

Bhalotra, Sonia R.; Fernández, Manuel

Working Paper

The distribution of the gender wage gap

ISER Working Paper Series, No. 2018-10

Provided in Cooperation with:

Institute for Social and Economic Research (ISER), University of Essex

Suggested Citation: Bhalotra, Sonia R.; Fernández, Manuel (2018) : The distribution of the gender wage gap, ISER Working Paper Series, No. 2018-10, University of Essex, Institute for Social and Economic Research (ISER), Colchester

This Version is available at:

<https://hdl.handle.net/10419/200374>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

The Distribution of the Gender Wage Gap

Sonia R. Bhalotra
ISER, University of Essex

Manuel Fernández
ISER, University of Essex

No. 2018-10
June 2018



INSTITUTE FOR SOCIAL
& ECONOMIC RESEARCH

Non-Technical Summary

Many modernizing economies are witnessing rapid growth in women's labour force participation, in the way that today's richer countries did. While this is often associated with women's empowerment, simple economic theory suggests that increases in the labour supply of women will tend to depress the wages of women, potentially widening the gender wage gap. The extent to which this happens will depend upon the degree to which men and women are substitutes in the workplace. One of our key observations is that this varies across the wage distribution, in line with the task content of jobs. Recent research has analysed the degree to which machines or technology substitute individuals, documenting a polarization of the labour market with jobs in the middle 'disappearing'. We focus instead on substitutability between men and women.

We use individual data on employment and earnings in Mexico that span more than a quarter of a century, 1989-2014, capturing one of the most rapid contemporary increases in women's labour supply. We find that the polarization of tasks is reflected in opposing effects on the gender wage gap. In particular, in high-paying occupations that are intensive in abstract and analytical skills, there is a high degree of substitutability between women and men. As a result, a large increase in (skilled) women joining the labour force may depress wages in general, but without increasing the gender wage gap. On the other hand, in low-paying occupations in which individuals do manual or routine tasks, we find that women and men are poor substitutes. Thus, increases in (low skilled) women joining the labour force will tend to widen the gender wage gap. While we have highlighted in this discussion the consequences of increases in women's labour supply, the gender wage gap will also depend on whether the demand for male vs female labour is growing. In fact demand trends have favoured women, attenuating the supply-driven downward pressure on women's wages in low-paid occupations, and fully counteracting it in high-paid occupations.

Pulling the supply and demand trends together and accounting for the education of workers and the skill-content of occupations, we find that the gender wage gap in Mexico narrowed between 5 and 18 percent among workers above the 80th percentile, and among workers below the median it widened by 10 to 22 percent.

The Distribution of the Gender Wage Gap*

Sonia R. Bhalotra[†] Manuel Fernández[‡]

June 26, 2018

Abstract

We analyse impacts of the rising labor force participation of women on the gender wage gap. We formulate and structurally estimate an equilibrium model of the labor market in which the elasticity of substitution between male and female labor is allowed to vary depending on the task content of occupations. We find that the elasticity of substitution is higher in high-paying occupations that are intensive in abstract and analytical tasks than in low-paying manual and routine occupations. Consistent with this we find a narrowing of the gender wage gap towards the upper end of the wage distribution and an increase in the gender wage gap at the low end. Demand side trends favoured women and this attenuated the supply-driven downward pressure on women's wages in low-paying occupations, and fully counteracted it in high-paying occupations. The paper contributes new evidence on the distribution of the gender wage gaps, and contributes to a wider literature on technological change, occupational sorting, wage inequality and polarization.

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

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

*We are grateful to Carlos Carrillo-Tudela, Marco Manacorda, Brian Bell, Climent Quintana-Domeque, Simon Quinn, John Muellbauer, James Fenske, and seminar participants at University of Essex, Universidad Javeriana, 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 increases in the labour force participation of women are one of the most salient features of labour markets in most countries over the last century (Killingsworth and Heckman, 1987; Goldin, 2006; Fogli and Veldkamp, 2011; Goldin and Olivetti, 2013). But although the massive entry of women into the workforce has fundamentally changed the gender composition of labour supply, there is little understanding of what its impact has been on the wage structure. The relative scarcity of empirical studies on this topic reflects the complexity of the phenomena, since a variety of supply and demand factors will influence the decision of women to participate in the labour market, and natural experiments that could provide exogenous variation in female labour supply have been elusive. Moreover, the few studies that have tackled this question provide mixed – and at times contradictory – conclusions, mostly because there is no agreement in the literature of whether male and female labour are close substitutes or not, a key factor underlying the debate.

This paper studies the effect of the rapid rise of female labour force participation on the wage structure in Mexico. Using an equilibrium model of the labour market, we structurally estimate the elasticities of substitution between male and female labour that are consistent with observed patterns of the gender earnings gap since 1989. Our model differs from other work on the subject in two significant ways: First, the model incorporates the ideas of the task-based approach (Autor et al., 2003), allowing the elasticity of substitution between male and female labour to vary as a function of the task content of occupations. This feature implies that the effect of female labour supply on relative earnings can vary throughout the pay distribution, a possibility that is suggested by the descriptive evidence. Second, following the work of Johnson and Keane (2013), labour supply is endogenously determined in the model, a feature that addresses one of the major criticisms of the canonical framework in the past.

The focus on Mexico is motivated by the extent to which Mexican women have joined the workforce over the last quarter century, which was the largest in the Latin American region (Ñopo, 2012), and one of the largest in the world (The World Bank, 2012): between 1989 and 2014, labour force participation among prime-age women in Mexico increased from 36 to 58 percent, going from 4.7 to 14.7 million female workers in a span of 25 years – it took women in the U.S. 37 years, from 1956 to 1993, to cover the same ground. Moreover, had the female participation rate remained at the levels of 1989, there would be approximately 5.6 million fewer

prime-age females in the labour force.

At the same time that women started moving into the workforce, growth in labour earnings in Mexico diverged significantly between men and women at the tails of the earnings distribution, but remained constant at the mean. Growth in labour earnings among workers in low-paying occupations (e.g. agriculture, services, transportation, and repetitive production occupations) was higher for males relative to females, while the opposite was true within high-paying occupations (e.g. professional, education, technical and managerial occupations).

The main proposition of this paper is that these two phenomena are causally connected, so that the rapid increase of female labour supply has fundamentally changed the wage structure in the Mexican economy. In particular, we argue that the influx of women into the labour market has generated downward pressure on the wages of female workers in low-paying occupations, but on *both* male *and* female workers in high-paying occupations. The hypothesis we put forward is that this heterogeneity arises because the elasticity of substitution between male and female labour is not uniform throughout the occupational distribution: male and female labour are closer substitutes in occupations that rely heavily on abstract analytical skills than in those occupations in which physical or other non-cognitive skills are more prominent in the jobs.

To test this hypothesis we propose an extension to the canonical supply and demand framework popularized by Katz and Murphy (1992), Murphy and Welch (1992), and Juhn et al. (1993), which incorporates the ideas of the task-based approach of Autor et al. (2003). In the original formulation of the supply and demand framework, the degree to which changes in the composition of labour supply are associated with changes in the wage structure depends on the magnitude of the elasticity of substitution between the groups being compared, but these elasticities are usually taken to be constant across occupations. Starting from the work of Autor et al. (2003), a large strand of the literature¹ has emphasized that the complementary or substitutability between factors of production is determined by the type of tasks in which they are employed. In all occupations, cognitive, manual, physical, socio-emotional, and interpersonal skills are applied to specific tasks in different intensities, so the relative importance of any subset of skills is mostly determined by the nature of the activities being done. A logical implication of this

¹See, among others, Autor et al. (2003); Goos and Manning (2007); Autor et al. (2008); Gathmann and Schoenberg (2010); Acemoglu and Autor (2011); Firpo et al. (2011); Autor and Dorn (2013); Altonji et al. (2014); Goos et al. (2014); Michaels et al. (2015).

theory is that as long as there is some difference in the bundle of skills that men and women are supplying to the labour market, compositional changes by gender will have heterogeneous impacts on relative earnings that depend on the task content of occupations.

This paper contributes to the relatively scarce empirical literature studying the relation between female labour supply and changes in the wage structure. Arguably, the most influential paper on this topic is that of Acemoglu et al. (2004), who exploited state level variation in U.S. military mobilizations for World War II to investigate the effects of a rise of female labour force participation on the wage structure during the 1940s and 1950s. The authors found that rising female labour supply reduced both male and female wages. They report estimates of the short-run elasticity of substitution between male and female labour of around 3. The authors qualify this finding arguing that this elasticity potentially varies across skill groups, but they do not properly test that hypothesis.

In earlier studies, both Topel (1994) and Juhn and Kim (1999) examined the relation between the rise in female labour supply and rising inequality in the U.S. during the 1970s and 1980s, reaching opposite conclusions: while Topel finds that the rise in the supply of female labour negatively affected the wages of low-skilled male workers, Juhn and Kim challenged this view arguing that, once changes in relative demand are accounted for, there is little evidence that women are close substitutes for men or that the entry of educated women into the labour force contributed to male wage inequality growth in the 1980s. More recently, Johnson and Keane (2013) estimated a dynamic equilibrium model of the labour market fitted to replicate the patterns of the U.S. wage structure from 1968 to 1996. They report a relatively high elasticity of substitution between male and female labour of between 4.76 and 5.6.

The lack of agreement about the range of values that the elasticity of substitution between male and female labour can take is surprising, especially since this is a structural parameter that can be of first-order importance in contexts in which women are transitioning rapidly into the workforce. We argue that the application of the task-based approach provides the right framework to bridge together the mixed evidence on this subject. In the particular case of Mexico, we estimate these elasticities to be between 1.2 in manual and routine task-intensive occupations, and 2.7 in abstract task-intensive occupations. These numbers imply that the change in the gender composition of the workforce has had a significant effect on the gender earnings gap in Mexico, especially at the lower end of the pay distribution.

