in the data (line) and in the model (bars).

REFERENCES

- , and ———, “Skills, Tasks and Technologies: Implications for Employment and Earnings,” in *Handbook of Labor Economics*, Volume 4, Part B (Amsterdam: Elsevier, 2011), 1043–171.
- , “What Does Human Capital Do? A Review of Goldin and Katz’s *The Race between Education and Technology*,” *Journal of Economic Literature* 50 (2012), 426–63.
- , and E. HURST, “Measuring Trends in Leisure: The Allocation of Time over Five Decades,” *Quarterly Journal of Economics* 122 (2007), 969–1006.
- , ———, and L. KARABARBOUNIS, “Recent Developments in the Economics of Time Use,” *Annual Review of Economics* 4 (2012), 373–97.
- ALBANESI, S., “Changing Business Cycles: The Role of Women’s Employment,” Discussion Paper, National Bureau of Economic Research, 2020.
- AUTOR, D. H., and D. DORN, “The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market,” *American Economic Review* 103 (2013), 1553–97.
- , L. F. KATZ, and M. S. KEARNEY, “Trends in US Wage Inequality: Revising the Revisionists,” *Review of Economics and Statistics* 90 (2008), 300–23.
- BÁRÁNY, Z. L., and C. SIEGEL, “Job Polarization and Structural Change,” *American Economic Journal: Macroeconomics* 10 (2018), 57–89.
- , and ———, “Explaining Women’s Success: Technological Change and the Skill Content of Women’s Work,” Discussion Paper, National Bureau of Economic Research 13116, 2007.
- , “Explaining Women’s Success: Technological Change and the Skill Content of Women’s Work,” *Review of Economics and Statistics* 92 (2010), 187–94.
- BRIDGMAN, B. “Home Productivity,” *Journal of Economic Dynamics and Control* 71 (2016), 60–76.
- , A. DUGAN, M. LAL, M. OSBORNE, and S. VILLONES, “Accounting for Household Production in the National Accounts, 1965–2010,” *Survey of Current Business* 92 (2012), 23–36.
- , and J. P. KABOSKI, “The Rise of the Service Economy,” *American Economic Review* 102 (2012), 2540–69.
- , ———, R. ROGERSON, and J. I. VIZCAINO, “Skill-Biased Structural Change,” *Review of Economic Studies* (forthcoming).
- , ———, and M. Q. ZHAO, “The Rise of Services: The Role of Skills, Scale, and Female Labor Supply,” Discussion Paper, National Bureau of Economic Research, 2013.
- CERINA, F., E. DIENESCH, A. MORO, and M. RENDALL, “Spatial Polarization,” Discussion Paper 201909, CRENoS, 2019.
- , A. MORO, and M. RENDALL, “A Note on Employment and Wage Polarization in the U.S.,” Discussion Paper, 2020.
- CHANG, Y., and F. SCHORFHEIDE, “Labor-Supply Shifts and Economic Fluctuations,” *Journal of Monetary Economics* 50 (2003), 1751–68.
- COMIN, D., A. DANIELI, and M. MESTIERI, “Income-Driven Labor-Market Polarization,” Discussion Paper, National Bureau of Economic Research, 2020.
- FALVEY, R. E., and N. GEMMELL, “Are Services Income-Elastic? Some New Evidence,” *Review of Income and Wealth* 42 (1996), 257–69.

