Female Employment and Job Polarization: The Case of Germany*

Lara Vivian†

June 4, 2018

Abstract

Job polarization refers to the disproportionate increase in employment in top and bottom-paying occupations, at the expense of middling-paying jobs that most Western economies have witnessed since 1990. At the same time, a large increase in female labour market participation took place in the same economies, raising the question of whether female employment may explain part of the observed polarization. This paper uses German data to investigate the role of high-skilled females in the development of a market for home production substitutes. To provide causal evidence on the demand for services, I exploit exogenous variation in top paying-occupations' labour market incentives, captured by the dispersion of hourly wages at the occupational level. I find that when top-employed females work more hours, low-skilled women are more likely to be employed.

JEL Classification: *J22, E24*

Key words: Job Polarization, female employment, occupational wage inequality, working hours.

*Ackowledgements: I am grateful to Cecilia García-Peñalosa and to audience participants at AMSE and at the 2018 AFSE annual meeting.

†Aix-Marseille Univ., CNRS, EHESS, Centrale Marseille, AMSE

1 Introduction

During the period 1990-2010, most Western economies have experienced increasing job polarization: the relative employment share of managers, professionals and low-skilled service jobs has increased over time, at the expense of middling-paying occupations, such as clerks or assemblers (Goos et al., 2009). Following the paper by Autor et al. (2003), the literature has shown technology to complement (substitute) workers in top (middling)-paying occupations, and therefore to be a major driver of job polarization. However, more recent papers highlight the contribution of female labour force participation to the polarization of employment via home production outsourcing (Cerina et al., 2017). The development of a market for home production substitutes allows educated females to shift hours from housework to the workplace, while fostering female employment in low-skilled service occupations. Nevertheless, the empirical evidence on the relevance of home production outsourcing for job polarization is mixed (David and Dorn, 2013, Cerina et al., 2017). This paper uses German data to explore the role of educated females in the development of a market for home production substitutes and how it relates to the polarization of employment.

The main motivation behind my work is the different employment distribution of males and females. Figure 1 plots the employment shares of both genders in top, middling and bottom-paying occupations for the years 1990 and 2012 in Germany. While certain features are common to both genders, such as the reduction of the employment share of middling-paying occupations, or the increase of the relative importance of workers at the top, the overall distribution of female employment in 2012 appears more "U-shaped" than the male counterpart, suggesting that female employment is more polarized. Although Mazzolari and Ragusa (2013) provide evidence on high-skill consumption spillovers to low-skill employment, the implications of outsourcing home production for female employment are less clear. David and Dorn (2013) find no evidence in support of a link between female employment at the top and the bottom of the wage distribution in the U.S., while Cortes and Tessada (2011) find that low-skilled migration to the U.S., used as a proxy for cheaper home services' substitutes, increases time spent in the workplace by high-skilled females. Recently, Cerina et al. (2017) have argued that the increase in employment in service occupations is partly explained by females at the top of the wage distribution working more hours in the U.S.

My paper follows the theoretical mechanism proposed by Cerina et al. (2017) and uses German individual data to provide evidence on the causal effect of an increase in the hours worked by top-employed women and the subsequent employment change of low-skilled female workers. To this end, my framework addresses two problems that make investigating whether the demand for services creates its own supply not obvious. The first concerns the spatial dimension of the labour market; personal services that substitute home production require proximity, meaning that low-skilled females are more likely to provide services in areas where top-employed females work more. I deal with the proximity argument by aggregating at the state level hours spent in the workplace by top-paid females, therefore exploiting cross-state and time variation in the demand for services.

Endogeneity in the econometric specification represents the second challenge to the empirical strategy. Longer hours worked at the top could be caused by increased employment in personal service occupations, implying reverse causality and, in such a case, the estimates of the econometric model are biased. A solution is to use an instrumental variable approach, as in Cortes and Tessada (2011). To provide causal evidence on the demand for services and how it relates to the employment of low-skilled females, I exploit exogenous variation in top paying-occupation's labour market incentives, captured by the dispersion of hourly wages at the occupational level. Choosing occupational inequality as an instrument is likely to solve the endogeneity issue because it fills both the exogenety and the relevance requirements that are necessary in order to have consistent estimates. Since the instrument is computed over top-paying occupations, it is unlikely to be correlated with the participation decision of low-skilled females that are unqualified for such jobs, while I show that it is correlated with the intensive margin changes of top-employed females.

My underlying hypothesis is that the changes in female employment in top positions is partly explained by higher wage incentives associated to top occupations. My framework builds on Bell and Freeman (2001) and tests whether the possibility of reaching higher wages, captured by wage dispersion at the occupational level, acts as an incentive for workers to supply more hours. While their findings show that in the U.S. more hours are supplied when inequality is higher, they find no significant effect for Germany. However, they focus on the period 1985-1995, which is the decade prior to the well-documented increase in wage inequality in Germany.[1] By considering a longer time period for Germany, I show that occupational wage inequality plays a role in explaining the changes in the hours worked by top-employed females, while I find that wage dispersion does not motivate other categories of workers to supply more hours. This difference is likely to be due to the importance of the wage level for women who consider outsourcing

[1] See, for instance, Card et al. (2013).

home-production.

The contribution of this paper to the literature on the drivers of job polarization is two-fold. First, I provide empirical evidence on the relevance of home production outsourcing for job polarization, finding that when top-employed females work more hours, low-skilled women are more likely to be employed. Secondly, I show that high-skilled females partly increased their supply of hours as a result of changing labour market incentives, therefore providing evidence on a mechanism other than technology in explaining the variation of employment in top-paying occupations. My results suggests that the development of a market for home production substitutes is partly a result of the changes in the incentive-responsiveness of top-paid females. I find that the increase in inequality explains around 15% of the changes in the female hour share at the bottom of the distribution, and 2% of female employment variation at the top.

The paper is organized as follows: Section 2 describes the data and introduces job polarization in Germany, section 3 presents the empirical strategy including how I deal with endogeneity, while section 4 discusses the results, and I conclude in Section 5.

2 Data and Key Magnitudes

2.1 Data

The German Socio-Economic Panel (SOEP) provides individual data collected from the German Institute for Economic Research (DIW) for the period 1990-2012. Germany offers good quality data and an interesting setting to study questions related to gender and the labour market because of its rapid, and relatively recent, changes in both female employment and labour market regulation. For instance, table 1 shows that female workers went up of almost 8pp, going from 29% of the sample, to representing 37%.[2]

I define occupations using the 2-digit International Standard Classification of Occupation (ISCO) variable and follow the ranking of occupations presented by Goos et al. (2009) to identify job polarization. In their paper, Goos et al. (2009) group occupations into three categories according to the classification based on the mean wage by job type in 1993: top, middling and bottom-paid jobs that, respectively, include 8, 9 and 4 occupations, and show that employment increased at the top and at the bottom of the wage distribution, while it decreased in middling-paying occupations. The detailed list of occupations by category is available in column (1) of

[2]The reader can refer to Burda and Hunt (2011) for a detailed explanation of the recent German labour market reforms and their impact on the economic outputs of the country.

table A.1, while column (2) shows the ranking that results from classifying occupations according to the German 1990 mean wage and few differences, highlighted in bold, can be noted. Since employment is defined as total hours worked, in this paper I refer to employment or hours worked indifferently.

My analysis focuses on the main current job of the survey respondent and uses information about both hours worked and earnings. The definition of hours worked has a central role in my analysis since hours are used both as a proxy for labour supply and to compute hourly wages that, in turn, will constitute the base for my measure of labour market incentives. Therefore, the computation of hours require to be as accurate as possible. Two notions of hours are usually employed in the literature: actual and hours agreed upon. The first notion refers to the hours worked during an average week,[3] while hours agreed upon is a variable meant to indicate hours agreed upon by contract. The latter definition is truncated at 80 because values higher than that are considered implausible. While the notion of hours worked on average is available for the entire working sample, hours agreed upon is asked solely to these employed with contracts that specify a set weekly amount of hours.

Bell and Freeman (2001) use the notion of "actual hours", instead of hours agreed upon. They are concerned with understating the hourly wage since more than 30% of Germans report using compensatory time, meaning that overtime in one week equals less hours worked during the following weeks. The second issue with using actual hours is that, in certain jobs, overtime is likely not to be remunerated. This may be due to several reasons, such as a decision to signal hard work in order to increase future earnings. In that case, hours would increase before a possible consequent wage increase, such as a promotion, therefore also pushing the bias of the computation of current hourly wages downward. I deal with this problem as follows: I consider hours to be equal to actual hours worked for all the workers that report to be remunerated for overtime, while I assign hours agreed upon to the rest of the workers as a measure of labour supply. Figure A.1 reports the distribution of the three possible definitions of hours: my definition of hours, actual, and agreed upon, divided by gender. The definition of hours that I obtain with my procedure is indeed much closer to the definition of agreed hours, confirming that considering actual hours would lead to an underestimation of hourly wages. Furthermore the miscomputation would be more important for males on average, therefore generating a gender bias in the computation of hourly wages.