The paper also contributes to the literature studying changes in the wage structure in Latin American countries. Most countries in Latin America experienced a sharp fall in earnings inequality since the late 1990s (Ferreira et al., 2008; Kahhat, 2010; López-Calva and Lustig, 2010; Gasparini and Lustig, 2011; Gasparini et al., 2011; Levy and Schady, 2013; Lustig et al., 2013; Fernandez and Messina, 2017). We show that Mexico also experienced a significant contraction of the male earnings distribution over the last 25 years, driven by higher wage growth among low-skilled workers, but that this was not the case for women: female wage growth was u-shaped, a phenomena that is partly attributable to the impact of rising female labour supply. Moreover, our results indicate that the sluggish wage growth among high-skilled workers in Latin American countries could be partly explained by the the incorporation of college educated women into the workforce. This specific 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, sample selection, and treatment of the main variables used in the paper. Section 3 presents the main stylized facts, reviewing the evolution of male and female labour force participation in Mexico, and documenting the changes in the wage and occupational structure over the last quarter century. In Section 4 we formulate an equilibrium model of the labour market, and describe the empirical strategy that will be used to estimate its main parameters. Results are presented in Section 5, while robustness exercises using alternative specification of the model and different measures of labour supply are shown in Section 6. Finally, Section 7 concludes.

2 Data, Choice of Variables, and Sample Selection

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).² The first three waves of ENIGH were collected intermittently (1984, 1989, and 1992), but from 1992 onwards the

²The main labour 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.

survey was collected biannually, with an exceptional gathering in 2005. we use the 13 waves of ENIGH between the years 1989 and 2014. As a check for consistency, we replicate part of the results using data from the 1960, 1970, 1990, 2000 and 2010 Mexican CENSUS.³

Our measure of individual incomes includes only remuneration from labour. In particular, we use the monthly monetary remuneration from labour in all occupations,⁴ 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.⁵ 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 males and females: 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.⁶ 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 \$150.⁷

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 through 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

³The CENSUS data corresponds to the extended questionnaire applied to a 10 percent sample of the population in each year.

⁴Only 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 labour in all jobs in every year.

⁵The extended questionnaire of the CENSUS only asks about overall remuneration from labour, which can include income from self-employment. Hence, the two series are not fully comparable.

⁶Results 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.

⁷In 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.

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.⁸ The principal group division has 18 distinct occupational groups that can be consistently followed throughout the period of analysis.

Following the task-based framework, occupations are classified in three groups defined by whether the activities performed in the jobs are predominantly manual, routine, or abstract/analytical. 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 construction of the groups are presented in Appendix 8.2, while the final division is shown in Table 1.

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 labour market attachment, and to ameliorate selection problems arising from changes in educational and retirement choices of different cohorts.

3 Trends in Female labour Supply and Relative Earnings by Gender

In 1989 prime-age population in Mexico was just above 25 million people, out of which 64.2 percent were either working or actively searching for a job. Female labour force participation was as low as 36 percent, and women accounted for only 29 percent of the workforce (4.7 million). By 2014 this picture had changed drastically: prime-age population had grown to 48 million people, and women already represented 41 percent of the workforce, with close to 14.7 million females active in the labour market (Panels (a) and (b) of Figure 1). Moreover, overall participation rates had reached 76 percent, with all of the increase during this 25 year period accounted by the sharp rise of female labour force participation (58% in 2014) (see Panels (c) and (d) of Figure 1).

The extent to which Mexican women have joined the workforce over the last quarter century is remarkable: had the female participation rate remained at the levels of 1989, there would be 5.6 million fewer prime-age females in the labour

⁸<http://www.inegi.org.mx/est/contenidos/proyectos/aspectosmetodologicos/clasificadoresycatalogos/sinco.aspx>

force, which is close to 12 percent of its actual size in 2014. But the pace of this transition was also exceptional: in 1956, female labour force participation in the U.S. was similar to that of Mexico in the late 1980's, but it took the U.S. until 1993, 37 years later, to reach 58 percent, a rate at which economy stabilized (U.S. Bureau of Labor Statistics, 2017). Cross-country comparisons suggest that the rise of female labour force participation 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).

Can such a rapid change in the gender composition of labour supply affect the gender earnings gap, and, more generally, the wage structure in a country? The main proposition of this paper is that it does, especially among female workers labouring in low-paying occupations.

Panels (a) and (b) of Figure 2 show that, on average, Mexican women earn between 16 and 22 percent less than their male counterparts, even after taking into account differences in the levels of education, age, and occupational choice.⁹ Interestingly, the mean gender earnings gap in Mexico has remained more or less constant over the past 25 years, a fact that has generated significant debate over the reasons for such lack of convergence. But focusing only on mean relative earnings conceals the substantial heterogeneity in the earnings dynamics between men and women across the distribution. Panels (c) and (d) of Figure 2 show the change between C.1990 and C.2013 (1990 and 2010 in the CENSUS)¹⁰ of log hourly earnings at different percentiles of the earnings distribution, calculated separately for men and women. 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 these divergent dynamics (Panels (e) and (f) of Figure 2) is that the gender earnings gap increased between 10 and 22 percent among workers below the median, but declined between 5 and 18 percent among workers above the 80th percentile, with the average gap remaining relatively stable across the years.

We postulate that these diverging dynamics of relative earnings across the

⁹The 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.

¹⁰To increase sample size in the ENIGH survey, we joined together surveys from 1989 and 1992, and from 2012 and 2014.

distribution are associated to the sharp gender compositional shift that the Mexican economy experienced during this period. In particular, that the influx of women into the labour market generated downward pressures to female earnings among workers in low-paying occupations, but to *both* male and female workers in high-paying occupations. The reason for this heterogeneity is that the elasticity of substitution between male and female labour is not homogeneous throughout the occupational distribution, so similar changes in male/female relative supplies in high and low-paying occupations can have very different affects on the wage structure.

The underlying theoretical base for this proposition is the canonical supply and demand framework popularized by Katz and Murphy (1992), Murphy and Welch (1992), and Juhn et al. (1993); and the more recent task-based approach proposed by Autor et al. (2003). Within the supply and demand framework, the degree to which changes in the composition of labour supply are associated with changes in the wage structure fundamentally depends on the magnitude 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:

$$Y_t = [\alpha_t L_{k,t}^\rho + (1 - \alpha_t) L_{f,t}^\rho]^{1/\rho}, \quad (3.1)$$

where Y_t is total output at time t ; $L_{k,t}$ is the total supply of labour from male workers; $L_{f,t}$ is the total supply of labour from female workers; α_t is a time-varying “share” parameter that captures differences in the intensity of labour demand between male and female labour, 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 labour: $\sigma_\rho \equiv \frac{1}{1-\rho}$. 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:

$$\log \left(\frac{W_{k,t}}{W_{f,t}} \right) = \log \left(\frac{\alpha_t}{1 - \alpha_t} \right) - \frac{1}{\sigma_\rho} \log \left(\frac{L_{k,t}}{L_{f,t}} \right), \quad (3.2)$$

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

Equation (3.2) shows that the evolution of the log (male/female) earnings ratio depends on two factors: (i) how the relative supply of labour between males and females is evolving $\left(\frac{L_{k,t}}{L_{f,t}}\right)$, scaled by the inverse of the elasticity of substitution (σ_ρ); and (ii) how relative demands are moving, as captured by the variation in time of the log ratio of the α_t share. If male and female labour are imperfectly substitutable in production, that is, if σ_ρ has a “low” value, and relative demands trends are constant, large shifts in female labour 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: as the supply of female labour becomes more abundant its relative price falls. If, on the other hand, male and female workers are close to perfect substitutes – “high” values of σ_ρ –, a rise in female labour force participation should depress *both* male and female wages, with relative earnings remaining fairly constant.

A direct implication of this simplified framework is that if there is some variation in the substitutability between male and female labour across the earnings distribution, relative changes in the gender composition of labour supply will not have a symmetric impact on the wage structure. But should we expect this type of heterogeneity in the labour market? Starting from the influential work of Autor et al. (2003), a rapidly growing strand of the literature has emphasized that the complementary or substitutability between factors of production is determined by the type of tasks in which they are employed.¹¹ 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. Hence, as long as there is some difference in the bundle of skills that men and women supply to the labour market, compositional changes by gender will have an affect on relative earnings, and the extent to which this mechanism matters will depend on the task content of occupations.

Panels (a)-(c) of Figure 3 present some suggestive evidence that the sharp change in the gender composition of labour supply in Mexico was negatively correlated with the change in the gender earnings gap – something we might expect from the supply and demand framework –, but only in low-paying occupations. Each panel, corresponding to a specific occupational group, depicts two series: one has the evolution of log (male/female) earnings ratio, and the other has the evolution of

¹¹This 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.

log (male/female) relative supply. The three groups are chosen following the standard practice in the task-based framework, where occupations are divided based on whether the activities in the jobs are more intensive in manual, routine, or abstract tasks (see Table 1).¹²

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 subsequent 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 labour, 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 where 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.

The first three Panels of Figure 3 show that between 1989 and 2014, log (male/female) relative supply declined by 44 log points in both manual and routine task-intensive occupations, and by 69 log points in abstract task-intensive occupations. During the same period, the log (male/female) earnings ratio only increased in the lower-paying jobs: 11 log points in manual task-intensive occupations and 3 log points in routine task-intensive occupations. In the abstract task-intensive group the gender earnings gap fell by 10 log points.

The negative comovement between relative quantities and relative prices suggest that the rise of female labour force participation is potentially a factor behind the rise of the gender earnings gap at the low end of the earnings distribution, but the argument is less clear for high-paying abstract task-intensive jobs, even though the relative shift in supply was of an even larger magnitude. Under the canonical framework, the fact that relative quantities and relative prices are moving in the same direction indicate that men and women are more closely substitutable in those

¹²Section 2 discusses the details of the division of the principal level occupations along the three groups.

jobs, or that relative demand trends must be favouring female workers at the top. The equilibrium model we propose in Section 4 will shed light and quantify the relative importance of these two alternatives.

Panel (d) of Figure 3 provides further suggestive evidence that the supply and demand framework, in combination with the task-based approach, might be useful to understand the changes in the gender earnings ratio in Mexico. The figure depicts the change in log mean (male/female) earnings ratio as a function of the change of log mean (male/female) relative supply between C.1990 and C.2013, where each circle represents one of the principal group occupations of the ENIGH. Although some occupations, mostly in the abstract task-intensive group, saw a positive comovement between relative quantities and relative prices, there is a clear negative relation between the magnitude of the compositional shift in labour supply and the magnitude of the change in the gender earnings gap.

3.1 Did the Wage Structure Changed? Decomposing the Gender Earnings Gap Across Quantiles

The underlying assumption behind the previous descriptive evidence is that the wage structure did in fact change, which is not necessarily the case. The rise of female labour force participation was not the only major change in the Mexican labour market over this period, so pure compositional shifts, absent of any changes in relative pay, could also be consistent with the observed trends.¹³ For example, as shown in Table 2, the share of prime-age women in the workforce with at least some college education increased from 14.5 to 24.0 percent between C.1990 and C.2013, 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 observed convergence in earnings at the top of the distribution. Also, the average Mexican worker is becoming older, with the share of workers between 25 and 34 years of age falling between C.1990 and C.2013 from 46.5 to 34.9 percent in the case of women, and from 43.5 to 36 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

¹³A 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).

behind the divergence in earnings at the bottom of the distribution. Finally, women saw a slight relative shift towards the low-paying manual task-intensive occupations (from 34.4% to 35.3%), which is mostly explained by an increase of female workers in service occupations, while the participation of males in that occupational group declined (from 40.4% to 37.1%), which is mostly explained by a movement away from agricultural occupations.

Has the wage structure changed or is there a simple re-composition of the workforce? A simple 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 ratio at different percentiles under the observed compositional changes. Alternatively, we can keep the composition of the labour force fixed at a given point in time and construct counterfactual earning ratios to evaluate how changes in schooling, age and occupational premiums have affected the observed dynamics. Again, 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 8.3. As a starting point, consider a transformed wage-setting model of the form:

$$RIFq_{\tau,gen,t} = X'_{gen,t}\gamma_{gen,t} + \epsilon_{gen,t}, \quad (3.3)$$

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 quadratic terms in education and age, and a set of indicator variables for the 18 principal group level occupations of the ENIGH;¹⁴ and $\epsilon_{gen,t}$ is the error term assumed to have zero conditional mean. We can estimate Equation (3.3) for each gender and period separately by OLS, and then express the estimated

¹⁴As in the Oaxaca-Blinder decomposition for the mean, this decomposition is not invariant to the reference variable chosen when covariates are categorical. We limit this problem by using quadratic polynomials in years of education and age instead of levels. In the case of occupations, we use the more disaggregate principal group level, and repeat the decomposition changing the reference groups, but results were qualitatively similar.

difference over time of the expected value of the earnings quantile \hat{q}_τ as:

$$\Delta_t \hat{q}_{\tau,gen} = \underbrace{(\overline{X'}_{gen,C.2013} - \overline{X'}_{gen,C.1990}) \hat{\gamma}_{gen,P}}_{\Delta_t \hat{q}_{X,\tau,gen}} + \underbrace{\overline{X'}_{gen,P} (\hat{\gamma}_{gen,C.2013} - \hat{\gamma}_{gen,C.1990})}_{\Delta_t \hat{q}_{S,\tau,gen}}, \quad (3.4)$$

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.¹⁵ 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:

$$\underbrace{\Delta_t \hat{q}_{\tau,k} - \Delta_t \hat{q}_{\tau,f}}_{\text{Overall}} = \underbrace{(\Delta_t \hat{q}_{X,\tau,k} - \Delta_t \hat{q}_{X,\tau,f})}_{\text{Composition}} + \underbrace{(\Delta_t \hat{q}_{S,\tau,k} - \Delta_t \hat{q}_{S,\tau,f})}_{\text{Wage Structure}}. \quad (3.5)$$

The results of the decompositions for 5 selected percentiles are shown in Table 3. Each cell reports the contribution of the row variable to the observed change in the log (male/female) earnings ratio between C.1990 and C.2013 at the given percentile. The first conclusion from the decompositions is that wage structure effects are quantitatively more important than pure compositional effects in explaining the change in the gender earnings gap across the distribution. Estimated wage structure effects contribute 72 percent of the observed rise in the gender earnings gap at the 10th percentile, and close to 98 percent at the median. Moreover, they over-predict the fall in the gap at the 90th percentile (-0.097 log points observed vs. -0.119 log

¹⁵This specific counterfactual allows us to analyze composition and wage structure effects relative to a baseline defined by both the (weighted) mean returns and (weighted) mean characteristics over the two periods.

points attributed to the wage structure).

The results also show that if the wage structure had remained constant at the average levels over the two periods, compositional effects would lead to a larger counterfactual gender earnings gap at all points of the distribution, which 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.

The second conclusion from the decomposition is that, out of the three set of variables in the wage setting model, changes in relative pay across occupations are the most important factor behind the overall wage structure effects, particularly at the two tails of the distribution. Wage structure effects associated with the occupation variables account for 83 percent of the overall wage structure effects at the 10th percentile, and also over-predict the contribution of changes in relative pay at the top (-0.151 log points attributed to occupations vs. -0.119 log points attributed to the overall wage structure). This result can be seen more clearly in Figure 4, where we show both the overall change in the log (male/female) earnings ratio between C.1990 and C.2013, and the estimated amount of this change that is attributed to changes in the occupation wage structure. It is clear from the figure that the trends in the two series mimic each another.

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 question remains if this phenomenon was driven by the rise in female labour force participation. In the next section we develop an equilibrium model of supply and demand for labour to try to answer this question. 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 labour across the occupational distribution; finally, to run counterfactual exercises and get quantitative estimates of the effect of the rise of female labour 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 labour 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.¹⁶ Each type of agent chooses 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 3 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 labour in the three types of market occupations:

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

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 labour 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 labour 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 labour in routine and manual task-intensive occupations

¹⁶In order to maintain a tractable number of parameters in the model, we omit the age dimension from this analysis. The decision is also supported by the fact that the age dimension played a secondary role for explaining the patterns of the gender earnings gap in the decomposition.

($\sigma_{\rho_2} \equiv \frac{1}{1-\rho_2}$).¹⁷ Finally, $\alpha_{1,t}$ is a time-varying share parameter that captures both differences in the intensity of labour 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 labour 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-labour 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, labour in each occupation is divided in two groups based on the schooling level of the workers. In particular, labour in each occupation consists of a productivity weighted CES combination of labour from skilled workers, indexed by s , and labour from unskilled workers, indexed by u . That is:

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

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

Finally, in the third level of the production technology labour 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:

¹⁷Note that by assumption, the elasticity of substitution between labour 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 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.

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

where the parameters have an analogous interpretation to those in Equation (4.1). Note that elasticities of substitution between male and female labour are allowed to vary between occupations, but not within occupation and education. Any differences in substitutability between labour by education group is captured by the relevant parameters in the second level of the production technology.¹⁸

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 ($\rho_1, \rho_2, \rho_{3,a}, \rho_{3,r}, \rho_{3,m}, \rho_{4,a}, \rho_{4,r}$, and $\rho_{4,m}$); and a set 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.¹⁹ For example, the parameter $\alpha_{1,t}$ is allowed to change according to:

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

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

¹⁸To 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.

¹⁹The 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.

labour within a given occupation and education group can be expressed as:

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

$$\text{and } occ = a, r, m,$$

$$(4.5)$$

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; Welch, 2000; Rendall, 2013; Baker and Cornelson, 2016). The model attempts to incorporate both individual preferences and equilibrium returns to labour 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

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

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 labour 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 labour market – or, more specifically, into market occupations – 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 labour 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:

$$U(h \mid gen, edu, t) = \pi_{1,gen} + \pi_{2,gen}t + \pi_{3,gen,edu}Pr(CHL5 = 1 \mid gen, edu, t) + \pi_{4,gen,edu}Pr(MARR = 1 \mid gen, edu, t) + \epsilon_{gen,edu,h,t}, \quad (4.7)$$

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; $Pr(\text{CHL5} = 1 \mid \text{gen}, \text{edu}, t)$ 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 $Pr(\text{MARR} = 1 \mid \text{gen}, \text{edu}, t)$ is the probability that the worker is married or has permanent partner. Finally, $\epsilon_{\text{gen}, \text{edu}, h, t}$ is a idiosyncratic taste shock assumed to be independent and identically distributed extreme value.

Panel (a) and (b) of Figure 5 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.²⁰ 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 labour 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

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

We can use these probabilities to find the total labour supply of each type

²⁰The 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.

in each occupation. For example, the total supply of female workers with college education in abstract task-intensive occupations is

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

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

$$W_{gen,edu,occ,t} = \frac{\partial Y_t}{\partial L_{gen,edu,occ,t}^d}. \quad (4.10)$$

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

$$L_{gen,edu,occ,t}^d = L_{gen,edu,occ,t}^s \quad \text{for } occ = a, r, m, \quad (4.11)$$

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 set of parameters, this can be accomplished 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 labour 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 total²¹ that are a function of the 81 parameters. We fit the model to the ENIGH data using the method of moments, targeting observed labour 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 8.4.

5 Results

5.1 Model Fit to the Data

Figure 6 and Table 4 show the fit of the model with respect to the targeted moments. The figure shows that the dynamics predicted by the model for both log (male/female) relative earnings and log (male/female) relative supply follow closely what is observed in the ENIGH data. The model is not able to accurately capture short-term variations in these series, but it is able to replicate the long-term trends by occupation groups across the 25 year period. Importantly, the model also does a good job capturing the level and changes in labour force participation for both males and females.

Table 4 shows the observed and predicted occupation shares and mean wages for each type of agent in the model in two periods: C.1990 and C.2013. There is a close fit of the model with respect to the occupational distribution, with no major differences between model predictions and data in any of the cells. In the case of

²¹Wages: 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.

mean wages, the model predictions are also close to the data, except in the cases in which the occupational shares of the groups being compared are relatively small. 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.

Having established that the model can recreate the level and changes of the occupational and wage structure of the Mexican economy over the last 25 years, we now turn to the estimates of the structural parameters that give rise to these patterns.

5.2 Parameter Estimates

Table 5 presents the point estimates and standard errors of the different elasticities of substitution in the production technology. The elasticities of substitution between male and female labour are estimated to be around 1.2 in the low-paying manual and routine task-intensive occupations, and close to 2.6 in the high-paying abstract task-intensive occupations. These results provide support to the hypothesis that male and female labour are closer substitutes in occupations with stronger requirements for abstract analytical skills.

To get a sense of what these values represent in the Mexican context, note that the log (male/female) relative supply fell by 44 log points between 1989 and 2014 in both manual and routine task-intensive occupations, so an elasticity of 1.2 implies that the log (male/female) earnings ratio should have increased – holding everything else constant – by 36 log points (43%) in both cases. These numbers are much higher than the observed 3 and 11 log point increase of the gender earnings gap in the respective occupations, so that relative demand trends favouring women in those occupations must be part of the story. For abstract task-intensive occupations, a similar calculation shows that the 69 log points decline in the relative supply should have resulted in a larger earnings gap by close to 26 log points (30%), instead of the 10 log points decline observed in the data. Again, these results suggest that relative demand trends in high-paying occupations were strongly favourable to women.

To my knowledge, these are the first estimates of the elasticity of substitution between male and female labour for any country in Latin America, so there are no

points of comparison. For the case of the U.S., Acemoglu et al. (2004) estimate this elasticity to be around 3, while Johnson and Keane (2013) report values for this parameter in a range between 4.76 and 5.6. Although not directly comparable, the differences in the estimates between the U.S. and Mexico are substantial. It is possible that the technologies of production in less developed economies are such that the differences in the skills supplied in the labour market by men and women are accentuated (e.g. due to a lower technification of tasks and higher reliance on “brawn” instead of “brain”). Moreover, social norms might be such that women are restricted from entering certain occupations, so that the elasticity of substitution between men and women not only captures technical possibilities, but also cultural traits. It is not possible to disentangle between these and other alternative explanations, but this is a fertile area for future research.

On the educational dimension, the point estimates of the elasticities of substitution between skilled and unskilled labour are 1.4 in abstract, 1.6 in routine, and 3.6 in manual task-intensive occupations. The fact that the estimated elasticities of substitution between these groups are lower in occupations that are more intensive in cognitive than in physical labour is encouraging. Moreover, these results suggest that the sharp educational upgrading in Mexico has contributed to the fall of overall earnings inequality by generating downward pressures on wages of workers with more schooling, an outcome that is consistent with the recent inequality trends in the Latin American region.

The three point estimates of the elasticities of substitution between skilled and unskilled labour are in a range that incorporates the values that other studies have reported for Latin American. For example, using data from five Latin American countries during the 1990s, Manacorda et al. (2010) report estimates of this elasticity of around 3. Fernandez and Messina (2017) expanded the model used by Manacorda et al. (2010) finding a value close to 2.1. For the United States, there is a general convergence among different studies showing that the elasticity of substitution between skilled and unskilled labour is close to 1.5 (Katz and Murphy, 1992; Ciccone and Peri, 2005; Johnson and Keane, 2013).

The previous discussion shows that the rise of female labour force participation is having a substantial influence on the Mexican wage structure, but they also suggest that the demand for female labour must have increased considerably during the period, mitigating the compositional effects, otherwise we should have observed a much larger rise in the gender earnings gap. In Figure 7 we present the model

predictions of the evolution of both the relative demand trends and total labour productivity.²² To facilitate interpretation, each series capturing changes in relative demands is normalized to take a value of zero in 1989. Panels (a) and (b) show that during the last 25 years, the demand for female labour in Mexico increased relative to that of males. This result holds within the three occupational groups, and also among workers with low and high levels of schooling. For example, 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 tertiary 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. But the predictions of the model are consistent with previous studies that have shown that the Mexican labour market has experienced an increase in the relative demand for female labour. Juhn et al. (2014) show that as a consequence of the reduction of import tariffs, the more productive firms in the Mexican economy were induced to enter the export markets and modernize their technologies. The authors argue that these new technologies lowered the need for physically demanding skills in blue-collar occupations, which translated in a rise of the relative demand for female labour. A similar argument emphasizing how structural change can lower the need for physically intensive tasks, favouring female labour, has been recurrent in other contexts (Galor and Weil, 1996; Rendall, 2010; Black and Spitz-Oener, 2010; Lup Tick and Oaxaca, 2010; Akbulut, 2011; The World Bank, 2012; Aguayo-Tellez et al., 2013; Rendall, 2013). It is also possible that social attitudes regarding female labour are changing, and that employers are more willing to hire women in positions that were previously exclusive to men.

Relative demand trends are also favouring workers with more schooling and labour in abstract task-intensive occupations (see Panels (c) and (d)).²³ 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

²²See Equations (4.4) and (4.5) for an example of how the relative demands trends are specified.

²³Other studies showing that relative demand for skilled labour has been growing in Mexico include Sánchez-Páramo and Schady (2003); Binelli and Attanasio (2010); Manacorda et al. (2010); Gasparini and Lustig (2011)

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). Interestingly, the results provide evidence that non-neutral technical change – the most widely accepted explanation for the rise in income inequality in developed economies – is also a factor in developing countries, but that the educational upgrading of the workforce is strong enough to counteract this force, so we don’t observe such a sharp rise in the wages of high-skilled workers. Quoting the famous Tingbergian analogy, the race between education and technology is being won by the former. Finally, the evolution of total labour productivity provide a grim picture for the Mexican economy, with the estimated trends falling by as much as 27 log points between 1989 and 2014.

The estimates of the parameters from the supply side of the model are shown in Table 6. Each row reports the point estimate and standard error of a different parameter, and 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 labour type chooses a given market occupation with respect to the wage. we calculate this derivative in each year and for every possible labour type and occupation combination separately, and then take the average across all the values.

As expected, 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 labour 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 increase 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 there is substantial persistence in the part of workers when choosing an occupational group, and that there has to be very large changes in relative earnings across occupations for workers to switch between them.

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. This decline is associated with a rise in labour force participation of 3.3 percentage points, which is close to 17 percent of the total observed change. The percentage of females in the same educational group that are either married or with a permanent partner fell by 6 percentage points, and this decline is associated with a rise of labour force participation of 1 percentage point, close to 5 percent of the observed change. Together, the fertility transition and changes in marital status explain close to 22 percent of the overall rise in female labour force participation among unskilled females, the group for which we observe the largest movement into the workforce. The other 78 percent is explained by either changes in the demand side of the economy, or by changes in preferences for home production and home production technology as captured by the linear trends in Equation (4.7).

Among men with at most a high school degree, both marital status and fertility decisions influence the decision to participate in the labour market, but 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 labour force participation, although the average marginal effects are small. Interestingly, neither of the two variables is statistically significant in the case of college educated men and women.

5.3 Counterfactual Exercises

To quantify the impact of the different factors behind the decision to participate in the labour market on the occupation and wage structure, we do a series of counterfactual exercises. The idea is to “turn off” the different drivers, one at a time, and then compute the model predictions of occupational shares and mean wages, taking advantage of the general equilibrium set-up. In particular, we construct counterfactual distributions under four scenarios: (CF1) setting the linear, quadratic, and cubic coefficients of the demand trends (α shares) to zero; (CF2) setting the share of each group with children under the age of 5 to be equal to the level in 1989, and constant across the years; (CF3) setting the share of each group that is married or with a permanent partner to be equal to the level in 1989, and constant across the

years; and (CF4) setting the coefficient of the linear trends $(\pi_{2,f}, \pi_{2,k})$ in the home production utility function to zero.

Table 7 presents the results from these experiments, with each column corresponding to a different scenario and each cell reporting the difference between males and females of the difference between C.1990 and C.2013 of the occupational shares and wages, that is, the change over time in occupation and wage gaps. The results from the first counterfactual (CF1) reinforce our previous discussion about the role of changes in the demand side of the economy: relative demand trends in Mexico tended to favour female labour, and they helped mitigate the impact of the gender compositional change in labour supply. In the absence of these shifts in demand, the model predicts that the (male/female) gender earnings gap would be significantly higher in all occupations and skill groups. The differences in occupational choices between men and women were mostly unaffected by how the demand evolved, which is probably due to the fact that the wage elasticity of labour supply is low.

The second (CF2) and third counterfactuals (CF2) 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 meaningful impact on the occupational or wage structure. Even though we found that both variables were significant factors behind 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.

But the story is different when we look at the final counterfactual (CF4). The linear trends in the home production utility, which can be capturing either changes in preferences for this choice or improvements in the technology used in home production activities, have a strong impact on the occupational and wage structure. In particular, once we switch off the trends in the utility for home production, the rise in female labour force participation is much lower (6% predicted instead of the 22% observed). This in turn implies that the gender compositional shift in the workforce is small, 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, including in the manual and routine task-intensive occupations.

6 Robustness Exercises

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 8 reports the point estimates and standard errors of the different elasticities of substitution in the production technology using alternative earnings series and measures of labour supply. The table reports the baseline estimates; the estimates once we include income from part-time workers; and the estimates if we measure labour 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 labour 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 labour 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 labour supply is having on the wages of female workers in low-paying occupations. Using a labour 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 9. 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, labour 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 10 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 labour 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 11. Results are essentially unchanged compared to the baseline.

7 Conclusions

This paper studies the effect of the rapid rise of female labour force participation on the Mexican occupational and wage structure over the last 25 years. It argues that the increase in female labour 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.

The paper shows that the change in the gender earnings gap in Mexico over the last quarter century has not been uniform throughout the pay distribution. In particular, the gender earnings gap increased in low-paying manual and routine task-intensive occupations, but it fell in high-paying abstract task-intensive occupations. Using recently developed decomposition techniques, we provide evidence that the shift in relative earnings was not the result of a recomposition of the socio-

demographic characteristics of the workforce, but that it reflects actual changes in relative pay between men and women.

The paper develops a equilibrium model of the labour market to study if rising female labour supply can explain the peculiar evolution of the wage structure. The model follows the task-based approach, allowing the elasticity of substitution between male and female labour to vary depending on the task content of occupations. This feature implies that it is possible for changes in female labour supply to have heterogeneous impacts on relative earnings throughout the pay distribution.

we structurally estimate the model using Mexican household survey data covering the last 25 years. The estimates from the model support the hypothesis that the elasticity of substitution between male and female labour is lower in low-paying manual and routine task-intensive occupations, than in high-paying abstract task-intensive occupations. In particular, in the preferred specification, we find that these elasticities are between 1.2 in manual and routine task-intensive occupations, and 3 in abstract task-intensive occupations. The implication is that the movement of women into the labour market has imposed a large burden on the wage growth of female workers, especially those with low levels of schooling.

The estimates from the model also show that the demand for female labour relative to that of males has been increasing during the last 25 years in Mexico. we find that this was the case across the board: both within broadly defined occupational groups, and for workers with high and low levels of schooling. The counterfactual exercises show that these favourable trends in the demand side of the economy have attenuated the supply-driven downward pressure on female wages at low-paying occupations, and fully counteract it at 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., katz, L. and Kearney, M. (2006). The Polarization of the U.S. Labor Market. *American Economic Review* 96(2): pp. 189–194.
- 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. (2008). Trends in U.S. Wage Inequality: Revising the Revisionists. *The Review of Economics and Statistics* 90: 300–323.

- 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.
- 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.
- Brown, R., Moon, M. and Zoloth, B. (1980). Occupational Attainment and Segregation by Sex. *Industrial and Labor Relations Review* 33(4): pp. 506–517.
- 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.
- 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, R. (2013). Cultural Change as Learning: The Evolution of Female Labor Force Participation over a Century. *American Economic Review* 103: 472–500.
- Fernández, R., Fogli, 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.
- Ferreira, F., Leite, P. and Litchfield, J. (2008). The Rise and Fall of Brazilian Inequality. *Macroeconomic Dynamics* 12(2): pp. 199–230.
- 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.
- Firpo, S., Fortin, N. and Lemieux, T. (2011). Occupational Tasks and Changes in the Wage Structure. IZA Discussion Paper No. 5542.
- Fogli, A. and Veldkamp, L. (2011). Nature or Nurture? Learning and the Geography of Female Labor Force Participation. *Econometrica* 79: 1103–1138.

- Galor, O. and Weil, D. (1996). The Gender Gap, Fertility, and Growth. *The American Economic Review* 86(3): pp. 374–387.
- Gasparini, L., Cruces, G. and Tornarolli, L. (2011). Recent Trends in Income Inequality in Latin America. *Economía* 11(2): pp. 147–201.
- 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.
- Gathmann, C. and Schoenberg, U. (2010). How General Is Human Capital? A Task-Based Approach. *Journal of Labor Economics* 28: 1–49.
- 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.
- 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., Murphy, K. and Pierce, B. (1993). Wage Inequality and the Rise in Returns to Skill. *The Journal of Political Economy* 101(2): pp. 410–442.
- 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.
- Kahhat, J. (2010). Labor Earnings Inequality: The Demand for and Supply of Skills. In López-Calva, L. F. and Lustig, N. (eds), *Declining Inequality in Latin America: A Decade of Progress?*. Brookings Institution Press, chap. 2.
- 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.
- Lemieux, T. (2006). Increasing 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.
- 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 Economics* 107(1): pp. 285–326.
- 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.
- 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*. 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.
- U.S. Bureau of Labor Statistics (2017). Civilian Labor Force Participation Rate: Women [LNS11300002], Retrieved from FRED, Federal Reserve Bank of St. Louis; <https://fred.stlouisfed.org/series/LNS11300002>, May 2, 2017.
- 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.
- Welch, I. (2000). Herding among security analysts. *Journal of Financial Economics* 58: 369–396.

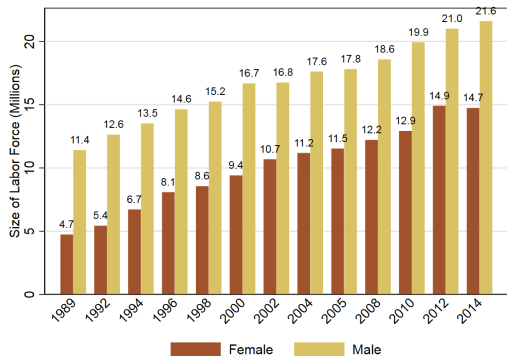
8 Appendix

8.1 Tables and Figures

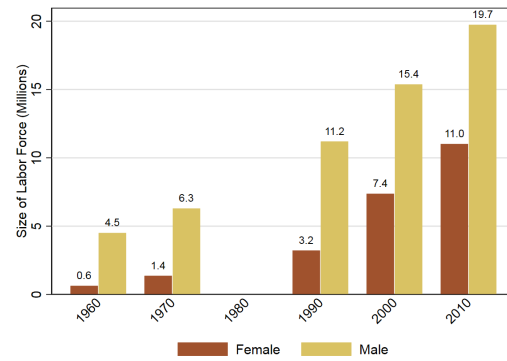
Figure 1: Labor Force Participation by Sex

Absolute Numbers

(a) ENIGH

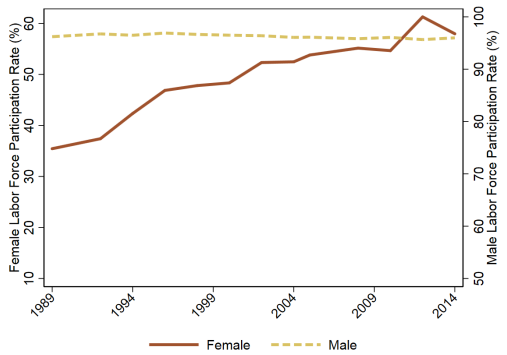


(b) CENSUS

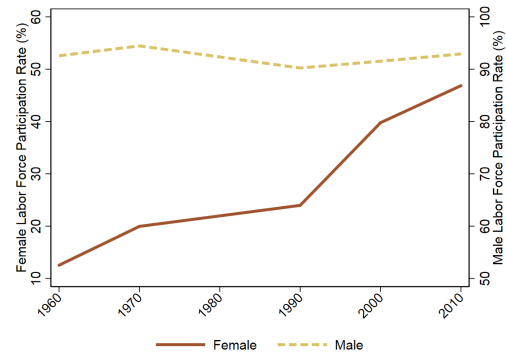


Participation Rates

(c) ENIGH

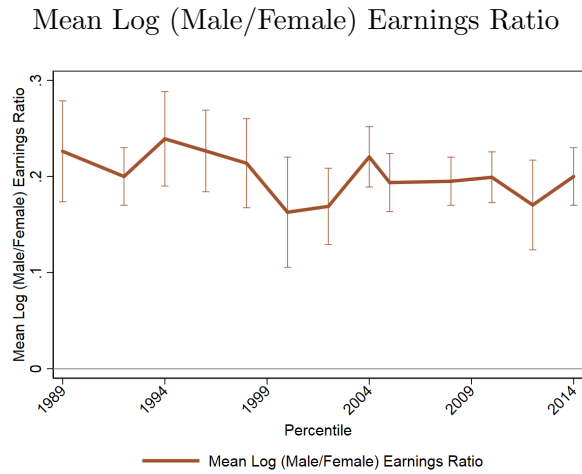


(d) CENSUS

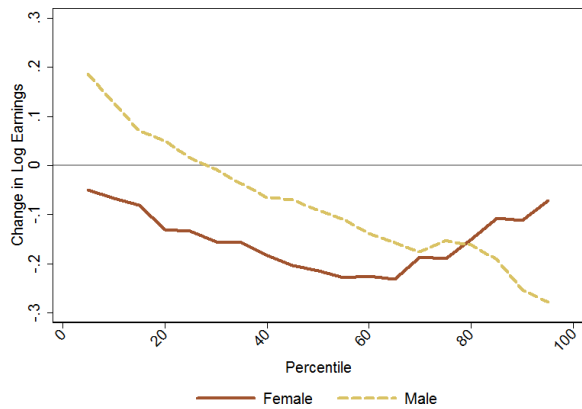


Notes: Panels (a) and (b) show the total number of prime-age population that are either working or actively searching for a job by sex. The series in Panels (c) and (d) show the share of prime-age males/females that are either working or actively searching for a job by sex. The differences between the 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.

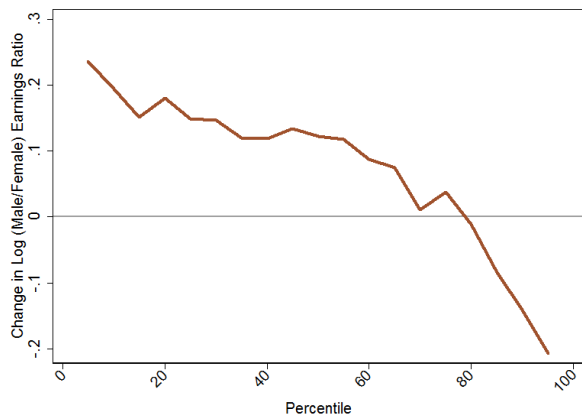
Figure 2: Changes in the Log (Male/Female) Earnings Ratio



Change in Log Hourly Earnings by Percentile: C.1990-C.2013



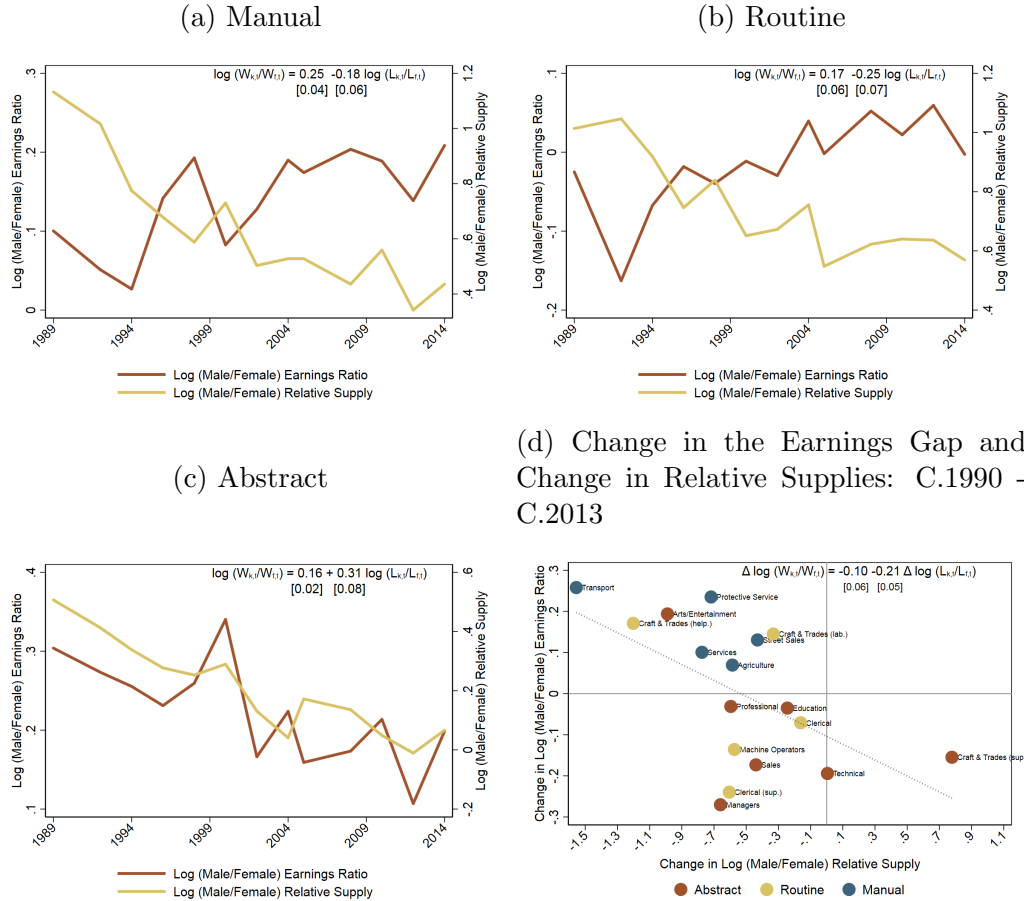
Change in Log (Male/Female) Earnings Ratio by Percentile: C.1990-C.2013



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 10 percent confidence intervals. Sample weights used in all calculations.

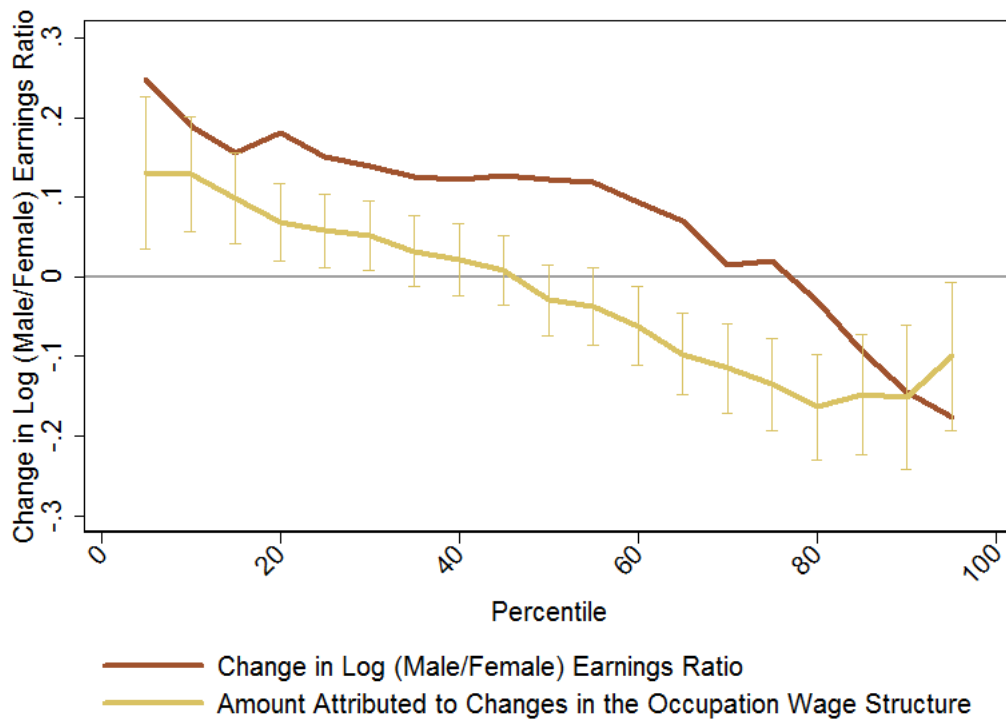
SUPPLY AND DEMAND BY OCCUPATIONS

Figure 3: Comovement Between Log (Male/Female) Earnings Ratio and Log (Male/Female) Relative Supply by Occupation Groups



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.

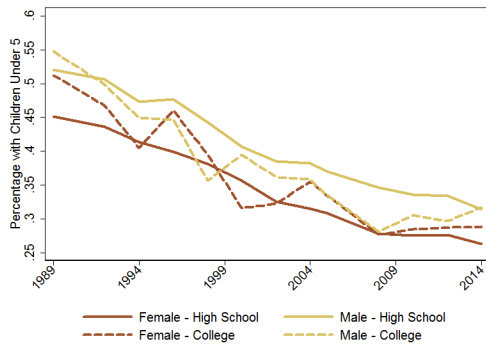
Figure 4: Decomposition Results: Change in Log (Male/Female) Earnings Ratio Between C.1990 and C.2013 Attributed to Changes in Occupational Wage Structure



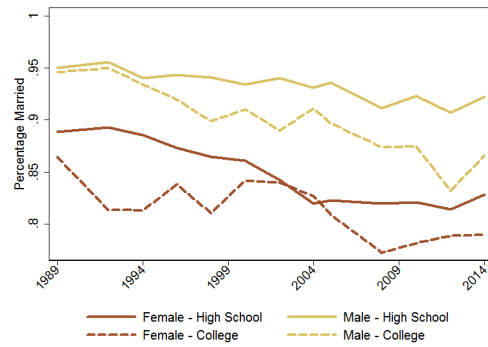
Notes: The occupation wage structure series depicts the unexplained component of the detailed Oaxaca-Blinder decomposition – defined in Equation (3.5) – that is attributed to the 18 principal group occupations of the ENIGH. The estimation is done separately for 19 percentiles. Confidence intervals are estimated via bootstrap with 500 replications. Sample weights used in all calculations.

Figure 5: Determinants of Movements into the Labor Market

(a) Changes in Fertility



(b) Changes in the Marriage Market

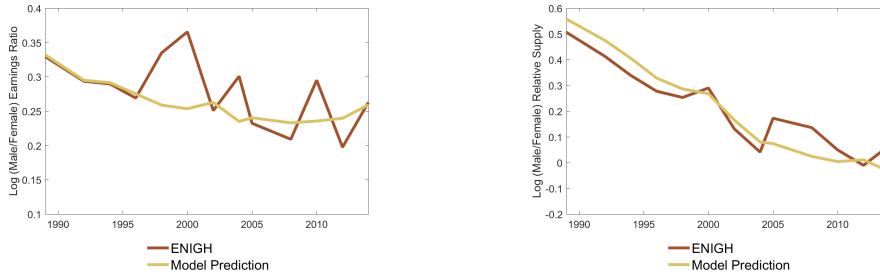


Notes: Panel (a) depicts the proportion of the sample with children under the age of 5, and Panel (b) depicts the proportion of the sample that is married, both conditional on education and gender. The ENIGH survey started asking the question on marital status to all member of the household in 1996, 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 6: Log (Male/Female) Relative Earnings, Log (Male/Female) Relative Supplies, and Participation Rates: ENIGH and Model Predictions

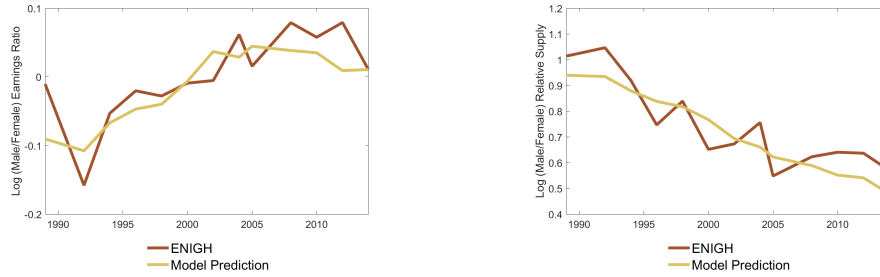
Analytical Occupations

(a) Log (Male/Female) Earnings Ratio (b) Log (Male/Female) Relative Supply



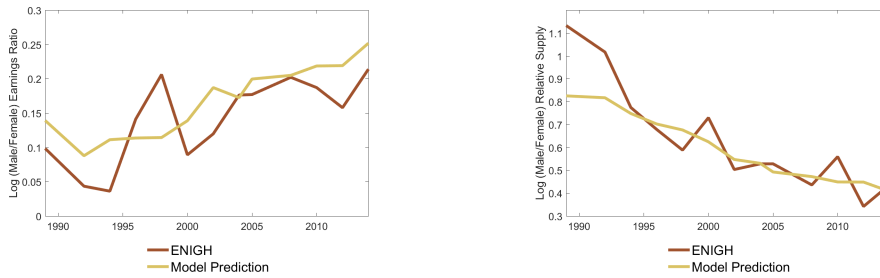
Routine Occupations

(c) Log (Male/Female) Earnings Ratio (d) Log (Male/Female) Relative Supply



Manual Occupations

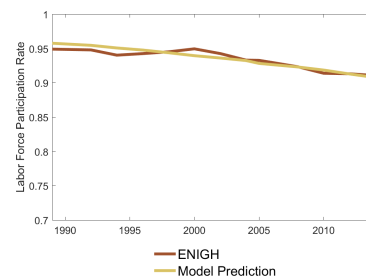
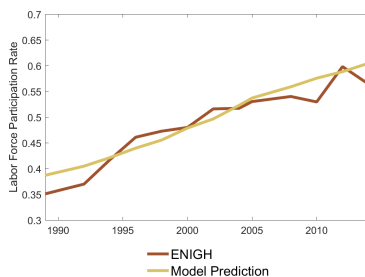
(e) Log (Male/Female) Earnings Ratio (f) Log (Male/Female) Relative Supply



Participation Rates

(g) Female

(h) Male

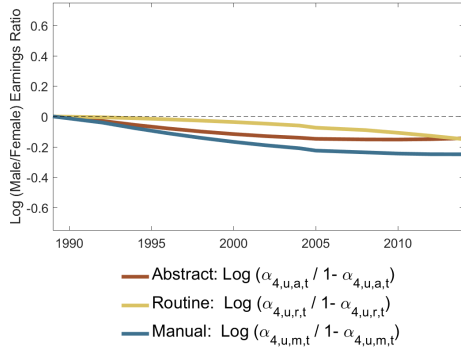


Notes: The different panels depict the series of (male/female) relative earnings, (male/female) relative supplies, and labor force participation rates, both from the ENIGH and predicted from the model.

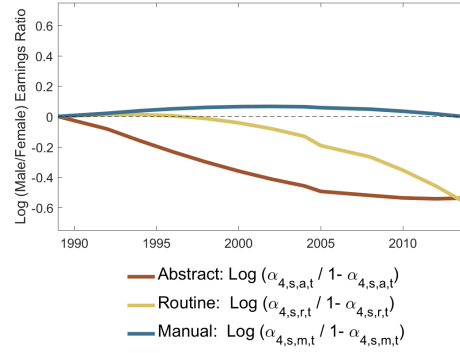
Figure 7: Estimates of the Relative Demand Indexes and Total Factor Productivity

Production Technology: Level III

(a) Male vs. Female (High School)

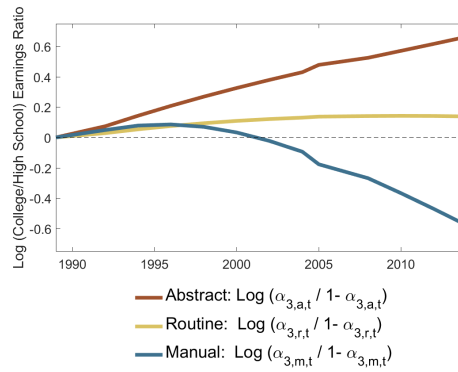


(b) Male vs. Female (College)



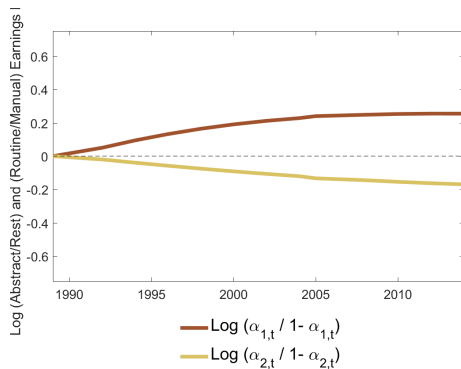
Production Technology: Level II

(c) College vs. High School

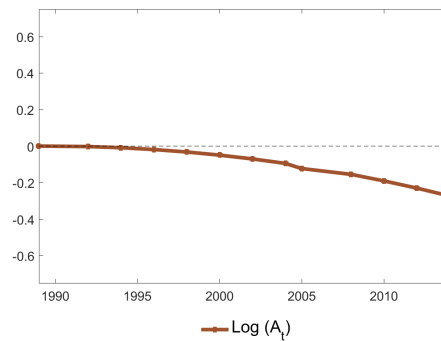


Production Technology: Level I

(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: Occupation Groups

ENIGH Principal Group	Median Percentile of the Task Measure			Group
	Abstract	Routine	Manual	
Managers	90.0	17.0	27.5	Abstract
Crafts and Trades (Supervisors)	84.0	42.0	62.0	Abstract
Education	83.0	11.0	65.0	Abstract
Professional	83.0	42.0	46.0	Abstract
Technical	71.0	69.0	43.0	Abstract
Arts/Entertainment	66.0	35.0	48.0	Abstract
Sales	61.0	22.5	15.0	Abstract
Crafts and Trades (Laborers)	40.0	82.0	73.0	Routine
Clerical (Supervisors)	61.0	63.0	51.5	Routine
Crafts and Trades (Helpers)	10.5	62.0	60.5	Routine
Machine Operators	16.0	62.0	51.0	Routine
Clerical (Laborers)	41.5	53.0	12.0	Routine
Transport	19.5	21.0	96.0	Manual
Agriculture	32.0	27.0	82.0	Manual
Protective Services	24.5	5.5	76.5	Manual
Domestic Service	9.0	8.0	76.0	Manual
Street Sales	38.0	13.0	64.0	Manual
Service	28.0	25.0	63.0	Manual

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, I 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.

Table 2: Changes in Composition between C.1990 and C.2013

	C.1990			C.2013			Dif. in Dif.
	Female	Male	Δ	Female	Male	Δ	
Overall Population							
<i>Participation Rate</i>	38.59	96.49	57.89	59.59	95.82	36.24	-21.66
Education							
<i>High School</i>	92.09	84.27	-7.82	81.24	79.05	-2.18	5.64
<i>College</i>	7.91	15.73	7.82	18.76	20.95	2.18	-5.64
Age							
<i>25-34</i>	44.63	43.61	-1.02	35.74	36.36	0.62	1.64
<i>35-44</i>	32.34	32.84	0.50	34.16	33.97	-0.20	-0.69
<i>≥ 45</i>	23.02	23.55	0.52	30.09	29.67	-0.42	-0.95
Workforce							
Education							
<i>High School</i>	85.48	84.37	-1.11	76.00	79.19	3.19	4.29
<i>College</i>	14.52	15.63	1.11	24.00	20.81	-3.19	-4.29
Age							
<i>25-34</i>	46.48	43.58	-2.90	34.96	36.03	1.08	3.97
<i>35-44</i>	33.31	33.38	0.07	36.18	34.72	-1.46	-1.53
<i>≥ 45</i>	20.21	23.04	2.83	28.87	29.25	0.39	-2.44
Occupations (Emp. Share by Gender (%))							
<i>Abstract</i>	38.19	26.27	-11.93	37.45	27.47	-9.98	1.95
<i>Routine</i>	27.37	33.35	5.98	27.22	35.43	8.21	2.23
<i>Manual</i>	34.44	40.39	5.95	35.32	37.09	1.77	-4.18
Occupations (Emp. Share Overall (%))							
<i>Abstract</i>	11.91	18.08	6.17	15.59	16.04	0.45	-5.72
<i>Routine</i>	8.53	22.95	14.42	11.33	20.69	9.36	-5.06
<i>Manual</i>	10.74	27.79	17.06	14.70	21.66	6.96	-10.10
Overall	31.18	68.82	37.65	41.62	58.38	16.77	-20.88

Notes: The table shows the composition by sex, education, age, and occupation of both overall prime-age population and of prime-age population in the workforce. Each cell reports the share of the respective column group. In the specific case of the overall employment by occupations, the cells report the shares including both males and females. Any individual that has at least completed one year of post-secondary education is included in the college group. We joined together surveys from 1989 and 1992, and from 2012 and 2014 to increase sample size of the ENIGH survey. The sample is restricted to prime-age population. Sample weights used in all calculations.

Table 3: Compositional Changes and The Gender Earnings Gap: Oaxaca-Blinder Decomposition Results by Selected Percentiles

	P10		P25		P50		P75		P90	
	Est.	[S.E.]	Est.	[S.E.]	Est.	[S.E.]	Est.	[S.E.]	Est.	[S.E.]
Observed Change	0.214	[0.034]	0.131	[0.020]	0.150	[0.023]	0.022	[0.025]	-0.097	[0.032]
Overall Wage Structure	0.155	[0.036]	0.108	[0.021]	0.148	[0.021]	0.006	[0.023]	-0.119	[0.033]
Occupation	0.129	[0.037]	0.058	[0.023]	-0.029	[0.023]	-0.135	[0.029]	-0.151	[0.046]
Education	0.007	[0.040]	0.021	[0.020]	0.019	[0.012]	0.004	[0.014]	0.052	[0.019]
Age	0.009	[0.014]	-0.004	[0.008]	0.002	[0.008]	0.029	[0.009]	0.049	[0.014]
Constant	0.009	[0.071]	0.033	[0.039]	0.157	[0.033]	0.108	[0.037]	-0.069	[0.059]
Overall Composition	0.059	[0.019]	0.023	[0.013]	0.002	[0.014]	0.016	[0.016]	0.022	[0.022]
Occupation	0.001	[0.002]	-0.001	[0.002]	-0.001	[0.004]	-0.003	[0.005]	-0.011	[0.006]
Education	0.079	[0.017]	0.037	[0.011]	0.022	[0.011]	0.040	[0.013]	0.055	[0.018]
Age	-0.021	[0.005]	-0.013	[0.003]	-0.019	[0.003]	-0.020	[0.004]	-0.022	[0.007]

Notes: The table shows results of the aggregate and detailed Oaxaca-Blinder decomposition defined in Equation (3.5). The Standard errors in brackets are calculated via bootstrap with 500 replications. Sample weights used in all calculations.

Table 4: Model Fit: ENIGH and Model Predictions of the Occupation Shares and Wages

	Female C.1990		Male C.1990		Female C.2013		Male C.2013		Dif. in Dif.	
	ENIGH	Model	ENIGH	Model	ENIGH	Model	ENIGH	Model	ENIGH	Model
Occupation Shares										
<i>College</i>										
<i>Abstract</i>	2.26	1.72	5.30	4.58	5.22	5.32	6.14	6.18	-2.13	-2.01
<i>Routine</i>	0.63	0.77	1.27	1.38	1.66	1.81	1.90	1.67	-0.39	-0.76
<i>Manual</i>	0.05	0.61	0.39	1.12	0.49	0.83	0.92	1.05	0.09	-0.29
<i>Home Production</i>	1.26	1.09	0.45	0.33	2.55	1.95	0.93	0.99	-0.81	-0.20
<i>High School</i>										
<i>Abstract</i>	5.45	5.36	6.39	6.77	6.27	6.44	5.69	5.43	-1.53	-2.42
<i>Routine</i>	4.89	5.64	13.58	14.63	6.69	7.00	13.34	12.98	-2.04	-3.00
<i>Manual</i>	6.90	7.34	17.59	16.47	10.35	10.08	15.04	15.68	-5.99	-3.53
<i>Home Production</i>	31.49	30.37	2.12	1.83	19.59	19.35	3.23	3.25	13.02	12.44
Mean Wages										
<i>College</i>										
<i>Abstract</i>	6.99	6.44	10.04	9.51	6.30	6.15	8.11	7.93	-9.39	-12.13
<i>Routine</i>	5.98	6.98	7.98	7.70	4.90	4.81	5.25	5.46	-15.99	2.06
<i>Manual</i>	3.17	8.06	5.11	6.37	3.16	2.31	3.30	2.45	-35.05	9.84
<i>High School</i>										
<i>Abstract</i>	3.89	3.58	4.54	4.26	2.67	2.68	3.12	3.29	0.20	2.63
<i>Routine</i>	3.33	3.14	3.06	2.91	2.41	2.44	2.52	2.60	9.74	11.94
<i>Manual</i>	1.89	1.78	2.14	2.09	1.73	1.74	2.18	2.24	10.52	9.55

Notes: The Table shows the occupation shares and average wages in C.1990 and C.2013, both from the ENIGH and predicted by the model. The last two columns report the difference between males and females of the difference between C.1990 and C.2013 of occupational shares and wages. In the case of wages the difference in difference corresponds to the percent change.

Table 5: Parameter Estimates: Production Technology

	Elasticities of Substitution		
	Estimate	[SE]	Implied Elasticity ($1/(1 - \rho)$)
Gender			
$\rho_{4,m}$: female, male (manual)	0.175	[0.181]	1.212
$\rho_{4,r}$: female, male (routine)	0.179	[0.129]	1.219
$\rho_{4,a}$: female, male (abstract)	0.622	[0.099]	2.646
Education			
$\rho_{3,m}$: college, high school (manual)	0.722	[0.067]	3.594
$\rho_{3,r}$: college, high school (routine)	0.355	[0.041]	1.549
$\rho_{3,a}$: college, high school (abstract)	0.276	[0.121]	1.382
Occupation			
ρ_1 : abstract, routine and manual	0.031	[0.094]	1.032
ρ_2 : routine, manual	-0.141	[0.183]	0.877

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

Table 6: 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. I 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 respect to the wage. I 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 7: Counterfactual Exercises

	Change in Occupation and Wage Gaps: Difference in Difference.					
	ENIGH	Model	CF1 (α 's)	CF2 (CHL5)	CF3: (MARR)	CF4: ($\pi_{2,f}, \pi_{2,k}$)
Occupation Shares						
<i>College</i>						
<i>Abstract</i>	-2.13	-2.01	-0.40	-2.06	-2.00	-0.95
<i>Routine</i>	-0.39	-0.76	-0.65	-0.76	-0.74	-0.46
<i>Manual</i>	0.09	-0.29	-0.34	-0.29	-0.29	-0.10
<i>Home Production</i>	-0.81	-0.20	-0.35	-0.14	-0.22	-1.75
<i>High School</i>						
<i>Abstract</i>	-1.53	-2.42	-2.09	-2.06	-2.30	-0.48
<i>Routine</i>	-2.04	-3.00	-2.94	-2.62	-2.88	-0.88
<i>Manual</i>	-5.99	-3.53	-3.75	-2.96	-3.34	-0.40
<i>Home Production</i>	13.02	12.44	12.40	11.13	12.01	5.26
Mean Wages (% Change)						
<i>College</i>						
<i>Abstract</i>	-9.39	-12.13	20.01	-11.81	-12.13	-18.47
<i>Routine</i>	-15.99	2.06	34.51	2.06	1.68	-7.58
<i>Manual</i>	-35.05	9.84	18.45	10.04	9.79	3.87
<i>High School</i>						
<i>Abstract</i>	0.20	2.63	10.94	0.97	2.07	-7.91
<i>Routine</i>	9.74	11.94	19.78	8.50	10.79	-9.94
<i>Manual</i>	10.52	9.55	21.78	5.03	8.03	-19.37

Notes: The Table shows the difference between males and females of the difference between C.1990 and C.2013 of the occupational shares and wages 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 8: 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 9: 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 respect 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 10: 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 11: 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 respect 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.

8.2 Division of Occupations into Manual, Routine, and Abstract Task-Intensive Groups

Following the task-based framework, 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. (2003) 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.²⁴

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 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 1).

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 frame-

²⁴See 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).

work. 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.

8.3 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 τ th quantile of the distribution of wages, expressed in terms of the cumulative distribution $F_W(w)$. Define the following mixture distribution:

$$G_{W,\epsilon} = (1 - \epsilon)F_W + \epsilon H_W \quad \text{for } 0 \leq \epsilon \leq 1 \quad (8.1)$$

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:

$$IF(w; q_\tau, F_W) = \lim_{\epsilon \rightarrow 0} \frac{q_\tau(G_{W,\epsilon}) - q_\tau(F_W)}{\epsilon} \quad (8.2)$$

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.²⁵ The influence function in the case of the τ th quantile takes the form:

$$IF(w; q_\tau, F_W) = \frac{\tau - \mathbb{1}[w \leq q_\tau]}{f_W(q_\tau)} \quad (8.3)$$

where $\mathbb{1}[\cdot]$ is an indicator function and f_W is the PDF.²⁶ 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.

$$\int_{-\infty}^{+\infty} IF(w; q_\tau, F_W) dF(w) = 0 \quad (8.4)$$

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).

$$RIF(w; q_\tau, F_W) = q_\tau + IF(w; q_\tau, F_W) \quad (8.5)$$

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.

$$E[RIF(w_t; q_\tau, F_{W,t}|X_t)] = X_t' \gamma_t \quad (8.6)$$

Moreover, if we apply the law of iterated expectations to Equation 8.6, 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 (8.6).

²⁵Essama-Nssah and Lambert (2011) provides a comprehensive list of influence functions for different distributional statistics.

²⁶Note 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.

Estimation of Equation (8.6) 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 $\hat{\gamma}_t$ 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).

8.4 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 and supply predictions of the model as function of the parameters. Finally, let q be the observed vector of wages and labour 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,

$$e_i = q_i - p_i(\Theta^*), \quad (8.7)$$

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

The log-likelihood function takes the form

$$\log \mathcal{L}(\Theta) = \sum_i \log f(e_i) = \sum_i \log f(q_i - p_i(\Theta)), \quad (8.8)$$

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

$$s(\Theta) = \frac{\partial \log \mathcal{L}(\Theta)}{\partial \Theta} = \sum_i \frac{\partial \log f(q_i - p_i(\Theta))}{\partial \Theta} = \sum_i \frac{1}{\sigma_i^2} \frac{\partial p_i(\Theta)}{\partial \Theta} (q_i - p_i(\Theta)), \quad (8.9)$$

which we can write more compactly in vector form as

$$s(\Theta) = W'(\Theta)(q - p(\Theta)). \quad (8.10)$$

Here, $W(\Theta)$ is 364×81 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:

$$s(\hat{\Theta}_{ml}) = W'(\hat{\Theta}_{ml})(q - p(\hat{\Theta}_{ml})) = 0. \quad (8.11)$$

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

$$g(\hat{\Theta}_{gmm}) = W'(q - p(\hat{\Theta}_{gmm})) = 0, \quad (8.12)$$

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 (8.10). 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 an 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 $\frac{\partial p_i(\Theta_0)}{\partial \Theta_0}$. The estimates of the variance of each error, $\hat{\sigma}_{i,0}^2$, 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.²⁷ The estimates obtained after this first iteration²⁸ are used to update the weight matrix, and the process continues until the parameter vector converges to a stable point.

²⁷The parameter search is done using the interior-point algorithm in Matlab.

²⁸Note that even though the weight matrix is a function of the parameters, it remains fixed during the parameter search.

Since it is usually not possible to satisfy Equation (8.12), we estimate the parameters of the model using the quadratic form:

$$\hat{\Theta}_{gmm} = \operatorname{argmin}[q - p(\Theta)]'W(\Theta)W'(\Theta)[q - p(\Theta)]. \quad (8.13)$$

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 i th row corresponds to:

$$\Gamma_i(\hat{\Theta}_{gmm}) = \frac{\partial \bar{m}_i(\hat{\Theta}_{gmm})}{\partial \hat{\Theta}_{gmm}}, \quad (8.14)$$

so the variance-covariance matrix can be calculated using:

$$\hat{V}ar(\hat{\Theta}_{GMM}) = [\Gamma(\hat{\Theta}_{gmm})'\hat{V}ar[\bar{m}(\Theta_{gmm})]^{-1}\Gamma(\hat{\Theta}_{gmm})]^{-1}. \quad (8.15)$$