- GELBER, A. M., and J. W. MITCHELL, "Taxes and Time Allocation: Evidence from Single Women and Men," *Review of Economic Studies* 79 (2012), 863–97.
- GOLDENBERG, K., and J. RYAN, "Evolution and Change in the Consumer Expenditure Surveys: Adapting Methodologies to Meet Changing Needs," in *NBER Summer Institute 2009—Conference on Research on Income and Wealth*, vol. 13 (2009).
- GOLDIN, C., "The Quiet Revolution that Transformed Women's Employment, Education, and Family," *American Economic Review* 96 (2006), 1–21.
- HAZAN, M., and H. ZOABI, "Do Highly Educated Women Choose Smaller Families?" *Economic Journal* 125 (2015), 1191–226.
- HEATHCOTE, J., K. STORESLETEN, and G. L. VIOLANTE, "The Macroeconomic Implications of Rising Wage Inequality in the United States," *Journal of Political Economy* 118 (2010), 681–722.
- KATZ, L. F., and K. M. MURPHY, "Changes in Relative Wages, 1963–1987: Supply and Demand Factors," *Quarterly Journal of Economics* 107 (1992), 35–78.
- KNOWLES, J. A., "Why are Married Men Working So Much? An Aggregate Analysis of Intra-Household Bargaining and Labour Supply," *Review of Economic Studies* 80 (2013), 1055–85.
- LANDEFELD, J. S., and S. H. McCULLA, "Accounting for Nonmarket Household Production within a National Accounts Framework," *Review of Income and Wealth* 46 (2000), 289–307.
- LEE, D., and K. I. WOLPIN, "Intersectoral Labor Mobility and the Growth of the Service Sector," *Econometrica* 74 (2006), 1–46.
- MANNING, A., "We Can Work It Out: The Impact of Technological Change on the Demand for Low-Skill Workers," *Scottish Journal of Political Economy* 51 (2004), 581–608.
- MAZZOLARI, F., and G. RAGUSA, "Spillovers from High-Skill Consumption to Low-Skill Labor Markets," *Review of Economics and Statistics* 95 (2013), 74–86.
- McFADDEN, D., "A Method of Simulated Moments for Estimation of Discrete Response Models without Numerical Integration," *Econometrica: Journal of the Econometric Society* 57(5) (1989), 995–1026.
- McGRATTAN, E. R., R. ROGERSON, and R. WRIGHT, "An Equilibrium Model of the Business Cycle with Household Production and Fiscal Policy," *International Economic Review* 38 (1997), 267–90.
- MORO, A., S. MOSLEHI, and S. TANAKA, "Does Home Production Drive Structural Transformation?" *American Economic Journal: Macroeconomics* 9 (2017), 116–46.
- MULLIGAN, C. B., and Y. RUBINSTEIN, "Selection, Investment, and Women's Relative Wages over Time," *Quarterly Journal of Economics* 123 (2008), 1061–110.
- , and B. PETRONGOLO, "Gender Gaps and the Rise of the Service Economy," *American Economic Journal: Macroeconomics* 9 (2017), 1–44.
- , and ———, "Trends in Hours and Economic Growth," *Review of Economic Dynamics* 11 (2008), 239–56.
- , "Taxes, Social Subsidies, and the Allocation of Work Time," *American Economic Journal: Macroeconomics* 3 (2011), 1–26.
- OLIVETTI, C., "Changes in Women's Hours of Market Work: The Role of Returns to Experience," *Review of Economic Dynamics* 9 (2006), 557–87.
- RAGAN, K. S., "Taxes and Time Use: Fiscal Policy in a Household Production Model," *American Economic Journal: Macroeconomics* 5 (2013), 168–92.
- RENDALL, M., "Female Market Work, Tax Regimes, and the Rise of the Service Sector," *Review of Economic Dynamics* (2017).
- and F. J. WEISS, "Employment Polarization and the Role of the Apprenticeship System," *European Economic Review* 82 (2016), 166–86.
- ROGERSON, R., "Taxation and Market Work: Is Scandinavia an Outlier?" *Economic Theory* 32 (2007), 59–85.
- RUPERT, P., R. ROGERSON, and R. WRIGHT, "Estimating Substitution Elasticities in Household Production Models," *Economic Theory* 6 (1995), 179–93.
- TÜZEMEN, D., and J. L. WILLIS, "The Vanishing Middle: Job Polarization and Workers Response to the Decline in Middle-Skill Jobs," *Economic Review* 98 (2013), 5–32.