Gender Barriers, Structural Transformation, and Economic Development*

Gaurav Chiplunkar
University of Virginia

Tatjana Kleineberg
World Bank

June 2023

Abstract

Using nationally representative data across multiple countries spanning over four decades, we document distinct gender patterns in the process of structural transformation. Despite improvements, gender employment and wage gaps continue to persist even today, even in high-income countries. Interpreted through the lens of a general equilibrium Roy model with heterogeneous agents, we find that a large share of the observed changes in sectoral employment and output are driven by a decline in non-economic gender barriers like gender norms and wage discrimination. Effects are heterogeneous across sectors and countries.

Keywords: Gender, Structural Transformation, Economic Development

*We would like to thank Michael Peters, and participants in conferences and seminars at the World Bank, Northwestern University, Boston University, University of Virginia, CEPR STEG Annual Workshop, NEUDC, ASSA, SED, 100 Years of Economic Development Conference at Cornell University, and the Washington Area Development Economics Symposium for helpful suggestions and comments. Anushka Chawla, Sebastian Cubillos, and Facundo Ulivarri provided excellent research assistance. Chiplunkar: Darden School of Business, University of Virginia, 100, Darden Blvd., Charlottesville, VA 22905. Email: ChiplunkarG@darden.virginia.edu. Kleineberg: Development Research Group, World Bank, 1818 H Street NW, Washington, DC 20433. Email: tkleineberg@worldbank.org.

1 Introduction

Many countries, at different levels of economic development, have witnessed a notable shift towards greater gender parity in employment choices over the last few decades. For example, women have achieved more equitable representation in professional occupations and in the service sector. However, economic progress has not been able to eliminate gender disparities for other indicators, such as wages. Despite substantial reductions over the last decades, gender wage gaps continue to persist even in the most developed countries today. Examining the implications of these gender differences in the labor market are important to further understand their impact on the aggregate economy, productivity, and economic development. For example, Hsieh, Hurst, Jones and Klenow (2019) show that a reduction in gender norms and gender wage discrimination over time substantially improved the allocation of talent in the U.S. labor market, and can explain 20-40% of the aggregate growth in the U.S. since 1960. A key distinction that makes this exercise important to study from the broader perspective of economic development is its implications on the process of structural transformation i.e., its impact on the reallocation of workers (and economic activity) away from agriculture and towards manufacturing and services as countries develop.

Taking this as the starting point, this paper begins by documenting stylized facts on the gender dimension of structural transformation across multiple (91) countries and over a long period of time spanning over four decades (1970-2015). We find that with economic development, men follow the standard employment transitions of structural transformation– moving from agriculture into manufacturing and services. Women on the other hand, follow a very different pattern: at low levels of economic development (GDP per capita), they move out of agriculture, but exit the labor force all together. Women re-enter the workforce at higher levels of development, working mostly in the services sector. The share of women working in the manufacturing sector, unlike their male counterparts, is small and constant across all levels of economic development. Furthermore, there are important gender differences across occupations. For example, women are relatively more likely to work as clerks, secretaries, cashiers, and librarians, as opposed to the more managerial and professional occupations held by men.

Studying gender-specific transitions across both occupations and sectors is important. For example, we compute the Theil index, which (similar to the Gini index) measures the extent of segregation (gender, in our case) in an economy when compared to an egalitarian allocation. Across the 91 countries in our sample, we find that gender segregation across all occupation-sector pairs follows an inverted-U shape pattern with economic development.

This is perhaps unsurprising, given that most individuals of both gender work in agriculture in poor countries. With economic development, women transition out of the labor force, while men move into manufacturing and services, which worsens gender segregation in middle-income countries. At high levels of economic development, the Theil index declines again as both men and women work mostly in the service sector and sectoral employment choices converge across gender. A key advantage of the Theil index is that it can be decomposed into the share of segregation that is due to across-sector variation as compared to the part that is due to within-sector (across-occupations) variation. We find that the latter accounts for 40% of the overall gender segregation in low- and middle-income countries, and more than 60% of the overall gender segregation in high-income countries. Put together, this suggests that incorporating the occupation dimension along with the sectoral one is important to be able to understand the allocation of talent, and hence its impact on aggregate productivity and income in the economy. Apart from an egalitarian point of view, understanding sectoral and occupational choices is important from a policy perspective as well. For example, while men and women both move in to service sector jobs, women are over-represented in clerk occupations (such as secretaries, librarians, clerks, and cashiers), while men are more likely to work in managerial, professional and technical jobs. Understanding the drivers behind these choices (both economic and social) can help design more robust and inclusive policies.