Earnings refer to the gross monthly wage in the main current job divided by 4.33 to obtain a

[3]The data are precisely obtained by asking respondents how many hours they work on average per week.

weekly estimation. Hourly wages are computed by dividing weekly earnings by hours worked, and wages below the 1st and above the 99th percentile of the hourly wage distribution calculated for each year are considered implausible and, therefore, discarded. I capture labour market incentive following Bell and Freeman (2001) notion of occupational inequality, therefore computing, for each year and each occupation, the standard deviation of the natural log of hourly wages.[4] I reckon this measure to be informative of the possible wage attainable for workers that have comparable characteristics, such as education or experience, and I argue that it has the potential of motivating ambitious workers to supply more hours. The identification strategy will be detailed in the next section.

My sample is restricted to prime-aged workers (25-54 years old) that are likely to be more attached to the labour market than younger or older workers, meaning that they might be more responsive to the possibility of career advancement or promotions. The self-employed are excluded due to the non-trivial issue of computing a representative hourly wage for non-standard workers. I include both West and East Germany in my sample, that goes from 1990, the first date for which observations from both regions are available, to 2012. Including both regions makes sure that I do not capture an increase in job polarization that is solely the result of occupational migration within Germany. For instance, Brücker and Trübswetter (2007) show that a positive selection bias of East-West migrants is in place and that it mainly concerns skilled workers, meaning that internal migration is not equally spread over occupations.

Although the West- and East-German labour markets present different features at the beginning of the period of this study, wages and inequality between the two regions have converged.[5] Fig 2 plots hourly wage inequality, measured by the standard deviation of the log of hourly wages, for Germany and separately for each region, and shows that wage inequality converged after 1995, consistently with the results of Gernandt and Pfeiffer (2008). As a result of the pre-unification regional wage gap, Fig 2 shows that wage inequality reaches a peak in 1990 and comes to comparable levels of the western region in the following 5 years, as soon as the regional wage gap narrows down.[6] Since this paper makes use of the same measure of wage dispersion, but taken at the occupational level, such a behaviour of inequality needs to be accounted for; I do so by including time$\times$region trends in all my econometric specifications. The results are robust, but elasticities are lower, in the case of a sample restricted to West Germany.[7] I present

[4] I have also considered the Gini index in my analysis and the results do not show sensitivity to the definition of labour market incentives.

[5] See for instance a discussion on wages in Burda (2000), and on inequality in Gernandt and Pfeiffer (2008).

[6] Refer to Gernandt and Pfeiffer (2008) and Burda (2000) for a detailed discussion on this.

[7] Results obtained considering only residents in West-Germany are presented in table A.4.

descriptive statistics of my sample for the year 1995 and 2012 in table 1.

2.2 Job Polarization

Fig 3 shows the percentage point changes in the employment shares that occurred in Germany over the period 1990-2012 by occupation and gender. Consistently with the results of the previous literature, figure 3 shows that employment in top and bottom-paying occupations has increased over time at the expenses of middling-remunerated jobs. Relative employment in top-paying occupations went from 35% to 48% and female workers account for an important share of it. They went from representing 14% of employment in 1990 to 22% in 2012, while males increased their relative hour share by 5 percentage points, going from 21% to 26% of total employment.

Employment changes in bottom paying occupations are less sizeable, but the hours worked in personal service related jobs, such as housekeepers, personal care and related workers, went up by almost 2 percentage points. Female employment changes explains most of the variation in the least-paid occupations since female employment at the bottom went from 8.5% to 10%, while male hours used to represent 6% of total employment in 1990 and were at a very comparable 6.7% in 2012.

One may ask at this point which part of the changes in employment is due to the variation in the relative share of a group in the population and how much of the total hour change is instead attributable to the changes in employment by group. One way to answer the question is to equate the changes in employment over the period t-$t+1$ to the sum of the hour changes of each group j weighted by their population shares. In order to decompose total employment changes into a term that captures the variation due to the difference in the hours worked by the group, and a second component that indicates how much of the changes in polarization are attributable to population changes, I perform the following decomposition over the years 1990 and 2012:

$$\begin{aligned} \frac{\Delta H}{H_t} &= \frac{\sum_{j=1}^{J} p_{j,t+1} H_{j,t+1} - \sum_{j=1}^{J} p_{j,t} H_{j,t}}{H_t} \\ &= \underbrace{\frac{\sum_{j=1}^{J} p_{j,t} \Delta H_j}{H_t}}_{\Delta Hours} + \underbrace{\frac{\sum_{j=1}^{J} H_{j,t+1} \Delta P_j}{H_t}}_{\Delta Population}, \end{aligned} \tag{1}$$

where H refers to employment, that is the weighted sum of hours worked, and P indicates the population weight associated to each group. I decompose the total hour change over six groups that result as a combination of gender and the three levels of occupations based on their average wage: top, middling and bottom-paying jobs. The term "$\Delta Hours$" refers to the changes coming from the variation in the total of hours worked by each group, while the second term "$\Delta Population$" denotes the changes in the group weights once hours are kept constant at their $t+1$ level.

I report the results of the decomposition for Germany over the period 1990-2012 in table 2. Overall employment (i.e. total hours) decreased by 22%. Consistently with the previous literature, I find that most of the changes are attributable to middling-paid males, but, also, that the changes in the hours worked by top-paid females and the variation in the population share of bottom-employed women account, respectively, for 25 and 4% of the overall employment changes.[8] This exercise suggests that the changes in the hours worked on average by top-employment females are relevant in explaining the changes in employment at the top, while the variations in the hours worked of bottom-paid females are entirely due to the increase in the importance of their relative weight in the population.

3 Empirical Strategy

This section presents the empirical specification that focuses on the causal effect of an increased demand for services, fostered by top-paid females, on the employment of low-skilled women. I discuss how I construct the variable that I use as a proxy for the local demand for services in section 3.1, while I describe the use of labour market incentives as an instrument in section 3.2. Lastly, I present an individual level regression to test the relevance of wage incentives for top employment in section 3.3.

3.1 The Market for Home Production

The core of my analysis is the "Market for Home Production" equation and it captures the elasticity of the employment decision of low-skilled females with respect to the local labour supply of top-employed women. I consider low-skilled all the prime-aged individuals who have at most completed high-school, and whenever they are in paid positions the dummy associated

[8]See Goos et al. (2009) on the relationship between skilled-bias technology and the loss of employment in middling-paying jobs.

to their employment status equals one. The "Market for Home Production" equation is then

$$Prob(Emp = 1)_{i(s,r),t} = \alpha X_{i(s,r),t} + \beta ln(H)_{s,t} + \gamma_{r\times t} + \delta_s + u_{i(s,r),t} \tag{2}$$

This specification regresses the probability of being employed of low-skilled female i belonging to state s, region r, at time t, on different controls, such as age, age squared, education, and marital status×number of children, captured by the vector X, and on the average of the log of hours worked by top-employed females in state s and year t. The equation is estimated using a linear probability model.[9] Equation 2 also includes region×year dummies, $\gamma_{r\times t}$, in order to account for possible differences in the converge patterns of female labour market participation by region, while state dummies, δ_s, are a proxy for local unobservables that may impact female labour market participation, such as religion.

My regressor of interest is $ln(H)_{s,t}$ and it corresponds to the log of the average of hours worked by top-employed females in state s and year t. Since the services that are demanded in order to outsource home production require physical proximity, I consider the variation in the labour supply of top-paid females at the state level and for each year, the narrower dimension allowed from the data. In order to construct a proxy for the local demand for services fostered by top-employed females, I follow three steps:

1. I follow the strategy proposed by Bell and Freeman (2001), but instead of computing an average of the log of hours worked, as the authors do, I take the average of the hours worked by females in each top-occupation, state and year using individual survey weights w for each observation to ease the interpretation of the results, that is $\bar{H}_{o,s,t} = \dfrac{w_{i(o,s),t} \times H_{i(o,s),t}}{sum(w_{i(o,s),t})} \quad \forall o, s, t.$
2. I then calculate the relative importance of females in each top-occupation, for every state and year, $w^f_{o,s,t} = \dfrac{w_{i\in f(o,s),t} \times H_{i\in f(o,s),t}}{sum(w_{i(o,s),t} \times H_{i(o,s),t})} \quad \forall o, s, t.$
3. Lastly, I use these weights to compute the log of the weighted average of top-employed female hours that varies by state and year $ln(\bar{H})_{s,t} = ln(\dfrac{w^f_{o(s),t} \times \bar{H}_{o(s),t}}{sum(w^f_{o(s),t})}) \quad \forall s, t.$

I consider the hours worked on average by females in each top occupations jointly with the relative representation of females in top jobs so to assign more weight to the hours worked on average in occupations where females represent a larger share. Since my variable of interest only

[9] As a robustness check, I provide estimates of the model obtained following the alternative "Control Function Approach" proposed in Blundell and Powell (2004), and results, presented in table A.3, stay unchanged.

varies by state×year, I two-way cluster the standard errors. Although the choice of considering few clusters (well below 40) in the first level of clustering is likely to increase the risk of high acceptance of the null, therefore possibly reducing the significance of the coefficients, my specification still delivers the expected results.[10]

3.2 Dealing with endogeneity

The previous literature provides little to no evidence on the causal effect of an increase in the hours worked by top-employed females and the consequent positive change in employment for low-skilled women.[11] Instead, Cortes and Tessada (2011) show that high-skilled females are likely to supply more hours of work in locations where migrants represent a larger share of the unskilled workers, i.e. where the supply of services is larger. This means that if longer hours worked at the top are caused by higher employment in personal service occupations, the regression is subject to reverse causality and, in such a case, the estimates are biased. One solution to this problem is to use an instrumental variable approach.

In order to deal with endogeneity in the "Market for Home Production" equation, I build on the "Hour-Inequality hypothesis" of Bell and Freeman (2001). Their argument is that individuals work more hours whenever they have the possibility of attaining higher wages. Wage dispersion is captured by taking the standard deviation of the log of hourly wages (W) for each occupation and each year, and is used as a proxy for labour market incentives.[12] Building on Bell and Freeman (2001), I use an aggregate measure of wage inequality for top-occupations as an instrument for hours worked on average by top-employed females in state s and year t. The construction of the instrument relies on two implicit assumptions. The first one is that females do not include gender differences in their possible wage computation, meaning that the measurement of wage incentives is computed over both males and females. That is, I rule out that women only look at the wages perceived by other females whenever they make predictions about their future wages. The second assumption is that workers compare their wages to those of other individuals in the same position and independently of the location, meaning, for instance, that managers are aware of the wage differential within the occupation and are in the position of acting accordingly, either by moving or negotiating with their superiors. The implications of the two assumptions for my identification are the following:

[10] Refer to Cameron and Miller (2015) for a detailed discussion on the implications of considering few clusters for significance of the coefficients of the model.

[11] The only exception is the paper by Cerina et al. (2017) that provide evidence of this mechanism in the US.

[12] Checks have been performed using the Gini index and the results are robust to such changes.

1. Hourly wage inequality is computed for each top occupation and year, not distinguishing between genders or states $Ineq(W)_{o,t} = \sqrt{Var(ln(W)_{i(o),t})} \quad \forall o, t.$

2. I compute the relative importance of each occupation in each state and year $w_{o,s,t} = w_{i(o,s),t} \times H_{i(o,s),t} \quad \forall o, s, t.$

3. Finally, I calculate a weighted mean by state and year using the above computed weights[13] $Ineq_{s,t} = \dfrac{w_{o(s),t} \times Ineq_{o(t)}}{sum(w_{o(s),t})} \quad \forall s, t.$

I argue that my set of instruments fulfils both the relevance and the exogeneity requirements, since it explains part of the increase in hours worked by top-employed females and concerns only occupations that require a higher level of education than the one attained by low-skilled females. Although this last claim is reasonable, table 3 shows that there is not a one-to-one match between education and occupation, meaning that low-skilled females are represented in top-occupations and therefore my set of instruments may not be exogenous. The Sargan test allows me to check the validity of the instrument set in the case of a model that includes more instruments than endogenous regressors. In order to conduct the test, I add a second instrument to the analysis, this one being an aggregate measure of hourly wages by occupation that also varies by state and year.[14] The results of the test show that my set is valid and therefore the IV model delivers consistent estimates.[15]

3.3 The Hour-Inequality Hypothesis

The "Hour-Inequality hypothesis" regression provides further evidence on the responsiveness to labour market incentives of high-skilled females at the individual level, therefore showing that the results hold when personal characteristics are accounted for. The "Hour-Inequality hypothesis" regression takes the form

$$ln(H)_{i(o,s,r),t} = \alpha X_{i(o,s,r),t} + \beta Ineq(W)_{o,t} + \gamma_{r \times t} + \delta_s + u_{i(o,s,r),t} \tag{3}$$

where the log of hours of the top-employed female i in occupation o, state s, region r, and year t is regressed on different controls, such as age, age squared, education, marital status$\times$number of children, the log of hourly wages, and on our main regressor of interest: wage inequality taken

[13]Checks where weights are kept fixed at their 1990's values are presented in table A.5 and results show robustness to such changes.

[14]The computation of such a variable follows the aggregation presented in section 3.1.

[15]Results of the Sargan test, as well as those of the F-test, will be presented in section 4.

at the occupational level. Consistently with section 3.2, wage dispersion is captured by taking the standard deviation of the log of hourly wages for each occupation and each year and it is computed over both males and females, and indifferently across states. In order to rule out the concern that higher inequality is a consequence of more females working part-time, I present a robustness check where I compute wage dispersion taking into account only male workers, and find comparable results.[16]

Finally, although I use survey weights in the computation of indices and other magnitudes, I follow the recommendation of Solon et al. (2015) and do not perform weighted least square for the other identifications. According to the results of Solon et al. (2015), whenever the survey weights are computed considering some of the controls that are also included in the model (in my case the overlapping variables are age and region), then weighted least square is less precise than unweighted ordinary least square.

4 Results

4.1 The Hour-Inequality Hypothesis

I focus first on the relevance of labour market incentives, captured by wage dispersion at the occupational level, and their impact on the intensive margin of top-employed females, and secondly, on the responsiveness of the employment decision of low-skilled females to the hours worked on average by top-employed females. Results for equation 3 are presented in table 4, where the econometric specifications take a log-log form to ease the interpretation of the coefficient associated with inequality, the regressor of interest. The corresponding coefficient is therefore an elasticity. Column (1) shows the regression performed over the overall working sample, column (2) shows the results for top-employed females, while in column (3) I report the coefficients associated to top-employed males. Finally, columns (4) and (5) show, respectively, the same regression run over middle and bottom-employed females. All specifications include year$\times$region and state dummies, and the standard errors are clustered at the state$\times$year level.

Table 4 shows that inequality does not motivate workers to supply more hours of work when we consider the full sample, but when I take into account the relative position in the wage distribution I unravel different patterns. Inequality plays no role in explaining the labour supply of middling or bottom-paid workers, while it acts as an incentive for the hours worked

[16] The results of the "Hour-Inequality hypothesis" regression obtained using only male wages in the computation of the wage dispersion are presented in table A.2.

by top-employed. This difference is likely to be explained by the higher opportunity cost of time of top-workers. High-skilled females, in particular, face a trade-off between paid-employment and home-production, and higher inequality makes outsourcing home production relatively more convenient for top-paid females than their less-paid workers. In table A.2, I present the results of the "Hour-Inequality hypothesis" regression obtained using male wages in the computation of the wage dispersion. Although the magnitudes of the coefficients are lower, inequality motivates top-employed females to supply more hours, ruling out the concern that wage dispersion increased as a result of a higher number of females employed with part-time contracts.

Although I am only interested in the hours worked by top-employed females, I also opted for pulling the working population together and studying the effect of occupational inequality on the hours supplied by top-paid females by taking advantage of interaction terms, the results are shown in table 5. In particular, I interact the log of inequality with a gender dummy and a categorical variable, named top, that indicates whether the individual is top, middle or bottom employed. This choice allows me to keep a larger number of occupations than I could when I previously focussed on top-employed females only, therefore providing more precise standard errors once I cluster at the occupation-year level. Although 21 occupation-clusters are still less than the usual recommendation, I want to avoid considering only top-workers since doing so would cut the clusters by more than half.[17]

My regression of choice, and therefore the one that I comment on, is reported in column (3) and it includes region$\times$year, state and occupation dummies. Region$\times$year dummies capture the differences in hour convergence between the two regions, the state dummies accounts for unobservables that may influence hours at the state level, such as religion, while the occupation dummies allow to rule out the option that inequality may be capturing occupation-specific requirement of hours.[18] Controls not shown include age, age squared, education, and marital status$\times$number of children.

In order to conclude on the elasticity of hours worked by top-employed females to inequality in the case of multiple interaction terms, I need to sum up different coefficients. The dummy gender is equal to one whenever the individual is a female, while workers employed in middling-paying occupations are the reference for the categorical variable top. The log of inequality is a continuous variable and it is interacted with a dummy (gender) and another categorical variable (top), therefore the coefficient associated to the elasticity of hours with respect to inequality

[17] Refer to Cameron and Miller (2015) for a detailed discussion on the implications of considering few clusters for significance of the coefficients of the model.

[18] In the split samples occupations may be too homogeneous to show differences in behaviour within occupations.

for top-employed females is equal to the sum of the elasticity associated to inequality, plus the coefficient corresponding to the interaction of inequality with top-workers and the elasticity associated to the interaction of inequality and top-employed females. Such a computation allows to conclude that an increase of 10% in occupational wage inequality corresponds to a change of 1.4% in the hours worked by top-employed females. In order to check whether my choice of the index is driving the results, I also performed the same analysis with a Gini index and found similar results that can be checked in table 6.

The results suggest that top-employed females have a stronger substitution than income effect and therefore that they respond positively to labour market incentives than less-paid same-sex workers. Such a finding can be explained by the higher opportunity cost of leisure of high-skilled females.

4.2 The Market for Home Production

The main interest of the paper is to show that changes in female top-employment create a demand for home production substitutes, that, in turn, generate an increase in low-skilled female-biased employment. In table 7 column (1), I provide evidence of the "Market for Home Production" at the aggregated state-year level, by regressing the log of weighted hours worked by females employed in personal service related occupations over the same measure of employment for females in top-paying occupations. Column (2) regresses female employment in service jobs over hours worked by top-employed males and find no effect, while a positive employment change of 10% of females at the top generates an increase of 4.2% in the hours worked by the females employed in service jobs.

Table 8 shows the results of regression 2 that exploits the individual level participation decision of low-skilled females, therefore allowing to control for personal characteristics that might play a role in the participation decision. The sample is restricted to low-skilled females and looks at whether they are more likely to be employed in states where hours worked by top-employed females are higher on average. Since the variable of interest is a log of an average, the interpretation of the elasticity refers to the mean of hours worked.

My regression of choice is depicted in column (2) of table 8 and it includes controls and both region$\times$year and state dummies. The coefficient associated to the average of hours worked by top-and locally-employed females allows to conclude that an increase in 1% in the mean of the hours worked by top-employed females will generate an increase in employment of low-skilled same-sex workers by 0.12 percentage points.

As I discuss in section 3.2, endogeneity is an issue in this context, therefore I also consider an instrumental variable model. To construct my set of instruments, I build on the results of the previous section and use an aggregate measure of occupational wage mean and inequality computed at the state and year level.

Table 8 shows both the second and the first stage of the IV (2SLS) strategy. Here as well, I comment on the regressions that includes controls and both region×year and state dummies. The first stage is presented in log-log form to make the interpretation straightforward and it suggests that a positive change of 10% in the average occupational wage inequality index in top occupations increases the mean of hours worked by top-employed females by 5.1%. The coefficient associated with the hours worked on average by top-employed females in the second stage is equal to 1.07, implying that an increase in 1% in the mean of hours worked by females at the top increases the probability of being employed for low-skilled females by 1.07 percentage points.

The coefficient of interest in the IV model is higher than the one in the LPM regression. Ex-ante, it is difficult to anticipate the direction of the bias. If, for instance, top-employed females work more in states where low-skilled same-sex participation is rather high, I would expect the LPM estimates to be upward biased. If, instead, states are characterized by occupational segregation, meaning that low and high-skilled females are not likely to find employment in the same state, then the LPM coefficients should be smaller.

In table 9, I conduct the same investigation discussed above, but instead of looking at the probability of being employed for a low skilled female, I focus on the probability of being employed in low-skill service occupations. Although the underling assumption of this specification is that being unemployed and working in a low-skill non-service related occupation are equally substitutable to services, focusing on the probability of being employed in a low-skill service job allows to disentangle the channel of the development of a market for home production substitutes from the other possible drivers of the rise in employment for low-skilled females. The change in the coefficient associated to the log of the hours worked on average by low-skilled females becomes larger and significant when passing from a LPM to the 2SLS specification. Column (4) suggests that an increase in 1% in the mean of hours worked by females at the top increases the probability of being employed in service occupations for low-skilled females by 0.25 percentage points.

The magnitudes of the IV coefficients suggest that the incentive responsiveness of top-employed females is likely to drive part of the hour increase in top occupations and, also, to be responsible for part of the increase in employment of low-skilled females. Results are unchanged in table A.3, where I estimate equation 2 using a probit model and I deal with endogenous regressors following the "Control Function Approach" proposed in Blundell and Powell (2004). In table A.4 I focus on residents of West-Germany, while table A.5 reports the same analysis keeping the weights of occupations by state fixed at their 1990 level so that the effect of inequality can be disentangled from a possible change in the relative importance of occupations. I find that results hold in both cases.[19] To make sure that high-skilled women are fostering a demand for services, I also perform a placebo test where I instrument the hours worked by top-paid males and I find no effect on the probability of being employed for low-skilled females.[20]

4.3 Testing the IV-Model

In this section I provide the results of the tests performed in order to check whether the IV model delivers consistent estimates. The first assumption that I test through an F-test is the instrument relevance. This examination requires checking if my instruments are sufficiently strongly correlated with the endogenous variable, meaning that inequality has explanatory power in the labour supply decisions of top workers. The second test is a Sargan test and I employ it to demonstrate that the instrument set is indeed exogenous, meaning that top-occupational wage inequality does not play any role in the decision participation of low-skilled females. Although low-skilled females are not supposed to be qualified for top-jobs, and therefore their participation decision is likely to be independent from labour market incentives in top-occupations, table 3 shows that low-skilled workers are likely to be employed in top-occupations as well, raising concerns about the exogeneity of the instruments.[21] The Sargan test can be performed in my framework because the number of instruments that I consider is larger than the one of the endogenous variables; my instrument set includes both an aggregated measure of occupational inequality and mean wages.

The results of the tests are reported in table 8 and show that both the relevance and the exogeneity requirements are met. In particular, The F-test confirms the relevance of the set of instruments and we cannot reject the validity of the IV-set according to the Sargan test. I also ask whether the endogenous regressor is exogenous in the first place using an Hausman test,

[19] In the latter identification, I do not include State dummies in the estimation presented in table A.5 since I compute inequality keeping occupational weights fixed at the 1990's level for each state.

[20] Results for the placebo test are presented in table A.6.

[21] In the context of instruments constructed using relative shares of occupations at the local level, Goldsmith-Pinkham et al. (2018) also stress the importance of overidentification tests, like the Sargan test, to conclude on the validity of the instrument set.

the results are also reported in table 8 and suggests that the exogeneity of the regressor can be rejected.

5 Counterfactual

Using my empirical specifications, I can compute counterfactual scenarios for job polarization. In particular, I can investigate how employment would be distributed if inequality stayed at its 1995 level.[22] If wage dispersion taken at the occupational level is kept fixed at its value in 1995, top-employed females work on average 26 minutes less per week than when inequality is equal to its 2012 level. Low-skilled female participation is also lower if the aggregate measure of inequality is equal to its 1995 level. In that case 58% of the low-skilled are in paid employment, while 61% of them participate to the labour market if I consider wage dispersion in 2012. Those predicted values can be combined so to obtain counterfactual employment shares for 2012. Table 10 shows in column (3) the counterfactual difference in employment shares between 1995 and 2012, while column (4) reports the actual changes in employment shares for the same period, and I show their difference in column (5). My interest lies in column (6) that shows the ratio of the difference between the counterfactual and the actual polarization scenario over the actual changes in the employment distribution. According to this ratio, the changes in inequality explain around 15% of the changes in the hour share at the bottom of the distribution, and 2% of female employment variation at the top.

6 Discussion and Conclusion

This paper investigates the role of female employment in the context of job polarization. I provide descriptive evidence on the relevance of the variation in the share of hours worked by females in explaining most of the increase in top and bottom-paying occupations. I also highlight the importance of the changes in hours worked on average by top-employed females in explaining the increase in the relative employment share of top occupations, while the increasing representation of bottom-paid women is mainly due to the change in the participation decision of low-skilled females.

I contribute to the literature on the drivers of job polarization by considering the role of labour market incentives in explaining the changes in female employment. In particular, the paper shows that (i) the intensive margin of high-skilled females responds positively to labour

[22] I use 1995 levels instead of 1990, because of the post-unification inequality spike presented in figure 3.

market incentives, captured by the dispersion of hourly wages at the occupational level, and that (ii) low-skilled females are more likely to be employed in states where educated same-sex workers supply more hours. These results highlight the role of female employment in explaining job polarization and suggest that the development of a market for home production substitutes is partly a result of the changes in the incentive-responsiveness of top-paid females.

Within the multiple questions that are still left unanswered, my future agenda will focus on the relevance of the recent labour market reform in explaining part of the observed changes and the implications of the findings for migration.

Variable	Statistic	1995	2000	2005	2012
Hours	Mean	35.824 (0.155)	35.583 (0.118)	35.092 (0.168)	34.936 (0.162)
	sd	7.674 (0.18)	8.886 (0.128)	9.098 (0.164)	9.099 (0.157)
	Min	2 (0.764)	1.5 (0.649)	1.5 (0.733)	2 (0.331)
	Max	72 (3.972)	80 (3.505)	80 (1.688)	72 (1.604)
	Mean High-Skilled Females	33.298 (0.56)	32.804 (0.446)	33.012 (0.479)	34.016 (0.431)
	Mean Low-Skilled Females	31.347 (0.34)	30.159 (0.247)	29.414 (0.338)	29.455 (0.338)
Wages	Mean	12.823 (0.119)	13.885 (0.089)	15.233 (0.12)	16.495 (0.134)
	sd	5.462 (0.149)	5.849 (0.09)	6.429 (0.108)	7.509 (0.11)
	Min	3.106 (0.024)	2.964 (0.018)	2.382 (0.026)	3.085 (0.028)
	Max	41.097 (0.835)	37.286 (0.015)	40.191 (0.109)	46.101 (0.166)
Ineq	eng. profess.	0.337 (0.015)	0.368 (0.015)	0.369 (0.018)	0.401 (0.031)
	small ent. managers	0.359 (0.044)	0.49 (0.068)	0.496 (0.079)	0.421 (0.046)
	life sc. ass. profess.	0.312 (0.028)	0.285 (0.016)	0.349 (0.035)	0.328 (0.024)
	Metal and trade workers	0.323 (0.015)	0.342 (0.014)	0.331 (0.015)	0.359 (0.02)
	Customer service clerks	0.381 (0.068)	0.335 (0.027)	0.371 (0.031)	0.38 (0.032)
	Personal Services	0.331 (0.019)	0.357 (0.014)	0.392 (0.019)	0.375 (0.018)
	Sales and elementary occ.	0.383 (0.034)	0.386 (0.018)	0.378 (0.023)	0.377 (0.026)
Shares	Working females	28.862 (0.751)	31.704 (0.556)	32.231 (0.743)	36.949 (0.784)
	Working within high-skilled females	65.794 (2.607)	67.377 (1.663)	63.381 (2.314)	72.893 (1.978)
	Working within low-skilled females	55.645 (2.55)	60.676 (1.808)	60.487 (1.733)	69.331 (1.759)
	Low-skilled	77.562 (1.289)	76.248 (0.907)	75.849 (1.24)	72.939 (1.206)
	Working part-time within high-skilled females	19.423 (0.675)	19.176 (0.507)	14.577 (0.684)	14.727 (0.711)

Table 1: Descriptive Statistics

group	gender	ΔH	Δpop	$\Delta\bar{H}$	Δpop %	$\Delta\bar{H}$ %
		-0.22				
top	males		0.03	0.02	14.94	8.77
	females		0.05	0.06	25.06	26.32
middle	males		-0.19	-0.15	-88.53	-69.85
	females		-0.02	-0.03	-9.61	-12.10
bottom	males		0.00	-0.00	0.52	-0.06
	females		0.01	0.00	4.38	0.17

Table 2: Changes in hours by group, Germany 1990-2012

group	high-skilled %	low-skilled %
top	41	59
middle	10	90
bottom	10	90

Table 3: Top, middle and bottom female employment by education, Germany 1990-2012

	Dependent variable:				
	lnH				
	All workers	Top-F	Top-M	Middle-F	Bottom-F
lnW	0.098***	0.107***	−0.002	0.144***	0.123***
	(0.029)	(0.038)	(0.008)	(0.029)	(0.027)
Ineq_o	−0.137***	0.124**	0.059**	−0.241	−0.277
	(0.049)	(0.054)	(0.027)	(0.209)	(0.193)
t×r dummies	✓	✓	✓	✓	✓
s dummies	✓	✓	✓	✓	✓
Observations	127,929	25,711	25,105	17,215	15,709
R^2	0.338	0.295	0.041	0.266	0.233
Adjusted R^2	0.337	0.293	0.038	0.262	0.228
Residual Std. Error	0.291 (df = 127797)	0.320 (df = 25615)	0.136 (df = 25007)	0.355 (df = 17120)	0.480 (df = 15614)

*p<0.1; **p<0.05; ***p<0.01

Note: log-log model, controls include age, age squared, education, marital status×number of children. Standard errors are clustered at the occupation×year level.

Table 4: Hour-Inequality Hypothesis, Split Samples

	Dependent variable:		
	lnH		
	(1)	(2)	(3)
lnW	0.090***	0.091***	0.083***
	(0.025)	(0.025)	(0.024)
Ineq_o	−0.091*	−0.091*	−0.059
	(0.049)	(0.049)	(0.063)
gender=f	−0.355***	−0.354***	−0.336***
	(0.091)	(0.091)	(0.086)
Ineq_o:gender=f	−0.106	−0.106	−0.122
	(0.099)	(0.099)	(0.086)
Ineq_o×top=top	−0.006	−0.007	−0.062**
	(0.014)	(0.013)	(0.025)
Ineq_o×top=bottom	0.013	0.014	0.067***
	(0.029)	(0.029)	(0.021)
gender=f×top=top	0.299***	0.300***	0.241***
	(0.086)	(0.086)	(0.066)
gender=f×top=bottom	−0.192***	−0.193***	−0.224***
	(0.050)	(0.050)	(0.077)
Ineq_o×gender=f×top=top	0.267***	0.268***	0.227***
	(0.095)	(0.095)	(0.060)
Ineq_o×gender=f×top=bottom	−0.021	−0.022	−0.046
	(0.050)	(0.050)	(0.053)
t×r dummies	✓	✓	✓
s dummies		✓	✓
o dummies			✓
Observations	130,527	130,527	130,527
R^2	0.282	0.282	0.309
Adjusted R^2	0.281	0.282	0.308
Residual Std. Error	0.308 (df = 130427)	0.308 (df = 130412)	0.302 (df = 130394)

*p<0.1; **p<0.05; ***p<0.01

Note: log-log model, controls include age, age squared, education, marital status×number of children. The reference for gender is male and the one for top is being middle paid. Standard errors are clustered at the occupation×year level.

Table 5: Hour-Inequality Hypothesis

	Dependent variable:		
	lnH		
	(1)	(2)	(3)
lnW	0.089***	0.090***	0.083***
	(0.023)	(0.023)	(0.024)
Ineq_o	−0.108*	−0.108*	−0.068
	(0.058)	(0.058)	(0.070)
gender=f	−0.495**	−0.494**	−0.520***
	(0.208)	(0.208)	(0.166)
Ineq_o×gender=f	−0.152	−0.151	−0.187*
	(0.139)	(0.139)	(0.103)
Ineq_o×top=top	0.008	0.006	−0.075**
	(0.020)	(0.018)	(0.031)
Ineq_o×top=bottom	0.004	0.006	0.222*
	(0.039)	(0.038)	(0.124)
gender=f×top=top	0.632***	0.636***	0.564***
	(0.202)	(0.202)	(0.148)
gender=f×top=bottom	−0.539*	−0.541*	−0.536*
	(0.311)	(0.312)	(0.317)
Ineq_o×gender=f×top=top	0.364***	0.366***	0.335***
	(0.135)	(0.135)	(0.090)
Ineq_o×gender=f×top=top	−0.239	−0.240	−0.232
	(0.185)	(0.185)	(0.159)
t×r dummies	✓	✓	✓
s dummies		✓	✓
o dummies			✓
Observations	130,527	130,527	130,527
R^2	0.286	0.287	0.310
Adjusted R^2	0.286	0.286	0.309
Residual Std. Error	0.307 (df = 130427)	0.307 (df = 130412)	0.302 (df = 130394)

*p<0.1; **p<0.05; ***p<0.01

Note: log-log model, controls include age, age squared, education, marital status×number of children. The reference for gender is male and the one for top is being middle paid. Standard errors are clustered at the occupation×year level.

Table 6: Hour-Inequality Hypothesis, Gini Index

	Dependent variable:	
	emp females in personal services	
	(1)	(2)
empTOP-females	0.425*	
	(0.236)	
empTOP-males		0.355
		(0.385)
t×r dummies	✓	✓
s dummies	✓	✓
Observations	369	369
R^2	0.779	0.776
Adjusted R^2	0.735	0.731
Residual Std. Error (df = 307)	0.544	0.548

*p<0.1; **p<0.05; ***p<0.01

Note: log-log model. Standard errors are clustered at the state×year level.

Table 7: Market for Home Production Aggregated at the State Level

	Dependent variable:					
	participation				First Stage	
	LPM	LPM	2SLS	2SLS		
	(1)	(2)	(3)	(4)	(5)	(6)
lnH_g	0.142***	0.126**	1.038***	1.077**		
	(0.051)	(0.055)	(0.374)	(0.426)		
lnW_g					−0.057	−0.057
					(0.038)	(0.038)
log(Ineq_o)					0.518***	0.518***
					(0.163)	(0.163)
controls		✓		✓		✓
t×r dummies	✓	✓	✓	✓	✓	✓
s dummies	✓	✓	✓	✓	✓	✓
F-test				11.3***		
Sargan-test				1.2		
Hausman-test				5.1**		
Observations	71,089	71,089	71,089	71,089	71,089	71,089
R^2	0.019	0.123	0.010	0.113	0.657	0.658
Adjusted R^2	0.018	0.122	0.009	0.112	0.657	0.657
Residual Std. Error	0.483 (df = 71027)	0.456 (df = 70985)	0.485 (df = 71027)	0.459 (df = 70985)	0.051 (df = 71026)	0.051 (df = 70984)

*p<0.1; **p<0.05; ***p<0.01

Note: First-stage is a log-log model. Controls include age, age squared, education and marital status×number of children. Standard errors are clustered at the state×year level.

Table 8: Market for Home Production-IV

	Dependent variable:					
	emp in services				First Stage	
	LPM	LPM	2SLS	2SLS		
	(1)	(2)	(3)	(4)	(5)	(6)
lnH_g	0.016	0.015	0.249^{*}	0.257^{*}		
	(0.020)	(0.021)	(0.148)	(0.145)		
lnW_g					−0.057	−0.057
					(0.038)	(0.038)
log(Ineq_o)					0.518^{***}	0.518^{***}
					(0.163)	(0.163)
controls		✓		✓		✓
t×r dummies	✓	✓	✓	✓	✓	✓
s dummies	✓	✓	✓	✓	✓	✓
Observations	71,089	71,089	71,089	71,089	71,089	71,089
R^2	0.005	0.010	0.002	0.008	0.657	0.658
Adjusted R^2	0.004	0.009	0.002	0.007	0.657	0.657
Residual Std. Error	0.256 (df = 71027)	0.255 (df = 70985)	0.256 (df = 71027)	0.256 (df = 70985)	0.051 (df = 71026)	0.051 (df = 70984)

$^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$

Note: First-stage is a log-log model. Controls include age, age squared, education and marital status×number of children. Standard errors are clustered at the state×year level.

Table 9: Market for Home Production-Services

	Gender	Counterfactual changes p.p.	Actual changes p.p.	Difference	$\frac{Difference}{Actual}\%$
Top	Males	5.15	4.99	-0.16	-3.18
	Females	8.11	8.30	0.20	2.40
Middle	Males	-11.85	-12.01	-0.16	1.32
	Females	-3.45	-3.52	-0.07	1.86
Bottom	Males	0.85	0.81	-0.04	-5.35
	Females	1.20	1.43	0.23	15.88

Table 10: Counterfactual scenario for Job Polarization

Figure 1: Employment share by gender and occupation in 1990 and 2012.

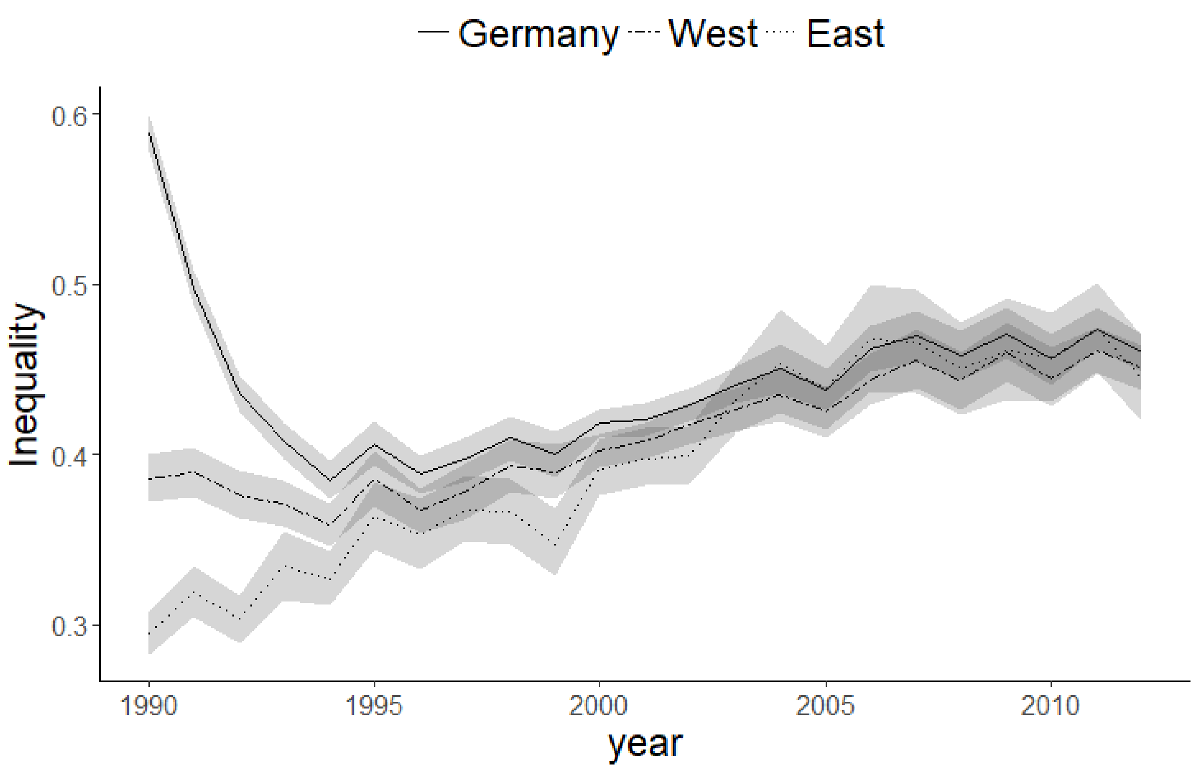

Figure 2: Inequality in Germany and by region, 1990-2012.

Figure 3: Job Polarization by gender, Germany 1990-2012.

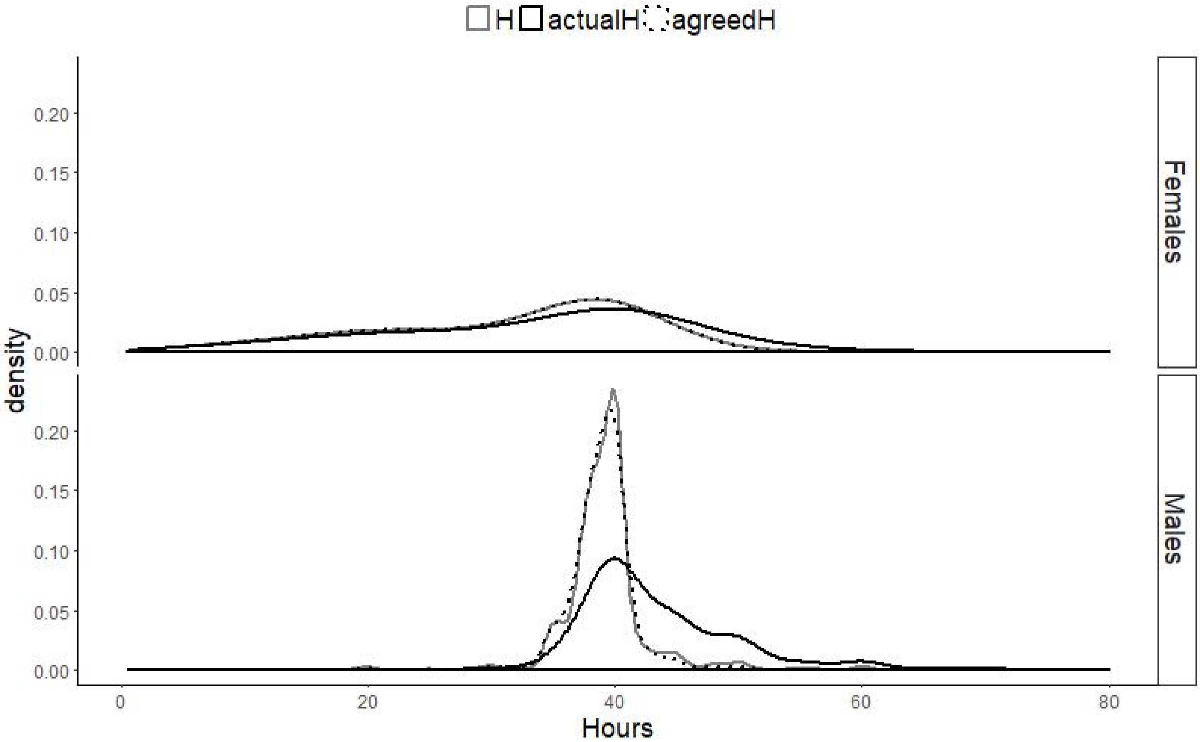

Figure A.1: Different Hour definitions, Germany 1990-2012.

A.1 Data

The following tables and graphs provide further detail on the classification of occupations and the variable on hours worked. In particular, table A.1 defines occupations using the 2-digit International Standard Classification of Occupation (ISCO) and reports the ranking of occupations based on the European average wage in 1993 that first appeared in Goos et al. (2009) and that I also use throughout the paper in column (1). Instead, column (2) provides the classification of occupations that results from computing the average occupational wage in Germany in 1990, the ranking difference concerns few occupations and most of them still fall into the three main groups: top, middling and bottom-paid, those that do not are highlighted in bold.

Figure A.1 plots the distribution of my definition of Hours (H), actual hours ($actualH$), and hours agreed upon ($agreedH$). In order to compute my definition of hours, I assign actual Hours to the workers that declare to have done paid overtime and hours agreed upon to the others. The purpose of this exercise is to show that my definition of hours is much closer to hours agreed upon than it is to actual hours, suggesting that using actual hours could downward bias the hourly wage computation.

Classification	Goos et al. (2009) ref.year=1993	Germany ref.year=1990
8 top-paying occupations		
	corp. managers	small ent. managers
	eng. profess.	eng. profess.
	life science profess.	life science profess.
	other profess.	other profess.
	small ent. managers	corp. managers
	eng. associate profess.	eng. associate profess.
	other associate profess.	other associate profess.
	life sc. ass. profess.	**stat. plant oper.**
9 middling-paying occupations		
	drivers plant oper.	office clerks
	stat. plant oper.	metal and trade workers
	metal and trade workers	**life sc. ass. profess.**
	precision trade workers	precision trade workers
	office clerks	machine operators
	customer service clerks	customer service clerks
	extraction workers	extraction workers
	machine operators	drivers plant oper.
	other craft workers	**constr. transports**
4 bottom-paying occupations		
	personal services	**other craft workers**
	constr. transports	personal services
	models, demonstrators	models, demonstrators
	sales and elementary occ.	sales and elementary occ.

Note: Occupations that present differences in the rankings are reported in bold.

Table A.1: Classification of occupations

A.2 Between and Within Occupation Employment Decomposition

Defining employment as the sum of hours worked in the economy, allows to decompose the hour trend of a certain group in the economy, for instance top-employed females, into a between- and within-occupation components. In other words, the difference in the hour share of a group between 1990 and 2012 can be decomposed into a term reflecting the change in the share of that group occupations, and a term reflecting changes in the group intensities within occupations. I follow Ngai and Petrongolo (2017) and perform a standard shift-share decomposition, therefore the change in the hours share of group j between year t and year $t+1$ can be expressed as

$$\Delta shH_j = shH_j^{t+1} - shH_j^t \approx \underbrace{\sum_{k=1}^{K} \alpha_{j,k} \Delta shH_k}_{\Delta OccPolarization} + \underbrace{\sum_{k=1}^{K} \alpha_k \Delta shH_{j,k}}_{\Delta IntensityWithin} \tag{4}$$

Where k is occupation, and α is the value of the share of hours averaged over the two periods. I perform this decomposition over six groups that result as a combination of gender and the three levels of occupations based on their average wage: top, middling and bottom-paying jobs.

The first term of the decomposition "$\Delta OccPolarization$" refers to the changes coming from the variation in the relative importance of the group occupations in the economy, while the second term "$\Delta IntensityWithin$" denotes changes in the group intensity within occupations.

Figure A.2 plots the results of the decomposition performed over the two main groups of interest of the analysis; top-employed females and same-sex bottom workers, plus the male counterparts. The top two panels depict the results in absolute values, while the bottom ones have been divided by the change in the share of hours for the group so that the two terms resulting from the shift-share analysis sum up to 100. Occupation polarization, or, in other words, the change in the hour share of each group that is attributable to the variation in the relative importance of the occupations in the economy, explains most of the employment share changes for the four groups considered, suggesting that there has been an increasing demand of such jobs in the economy. Hour-intensity within occupations, instead, contributes positively to the increase in total employment share only when considering top-employed females, therefore suggesting that the intensive margin of high-skilled females is likely to play a role in explaining part of the job polarization process.

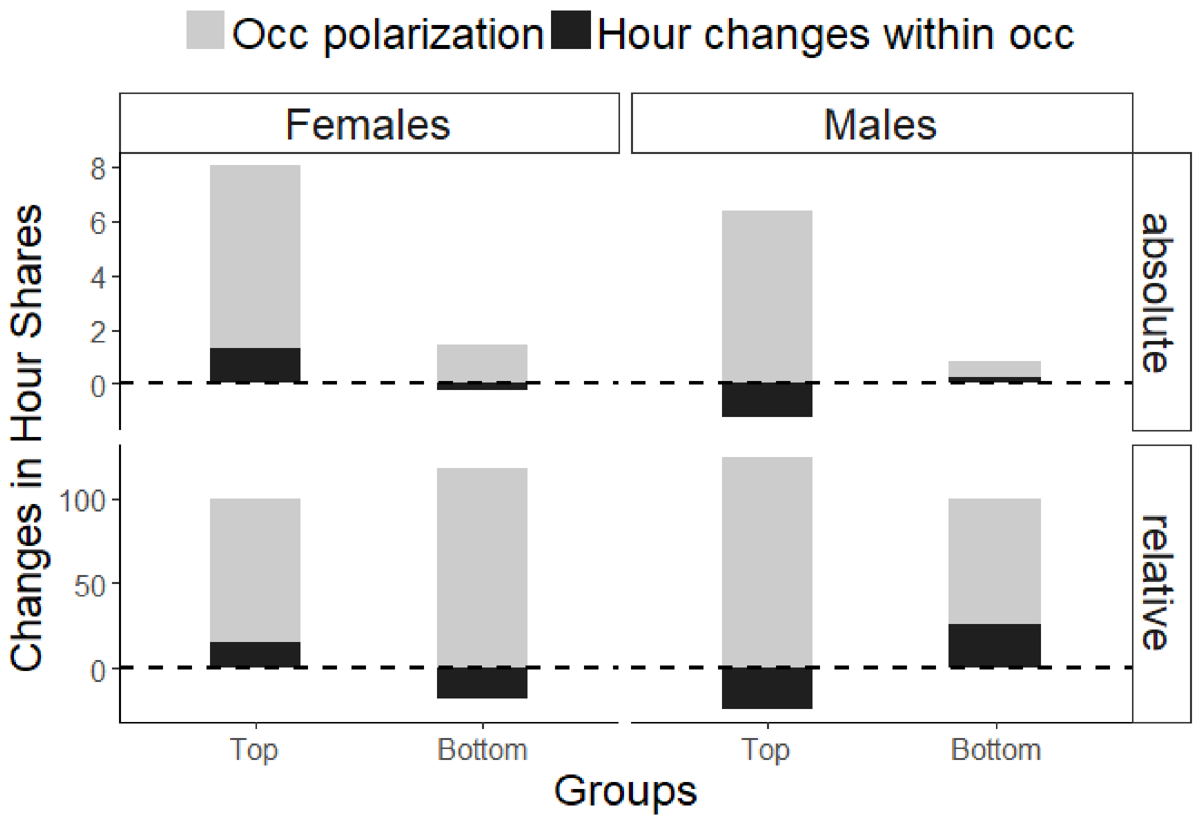

Figure A.2: Changes in hour shares by group, Germany 1990-2012.

A.3 Appendix

	Dependent variable:				
	lnH				
	All workers	Top-F	Top-M	Middle-F	Bottom-F
lnW	0.098***	0.108***	−0.001	0.133***	0.121***
	(0.029)	(0.037)	(0.008)	(0.033)	(0.026)
Ineq_o	−0.058	0.040***	0.020	0.077	−0.278**
	(0.043)	(0.015)	(0.015)	(0.102)	(0.131)
t×r dummies	✓	✓	✓	✓	✓
s dummies	✓	✓	✓	✓	✓
Observations	127,929	25,711	25,105	17,215	15,709
R^2	0.337	0.295	0.041	0.266	0.235
Adjusted R^2	0.337	0.292	0.037	0.262	0.230
Residual Std. Error	0.291 (df = 127797)	0.320 (df = 25615)	0.136 (df = 25007)	0.356 (df = 17120)	0.479 (df = 15614)

*p<0.1; **p<0.05; ***p<0.01

Note: log-log model, controls include age, age squared, education, marital status×number of children. Standard errors are clustered at the occupation×year level. The inequality index is computed excluding females.

Table A.2: Hour-Inequality Hypothesis with Inequality computed excluding females, Split Samples

	Dependent variable:					
	participation				First Stage	
	Probit	Probit	IV Probit	IV Probit		
	(1)	(2)	(3)	(4)	(5)	(6)
lnH_g	0.14**	0.14*	1.04**	1.18*		
	(0.05)	(0.06)	(0.39)	(0.50)		
residuals			−0.92*	−1.06*		
			(0.37)	(0.47)		
lnW_g					−0.06	−0.06
					(0.04)	(0.04)
log(Ineq_o)					0.52**	0.52**
					(0.16)	(0.16)
controls		✓		✓		✓
t×r dummies	✓	✓	✓	✓	✓	✓
s dummies	✓	✓	✓	✓	✓	✓
Num. obs.	71089	71089	71089	71089		
Log Likelihood	-46744.95	-42816.97	-46737.86	-42808.14		
Deviance	93489.89	85633.94	93475.72	85616.28		
AIC	93613.89	85841.94	93601.72	85826.28		
BIC	94182.54	86795.80	94179.54	86789.31		
Num. obs.					71089	71089
R^2 (full model)					0.66	0.66
R^2 (proj model)					0.02	0.02
Adj. R^2 (full model)					0.66	0.66
Adj. R^2 (proj model)					0.02	0.02
F statistic (full model)					2197.85	1311.13
F (full model): p-value					0.00	0.00
F statistic (proj model)					7.20	136.82
F (proj model): p-value					0.00	0.00
				*p<0.1; **p<0.05; ***p<0.01		

Note: First-stage is a log-log model. IV-Probit is computed using the "Control Function Approach" developped in Blundell and Powell (2004). Second Stage is a Probit model where residuals from the first OLS-stage have been added as a covariate to account for endogeneity. Controls include age, age squared, education and marital status×number of children. Standard errors are clustered at the state×year level.

Table A.3: Market Home Prod - IV Probit

	Dependent variable:					
	participation				First Stage	
	LPM	LPM	2SLS	2SLS		
	(1)	(2)	(3)	(4)	(5)	(6)
lnH_g	0.140**	0.123*	0.745**	0.740*		
	(0.058)	(0.063)	(0.347)	(0.390)		
lnW_g					−0.083**	−0.083**
					(0.041)	(0.041)
log(Ineq_o)					0.622***	0.621***
					(0.232)	(0.231)
controls		✓		✓		✓
t×r dummies	✓	✓	✓	✓	✓	✓
s dummies	✓	✓	✓	✓	✓	✓
Observations	56,012	56,012	56,012	56,012	56,012	56,012
R^2	0.017	0.136	0.012	0.131	0.384	0.385
Adjusted R^2	0.016	0.135	0.012	0.130	0.384	0.384
Residual Std. Error	0.484 (df = 55978)	0.454 (df = 55938)	0.485 (df = 55978)	0.455 (df = 55938)	0.054 (df = 55977)	0.054 (df = 55937)

*p<0.1; **p<0.05; ***p<0.01

Note: Only States in the West region have been selected. First-stage is a log-log model. Controls include age, age squared, education and marital status×number of children. Standard errors are clustered at the state×year level.

Table A.4: Market Home Prod - West Germany

	Dependent variable:					
	participation				First Stage	
	LPM	LPM	2SLS	2SLS		
	(1)	(2)	(3)	(4)	(5)	(6)
lnH_g	0.180**	0.143*	0.775**	0.873*		
	(0.082)	(0.079)	(0.391)	(0.476)		
lnW_g					−0.066	−0.066
					(0.056)	(0.056)
log(Ineq_o)					0.683**	0.683**
					(0.312)	(0.310)
controls		✓		✓		✓
t×r dummies	✓	✓	✓	✓	✓	✓
Observations	71,089	71,089	70,195	70,195	70,195	70,195
R^2	0.016	0.121	0.011	0.113	0.540	0.540
Adjusted R^2	0.015	0.120	0.010	0.112	0.539	0.540
Residual Std. Error	0.483 (df = 71042)	0.457 (df = 71000)	0.485 (df = 70148)	0.459 (df = 70106)	0.059 (df = 70147)	0.059 (df = 70105)

*p<0.1; **p<0.05; ***p<0.01

Note: First-stage is a log-log model. Controls include age, age squared, education and marital status×number of children. Standard errors are clustered at the state×year level.

Table A.5: Market for Home Production-IV, occupational weights kept fixed at 1990's level

	Dependent variable:					
	participation				First Stage	
	LPM	LPM	2SLS	2SLS		
	(1)	(2)	(3)	(4)	(5)	(6)
lnH_g	0.131	0.131	−3.278	−3.968		
	(0.120)	(0.133)	(12.990)	(15.290)		
lnW_g					−0.008	−0.008
					(0.030)	(0.030)
log(Ineq_o)					−0.027	−0.027
					(0.137)	(0.136)
controls		✓		✓		✓
t×r dummies	✓	✓	✓	✓	✓	✓
s dummies	✓	✓	✓	✓	✓	✓
Observations	71,089	71,089	71,089	71,089	71,089	71,089
R^2	0.019	0.123	−0.008	0.084	0.412	0.412
Adjusted R^2	0.018	0.122	−0.009	0.083	0.411	0.411
Residual Std. Error	0.483 (df = 71027)	0.456 (df = 70985)	0.489 (df = 71027)	0.466 (df = 70985)	0.024 (df = 71026)	0.024 (df = 70984)

*p<0.1; **p<0.05; ***p<0.01

Note: First-stage is a log-log model. Controls include age, age squared, education and marital status×number of children. Standard errors are clustered at the state×year level.

Table A.6: Market for Home Production-IV-Male hours

References

Autor, D. H., F. Levy, and R. J. Murnane (2003). The skill content of recent technological change: An empirical exploration. *The Quarterly journal of economics 118* (4), 1279–1333.

Bell, L. A. and R. B. Freeman (2001). The incentive for working hard: explaining hours worked differences in the us and germany. *Labour Economics 8* (2), 181–202.

Blundell, R. W. and J. L. Powell (2004). Endogeneity in semiparametric binary response models. *The Review of Economic Studies 71* (3), 655–679.

Brücker, H. and P. Trübswetter (2007). Do the best go west? an analysis of the self-selection of employed east-west migrants in germany. *Empirica 34* (4), 371–395.

Burda, M. (2000). East-west german wage convergence after reunification: Migration or institutions? *Department of Economics Humboldt University Berlin Unpublished Working Paper*.

Burda, M. C. and J. Hunt (2011). What explains the german labor market miracle in the great recession? Technical report, National bureau of economic research.

Cameron, A. C. and D. L. Miller (2015). A practitioners guide to cluster-robust inference. *Journal of Human Resources 50* (2), 317–372.

Card, D., J. Heining, and P. Kline (2013). Workplace heterogeneity and the rise of west german wage inequality. *The Quarterly Journal of Economics 128* (3), 967–1015.

Cerina, F., A. Moro, and M. Rendall (2017). The role of gender in employment polarization. University of Zurich, Department of Economics, Working Paper No. 250.

Cortes, P. and J. Tessada (2011). Low-skilled immigration and the labor supply of highly skilled women. *American Economic Journal: Applied Economics 3* (3), 88–123.

David, H. and D. Dorn (2013). The growth of low-skill service jobs and the polarization of the us labor market. *The American Economic Review 103* (5), 1553–1597.

Gernandt, J. and F. Pfeiffer (2008). Wage convergence and inequality after unification:(east) germany in transition. Mannheim, ZEW.

Goldsmith-Pinkham, P., I. Sorkin, and H. Swift (2018). Bartik instruments: What, when, why, and how. Technical report, National Bureau of Economic Research.

Goos, M., A. Manning, and A. Salomons (2009). Job polarization in europe. *The American Economic Review 99*(2), 58–63.

Mazzolari, F. and G. Ragusa (2013). Spillovers from high-skill consumption to low-skill labor markets. *Review of Economics and Statistics 95*(1), 74–86.

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

Solon, G., S. J. Haider, and J. M. Wooldridge (2015). What are we weighting for? *Journal of Human Resources 50*(2), 301–316.