Women's Labor Force Participation and the Distribution of the Gender Wage Gap*

Sonia R. Bhalotra† Manuel Fernández‡

November 11, 2018

Click here for latest version

Abstract

We analyse how the rising labor force participation of women influences the distribution of the gender pay gap and inequality. We formulate an equilibrium model of the labor market in which the elasticity of substitution between male and female labor varies with the task content of occupations. We structurally estimate the parameters using individual data from Mexico between 1989 and 2014, when women's labor force participation increased by fifty percent. We provide novel evidence that male and female labor are closer substitutes in high-paying abstract task-intensive occupations than in lower-paying manual and routine task-intensive occupations. Consistent with this, we find a widening of the gender pay gap at the lower end of the distribution, alongside a narrowing towards the top. Demand side trends favoured women, attenuating the supply-driven negative pressure on their wages, and more so among college-educated workers. The paper contributes new evidence on the distribution of the gender wage gap, and contributes to a wider literature on technological change, occupational sorting and wage inequality.

JEL classifications: J16, J21, J24, J31, O33

Keywords: Female labor force participation, Gender wage gap, Technological change, Supply-Demand framework, Task-based approach, Wage distribution, Wage inequality

*We are grateful to Brian Bell, Carlos Carrillo-Tudela, James Fenske, Simon Quinn, Climent Quintana-Domeque, Marco Manacorda, John Muellbauer and seminar participants at University of Essex, University of Oxford, Universidad del Rosario, LACEA, and SAEe for their comments, and to Atheendar Venkataramani for many helpful discussions.

†University of Essex: srbhal@essex.ac.uk

‡University of Essex: m.fernandezsierra@essex.ac.uk

1 Introduction

The secular increase in the labor force participation of women (FLFP) is widely regarded as one of the most salient features of the labor market over the last century (Killingsworth and Heckman, 1987; Costa, 2000; Goldin, 2006; Fogli and Veldkamp, 2011; Fernández, 2013; Goldin and Olivetti, 2013). Nevertheless, there is surprisingly limited evidence of how this massive change in the composition of the labor force has altered the wage distribution. Economic theory suggests that, as long as men and women are imperfect substitutes in production, increases in women's labor supply will (i) create downward pressure on the wages of both men and women and (ii) create greater downward pressure on the wages of women, and hence widen the gender wage gap. The extent to which this is the case will depend upon the elasticity of substitution between male and female labor.

We structurally estimate the elasticity and demand parameters in an equilibrium model that extends the canonical labor demand-supply model discussed in Katz and Autor (1999) (also see Katz and Murphy (1992); Murphy and Welch (1992) and Card and Lemieux (2001)). Aggregate production is described by a nested-CES production function in which types of labor defined by their gender, skill and occupation are allowed to be imperfect substitutes. Following Johnson and Keane (2013), we model occupational choice using a random utility framework, incorporating both individual preferences and equilibrium returns to labor in the decision problem. The four types of labor, male vs female and unskilled vs skilled, either choose home production or select among three occupational categories. Following Autor et al. (2003), these are occupations intensive in abstract, routine or manual tasks.

Our first innovation is to allow the elasticity of substitution between male and female labor to be specific to the occupation group they are in. As mean wages in manual, routine and abstract task-intensive occupations follow a clear ordering, our model allows the elasticity to vary across the earnings distribution. In line with this, impacts of rising FLFP on both the gender wage gap and wage inequality within gender will tend to vary across the wage distribution. A second contribution we make is to endogenize labor force participation. This is relevant because the standard labor supply and demand model assumes a vertical (inelastic) short-run relative labor supply (Katz and Autor, 1999), making it unsuitable to study the dynamics of FLFP. The observed evolution of the wage structure will of course depend not only on relative (female vs male) changes in labor supply, but also on relative demand trends. A third contribution of this paper is that we allow for differential demand trends across occupation, education, and gender groups. This feature allows us to quantify the importance for the wage structure of gender, skill, and occupation-biased technical change independently, but within a unified model. We show that this flexibility is important to account for the patterns in the data.

We apply this framework to investigate impacts of the rapid rise of women's labor force participation on the wage structure in Mexico. Starting from about 1990, Mexico has experienced one of the largest increases in FLFP in the world during the last quarter century (Ñopo, 2012; The World Bank, 2012). FLFP among women aged 25-55 increased 21 percentage points (50%), from close to 40 percent in 1990 to close to 60 percent in 2013, rising from 4.7 to 14.7 million. In this same period, the average real wage fell, and the mean (median) gender earnings gap increased (by close to 5 (9.4) percentage points). That the gender earnings gap increased in favour of men is consistent with the increase in the relative supply of women exerting more downward pressure on women's wages than on men's wages.[1]

However, what motivates this analysis is that changes in the gender earnings gap varied dramatically across the earnings distribution. The unconditional earnings gap widened by 39 percentage points at the 5th percentile of the wage distribution, while narrowing by 18 percentage points at the 95th percentile of the distribution. Since a large share of the additional women in the workforce were unskilled, the skill composition of the labor force changed in line with the increase in FLFP. However, the skill (and age) adjusted earnings gap evolves similarly. It widened by about 13 percentage points at the 5th percentile, and narrowed by 11 percentage points at the 95th percentile. Following Firpo et al. (2007, 2009), we conducted a decomposition of the gender earnings gap at every percentile of the distribution. These estimates indicate that changes in the gap are not readily explained by composition changes in education, experience and interactions between these characteristics. The decomposition results are, as is always the case, based on partial equilibrium counterfactuals, which do not take into account that changes in relative supplies may affect relative wages. We therefore rely upon the structural estimates, and these are obtained conditional on skill.

Our structural estimates are able to explain the distributional patterns in the data. Our first result is that male and female labor are closer substitutes in high-wage abstract task-intensive occupations (elasticity of 2.6) than in the lower-wage manual or routine task-intensive occupations (elasticity of 1.2 in each case). This implies that increasing FLFP can have had distributional effects. In particular, it can explain why the increase in FLFP in Mexico exerted greater downward pressure on earnings at the lower end of the distribution. Our second result is that demand trends favored female workers in general, and more so among college-educated workers in abstract task-intensive occupations. Demand trends acted to mitigate the downward pressure on female wages arising from the increase in FLFP. This can explain why there was a narrowing of the gap at the upper end of the distribution (in fact, in the top 30% of the distribution). Our third finding is that the strong downward trend in fertility in Mexico, together with the shallower negative trend in marriage, can

[1] The gender earnings gap is defined here as the logarithm of male earnings per hour minus the logarithm of female earnings per hour.

together account for about 22% of the increase in FLFP, the rest being explained by preferences.

Although this was not our central focus, we also estimate the college premium and the elasticity of substitution between skilled and unskilled labor by occupation group. As may be expected, we find that unskilled labor more easily substitutes skilled labor in manual than in abstract-intensive occupations, with the elasticity for routine lying closer to that for abstract tasks. We use the estimated model parameters to conduct counter-factual exercises that allow us to specify how the earnings wage gap by occupation and skill group would have evolved if any of the female labor supply, demand, fertility, marriage or home-production preference channels were shut down. The counterfactual exercises show that, if FLFP were to have remained constant at its 1990 level, the mean gender earnings gap among unskilled workers would have declined rather than increased, and that the average gap among skilled workers would have declined even more than it actually did. So, clearly, the increase in FLFP acted to widen the gap in favour of men for both skill groups. We obtain a similar result when we hold demand at its 1990 level, consistent with demand trends having favoured women across the skill distribution.

The rest of this section places our work in the context of existing work. This paper contributes to the literature on the task-based approach of Autor et al. (2003); Acemoglu and Autor (2011); Autor and Dorn (2013); Altonji et al. (2014); Goos et al. (2014); Michaels et al. (2015) in being the first to introduce gender into the frame. Previous studies in this literature have emphasized that complementary or substitutability between the factors of production is determined by the type of tasks in which they are employed, but they have typically looked at technology or capital substituting labor, rather than at female labor substituting male labor. This paper provides the first estimates of the elasticity of substitution between (skilled vs unskilled) male and female labor by task-based occupation (and, implicitly, across the earnings distribution).

Identifying the extent to which male and female labor are substitutable in different tasks is pertinent. First, there is recent evidence of a growing demand for (non-cognitive) social skills (Deming, 2017), and evidence that women have stronger social skills (Jaimovich et al., 2017). Second, the importance of manual (brawn-intensive) skills relative to cognitive skills has shown a global decline. Given that men have a comparative advantage in brawn-intensive tasks but not in cognitive tasks, this too has led to an increase in the relative demand for female labor, tending to result in a narrowing of the gender earnings gap. (Galor and Weil, 1996; Blau and Kahn, 1997; Weinberg, 2000; Rendall, 2010; Black and Spitz-Oener, 2010; Pitt et al., 2012; Aguayo-Tellez et al., 2013; Rendall, 2013). Another line of work has emphasized the marketization of home production and the growth of service industries to the same effect. (Lup Tick and Oaxaca, 2010; Akbulut, 2011; Olivetti and

Petrongolo, 2014; Ngai and Petrongolo, 2017).[2] A further contribution we make, relative to the partial equilibrium studies of factors leading women to join the workforce, is to provide estimates of how both female and male LFP in the unskilled and skilled groups respond to changes in fertility and marriage, and to estimate the role of preferences. Our finding that preferences played a large role is in line with Fogli and Veldkamp (2011) and Fernández (2013).

As discussed earlier, an alternative way of framing our contribution, starting from a different point of departure, is that it extends the often used equilibrium model of Katz and Murphy (1992); Murphy and Welch (1992); Katz and Autor (1999); Card and Lemieux (2001) in three ways. Here we delineate more carefully how our approach differs from the existing literature. First, by endogenizing labor supply, we provide a new modelling framework for analysis of increases in women's LFP. Second, by allowing the elasticity of substitution between different factors of production to vary across occupations, we effectively merge two strands of the literature (Katz and Murphy (1992) and Autor et al. (2003)). The closest relative to this aspect of our work is Johnson and Keane (2013), whose analysis is motivated to explain the evolution of the entire wage structure, while we focus upon the male-female wage "premium" because we are looking at a massive increase in FLFP.[3] Third, by allowing demand trends to vary by occupation, education, and gender, we take forward a literature (cited earlier) in which the gender wage gap or the skill premium has been analysed with respect to either task-biased or skill-biased or gender-biased technical change, rather than allowing that all are at play at once.[4]

Only a few previous studies have investigated impacts of female labor supply on changes in the wage structure. Topel (1994) examined whether the rise in female labor supply contributed to rising inequality in the U.S. during the 1970s and 1980s, concluding that it did, by depressing the wages of low-skilled male workers. Juhn and Kim (1999)challenged this result, arguing that it was dissipated by accounting for changes in relative demand. Instead, they argued, in line with our findings, that college-educated women are close substitutes for college-educated men, so that their entry into the labor market may have tempered the growth in male wage inequality in the 1980s.

[2]The cited studies tend to analyse the partial equilibrium; we differ in simultaneously analysing supply and demand.

[3]As noted above, the unique feature of our analysis is that we allow the elasticity of substitution between male and female labor to vary by occupation. Much of the literature on the wage structure is descriptive or partial equilibrium in nature (see the discussion in Johnson and Keane (2013)). In equilibrium models previous to Johnson and Keane, following Heckman et al. (1998a,b) who distinguish labor only by skill (college and high school), Lee (2005) differentiates labor by both education and occupation (white- vs. blue-collar) to get four types. Lee and Wolpin (2006, 2010) allow further differentiation of labor by occupation, education, gender, and age, but they assume they are perfect substitutes in production.

[4]Many of these studies were cited earlier. For instance, Bound and Johnson (1992) and Katz and Autor (1999) focus on skill-biased technological change, Pitt et al. (2012) allow for gender-biased technical change, Goos et al. (2014) allow for routine-biased technical change, but no previous work allows for all.

Only two studies appear to have attempted to directly estimate the elasticity of substitution between male and female labor. Exploiting state level variation in U.S. military mobilizations for World War II, Acemoglu et al. (2004) report estimates of the elasticity of substitution between male and female labor of around 3. The authors qualify this finding arguing that this elasticity potentially varies across skill groups, but they do not test that hypothesis. Johnson and Keane (2013) estimate a dynamic equilibrium model of the US labor market fitted to assess factors driving changes in the wage structure from 1968 to 1996. They report an elasticity of substitution between male and female labor of between 1.85 and 2.2. No previous work estimates elasticities by task-based occupation. The range of the elasticities we estimate across the distribution is consistent with existing estimates but the heterogeneity we find is quantitatively important and has significant distributional consequences.[5]

The speed with which FLFP has recently grown in Mexico is rather unusual, but there were similarly large increases in female labor force participation rates (of married women) between about 1940 or 1950 and 1980 in the OECD countries. For instance, FLFP increased by about 50% (from around 25% to 52%) between 1940 and 1980 in the US. The average gender pay gap increased in favour of men in this period. It only started to narrow in the late 1970s and this convergence has slowed since 1990 (Bailey and DiPrete, 2016; Blau and Kahn, 2017). We are unaware of any analysis of how increases in FLFP in OECD countries influenced the relative wage of women across the wage distribution, the distribution of gender-specific skill premia and, by implication, the evolution of gender-specific inequality within and between women and men.

It was not our motivating focus, but this paper also contributes to a growing literature analysing the sharp fall in earnings inequality in Latin America since the late 1990s, which has been attributed to a fall in skill and experience premia (López-Calva and Lustig, 2010; Levy and Schady, 2013; Lustig et al., 2013; Galiani et al., 2017; Fernández and Messina, 2018). Existing work concentrates on the male earnings distribution to avoid dealing with selection issues arising from the increase in FLFP. However, increases in FLFP have been of such a magnitude that their potential distributional effects are large. This is the first study to endogenize this and provide estimates of changes in skill premia for both men and women in a general equilibrium framework.

Our results confirm that Mexico experienced a significant contraction of the male earnings distribution over the last 25 years, driven by higher wage growth among low-skilled workers than among the high-skilled. Our analysis adds two further insights: First, there was no corresponding decline in inequality among women, the wages of low skilled women being depressed by rising female labor supply. Sec-

[5] Available estimates are all for the US, we know of no estimates for other high income countries or for low and middle income countries.

ond, sluggish wage growth among high-skilled men in Latin American countries may be partly explained by the incorporation of college educated women into the workforce. This channel has been mostly overlooked by economists working on earnings inequality in this region.

The rest of the paper is organized as follows: Section 2 discusses the data, Section 3 presents the main stylized facts, reviewing the evolution of male and female labor force participation rates, and documenting changes in the wage and occupational structure over the last quarter century in Mexico. It also presents a decomposition of the change over time in the gender earnings gap across the earnings distribution. In Section 4 we formulate an equilibrium model of the labor market, and describe the empirical strategy used to estimate its parameters. The results are presented in Section 5, and robustness exercises using alternative specification of the model and different measures of labor supply are in Section 6. Section 7 concludes.

2 Data

We use the Mexican Household Income and Expenditure Survey (ENIGH), which is a nationally representative household survey carried out by the Mexican National Institute of Statistics and Geography (INEGI).[6] There are 13 waves of the ENIGH in our sample, covering the years between 1989 and 2014.[7] As a check for consistency, we replicate part of the results using data from the 1960, 1970, 1990, 2000 and 2010 Mexican CENSUS.[8]

Our measure of individual incomes only includes remuneration from labor. In particular, we use the monthly monetary remuneration from labor in all occupations,[9] which include wages, salaries, piecework, and any overtime pay, commissions, or tips usually received, but excludes income received from government transfers or profits from self-employment work.[10] Monthly earnings are converted into hourly rates dividing by the worker's total hours of work per week in all jobs multiplied by the usual number of weeks in a month.

Each year, between 77 and 83 percent of workers reported having worked full time (35 hours or more in the previous week), but there are clear differences between

[6]The main labor force survey in Mexico is the Encuesta Nacional de Ocupación y Empleo (ENOE), which replaced the Encuesta Nacional de Empleo (ENE) and the Encuesta Nacional de Empleo Urbano (ENEU) in 2005. The transition from the ENE/ENEU to the ENOE led to methodological changes that difficult comparability over time. To guarantee a time consistent series of employment and earnings we use the ENIGH survey.

[7]The first three waves of ENIGH were collected intermittently (1984, 1989, and 1992), but from 1992 onwards the survey was collected biannually, with an exceptional gathering in 2005

[8]The CENSUS data corresponds to the extended questionnaire applied to a 10 percent sample of the population in each year.

[9]Only after the year 2008 it becomes possible to separate the remuneration arising from the main job from that of any secondary jobs. To keep a consistent series, we add up remuneration from labor in all jobs in every year.

[10]The extended questionnaire of the CENSUS only asks about overall remuneration from labor, which can include income from self-employment. Hence, the two series are not fully comparable.

men and women: on average, full-time female workers represent between 62 and 67 percent of all female workers, while the share of full-time male workers ranges between 87 and 90 percent. In order to have groups that are more closely comparable, the earnings series is calculated using incomes from full-time workers only.[11] All earnings are transformed into real U.S. Dollars of 2012 using the Mexican Consumer Price Index and the purchasing power parity adjusted exchange rate estimated by the IMF. There were exceptionally low and high values of hourly earnings in the data, so we only use earnings from workers that reported having hourly rates above \$0.1 and below \$150.[12]

We classify jobs using the principal group level of the Mexican occupation classification (CMO), which has 18 groups that can be consistently followed throughout the period of analysis. Occupations are classified in three groups defined by whether the activities performed in the jobs are predominantly manual, routine, or abstract. The division is based on the measures constructed by Autor et al. (2003) from different sets of variables from the U.S. 1977 Dictionary of Occupational Titles (DOT). Details of the homogenization of the occupation classification and the construction of the three aggregate groups are presented in Appendix A.2. The final division is shown in Table 3.

Finally, the sample is limited to workers between the ages of 25 and 55 (henceforth prime-age workers). This is done to guarantee we are working with a group that has a strong labor market attachment, and to ameliorate selection problems arising from changes in educational and retirement choices of different cohorts.

3 Trends in Women's Labor Supply and Relative Earnings

In 1989, the labor force participation rate amongst prime-age population in Mexico was close to 64.2 percent, the FLFP rate was as low as 36 percent, and women accounted for only 29 percent of the workforce. By 2014 this picture had changed drastically: overall participation rates had reached 76 percent, the FLFP rate was close to 58 percent, and women represented 41 percent of the workforce (see Panels (a) and (c) Figure 1).

Although not directly comparable to the ENIGH, the labor force participation numbers from the decadal Mexican CENSUS[13] corroborate this trend, and allow us to go back in time as far as 1960 (see Panels (b) and (d) of Figure 1). Between 1960 and 1990, the FLFP rate increased 11 percentage points, going from 12 to 23

[11]Results of the main estimates including the earnings of part-time workers, or alternative measures that account for differences in the intensive margin, are presented in the robustness section of the paper.

[12]In no year did the percentage of the observations trimmed exceeded one percent. Moreover, the results in the paper are qualitatively unchanged if we include the outliers in the sample.

[13]The differences between the CENSUS and the ENIGH are due to the fact that the CENSUS only includes as economically active those individuals whose primary activity was either working or looking for a job. For example, part-time workers whose primary activity was studying are categorized as outside the labor force, leading to an underestimation.

percent. But in the two following decades the rise of the FLFP rate accelerated: between 1990 and 2010 the rate increased by 22 percentage points, reaching 45 percent in 2010. Cross-country comparisons show that the rise of FLFP in Mexico during the last 25 years was the largest in the Latin American region (Ñopo, 2012), and one of the largest in the world (The World Bank, 2012).

Three stylized facts characterize the evolution of labor force participation during this period. First, most of the gains came from low skilled women joining the workforce: the FLFP rate amongst prime-age women with at most high school education went from 35.7 to 55.4 percent between C.1990 and C.2013,[14] while that of college educated women went from 71.7 to 77.4 percent (see Table 1). Second, the increase in participation happened across all age groups: between C.1990 and C.2013, the FLFP rate increased by 18.7 percentage points amongst women between the ages of 25 and 34, and by 23.1 percentage points amongst women in both the 35 to 44 and the 45 to 55 age groups. Finally, the LFP rate amongst prime-age male population was above 94 percent and stable across the period.