After documenting the gender-specific patterns of structural transformation more broadly, we then then focus our attention on a smaller, but non-trivial group of six countries (India, Indonesia, Brazil, Mexico, Canada, and U.S.), where we have good measures of hourly wages at the gender-occupation-sector level, data on sectoral value-added, etc. over this time period. We then study the prevalence and evolution of gender gaps in employment and wages across sectors and occupations, and over the development process. Employment gaps i.e., the probability that a women works in an occupation-sector as compared to a man, tends to improve (become more equal) with economic development. On the other hand, gender wage gaps i.e., the ratio of female to male hourly wages, are strikingly similar in both poor and rich countries. Both gaps close over time for most countries, but sizable gender inequality still persists in all countries even today. A good gender balance in employment can co-exist with large gender pay gaps, which is most salient for example, in professional occupations and to a lesser degree within the service sector.

Several economic and "non-economic" mechanisms can explain the prevalence of these patterns across countries, and their evolution over time. For example, potential economic drivers could comprise of sector-, occupation-, or skill-biased technological change, or an expansion in education, or income effects that change consumption and sectoral expenditure shares, or the occupation mix in the sectoral production function, etc. Gender-specific gaps in wages and employment choices can on the other hand, also be driven by "non-economic" channels such as restrictive gender norms, or wage discrimination in the labor market, which would make it more or less attractive for women (relative to men) to work in certain occupations and/or sectors. Understanding the importance of both sets of channels is crucial to gaining insights on the drivers behind these gender-specific transitions in employment, and its resulting implications on wages, productivity and welfare.

To make progress, we develop a theoretical framework that incorporates both these channels and provides a parsimonious way of studying their implications on the sectoral and aggregate economy. We build the model in a way that can later in the calibration, flexibly allow for differences in the fundamental parameters across countries, as well as their evolution within a country over time. Each individual in our model is indexed by their gender and ability and in the spirit of a Roy model, choose: (i) whether to participate in the labor force or not; and (ii) conditional on working, choose their occupation and sector of work. Several key economic channels discussed above influence these individual decisions. For example, production technology is sector-specific, and the human capital units used as an input from each occupation in the sectoral production function varies across sectors. We allow for gender-specific ability distributions for each country-year, as well as gender-specific returns to ability in each occupation-sector (that proxies for gender-specific comparative advantage). On the demand side, we relax the standard assumption of homothetic preferences. A key implication of homothetic preferences are constant expenditure shares, which is inconsistent with empirical patterns of economic development (Alder et al., 2022; Fan et al., 2021). Of relevance to us is the fact that consumption baskets tend to skew more towards the services in richer countries (as compared to poorer ones) and women tend to have a comparative advantage in services.

Apart from these economic channels, we model gender-specific “non-economic” channels or barriers that can distort workers’ employment choices, thus generating talent misallocation in the economy, and potentially lowering productivity and growth. In particular, we consider two channels, namely “gender norms” and “wage discrimination”. We model gender norms as a utility cost incurred by individuals, both men and women, for working in a given occupation and sector. This cost captures general occupation-sector-specific amenities, or entry costs that are not directly reflected in the wage. We interpret differences in these utility costs between men and women as normative gender barriers and refer to them as "gender norms". Wage discrimination on the other hand, is modeled as a difference between the wage received by a woman compared to a man, after accounting for differences in their respective marginal product in an occupation-sector.

Using rich data spanning multiple decades across the six countries in our sample (34 country-years), we then take the model to the data. Our calibration strategy follows in the spirit of a growth accounting exercise, where we impose minimal restrictions on the fundamental economic parameters i.e., we allow them to evolve independently across countries and over time, and allow them to be positive or negative (i.e., we do not impose that one gender faces larger distortions or barriers than the other). The aim of the calibration exercise is twofold: first, to quantify the magnitude of (and change in) the gender norms and wage discrimination faced by women across countries and over time; and second, decomposing the empirical employment transitions, structural transformation, and economic growth in these countries that can be attributed to a reduction in the non-economic gender barriers.

Turning to the results, we find that middle- and high-income countries (Brazil, Mexico, Canada, and USA) had lower gender norms in the beginning of our sample period (1970s) as compared to low-income ones (India and Indonesia). Moreover, norms are lower in the service sector and in clerk and professional occupations relative to other jobs in the middle- and high-income countries, while the opposite is true for low-income ones. Lastly, middle- and high-income countries have also experienced a larger reduction in gender barriers over the last four decades, while gender norms have not declined much in India and Indonesia.

Turning to wage discrimination, we find that most countries had high wage discrimination in the 1970s, paying women only between 40-60% of their marginal product. While Wage discrimination has decreased over time in most occupations and sectors, it continues to be substantive even in recent decades (2010 onwards) in many occupations and sectors. In Brazil, Mexico, Canada and the US for example, women face more wage discrimination in professional occupations (where they are paid between 65-75% of their marginal value) as opposed to India and Indonesia, where discrimination is highest in agriculture. Interestingly, unlike gender norms, levels in female wage discrimination are strikingly similar across countries irrespective of their economic development i.e., today's developed countries do not perform better in fair remuneration for women than today's emerging countries.

Lastly, we turn to examining the importance of these non-economic gender barriers in explaining the growth trajectories of these countries over time. To do so, we hold wage discrimination τ_{oj} and gender norms ