Spatial Polarization*

Fabio Cerina[†] Elisa Dienesch[‡] Alessio Moro[§] Michelle Rendall[¶]

May 2022

Abstract

We document the emergence of *spatial polarization* in the U.S. during the 1980-2008 period. This phenomenon is characterized by stronger employment polarization in larger cities, both at the occupational and the worker level. We quantitatively evaluate the role of technology in generating these patterns by constructing and calibrating a spatial equilibrium model. We find that faster skill-biased technological change in larger cities can account for a substantial fraction of spatial polarization in the U.S. Counterfactual exercises suggest that the differential increase in the share of low-skilled workers across city size is due mainly to the large demand by high-skilled workers for low-skilled services and to a smaller extent to the higher complementarity between low- and high-skilled workers in production relative to middle-skilled workers.

JEL Classification: J21, O14, R12, R23.

Keywords: Employment Polarization, Spatial Sorting, City Sizes.

---

*We thank Nicola Fuchs-Schündeln, Nezih Guner, Nir Jaimovich, Joe Kaboski, Jordy Meeks, Eric Mengus, Andrew Rendall, Frederic Robert-Nicoud, Andrii Parkhomenko, Kurt Schmidheiny, Benoit Schmutz and seminar participants at Universitat Autonoma de Barcelona, Universitat de Barcelona, Universitat Rovira y Virgili, University of Pau, the Workshop of the Australasian Macroeconomics Society (WAMS) in Queenstown, the Workshop on Macroeconomics and Development in Dalian, the VII Workshop on Structural Transformation and Macroeconomic Dynamics in Cagliari, the European and American Meetings of the Urban Economics Association in Amsterdam and Philadelphia, the 18th Journees Louis-André Gerard-Varet in Aix-en-Provence and the 60th meeting of the Society of Italian Economists in Palermo for the useful comments and discussions. Cerina and Moro thank the Region of Sardinia for financial support, research grant RASSR89213. Dienesch thanks the support by the French National Research Agency Grant ANR-17-EURE-0020, and by the Excellence Initiative of Aix-Marseille University - A*MIDEX. We also thank Sara Pau for excellent research assistance. The usual disclaimers apply.

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

[‡] Sciences Po Aix and Aix Marseille University, CNRS, AMSE. E-mail: elisa.dienesch@sciencespo-aix.fr.

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

[¶] Monash University and CEPR. E-mail: michelle.rendall@monash.edu.

1 Introduction

The polarization of the U.S. labor market in the last forty years has been extensively documented (Acemoglu and Autor, 2011). This empirical observation refers to the contemporaneous increase of the hours employment shares of occupations pertaining to the bottom and the top of the occupational skill distribution, with a consequent decline of shares in the middle. The empirical work is accompanied by a large literature investigating the theoretical causes of the phenomenon. The workhorse theory focuses on the fact that occupations in the middle of the skill distribution are those which mostly involve routinary tasks, and for this reason they are more exposed to substitutability with new forms of capital that can be *programmed* to perform those tasks (Autor and Dorn, 2013). Recent work, instead, suggests that skill-biased technological change (SBTC) played a key role in the emergence of employment polarization through *consumption spillovers* (Cerina et al., 2021b): it increased the opportunity cost of working at home of high-skilled individuals, who reacted by increasing participation in the labor market, reducing home production time, and demanding more services in the market that represent goods substitutes for home production. As these services are typically produced by low-skilled workers, this process generates aggregate employment polarization. By their nature, consumption spillovers have a local dimension, so that in locations in which the impact of SBTC has been stronger - typically larger cities (Baum-Snow et al., 2018) - stronger employment polarization should be observed. In this paper we study the emergence of spatial employment polarization in the U.S. and investigate to what extent the spatial heterogeneity of SBTC contributed to generate this phenomenon.

We first document that during the 1980-2008 period the observed increase of employment shares of high-skilled and low-skilled occupations relative to middle-skilled ones (i.e. aggregate occupational polarization) has been stronger in larger cities. In addition, we provide evidence that this difference, as well as aggregate employment polarization, is largely accounted for by a change in the extensive margin (i.e. number of workers) rather than in the intensive margin (i.e. hours worked by each worker). This is consistent with the idea that over time larger cities attracted not only a larger fraction of high-skilled workers (Glaeser and Resseger, 2010, Diamond, 2016), but also a larger fraction of low-skilled workers relative to medium-skilled ones.

Next, we investigate whether consumption spillovers have empirical relevance at the spatial level. To do this, we split the economy into two sectors: one producing services that can be considered good substitutes for home produced activities (the *substitutable services sector*), and one encompassing the remaining goods and services (the *rest of the economy*). We find that the aggregate increase of employment shares at the top of the skill distribution is exclusively driven by the rest of the economy. At the spatial level, the change is more pronounced in larger cities, which experience a faster growth of employment shares at the top. Instead, the aggregate increase of employment shares at the bottom of the skill distribution is driven by both sectors, although the sector producing substitutable services accounts for a larger fraction of the increase. At the spatial level, the increase at the left tail is more pronounced in larger cities for both sectors. This evidence suggests that consumption spillovers have empirical relevance at the local level, as cities with a larger growth of employment shares at the top of the distribution also experience a faster growth of employment shares at the bottom because of the substitutable services sector. However, the evidence also suggests that spatial employment polarization could be potentially generated without the sector producing substitutable services.

To investigate whether the differential evolution of SBTC across space displays quantitative relevance in generating spatial employment polarization, we construct a spatial equilibrium model with workers belonging to one of three skill levels (low, middle and high), a home/market labor time decision and a multi-sector environment in which agents consume, in addition to housing and a tradable good, services produced at home and services produced in the market, which are imperfect substitutes. Market services are assumed to be locally produced, non-tradable across locations and intensive in low-skilled labor.

Motivated by the evidence discussed above, we consider two mechanisms through which SBTC at the local level can potentially foster the emergence of local employment opportunities for low-skilled individuals. The first is that of consumption spillovers, as in Cerina et al. (2021b). The second is that of *extreme-skill complementarity* in production, as in Eeckhout et al. (2014). This second mechanism implies that the productivity of workers at one end of the skill distribution is enhanced by workers at the other end of it (i.e. the productivity of high-skilled workers is enhanced by low-skilled workers and vice-versa). In this view, for instance, the opening of a new investment bank, a law office or a hi-tech ICT company would generate new demand for security, janitors, reception services, etc. This mechanism allows the model to potentially account for the increase of employment shares at the bottom of the skill distribution which are generated regardless of the existence of consumption spillovers.

Given the above mechanisms, faster SBTC in larger cities, as suggested by the data, implies that a a larger fraction of both high- and low skilled workers is attracted to those cities with respect to smaller ones. In addition, we allow for two other types of technological change in the model: total factor productivity (TFP) growth in both the tradable and the non-tradable sector. As in Eeckhout et al. (2014), TFP growth differentials across cities in the tradable sector can also interact with extreme-skill complementarity and thus potentially generate a larger increase in the share of high- and low-skilled workers in the city with larger

TFP growth.

For the quantitative exercise, we consider a version of the model with two locations and two equilibria, representing the years 1980 and 2008. We calibrate the model to match the observed differences in the change of the employment shares of the three types of workers between large and small cities during the 1980-2008 period. The calibration accounts well for the targets and shows that to match them, both SBTC and TFP growth must be stronger in the larger city. This result suggests that both types of technological change are relevant for the model to produce spatial employment polarization. We then run a series of counterfactuals to assess the role of the different types of technological change in generating this phenomenon. To do this, we impose that the two cities differ only with respect to one type of technological change at a time. By only allowing for differences in SBTC, the residual difference in the change in the share of the three types of workers between the two cities with respect to the benchmark calibration is 32% for the low-skilled, 67% for the middle-skilled and 80% for the high-skilled. By only allowing for differences in the TFP growth of tradables in the two cities, instead, the corresponding figures are roughly equal across skills and less than 20% for each of the three types of workers. This suggests that the SBTC channel is quantitatively more relevant than the TFP one to explain spatial employment polarization.

To disentangle the roles of consumption spillovers and extreme-skill complementarity in association with each kind of technological change, we perform the same counterfactuals described above but now removing extreme-skill complementarity in the production function of the tradable sector between high and low-skilled workers. In this case, allowing only for spatial differences in TFP growth in tradables does not generate any difference in the skill distributions, confirming that the extreme-skill complementarity is needed for the TFP channel to generate spatial polarization. By contrast, allowing only for city-specific SBTC, the residual spatial difference in the change in the share of the three types of workers amounts to 28% for the low-skilled, 80% for the middle skilled and 102% for the high-skilled relative to the benchmark with no extreme-skill complementarity. These results suggest that consumption spillovers have a sizable role in explaining the faster growth in the employment shares of low-skilled workers in large cities generated by faster SBTC therein, accounting for more than 85% of the total change.

Quantifying the relative importance of the two mechanisms that in our model connect the upper and the lower part of the skill distribution is also relevant for policy reasons. For instance, if local policy makers in large cities are interested in improving the welfare of low-wage workers, they might subsidize these agents in those locations. However, different ways through which the subsidy is financed can affect the spatial allocation of workers differently, depending on their interaction with consumption spillovers and extreme-skill complementarity. We use the calibrated model to implement different tax schemes that give in equilibrium an equivalent government revenue that is used to finance the subsidy, and study how they affect spatial polarization. We consider first a tax on wages of high-skilled workers and then a tax on consumption of non-tradables of high-skilled workers. The two policies have a similar impact on high-skilled workers, who are less attracted to large cities. However, they have an opposite impact on the spatial sorting of low-skilled workers. Large cities become more attractive for them with respect to the benchmark case when the subsidy is financed through the wage tax. Instead, when the subsidy is financed through the consumption tax, large cities lose attractiveness for low-skilled workers, such that the spatial difference in the change of their employment shares between large and small cities becomes negative. The intuition for such opposite effects relies on how the different policies affect the main quantitative mechanism generating spatial polarization, that of consumption spillovers. The consumption tax has a direct effect on the price of services, which is sufficient to neutralize the role of consumption spillovers in large cities. Thus, in this context, the subsidy is not enough for large cities to increase (or even retain) attractiveness for low-skilled workers, and the policy has the unintended effect of reducing welfare of the low-skilled and so their employment share in large cities. Instead, consumption spillovers are only mitigated by the wage tax, as the reduction of purchasing power of high-skilled workers affects all goods and services in their consumption bundle. As a result, the subsidy increases the attractiveness of large cities for low-skilled workers.

Finally, we use the model to show that the spatial polarization of the *occupational* skill distribution emerges together with the spatial polarization of the *workers* skill distribution. While related, the two phenomena do not necessarily imply each other. For instance, we could observe spatial employment polarization in occupations without concurrently observing spatial polarization in individual skills if, for example, a high- and a low-skilled vacancy (where the skill is identified by the mean wage of those occupations) are opened in a large city and filled by two middle-skilled workers (where the skill is identified by the worker's individual characteristics) who abandon their middle-skilled occupations. We use the theory to construct a model-based measure of skill to document that faster growth in the employment shares of high- and low-paid *occupations* in large cities is associated with a relatively stronger attraction for respectively high- and low-skilled *workers* in those locations. Crucially, we also document that the process of spatial workers' polarization starts emerging after 1980, as before that year small and large cities display a remarkably similar skill distribution. These results reinforce the view that spatial employment polarization is associated with a change in the spatial sorting of heterogenously skilled workers.

The reminder of the paper is organized as follows. In Section 2 we discuss the background literature and in Section 3 we present the evidence on employment polarization by city size; in Section 4 we set out the model while in Section 5 we present the calibration and the quantitative exercises; in Section 6 we discuss the evidence on the spatial polarization of the individual skill distribution. Finally Section 7 concludes.

2 Related Work

From an empirical perspective, the geography of employment polarization in the U.S. is studied in Autor (2019), who finds that the employment share of middle-skilled occupations shrink faster in denser areas. Our empirical results differ along two dimensions. First, we provide an analysis of employment polarization by city size based on a more disaggregated definition of occupations. This confirms that the disappearance of middle-skilled and the rise of high-skilled occupations are more pronounced in areas with a larger population.[1] Second, the classification of low-skilled occupations in this paper is driven by the theory. Thus, our low-skilled jobs only include service occupations while Autor (2019) also considers in this category transport, laborers and construction workers, which are included in middle-skilled occupations in our classification.[2]

The spatial dimension of labor market polarization in the U.S. is empirically investigated also in the state-level analysis of Lindley and Machin (2014). They find that between 1980 and 2010 employment polarization has been stronger in states where there was more education sorting and where both college premium and housing/amenities prices increased faster. Such high-polarization states also experienced bigger increases in the numbers of eating and drinking places, apparel stores, and hair and beauty salons. This observation, coupled with the finding in Moretti (2013), who reports that house prices are higher and have risen faster in cities where wage inequality has risen by more, is in line with our idea that large cities are becoming increasingly polarized due to a rising concentration of more educated workers who demand more services which are supplied by low-skilled labor.

On the theoretical side, Autor and Dorn (2013) argue that the canonical model of skill-

[1] The running variable in Autor (2019) is urban density in 1970 while for us it is urban population in 1980. Moreover, his location units are 722 commuting zones in the US., while our analysis is based on 218 metropolitan areas, which are on average significantly larger than the typical Commuting Zone.

[2] Autor (2019) argues that the decline in middle-skilled occupations in urban areas is driven by the fact that in large cities non-college workers move from increasingly disappearing clerical/administrative/manufacturing occupations to rising low-skilled service occupations. Finding direct evidence for this hypothesis requires rich longitudinal data keeping track of the job history of workers' cohorts and represents an intriguing future research agenda. Our model abstracts from occupational choice as we use occupation groups as invariant proxies for skills. For this reason, the faster increase in low-skilled occupations in large cities are more naturally interpreted as sorting of workers with innate low skills into large cities rather than a degrading of non-college workers into low-skilled occupations.

biased technological change cannot account for the U-shaped pattern of changes of employment shares along the U.S. occupational skill distribution, because of the absence of a theoretical distinction between skills and tasks. They present a spatial equilibrium model in which the declining price of computer capital induces firms to substitute low-skill *workers* performing middle paid routinary *occupations* with capital, a process typically referred to as *routinization*. Their key assumption in the spatial equilibrium is that local labor markets have different degrees of specialization in routine-intensive industries. In this paper, instead, we build on the results in Cerina et al. (2021b), who show that a model of skill-biased technological change, augmented with home production and a market sector producing services substitutable to it, can account for aggregate employment polarization in the U.S. through consumption spillovers. They show that aggregate employment polarization in the U.S. is largely generated by rising SBTC after 1980, which fostered women's participation, directly, at the top and, indirectly, at the bottom of the skill distribution, due to a larger demand for low-skilled services by skilled women. As low-skilled services are produced and consumed locally, the mechanism has its main empirical relevance at the level of metropolitan areas. Given the results in Baum-Snow et al. (2018), who find that the skill bias of agglomeration economies magnifies the impact of skill-biased technical change in larger cities, the mechanism of consumption spillovers should have generated stronger employment polarization in larger cities.[3] The idea of consumption spillovers is also supported by the results in Leonardi (2015).[4]

By using a model-based measure of skill, Eeckhout et al. (2014) find that larger U.S. cities in 2009 display fatter tails of the skill distribution and argue that extreme-skill complementarities in production are the main driver of the spatial sorting of both high- and low-skilled workers in more populated areas. The main differences between their work and this paper span three areas. First, we consider the role of SBTC, which is a key type of technological change over the period considered (1980-2008) and study its role in generating consumption spillovers when coupled with a home sector that allows substitutability with

[3]Baum-Snow et al. (2018) find that skill-bias of agglomeration economies, by boosting the impact of skill-biased technical change in larger cities, can account for most of the increase in urban inequality by city size since 1980. Their analysis focuses on *wage* inequality (measured as the log of the college wage premium at the CBSA level) while we study the impact of spatial differences in the pace of SBTC on spatial differences in the occupation and skill structure, focusing on the connection between the top and the bottom tail of the occupational skill distribution. The spatial effects of SBTC have also been studied by Giannone (2017) who find that, had SBTC not taken place in 1980, the convergence process across U.S metropolitan areas between 1940-80 would not have reverted but instead been only slightly mitigated.

[4]In Appendix B we also provide empirical evidence showing that, consistent with the theory in Cerina et al. (2021b), the spatial differences in employment polarization across city size are mainly driven by women, whose changes in employment shares of the three occupational groups (positive for the high- and low-skill occupations and negative for the middle-skilled ones) are substantially more pronounced in larger cities.

the non-tradable sector in the market. Second, we calibrate the model using U.S. data to perform a quantitative analysis to assess i) the role of technology and ii) the contribution of each of two channels, extreme-skill complementarity and consumption spillovers, in generating the emergence of fatter tails in larger cities over time. This exercise indicates that, while faster TFP growth has a non negligible role in generating the emergence of fatter tails in larger cities, the impact of skill-biased technological change is quantitatively more relevant.[5] Third, we extend their result by showing that in 1980 large U.S. cities did not display fat tails in the individual skill distribution relative to small ones.

Finally, two other recent papers study the changes of employment shares along the skill distribution in a spatial dimension. Davis et al. (2019) build a model based on elements of Autor and Dorn (2013) and Davis and Dingel (2019) which predicts, for larger cities, a faster increase in employment shares for the high-skilled, a faster decrease in employment shares for the middle skilled, and *a slower* increase in employment shares for the low-skilled workers. They document that the evidence for France supports these theoretical predictions. Parkhomenko (2021), instead, suggests that middle-income households are more likely than low- and high-income households to move for housing-related reasons to more affordable housing markets, i.e. small cities.

3 Employment Polarization and City Size

Employment polarization in the U.S., i.e. the relative disappearance of middle-skill occupations in favor of both high- and low-skilled ones since the beginning of the 1980s, is a well documented fact.[6] Based on individual data from the 1980 U.S. Census and the 2008

[5]Appendix C in Eeckhout et al. (2014) also considers the role of home produced services in a model with extreme-skill complementarity in production. In contrast to the results in this paper, their main conclusion is that the expenditure share on non-tradables must be unlikely high in order for consumption spillovers to generate fatter tails in larger cities. We note here that there are key differences between the two approaches, that allow us to reach the opposite conclusion, that consumption spillovers can produce fatter tails in larger cities with an empirically relevant expenditure share of non-tradables. First, we explicitly model home-production and low-skilled market services as two distinct sectors, while they assume that services are produced only at home and can be traded in the market. Thus, our approach allows us to control for the value of the elasticity of substitution between home production and non-tradables, whose value turns out to be key to assess quantitatively the role of consumption spillovers. Second, we discipline the quantitative role of consumption spillovers in a calibration exercise which targets the observed differential emergence of fat tails between large and small cities in the U.S. *between 1980 and 2008*, together with a number of other moments, such as the change in wage premia and hours worked by different types of workers. Thus, our calibration identifies the role of consumption spillovers by using both spatial and time patterns. In contrast, Eeckhout et al. (2014) search for parameter values under which the observed difference in thick tails between large and small cities *in 2009* can be obtained through consumption spillovers, and discuss the empirical plausibility of those values.

[6]Following the literature on employment polarization, the skill of an occupation is identified by the mean wage of the occupation in the initial year (1980). See Acemoglu and Autor (2011).

Table 1: Employment polarization in the U.S. in the period 1980-2008

Occupation Group	**Avg hourly wage 1980**	**Emp. Share 1980**	**Change 1980-2008**
Services	4.80	11.61%	+3.12%
Admin, Tech, etc.	6.79	62.72%	-11.66%
Prof. and Manag.	9.63	25.68%	+8.54%

Note: Shares are computed according to the number of hours worked.

American Community Survey, we start our investigation by providing novel evidence showing that employment polarization is more pronounced in larger cities, so that there is a spatial dimension to this phenomenon.[7] We adopt the same occupation classification used in Autor and Dorn (2013) which harmonizes U.S. Census codes overtime. Next, we divide occupations into three broad skill groups. Guided by the theory we present in Section 4, in which a key role is attributed to market services representing good substitutes for home produced services, the group of low-skill occupations is that of Services (1990 Census codes 405-472). This group accounts for the increase of employment shares at the bottom of the occupational skill distribution at the aggregate level in the U.S. between 1980 and 2008, as documented in Autor and Dorn (2013). On the other spectrum of the occupational skill distribution, we define as high-skilled all Managerial and Professional Specialty Occupations (codes 004-199). All remaining occupations are in the middle-skill group (codes 203-889 except 405-472).[8] Table 1 reports the well known pattern for the U.S. economy. The employment shares of occupations at the extremes of the distribution increase over time, while that of occupations in the middle shrink.

To perform the analysis by city size we consider 218 metropolitan statistical areas which we rank according to their population in 1980. Using this ranking, we define large and small cities, by splitting the sample in three groups of cities with equal total population. We then consider large cities as those in the top 33% and small cities those in the bottom 33%.[9] We emphasize that, consistently with previous work (Autor, 2019 and Baum-Snow et al., 2018) we fix the classification of city size in the initial year. This treatment of the data is consistent with the theory we present in the next section, in which we assume that cities that are *initially* large display different technological trajectories over time, and that these exogenous technological processes generate spatial polarization.

The left panel of Figure 1 reports the change in occupational shares and the pattern of

[7] Data and sample description can be found in Appendix A.

[8] As in Autor and Dorn (2013) we exclude agriculture and military occupations.

[9] This implies that the number of large cities is smaller than the number of small cities. Specifically, there are 10 large and 174 small cities (Miami with 2.6 millions people and Toledo with 791,000 being the marginal cities). A detailed description of city size and grouping definitions is provided in Appendix A.

employment polarization for large and small cities. Employment shares of low-skill occupations increase by 3.05 percentage points in small cities compared to 3.72 in large cities. For middle-skill occupations the figures are -10.90 for small cities and -13.39 for large cities and for high-skill occupations 7.85 in small cities versus 9.66 in large ones. In the right panel of Figure 1 we report the difference in difference between the two groups (i.e. the difference between the two groups of cities in the change of employment shares of each group of occupations) by also considering two additional splits of the sample: one in which we define as large the highest ranked cities where 50% of the population is concentrated - small cities being the remaining ones - and another in which large cities are the highest ranked ones where 25% of the population is concentrated, while small cities consist of the group of the lowest ranked cities where 25% of the population is concentrated.[10] Differences in employment polarization between small and large cities increase monotonically with the difference in city size. Thus, the results for broad occupation categories confirm the well documented existence of employment polarization at the aggregate level, but suggest a spatial dimension of the phenomenon, which is more pronounced in large cities than in small ones.

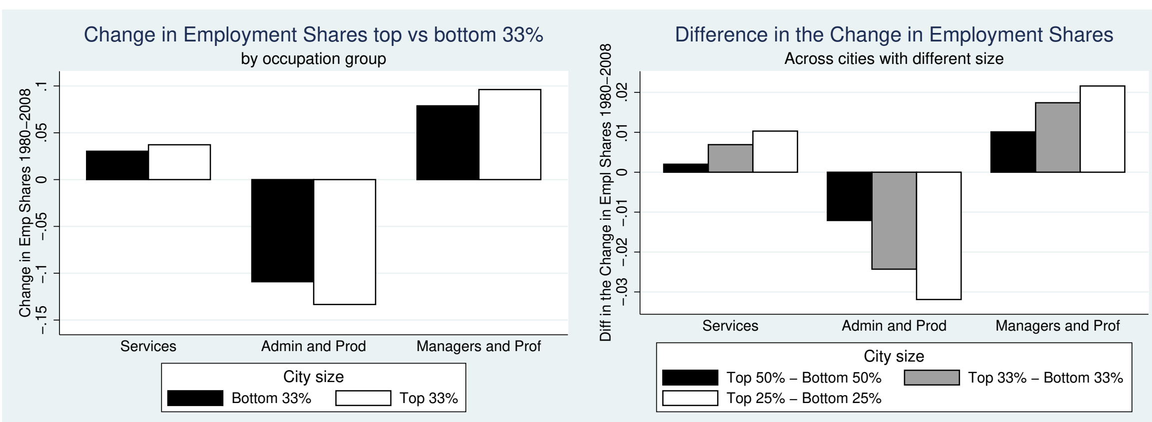

Figure 1: Employment polarization by city size. The left panel compares metropolitan areas belonging to the *top vs bottom 33%* grouping. The right panel reports the difference in the change in employment shares across cities with different size for three groupings: *top vs bottom 50%*, *top vs bottom 33%*, and *top vs bottom 25%*.

To provide further evidence on spatial employment polarization, we produce employment polarization graphs for each group of cities (i.e. small and large) by using the same methodology as in Acemoglu and Autor (2011). This procedure does not require us to classify occupations into pre-determined groups, and so it allows us to show that spatial employment polarization does not depend on the specific grouping of occupations that we employ

[10]Thus, within each categorization, the groups of large and small cities are equally populated.

in Figure 1. More precisely, we compute the average wage in 1980 of each occupation at the three digit level according to the 1990 occupational classification used by Autor and Dorn (2013). Then, we rank these occupations according to their average wage and construct occupation percentiles. By keeping the same ranking in 2008 we construct employment polarization graphs by measuring the change in employment share of each 1980 percentile and using a locally weighted smoothing regression. Results appear in Figure 2. As for the broad occupation categories in Figure 1, employment polarization is more pronounced in larger cities than in smaller ones.[11]

An important caveat here is that changes in employment shares include both the intensive and the extensive margin of employment. Thus, changes of employment shares across occupations can be due to either (a) workers who change their working time in the market while performing the same occupation, (b) workers switching occupations (within cities or across them) or (c) both channels. Measuring to what extent the two channels contribute to spatial employment polarization provides information on the role of the sorting of workers in producing the phenomenon. In Appendix B.1.1. we modify the graph in Figure 2 to consider only the change in the number of workers along the skill distribution, rather than the change in hours. Even with this alternative measure large cities appear more polarized than small ones, suggesting that the observed spatial employment polarization is driven by a larger increase in the proportion of *individuals* working in high- and low-skilled occupations in large cities than in small ones.[12] This, in turn, is consistent with the view that the higher employment polarization in larger cities is associated with a spatial sorting of workers that move across space, to fill the different portfolios of occupations offered in cities of different size over time.[13]

[11]We report the results for the *top vs bottom 33%* grouping. Results for the *top vs bottom 50%* and *top vs bottom 25%* groupings deliver similar results and confirm that differences in employment polarization between small and large cities increase monotonically with the difference in city size. Results are available upon request.

[12]This result holds also at the aggregate level, i.e. overall employment polarization is driven by the extensive margin rather than the intensive one. Results are available upon request.

[13]In principle, occupational sorting can also be due to a change in the way workers sort across occupations *within cities*, such that the spatial sorting of workers *across cities* does not change over time. However, the increase in the spatial sorting of the high-skilled after 1980 is a well documented fact, as discussed in Diamond (2016), which suggests a key role of this phenomenon in the emergence of spatial employment polarization. We devote Section 6 to show that spatial employment polarization emerges together with a change in the spatial allocation of both high- and low-skilled workers.

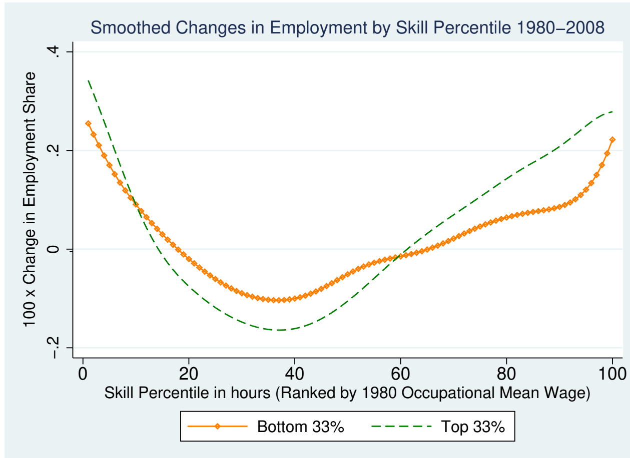

Figure 2: Employment polarization by city size. *Top vs bottom 33%* grouping.

To rationalize the emergence of spatial employment polarization, in the next section we present a theory that considers a specific partition of the production structure: the economy is split in two sector, one producing non-tradable services which represent market substitutes to home produced services, and one producing tradable goods and services (i.e. the rest of the economy).[14] In Figure 3 we document the spatial dimension of employment polarization when partitioning the economy into these two sectors. The figure provides several insights that support the existence of the two theoretical mechanisms proposed in this paper. First, the increase of employment shares at the top of the skill distribution is driven by tradables, and the change is more pronounced in large cities. The drop in middle-skilled occupations is also driven by the tradable sector, and such drop is more pronounced in large cities. Instead, both sectors contribute to the change of employment shares at the bottom. Also in this case, the change is more pronounced in both sectors of large cities. When comparing sectors, the largest contribution to changes in employment shares at the bottom comes from the non-tradable sector, while this sector does not contribute at the top of the skill distribution. Thus, Figure 3 indicates that at least two mechanisms are potentially at work in generating spatial employment polarization. One is that of consumption spillovers (Cerina et al., 2021b), driven by a strong increase in the demand for substitutable services from agents increasing employment shares at the top of the occupational skill distribution. The second can be linked to extreme-skill complementarity in the tradable sector (Eeckhout et al., 2014), which induces an increase of employment shares at the top and at the bottom of the skill distribution because of the higher complementarity in production of these two types of

[14]The list of sectors included in non-tradable services is the same as in Moro et al. (2017). See Appendix A for details.

occupations relative to middle-skilled ones. In Section 5 we use the theory to quantify the role of the two mechanisms in generating spatial employment polarization.

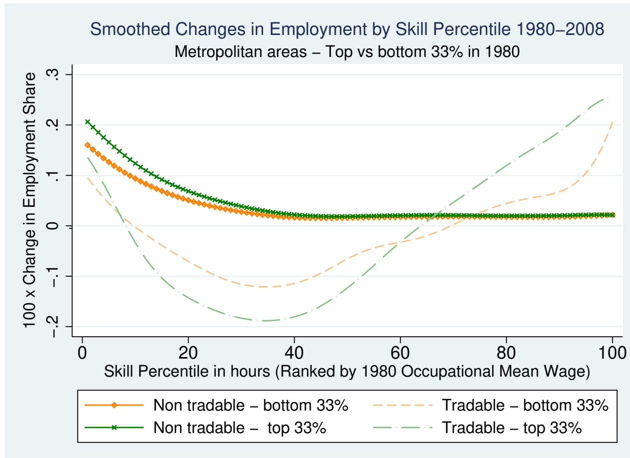

Figure 3: Employment polarization by city size and sectors. We compare metropolitan areas belonging to the *top vs bottom 33%* grouping.

4 Theoretical Framework

In this section we develop a spatial theory of employment polarization that can account for the data patterns reported in Section 3. Workers make a location decision based on their skill level, the wage rate paid to their skill type in each location and the cost of living. Cross-city differences in the cost of living arise as housing and non-tradable services have different prices across locations. In equilibrium, the utility of two workers with the same skill level but living in two different cities is equalized. The distributions of the different types of workers across locations and time are determined by the state of technology that we allow to vary in space and time due to total factor productivity growth and skill-biased technological change. In a nutshell, the model builds on elements from the multi-sector environment with a home production sector in Cerina et al. (2021b) and the spatial setting in Eeckhout et al. (2014).

4.1 The Environment

The economy consists of K locations (cities) indexed by $k \in (1, 2, ...K)$. In each location there is a fixed amount of housing H^k whose unit-price is location-specific and defined by p_H^k. As in Eeckhout et al. (2014) the expenditure on housing is the flow value that compensates for the depreciation and interest on capital. In a competitive rental market, the flow payment equals the rental price. To highlight the main mechanisms at work we restrict the number of cities to $K = 2$ but the model can be generalized to any number of cities.

Both cities are populated by workers with heterogeneous skills indexed by $i \in (1, 2....I)$ and associated with this skill order is a level of productivity a^{ik}. We focus on the case of three skills, $i = h,\ m,\ l$. At the economy wide level, there is a fixed amount of workers for each skill N^i. Following the literature on skill-biased technical change, we consider skills and occupations as isomorphic, such that each type of worker performs the unique occupation associated with her skill. The worker, however, can potentially perform such an occupation in any city, although in the equilibrium she chooses a specific location in which to work and consume.[15]

By n_j^{ik} we define the number of workers of skill i working in sector $j = g,\ s$ in location k. Hence $S_k = \sum_i n^{ik} = \sum_i \sum_j n_j^{ik}$ is the population size of city k. Workers of each skill move towards the city where their utility is higher so that the size of city k is an endogenous equilibrium outcome pinned down by the equalization of utilities across cities for the same skill. Total population of the economy is then exogenously given by $S = \sum_k S^k = \sum_k \sum_i n^{ik} = \sum_i N^i$.

4.2 Production

On the production side there are two sectors: the tradable sector, which produces in all cities *goods* that can be traded across locations; and the non-tradable sector which produces *market services* that can only be consumed in the same location where they are produced. Also, there exists a non-marketable service h which is produced within the household and interpreted as home production, which we describe in Section 4.3 together with the demand side.

4.2.1 The Tradable Sector

There is a representative firm in each location which employs three kinds of labor, h, m and l. The production function of the representative firm in city k in the g sector is

$$Y_g^k = A_g^k F\left(e_g^{hk}, e_g^{mk}, e_g^{lk}\right),$$

where e_g^i is the amount of hours worked by workers of skill i. In equilibrium, this amount of time is the product of an intensive margin - the individual labor supply $1 - l^{ik}$, and an extensive margin - the number of workers employed by the firm, n_g^{ik}. Since labor supply is

[15] In Section 6 we drop the three skills/occupation paradigm, and assume a large number of skills in the economy. This allows us to construct detailed individual skill distributions.

chosen by the individual worker who maximizes utility, the equilibrium number of workers of each skill employed by the firm is pinned-down by the relationship $n_g^{ik} = e_g^{ik}/(1 - l^{ik})$. A_g^k is the location-specific TFP in the tradable sector. We follow Eeckhout et al. (2014) in assuming that the production function of the representative firm has the following functional form:

$$Y_g^k = A_g^k \left[\left(\left(a^{hk} e_g^{hk} \right)^{\eta} + \left(a^l e_g^{lk} \right)^{\eta} \right)^{\lambda} + \left(a^m e_g^{mk} \right)^{\eta} \right]. \tag{1}$$

We assume $\eta < 1$ so that there are decreasing returns to scale. We also assume that the firm is owned by absentee capitalists, such that the profits of the firm do not enter the budget constraint of the workers. The parameters a^m and a^l are economy wide productivities of middle- and low-skilled workers, respectively. In the quantitative exercises in Section 5, we allow both parameters A_g^k and a^{hk} to change over time, potentially at a different pace across cities. We interpret the time changes in a^{hk} as *skill-biased technological change* (SBTC).[16] Also, as in Eeckhout et al. (2014), we allow $\lambda > 0$ to be potentially different from one. With $\lambda > 1$ there is extreme-skill complementarity and when $\lambda < 1$ there is extreme-skill substitutability.

The representative firm in city k in the g sector solves the following problem

$$max_{\{e_g^{hk}, e_g^{mk}, e_g^{lk}\}} \pi^k = Y_g^k - w^{hk} e_g^{hk} - w^{mk} e_g^{mk} - w^{lk} e_g^{lk},$$

where w^{ik} is the wage per unit of time worked by a worker of skill i in location k. Note that, despite workers' spatial mobility, wages are not equalized across cities because workers decide their location according to their utility, which depends both on wages and on local prices of housing and services.

[16] As described in note 4, Baum-Snow and Pavan (2013) find that agglomeration economies create a stronger impact of economy level SBTC in larger cities with respect to smaller ones. Differential SBTC at the spatial level in our model can be interpreted as a reduced-form version of the mechanism proposed and estimated by Baum-Snow et al. (2018).

First order conditions imply the following input demands:

$$
\begin{aligned}
e_g^{lk} &= \left(\frac{w_g^{lk}}{A_g^k \lambda \eta \left(\phi^l\right)^{\lambda}}\right)^{\frac{1}{\eta\lambda-1}} \left(1+\left(\frac{\phi^l}{\phi^h}\right)^{\frac{1}{\eta-1}} \left(\frac{w_g^{hk}}{w_g^{lk} a^{hk}}\right)^{\frac{\eta}{\eta-1}}\right)^{\frac{1-\lambda}{\eta\lambda-1}}, && (2)\\
e_g^{mk} &= \left(\frac{w_g^{mk}}{A_g^k \eta \left(a^m\right)^{\eta} \phi^m}\right)^{\frac{1}{\eta-1}}, && (3)\\
e_g^{hk} &= \left(\frac{w_g^{hk}}{w_g^{lk}} \frac{\phi^l}{\phi^h}\right)^{\frac{1}{\eta-1}} \left(a^{hk}\right)^{\frac{\eta}{1-\eta}} e_g^{lk}. && (4)
\end{aligned}
$$

Theorem 2 in Eeckhout et al. (2014) states that when $\lambda > 1$ and $\eta\lambda < 1$, a city with a larger A_g^k displays a larger fraction of high- and low-skilled workers relative to a city with a smaller level of A_g^k. Thus, due to the extreme-skill complementarity in technology, cities with faster TFP growth experience a larger increase in the proportions of high- and low-skilled workers with respect to cities with slower TFP growth. Here we show that in our model also SBTC induces a similar pattern. By deriving (2) with respect to a^{hk} we obtain the effect of SBTC on the demand for low-skilled workers:

$$
\frac{\partial e_g^{lk}}{\partial a^{hk}} = \frac{(\lambda-1)\,\eta}{(1-\eta\lambda)\,(1-\eta)} \frac{\left(\frac{\phi^l}{\phi^h}\right)^{\frac{1}{\eta-1}} \left(\frac{w_g^{hk}}{a^{hk} w_g^{lk}}\right)^{\frac{\eta}{\eta-1}}}{1+\left(\frac{\phi^l}{\phi^h}\right)^{\frac{1}{\eta-1}} \left(\frac{w_g^{hk}}{a^{hk} w_g^{lk}}\right)^{\frac{\eta}{\eta-1}}} \frac{e_g^{lk}}{a^{hk}}.
$$

This derivative is strictly positive if high- and low-skilled labor are complements, i.e. $\lambda > 1$, and if $1 > \eta\lambda$.[17] If these conditions hold, for a given wage premium between high- and low-skilled workers, a joint increase in the labor demand of both high- (as Equation 4 shows) and low-skilled workers occurs with SBTC. Instead, this type of technological change does not induce a change in the demand for middle-skilled workers, as shown by Equation (3). Thus, SBTC interacts with extreme-skill complementarity in a similar fashion as TFP growth. This implies that, *ceteris paribus*, cities with faster SBTC are expected to display a larger increase in the proportion of high- and low-skilled workers, with respect to cities with a slower pace of this type of technological change.

[17]Note that these are the same restrictions of Theorem 2 in Eeckhout et al. (2014).

4.2.2 The Non-Tradable Service Sector

The representative firm in the non-tradable service sector operates with the following production function

$$Y_s^k = A_s^k e_s^{lk}, \tag{5}$$

where A_s^k is the location-specific TFP in the non-tradable sector. Profit maximization implies equality between prices and marginal costs.

$$p_s^k = \frac{w_s^{lk}}{A_s^k}, \tag{6}$$

where w_s^{lk} is the wage per unit of time worked by a low-skilled worker in the non-tradable sector in location k. Note that low-skilled workers are perfectly mobile across sectors so that, in a given location the wage rate is equal across sectors and therefore $w_g^{lk} = w_s^{lk} = w^{lk}$ holds.

The assumption that only low-skilled workers are employed in the services sector is motivated by the fact that in the data the hours share of this type of worker (i.e. individuals employed in service occupations, as defined in Section 3) in this sector is substantially larger (52.44% in 1980 and 51.25% in 2008) than in the overall economy (11.16% in 1980 and 14.73% in 2008). Also, conditional on being employed in a service occupation, the probability of working in the non-tradable sector is substantially larger (36.75% in 1980 and 39.58% in 2008) than the same probability computed for the overall economy (8.24% in 1980 and 11.37% in 2008).[18]

4.3 Workers' Problem

Citizens of skill type i who live in city k have preferences over consumption of the tradable good c_g^{ik}, the amount of housing H^{ik} and consumption of services c_n^{ik}. We assume the latter is a CES bundle of home services c_h and market services c_s, which are assumed to be imperfect substitutes with an elasticity of substitution equal to $\gamma > 1$.[19] More precisely, a worker of

[18] While in 1980 cities belonging to the top and bottom 33% are in this respect almost symmetric (with 0.05% difference between city groups), we highlight that employment shares of the non-tradable sector increase faster in large cities. That is, the spatial difference in the growth of non-tradables increases with the difference in city size: while it is 0.83% when comparing top versus bottom 33%, it reduces to 0.40% when comparing top versus bottom 50% and increases to 1.13% when comparing top vs bottom 25%. Thus, the spatial and overtime pattern of the employment shares in the non-tradable sector follows closely that of the low-skilled occupations as documented in Section 3.

[19] See Rogerson (2007) and Ngai and Pissarides (2011).

skill i living in city k has the following preferences

$$\begin{aligned} U^{ik} &= \left(H^{ik}\right)^{\alpha}\left(c_g^{ik}\right)^{\omega}\left(c_n^{ik}\right)^{1-\omega-\alpha} \\ c_n^{ik} &= \left(\psi\left(c_s^{ik}\right)^{\frac{\gamma-1}{\gamma}}+(1-\psi)\left(c_h^{ik}\right)^{\frac{\gamma-1}{\gamma}}\right)^{\frac{\gamma}{\gamma-1}}, \end{aligned} \tag{7}$$

where c_j, with $j = g,\ n,\ s,\ h$, represents consumption of goods, services, market services and home services, respectively. We impose $\alpha + \omega < 1$ and $\psi \in (0, 1)$.

Home services are produced within the household according to the technology

$$c_h^{ik} = A_h l^{ik}, \tag{8}$$

where $l^{ik} \in (0, 1)$ is the fraction of time an agent of skill i in city k devotes to work at home, thus, $1 - l^{ik}$ being the fraction of time dedicated to work in the market. We assume that home productivity is invariant across skills and locations. The budget constraint for workers of ability i living in city k is

$$p_g c_g^{ik} + p_s^k c_s^{ik} + p_H^k H^{ik} = w^{ik}(1 - l^{ik}), \tag{9}$$

where p_s^k and p_H^k are, respectively, the price of market services and housing, which are both location-specific and, therefore, indexed by k. Instead, the price of the tradable good, p_g, is the same in the whole economy. In what follows, we choose good g as the numeraire and, therefore, set $p_g = 1$. Workers of skill i living in city k solve the following problem

$$\begin{aligned} \max_{c_g^{ik}, c_s^{ik}, c_h^{ik}, l^{ik}} U^{ik} &= \left(H^{ik}\right)^{\alpha}\left(c_g^{ik}\right)^{\omega}\left(\left(\psi\left(c_s^{ik}\right)^{\frac{\gamma-1}{\gamma}}+(1-\psi)\left(c_h^{ik}\right)^{\frac{\gamma-1}{\gamma}}\right)^{\frac{\gamma}{\gamma-1}}\right),^{1-\omega-\alpha} \\ s.t. &: c_g^{ik} + p_s^k c_s^{ik} + p_H^k H^{ik} = w^{ik}(1 - l^{ik}), \\ c_h^{ik} &= A_h l^{ik}. \end{aligned}$$

First-order conditions imply the following demand functions

$$l^{ik} = \frac{1-\omega-\alpha}{1+\left(\frac{w^{ik}}{A_h p_s^k}\right)^{\gamma-1}\left(\frac{\psi}{1-\psi}\right)^{\gamma}}, \tag{10}$$

$$c_h^{ik} = A_h \frac{1-\omega-\alpha}{1+\left(\frac{w^{ik}}{A_h p_s^k}\right)^{\gamma-1}\left(\frac{\psi}{1-\psi}\right)^{\gamma}}, \tag{11}$$

$$c_s^{ik} = A_h \frac{(1-\omega-\alpha)\left(\frac{w^{ik}}{A_h p_s^k}\frac{\psi}{1-\psi}\right)^{\gamma}}{1+\left(\frac{w^{ik}}{A_h p_s^k}\right)^{\gamma-1}\left(\frac{\psi}{1-\psi}\right)^{\gamma}}, \tag{12}$$

$$c_g^{ik} = \frac{\omega w^{ik}}{p_g}, \tag{13}$$

$$H^{ik} = \frac{\alpha w^{ik}}{p_H^k}. \tag{14}$$

Equation (10) shows that when $\gamma > 1$, that is, when non-tradables and home production are imperfect substitutes, labor supply at home is a negative function of $\frac{w^{ik}}{A_h}$, which is the implicit price of home production for the agent. As long as wages are increasing in skills, high-skilled workers devote less time to home work and, given the common home production technology across agents, produce and consume less home services with respect to workers with lower skills levels. In addition, by differentiating Equation (12) with respect to w^{ik} we can compute the income elasticity of non-tradable services for agent of type i

$$\frac{\partial c_s^{ik}}{\partial w^{ik}}\frac{w^{ik}}{c_s^{ik}} = \frac{\gamma+\left(\frac{w^{ik}}{A_h p_s^k}\right)^{\gamma-1}\left(\frac{\psi}{1-\psi}\right)^{\gamma}}{1+\left(\frac{w^{ik}}{A_h p_s^k}\right)^{\gamma-1}\left(\frac{\psi}{1-\psi}\right)^{\gamma}} > 1. \tag{15}$$

This expression is larger than one when $\gamma > 1$, meaning that, when wages are increasing in skills, high-skilled workers consume a larger fraction of their income on this type of consumption with respect to workers of lower skill levels.[20] This implies that, *ceteris paribus* (i.e. for given skill premia), an increase in the spatial concentration of high-skilled workers results into an increase in the local demand for non-tradable services larger than for the other goods, because these agents display the largest consumption share of substitutable services, as suggested by Equation (15). Due to market clearing and the production function (5), such an increase in the spatial concentration of the high-skilled results in an increase of labor demand for low-skilled workers larger than for the other types of workers. This is the channel of *consumption spillovers*, which contributes to the emergence of spatial polarization in the model. Also, note that the income elasticity in Equation (15) grows with the value of

[20] Technically, Equation (15) provides the elasticity of non-tradable services to the *total income* of the agent, given by the market wage $w^{ik}(1-l^{ik})$ plus the implicit home wage, which, due to competitive markets, is given by $w^{ik}l^{ik}$. Note that we obtain an income elasticity larger than one in equilibrium regardless of the fact that the model displays homothetic preferences. This is because an increase in the market wage w^{ik} per unit of time also increases the opportunity cost of working at home, and so the implicit price of home production. Thus, there is an increase in substitutable services more than proportional to the increase in w^{ik}. Note also, from Equations (13) and (14), that the income elasticity of tradables and housing is always equal to one.

γ, such that the larger the elasticity of substitution between market and home services, the stronger the channel of consumption spillovers.

4.4 Equilibrium

The equilibrium is defined as a set of prices $\left(p_s^k, p_H^k\right)$, a set of wages $\left(w^{lk}, w^{mk}, w^{hk}\right)$, a choice k of a living location for each agent, a set of hours worked in the g sector $\left(e_g^{lk}, e_g^{mk}, e_g^{hk}\right)$, hours worked in the s sector $\left(e_s^{lk}\right)$, a set of time allocations for each type of consumer $\left(l^{lk}, l^{mk}, l^{hk}\right)$, a set of tradable good consumption for each type of consumer $\left(c_g^{lk}, c_g^{mk}, c_g^{hk}\right)$, a set of substitutable services consumption for each type of consumer $\left(c_s^{lk}, c_s^{mk}, c_s^{hk}\right)$ and a set of consumption vector for housing for each type of consumer $\left(H^{lk}, H^{mk}, H^{hk}\right)$ for each city $k = 1,\ 2$, such that:

- given wages and prices, and her choice of living location, the agent of skill i maximizes utility (7) subject to her budget constraint (9) in the chosen living location k and her home production technology constraint (8);
- given wages and prices, the representative firm in sector j in location k maximizes profits;
- labor markets for each skill i in each city k clear;
- the housing market in each city k clears;
- the market for substitutable services in each city k clears;
- the economy-wide market for tradable goods clears.

5 Quantitative Analysis

In Section 3 we document that employment polarization has been stronger in larger than in smaller cities. Hence, our analysis supports the emergence of spatial employment polarization after 1980. The aim of this section is to use a calibrated version of the model to investigate the role of technological change in generating this phenomenon.

There are two types of technological change that can generate spatial employment polarization in the model. The first is SBTC. Cerina et al. (2021b) note that the take-off in the skill premium coincides with the timing of employment polarization in the U.S. They show that SBTC, a typical driver of the increasing skill premium, can generate employment polarization in a general equilibrium setting through consumption spillovers. SBTC increases the productivity and so the wage of the high-skilled, who work little at home and purchase a substantial amount of market services. In the spatial equilibrium model presented in this paper, faster SBTC in a city relative to another implies that the first city attracts more high-skilled workers who, through consumption spillovers and extreme-skill complementarity, also attract more low-skilled workers to that location. The second type of technological change that can potentially generate spatial employment polarization is TFP growth in the tradable sector. Eeckhout et al. (2014) show that with extreme-skill complementarity, a city with a larger TFP displays a skill distribution with a larger fraction of high- and low-skilled workers with respect to a city with a smaller TFP. Thus, we allow for a differential evolution of TFP in the two cities coupled with a value of λ different from one. Lastly, for completeness we also allow for spatial changes in the TFP growth of non-tradables in the quantitative analysis.

Note that, while allowing for a *potentially different* evolution of technology in the two cities, we are not imposing any restriction of the growth of SBTC and TFP across cities. Thus, the calibration itself provides a measure of the differential technological change needed to generate stronger employment polarization in larger cities in the model. Next, by using the calibrated model we run counterfactual exercises to quantify the role of each type of technological change in generating the phenomenon.

5.1 Calibration

The quantitative exercise is set up as a horse-race between different types of technical change in explaining the spatial differences in employment polarization. We thus calibrate the model such that, given the types of technological change that we allow, it replicates two spatial equilibria at different points in time, namely the 1980 and the 2008 U.S. economies. In the two equilibria all preference and technology parameters are imposed to be the same except for the levels of SBTC and TFP in the two market sectors, which are allowed to grow differentially across cities according to the growth rates $g_{a^{hk}}$, $g_{A_g^k}$, $g_{A_s^k}$, where g indicates the total growth rate between 1980 and 2008 in city k of the variable at the subscript. We impose the two cities in the model to be symmetric in the 1980 equilibrium. This implies that all technological parameters are the same in the two cities in 1980.[21] The symmetry assumption reduces the number of parameters to be calibrated, by exploiting the fact that the quantitative exercise is designed to study the differential evolution of employment shares across cities since 1980. In this light, we do not require the calibrated model to account for the differences in the initial conditions between the two representative cities.

We first adopt the following normalizations/restrictions:

[21] In particular, we have $a^{h1} = a^{h2}$, $A_g^1 = A_g^2$, $A_s^1 = A_s^2$. All other technology parameters are the same both across cities and over time.

- Productivity of low-skilled workers is normalized to one, $a^l = 1$;
- The amount of land in each location is normalized to one, $H = 1$;
- Following the evidence in Bridgman (2016) there is no home productivity change between 1980 and 2008, and we normalize it to one in both periods, $A_{h,1980} = A_{h,2008} = A_h = 1$;
- We do not allow market TFP to decline in any sector, as the calibration could, in principle, deliver negative TFP growth in low-skilled services to better match the allocation of low-skilled workers across cities.

The elasticity of substitution between home production and non-tradable market services γ is key for the emergence of consumption spillovers. Following the discussion in Rogerson (2007), we set its value to 5.[22] Next, we obtain the values of α and ω by computing average consumption shares in housing and tradable goods between 1980 and 2008 using NIPA data and re-scaling them to take into account that, by introducing home produced services in the utility function, we have to consider the concept of *extended* total consumption expenditure in the data, i.e. the value of market consumption plus the market value of home production.[23] The nominal value of home production is taken from estimates in Bridgman (2016). This procedure gives a value of ω equal to 0.52, and of α equal to 0.13. The relative supply of skills (i.e. the aggregate skill distribution) in 1980 and 2008 is taken from U.S. Census data. The definition of the low-, middle- and high-skilled is the same as in Section 3. Low-skilled workers are those working in service occupations, high-skilled workers those in professional or managerial occupations and middle-skilled workers those in all remaining occupations.[24] Hence, following these definitions, we first normalize to one total population in 1980 ($\sum_i N^i_{1980} = 1$). Then, we compute the population growth rate g_N between 1980 and 2008 and impose that $\sum_i N^i_{2008} = 1 + g_N$. Finally, we use these restrictions, together with

[22]There are several studies providing estimates of the elasticity of substitution between home services and *total market consumption* (these estimates range from 1.8 as in Aguiar and Hurst, 2007, up to 2.5 as in Rupert et al., 1995, and McGrattan et al., 1997). In contrast, we are aware of only two other papers that calibrate the elasticity of substitution between home services and market substitutes exclusively. The first is Olivetti (2006) who finds a value of 4, the second is Ragan (2013) who uses a value of 6.66. We run alternative calibrations with these values of the elasticity of substitution which, as expected, deliver a smaller and larger role of consumption spillovers in generating spatial polarization. However, even in the most conservative case of $\gamma = 4$ consumption spillovers remain quantitatively relevant. Results are available upon requests. In general, several works discuss how the elasticity of substitution between home production and market services should be substantially higher than the one between home production and total market consumption. See for instance the discussion in Rogerson (2007), Ngai and Pissarides (2011) and Moro et al. (2017).

[23]See Moro et al. (2017) for a discussion of the concept of extended total consumption expenditure.

[24]As detailed in Appendix A, we exclude agriculture and military occupations.

the aggregate shares of low-, middle- and high skilled workers in 1980 and 2008 to obtain the values of $\{N^i_{1980}\}_{i=l,m,h}$ and $\{N^i_{2008}\}_{i=l,m,h}$. In doing so we are taking aggregate polarization as given.[25]

There are thirteen parameters left: (1) weight in preferences $\{\psi\}$, (2) productivity parameters $\{a^m,\ a^h,\ A_g,\ A_s\}$, (3) production parameters $\{\eta, \lambda\}$ and (4) technological change $\{g_{a^{hk}},\ g_{A^k_g},\ g_{A^k_s}\}_{k=1,2}$.[26] To calibrate these parameters we require the model to match thirteen data moments. Below we outline the general strategy by discussing how the various moments inform on the identification of specific parameters:

- Weight in preferences $\{\psi\}$ (1 target): The parameter ψ governs the relative weight of home and market services in preferences. We use average market hours (which equal total hours less home hours) of high-skilled workers in the economy in 1980 to identify this parameter.[27]

- Factor-specific productivity parameters $\{a^m,\ a^h\}$ (2 targets): The aggregate wage premium of high- to low-skilled workers in 1980 pins down a^h. As the middle-skilled productivity a^m is fixed over time and equal across cities, we use the wage premium of middle- to low-skilled workers in 1980 to pin down this productivity.

- Total-factor productivity parameters $\{A_g,\ A_s\}$ (2 targets): TFP in non-tradables, A_s, determines how many hours the economy needs in the production of non-tradable services. If A_s is low, more workers need to be allocated to non-tradable production for a given level of demand. Wages of low-skilled workers equalize across sectors, therefore the 1980 aggregate employment share of the low-skilled in tradables pins down A_s. Relative prices of tradables to non-tradables are governed by A_g and A_s, and these prices are crucial in determining the expenditure share allocated to non-tradable consumption. Thus, having pinned down A_s, the consumption share of non-tradables in 1980 identifies A_g.

[25] Note that this is consistent with the aim of our quantitative exercise, which is that of accounting for the differential patterns in employment polarization across cities. We stress that one could extend the current model by allowing aggregate shares of high-, middle- and low-skilled workers to be endogenized through an education and/or occupational decision, and account for the emergence of aggregate polarization through the same mechanisms at work in our model. For a model in which SBTC can generate employment polarization in a multi-sectoral environment with a home/market work decision see Cerina et al. (2021b).

[26] The value of a^h represents the common value of a^{hk} across cities in 1980, as we impose a symmetric equilibrium in that year. In 2008, the corresponding values in the tradable production functions of the two cities are $a^h(1 + g_{a^{h1}})$ and $a^h(1 + g_{a^{h2}})$ respectively for city 1 and city 2. The same reasoning applies to A_g and A_s.

[27] We use data from Aguiar and Hurst (2007) to determine total available working hours. Using their aggregate time us data, we find total hours worked are roughly constant from 1985 to 2003 at 52 hours per week.

- Production parameters $\{\eta, \lambda\}$ (2 targets): From the first order conditions of the firm producing the tradable good we can express the relative wage of high- to low-skilled workers as a function of η in which λ does not appear. Therefore, the wage premium of high- to low-skilled workers in 2008 can be used to pin down η. The first order conditions of the firm also show that the relative wage of middle- to low-skilled workers is a function of both η and λ. With η pinned down, the wage premium of middle- to low-skilled workers in 2008 allows us to identify λ.

- Technological change in the tradable sector $\{g_{a^{hk}},\ g_{A_g^k}\}_{k=1,2}$ (4 targets): differential changes in employment shares by cities directly capture differences across cities in terms of technological change. That is, the differential change in employment shares of the high-skilled in the two cities pins down the difference in growth rates of high-skilled productivity, $(g_{a^{h1}} - g_{a^{h2}})$, and the differential change in employment shares of middle-skilled in the two cities pins down the difference in growth rates of TFP, $\left(g_{A_g^1} - g_{A_g^2}\right)$. To pin down individual growth rates, we note that a change in the demand for high-skilled workers affects the share of tradable to non-tradable low-skilled workers. Therefore, for given elasticity parameters η and λ, the change in the share of low-skilled workers in the tradable sector across cities informs the level of one of the $g_{a^{hk}}$ values. Similarly, aggregate growth in tradable consumption is informative of $g_{A_g^k}$ in one of the two cities.

- Technological change in the non-tradable sector $\{g_{A_s^k}\}_{k=1,2}$ (2 targets): The change in the relative price of housing pins down the differential growth in non-tradable TFP, $\left(g_{A_s^1} - g_{A_s^2}\right)$, as utility of low-skilled agents equalizes across cities in equilibrium. Given the differential between the two sectors in TFP growth, aggregate growth in non-tradable consumption identifies growth of TFP in non-tradables in one of the two cities.

We use a method of simulated moments (MSM) to match the thirteen parameters to the thirteen moments concurrently, by minimizing the distance between data targets and model moments.[28] Parameter values of the benchmark calibration are reported in Table 2. Table 3 reports the data targets and the corresponding values produced by the calibrated model.

5.2 Results

Despite its parsimonious structure, the model does a good job at replicating the data targets. In particular, the calibration matches perfectly the difference between the two cities in the

[28]See McFadden (1989).

Table 2: Model Parameters

Preferences			
α	ω	γ	ψ
0.13	0.52	5	0.19

Technology											
λ	η	a^h	a^m	$g_{a^{h1}}$	$g_{a^{h2}}$	A_g	$g_{A_g^1}$	$g_{A_g^2}$	A_s	$g_{A_s^1}$	$g_{A_s^2}$
1.05	0.76	3.94	3.13	22.6%	25.1%	1.32	76.9%	85.4%	2.19	0%	1.4%

Note: The reported values of a^h, A_g and A_s correspond to the common value across cities in 1980.

change in the shares of the three types of workers between 1980 and 2008 (i.e. stronger polarization in city 2 relative to city 1). Thus, the values of the calibrated parameters in Table 2 provide an assessment of the role of technology in generating spatial polarization in the model. First, we note that both SBTC and TFP in tradables grow over time in both cities. This suggests that both types of technological change are key for the model to match the data targets. Second, there is faster growth of both SBTC and TFP in tradables in larger cities over time. This suggests that both types of technological change are important to generate stronger polarization in larger cities. Third, the predicted value of λ is larger than 1, which suggests that extreme-skill complementarity contributes to explain the spatial employment polarization observed in the data.

5.3 Counterfactuals

5.3.1 SBTC versus TFP growth

We now describe three counterfactuals to disentangle the effect of SBTC and that of TFP growth in generating spatial employment polarization. In order to do this, we allow spatial differences of one type of technological change at a time. In the first counterfactual, we impose the same growth of TFP in tradables and non-tradables between 1980 and 2008 in both cities, which is set to the average growth between the two cities in the benchmark calibration, so that the only source of spatial employment polarization is city-specific SBTC. More precisely, unlike the benchmark calibration reported in Table 2, we set $g_{A_g^1} = g_{A_g^2} = 81.2\%$ and $g_{A_s^1} = g_{A_s^2} = 0.7\%$ while we keep the value of $g_{a_h^1} = 22.6\%$ and $g_{a_h^2} = 25.1\%$. In the second and in the third exercise we apply the same approach for TFP growth in tradables (setting $g_{A_s^1} = g_{A_s^2} = 0.7\%$ and $g_{a_h^1} = g_{a_h^2} = 23.9\%$ but $g_{A_g^1} = 76.9\%$ and $g_{A_g^2} = 85.4\%$) and

Table 3: Model's fit

Moment		Data	Model
Diff. in change in emp. shares by cities	Low-skilled	0.70%	0.70%
	Middle-skilled	-2.43%	-2.43%
Aggregate wage premia	Medium/Low 1980	1.39	1.34
	Medium/Low 2008	1.44	1.32
	High/Low 1980	1.97	2.02
	High/Low 2008	2.49	2.14
Change in relative price of housing	$\frac{\left(p^2_{h,2008}/p^2_{h,1980}\right)}{\left(p^1_{h,2008}/p^1_{h,1980}\right)}$	1.15	1.16
Aggregate growth in consumption	Trad: $\frac{\sum_j \sum_k n^{jk}_{2008} c^{jk}_{g,2008}}{\sum_j \sum_k n^{jk}_{1980} c^{jk}_{g,1980}}$	2.71	2.71
	Non-trad: $\frac{\sum_j \sum_k n^{jk}_{2008} c^{jk}_{s,2008}}{\sum_j \sum_k n^{jk}_{1980} c^{jk}_{s,1980}}$	2.10	2.15
Aggr. consumption share non-trad in 1980	$\frac{\sum_j \sum_k n^{jk}_{2008} p^k_{s,2008} c^{jk}_{s,2008}}{\sum_j \sum_k n^{jk}_{2008} \left(p^k_{s,2008} c^{jk}_{s,2008} + c^{jk}_{g,2008}\right)}$	10.60%	6.90%
Diff. change emp. share of low-skilled in tradables/low-skilled workers	$\left(\frac{e^{l2}_{g2008}}{e^{l2}_{2008}} - \frac{e^{l2}_{g1980}}{e^{l2}_{1980}}\right) - \left(\frac{e^{l1}_{g2008}}{e^{l1}_{2008}} - \frac{e^{l1}_{g1980}}{e^{l1}_{1980}}\right)$	-1.86%	-1.86%
High-skilled market hours in 1980	$l^{hk} = l^h$	0.72	0.73
Aggregate employment share of low-skilled in tradables in 1980	$\frac{\sum_k e^{lk}_{g1980}}{\sum_i \sum_j \sum_k e^{ik}_{j1980}}$	7.50%	6.09%

non-tradables ($g_{a_h^1} = g_{a_h^2} = 23.9\%$ and $g_{A_g^1} = g_{A_g^2} = 81.2\%$ but $g_{A_s^1} = 0\%$ and $g_{A_s^2} = 1.4\%$) respectively. We focus on the three moments we are interested in: the difference in the change in employment shares in high-, middle- and low-skilled workers between large and small cities.

The SBTC counterfactual is displayed in the left panel of Figure 4.[29] When only SBTC differs between large and small cities, the difference in the change in the share of the three types of workers between the two cities is 32% for the low-skilled, 67% for the middle-skilled and 80% for the high-skilled with respect to the benchmark calibration.[30] This suggests that the existence of this type of technological change alone produces a large fraction of the asymmetry between the two cities. A key point here is that, while SBTC has a direct effect on the productivity of the high-skilled, it has a substantial impact also on the difference in the fraction of middle- and low-skilled across cities.

The right panel of Figure 4 reports the effect of allowing only differences in the TFP growth in tradables across cities. In this case, the difference in the change in the share of the three types of workers between the two cities is 18% for both high- and middle-skilled workers and 19% for low-skilled workers with respect to the benchmark calibration. Thus, with respect to SBTC, spatial differences in the growth of TFP in tradables have both a substantially smaller and a more homogeneous effect on the difference in the change in the share of the three types of workers between the two cities.

In addition to the above counterfactuals on SBTC and TFP in tradables, the bottom panel of Figure 4 also reports the effect of allowing only for differential growth of TFP in non-tradables in the two cities. In this case, the effect is mostly on low-skilled workers, with a difference in employment shares which is 48% of the benchmark. A small effect also remains for the middle-skilled, for whom the difference in employment shares generated is 15%. However, the effect on the upper tail is almost null (1%). Thus, this type of technological change alone cannot generate spatial polarization.

Taken together, these counterfactual exercises suggest that faster SBTC in large cities is the main driver of spatial polarization.

[29] For all counterfactuals the black bars represent these moments in the benchmark calibration while the white ones represent those in the counterfactual.

[30] For each counterfactual, we report the percentage of each bar accounted for by the counterfactual with respect to the benchmark.

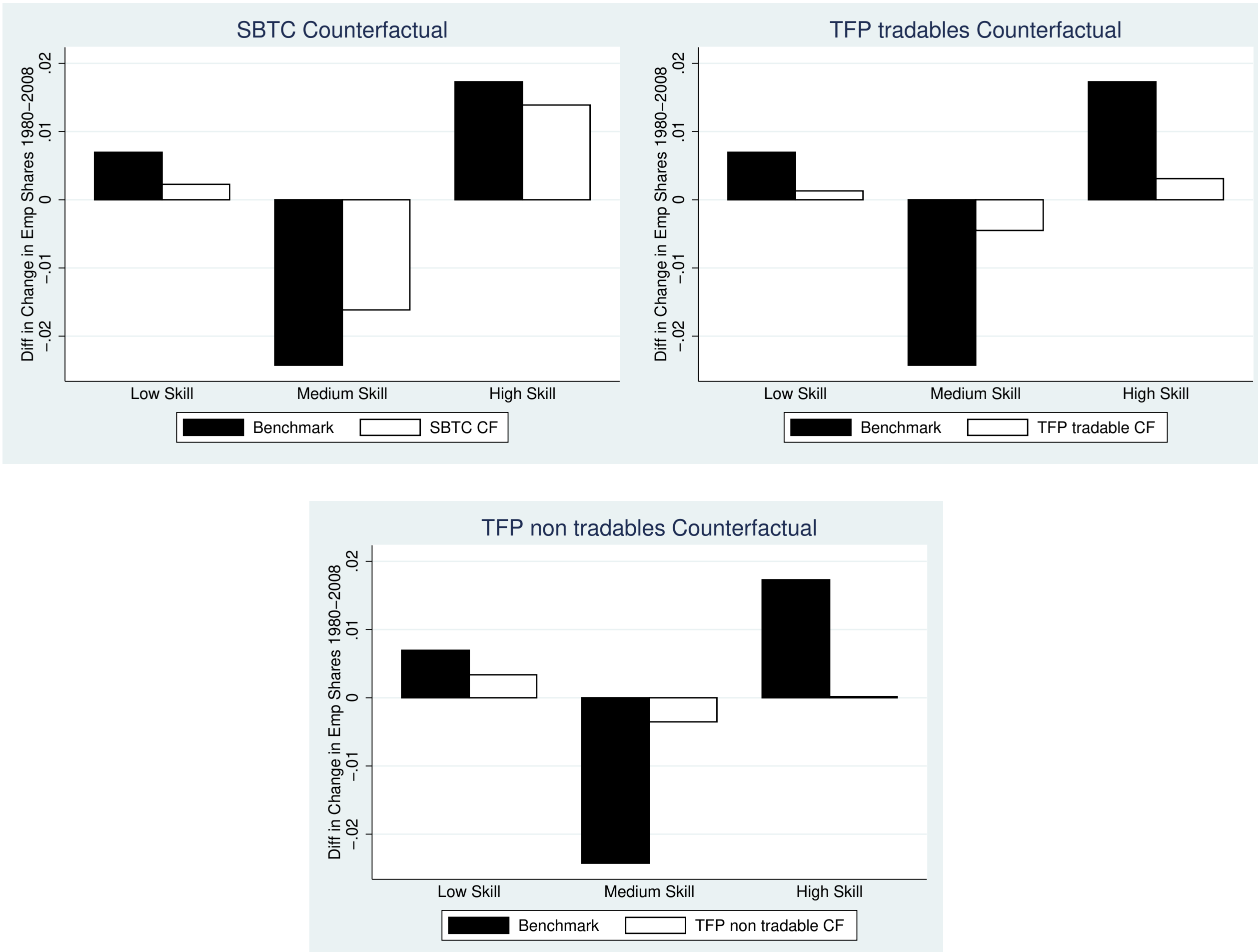

Figure 4: Counterfactual exercises. Black bars represent the benchmark calibration and white bars represent the counterfactual. Left panel: SBTC counterfactual. Right panel: tradables TFP counterfactual. Bottom panel: non-tradables TFP counterfactual.

5.3.2 Extreme-skill complementarity versus consumption spillovers

The value of λ which minimizes the distance between the data targets and the model moments is larger than one in our benchmark calibration. This value supports the existence of extreme-skill complementarity in the tradable sector, as posited by Eeckhout et al. (2014). In the latter, this mechanism - triggered by spatial differences in TFP - is the only driver of the relative increase in the share of high- and low-skilled individuals in larger cities. In our model we allow for another mechanism connecting the upper and bottom tail of the skill distribution, that of consumption spillovers. In this section, we ask what is the relative contribution of these two mechanisms in generating spatial differences in changes in the employment shares of low-skilled workers. To answer this question, our first step is to shut down the extreme-skill complemetarity channel by setting λ equal to one in 2008. This enables us to interpret the residual spatial polarization as generated by the joint effect of technological change and the existence of the non-tradable sector.[31]

The results are reported in Figure 5, in light gray bars. Relative to the benchmark calibration (black bars), the difference in the change in the share of the three types of workers between the two cities is reduced by 31% for the low-skilled, 32% for the middle-skilled and 32% for the high-skilled. A first conclusion that can be drawn from this exercise is that the mechanism of extreme-skill complementarity in the tradable sector, while being quantitatively relevant, is not a necessary condition for spatial polarization to emerge. This result is in contrast with that of Eeckhout et al. (2014) who suggest that a channel based on low-skilled services in combination with home production is not quantitatively relevant in accounting for more dispersed skill distributions in larger cities.[32]

However, not all the residual spatial disparity at the bottom of the skill distribution (69% of the benchmark case) is due to consumption spillovers. As described above, employment shares at the bottom of the skill distribution grow faster in large cities also because of faster TFP growth in non-tradables therein. Hence, to quantify the role of consumption spillovers in generating the rise of employment shares at the bottom of the skill distribution, we perform the same counterfactuals described in Section 5.3.1 on SBTC and TFP growth in tradables, but now imposing $\lambda = 1$, and thereby removing extreme-skill complementarity. This exercise, by removing the direct role of TFP in non-tradables, allows us to interpret the residual spatial disparity in the bottom tail (with respect to the new benchmark in which $\lambda = 1$) as generated by the consumption spillovers mechanism only. Results are reported in Figure 5 in white bars for the SBTC counterfactuals (left panel) and the TFP counterfactual (right panel). To allow for an easy comparison, we also report in the same figure the results for the counterfactuals with $\lambda > 1$ (dark gray bars).

[31] We also performed a calibration exercise where we remove the target which identifies λ (the wage premium of middle- to low-skilled workers in 2008) and we exogenously fix this parameter to one. This exercise allows us to study the case in which only one channel is present in the model, that of consumption spillovers. The fit of this calibration is worse with respect to that of a calibration where we fix λ to its value in the benchmark calibration (1.05), which confirms that extreme-skill complementarity plays an empirically relevant role. Results are available upon request.

[32] We refer to footnote 5 for the list of differences between their approach and ours that lead to opposite conclusions regarding the role of consumption spillovers.

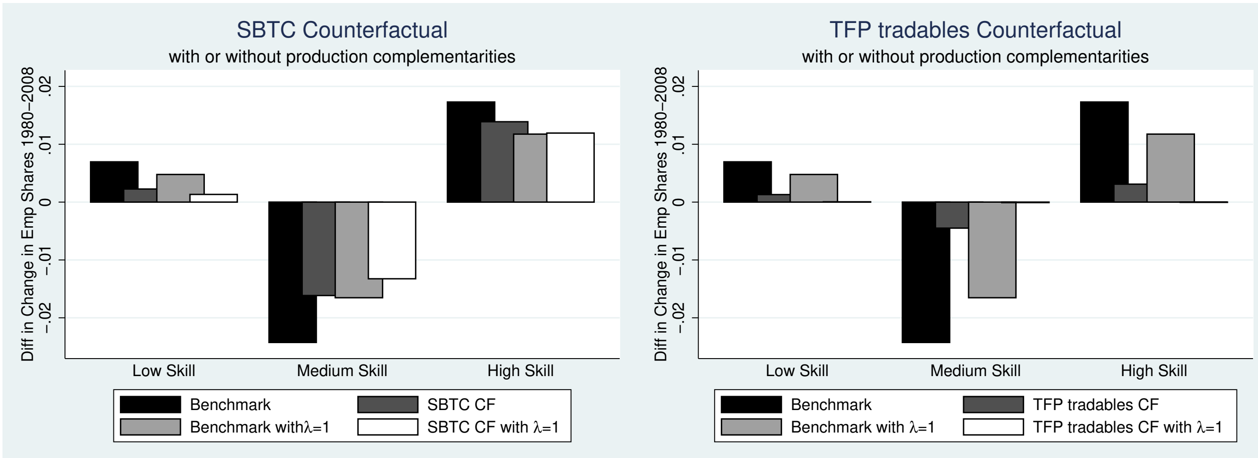

Figure 5: Counterfactual exercises. Black bars represent the benchmark calibration, dark gray bars represent the counterfactual allowing for spatial differences in SBTC (left panel) or TFP growth (right panel) only, light gray bars represent the counterfactual imposing $\lambda = 1$, and white bars represent the counterfactual allowing for spatial differences in SBTC (left panel) or TFP growth (right panel) only *and* imposing $\lambda = 1$.

As expected, the role of spatial disparities in TFP growth in tradables is canceled out without extreme-skill complementarity. By looking at the white bars in the right panel, it appears that when $\lambda = 1$ and only differences in TFP growth in tradables across cities are allowed, there is no spatial polarization, as the white bars are virtually zero for each skill. This is not the case for SBTC: imposing $\lambda = 1$, allowing only for differences in SBTC growth across cities as suggested by the benchmark calibration (25.1.0% in large cities, 22.6% in small cities), and imposing the same TFP growth both in tradables (81.2%) and non-tradables (0.7%), the remaining amount of spatial polarization is quantitatively relevant. More precisely, the residual spatial difference in the change of employment shares of high-, middle- and low-skilled workers is respectively 86%, 82% and 59% of the corresponding values in the SBTC counterfactual with the value of λ predicted by the benchmark calibration ($\lambda = 1.05$).[33] The counterfactual results suggest that the key type of technological change needed to generate consumption spillovers is the skill-biased one. Another way to look at the results is by considering that in a world in which consumption spillovers is the only mechanism linking the top and the bottom part of the skill distribution (i.e. when $\lambda = 1$), spatial differences in SBTC alone still explain around 28% of the difference at the bottom generated by all types on technological change together, 80% for the middle skilled and 102%

[33] Note that the percentages here refer to the fraction of the bar in the counterfactual "SBTC CF with $\lambda = 1$" relative to the counterfactual "SBTC CF".

for the high-skilled.[34] Since with extreme-skill complementarity ($\lambda > 1$) spatial differences in SBTC alone explain around 32% of the difference in the bottom tail (see Section 5.3.1), we conclude that consumption spillovers have a sizeable role in explaining the faster growth in the employment shares of low-skilled workers in large cities generated by faster SBTC therein, accounting for more than 85% of the total effect. This results points to the importance of the demand channel associated with the substitution between home and market services in generating spatial polarization.

5.3.3 Spatial policies

This section aims at further exploring the quantitative features of the benchmark calibration by implementing counterfactual policy exercises. We are interested in analyzing how different tax schemes across locations can affect the spatial sorting of workers of different skill levels. We imagine that a local policy-maker in the large city is willing to give a lump-sum subsidy to low-skilled workers therein, financing this program with the revenues generated by taxing high-skilled workers in that location.[35] We consider two kind of taxes through which this can be achieved: 1) a tax on wages of high-skilled workers; and 2) a tax on consumption of non-tradables of high-skilled workers.[36] To make the comparison meaningful, we set the two tax rates in order for the tax revenues to be equal to 3% of low-skilled workers' income in large cities in 2008. Results are reported in Figure 6. The two policies have similar impacts on high-skilled workers who, as expected, are less attracted to large cities. These cities remain relatively more attractive for high-skilled workers in the 2008 equilibrium, but the spatial difference in the change of the employment shares is reduced from 1.73% of the benchmark calibration to 0.95% and 1.10%, respectively.

[34] Compare "SBTC CF with $\lambda = 1$" and "$\lambda = 1$" in Figure 5.

[35] An alternative interpretation is that of a government willing to improve the equity-efficiency tradeoff of a taxation policy by implementing place-based redistribution policies. Gaubert et al. (2021) show that place-based taxation schemes can lower the efficiency cost of redistribution across heterogenous households and generate welfare gains that income-based redistribution alone cannot attain.

[36] In a recent paper, Bastani et al. (2020) use a quantitative model calibrated to the U.S. economy to argue that taxing child care expenditures of high-skilled workers can be welfare-improving.

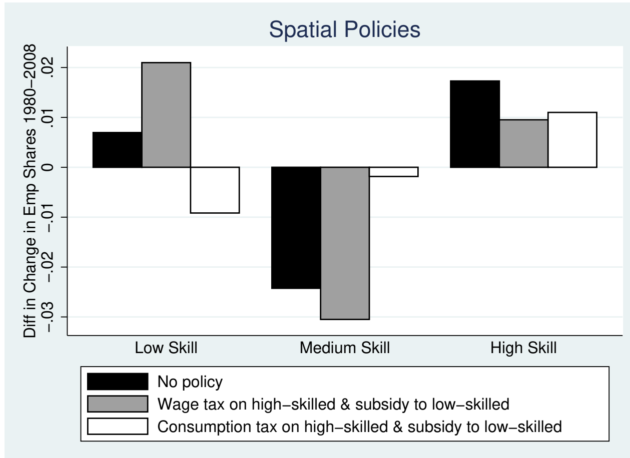

Figure 6: Spatial policies. Black bars represent the performance of the model in the benchmark calibration without policy, gray bars represent the policy with a wage tax for the high-skilled in the large city, white bars represent the policy with a tax on non-tradable consumption for the high-skilled in the large city.

The two local policies have instead an opposite impact on low-skilled workers. Large cities become more attractive for them with respect to the benchmark case when the subsidy is financed through the wage tax. In this case, the spatial difference in the change of the employment share of low-skilled workers increases from 0.70% to 2.10%. By contrast, when the subsidy is financed through the consumption tax, the share of low-skilled workers is skewed towards the small city in the 2008 equilibrium, such that the spatial difference in the change of their employment shares becomes negative (-0.92%). Also for middle-skilled, spatial differences are amplified with the wage tax while they are substantially dampened with the consumption tax. However, in contrast with the low-skilled, these differences retain the same sign as in the benchmark with both types of policies.

The opposite patterns of the spatial allocation of low-skilled workers are due to the way the two policies affect the two different channels linking the tails of the occupational skill distribution (extreme-skill complementarity and consumption spillovers). The wage tax on high-skill workers reduces the incentive for the latter to move to the large city, thereby indirectly reducing, *ceteris paribus*, both the marginal productivity of low-skilled workers in the tradable sector (by extreme-skill complementarity) and the labor demand for low-skilled workers in the non-tradable sector (by consumption spillovers). However, these worsened market conditions for low-skilled workers are more than compensated by the subsidy, making the large city even more attractive for them than in the benchmark case.

When the tax on non-tradable consumption is implemented, the negative effect on low-skilled workers is larger because the policy directly impacts the price of services. The tax on non-tradable consumption raises the cost of non-tradables, which leads to a sizeable reduction in the expenditure for these services by high-skilled workers (-23%). This fact, in turn, dampens labor demand for low-skilled workers in the non-tradable sector and induces them to move to the small city. As consumption spillovers represent the channel which is quantitatively more relevant in generating spatial polarization, as we discuss in Section 5.3.2, this tax policy offsets the attractiveness of the subsidy for low-skilled workers in the non-tradable sector: the negative value of the difference in the change of employment shares is entirely driven by them (-1.79%) while it remains positive for low-skilled workers in the tradable sector (+0.87%). To conclude, while consumption spillovers in the large city are only mitigated by the wage tax, they are entirely washed out by the consumption tax, such that the latter policy has the unintended effect of lowering welfare of the low-skilled in large cities and so their employment share in those locations.

6 Spatial workers polarization

Section 3 documents that the joint increase in the employment shares of high- and low-skilled occupations observed at the aggregate level in the U.S. between 1980 and 2008 has been stronger in larger cities, suggesting a spatial dimension of the phenomenon. In this section we ask whether spatial polarization at the *occupation* level is associated with another phenomenon, that of spatial polarization at the *worker* level. Specifically, we investigate whether the faster growth in the employment shares of high- and low-paid *occupations* in large cities is associated with a relatively stronger pull of high- and low-skilled *workers* respectively in those locations. This would suggest that spatial employment polarization induces a change in the spatial sorting of workers with heterogenous skills over time. In this section, we refer to the latter phenomenon as *spatial workers polarization*.

Indeed, while related, the two phenomena do not necessarily imply each other. To construct the *occupational skill distribution*, the employment polarization literature ranks occupations by their mean wage, and constructs occupational shares accordingly. By contrast, the *workers skill distribution* is constructed by considering some characteristics (either observable - like educational attainments or individual wages as in Hunt and Nunn, 2022 - or unobservables as in Eeckhout et al., 2014) at the worker level. This difference has two main implications: first, the dispersion in the skills of workers within an occupation does not play a role in generating evidence of employment polarization, because only the wage of the average worker matters for the skill rank of the occupation; second, two workers located in a similar quantile of the workers skill distribution (having for instance similar educational attainments or similar wages) might be located in substantially different positions of the occupational skill distribution if they are employed in occupations displaying large differences in mean wages.[37]

Being two intrinsically distinct concepts, the occupational and the workers skill distribution might evolve over time and across space according to different patterns. For instance, we could observe spatial employment polarization (as documented in Section 3) without concurrently observing spatial workers polarization. Consider for instance two vacancies opened in a large city, one for a high-skilled occupation and one for a low-skilled occupation. These could potentially be filled by two workers who abandon their middle-skilled occupations in the same large city. In this case, we would observe faster employment polarization in the large city relative to the small city, where no vacancies have been opened, but no spatial workers polarization. Instead, if the two vacancies are filled by two workers previously located in a small city, a low-skilled worker performing a low-skilled occupation, and a high-skilled worker performing a high-skilled occupation, then we would observe that spatial employment polarization is associated with spatial workers polarization, as the individual skill distribution of large cities would display a larger increase in the employment shares of high- and low-skilled workers. To put it differently, if spatial employment polarization is entirely driven by occupational sorting *within* cities (rather than by spatial sorting across cities) there should be no spatial polarization at the worker level. By contrast, if spatial employment polarization is jointly observed with spatial workers polarization, then the evidence would suggest that the change in the occupational structure across cities of different size gives rise to a change in the allocation of heterogeneously skilled workers across space.[38]

To investigate whether spatial employment polarization at the occupation and at the worker level are associated, we use the model's equilibrium, which provides a measure of skill at the worker level given by their indirect utility. The biunivocal relationship between indirect utility and individual skill is ensured by the equilibrium condition according to which workers of each skill are free to choose the location which ensures the highest utility. This strategy has several advantages. First, the individual measure of skill is derived from the same theory used to study spatial employment polarization in Sections 4 and 5. The only difference is that we now drop the three occupations/skills assumption, and allow for

[37]For instance, a relatively high-paid professional occupation like *dentists* (mean hourly wage 14.26 dollars in 1980), exhibits a high dispersion in its wage distribution such that the worker at the 90th percentile of the wage distribution of dentists earns almost 12 times more than the worker at the 10th percentile (36.05 versus 3.10). By contrast, a middle-skilled occupation like *postal clerks* (mean hourly wage 8.56) displays a wage distribution which is relatively concentrated around the mean, with the worker at the 90th percentile earning less than twice as much the worker at the 10th percentile (10.10 versus 6.04).

[38]Note that to formally relate spatial polarization at the occupation and at the worker level in the data would require micro data that allows us to track individuals across time and space. Such analysis is unfeasible using U.S. Census data.

a large number of skills in the economy, in such a way that the occupation dimension does not play a role, as typical for individual skill distributions based on workers' characteristics. Second, despite being model based, the measure allows us to easily exploit the whole micro evidence at the worker level, to derive detailed individual skill distributions at the city level at different points in time. Third, such measure allows for a direct comparison with the results in Eeckhout et al. (2014), who construct similar individual skill distributions for large and small cities in 2009 in a model without non-tradable goods.[39]

By using the first order conditions of the household's problem we obtain the indirect utility for a worker of skill i in city k, which is given by

$$U^{ik} = \Omega \left(p_H^k\right)^{-\alpha} \left(w^{ik}\right)^{\alpha+\omega} \left(1 + \left(\frac{\psi}{1-\psi}\right)^{\gamma} \left(\frac{w^{ik}}{A_h p_s^k}\right)^{\gamma-1}\right)^{\frac{1-\omega-\alpha}{\gamma-1}} \tag{16}$$

and where $\Omega = \alpha^{\alpha}\omega^{\omega}\left(1-\omega-\alpha\right)^{(1-\omega-\alpha)}\left(1-\psi\right)^{\frac{\gamma(1-\omega-\alpha)}{\gamma-1}}\left(A_h\right)^{(1-\omega-\alpha)}$.

The assumption of workers mobility ensures that utility of two workers of the same type is the same across locations ($U^{i1} = U^{i2}$). Thus, there is a one-to-one mapping between equilibrium utility and skill level for the worker of type i in any city k. We can interpret (16) as the measure of skill implied by the model and use it to construct a model-based distribution of skills in a particular year by using data on p_H^k, p_s^k and w^{ik}. The model-based measure of skills (16) only requires a subset of model parameters to be computed, which we take from the calibration in Section 5.

[39]When $\alpha + \omega = 1$ our setting coincides (except for location-specific productivity of high-skilled workers) with that of Eeckhout et al. (2014), in which there is no home production and no market production of services.

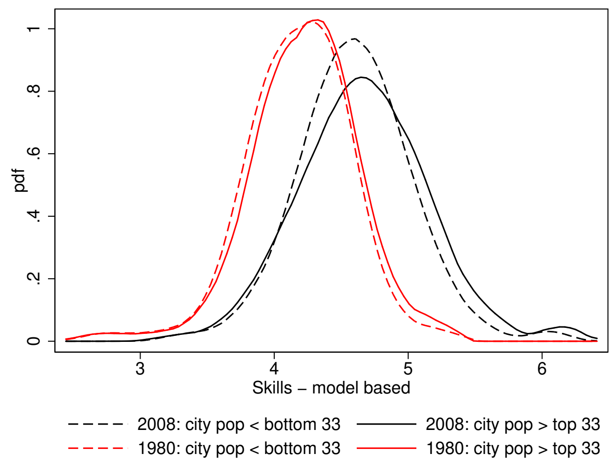

Figure 7: Skill distribution (logarithm of Equation 16) in 1980 (red) and 2008 (black) in small (dashed lines) and large (thick lines) cities. *Top vs bottom 33%* grouping.

Figure 7 reports the skill distribution across time and space showing cities in the *top vs bottom 33%* categorization used in Section 3 and 5. In 1980 (red lines) the skill distributions of large and small cities are remarkably similar, although first order stochastic dominance for large cities emerges. For 2008 (black lines), the evidence is substantially different, with large cities displaying a larger dispersion of the skill distribution with respect to small cities. To check whether this finding is affected by the definition of "large" and "small", we perform a set of quantile regressions aiming at analyzing to which extent the divergence of the individual skill distribution is affected by relative city size in 1980 and 2008. Formally, assuming a linear relation between the individual characteristic x^{ik} (representing skill U^{ik}), and population (S^k) in location k, we estimate the following specification for each quantile τ:

$$Q_\tau(x^{ik}|S^k) = \beta_0(\tau) + \beta_1(\tau)S^k,$$

where consistent estimators of $\beta_0(\tau)$ and $\beta_1(\tau)$ are obtained by minimizing an asymmetrically weighted sum of absolute errors. We perform this exercise for the skill distribution in 1980 and 2008. Both exercises are represented in a figure with two panels: on the left one we plot five quantiles of the distribution (the 10th, the 25th, the median, the 75th and the 90th) against city size, while in the right panel we plot the coefficient of each quantile against its quantile rank. This procedure shows how the effect of city size on the shape of the skill distributions changes from 1980 to 2008.

Figure 8 reports the result for the 1980 skill distribution. There is a clear first order stochastic dominance of large versus small cities. However, there is no divergence across city size in 1980. Coefficients of the quantile regressions are slightly positive and similar for each quantile (except the very last quantiles), suggesting that quantiles increase proportionally as city size increases. Thus, the quantile regression confirms that in 1980 large cities do not display a more dispersed workers skill distribution.

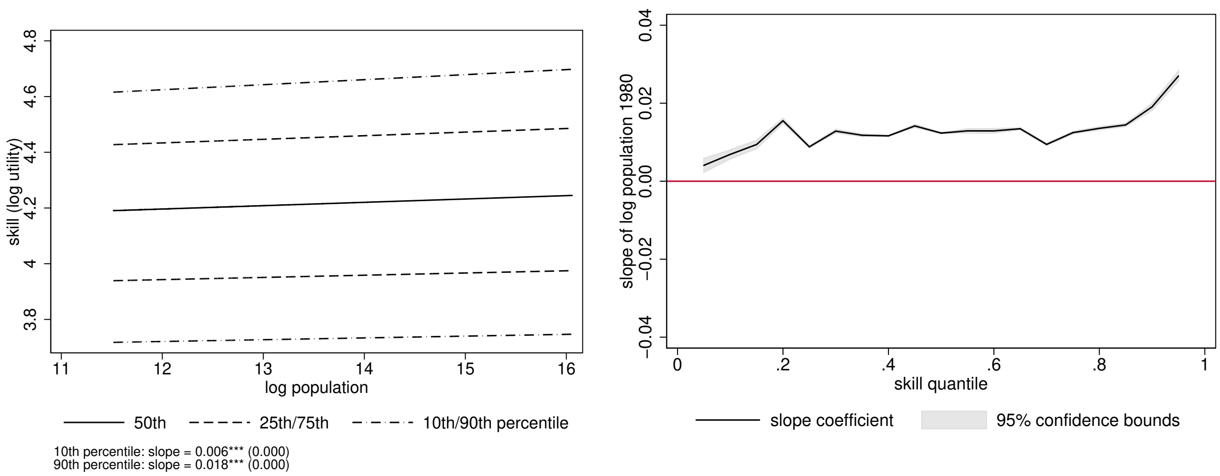

Figure 8: Quantile regression of utility on population in 1980 (i.e. model-based skill measure): left, five selected quantiles; right, estimated slope for all quantiles.

Figure 9 reports the results of quantile regressions for the skill distribution in 2008. The right panel shows that slopes are increasing with the quantile rank, being negative up to the 25th percentile and positive otherwise. This confirms the visual result of Figure 7 for the year 2008: lower quantiles decrease with city size while the opposite happens for higher quantiles (left panel).

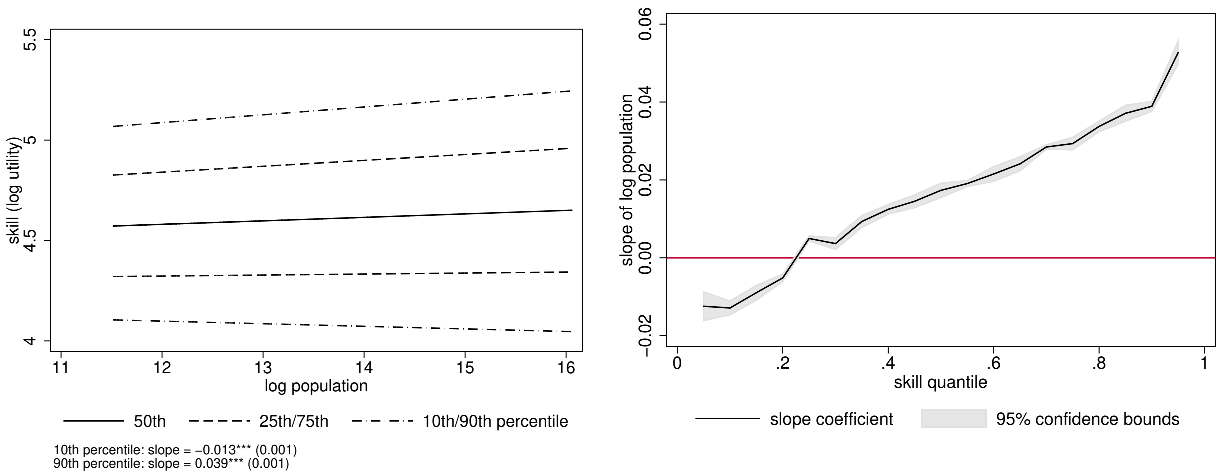

Figure 9: Quantile regression of utility in 2008 (i.e. model-based skill measure) on population in 1980: left, five selected quantiles; right, estimated slope for all quantiles.

To sum-up, the evidence in this section confirms that spatial employment polarization occurred together with spatial workers polarization in the U.S. The evolution in the spatial patterns of the individual workers distribution therefore mimics the evolution in the spatial patterns of the occupational skill distribution presented in Section 3, suggesting that spatial polarization at the occupational and at the worker level occurred together and that, in both cases, spatial polarization increases with the difference in city size. In addition, our results confirm the findings in Eeckhout et al. (2014), who show that in 2009 the average and the median worker have similar skill in large and small cities but, crucially, the skill distribution in larger cities has fatter tails both at the top and at the bottom of the distribution. Thus, our observation for the year 2008 is consistent with their results for 2009 in a model without home production and substitutable services.[40]

Finally, we stress that our results suggest that the emergence of spatial workers polarization in large cities is a relatively recent phenomenon, that emerged during the employment polarization era (i.e. post-1980). This is confirmed by the analysis of the skill distribution in 1960. In Figure 10 we document that, as for 1980, the skill distribution in 1960 is similar in small and large cities.[41] The larger dispersion in 1980 relative to 1960 is an aggregate phe-

[40] In Appendix B we provide further evidence on the spatial polarization at the worker level by focusing on an observable measure of skills, that of educational attainment. Specifically, we document that large cities display a relatively faster increase in the share of both highly educated (college+) and poorly educated (less than high-school) workers relative to small cities.

[41] We use city-level prices for non-tradables from Carrillo et al. (2014) as a measure of p_s^k in constructing the skill distributions of 1980 and 2008, but a similar procedure cannot be applied to the year 1960 due to a lack of data. To overcome this problem we use the first order condition of the model $p_s^k = w^{lk}/A_s^k$, which implies that the price of non-tradables in city k is proportional to the local wages in the non-tradable sector.

nomenon unrelated to size. Thus, spatial workers polarization should be related to changes in the economic structure that occurred after 1980. Crucially, the year 1980 is typically documented as the starting point of employment polarization in the U.S. (Acemoglu and Autor, 2011). To conclude, the evidence presented in this section suggests that spatial polarization in the workers skill distribution emerges since the 1980s. Coupled with the analysis of Section 3, it reinforces the view that spatial employment polarization induces a change in the spatial sorting of heterogenously skilled workers (i.e. spatial workers polarization), as the timing of the two phenomena is the same.

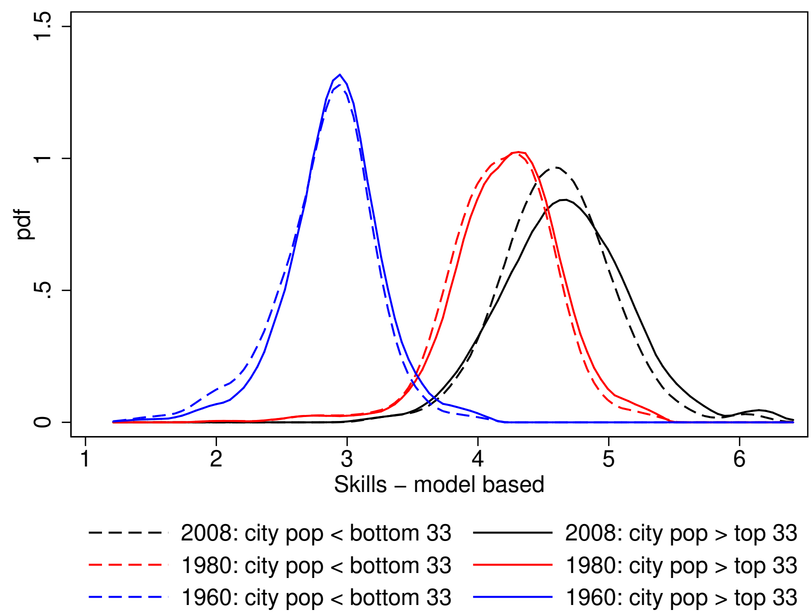

Figure 10: Skill distribution in 1960 (blue), 1980 (red) and 2008 (black) in small (dashed lines) and large (thick lines) cities. The figure compares metropolitan areas with population belonging to the top and bottom 33% of the population distribution in 1980.

7 Conclusions

In this paper we document that employment polarization is stronger in cities whose size is larger in 1980, and that the intensity of this phenomenon increases with city size. Importantly, we document that this pattern is driven by the extensive (heads) rather than

We then compute the average of the wages of all workers in the non-tradable sector (weighted by hours worked) for each of the k =218 metropolitan areas in the sample for the years 1960, 1980 and 2008. As we do not have a measure for A_s^k across cities in 1960, we choose to set $A_{s,1960}^k = 1$ for all cities. While this is an arbitrary choice, we use the same assumption, that is $A_{s,1980}^k = A_{s,2008}^k = 1$ for each city k, to compute the skill distributions for 1980 and 2008 appearing in Figure 7.

the intensive (hours) margin, and that the increase of employment shares at the bottom of the skill distribution is driven to a large extent by the sector producing services that are substitutable to home production.

To account for the patterns observed in the data, we build a spatial equilibrium model with location-specific skilled-biased technical change in the tradable sector, a low-skill intensive non-tradable sector and a home versus market labor decision. We calibrate the model using two groups of cities and three groups of skills. The benchmark calibration suggests that the role of both unbiased and biased technological change are quantitatively important and supports the existence of both consumption spillovers and extreme-skill complementarity in production of the tradable sector. We then perform a series of counterfactuals which show that faster skilled-biased technological change experienced in larger cities is responsible for a large fraction of spatial employment polarization.

Finally, we use the model to show that spatial polarization at the occupational level occurs together with spatial polarization at the worker level. This finding supports the idea that the increasingly different occupational structure of more versus less urban areas has been fueled by the sorting of both low- and high-skilled workers who have been largely attracted to large cities due to the relative increase in the labor demand for their skills.

References

Acemoglu, D. and Autor, D. (2011). Chapter 12 - skills, tasks and technologies: Implications for employment and earnings. volume 4, Part B of *Handbook of Labor Economics*, pages 1043 – 1171. Elsevier.

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

Autor, D. H. (2019). Work of the past, work of the future. *AEA Papers and Proceedings*, 109:1–32.

Autor, D. H. and Dorn, D. (2013). The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market. *American Economic Review*, 103(5):1553–97.

Bastani, S., Blomquist, S., and Micheletto, L. (2020). Child care subsidies, quality, and optimal income taxation. *American Economic Journal: Economic Policy*, 12(4):1–37.

Baum-Snow, N., Freedman, M., and Pavan, R. (2018). Why has urban inequality increased? *American Economic Journal: Applied Economics*, 10(4):1–42.

Baum-Snow, N. and Pavan, R. (2013). Inequality and city size. *Review of Economics and Statistics*, 95(5):1535–1548.

Bridgman, B. (2016). Home productivity. *Journal of Economic Dynamics and Control*, 71:60 – 76.

Carrillo, P. E., Early, D. W., and Olsen, E. O. (2014). A panel of interarea price indices for all areas in the United States 1982-2012. *Journal of Housing Economics*, 26:81–93.

Cerina, F., Moro, A., and Rendall, M. (2021a). A note on employment and wage polarization in the u.s.

Cerina, F., Moro, A., and Rendall, M. (2021b). The role of gender in employment polarization. *International Economic Review*.

Davis, D. R. and Dingel, J. I. (2019). A spatial knowledge economy. *American Economic Review*, 109(1):153–70.

Davis, D. R., Mengus, E., and Michalski, T. (2019). Labor Market Polarization and The Great Divergence: Theory and Evidence. *mimeo*.

Diamond, R. (2016). The determinants and welfare implications of us workers' diverging location choices by skill: 1980-2000. *American Economic Review*, 106(3):479–524.

Eeckhout, J., Pinheiro, R., and Schmidheiny, K. (2014). Spatial sorting. *Journal of Political Economy*, 122(3):554 – 620.

Gabaix, X. and Ioannides, Y. (2004). The evolution of city size distributions. In Henderson, J. V. and Thisse, J. F., editors, *Handbook of Regional and Urban Economics*, volume 4, chapter 53, pages 2341–2378. Elsevier, 1 edition.

Gaubert, C., Kline, P. M., and Yagan, D. (2021). Place-based redistribution. Technical report, National Bureau of Economic Research.

Giannone, E. (2017). Skill-Biased Technical Change and Regional Convergence. *mimeo*.

Glaeser, E. L. and Resseger, M. G. (2010). The complementarity between cities and skills*. *Journal of Regional Science*, 50(1):221–244.

Hunt, J. and Nunn, R. (2022). Has us employment really polarized? a critical reappraisal. *Labour Economics*, page 102117.

Leonardi, M. (2015). The effect of product demand on inequality: Evidence from the united states and the united kingdom. *American Economic Journal: Applied Economics*, 7(3):221–47.

Lindley, J. and Machin, S. (2014). Spatial changes in labour market inequality. *Journal of Urban Economics*, 79(C):121–138.

McFadden, D. (1989). A method of simulated moments for estimation of discrete response models without numerical integration. *Econometrica: Journal of the Econometric Society*, pages 995–1026.

McGrattan, E. R., Rogerson, R., and Wright, R. (1997). An equilibrium model of the business cycle with household production and fiscal policy. *International Economic Review*, 38(2):267–90.

Moretti, E. (2013). Real wage inequality. *American Economic Journal: Applied Economics*, 5(1):65–103.

Moro, A., Moslehi, S., and Tanaka, S. (2017). Does home production drive structural transformation? *American Economic Journal: Macroeconomics*, 9(3):116–46.

Ngai, L. R. and Pissarides, C. A. (2011). Taxes, social subsidies, and the allocation of work time. *American Economic Journal: Macroeconomics*, 3(4):1–26.

Olivetti, C. (2006). Changes in women's hours of market work: The role of returns to experience. *Review of Economic Dynamics*, 9(4):557–587.

Parkhomenko, A. (2021). Homeownership, inequality, and polarization. *Available at SSRN 3854352*.

Ragan, K. S. (2013). Taxes and time use: Fiscal policy in a household production model. *American Economic Journal: Macroeconomics*, 5(1):168–92.

Rogerson, R. (2007). Taxation and market work: is Scandinavia an outlier? *Economic Theory*, 32(1):59–85.

Rupert, P., Rogerson, R., and Wright, R. (1995). Estimating substitution elasticities in household production models. *Economic Theory*, 6(1):179–193.

A Appendix: Data treatment

This appendix discusses the data used in the paper with focus on comparability issues, as spatial boundaries of geographical statistical areas change over time.

A.1 Individual data

To present evidence of employment polarization and to construct information about workers of different skills (Section 6), we use the national 5-percent public-use micro data samples for the 1960 and 1980 Census of Population and the 1-percent American Community Survey for 2008. When constructing employment polarization figures, we use data for all individuals who report positive wages and salary income, considering both full and part-time workers. However, turning to the individual skill distribution analysis, in order to avoid any data mismeasurement on wages, and consistently with the literature, we restrict the sample to individuals that work at least 35 hours per week and 40 weeks per year. As in Autor and Dorn (2013) we exclude farmers and military occupations.[42] Also, following Eeckhout et al. (2014), we drop the lowest 0.5 percent of wages to eliminate likely misreported of wages close to zero. Instead of using the IPUMS version of the 1990 Census Bureau occupational classification scheme, we work with the balanced set of occupations for 1980 and 2008 used in Autor and Dorn (2013). As a result, the total number of full-time workers considered is $1,674,247$ in 1980 and $533,021$ in 2008 while, when dealing with employment polarization, total observations rise to $3,093,320$ in 1980 and $705,536$ in 2008.

A.2 Definition of non-tradables

To identify the non-tradable sector in the data, we follow Moro et al. (2017). Accordingly, from the 1990 Census classification (3 digits) we select the following industries: Bakery products; Miscellaneous personal services; Beauty shops; Eating and drinking places; Laundry, cleaning, and garment services; Taxicab service; Food stores, n.e.c.; Private households; Child day care services; Retail bakeries; Nursing and personal care facilities; Miscellaneous repair services; Educational services, n.e.c.; Residential care facilities, without nursing; Bus service and urban transit; Personnel supply services; Liquor stores; Barber shops.

[42] We exclude agricultural occupations because we are interested in two types of low-skilled occupations: i) those in the non-tradable services sector, which are the ones that are created through the consumption spillover mechanism; and ii) those in the tradable sector, that complement high-skilled occupations in production, like security, janitors, and reception services. As agricultural occupations do not appear to belong to these two groups, we drop them from the sample. For a discussion of the role of agricultural occupations in generating employment polarization see Cerina et al. (2021a).

A.3 Spatial boundaries

To analyze how the patterns of the distributions differ across city size, we need to match Census micro data to metropolitan areas. The main issue is that the variable "metro area" reports a combination of metropolitan area codes (MSA, primary MSA, central city or county) which has evolved considerably over time, and thus leads to difficulties in matching with PUMA codes or any other harmonized classification of cities. Thus, one issue is to define spatial boundaries of locations which are consistent over time. The most common way to proceed is to use allocation factors between PUMA (or CBSA) codes in 2008 and metro areas in 1980. This step requires some manual correction when the county composition of each metro area has changed between 1980 and 2008. For this purpose, population data at the county level is useful in order to check the consistency of geographical composition. Once this consolidation of spatial boundaries is done, it is possible to merge individual data with population data coming from the 1960, 1980 and 2008 National Censuses. We obtain a subset of 218 metro areas, representing 63% of the 1980 U.S. population and 71% of the 2008 U.S. population.[43]

To construct information about workers of different city size, we first rank these 218 metro areas according to their population in 1980 and then, using this ranking, we define large and small cities by splitting the sample into three groups of cities with equal total population each. We then consider large cities as those in the top 33% and small cities those in the bottom 33% of the 1980 population distribution. We also adopt different definitions of "large" and "small" by splitting the sample in two equal groups and defining large (small) cities as those belonging to the top (bottom) 50% of the U.S. population distribution in 1980 and splitting the sample in four equal groups and defining large (small) cities as those belonging to the top (bottom) 25% of the U.S. population distribution in 1980.

Providing a sense of how population is concentrated across cities in 1980, consider that 50% of the total population concentrates in the 22 largest cities (Phoenix with around 1.5 millions inhabitants being the marginal city), 33% of the total population concentrates within the largest 10 cities (Miami-Hileah with around 2.6 millions inhabitants being the first marginal city) and within the 174 smallest cities (Toledo with around 791,000 inhabitants being the second marginal city), 25% of the total population concentrates in the largest 6 cities (San Francisco-Oakland-Vallejo with around 3.2 millions being the first marginal city) and within the smallest 155 cities (Fresno with around 515,000 inhabitants being the second marginal city).

[43] Diamond (2016) uses the same number of MSA for the period 1980-2000.

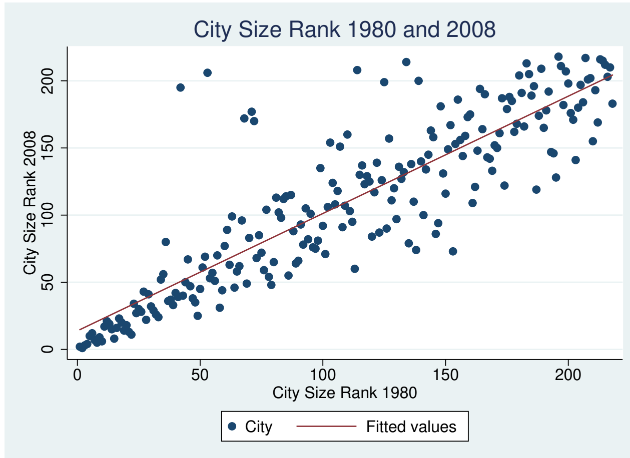

Figure 11: Cities population ranks in 1980 and 2008

In order to ensure that city groups are of equal size, we re-weight the contribution of marginal cities, such that part of their population is accounted to be in the group of small cities and the other part in the group of large cities. Specifically, the total population of the 218 cities in 1980 is 142,668,593 so that 33% of the population is 47,556,198 (rounded up). For instance, the 50th decile of the U.S. population divides the city of Phoenix in two: 70.7% of its population is part of the top 50% and 29.3% is part of the bottom 50%. The same approach is applied for the other marginal cities of the other groups. Technically, we double each individual belonging to each marginal city and we split her weight according to the share of the city population that enters in each group.

A possible concern is that our city size classification is not stable over time, so that cities that are labeled small in 1980 grow faster than average, and should then be labeled large in 2008, while our ranking considers them small also in 2008. However, the size distribution of cities is stable over time as documented by Gabaix and Ioannides (2004), so that the majority of cities that we label large in 1980 would be labeled large even if re-assessing the ranking in 2008. We confirm this finding in Figure 11, in which we plot the rank of city size in 1980 against the rank in 2008. The correlation between the two is positive and large (0.87) and the noise is limited (R-squared 0.76). We remark that such evidence is consistent with the performance of the model: cities that are large in the initial period in the model display faster SBTC and TFP growth over time, which induces them to become more polarized and also display faster population growth over time. Finally, we report for completeness the rank and population in 1980 of a selected group of cities (see Table 4).

Table 4: A sample of US Cities ranked by their population in 1980

Rank 1980	**Metro Area**	**Pop. 1980**
1	Los Angeles-Long Beach, CA	9.410.212
2	New York, NY-Northeastern NJ	9.120.346
3	Chicago-Gary-Lake, IL	7.226.761
4	Philadelphia, PA/NJ	4.716.818
5	Detroit, MI	4.353.413
6	San Francisco-Oakland-Vallejo, CA	3.250.630
7	Washington, DC/MD/VA	3.060.922
8	Houston-Brazoria, TX	2.905.353
9	Boston, MA	2.763.357
10	Miami-Hialeah, FL	2.643.981
11	St. Louis, MO/IL	2.356.460
12	Pittsburgh-Beaver Valley, PA	2.263.894
13	Baltimore, MD	2.174.023
14	Minneapolis-St. Paul, MN	2.113.533
15	Atlanta, GA	2.029.710
16	San Diego, CA	1.861.846
17	Cincinnati, OH/KY/IN	1.660.278
18	Denver-Boulder-Longmont, CO	1.620.902
19	Seattle-Everett, WA	1.607.469
20	Tampa-St. Petersburg-Clearwater, FL	1.569.134
21	Riverside-San Bernardino, CA	1.558.182
22	Phoenix, AZ	1.509.052
	.....	.....
213	Kokomo, IN	103.715
214	Gadsden, AL	103.057
215	Kankakee, IL	102.926
216	St. Joseph, MO	101.868
217	Sheboygan, WI	100.935
218	Columbia, MO	100.376

A.4 Skill distributions

In Section 6 we construct the skill distributions using a price-theoretic measure of skills formally represented by Equation (16), which we report here for convenience

$$U^{ik} = \Omega \left(p_H^k\right)^{-\alpha} \left(w^{ik}\right)^{\alpha+\omega} \left(1 + \left(\frac{\psi}{1-\psi}\right)^{\gamma} \left(\frac{w^{ik}}{A_h p_s^k}\right)^{\gamma-1}\right)^{\frac{1-\omega-\alpha}{\gamma-1}}, \tag{17}$$

and where $\Omega = \alpha^{\alpha}\omega^{\omega}\left(1-\omega-\alpha\right)^{(1-\omega-\alpha)}\left(1-\psi\right)^{\frac{\gamma(1-\omega-\alpha)}{\gamma-1}}\left(A_h\right)^{(1-\omega-\alpha)}$. To quantify this measure using individual wages w^{ik}, we need to provide values for the prices p_H^k and p_s^k.

For the price of housing, following the methodology in Eeckhout et al. (2014), we compute location-specific housing price indices using a hedonic regression model. While housing is a homogeneous good in the model, in the data housing differs in many characteristics that may affect prices. Thus, by relating the log of rent against a number of housing characteristics (number of rooms, age and size of the structure, etc.) and with *city-specific fixed effects*, we isolate the location-specific component of housing prices that can be used to index the difference in housing values across cities. Data on dwelling features comes from the American Community Survey (ACS) and are reported in the IPUMS database at the public use metropolitan area level (PUMA codes) after 2000 and at the metropolitan area level (METAREAD) before 1990. Metro areas are "regions consisting of a large urban core together with surrounding communities that have a high degree of economic and social integration with the urban core".

For the price of non-tradables p_s^k, we rely on the price indexes at the metropolitan area level for the period 1982-2012 provided by Carrillo et al. (2014). Since this paper provides only aggregate prices for goods and services, we use the value of the consumption share of non-tradable from the benchmark calibration ($1-\alpha-\omega=0.35$) to impute the variation of prices across location only to the non-tradable services assuming that for tradable goods the law of one price holds. We stress, however, that the measure of skill distribution obtained is very robust to different value of non-tradable prices.[44]

B Appendix: Robustness

In this appendix we provide some robustness checks concerning both the empirical and the quantitative analysis. In the first subsection we document that spatial and aggregate employment polarization are mainly driven by the extensive margin of employment. Next,

[44] Results are available upon request.

we provide evidence according to which spatial employment polarization is mainly driven by women. Finally, we provide evidence of spatial educational polarization, i.e. the fact that from 1980 and 2008 larger cities became more attractive for both highly educated (at least college) and poorly educated (less than high school) workers. By complementing the analysis of Section 6 with an observable measure of skills, this finding reinforces the view according to which the spatial change in the occupational structure (spatial employment polarization) is associated with the spatial sorting of heterogenously skilled workers (spatial polarization at the worker level).

In the second subsection we present further results on the distribution of the model-based skill measure presented in Section 6. In particular, to facilitate the comparison between the result on the occupational skill distribution and the workers skill distribution, we split the latter in three groups (as for occupations) according to different categorizations and we show that the pattern of spatial polarization at the worker level is robust to these changes.

B.1 Additional evidence

B.1.1 Employment Polarization of Workers

The measures presented in Figure 2 suggests that the employment shares in terms of hours of high- and low-skill occupations increase more in large cities than in small ones. As discussed in Section 3, this implies that changes in employment shares include both the intensive and the extensive margin of employment. Thus, changes of employment shares across occupations can be due to either (a) workers who change their working time in the market while performing the same occupation, (b) workers switching occupations (within cities or across them) or (c) both channels. Measuring to what extent the two channels contribute to spatial employment polarization provides information on the role of the sorting of workers in producing the phenomenon. In this appendix we modify the graph in Figure 2 to consider only the change in the number of workers along the skill distribution, rather than the change in hours. To put it differently, we reconstruct Figure 2 by assuming that there is no change in hours worked between 1980 and 2008 in any of the occupations used to construct Figure 2. Formally, we retain the same percentiles classification as in Figure 2 and measure, for each percentile, the change in the share of *workers* from 1980 to 2008.[45] The results are reported in Figure 12 and show that the U-shape is driven by a change in the number of workers along the skill distribution. The similarity with Figure 2 suggests that the observed spatial employment polarization is driven by a larger increase in the proportion of *individuals*

[45]Thus, 1980 percentiles used to construct Figure 2 can be considered as bins of occupations that are kept constant over time. Using these bins we construct Figure 12.

working in high- and low-skilled occupations in large cities than in small ones.

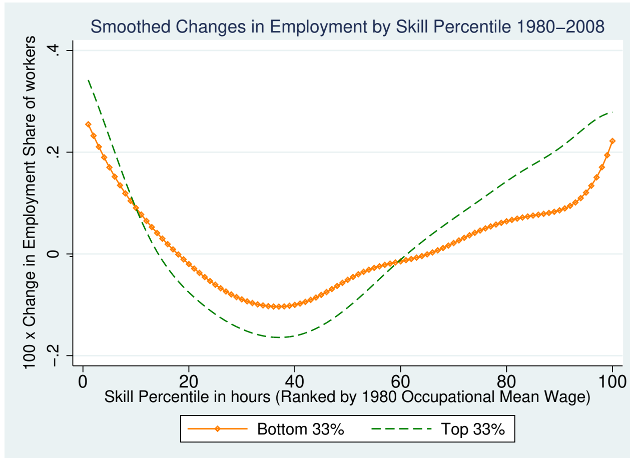

Figure 12: Employment polarization by city size in terms of workers. *Top vs bottom 33* grouping. The ranking of occupations and the bins of occupations are the same as in Figure 2. The variable on the vertical axis is the change in the share of workers in each bin.

B.1.2 The role of gender in spatial polarization

Cerina et al. (2021b) document that a main driver of employment polarization in the U.S. is the reallocation of hours from home production to market work experienced by women since 1980s. They show how the sharp increase in the education premium in the 1980s increased women's participation, directly, at the top and, indirectly, at the bottom of the occupational skill distribution, due to a larger demand for low-skilled services by skilled women. By its nature, this mechanism should emerge at the level of metropolitan areas, because low-skilled services are produced and consumed locally. Also, it should be more evident in large cities, where SBTC was stronger - Baum-Snow et al. (2018)- so that the results in Cerina et al. (2021b) suggest that a large fraction of the spatial differences in employment polarization should be driven by women. In this section we investigate to what extent this is the case.

Figure 13 decomposes the overall spatial difference in the change of employment shares in our three main occupational categories (bar in gray) between men (bar in white) and women (bar in black). The first panel presents the spatial differences in employment polarization between cities belonging to the top vs bottom 50% grouping in 1980, while the second and the third present the same differential pattern for cities in the top vs bottom 33% grouping and in the top vs bottom 25% grouping, respectively. The graphs reveal that the differential pattern in employment polarization between large and small cities is mainly driven by women, especially at the bottom and in the middle of the occupational skill distribution. When comparing the top vs bottom 50% (first panel), women display 230%, 166% and 49% of the difference in the change of employment shares for low-, middle- and high-skilled occupations, respectively. When comparing cities belonging to the top and bottom 33%, the corresponding figures are 136%, 123% and 42% (second panel), while for cities in the top and bottom 25% they are 111%, 101% and 43% (third panel). Thus, irrespective of the definition of small and large cities, women are responsible for the majority of the change at the bottom and in the middle of the skill distribution and for slightly less than half of the difference at the top. In low- and middle-skilled occupations men display a pattern which is either similar across city size or opposite with respect to the aggregate one. In particular, the increase in employment shares in low-skilled occupations for men is always higher in small than in large cities.

This section shows that the contribution of women to employment polarization is particularly strong in large cities, in which there has been more employment polarization and where SBTC has been stronger (Baum-Snow et al., 2018).

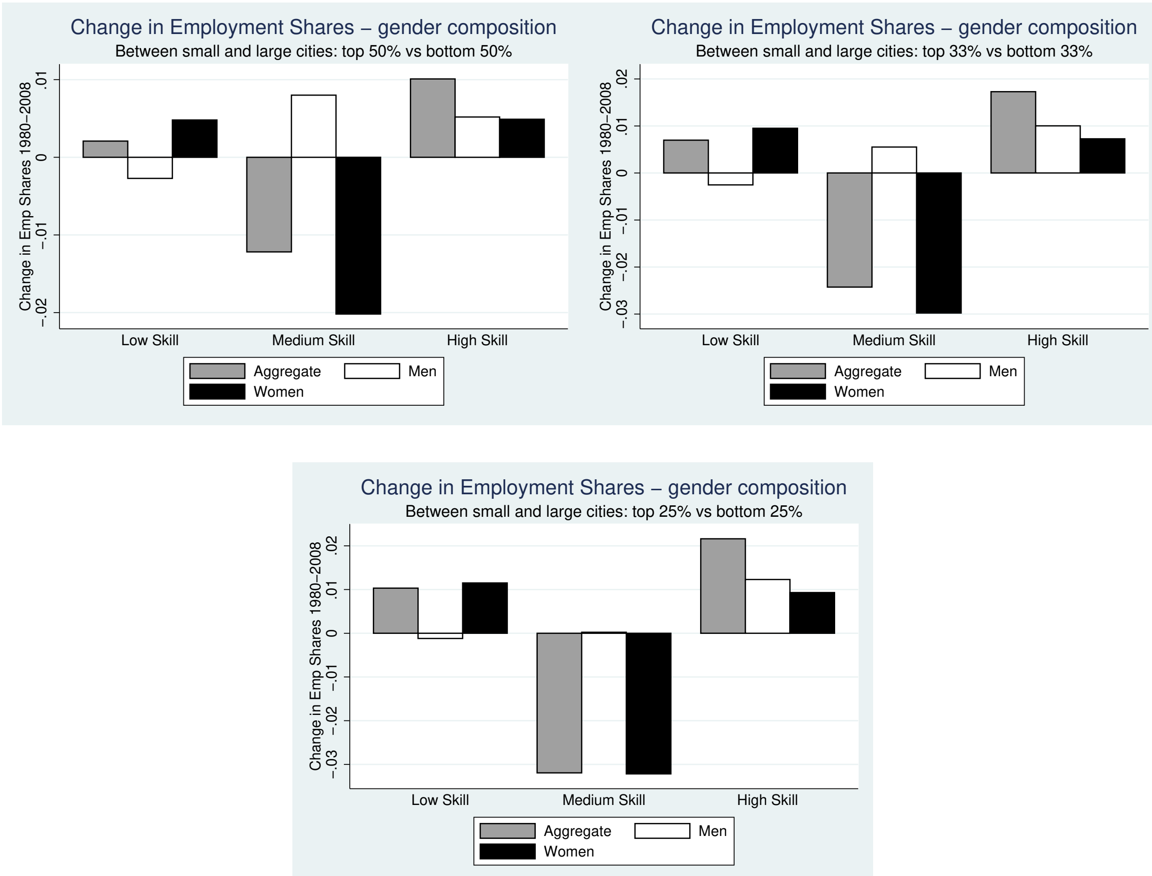

Figure 13: Difference in the change in employment shares between large and small cities in low-, middle- and high-skilled occupations across gender. The left panel compares metropolitan areas belonging to the top vs bottom 50% grouping in 1980, the right panel compares metropolitan areas belonging to the top vs bottom 33% in 1980 and the bottom panel compares metropolitan areas belonging to the top vs bottom 25% in 1980.

B.1.3 Spatial educational polarization

Here we study spatial polarization by using an observable education measure as a proxy of skills. Figure 14 shows how the distribution of educational attainments evolved differently in large and small cities between 1980 and 2008. Based on the sample of workers used to analyze employment polarization, we observe that larger cities display a relative increase in the shares of both low-skilled workers (high school dropouts) and high-skill workers (college degree or more) and a relative decrease in middle-skilled workers (less than college). Specifically, focusing on the top vs bottom 33% grouping, we observe that the share of high-school dropouts is very similar in large and small cities in 1980 (respectively 20.75% and 20.77%) but, while decreasing in both groups of cities, it decreases faster in small (-12.95 percentage points) than in large cities (-10.85 pps), giving rise to a spatial difference in the change between the latter and the former of +2.10 pps. Just as high-school dropouts become over-represented in large versus small cities, so do individuals with college degree or more: their share increase by 16.09 pps in large cities while only by 11.20 pps in small cities, resulting in a spatial difference of 4.89 pps. As a consequence, middle-educated individuals (college dropouts) becomes remarkably under-represented in large (-5.24 pps) as compared to small cities, where their market working hours share increases by 1.75 pps, resulting in a spatial difference of 6.99 pps.

We also observe, by analyzing the bottom-right panel of Figure 14 how the pattern of spatial educational polarization increases with more extreme definitions of large and small cities thereby confirming the pattern observed for both occupations and the model-based measure of individual skill. We conclude that this evidence on observable skill measure complements and reinforces the one presented in the main text on spatial polarization.

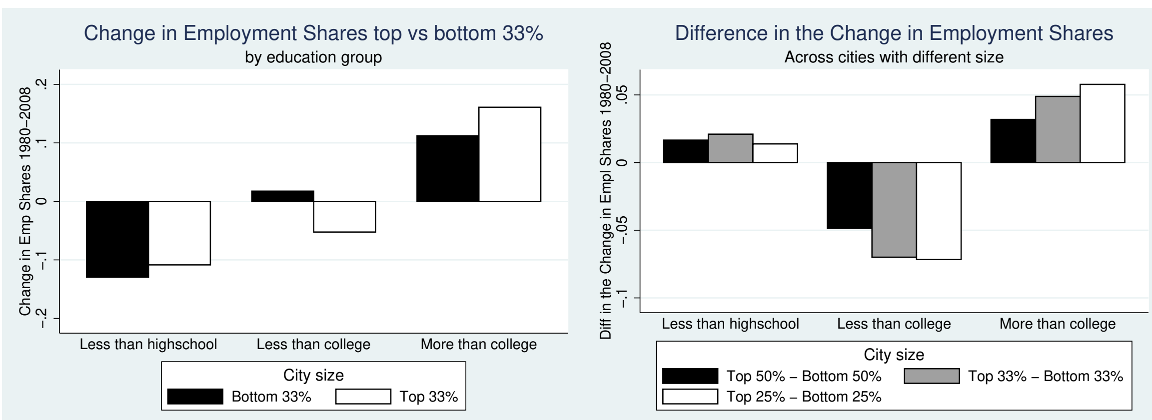

Figure 14: Educational polarization by city size. The left panel compares metropolitan areas belonging to the *top vs bottom 33%* grouping. The right panel reports the difference in the change in educational shares across cities with different size for three groupings: *top vs bottom 50%*, *top vs bottom 33%*, and *top vs bottom 25%*

B.2 Model-based individual skill distribution with discrete bins

In Section 6 we generate continuous individual skill distributions taking advantage of the theory developed in Section 4 and in particular of the mobility assumption as in Eeckhout et al. (2014). Our main aim here is to directly compare changes across space and time of the individual skill distribution with the changes in the occupational skill distribution documented in Section 3, Table 1. This aim requires us to identify three groups of skills. To do this we proceed according to the following steps: 1) we compute each individual's utility resulting from Equation (16) using wages and prices from the data and using the same parametrization of the benchmark calibration (see Table 2); 2) we rank individuals in 1980 and 2008 according to this utility and we compute the resulting cumulative distribution of hours worked for each year in the overall economy; and 3) is that of splitting this cumulative distribution according the two cutoffs of the skill distribution which identify the three skill groups. Every cutoff choice implies a certain degree of arbitrariness. In order to limit the latter, we provide here four different sets of cutoffs (and, therefore, skill groupings). In the first (33 33 33) we split the sample in three equally populated bins, so that the low-skilled group is made up of individuals belonging to the 33% with the lowest indirect utility, the high-skilled group consists of individuals belonging the top 33% with the highest indirect utility and the middle-skilled group is the remaining one. In the second (25 75 25) and in the third (20 60 20) we split the sample at, respectively, the 25th and the 75th percentiles and at the 20th and the 80th percentiles. So that in the second (third) grouping, the low-skilled group is made of individuals belonging to the 25% (20%) with the lowest indirect utility, the high-skilled group consists of individuals belonging the 25% (20%) with the highest indirect utility and the middle-skilled group is the remaining one. Finally, in the last grouping (Occ) we spilt the sample in order for the share of cumulative hours worked by each group to match the corresponding share of hours worked in each occupational group in 1980 and 2008 as in Table 1. More precisely: low-skilled workers in 1980 (2008) are defined as those who fill the first 11.61% (14.73%) of the cumulative distribution of hours; high-skilled workers in 1980 (2008) are defined as those who fill the last 25.68% (34.22%) of the cumulative distribution of hours; middle-skilled workers are defined as all the remaining workers in both year. In this latter classification, unlike the former, the size of bins are not time invariant but by matching occupational shares it allows for a more direct quantitative comparison with the latter.

Finally, once we have defined the three skill groups at the aggregate level, the last step is to perform the same analysis performed in Section 3, for the occupational groups and, therefore, study how their share of market working hours are allocated between large and small cities in 1980 and 2008. We use here the same definition of large and small cities used in Section 3, which are, respectively, those belonging to the top and bottom 33% of the population distribution in 1980. Results are presented in Figure 15 where we report the spatial difference in the change overtime in market working hours for the above defined four skill groupings.

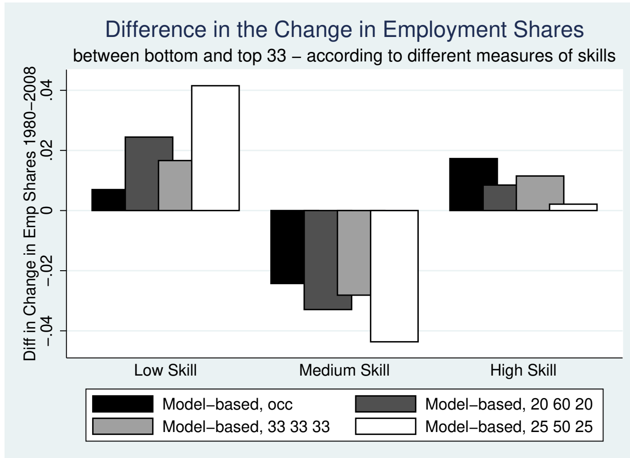

Figure 15: Employment polarization at the worker level and city size. The graph compares the difference in the change in market working hours shares of high-, middle- and low-skilled workers between 1980 and 2008 between top and bottom 33% cities according to four alternative measures: black bars report the values of the three skill groups obtained by splitting the model-based skill distribution in order to match occupational shares in 1980 and 2008 (Occ): dark-gray bars report the values of three skill groups obtained by splitting model-based skill distribution at the 20th and 80th percentiles in both 1980 and 2008; light-grey bars report the values of three skill groups obtained by splitting model-based skill distribution at the 33.33th and 66.66th percentiles in both 1980 and 2008; white bars report the values of three skill groups obtained by splitting model-based skill distribution at the 25th and 75th percentiles in both 1980 and 2008.

These results confirm that spatial employment polarization occurs both at the occupational and at the worker level and point to a primary role of spatial sorting of heterogenously skilled individuals across cities of different size in explaining the stronger pattern of spatial polarization in large cities.