At the same time that women were increasingly joining the worforce, the wage structure in Mexico changed substantially. Figure 2 shows the change between C.1990 and C.2013[15] of log hourly earnings at different percentiles of the earnings distribution, calculated separately for men and women. There are four main takeaways from this figure: First, real wages tended to fall during the last quarter century in Mexico. This is mostly explained by the two major crisis the country faced since the early 1990s: the 'Tequila Crisis' of 1994 and the Great Recession of the late 2000s. Since we are focusing on relative wages, we will not concentrate in this stylized fact. Second, the male wage distribution contracted sharply over the period: wage growth at the lower-tail of the distribution was higher than at the upper-tail. A similar pattern was also observed in most countries in Latin America during the same period (López-Calva and Lustig, 2010; Levy and Schady, 2013; Lustig et al., 2013; Galiani et al., 2017; Fernández and Messina, 2018). Third, wage growth for females was slightly u-shaped across the distribution: wages at the bottom and the top performed better than at the middle. Finally, growth (decline) in earnings among workers at the bottom was significantly higher (lower) for males than for females, while the exact opposite is true among workers at the top.

The implication of the last three takeaways is that the gender earnings gap changed in very different ways across the distribution. Panel (a) of Figure 3 shows the change in the gap between C.1990 and C.2013 at different percentiles of the distribution, which we calculate by subtracting the values of the male and female series in Figure 2. The gender earnings gap increased between 10 and 32 percent among workers below the median, but declined between 5 and 18 percent among

[14] We joined together surveys from 1989 and 1992 (C.1990), and from 2012 and 2014 (C.2013) to increase sample size of the ENIGH survey

[15] To increase sample size in the ENIGH survey, we joined together surveys from 1989 and 1992, and from 2012 and 2014.

workers above the 80th percentile. Panel (b) of Figure 4 shows a similar series, but this time using changes in the conditional wage distributions. We do this by running conditional quantile regression at each percentile in two periods: C.1990 and C.2013. We then report the change in the coefficient of the female dummy in the regression. The controls we use in the regression include: dummies for 7 education categories, dummies for 6 age categories in five year intervals, and all possible interactions. The negative monotonic change in the gender wage gap across percentiles is also observed after conditioning on relevant observable characteristics, but the magnitudes of the changes are lower and there is no longer convergence at the top.

3.1 Did the Wage Structure Change? Quantile Decomposition of the Gender Earnings Gap

The wage distribution can change because the characteristics of workers are changing, because the returns to those characteristics are changing, or by a combination of both. Since the rise of FLFP was not the only major change in the Mexican labor market over this period, other compositional shifts, absent of any changes in relative pay, could also be consistent with the observed trends.[16]

The lower panel of Table 2 shows how the age and educational attainment of male and female workers changed between C.1990 and C.2013. The share of prime-age women in the workforce with at least some college education increased from 14.5 to 24.0 percent, while that of men went from 15.6 to 20.8 percent: the larger share of women with more schooling can be an alternative explanation for the convergence in earnings at the top of the distribution. Also, the average Mexican worker is becoming older, more so if they are women: the share of workers between 45 and 55 years of age increased from 20.2 to 28.9 percent in the case of women, and from 23.0 to 29.2 percent in the case of men: if the gender earnings gap increases with age (Barth et al., 2017), this age compositional change can also be a factor behind the divergence in earnings at the bottom of the distribution.

Has the wage structure changed or is there a simple re-composition of the workforce? A decomposition exercise can help disentangle the importance of these two potential drivers. The idea is to exogenously fix the structure of relative earnings at the average level across the last 25 years, and then quantify the counterfactual levels of the gender earnings gap at different percentiles under the observed compositional changes. Alternatively, we can keep the composition of the labor force fixed at a given point in time and construct counterfactual earning gaps to evaluate how changes in schooling, age and occupational premiums have affected the observed

[16] A similar argument was made by Lemieux (2006) in the context of the debate about the rise of income inequality in the United States. Lemieux shows evidence that a substantial share of the the rise in residual earnings inequality in the U.S. – the largest component of overall earnings inequality – can be accounted by the fact that the earnings of workers that are older and have more schooling –both of which have increased their share in the workforce – tends to be more dispersed. This argument was later controverted by Autor et al. (2008).

dynamics. These are partial equilibrium counterfactuals, something that we will return to in the next section.

The decomposition methodology is based on Firpo et al. (2007, 2009), and the specific details can be found in Appendix A.3. As a starting point, consider a transformed wage-setting model of the form:

$$RIFq_{\tau,gen,t} = X'_{gen,t}\gamma_{gen,t} + \epsilon_{gen,t}, \tag{3.1}$$

where subscript gen indicates if the worker is male ($gen = k$) or female ($gen = f$); the subscript t indicates the period, either initial ($t =$C.1990) or final ($t =$C.2013); $RIFq_{\tau,gen,t}$ represents the value of the RIF corresponding to the τ'th quantile of the earning distribution at time t and for gender gen; X is a vector of socio-demographic characteristics including dummies for 7 education categories, dummies for 6 age categories in five year intervals, and all possible interactions; and $\epsilon_{gen,t}$ is the error term assumed to have zero conditional mean. We can estimate Equation (3.1) for each gender and period separately by OLS, and then express the estimated difference over time of the expected value of the earnings quantile $\hat{q}_\tau$ as:

$$\Delta_t\hat{q}_{\tau,gen} = \underbrace{\left(\overline{X'}_{gen,\text{C.2013}} - \overline{X'}_{gen,\text{C.1990}}\right)\hat{\gamma}_{gen,P}}_{\Delta_t\hat{q}_{X,\tau,gen}} + \underbrace{\overline{X'}_{gen,P}\left(\hat{\gamma}_{gen,\text{C.2013}} - \hat{\gamma}_{gen,\text{C.1990}}\right)}_{\Delta_t\hat{q}_{S,\tau,gen}}, \tag{3.2}$$

where overbars denote averages, and $\hat{\gamma}_{gen,P}$ and $\overline{X}_{gen,P}$ correspond to the estimated vectors of parameters and the explanatory variables of a wage-setting model in which observations are pooled across the two periods.[17] Here, $\Delta_t\hat{q}_{X,\tau,gen}$ corresponds to the composition effect, which captures the part of the change in the τ'th earnings quantile that is accounted for by changes in the average skill-demographic and occupational composition of workers, given that we set the returns at their (weighted) average over the two periods; and $\Delta_t\hat{q}_{S,\tau,gen}$ is the wage structure effect, and captures how changes in returns are affecting earnings at the quantile τ, given that the observable characteristics are fixed to be equal to their (weighted) average over time.

Since we are interested in the effects of composition and price changes on the gender earnings gap, we construct the following measures at 19 different percentiles:

[17] This specific counterfactual allows us to analyse composition and wage structure effects relative to a baseline defined by the (weighted) mean returns and (weighted) mean characteristics over the two periods, eliminating the interaction term present in other decompositions (Oaxaca and Ranson, 1994).

$$\underbrace{\Delta_t \hat{q}_{\tau,k} - \Delta_t \hat{q}_{\tau,f}}_{\text{Overall}} = \underbrace{\left(\Delta_t \hat{q}_{X,\tau,k} - \Delta_t \hat{q}_{X,\tau,f}\right)}_{\text{Composition}} + \underbrace{\left(\Delta_t \hat{q}_{S,\tau,k} - \Delta_t \hat{q}_{S,\tau,f}\right)}_{\text{Wage Structure}}. \tag{3.3}$$

The results of the decompositions for 5 selected percentiles are shown in Table 4. The table reports the observed change in the log (male/female) earnings ratio between C.1990 and C.2013, and the contribution of composition and wage structure effect to that change. The main conclusion from the decompositions is that wage structure effects are quantitatively more important than pure compositional effects. Estimated wage structure effects contribute 63 percent of the observed rise in the gender earnings gap at the 5th percentile, and close to 90 percent at the 25th percentile. Moreover, they over-predict the fall in the gap at the 95th percentile (-22.5. log points observed vs. -34.7 log points attributed to the wage structure).

The relative importance of composition and wage structure effects on the evolution of the gender wage gap can be seen clearly in Figure 7. Here we plot the results of the decomposition at the 19 percentiles. It is clear from the figure that wage structure effects are the dominant factor. The figure also shows that if the wage structure had remained constant at the average levels over the two periods, compositional effects would tend to lead to a larger counterfactual gender earnings gap. This is indicative that pure compositional shifts are contributing to the expansion of the gap at the lower tail, but have impede further convergence at the top of the distribution.

3.2 Explaining Changes in the Wage Structure

The evidence from the decompositions support the claim that the structure of relative pay between males and females in Mexico has in fact changed, but the explanation for this change remains an open question. Our hypothesis is that these diverging dynamics of relative earnings across the distribution are associated to the rise of FLFP. In particular, that the influx of women into the labor market generated downward pressures on female wages, especially among the less educated.

The underlying theoretical base for this hypothesis is the canonical supply and demand framework (Katz and Murphy, 1992; Murphy and Welch, 1992). In this framework, changes in the composition of labor supply can directly affect the wage structure. How important is this channel fundamentally depends on the value of the elasticity of substitution between the groups being compared. A simple example can help clarify the mechanism. Suppose that workers in an economy were similar in every respect except for their gender, and that aggregate production could be characterized by a Constant Elasticity of Substitution (CES) function of the form:

$$Y_t = \left[\alpha_t L_{k,t}^{\rho} + (1-\alpha_t) L_{f,t}^{\rho}\right]^{1/\rho}, \tag{3.4}$$

where Y_t is total output at time t; $L_{k,t}$ is the total supply of labor from male workers; $L_{f,t}$ is the total supply of labor from female workers; α_t is a time-varying 'share' parameter that captures differences in the intensity of labor demand between male and female labor, which is allowed to change in time; and $\rho \in (-\infty, 1]$ is a function of the elasticity of substitution (σ_ρ) between male and female labor: $\sigma_\rho \equiv \frac{1}{1-\rho}$. If the economy is operating along the demand curve, wages are equated to marginal productivities and the log (male/female) earnings ratio takes the form:

$$\log\left(\frac{W_{k,t}}{W_{f,t}}\right) = \log\left(\frac{\alpha_t}{1-\alpha_t}\right) - \frac{1}{\sigma_\rho}\log\left(\frac{L_{k,t}}{L_{f,t}}\right), \tag{3.5}$$

where $W_{k,t}$ and $W_{f,t}$ are the wages of male and female workers respectively.

Equation (3.5) shows that the evolution of the log (male/female) earnings ratio depends on two factors: (i) how the relative supply of labor between males and females is changing $\left(\frac{L_{k,t}}{L_{f,t}}\right)$, scaled by the inverse of the elasticity of substitution (σ_ρ); and (ii) how relative demands are changing, as captured by the variation in time of the log ratio of the α_t share. If male and female labor are imperfectly substitutable in production, that is, if σ_ρ has a 'low' value, and relative demands trends are constant, large shifts in female labor supply would impose downward pressure on female wages, and to a lesser extent to male wages, leading to an increase of the gender earnings gap. If, on the other hand, male and female workers are close to perfect substitutes – 'high' values of σ_ρ –, a rise in female labor force participation should depress *both* male and female wages, with relative earnings remaining fairly constant.

There is no general agreement in the literature about the value of σ_ρ, but at least three studies have provided estimates for the United States. Weinberg (2000) finds that, assuming a constant increase in the demand for female labor during the 1980s and 1990s, the value of this elasticity is close to 2.4. Acemoglu et al. (2004), exploiting state level variation in U.S. military mobilizations for World War II, report estimates of the short-run elasticity of substitution between male and female labor of around 3. Johnson and Keane (2013), using a dynamic equilibrium model of the labor market, report an elasticity of substitution between male and female labor of between 1.85 and 2.2. If we take these estimates at face value, and taking into account that the log (male/female) relative supply in Mexico declined by approximately 50 log points (see Table 5), the log (male/female) earnings gap should have increased between 16.7 and 27.7 log points, absent any changes in the demand side.

We further advance a second hypothesis: the effect of FLFP on the gender earnings gap varies throughout the pay distribution because the elasticity of substitution between male and female labor varies across occupation groups. The theoretical base for this hypothesis comes from the task-based approach popularized by Autor et al. (2003). In this framework, the complementary or substitutability between factors of production is determined by the type of tasks in which they are employed.[18] Any job requires, to a higher or lesser degree, cognitive, manual, physical, socio-emotional, and interpersonal skills. The relative importance of any subset of skills is then a function of the specific activities that workers are performing. As long as there is some difference in the bundle of skills that men and women supply to the labor market, the substitutability of male and female labor will vary across occupations.

We use the standard typology of the task-based framework[19] and aggregate the 18 occupations of the ENIGH into three groups: abstract, manual, and routine task-intensive (see Table 3).[20] This division has the advantage that it clearly separates occupations into groups defined by which skills and aptitudes are more relevant in the jobs, but also because it breaks the earnings distribution in three segments. For example, manual task-intensive occupations have more demands for physical than abstract analytical aptitudes, and require skills like strength and hand, eye, and foot coordination; these occupations are mostly found at the lower end of the pay distribution, and include examples like agriculture, services, and transportation. Abstract task-intensive occupations have more demands for analytical aptitudes than for physical or manual labor, and require skills like quantitative reasoning, direction, control, and planning of activities; these occupations are found at the top of the pay distribution, with the most prominent examples being professionals and managers. Finally, routine task-intensive occupations are in a middle ground, with a mixture of physical and analytical demands, and were aptitudes like adaptability to repetitive work and finger dexterity play a significant role; these occupations tend to be located at the middle of the pay distribution, with examples being clerical and crafts and trades jobs.

There are two additional points worth mentioning about the occupation aggregation we use (see Table 3). First, the division is balanced in size, with each group representing about one third of the workforce. Second, there is substantial occupational sorting by sex within the three groups. For example, in the abstract task-intensive group, 71 percent of managers and 84 percent of craft and trades

[18] This argument holds true when we think about the role that information and communication technologies have played in substituting workers in middle-income routine occupations, the phenomena that has received most attention from academics, but it also holds when we think about differences by gender or any other dimension.

[19] Other papers that have used this typology include Autor et al. (2006); Goos and Manning (2007); Dorn (2009); Rendall (2013); Autor and Dorn (2013); Adda et al. (2017).

[20] Appendix A.2 discusses the details of the division of the principal level occupations along the three groups.

supervisors are males, while 61.8 percent of workers in education are female. Segregation by sex is even stronger at low-paying manual task-intensive occupations. On average, 99 percent of workers in transport occupations and 93 percent in protective services are male, while 92.4 percent of workers in domestic services and 56.6 percent in general services are female.

In the labor supply and demand framework, we observe how supplies and wages are changing, but we don't see the demand trends nor the values of the elasticities of substitution. Panels (a)-(c) of Figure 5 show the time series of the two parts of equation that we do observe. Each panel corresponds to one of the three occupation groups. In each panel we report the evolution of the log (male/female) earnings ratio and the log (male/female) relative supply.

The figure shows some suggestive evidence that the sharp change in the gender composition of labor supply in Mexico was negatively correlated with the change in the gender earnings gap, but mostly in low-paying occupations. Between 1989 and 2014, log (male/female) relative supply declined by 44 log points in both abstract and routine task-intensive occupations, and by 69 log points in manual task-intensive occupations. During the same period, the log (male/female) earnings ratio only increased in the lower-paying jobs: 10 log points in manual task-intensive occupations and 2.33 log points in routine task-intensive occupations. In the abstract task-intensive group the gender earnings gap fell by 10 log points.

In the next section we develop an equilibrium model of supply and demand for labor to try to test the two hypothesis presented in this section. The objective of the model is threefold: first, to explore if a general equilibrium model combining the task-based approach with the canonical supply and demand framework is able to recreate the patterns and changes in the Mexican wage structure over the past 25 years; second, to estimate the key structural parameters to test the hypothesis of differential degrees of substitutability between male and female labor across the occupational distribution; finally, to run counterfactual exercises and get quantitative estimates of the effect of the rise of female labor force participation on the gender earnings gap.

4 Theoretical Model

4.1 Demand Side

The model assumes that aggregate production in the economy is a function of the amount of labor that imperfectly substitutable types of workers supply to the market. Agents are divided into 4 types according to their gender – male or female – and their level of schooling – secondary education at most, referred as unskilled, or at least some college education, referred as skilled.[21] Each type of agent chooses

[21] In order to maintain a tractable number of parameters in the model, we omit the age dimension from this analysis.

between entering the workforce in one of three possible market occupations: abstract, routine or manual task-intensive; or, alternatively, the agents can opt to stay in home production.

Furthermore, the model assumes that the aggregate production technology can be described by a three level nested constant elasticity of substitution (CES) function, where each nest of the production technology corresponds to a given dimension: occupation, education, and gender. At the top level, output is produced by a CES combination of labor in the three types of market occupations:

$$Y_t = Z_t \left[\alpha_{1,t} L_{a,t}^{\rho_1} + (1-\alpha_{1,t}) \Big(\alpha_{2,t} L_{r,t}^{\rho_2} + (1-\alpha_{2,t}) L_{m,t}^{\rho_2} \Big)^{\rho_1/\rho_2} \right]^{1/\rho_1}, \tag{4.1}$$

where Y_t is total output at time t; Z_t is a scale parameter that is allowed to vary in time to capture skill-neutral technological change; $L_{a,t}$, $L_{r,t}$, and $L_{m,t}$ are the total supply of labor in abstract, routine, and manual task-intensive occupations respectively; $\rho_1 \in (-\infty, 1]$ is a function of the elasticity of substitution (σ_{ρ_1}) between labor in non-abstract – manual and routine – and abstract task-intensive occupations ($\sigma_{\rho_1} \equiv \frac{1}{1-\rho_1}$); and $\rho_2 \in (-\infty, 1]$ is a function of the elasticity of substitution (σ_{ρ_2}) between labor in routine and manual task-intensive occupations ($\sigma_{\rho_2} \equiv \frac{1}{1-\rho_2}$).[22] Finally, $\alpha_{1,t}$ is a time-varying share parameter that captures both differences in the intensity of labor used between non-abstract – routine and manual – and abstract task-intensive occupations, as well as movements in relative demands between them. Similarly, $\alpha_{2,t}$ captures both differences in the relative intensity of labor used between routine and manual task-intensive occupations, and movements in relative demands between them. General examples of possible sources of shifts in relative demand include non-neutral technical change (e.g. routine-biased technical change as in Goos et al. (2014)); variations in non-labor input demands (e.g. through capital skill complementarity as in Krusell et al. (2000)); product market demand shifts (e.g. through changes in the external demand for commodities as in Fernandez and Messina (2017)); and trade and outsourcing (e.g. generating incentives to modernize the production processes as in Juhn et al. (2014), and increasing competition with local industries as in Autor et al. (2016)).

In the second level of the production technology, labor in each occupation is divided in two groups based on the schooling level of the workers. In particular, labor in each occupation consists of a productivity weighted CES combination of labor from skilled workers, indexed by s, and labor from unskilled workers, indexed

[22] Note that by assumption, the elasticity of substitution between labor in abstract and manual task-intensive occupations is the same as the elasticity of substitution between abstract and routine task-intensive occupations. We consider this a natural way of organizing the three occupational groups since we observe they are align this way in the earnings distribution – low vs. high paying occupations –, but we present results using alternative specifications in the robustness section of the paper.

by u. That is:

$$L_{occ,t} = \left[\alpha_{3,occ,t} L_{s,occ,t}^{\rho_{3,occ}} + (1-\alpha_{3,occ,t}) L_{u,occ,t}^{\rho_{3,occ}}\right]^{1/\rho_{3,occ}} \quad \text{for} \quad occ = a, r, m, \tag{4.2}$$

where the parameters have an analogous interpretation to those in Equation (4.1).

Finally, in the third level of the production technology labor is disaggregated in each occupation-education group according to the gender of the workers. This is done using a productivity weighted CES combination of female workers, indexed by f, and male workers, indexed by k. That is:

$$L_{edu,occ,t} = \left[\alpha_{4,edu,occ,t} L_{k,edu,occ,t}^{\rho_{4,occ}} + (1-\alpha_{4,edu,occ,t}) L_{f,edu,occ,t}^{\rho_{4,occ}}\right]^{1/\rho_{4,occ}} \quad \text{for} \quad edu = s, u, \quad \text{and} \quad occ = a, r, m, \tag{4.3}$$

where the parameters have an analogous interpretation to those in Equation (4.1). Note that elasticities of substitution between male and female labor are allowed to vary between occupations, but not within occupation and education. Any differences in substitutability between labor by education group is captured by the relevant parameters in the second level of the production technology.[23]

The demand side of the model has two types of relevant parameters that we need to estimate: 8 parameters that are functions of the elasticities of substitution (ρ_1, ρ_2, $\rho_{3,a}$, $\rho_{3,r}$, $\rho_{3,m}$, $\rho_{4,a}$, $\rho_{4,r}$, and $\rho_{4,m}$); and a group of time varying relative productivities/demand shifters parameters (Z_t, $\alpha_{1,t}$, $\alpha_{2,t}$, $\alpha_{3,a,t}$, $\alpha_{3,r,t}$, $\alpha_{3,m,t}$, $\alpha_{4,s,a,t}$, $\alpha_{4,s,r,t}$, $\alpha_{4,s,m,t}$, $\alpha_{4,u,a,t}$, $\alpha_{4,u,r,t}$, and $\alpha_{4,u,m,t}$). As argued by Johnson and Keane (2013), it is possible to fit the trends in relative wages perfectly if we did not impose any restrictions on the evolution of the relative demand parameters, but this would mean that we would not be able to identify the parameters capturing the elasticities of substitution. we then restrict these relative productivities to follow a cubic trend in their natural logarithm.[24] For example, the parameter $\alpha_{1,t}$ is allowed to change according to:

$$\log \alpha_{1,t} = \alpha_{1,0} + \alpha_{1,1} t + \alpha_{1,2} t^2 + \alpha_{1,3} t^3. \tag{4.4}$$

[23] To test how sensitive are our results to the selection of the ordering of the levels in the production technology, we report results using alternative model specifications in the robustness section.

[24] The cubic trends provided the best fit of the model to the data. Quadratic polynomials were not flexible enough, while the coefficients associated to the quartic polynomials were not statistically significant in most cases. Results using alternative specification of the polynomials are available upon request.

A simple example is useful to understand how the different set of parameters are identified in the model. Note that under the assumption that the economy is operating along the demand curve, log relative earnings between male and female labor within a given occupation and education group can be expressed as:

$$\log\left(\frac{W_{k,edu,occ,t}}{W_{f,edu,occ,t}}\right) = \log\left(\frac{\alpha_{4,edu,occ,t}}{1-\alpha_{4,edu,occ,t}}\right) - \frac{1}{\sigma_{\rho_{4,occ}}}\log\left(\frac{L_{k,edu,occ,t}}{L_{f,edu,occ,t}}\right) \quad \text{for} \quad edu = s, u, \quad \text{and} \quad occ = a, r, m, \tag{4.5}$$

where $W_{k,edu,occ}$ and $W_{f,edu,occ}$ are the wages of the respective groups. In the model, the elasticities of substitution are identified by the movement in relative supply, something we specify in the next section, but the demand trends, as captured by the log ratio of the third order polynomials, are identified residually. In particular, any change in relative wages that is not explained by movements in relative quantities is then absorbed by the relative demand parameters.

In total, the demand side of the model has 56 parameters that we need to estimate.

4.2 Occupational Choice

Male and female workers sort themselves into different market occupations based on preferences about job flexibility and earning's profiles (Goldin, 1984, 1986; Adda et al., 2017); responding to societal expectations and attitudes towards female work (Brown et al., 1980; Goldin, 1984, 2006); and as a function of gender specific comparative advantages associated to differences in physical, sensory, motor, and spatial aptitudes (Galor and Weil, 1996; Black and Juhn, 2000; Rendall, 2010, 2013; Baker and Cornelson, 2016). The model attempts to incorporate both individual preferences and equilibrium returns to labor in the decision problem of the agents.

Following the work of Johnson and Keane (2013), we model occupational choice using a random utility framework where agents of different type choose either to enter the workforce in any of the three market occupations or remain in home production. Each alternative generates an utility for the worker, and agents choose the alternative that provides the highest utility. we model these utilities as linear functions that depend only on pecuniary and nonpecuniary rewards from each choice. In particular, the utility that a worker of a given type receives from choosing to enter the workforce in one of the three market occupations at time t is

$$U(occ \mid gen, edu, t) = \psi_{gen,edu,occ} + \psi_1 W_{gen,edu,occ,t} + \epsilon_{gen,edu,occ,t}, \tag{4.6}$$

where $\psi_{gen,edu,occ}$ are time invariant parameters that capture nonpecuniary rewards that a worker gets from choosing occupation *occ* at time t; and ψ_1 measures the weight (in utility terms) that a worker gives to labor earnings ($W_{gen,edu,occ,t}$). $\epsilon_{gen,edu,occ,t}$ is an idiosyncratic taste shock assumed to be independent and identically distributed extreme value. The assumption about the distribution of the taste shock generates a tractable Multinomial Logit form to the choice probabilities.

The utility from home production is modelled in a symmetric way. The economic literature has linked movements of women into the labor market to changes in fertility and contraceptive technology (Katz and Goldin, 2000; Costa, 2000; Cruces and Galiani, 2007); changes in the marriage market (Grossbard-Shechtman and Neuman, 1988; Fernández and Wong, 2014; Greenwood et al., 2016); changes in social norms and attitudes towards women's work (Ronald R. Rindfuss, 1996; Costa, 2000; Fernández et al., 2004; Goldin, 2006; Fernández, 2013); and improvements in capital and technologies used for home production activities (Costa, 2000; Greenwood et al., 2005; de V. Cavalcanti and Tavares, 2008; Coen-Pirani et al., 2010). we do not specify how the underlying mechanism that explain the rise of female labor force participation interact with the demand side of the model. What we do is to condition the decision to remain in home production on variables linked to fertility choice, marriage patterns, changes in preferences, and technical change that is specific to home production activities (e.g. appliances). In particular, the utility from choosing home production, denoted by h, takes the form:

$$\begin{aligned} U(h \mid gen, edu, t) =& \pi_{1,gen} + \pi_{2,gen} t + \pi_{3,gen,edu} Pr(\text{CHL5} = 1 \mid gen, edu, t) \\ &+ \pi_{4,gen,edu} Pr(\text{MARR} = 1 \mid gen, edu, t) + \epsilon_{gen,edu,h,t}, \end{aligned} \tag{4.7}$$

where $\pi_{1,gen}$ and $\pi_{2,gen}$ are the intercept and slope of a gender specific linear trend that captures both changes in preferences for home production over time, and changes in the technology used in home production activities; $Pr(\text{CHL5} = 1 \mid gen, edu, t)$ is the probability that the agent has a child under the age of 5, which, since the level of aggregation is at the gender-education level, corresponds to the proportion of the population from a given group that has at least one children under the age of five at time t; and $Pr(\text{MARR} = 1 \mid gen, edu, t)$ is the probability that the worker is married or has permanent partner. Finally, $\epsilon_{gen,edu,h,t}$ is a idiosyncratic taste shock assumed to be independent and identically distributed extreme value.

Panel (a) and (b) of Figure 8 show the trends of the two variables we incorporate in the home production utility function – the proportion of individuals with children under the age of 5, and the proportion of individuals that are married or with a permanent partner – both conditional on education and gender. The two variables can only be calculated for a sample restricted to the household head, and, if applicable, to his/her spouse or partner, so the trends should be interpreted with this caveat in mind.[25] The percentage of each group with children under the age of 5 was between 45 and 55 percent in 1989, but these numbers fell in all cases by close to 20 percentage points, so the same probability in 2014 ranged between 25 and 35 percent. The change in the fertility decisions of the Mexican population is expected to be a significant driver behind the rise in female labor force participation. Changes in the marriage market were less pronounced: in 1989, the percentage of each group that was either married or had a permanent partner was between 86 and 94 percent; by 2014, these numbers had fallen between 3 and 8 percentage points, with the largest decline observed among the college educated (7.4 and 8 percentage points decline for males and females respectively).

Given the assumed distribution of the idiosyncratic taste shocks, the probability that a worker chooses one of the market occupations or home production is

$$Pr(d_O = 1 \mid gen, edu, t) = \frac{\exp(U(O \mid gen, edu, t))}{\sum_{occ=a,r,m,h} \exp(U(occ \mid gen, edu, t))} \quad \text{for} \quad O = a, r, m, h. \tag{4.8}$$

We can use these probabilities to find the total labor supply of each type in each occupation. For example, the total supply of female workers with college education in abstract task-intensive occupations is

$$L^s_{f,s,a,t} = L_{f,s,t} \times Pr(d_a = 1 \mid f, s, t) \tag{4.9}$$

Where $L_{f,s,t}$ is the total number of female workers with college education at time t, which we take as given. As the example shows, we condition on the schooling level of the agents, but we assume that the educational choice was taken before the age of 25, the starting age for an agent to become part of the sample.

The supply side of the model has a total of 25 parameters that we need to estimate.

4.3 Equilibrium and Estimation

The amount of labor demanded from agents of a given type in a specific market occupation and moment in time, denoted by $L^d_{gen,edu,occ,t}$, is fully determined by the equilibrium condition that wages are equated to marginal productivities:

[25]The ENIGH survey started asking the question on marital status to all member of the household in 1996, and for the number of children since 2004. For these reason, the two variables can only be constructed for the household head, and, if applicable, to his/her spouse or partner.

$$W_{gen,edu,occ,t} = \frac{\partial Y_t}{\partial L^d_{gen,edu,occ,t}}. \tag{4.10}$$

Also, in equilibrium, wages are set to equate the supply and demand for the different labor types:

$$L^d_{gen,edu,occ,t} = L^s_{gen,edu,occ,t} \quad \text{for} \quad occ = a, r, m, \tag{4.11}$$

while the total supply in home production can be recovered residually.

A solution to the model is then obtained by finding the vector of wages for which Equations (4.10) and (4.11) are satisfied by each type of agent in each occupation. For a given vector of parameters, this can be accomplish in an iterative way using a fixed-point algorithm: (i) start with an arbitrary wage vector W^0 and find the total supply of workers from each type in each occupation, which can be calculated using Equations (4.8) and (4.9). (ii) Plug the estimated supply of workers of each type into the marginal productivity function defined by Equation (4.10) and calculate the new vector of wages W^1. (iii) If $W^0 = W^1$, we have a solution for this set of parameters. If $W^0 \neq W^1$, set $W^0 = W^1$ and go back to step (i).

The model generates a prediction of the wage and labor supply of the four worker types in the four occupations (including home production) at every time period. With 13 years of data there are $(12 + 16) \times 13 = 364$ predictions in total[26] that are a function of the 81 parameters. We fit the model to the ENIGH data using the method of moments, targeting observed labor shares and wages. Given the dimensionality of the problem, and the different scales of the moments being targeted, we assume a simplified error structure to facilitate estimation. Details on the estimation technique are presented in Appendix A.4.

5 Results

5.1 Model Fit to the Data

Figure 9 show the fit of the model with respect to the data. Panels (a)-(f) in the Figure show the log (male/female) relative earnings series and the log (male/female) relative supply series by occupation groups, both predicted and observed. Panels (g) and (h) show the observed and predicted participation rates for females and males respectively. The model predictions follow closely the long-term trends in the data. This is true for the earnings and supply series within the three occupations, as well

[26] Wages: 4 types of workers in three market occupation leads to 12 predictions per year. Supplies: 4 types of workers in 4 possible occupations leads to 16 predictions per year.

as for the participation rates. The model is less successful at capturing short-term variations, especially those related to temporary spikes in the data.

Table 6 shows the observed and predicted wages and occupation shares for all groups C.1990 and C.2013. There is a close fit of the model with respect to the occupation shares, with no major differences between the predictions and the data in any of the cells. The model is also able to accurately predict the mean wages of the different worker types across the three occupations. In a few cases, the model predicts an average wage that differs significantly from what the data shows, but it usually corresponds to a 'marginal' group. For example, the model predicts higher wages for college-educated workers in manual and routine task-intensive occupations C.1990, especially among females, but these groups only represent between 0.05 and 0.63 percent of the prime-age population according to the ENIGH.

5.2 Parameter Estimates

5.2.1 Demand Side

Elasticity of substitution between male and female labor. Table 7 reports the point estimates, standard errors, and 95 percent confidence intervals of the elasticities of substitution in the production technology. The elasticities of substitution between male and female labor are estimated to be around 1.2 in manual and routine task-intensive occupations. In abstract task-intensive occupations the point estimate is 2.6. These results support the hypothesis that the elasticity of substitution between male and female labor varies across occupation groups. In particular, male and female labor are closer substitutes in occupations that rely more on abstract analytical skills, which tend to include the highest-paying jobs.

To get a sense of what these values represent quantitatively, we can do some back-of-the-envelope calculations. Within manual task-intensive occupations, the log (male/female) relative supply fell between C.1990 and C.2013 by 56.4 log points (see Table 5), so an elasticity of 1.2 implies that the log (male/female) earnings ratio should have increased – holding everything else constant – by 46.5 log points, significantly more than the observed 6.4 log point increase. In routine and abstract task-intensive occupations, the log (male/female) relative supply fell by 38.9 log points, so the implied elasticities would predict an increase in the gender earnings gap of 31.9 and 14.7 log points respectively, also significantly higher to what was observed. There are two main takeaways from this calculations: first, the increase in FLFP put substantial downward pressure on the wages of female workers, especially in the lowest-paying occupations. Second, relative demand trends must have been strongly favourable to women, otherwise the gender earnings gap would be much higher.

To our knowledge, there are no comparable estimates of the elasticity of substitution between male and female labor for this or other countries in Latin

America. In the literature for the U.S., Weinberg (2000) finds that, assuming a constant increase in the demand for female labor between the 1970s and early 1990s, the value of this elasticity is close to 2.4. Acemoglu et al. (2004), exploiting state level variation in U.S. military mobilizations for World War II, report estimates of the short-run elasticity of substitution between male and female labor of around 3. Johnson and Keane (2013), after estimating a dynamic equilibrium model of the labor market fitted to replicate the patterns of the U.S. wage structure from 1968 to 1996, report an elasticity of substitution between male and female labor of between 1.85 and 2.2. We find that the values of these elasticities are in line with our estimates for Mexico, but that they ignore important heterogeneities that have significant distributional effects.

In the education dimension, the point estimates of the elasticities of substitution between skilled and unskilled labor are 1.4 in abstract, 1.6 in routine, and 3.6 in manual task-intensive occupations. This implies that workers with college and those with at most high school education are closer substitutes in occupation where abstract analytical skills are less important. These results also support the hypothesis that the sharp educational upgrading in Mexico contributed substantially to the fall of earnings inequality, and is part of the explanation for the contraction of the male wage distribution. In the next section we quantify these effects using counterfactual exercises.

The three point estimates of the elasticities of substitution between skilled and unskilled labor are in a range that incorporates the values that other studies have found in similar contexts. Using data from five Latin American countries during the 1990s, Manacorda et al. (2010) report estimates of this elasticity of around 3. Fernández and Messina (2018) expanded the model used by Manacorda et al. (2010) finding a value close to 1.25. In the literature focusing on the U.S., there is a general convergence among different studies showing that the elasticity of substitution between skilled and unskilled labor is close to 1.5 (Katz and Murphy, 1992; Ciccone and Peri, 2005; Johnson and Keane, 2013). Our estimates are also in line with these studies, but we again show that there are important difference in the degree of substitutability between skill groups across occupations.

Our results suggest that changes in the gender composition of labor supply are having a strong impact on the gender earnings gap. But the magnitude of the increase in FLFP imply that the demand for female labor must have increased considerably during the period, otherwise we should have observed a much larger rise in the gap at all points of the pay distribution.

Demand trends. Figure 10 shows the model predictions of the evolution of the share parameters, which we use to capture changes in labor demand.[27] Panels (a) and (b) show the evolution of demand for male relative to female labor among workers with at most high school and those with college education respectively.

[27]See Equations (4.4) and (4.5) for an example of how the relative demands trends are specified.

Each panel has three series corresponding to the three occupation groups. Panel (c) shows the evolution of demand for college educated workers relative to those with at most high school education in the three occupations. Panel (d) and (e) show the evolution of the relative demand trends between occupations and total factor productivity respectively. To facilitate interpretation, each series capturing changes in relative demands is normalized to take a value of zero in 1989.

The demand for female relative to male labor increased substantially over the last quarter century. This result holds within the three occupation groups, and also among workers with low and high levels of schooling, but the effect is particularly strong at the top of the pay distribution (e.g. among college educated workers in abstract task-intensive occupations). In the case of workers with at most high school education, changes in the demand side of the economy imply that – holding everything else constant – the log (male/female) earnings ratio should have fallen by 14 log points in abstract and routine task-intensive occupations, and by 25 log points in manual task-intensive occupations. Among workers with at least some years of college education completed, these effects are even stronger: in abstract task-intensive occupations, the model predicts that log (male/female) earnings ratio should have fallen by 58 log points.

The coefficients of the relative demand polynomials are estimated residually, so we are not able to pinpoint exactly what are the main drivers behind the observed patterns. There is agreement in most of the literature that the structural change faced by most economies in the last decades has been favourable to female labor. Part of this literature has emphasized how labor reallocated from goods to service industries (Lup Tick and Oaxaca, 2010; Akbulut, 2011; Olivetti and Petrongolo, 2014; Ngai and Petrongolo, 2017). Other studies have argued that the skill requirements in the economy are changing, and there is less need for physical work (Galor and Weil, 1996; Blau and Kahn, 1997; Weinberg, 2000; Rendall, 2010; Black and Spitz-Oener, 2010; Aguayo-Tellez et al., 2013; Rendall, 2013; Juhn et al., 2014). In both cases, female labor benefits because jobs where they presumably have a comparative advantage are gaining ground in the economy.

Relative demand trends are also favouring workers with more schooling and labor in abstract task-intensive occupations (see Panels (c) and (d) of Figure 10).[28] Some potential drivers behind this skill specific demand shift in Mexico have been associated to the trade and investment liberalization of the Mexican economy since the late 1980s (Feenstra and Hanson, 1997; Hanson, 2003; Sánchez-Páramo and Schady, 2003; Behrman et al., 2007; Caselli, 2012); and to the growth of foreign direct investment (Feenstra and Hanson, 1997). The results provide evidence that skill and routine-biased technical change – the most widely accepted explanation

[28]Other studies showing that relative demand for skilled labor has been growing in Mexico include Sánchez-Páramo and Schady (2003); Binelli and Attanasio (2010); Manacorda et al. (2010); Gasparini and Lustig (2011)

for the rise in income inequality in developed economies – is also a factor in this developing economy, but that the educational upgrading of the workforce is strong enough to counteract this force, so we observe a decline of the college premium.

5.2.2 Supply Side

The estimates of the parameters from the supply side of the model are shown in Table 8. Each row reports the point estimate and standard error of a different parameter. In the more relevant cases, the table also reports the average marginal effects. Average marginal effects of the fertility and marital status variables are calculated by taking the numerical derivative of the probability of choosing home production with respect to the given variable. We calculate this derivative in each year separately and then take the average across all years. In the case of the pecuniary rewards (ψ_1), the reported average marginal effect is calculated by finding the numerical derivative of the probability that a labor type chooses a given market occupation with repect to the wage. We calculate this derivative in each year and for every possible labor type and occupation combination separately, and then take the average across all the values.

An increase in the wage in a given market occupation raises the probability that an agent will choose that occupation, but the wage elasticity of labor supply is relatively low. For example, take the case of a college-educated female worker choosing among the three market occupations or home production. If the hourly wage for her group in the abstract task-intensive occupation increased from 6.5 to 7.15 – a ten percent increase of the sample average of that group –, the probability of her choosing that occupation is predicted to increases by 1.43 percentage points. Since the individual probabilities are in one-to-one correspondence with the occupation shares, the share of college-educated female workers in abstract task-intensive occupation would also be predicted to increase by 1.43 percentage points. This result suggest that, conditional on education and sex, wage dynamics are not a major driver in the choice of occupation or FLFP.

The model predicts that having children in pre-school ages and being married or with a permanent partner are important determinants of the decision to participate in market occupations, but only among workers with lower levels of schooling. For example, the percentage of prime-age females with at most a high school degree and with children under the age of 5 fell by 18.9 percentage points from 1989 to 2014 (see Panel (a) of Figure 8). This decline is associated with a rise in labor force participation of 3.3 percentage points, which is close to 17 percent of the total observed change. We don't find that having a children under the age of five has a statistically significant effect on the decision to enter the labor market for college educated women.

The percentage of females with at most high school education that are either married or with a permanent partner fell by 6 percentage points (see Panel (b) of Figure 8), and this decline is associated with a rise of labor force participation of 1 percentage point, close to 5 percent of the observed change. Again, we don't find that being married has a statistically significant impact on the decision to enter the labor market for college educated women.

Together, the fertility transition and changes in marital status explain close to 22 percent of the overall rise in female labor force participation among women with at most high school education, the group for which we observe the largest movement into the workforce. The other 78 percent is explained by changes in preferences for home production and/or advances in home production technology, which we capture with the linear trend in Equation (4.7). This results suggest that changes in cultural and social norms regarding female work, as have also been stressed by recent studies (Fogli and Veldkamp, 2011; Fernández, 2013), are a major driver behind the increase in FLFP.

Finally, for men, both marital status and fertility influence the decision to participate in the labor market, but only among those with at most a high school degree. Interestingly, the direction of the effects is the opposite to that of females. Having children under the age of five and being married or with a permanent partner is associated with higher labor force participation, although the average marginal effects are small. Neither of the two variables is statistically significant in the case of college educated men.

5.3 Counterfactual Exercises

Using the parameter estimates of the model, we can can quantify the impact of relative supply and demand changes on the occupation and wage structure. We run four different counterfactual excercises. In each excercise, we 'turn off' a channel that directly affects either the participation and occupation choice, or the marginal product of labor. We then run the model under these alternative scenarios and compute the counterfactual equilibrium wages and occupation shares.

In the first scenario (CF1), we set the the linear, quadratic, and cubic coefficients of all the demand trends (α shares) to zero. This exercise allows us to quantify the estimated impact of demand side forces on the occupation and wage structure. In the second scenario (CF2), the share of individuals with children under the age of 5 in the four groups is set to be equal to the level in 1989 and constant across the years. This exercise allows us to quantify the estimated impact of changes in fertility decisions. In the third scenario (CF3), the share of individuals that are married or with a permanent partner in each group is set to be equal to the level in 1989 and constant across the years. This exercise allows us to quantify the estimated impact of changes in the marriage market. In the last scenario (CF4), the linear trends ($\pi_{2,f}, \pi_{2,k}$) in the home production utility function are set to be equal to zero. This exercise allows us to quantify the estimated impact of changes in preferences and home production technology.

Table 9 presents the results from these exercises. Each cell reports the difference between males and females of the difference between C.1990 and C.2013 of the occupation shares and log mean wages for all groups. The first two columns correspond to what is observed in the data and the prediction of the baseline model. Columns 4-6 report the results from the different counterfactual scenarios.

The results from the first counterfactual scenario (CF1) reinforce our previous discussion about the role of changes in the demand side of the economy: relative demand trends in Mexico were strongly favourable to female labor. If we hold the demand constant, the model predicts that the gender earnings gap would be significantly higher in all occupations and skill groups (see column CF1 of Table 9). For example, the gender earnings gap among college educated workers in abstract task-intensive occupation is predicted to increase by 29.3 log points in this scenario. At the bottom of the pay distribution, the gender earnings gap is predicted to increase by as much as 30.9 log points. These number support the idea that demand side trends help mitigate the impact of the gender compositional change in labor supply. Interestingly, the differences in occupational choices between men and women, especially participation rates, were mostly unaffected by how the demand evolved, which is due to the fact that the wage elasticity of labor supply is low.

The second (CF2) and third counterfactuals (CF3) capture the equilibrium effects of the changes in fertility and marital status of the Mexican population. The results show that these changes did not have a significant impact on the occupation or wage structure. Even tough we found that both variables had some explanatory power for the decision to enter the workforce among unskilled workers, the magnitude of the effect is not strong enough to generate a change in relative earnings or participation rates.

The story is quite different when we look at the final counterfactual (CF4). The linear trends in the home production utility have a strong impact on the occupation and wage structure. Once we switch off these trends, the rise in FLFP is much lower: 6 percent predicted instead of the 22 percent observed. This in turn implies that the gender compositional shift in labor supply is significantly smaller, and that downward pressures on female wages are attenuated. In fact, in this scenario there is strong convergence in (male/female) relative earnings in all occupations and skill groups. For example, in the manual task-intensive occupations, this counterfactual scenario predicts that the gender earnings gap among low-skilled workers would have declined, not increased: the log (male/female) earnings ratio falls by -18 log points, while the baseline predicts a 8.8 log points increase (see columns Model and CF4 of Table 9). In abstract task intensive occupations, the gender earnings gap among college educated workers would have decline more (-20.7 log points) than the baseline prediction (-13.5 log points). The main takeaway is that the rise in FLFP did have a quantitatively important effect on the gender wage gap across the distribution.

The upper block of Table 10 reports the results of the same counterfactuals but aggregating across the three occupations. We aggregate using a weighted average of the groups, where the weights are equal to the respective occupation shares. It is clear from counterfactuals CF1 and CF4 that demand and supply trends had strong impacts on the the gender earnings gap, but in opposite directions: Among higher-educated workers, demand side forces outpaced the impact of changes in relative supply, so there was convergence in earnings between men and women. Among lower-educated workers, supply side forces outpaced the impact of changes in relative demand, so there was divergence in earnings between men and women.

The lower block of Table 10 reports the results of the four counterfactual scenarios but looking at the effect on the college/high school premium by sex. We concentrate particularly on the impact of rising FLFP (column CF4) on the college premium, since this is the channel that most of the literature on wage inequality has ignored. We find that rising FLFP has also imposed downward pressure on the wages of college educated male workers. The baseline model predicts that the college gap declined by -20.4 log points, but once we limit the rise of FLFP, the fall of the gap halves to -9.8 log points. The implication is that the movement of women into the workforce is also a driver in the contraction of the male wage distribution. This is a direct consequence of the finding that the elasticity of substitution between male and female labor is higher in high-paying occupations. This results also suggest that the rise of FLFP is potentially an very important driver behind the documented fall in income inequality in most countries in Latin America during the past quarter century.

6 Robustness Checks

The earnings series used for the baseline estimates of the parameters of the model included only incomes from full-time workers – those that reported working 35 hours or more in the previous week. This was done to have groups that were more closely comparable, since the share of part-time workers differs markedly between males and females: part-time female workers represent between 33 and 38 percent of all female workers, while the share of part-time male workers ranged between 10 and 13 percent. Table 11 reports the point estimates and standard errors of the different elasticities of substitution in the production technology using alternative earnings series and measures of labor supply. The table reports the baseline estimates; the estimates once we include income from part-time workers; and the estimates if we measure labor supply by the total number of hours worked of each group instead of the head-count. Since we don't have a measure of hours worked for people that are in home production, we have to impute those values. We assign each person in home production the average number of hours worked by workers in market occupation with the same level of schooling, sex, and age.

There are no major difference between the baseline estimates of the elasticities of substitution and the ones obtained using the alternative earnings and labor supply series. The rank-size of the values in each of the levels is maintained. Including income from par-time workers leads to a lower estimate of the elasticity of substitution between male and female labor in manual and routine task-intensive occupations, which are now 0.8 and 0.97 respectively. This result only reinforces our conclusions from the previous section about the strong impact that female labor supply is having on the wages of female workers in low-paying occupations. Using a labor supply measure that captures changes in the intensive margin doesn't change our analysis in any meaningful way.

The corresponding estimates of the parameters from the supply side of the model under the alternative earnings and supply measures are shown in Table 12. Results are essentially unchanged compared to the baseline.

When modelling the structure of the production technology, two decisions were made that could influence the results but are not grounded on a solid theoretical basis: first, in the three nests, labor is first divided by education and then by gender. This division changes the number of relative demands that are estimated in each dimension, but it should not alter the main results in a significant way. Second, the model assumes that the elasticity of substitution between abstract and routine task-intensive occupations is the same to that between abstract and manual-task intensive occupations. The way the model is set-up implies that at least one occupational group will have a common elasticity with the other two; we choose the abstract task-intensive occupation since it leads to a natural division between high and low-paying jobs, but it did not have to be that way.

Table 13 presents point estimates and standard errors of the elasticities of substitution in the production technology under four different model specifications: (i) the baseline estimates; (ii) the estimates if we switch the order of second and third nests of the production technology; (iii) the estimates if the occupational group that has the common elasticity with the other two is the routine task-intensive; and (iv) the estimates if the occupational group that has the common elasticity with the other two is the manual task-intensive. Once again, we find that the rank-size of the values of the elasticities of substitution between male and female labor is maintained in all cases: in manual and routine task-intensive occupations these elasticities are between 0.7 and 1.2, while in abstract task-intensive occupations the number is between 1.9 and 2.6.

The corresponding estimates of the parameters from the supply side of the model under the alternative model specifications are shown in Table 14. Results are essentially unchanged compared to the baseline.

7 Conclusions

This paper estimates impacts of the rapid rise in female labor force participation in Mexico on the occupational and wage structure over the last 25 years. It argues that the increase in female labor supply in Mexico, one of the largest in the world during this period, has fundamentally changed the patterns of relative earnings between men and women. It shows that tracking changes in the mean (or median) gender earnings gap veils important - and divergent - changes across the earnings distribution. In particular, the gender earnings gap increased in manual and routine task-intensive occupations. However, it declined in the abstract task-intensive occupations at the higher end of the earnings distribution. Using recently developed techniques for decomposition of relative wage changes across quantiles, we provide evidence that the observed changes in relative earnings were not the result of changes in the skill or experience profile of the workforce but, rather, that the marginal return to female relative to male labor was changing.

We develop an equilibrium model of the labor market to investigate whether rising female labor supply can explain the evolution of the wage structure. In a departure from previous work on gender earnings gaps, the model follows the task-based approach, allowing the elasticity of substitution between male and female labor to vary depending on the task content of occupations. This feature allows female labor supply to have heterogeneous impacts on relative earnings throughout the pay distribution.

Structural estimates of the model parameters show that the elasticity of substitution between male and female labor is lower in manual and routine task-intensive occupations (1.2) than in abstract task-intensive occupations (2.6), and that demand trends have favoured women across the board but especially so for college-educated women in abstract-intensive occupations. Our analysis suggests that demand trends have attenuated the supply-driven downward pressure on female wages in low-paying occupations, while fully counteracting it in high-paying occupations.

References

Acemoglu, D. and Autor, D. (2011). *Skills, Tasks and Technologies: Implications for Employment and Earnings*. Elsevier, Handbook of Labor Economics *4*, chap. 12. 1043–1171.

Acemoglu, D., Autor, D. H. and Lyle, D. (2004). Women, War, and Wages: The Effect of Female Labor Supply on the Wage Structure at Midcentury. *Journal of Political Economy* 112: 497–551.

Adda, J., Dustmann, C. and Stevens, K. (2017). The Career Costs of Children. *Journal of Political Economy* 125: 293–337.

Aguayo-Tellez, E., Airola, J., Juhn, C. and Villegas-Sanchez, C. (2013). Did Trade Liberalization Help Women? The Case of Mexico in the 1990s. In Polachek, S. and Tatsiramos, K. (eds), *New Analysis of Worker Well-Being*. Research in Labor Economics. Bingley, UK: Emerald Group Publishing Limited, *vol. 38*.

Akbulut, R. (2011). Sectoral Changes and the Increase in Wome's Labor Force Participation. *Macroeconomic Dynamics* 15: pp. 240–264.

Altonji, J. G., Kahn, L. B. and Speer, J. D. (2014). Trends in Earnings Differentials across College Majors and the Changing Task Composition of Jobs. *American Economic Review* 104: 387–393.

Autor, D. and Dorn, D. (2013). The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market. *American Economic Review* 103(5): pp. 1553–1597.

Autor, D., Levy, F. and Murnane, R. (2003). The Skill Content of Recent Technological Change: An Empirical Exploration. *The Quarterly Journal of Economics* 118(4): pp. 1279–1333.

Autor, D. H., Dorn, D. and Hanson, G. H. (2016). The China Shock: Learning from Labor-Market Adjustment to Large Changes in Trade. *Annual Review of Economics* 8: 205–240.

Autor, D. H., Katz, L. F. and Kearney, M. S. (2006). The Polarization of the U.S. Labor Market. *American Economic Review* 96: 189–194, doi:10.1257/000282806777212620.

Autor, D. H., Katz, L. F. and Kearney, M. S. (2008). Trends in U.S. Wage Inequality: Revising the Revisionists. *The Review of Economics and Statistics* 90: 300–323.

Bailey, M. J. and DiPrete, T. A. (2016). Five decades of remarkable but slowing change in u.s. women's economic and social status and political participation. *RSF* 2: 1–32, doi:10.7758/RSF.2016.2.4.01.

Baker, M. and Cornelson, K. (2016). Gender Based Occupational Segregation and Sex Differences in Sensory, Motor and Spatial Aptitudes. NBER Working Papers 22248, National Bureau of Economic Research, Inc.

Barth, E., Kerr, S. P. and Olivetti, C. (2017). The Dynamics of Gender Earnings Differentials: Evidence from Establishment Data. Working Paper 23381, National Bureau of Economic Research, doi:10.3386/w23381.

Behrman, J., Birdsall, N. and Szekely, M. (2007). Economic Policy Changes and Wage Inequality in Latin America. *Economic Development and Cultural Change* 56(1): pp. 57–97.

Binelli, C. and Attanasio, O. (2010). Mexico in the 1990s: the Main Cross-Sectional Facts. *Review of Economic Dynamics* 13: 238–264.

Black, S. and Juhn, C. (2000). The Rise of Female Professionals: Are Women Responding to Skill Demand? *The American Economic Review* 90(2): pp. 450–455.

Black, S. and Spitz-Oener, A. (2010). Explaining Women's Success: Technological Change and the Skill Content of Women's Work. *The Review of Economics and Statistics* 92(1): pp. 187–194.

Blau, F. D. and Kahn, L. M. (1997). Swimming Upstream: Trends in the Gender Wage Differential in 1980s. *Journal of Labor Economics* 15: 1–42.

Blau, F. D. and Kahn, L. M. (2017). The Gender Wage Gap: Extent, Trends, and Explanations. *Journal of Economic Literature* 55: 789–865.

Bound, J. and Johnson, G. (1992). Changes in the Structure of Wages in the 1980's: An Evaluation of Alternative Explanations. *American Economic Review* 82: 371–392.

Brown, R., Moon, M. and Zoloth, B. (1980). Occupational Attainment and Segregation by Sex. *Industrial and Labor Relations Review* 33(4): pp. 506–517.

Card, D. and Lemieux, T. (2001). Can Falling Supply Explain the Rising Return to College for Younger Men? A Cohort-Based Analysis. *The Quarterly Journal of Economics* 116(2): pp. 705–746.

Caselli, M. (2012). Trade, Skill-Biased Technical Change and Wages in Mexican Manufacturing. Working Paper.

Ciccone, A. and Peri, G. (2005). Long-Run Substitutability Between More and Less Educated Workers: Evidence from U.S. States, 1950-1990. *The Review of Economics and Statistics* 87: 652–663.

Coen-Pirani, D., León, A. and Lugauer, S. (2010). The Effect of Household Appliances on Female Labor Force Participation: Evidence from Microdata. *Labour Economics* 17: 503–513.

Costa, D. L. (2000). From Mill Town to Board Room: The Rise of Women's Paid Labor. *Journal of Economic Perspectives* 14: 101–122.

Cowell, F. A. and Victoria-Feser, M.-P. (1996). Robustness Properties of Inequality Measures. *Econometrica* 64: 77–101.

Cruces, G. and Galiani, S. (2007). Fertility and Female Labor Supply in Latin America: New Causal Evidence. *Labour Economics* 14: 565–573.

Deming, D. J. (2017). The growing importance of social skills in the labor market*. *The Quarterly Journal of Economics* 132: 1593–1640, doi:10.1093/qje/qjx022.

Dorn, D. (2009). Essays on Inequality, Spatial Interaction, and the Demand for Skills. Ph.D. thesis, University of St. Gallen, Graduate School of Business Administration, Economics, Law and Social Sciences (HSG).

Essama-Nssah, B. and Lambert, P. (2011). Influence Functions for Distributional Statistics. ECINE Working Paper 236.

Feenstra, R. C. and Hanson, G. H. (1997). Foreign Direct Investment and Relative Wages: Evidence from Mexico's Maquiladoras. *Journal of International Economics* 42: 371–393.

Fernandez, M. and Messina, J. (2017). Skill Premium, Labor Supply and Changes in the Structure of Wages in Latin America. IZA Discussion Papers 10718, Institute for the Study of Labor (IZA).

Fernández, M. and Messina, J. (2018). Skill Premium, Labor Supply, and Changes in the Structure of Wages in Latin America. *Journal of Development Economics* 135: 555 – 573.

Fernández, R. (2013). Cultural Change as Learning: The Evolution of Female Labor Force Participation over a Century. *American Economic Review* 103: 472–500.

Fernández, R., i, A. and Olivetti, C. (2004). Mothers and Sons: Preference Formation and Female Labor Force Dynamics. *The Quarterly Journal of Economics* 119: 1249–1299.

Fernández, R. and Wong, J. C. (2014). Divorce Risk, Wages and Working Wives: A Quantitative Life Cycle Analysis of Female Labour Force Participation. *Economic Journal* 0: 319–358.

Firpo, S., Fortin, N. and Lemieux, T. (2007). Decomposing Wage Distributions using Recentered Influence Functions. Unpublished manuscript, PUC-Rio and UBC.

Firpo, S., Fortin, N. and Lemieux, T. (2009). Unconditional Quantile Regressions. *Econometrica* 77(3): pp. 953–973.

Fogli, A. and Veldkamp, L. (2011). Nature or Nurture? Learning and the Geography of Female Labor Force Participation. *Econometrica* 79: 1103–1138.

Galiani, S., Cruces, G., Acosta, P. and Gasparini, L. C. (2017). Educational Upgrading and Returns to Skills in Latin America: Evidence from a Supply-Demand Framework. Working Paper 24015, National Bureau of Economic Research, doi: 10.3386/w24015.

Galor, O. and Weil, D. (1996). The Gender Gap, Fertility, and Growth. *The American Economic Review* 86(3): pp. 374–387.

Gasparini, L. and Lustig, N. (2011). The Rise and Fall of Income Inequality in Latin America. In Ocampo, J. A. and Ros, J. (eds), *The Oxford Handbook of Latin American Economics*. Oxford University Press.

Goldin, C. (1984). The Historical Evolution of Female Earnings Functions and Occupations. *Explorations in Economic History* 21: 1–27.

Goldin, C. (1986). Monitoring Costs and Occupational Segregation by Sex: A Historical Analysis. *Journal of Labor Economics* 4: 1–27.

Goldin, C. (2006). The Quiet Revolution That Transformed Women's Employment, Education, and Family. *American Economic Review* 96: 1–21.

Goldin, C. and Olivetti, C. (2013). Shocking Labor Supply: A Reassessment of the Role of World War II on Women's Labor Supply. *American Economic Review* 103: 257–62.

Goos, M. and Manning, A. (2007). Lousy and Lovely Jobs: The Rising Polarization of Work in Britain. *The Review of Economics and Statistics* 89(1): pp. 118–133.

Goos, M., Manning, A. and Salomons, A. (2014). Explaining Job Polarization: Routine-Biased Technological Change and Offshoring. *American Economic Review* 104: 2509–2526.

Greenwood, J., Guner, N., Kocharkov, G. and Santos, C. (2016). Technology and the Changing Family: A Unified Model of Marriage, Divorce, Educational Attainment, and Married Female Labor-Force Participation. *American Economic Journal: Macroeconomics* 8: 1–41.

Greenwood, J., Seshadri, A. and Yorukoglu, M. (2005). Engines of Liberation. *The Review of Economic Studies* 72: pp. 109–133.

Grossbard-Shechtman, S. A. and Neuman, S. (1988). Women's Labor Supply and Marital Choice. *Journal of Political Economy* 96: 1294–1302.

Hanson, G. (2003). What has Happened to Wage in Mexico since NAFTA? Implications of Hemispheric Free Trade. National Bureau of Economic Research. NBER Working Paper 9563.

Heckman, J. J., Lochner, L. and Taber, C. (1998a). Explaining rising wage inequality: Explorations with a dynamic general equilibrium model of labor earnings with heterogeneous agents. *Review of Economic Dynamics* 1: 1 – 58, doi: https://doi.org/10.1006/redy.1997.0008.

Heckman, J. J., Lochner, L. and Taber, C. (1998b). General-equilibrium treatment effects: A study of tuition policy. *The American Economic Review* 88: 381–386.

Jaimovich, N., Siu, H. and Cortes, G. M. (2017). The End of Men and Rise of Women in the High-Skilled Labor Market. 2017 Meeting Papers 809, Society for Economic Dynamics.

Johnson, M. and Keane, M. (2013). A Dynamic Equilibrium Model of the US Wage Structure, 1969-1996. *Journal of Labour Economics* 31(1): pp. 1–49.

Juhn, C. and Kim, D. I. (1999). The Effects of Rising Female Labor Supply on Male Wages. *Journal of Labor Economics* 17: 23–48.

Juhn, C., Ujhelyi, G. and Villegas-Sanchez, C. (2014). Men, Women, and Machines: How Trade Impacts Gender Inequality. *Journal of Development Economics* 106: 179–193, doi:10.1016/j.jdeveco.2013.09.

Katz, L. and Autor, D. (1999). Changes in the Wage Structure and Earnings Inequality. In Ashenfelter, O. and Card, D. (eds), *Handbook of Labor Economics*. Amsterdam: North-Holland, *vol. 3A*, chap. 26.

Katz, L. and Murphy, K. (1992). Changes in Relative Wages, 1963-1987: Supply and Demand Factors. *The Quarterly Journal of Economics* 107(1): pp. 35–78.

Katz, L. F. and Goldin, C. (2000). Career and Marriage in the Age of the Pill. *American Economic Review* 90: 461–465.

Killingsworth, M. R. and Heckman, J. J. (1987). Female Labor Supply: A Survey. In Ashenfelter, O. and Layard, R. (eds), *Handbook of Labor Economics*. Elsevier, Handbook of Labor Economics *1*, chap. 2, 103–204.

Krusell, P., Ohanian, L., Ríos-Rull, J. and Violante, G. (2000). Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis. *Econometrica* 68(5): pp. 1029–1053.

Lee, D. (2005). An estimable dynamic general equilibrium model of work, schooling, and occupational choice. *International Economic Review* 46: 1–34.

Lee, D. and Wolpin, K. (2006). Intersectoral Mobility and the Growth of the Service Sector. *Econometrica* 74(1): pp. 1–46.

Lee, D. and Wolpin, K. I. (2010). Accounting for Wage and Employment Changes in the US from 1968-2000: A Dynamic Model of Labor Market Equilibrium. *Journal of Econometrics* 156: 68 – 85, doi:https://doi.org/10.1016/j.jeconom.2009.09.008, structural Models of Optimization Behavior in Labor, Aging, and Health.

Lemieux, T. (2006). Increasing Residual Wage Inequality: Composition Effects, Noisy Data, or Rising Demand for Skill? *The American Economic Review* 96(3): pp. 461–498.

Levy, S. and Schady, N. (2013). Latin America's Social Policy Challenge: Education, Social Insurance, Redistribution. *Journal of Economic Perspectives* 27: 193–218, doi:10.1257/jep.27.2.193.

López-Calva, L. F. and Lustig, N. (2010). Explaining the Decline in Inequality in Latin America: Technological change, Educational Upgrading, and Democracy. In López-Calva, L. F. and Lustig, N. (eds), *Declining Inequality in Latin America: A Decade of Progress?*. Brookings Institution Press, chap. 1.

Lup Tick, S. and Oaxaca, R. (2010). Technological Change and Gender Wage Gaps in the US Service Industry. *Annals of Economics and Statistics* 99/100: pp. 47–65.

Lustig, N., Lopez-Calva, L. and Ortiz-Juarez, E. (2013). Deconstructing the Decline in Inequality in Latin America. Tulane Economic Working Paper Series.

Manacorda, M., Sánchez-Paramo, C. and Schady, N. (2010). Changes in Returns to Education in Latin America: The Role of Demand and Supply of Skills. *Industrial and Labor Relations Review* 63(2): pp. 307–326.

Michaels, G., Natraj, A. and Van Reenen, J. (2015). Has ICT Polarized Skill Demand? Evidence from Eleven Countries over 25 Years. *The Review of Economic and Statistics* 86(1): pp. 60–77.

Murphy, K. and Welch, F. (1992). The Structure of Wages. *The Quarterly Journal of Economis* 107(1): pp. 285–326.

Ngai, L. R. and Petrongolo, B. (2017). Gender Gaps and the Rise of the Service Economy. *American Economic Journal: Macroeconomics* 9: 1–44, doi:10.1257/mac.20150253.

Oaxaca, R. and Ranson, M. (1994). On Discrimination and the Decomposition of Wage Differentials. *Journal of Econometrics* 61(1): pp. 5 – 21.

Olivetti, C. and Petrongolo, B. (2014). Gender Gaps Across Countries and Skills: Demand, Supply and the Industry Structure. *Review of Economic Dynamics* 17: 842 – 859, doi:https://doi.org/10.1016/j.red.2014.03.001.

Ñopo, H. (2012). *New Century, Old Disparities : Gender and Ethnic Earnings Gaps in Latin America and the Caribbean*. No. 11953 in World Bank Publications. The World Bank.

Pitt, M. M., Rosenzweig, M. R. and Hassan, M. N. (2012). Human Capital Investment and the Gender Division of Labor in a Brawn-Based Economy. *American Economic Review* 102: 3531–60.

Rendall, M. (2010). Brain Versus Brawn: The Realization of Women's Comparative Advantage. Institute for Empirical Research in Economics University of Zurich. Working Paper 1424-0459.

Rendall, M. (2013). Structural Change in Developing Countries: Has it Decreased Gender Inequality? *World Development* 45(1): pp. 1–16.

Ronald R. Rindfuss, A. L. K., Karin L. Brewster (1996). Women, Work, and Children: Behavioral and Attitudinal Change in the United States. *Population and Development Review* 22: 457–482.

Sánchez-Páramo, C. and Schady, N. (2003). Off and Running? Technology, Trade, and the Rising Demand for Skilled Workers in Latin America. The World Bank. Policy Research Working Paper 3015.

The World Bank (2012). *World Development Report 2012 - Gender Equality and Development*. No. 4391 in World Bank Publications. The World Bank.

Topel, R. H. (1994). Regional Labor Markets and the Determinants of Wage Inequality. *American Economic Review* 84: 17–22.

V. Cavalcanti, T. V. de and Tavares, J. (2008). Assessing the "Engines of Liberation": Home Appliances and Female Labor Force Participation. *The Review of Economics and Statistics* 90: 81–88.

Weinberg, B. A. (2000). Computer Use and the Demand for Female Workers. *Industrial and Labor Relations Review* 53: 290–308.

A Appendix

A.1 Tables and Figures

Figure 1: Labor Force Participation by Sex

Absolute Numbers

(a) ENIGH

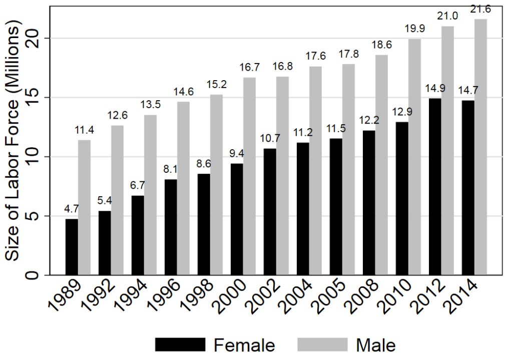

(b) CENSUS

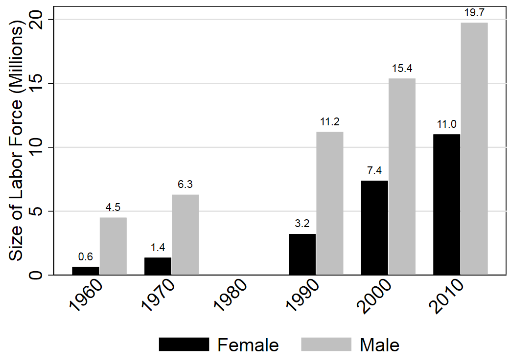

Participation Rates

(c) ENIGH

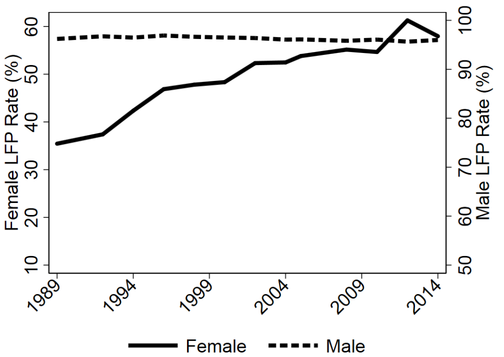

(d) CENSUS

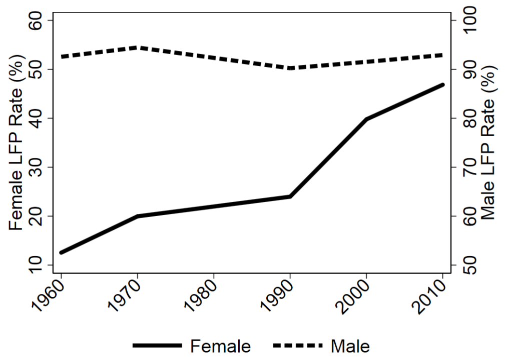

Notes: Panels (a) and (b) show the total number of prime-age population that are either working or actively searching for a job by sex. The series in Panels (c) and (d) show the share of prime-age males/females that are either working or actively searching for a job by sex. The differences between the CENSUS and the ENIGH are due to the fact that the CENSUS only includes as economically active those individuals whose primary activity was either working or looking for a job. For example, part-time workers whose primary activity was studying are categorized as outside the labor force, leading to an underestimation. Sample weights used in all calculations.

Figure 2: Changes in the Log Hourly Earnings by Sex between C.1990 and C.2013

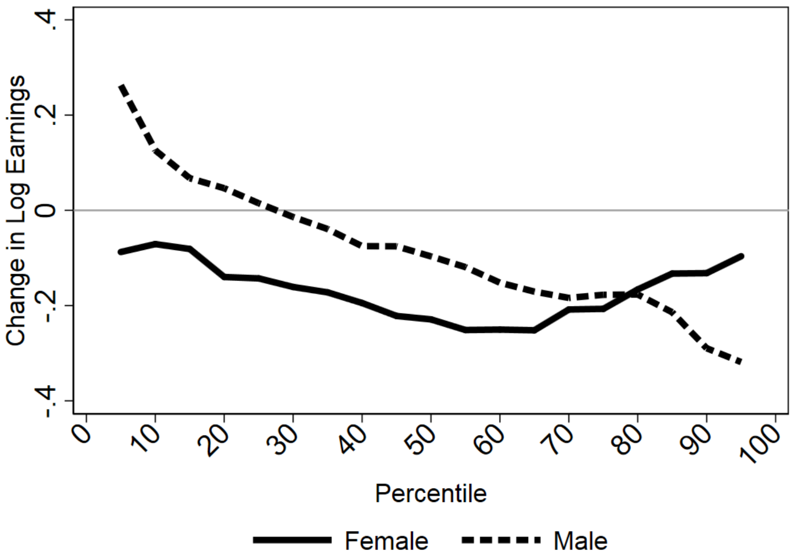

Notes: The series are constructed by computing the change in real log hourly earnings between C.1990 and C.2013 at each percentile of the distribution. We do this separetly for males and females. Sample is restricted to prime-age population that reported working for more than 35 hours a week. To increase sample size we joined together surveys from 1989 and 1992, and from 2012 and 2014.

Figure 3: Changes in the Log (Male/Female) Earnings Ratio

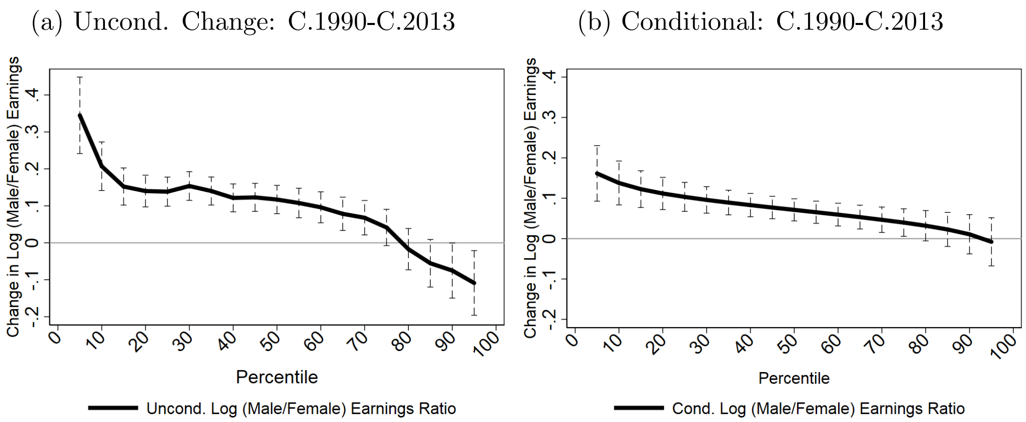

Notes: Panel (a) shows the change in log hourly earnings by percentile between C.1990 and C.2013. Panel (b) shows a similar statistic, but this time using changes in the conditional wage distributions. We do this by running conditional quantile regression at each percentile in two periods: C.1990 and C.2013. We then report the change in the coefficient of the female dummy in the regression. The controls in the regression include: dummies for 7 education categories, dummies for 6 age categories in five year intervals, and all possible interactions. Sample is restricted to prime-age population that reported working for more than 35 hours a week. To increase sample size we joined together surveys from 1989 and 1992, and from 2012 and 2014. Vertical bars correspond to 95 percent confidence intervals. Sample weights used in all calculations.

Figure 4: Changes in the Log (Male/Female) Earnings Ratio

(a) Unconditional Gap, Mean

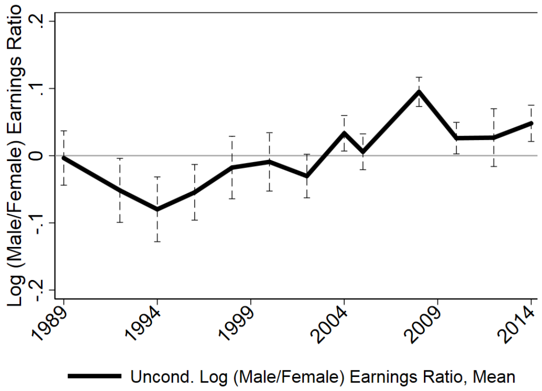

(b) Conditional Gap, Mean

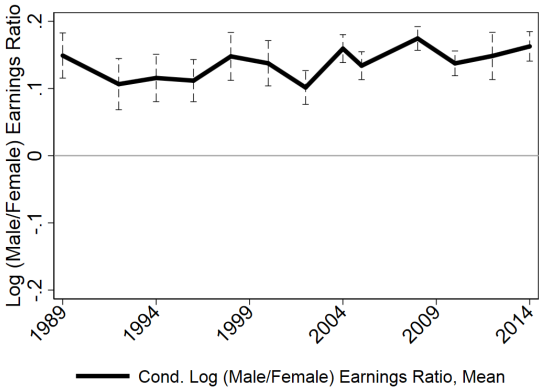

(c) Unconditional Gap, 5th Percentile

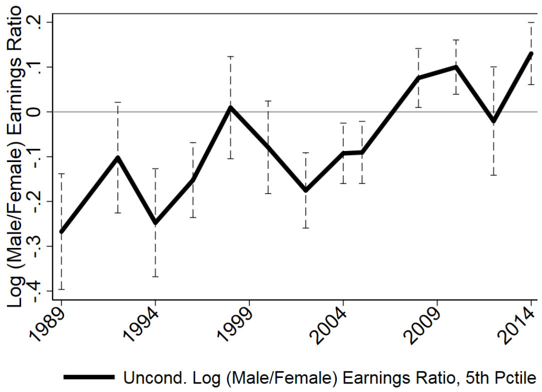

(d) Conditional Gap, 5th Percentile

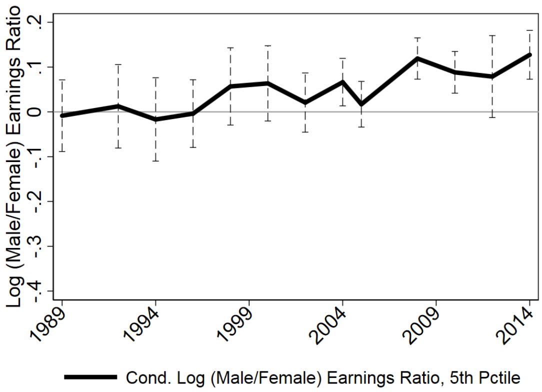

(e) Unconditional Gap, 95th Percentile

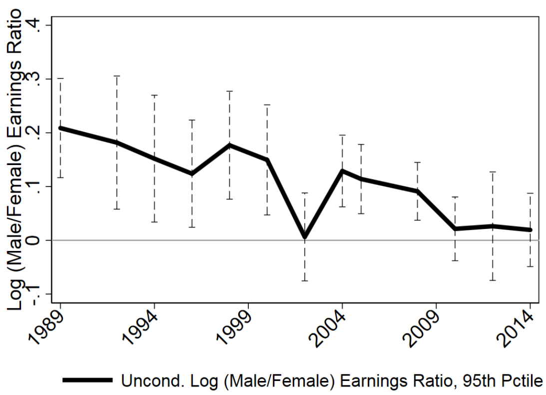

(f) Conditional Gap, 95th Percentile

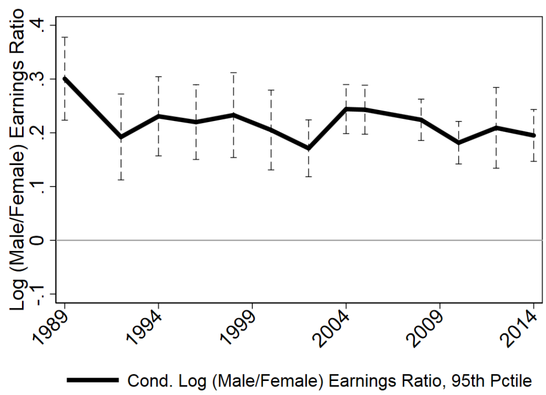

Notes: Panels (a)-(f) shows the evolution of the (male/female) gender earnings gap at different points of the distribution. The conditional gap is calculated by running OLS regressions (mean) or conditional quantile regression (percentiles) separately each year. We then report the coefficient of the female dummy associated to each year. The controls in the regression include: dummies for 7 education categories, dummies for 6 age categories in five year intervals and all possible interactions. Sample is restricted to prime-age population that reported working for more than 35 hours a week. To increase sample size we joined together surveys from 1989 and 1992, and from 2012 and 2014. Vertical bars correspond to 95 percent confidence intervals. Sample weights used in all calculations.

Figure 5: Comovement Between Log (Male/Female) Earnings Ratio and Log (Male/Female) Relative Supply by Occupation Groups

(a) Abstract Occupations

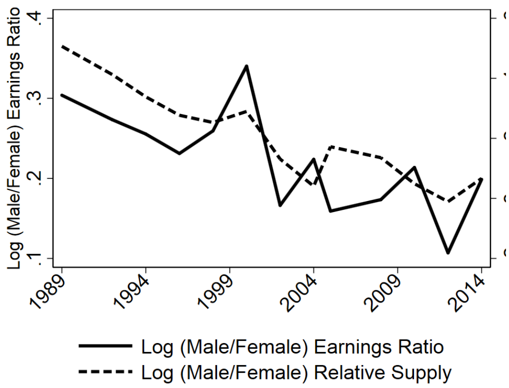

$d\ln(L_k/L_f) \approx -44$; $d\ln(W_k/W_f) \approx -10$

(b) Routine Occupations

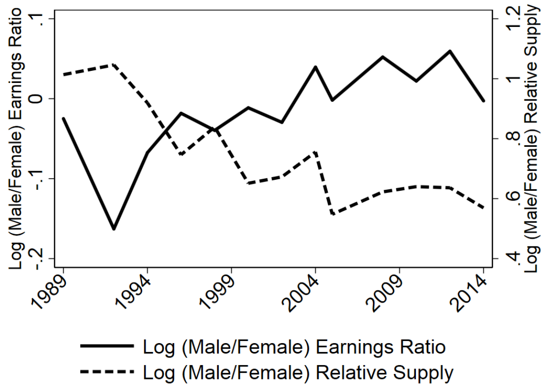

$d\ln(L_k/L_f) \approx -44$; $d\ln(W_k/W_f) \approx 2.3$

(c) Manual Occupations

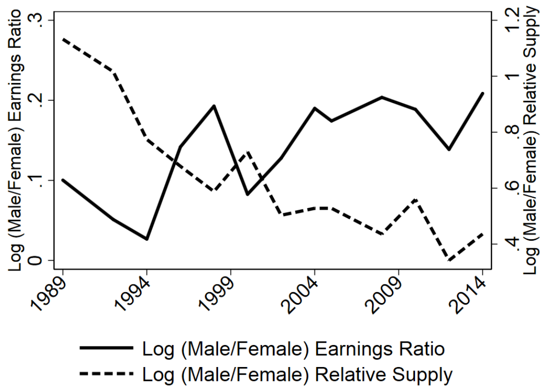

$d\ln(L_k/L_f) \approx -69$; $d\ln(W_k/W_f) \approx 10$

Notes: Panels (a)-(c) plot log (male/female) earnings ratios and log (male/female) relative supplies over the period of analysis, one panel for each occupational group. Panel (d) plots the change in log (male/female) earnings ratio and the change in log (male/female) relative supply between C.1990 and C.2013, using the principal group occupational division of the ENIGH as the unit of analysis. The regression line is estimated weighting each occupation by the average employment share between the two periods. The sample for the earnings series is restricted to full-time workers between the ages of 25 and 55.The sample for the relative supply series includes all workers in the labor force between the ages of 25 and 55. Sample weights used in all calculations.

Figure 6: Changes in the Log (Male/Female) Earnings Ratio. CENSUS

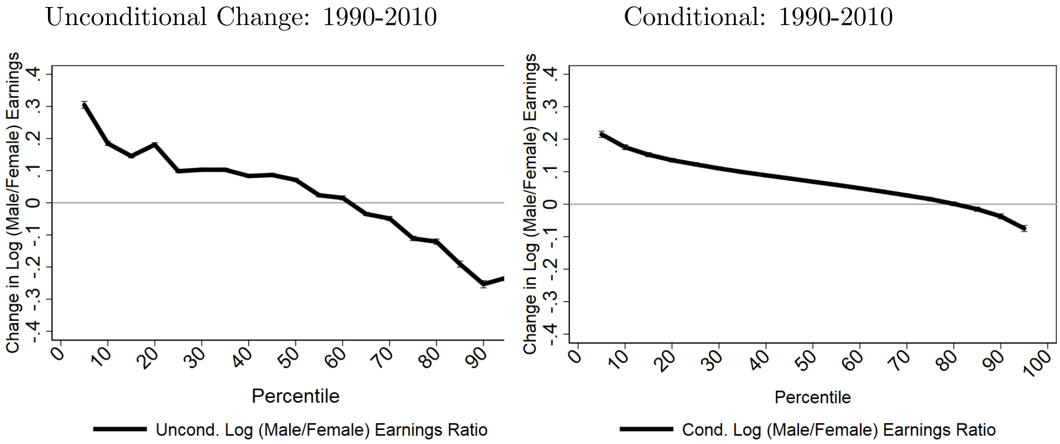

Notes: Panel (a) shows the estimated (male/female) gender earnings gap. The gender earnings gap corresponds to the estimated coefficient of a male dummy variable in a regression in which log hourly earnings is the dependent variable, and we include controls for education attainment (7 different categories), age (6 categories in five year intervals), occupation indicators (18 groups), and all possible interactions. The regressions are estimated separately in each year. Panel (b) shows the change in log hourly earnings by percentile between C.1990 and C.2013, calculated for each gender separately. To increase sample size we joined together surveys from 1989 and 1992, and from 2012 and 2014. Panels (c) shows the change in log (male/female) earnings ratio by percentile between C.1990 and C.2013. Sample is restricted to prime-age population that reported working for more than 35 hours a week. Vertical bars correspond to 95 percent confidence intervals. Sample weights used in all calculations.

Figure 7: Decomposition Results: Change in Log (Male/Female) Earnings Ratio Between C.1990 and C.2013

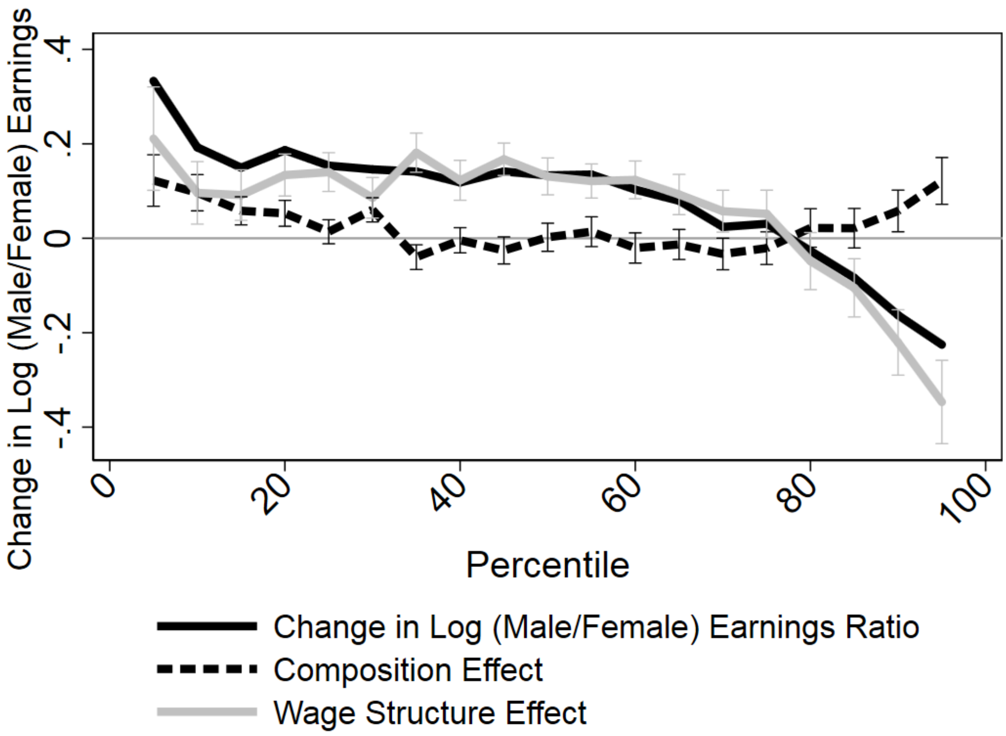

Notes: the Figure shows results of the Oaxaca-Blinder decomposition defined in Equation (3.3). The estimation is done separately for 19 percentiles. Confidence intervals are estimated via bootstrap with 500 replications. Sample weights used in all calculations.

Figure 8: Determinants of Movements into the Labor Market

(a) Changes in Fertility (b) Changes in the Marriage Market

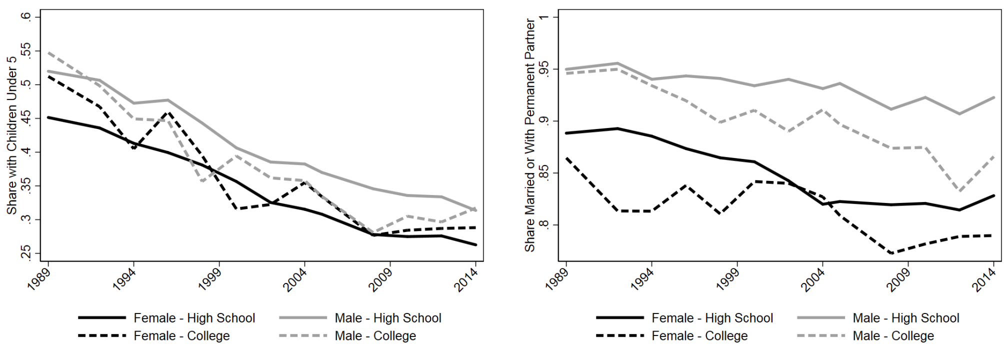

Notes: Panel (a) depicts the share of each group with children under the age of 5, and Panel (b) depicts the share of each group that is married or has a permanent partner. The ENIGH survey started asking the question on marital status to all member of the household in 1996, and for the number of children since 2004. For these reason, the two variables can only be constructed for the household head, and, if applicable, to his/her spouse or partner. The sample is restricted to prime-age population. Sample weights used in all calculations.

Figure 9: Log (Male/Female) Relative Earnings, Log (Male/Female) Relative Supplies, and Participation Rates: ENIGH and Model Predictions

Analytical Occupations

(a) Log (Male/Female) Earnings Ratio

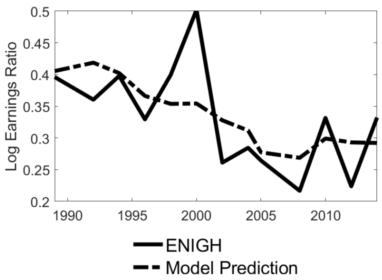

(b) Log (Male/Female) Relative Supply

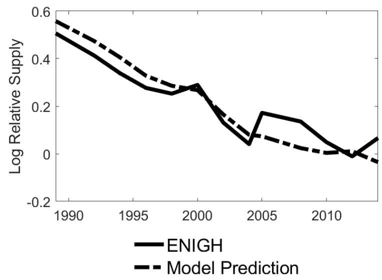

Routine Occupations

(c) Log (Male/Female) Earnings Ratio

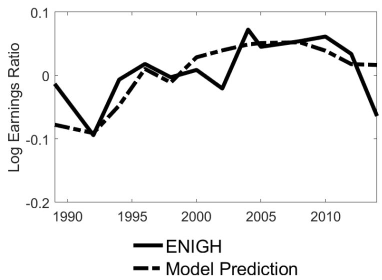

(d) Log (Male/Female) Relative Supply

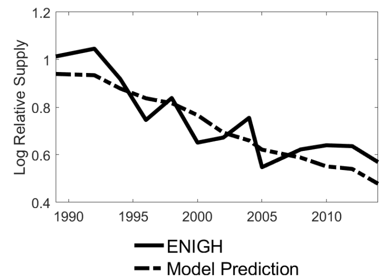

Manual Occupations

(e) Log (Male/Female) Earnings Ratio

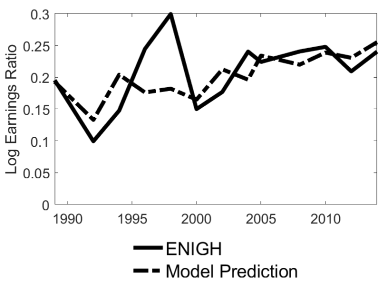

(f) Log (Male/Female) Relative Supply

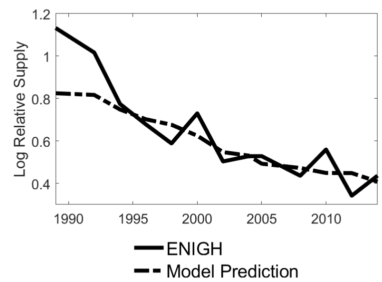

Participation Rates

(g) Female

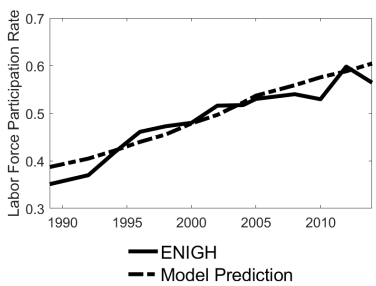

(h) Male

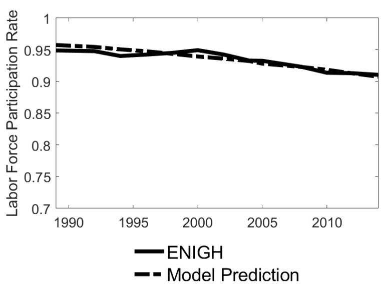

Notes: The different panels depict the series of (male/female) relative earnings, (male/female) relative supplies, and labor force participation rates, both from the ENIGH and predicted from the model.

Figure 10: Estimates of the Relative Demand Indexes and Total Factor Productivity

Production Technology: Level III

(a) Male vs. Female (High School)

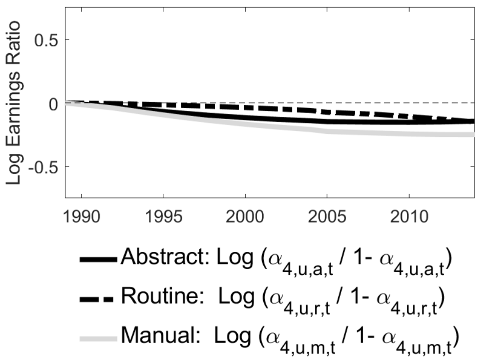

(b) Male vs. Female (College)

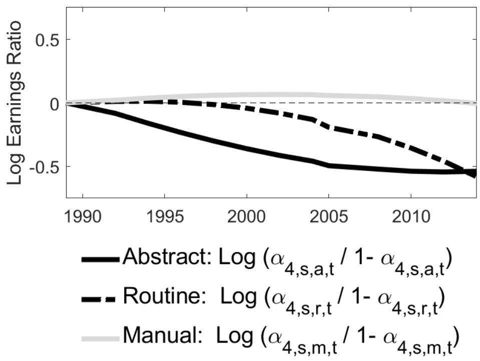

Production Technology: Level II

(c) College vs. High School

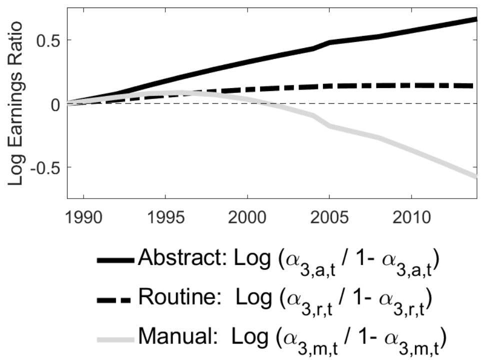

Production Technology: Level I

(d) Analytical vs. Routine and Manual; and Routine vs. Manual

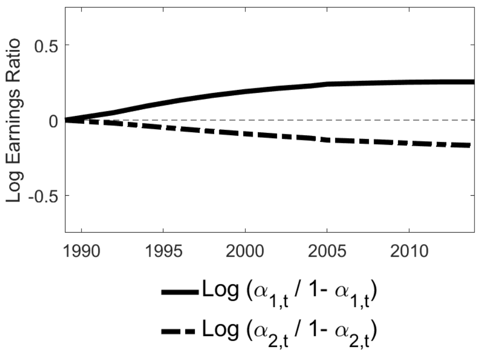

(e) Total Factor Productivity

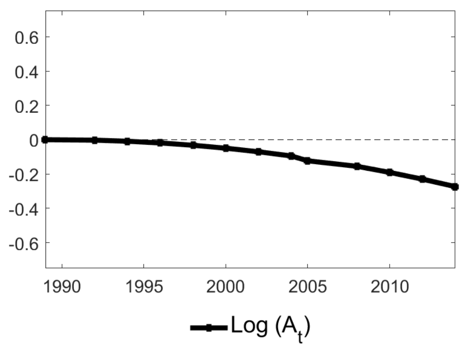

Notes: Panels (a)-(d) show the estimated change in the relative demand indexes captured by the log ratio of the α shares. Panel (e) shows the estimated change of log total factor productivity. The changes in total factor productivity and the α shares are estimated using a cubic time trend in their natural logarithm (see Equations (4.4) and (4.5) for an example). To facilitate interpretation, each series is normalized to zero in 1989.

Table 1: Labor Force Participation Rates by Selected Groups: C.1990 and C.2013

	C.1990		C.2013	
	Female Share (x100)	Male Share (x100)	Female Share (x100)	Male Share (x100)
Overall	38.59	96.49	59.59	95.82
Education				
High School	35.75	96.58	55.47	95.89
College	71.73	96.00	77.42	95.59
Age				
25-34	40.54	96.82	59.18	95.73
35-44	39.52	97.53	62.63	97.29
45-55	33.52	94.42	56.61	94.26

Notes: The table reports the labor force participation rate conditional on education, age and sex. Each cell reports the share of the respective column group. We joined together surveys from 1989 and 1992, and from 2012 and 2014 to increase sample size of the ENIGH survey. Sample weights used in all calculations.

Table 2: Changes in Composition by Education, Age, and Sex Between C.1990 and C.2013

	C.1990			C.2013			Dif. in Dif.
	Female Share (x100)	Male Share (x100)	$\Delta_{c.1990}$ (Male - Female)	Female Share (x100)	Male Share (x100)	$\Delta_{c.2013}$ (Male - Female)	$\Delta_{c.2013}$ - $\Delta_{c.1990}$
Prime-Age Population							
Participation Rate	38.59	96.49	57.89	59.59	95.82	36.24	-21.66
Education							
High School	92.09	84.27	-7.82	81.24	79.05	-2.18	5.64
College	7.91	15.73	7.82	18.76	20.95	2.18	-5.64
Age							
25-34	44.63	43.61	-1.02	35.74	36.36	0.62	1.64
35-44	32.34	32.84	0.50	34.16	33.97	-0.20	-0.69
45-55	23.02	23.55	0.52	30.09	29.67	-0.42	-0.95
Prime-Age Workforce							
Education							
High School	85.48	84.37	-1.11	76.00	79.19	3.19	4.29
College	14.52	15.63	1.11	24.00	20.81	-3.19	-4.29
Age							
25-34	46.48	43.58	-2.90	34.96	36.03	1.08	3.97
35-44	33.31	33.38	0.07	36.18	34.72	-1.46	-1.53
45-55	20.21	23.04	2.83	28.87	29.25	0.39	-2.44

Notes: The table reports the share of prime-age population and prime-age population in the workforce conditional on sex, education, and age. Each cell reports the share of the respective column group. We joined together surveys from 1989 and 1992, and from 2012 and 2014 to increase sample size of the ENIGH survey. Sample weights used in all calculations.

Table 3: Occupation Groups

ENIGH Principal Group	Median Percentile of the Task Measure: Abstract	Routine	Manual	Group	Av. Share (x100)	Av. Male Share (x100)	Av. Earnings Percentile
Managers	90.0	17.0	27.5	Abstract	2.9	71.3	85.4
Crafts and Trades (Supervisors)	84.0	42.0	62.0	Abstract	1.8	84.2	72.3
Education	83.0	11.0	65.0	Abstract	4.5	38.2	80.2
Professional	83.0	42.0	46.0	Abstract	4.1	62.4	82.3
Technical	71.0	69.0	43.0	Abstract	4.0	59.3	68.6
Arts/Entertainment	66.0	35.0	48.0	Abstract	0.6	76.4	70.4
Sales	61.0	22.5	15.0	Abstract	12.7	46.3	47.5
Crafts and Trades (Laborers)	40.0	82.0	73.0	Routine	14.3	76.4	47.4
Clerical (Supervisors)	61.0	63.0	51.5	Routine	2.5	65.0	77.9
Crafts and Trades (Helpers)	10.5	62.0	60.5	Routine	5.8	80.4	34.8
Machine Operators	16.0	62.0	51.0	Routine	3.6	62.4	48.4
Clerical (Laborers)	41.5	53.0	12.0	Routine	6.6	37.3	60.4
Transport	19.5	21.0	96.0	Manual	5.8	99.0	46.9
Agriculture	32.0	27.0	82.0	Manual	13.2	78.6	20.9
Protective Services	24.5	5.5	76.5	Manual	2.3	93.1	44.4
Domestic Service	9.0	8.0	76.0	Manual	4.1	7.6	27.0
Street Sales	38.0	13.0	64.0	Manual	3.4	44.0	30.3
Service	28.0	25.0	63.0	Manual	7.4	43.4	40.2

Notes: The three task measures were originally constructed for three-digit occupational codes of the U.S. CENSUS by Autor et al. (2003). For each measure, we first organize the three-digit occupations by percentiles, and then calculate the median percentile within the broader 18 occupational groups of the ENIGH. Each of the 18 occupations is assigned to the group in which the median percentile was highest. Further details are presented in Appendix A.2.

Table 4: Compositional Changes and The Gender Earnings Gap: Oaxaca-Blinder Decomposition Results by Selected Percentiles

	P5		P25		P50		P75		P95	
	Est.	[S.E.]	Est.	[S.E.]	Est.	[S.E.]	Est.	[S.E.]	Est.	[S.E.]
Observed Change	0.333	[0.054]	0.154	[0.021]	0.133	[0.021]	0.031	[0.027]	-0.225	[0.045]
Overall Wage Structure	0.211	[0.056]	0.140	[0.021]	0.131	[0.020]	0.052	[0.026]	-0.347	[0.045]
Overall Composition	0.122	[0.028]	0.014	[0.013]	0.002	[0.015]	-0.021	[0.018]	0.121	[0.025]

Notes: The table shows results of the Oaxaca-Blinder decomposition defined in Equation (3.3). The Standard errors in brackets are calculated via bootstrap with 500 replications. Sample weights used in all calculations.

Table 5: Levels and Changes in Real Hourly Earnings by Sex, Education, and Occupations. C.1990 and C.2013

	C.1990				C.2013				Dif. in Dif.	
	Female Earnings	Male Earnings	Log Gap Earnings	Log Gap Supplies	Female Earnings	Male Earnings	Log Gap Earnings	Log Gap Supplies	Δ Log Gap Earnings	Δ log Gap Supplies
Education										
College	7.03	9.81	33.31	86.54	5.74	7.01	20.08	19.62	-13.23	-66.92
	[0.17]	[0.16]	[2.97]	[0.00]	[0.10]	[0.13]	[2.54]	[0.00]	[3.91]	[0.00]
High School	3.11	2.95	-5.34	77.88	2.25	2.48	9.85	37.96	15.19	-39.92
	[0.04]	[0.02]	[1.44]	[0.00]	[0.02]	[0.02]	[1.39]	[0.00]	[2.09]	[0.00]
Occupation										
Abstract	5.34	7.59	35.12	41.75	4.70	5.84	21.58	2.87	-13.55	-38.88
	[0.11]	[0.12]	[2.59]	[0.00]	[0.09]	[0.11]	[2.47]	[0.00]	[3.68]	[0.00]
Rotuine	3.81	3.64	-4.61	98.93	3.00	2.89	-3.66	60.21	0.95	-38.73
	[0.07]	[0.04]	[2.10]	[0.00]	[0.05]	[0.03]	[2.16]	[0.00]	[2.94]	[0.00]
Manual	1.93	2.23	14.41	95.12	1.82	2.25	20.78	38.75	6.36	-56.37
	[0.04]	[0.03]	[2.37]	[0.00]	[0.03]	[0.02]	[1.82]	[0.00]	[3.01]	[0.00]
Educ.-Occ.										
College										
Abstract	7.35	10.37	34.40	85.46	6.30	8.15	25.86	16.15	-8.53	-69.31
	[0.21]	[0.19]	[3.36]	[0.00]	[0.13]	[0.18]	[2.96]	[0.00]	[4.50]	[0.00]
Rotuine	6.27	8.88	34.79	70.18	4.88	5.25	7.39	13.86	-27.40	-56.32
	[0.32]	[0.34]	[6.62]	[0.00]	[0.16]	[0.16]	[4.57]	[0.00]	[8.23]	[0.00]
Manual	3.75	5.27	35.14	214.06	3.15	3.31	5.15	63.92	-30.00	-150.14
	[0.65]	[0.41]	[18.96]	[0.01]	[0.23]	[0.14]	[8.31]	[0.00]	[20.50]	[0.01]
High School										
Abstract	3.95	4.66	16.64	16.00	2.67	3.12	15.72	-9.71	-0.92	-25.72
	[0.10]	[0.09]	[3.22]	[0.00]	[0.07]	[0.08]	[3.62]	[0.00]	[4.85]	[0.00]
Routine	3.42	3.10	-9.58	102.09	2.41	2.52	4.32	68.98	13.90	-33.11
	[0.06]	[0.03]	[1.90]	[0.00]	[0.04]	[0.02]	[1.86]	[0.00]	[2.63]	[0.00]
Manual	1.92	2.16	11.77	93.58	1.73	2.18	23.18	37.40	11.40	-56.19
	[0.04]	[0.03]	[2.37]	[0.00]	[0.03]	[0.02]	[1.72]	[0.00]	[2.96]	[0.00]

Notes: The table reports the average real hourly earnings, the average log (male/female) earnings gap, and the log (male/female) relative supply by education, occupation, and year. Sample is restricted to prime-age workers. The sameple for the construction of the earnings series is further restricted to include only full-time workers. We joined together surveys from 1989 and 1992, and from 2012 and 2014 to increase sample size of the ENIGH survey. The Standard errors in brackets are calculated via bootstrap with 500 replications. Sample weights used in all calculations.

Table 6: Model Fit: ENIGH and Model Predictions of the Occupation Shares and Wages

	C.1990				C.2013			
	Female ENIGH	Female Model	Male ENIGH	Male Model	Female ENIGH	Female Model	Male ENIGH	Male Model
Mean Wages								
College								
Abstract	7.35	6.45	10.37	9.51	6.30	6.14	8.15	7.92
Routine	6.27	6.80	8.88	7.68	4.88	4.80	5.25	5.46
Manual	3.75	8.00	5.27	6.39	3.15	2.27	3.31	2.45
High School								
Abstract	3.95	3.58	4.66	4.26	2.67	2.68	3.12	3.29
Routine	3.42	3.12	3.10	2.89	2.41	2.44	2.52	2.60
Manual	1.92	1.77	2.16	2.08	1.73	1.74	2.18	2.24
Occupation Shares (x100)								
College								
Abstract	2.26	1.72	5.30	4.58	5.22	5.32	6.14	6.18
Routine	0.63	0.77	1.27	1.38	1.66	1.81	1.90	1.67
Manual	0.05	0.61	0.39	1.12	0.49	0.83	0.92	1.05
Home Production	1.26	1.09	0.45	0.33	2.55	1.95	0.93	0.99
High School								
Abstract	5.45	5.36	6.39	6.77	6.27	6.44	5.69	5.43
Routine	4.89	5.64	13.58	14.63	6.69	7.00	13.34	12.98
Manual	6.90	7.34	17.59	16.47	10.35	10.08	15.04	15.68
Home Production	31.49	30.37	2.12	1.83	19.59	19.35	3.23	3.25

Notes: The Table reports average wages and occupations shares in C.1990 and C.2013, both from the ENIGH and predicted by the model.

Table 7: Parameter Estimates: Production Technology

	Elasticities of Substitution			
	Estimate	[SE]	Implied Elasticity $(1/(1-\rho))$	95% Conf. Int. $(1/(1-\rho))$
Gender				
$\rho_{4,m}$: female, male (manual)	0.175	[0.181]	1.212	[0.84, 2.12]
$\rho_{4,r}$: female, male (routine)	0.179	[0.129]	1.219	[0.93, 1.76]
$\rho_{4,a}$: female, male (abstract)	0.622	[0.099]	2.646	[1.75, 5.43]
Education				
$\rho_{3,m}$: college, high school (manual)	0.722	[0.067]	3.594	[2.44, 6.81]
$\rho_{3,r}$: college, high school (routine)	0.355	[0.041]	1.549	[1.38, 1.77]
$\rho_{3,a}$: college, high school (abstract)	0.276	[0.121]	1.382	[1.04, 2.05]
Occupation				
ρ_1: abstract, routine and manual	0.031	[0.094]	1.032	[0.87, 1.27]
ρ_2: routine, manual	-0.141	[0.183]	0.877	[0.67, 1.28]

Notes: The table reports the point estimates and standard errors of the elasticities of substitution from the production technology.

Table 8: Parameter Estimates: Occupational Choice

	Occupational Choice		
	Estimate	SE	Average Marginal Effect
Earnings			
ψ_1: earnings	0.154	[0.009]	0.022
Fertility/Children			
$\pi_{3,f,u}$: female, unskilled	0.712	[0.127]	0.174
$\pi_{3,f,s}$: female, skilled	-0.139	[0.229]	-0.025
$\pi_{3,k,u}$: male, unskilled	-0.367	[0.182]	-0.022
$\pi_{3,k,s}$: male, skilled	0.122	[0.219]	0.008
Marriage			
$\pi_{4,f,u}$: female, unskilled	0.589	[0.107]	0.144
$\pi_{4,f,s}$: female, skilled	0.267	[0.134]	0.047
$\pi_{4,k,u}$: male, unskilled	-0.524	[0.155]	-0.032
$\pi_{4,k,s}$: male, skilled	0.026	[0.196]	0.002
Non Pecuniary Rewards/Tastes			
$\pi_{1,f}$: female, home production	0.961	[0.078]	
$\pi_{2,f}$: female, home production trend	-0.059	[0.003]	
$\pi_{1,k}$: male, home production	-0.716	[0.127]	
$\pi_{2,k}$: male, home production trend	0.054	[0.005]	
$\psi_{f,u,m}$: female, unskilled, manual	0.036	[0.048]	
$\psi_{f,u,r}$: female, unskilled, routine	-0.437	[0.046]	
$\psi_{f,u,a}$: female, unskilled, abstract	-0.555	[0.044]	
$\psi_{f,s,m}$: female, skilled, manual	-0.756	[0.130]	
$\psi_{f,s,r}$: female, skilled, routine	-0.362	[0.088]	
$\psi_{f,s,a}$: female, skilled, abstract	0.508	[0.080]	
$\psi_{k,u,m}$: male, unskilled, manual	0.541	[0.063]	
$\psi_{k,u,r}$: male, unskilled, routine	0.296	[0.064]	
$\psi_{k,u,a}$: male, unskilled, abstract	-0.682	[0.066]	
$\psi_{k,s,m}$: male, skilled, manual	-0.347	[0.073]	
$\psi_{k,s,r}$: male, skilled, routine	-0.347	[0.092]	
$\psi_{k,s,a}$: male, skilled, abstract	0.581	[0.083]	

Notes: The table shows the point estimates, standard errors and average marginal effects of the main parameters from the supply side of the model. Average marginal effects of the fertility and marital status variables are calculated by taking the numerical derivative of the probability of choosing home production with respect to the given variable. We calculate this derivative in each year separately and then take the average across all years. In the case of the pecuniary rewards (ψ_1), the reported average marginal effect is calculated by finding the numerical derivative of the probability that a labor type chooses a given market occupation with repect to the wage. Ie calculate this derivative in each year and for every possible labor type-occupation combination separately and then take the average across all the values.

Table 9: Counterfactual Excercises I

	Change in Supply and Wage Gaps: C.2013 - C.1990					
	ENIGH	Model	CF1 α's	CF2 CHL5	CF3 Marr	CF4 $\pi_{2,f}$, $\pi_{2,k}$
100 × Δ Log (Male/Female)Earnings Ratio						
College						
Abstract	-8.5	-13.5	29.3	-13.2	-13.5	-20.7
Routine	-27.4	0.5	51.5	0.6	0.1	-13.2
Manual	-29.0	30.0	-11.0	30.6	29.9	12.1
High School						
Abstract	-0.9	3.1	12.9	0.9	2.4	-10.3
Routine	13.9	14.2	26.9	10.0	12.8	-11.3
Manual	11.4	8.8	30.9	4.4	7.3	-18.0
100 × Δ Male - Female Occ. Share						
College						
Abstract	-2.1	-2.0	-0.4	-2.1	-2.0	-0.9
Routine	-0.4	-0.8	-0.2	-0.8	-0.7	-0.5
Manual	0.1	-0.3	0.0	-0.3	-0.3	-0.1
Home Production	-0.8	-0.2	-0.3	-0.1	-0.2	-1.7
High School						
Abstract	-1.5	-2.4	-2.1	-2.1	-2.3	-0.5
Routine	-2.0	-3.0	-3.0	-2.6	-2.9	-0.9
Manual	-6.0	-3.5	-3.8	-3.0	-3.3	-0.4
Home Production	13.0	12.4	12.6	11.1	12.0	5.3

Notes: The Table reports the difference between C.1990 and C.2013 of i) the log (male/female) earnings ratio and i) the change in the occupational shares by sex under alternative scenarios. The first two columns correspond to the ENIGH and model predictions. Column CF1 corresponds to the counterfactual estimates once all the linear, quadratic, and cubic coefficients of the α shares are set to zero. Column CF2 correspond to the counterfactual estimates once we set the probability of having a children under the age of 5 to the values of 1989, and constant across the years. Column CF3 correspond to the counterfactual estimates once we set the probability of being married or having a permanent partner to the values of 1989, and constant across the years. Finally, column CF4 correspond to the counterfactual estimates once we set the coefficients of the linear trends ($\pi_{2,f}, \pi_{2,k}$) in the home production utility function to zero.

Table 10: Counterfactual Excercises II

	Change in Supply and Wage Gaps: C.2013 - C.1990					
	ENIGH	Model	CF1 α's	CF2 CHL5	CF3 Marr	CF4 $\pi_{2,f}$, $\pi_{2,k}$
100 × Δ Log (Male/Female)Earnings Ratio						
College	-13.2	-11.7	29.9	-11.4	-11.7	-20.4
High School	15.2	14.6	27.1	10.8	13.3	-8.2
100 × Δ Log (College/High School)Earnings Ratio						
Male	-16.3	-20.4	-25.7	-19.8	-20.2	-9.8
Female	12.1	5.8	-28.6	2.4	4.8	-6.6

Notes: The Table reports the difference between C.1990 and C.2013 of the i) the log (male/female) earnings ratio and ii) the the log (college/high school) earnings ratio under alternative scenarios.. The first two columns correspond to the ENIGH and and model predictions. Column CF1 corresponds to the counterfactual estimates once all the linear, quadratic, and cubic coefficients of the α shares are set to zero. Column CF2 correspond to the counterfactual estimates once we set the probability of having a children under the age of 5 to the values of 1989, and constant across the years. Column CF3 correspond to the counterfactual estimates once we set the probability of being married or having a permanent partner to the values of 1989, and constant across the years. Finally, column CF4 correspond to the counterfactual estimates once we set the coefficients of the linear trends $(\pi_{2,f}, \pi_{2,k})$ in the home production utility function to zero.

Table 11: Parameter Estimates: Production Technology. Alternative Supply Measures

	Full-Time Workers			Part-Time Workers			Hours Worked		
	Estimate	SE	Elasticity	Estimate	SE	Elasticity	Estimate	SE	Elasticity
Gender									
$\rho_{4,m}$: female, male (manual)	0.175	[0.181]	1.212	-0.258	[0.152]	0.795	0.161	[0.138]	1.192
$\rho_{4,r}$: female, male (routine)	0.179	[0.129]	1.219	-0.030	[0.110]	0.971	0.355	[0.146]	1.551
$\rho_{4,a}$: female, male (abstract)	0.622	[0.099]	2.646	0.607	[0.121]	2.543	0.666	[0.108]	2.990
Education									
$\rho_{3,m}$: college, high school (manual)	0.722	[0.067]	3.594	0.771	[0.083]	4.371	0.803	[0.120]	5.081
$\rho_{3,r}$: college, high school (routine)	0.355	[0.041]	1.549	0.364	[0.073]	1.572	0.342	[0.122]	1.519
$\rho_{3,a}$: college, high school (abstract)	0.276	[0.121]	1.382	0.151	[0.197]	1.177	0.173	[0.211]	1.209
Occupation									
ρ_1: abstract, routine and manual	0.031	[0.094]	1.032	0.688	[0.167]	3.206	0.621	[0.186]	2.639
ρ_2: routine, manual	-0.141	[0.183]	0.877	-0.519	[0.146]	0.658	-0.246	[0.192]	0.803

Notes: The table shows the point estimates and standard errors of the elasticities of substitution from the production technology using different labor supply measures. Columns 1-3 report the baseline estimates; columns 4-6 report the estimates if income from part-time workers is also included in earnings series; and columns 7-9 report the estimates if we measure labor supply by the total number of hours worked of each group instead of the head-count. Since there is no measure of hours worked for people that are in home production, those values are imputed. We assign each person in home production the average number of hours worked by workers in market occupation with the same level of schooling, sex, and age.

Table 12: Parameter Estimates: Occupational Choice. Alternative Supply Measures

	Full-Time Workers			Part-Time Workers			Hours Worked		
	Estimate	SE	Av. MFX	Estimate	SE	Av. MFX	Estimate	SE	Av. MFX
Earnings									
ψ_1: earnings	0.154	[0.009]	0.022	0.090	[0.009]	0.013	0.138	[0.012]	0.023
Fertility/Children									
$\pi_{3,f,u}$: female, unskilled	0.712	[0.127]	0.174	0.647	[0.214]	0.159	0.661	[0.227]	0.162
$\pi_{3,f,s}$: female, skilled	-0.139	[0.229]	-0.025	0.158	[0.172]	0.029	0.126	[0.207]	0.022
$\pi_{3,k,u}$: male, unskilled	-0.367	[0.182]	-0.022	-0.242	[0.259]	-0.015	-0.238	[0.206]	-0.015
$\pi_{3,k,s}$: male, skilled	0.122	[0.219]	0.008	0.022	[0.235]	0.002	0.023	[0.263]	0.001
Marriage									
$\pi_{4,f,u}$: female, unskilled	0.589	[0.107]	0.144	0.567	[0.146]	0.139	0.574	[0.17]3	0.141
$\pi_{4,f,s}$: female, skilled	0.267	[0.134]	0.047	0.127	[0.187]	0.023	0.294	[0.179]	0.051
$\pi_{4,k,u}$: male, unskilled	-0.524	[0.155]	-0.032	-0.537	[0.182]	-0.034	-0.535	[0.191]	-0.033
$\pi_{4,k,s}$: male, skilled	0.026	[0.196]	0.002	-0.043	[0.234]	-0.003	-0.045	[0.166]	-0.003
Non Pecuniary Rewards/Tastes									
$\pi_{1,f}$: female, home production	0.961	[0.078]		0.811	[0.116]		0.935	[0.167]	
$\pi_{2,f}$: female, home production trend	-0.059	[0.003]		-0.051	[0.004]		-0.057	[0.004]	
$\pi_{1,k}$: male, home production	-0.716	[0.127]		-0.707	[0.135]		-0.609	[0.165]	
$\pi_{2,k}$: male, home production trend	0.054	[0.005]		0.047	[0.004]		0.048	[0.004]	
$\psi_{f,u,m}$: female, unskilled, manual	0.036	[0.048]		-0.041	[0.057]		-0.126	[0.096]	
$\psi_{f,u,r}$: female, unskilled, routine	-0.437	[0.046]		-0.420	[0.059]		-0.361	[0.096]	
$\psi_{f,u,a}$: female, unskilled, abstract	-0.555	[0.044]		-0.504	[0.061]		-0.455	[0.094]	
$\psi_{f,s,m}$: female, skilled, manual	-0.756	[0.130]		-0.641	[0.099]		-0.723	[0.107]	
$\psi_{f,s,r}$: female, skilled, routine	-0.362	[0.088]		-0.263	[0.106]		-0.066	[0.082]	
$\psi_{f,s,a}$: female, skilled, abstract	0.508	[0.080]		0.721	[0.105]		0.622	[0.089]	
$\psi_{k,u,m}$: male, unskilled, manual	0.541	[0.063]		0.577	[0.062]		0.687	[0.124]	
$\psi_{k,u,r}$: male, unskilled, routine	0.296	[0.064]		0.385	[0.063]		0.404	[0.123]	
$\psi_{k,u,a}$: male, unskilled, abstract	-0.682	[0.066]		-0.497	[0.063]		-0.542	[0.129]	
$\psi_{k,s,m}$: male, skilled, manual	-0.347	[0.073]		-0.484	[0.116]		-0.355	[0.100]	
$\psi_{k,s,r}$: male, skilled, routine	-0.347	[0.092]		-0.378	[0.111]		-0.318	[0.102]	
$\psi_{k,s,a}$: male, skilled, abstract	0.581	[0.083]		0.873	[0.115]		0.727	[0.106]	

Notes: The table shows the point estimates, standard errors and average marginal effects of the main parameters from the supply side of the model using different labor supply measures. Average marginal effects of the fertility and marital status variables are calculated by taking the numerical derivative of the probability of choosing home production with respect to the given variable. We calculate this derivative in each year separately and then take the average across all years. In the case of the pecuniary rewards (ψ_1), the reported average marginal effect is calculated by finding the numerical derivative of the probability that a labor type chooses a given market occupation with repect to the wage. We calculate this derivative in each year and for every possible labor type-occupation combination separately and then take the average across all the values.

Table 13: Parameter Estimates: Production Technology. Alternative Model Specifications

	Baseline			Nests Order Swap			Routine			Manual		
	Estimate	SE	Elasticity	Estimate	SE	Elasticity	Estimate	SE	Elasticity	Estimate	SE	Elasticity
Gender												
$\rho_{4,m}$: female, male (manual)	0.175	[0.181]	1.212	-0.246	[0.098]	0.802	-0.427	[0.177]	0.701	-0.029	[0.104]	0.972
$\rho_{4,r}$: female, male (routine)	0.179	[0.129]	1.219	-0.278	[0.102]	0.782	-0.095	[0.153]	0.913	0.007	[0.023]	1.007
$\rho_{4,a}$: female, male (abstract)	0.622	[0.099]	2.646	0.466	[0.115]	1.872	0.529	[0.099]	2.121	0.551	[0.099]	2.225
Education												
$\rho_{3,m}$: college, high school (manual)	0.722	[0.067]	3.594	0.564	[0.045]	2.292	0.454	[0.104]	1.831	0.581	[0.058]	2.385
$\rho_{3,r}$: college, high school (routine)	0.355	[0.041]	1.549	0.382	[0.054]	1.618	0.380	[0.051]	1.614	0.189	[0.047]	1.233
$\rho_{3,a}$: college, high school (abstract)	0.276	[0.121]	1.382	0.012	[0.040]	1.012	0.008	[0.048]	1.008	0.446	[0.150]	1.805
Occupation												
ρ_1: abstract, routine and manual	0.031	[0.094]	1.032	0.441	[0.109]	1.788						
ρ_2: routine, manual	-0.141	[0.183]	0.877	-1.816	[0.135]	0.355						
ρ_1: routine, abstract and manual							-0.784	[0.231]	0.560			
ρ_2: abstract, manual							0.332	[0.171]	1.496			
ρ_1: manual, abstract and routine										0.411	[0.134]	1.697
ρ_2: abstract, routine										-0.714	[0.099]	0.583

Notes: The table shows the point estimates and standard errors of the elasticities of substitution from the production technology using different model specification. Columns 1-3 report the baseline estimates; columns 4-6 report the estimates after switching the order of the second (education) and third (gender) nests of the production technology; columns 7-9 show the estimates if the occupational group that has the common elasticity with the other two groups is the routine task-intensive; and columns 10-12 report the estimates if the occupational group that has the common elasticity with the other two groups is the manual task-intensive.

Table 14: Parameter Estimates: Occupational Choice. Alternative Model Specifications

	Baseline			Nests Order Swap			Routine			Manual		
	Estimate	SE	Av. MFX	Estimate	SE	Av. MFX	Estimate	SE	Av. MFX	Estimate	SE	Av. MFX
Earnings												
ψ_1: earnings	0.154	[0.009]	0.022	0.109	[0.009]	0.018	0.085	[0.008]	0.014	0.098	[0.010]	0.016
Fertility/Children												
$\pi_{3,f,u}$: female, unskilled	0.712	[0.127]	0.174	0.805	[0.118]	0.197	0.832	[0.162]	0.206	0.664	[0.050]	0.162
$\pi_{3,f,s}$: female, skilled	-0.139	[0.229]	-0.025	-0.042	[0.116]	-0.008	-0.126	[0.330]	-0.022	0.022	[0.150]	0.004
$\pi_{3,k,u}$: male, unskilled	-0.367	[0.182]	-0.022	-0.162	[0.098]	-0.010	-0.174	[0.240]	-0.009	-0.246	[0.164]	-0.015
$\pi_{3,k,s}$: male, skilled	0.122	[0.219]	0.008	0.073	[0.113]	0.005	0.061	[0.310]	0.004	-0.050	[0.196]	-0.003
Marriage												
$\pi_{4,f,u}$: female, unskilled	0.589	[0.107]	0.144	0.767	[0.115]	0.188	0.500	[0.107]	0.124	0.588	[0.056]	0.144
$\pi_{4,f,s}$: female, skilled	0.267	[0.134]	0.047	0.318	[0.105]	0.059	0.103	[0.176]	0.018	0.193	[0.093]	0.035
$\pi_{4,k,u}$: male, unskilled	-0.524	[0.155]	-0.032	-0.598	[0.115]	-0.037	-0.635	[0.244]	-0.034	-0.577	[0.144]	-0.035
$\pi_{4,k,s}$: male, skilled	0.026	[0.196]	0.002	-0.145	[0.124]	-0.009	-0.142	[0.299]	-0.010	-0.129	[0.154]	-0.009
Non Pecuniary Rewards/Tastes												
$\pi_{1,f}$: female, home production	0.961	[0.078]		0.822	[0.088]		0.712	[0.105]		0.848	[0.049]	
$\pi_{2,f}$: female, home production trend	-0.059	[0.003]		-0.051	[0.003]		-0.026	[0.003]		-0.054	[0.002]	
$\pi_{1,k}$: male, home production	-0.716	[0.127]		-0.646	[0.098]		-0.971	[0.221]		-0.672	[0.124]	
$\pi_{2,k}$: male, home production trend	0.054	[0.005]		0.045	[0.002]		0.081	[0.005]		0.045	[0.003]	
$\psi_{f,u,m}$: female, unskilled, manual	0.036	[0.048]		0.196	[0.050]		0.101	[0.076]		0.019	[0.021]	
$\psi_{f,u,r}$: female, unskilled, routine	-0.437	[0.046]		-0.228	[0.052]		-0.208	[0.076]		-0.391	[0.024]	
$\psi_{f,u,a}$: female, unskilled, abstract	-0.555	[0.044]		-0.356	[0.052]		-0.525	[0.078]		-0.492	[0.027]	
$\psi_{f,s,m}$: female, skilled, manual	-0.756	[0.130]		-0.699	[0.063]		-0.804	[0.148]		-0.720	[0.092]	
$\psi_{f,s,r}$: female, skilled, routine	-0.362	[0.088]		-0.235	[0.061]		-0.126	[0.120]		-0.265	[0.063]	
$\psi_{f,s,a}$: female, skilled, abstract	0.508	[0.080]		0.739	[0.058]		0.863	[0.123]		0.757	[0.064]	
$\psi_{k,u,m}$: male, unskilled, manual	0.541	[0.063]		0.625	[0.025]		0.717	[0.062]		0.631	[0.022]	
$\psi_{k,u,r}$: male, unskilled, routine	0.296	[0.064]		0.423	[0.029]		0.527	[0.063]		0.424	[0.023]	
$\psi_{k,u,a}$: male, unskilled, abstract	-0.682	[0.066]		-0.522	[0.026]		-0.403	[0.063]		-0.489	[0.033]	
$\psi_{k,s,m}$: male, skilled, manual	-0.347	[0.073]		-0.591	[0.069]		-0.582	[0.069]		-0.474	[0.064]	
$\psi_{k,s,r}$: male, skilled, routine	-0.347	[0.092]		-0.149	[0.063]		-0.269	[0.065]		-0.316	[0.058]	
$\psi_{k,s,a}$: male, skilled, abstract	0.581	[0.083]		0.783	[0.063]		0.895	[0.063]		0.749	[0.071]	

Notes: The table shows the point estimates, standard errors and average marginal effects of the main parameters from the supply side of the model using different model specifications. Columns 1-3 report the baseline estimates; columns 4-6 report the estimates after switching the order of the second (education) and third (gender) nests of the production technology; columns 7-9 show the estimates if the occupational group that has the common elasticity with the other two groups is the routine task-intensive; and columns 10-12 report the estimates if the occupational group that has the common elasticity with the other two groups is the manual task-intensive. Average marginal effects of the fertility and marital status variables are calculated by taking the numerical derivative of the probability of choosing home production with respect to the given variable. We calculate this derivative in each year separately and then take the average across all years. In the case of the pecuniary rewards (ψ_1), the reported average marginal effect is calculated by finding the numerical derivative of the probability that a labor type chooses a given market occupation with repect to the wage. We calculate this derivative in each year and for every possible labor type-occupation combination separately and then take the average across all the values.

A.2 Division of Occupations into Manual, Routine, and Abstract Task-Intensive Groups

The ENIGH survey uses the Mexican occupation classification system to categorize workers according to the type of tasks they perform in the main job. The system went trough two changes since 1989: first there was an update of the original Clasificación Mexicana de Ocupaciones (CMO) in 1992, and then a full change to the newly introduced Sistema Nacional de Clasificación de Ocupaciones (SINCO) in 2010. These changes make the series incompatible at high levels of disaggregation of the occupational groups, but it is possible to homogenize the SINCO classification to the principal group level of the CMO using the comparability tables produced by INEGI.[29] The principal group division has 18 distinct occupational groups that can be consistently followed throughout the period of analysis.

The 18 principal level occupations from the ENIGH are classified in three groups defined by whether the activities done in the jobs are predominantly manual, routine (repetitive and easily codifiable tasks), or abstract intensive. The division is based on the measures constructed by Autor et al. (2003) from different sets of variables of the 1977 Dictionary of Occupational Titles (DOT) of the U.S., and then linked to the three-digit occupation codes of the CENSUS. The DOT evaluated highly detailed occupations along 44 objective and subjective dimensions that include physical demands and required worker aptitudes, temperaments and interests. Autor et al. (2006) used a subset of those dimensions to generate a simple typology consistent of three aggregates for abstract, routine, and manual tasks. The abstract task measure corresponds to the average from two variables of the DOT: DCP, which measures direction, control, and planning of activities; and GED-MATH, which measures quantitative reasoning requirements. The routine task measure corresponds to an average from two variables of the DOT: STS, which measures adaptability to work requiring set limits, tolerances, or standards; and FINGDEX, measuring finger dexterity. Finally, the manual task measure uses a single variable, EYEHAND, which measures eye, hand, foot coordination.[30]

In practice, we first create a cross-walk between three-digit CENSUS codes in the U.S. and the 18 categories of the principal group occupational division of the ENIGH. This task is facilitated by the fact that both the ENIGH and the U.S. CENSUS follow similar international standards when constructing their own occupation classifications. Since the three task measures are ordinal, there is no direct way to use the actual magnitude of the variables to compare occupations across the three dimensions. For each task measure we first organize the three-digit

[29] `http://www.inegi.org.mx/est/contenidos/proyectos/aspectosmetodologicos/clasificadoresycatalogos/sinco.aspx`

[30] See the online Appendix in Dorn (2009) for further details. Other papers that have used this measures include Autor et al. (2006); Goos and Manning (2007); Dorn (2009); Rendall (2013); Autor and Dorn (2013); Adda et al. (2017).

occupations by percentiles, and then calculate the median percentile of the measure within the broader 18 occupational groups of the ENIGH. Each of the 18 occupations is assigned to the group in which the median percentile was highest (see Table 3).

This procedure generated a balanced division with respect to the overall employment share of each group, and it is also consistent with the broad classification of aggregate occupations used in the literature that follow the task-based framework. Two important caveats should be stressed: First, any attempt to homogenize occupation classification systems from different countries involves some subjective choices. In the cases where we found that an occupation didn't have an immediate correspondence between the two systems, we had to use my judgement, based on documentation about the description of the occupation, to select a corresponding match. Second, the task measures were created specifically for U.S. economy, and it is likely that there are differences in the intensity in which certain skills are used in given occupations between the U.S. and Mexico. Results should be interpreted with this two caveats in mind.

A.3 Using RIF to Decompose Changes in Distributional Statistics beyond the Mean

Firpo et al. (2007, 2009) allow extending the traditional Oaxaca-Blinder decomposition to distributional statistics beyond the mean. This is achieved through the use of influence functions (IF). Influence functions measure the effect that an infinitesimal amount of "errors" have on a given estimator (Cowell and Victoria-Feser, 1996), but they also have properties that allows us to model the sensitivity of a given unconditional wage quantile to a change in a set of covariates. To see this, let $q_\tau(F_W)$ be τth quantile of the distribution of wages, expressed in terms of the cumulative distribution $F_W(w)$. Define the following mixture distribution:

$$G_{W,\epsilon} = (1-\epsilon)F_W + \epsilon H_W \quad for \quad 0 \leq \epsilon \leq 1 \tag{A.1}$$

where H_W is some perturbation distribution that only puts mass at the value w. In that case, $G_{W,\epsilon}$ is a distribution where, with probability $(1-\epsilon)$, the observation is generated by F_W, and with probability ϵ, the observation takes the arbitrary value of the perturbation distribution. By definition, the influence function corresponds to:

$$IF(w; q_\tau, F_W) = lim_{\epsilon \to 0} \ \frac{q_\tau(G_{W,\epsilon}) - q_\tau(F_W)}{\epsilon} \tag{A.2}$$

where the expression is analogous to the directional derivative of q_τ in the direction of H_W. Analytical expressions for influence functions have been derived for many distributional statistics.[31] The influence function in the case of the τth quantile

[31] Essama-Nssah and Lambert (2011) provides a comprehensive list of influence functions for different distributional statistics.

takes the form:

$$IF(w; q_\tau, F_W) = \frac{\tau - \mathbb{1}[w \leq q_\tau]}{f_W(q_\tau)} \tag{A.3}$$

where $\mathbb{1}[\cdot]$ is an indicator function and f_W is the PDF.[32] Using some of the properties of influence functions, a direct link with the traditional Oaxaca-Blinder approach can be established. In particular, a property that is shared by influence functions is that, by definition, the expectation is equal to zero.

$$\int_{-\infty}^{+\infty} IF(w; q_\tau, F_W) dF(w) = 0 \tag{A.4}$$

Firpo et al. (2009) propose a simple modification in which the quantile is added back to the influence function, resulting in what the authors call the Recentered Influence Function (RIF).

$$RIF(w; q_\tau, F_W) = q_\tau + IF(w; q_\tau, F_W) \tag{A.5}$$

The importance of this transformation lies in the fact that the expectation of the RIF is precisely the quantile q_τ. With this result, Firpo et al. (2009) show that we can model the conditional expectation of the RIF as a linear function of the explanatory variables.

$$E[RIF(w_t; q_\tau, F_{W,t}|X_t)] = X_t' \gamma_t \tag{A.6}$$

Moreover, if we apply the law of iterated expectations to Equation A.6, the end result is an expression that directly relates the impact of changes in the expected values of the covariates on the unconditional quantile q_τ. Note that this result is all that is required to extend the Oaxaca-Blinder decomposition to quantiles, since the basic components of the method are all present in Equation (A.6).

Estimation of Equation (A.6) can be done by OLS, and only requires replacing the dependent variable, $\log w_t$ in the original wage setting model with the RIF of the quantile q_τ. The interpretation of the estimates $\widehat{\gamma}_t$ can be thought of as the effect of a small change in the distribution of X on q_τ, or as linear approximation of the effect of large changes of X on q_τ (Firpo et al., 2007).

A.4 Estimation of the Model: Error Structure, Weight Matrix, and Standard Errors

We assume a simplified error structure to facilitate estimation. Let Θ be the 81×1 vector of parameters to be estimated. Let $p(\Theta)$ by the 364×1 vector of wage

[32] Note that the influence function in this case depends on the density. In order to obtain the empirical density the authors propose non-parametric kernel density estimation.

and supply predictions of the model as function of the parameters. Finally, let q be the observed vector of wages and labor shares taken directly from the ENIGH data. For any given prediction i, we assume that the error term, e_i, at the true parameter vector, Θ^*, follows a normal distribution centred at zero that is independent across i. That is,

$$e_i = q_i - p_i(\Theta^*), \tag{A.7}$$

where $f(e_i) = \frac{1}{\sqrt{2\pi\sigma_i^2}} exp(\frac{e_i^2}{2\sigma_i^2})$.

The log-likelihood function takes the form

$$\log \mathcal{L}(\Theta) = \sum_i \log f(e_i) = \sum_i \log f(q_i - p_i(\Theta)), \tag{A.8}$$

and the respective score function, $s(\Theta)$, is:

$$s(\Theta) = \frac{\partial \log \mathcal{L}(\Theta)}{\partial \Theta} = \sum_i \frac{\partial \log f(q_i - p_i(\Theta))}{\partial \Theta} = \sum_i \frac{1}{\sigma_i^2} \frac{\partial p_i(\Theta)}{\partial \Theta}(q_i - p_i(\Theta)), \tag{A.9}$$

which we can write more compactly in vector form as

$$s(\Theta) = W'(\Theta)(q - p(\Theta)). \tag{A.10}$$

Here, $W(\Theta)$ is 364 $\times$ 81 weight matrix that depends on the derivatives of the vector of predictions with respect to each of the parameters, and the variance of each prediction error σ_i^2. At the maximum likelihood estimate, $\hat{\Theta}_{ml}$, the score vector of the log likelihood is set to zero:

$$s(\hat{\Theta}_{ml}) = W'(\hat{\Theta}_{ml})(q - p(\hat{\Theta}_{ml})) = 0. \tag{A.11}$$

We use $m = q - p(\Theta)$ as a vector of population moments such that $E(q - p(\Theta)) = 0$, and obtain a a consistent estimator of Θ^* by GMM:

$$g(\hat{\Theta}_{gmm}) = W'(q - p(\hat{\Theta}_{gmm})) = 0, \tag{A.12}$$

where W' is a fixed positive definite matrix of instruments. Efficient GMM estimator can be obtained by choosing instruments that are asymptotically equivalent to the weights $W'(\hat{\Theta}_{ml})$ in Equation (A.10). The problem is that we would need to have a consistent initial estimate of Θ^*. Given that we do not have those initial consistent estimates, we follow a iterative process. We start from a plausible set of initial values of the parameters (Θ_0), and use them to estimate the vector of partial derivatives $\frac{\partial \hat{p}_i(\Theta_0)}{\partial \Theta_0}$. The estimates of the variance of each error, $\hat{\sigma}^2_{i,0}$, are calculated as the square of the estimated error from this initial set of parameter values. Both of these estimates are then used to construct an initial weight matrix, which allows us to solve the minimization problem.[33] The estimates obtained after this first iteration[34] are used to update the weight matrix, and the process continues until the parameter vector converges to a stable point.

Since it is usually not possible to satisfy Equation (A.12), we estimate the parameters of the model using the quadratic form:

$$\hat{\Theta}_{gmm} = \text{argmin}[q - p(\Theta)]' W(\Theta) W'(\Theta)[q - p(\Theta)]. \tag{A.13}$$

Finally, the standard errors of the parameter estimates are calculated applying the standard method of moments formula. Let Γ be the matrix of partial derivatives of the sample moments $\bar{m}(\hat{\Theta}_{gmm})$ with respect to the parameters. The ith row correpsonds to:

$$\Gamma_i(\hat{\Theta}_{gmm}) = \frac{\partial \bar{m}_i(\hat{\Theta}_{gmm})}{\partial \hat{\Theta}_{gmm}}, \tag{A.14}$$

so the variance-covariance matrix can be calculated using:

$$\hat{Var}(\hat{\Theta}_{GMM}) = \left[\Gamma(\hat{\Theta}_{gmm})' \hat{Var}[\bar{m}(\Theta_{gmm})]^{-1} \Gamma(\hat{\Theta}_{gmm})\right]^{-1}. \tag{A.15}$$

[33] The parameter search is done using the interior-point algorithm in Matlab.

[34] Note that even though the weight matrix is a function of the parameters, it remains fixed during the parameter search.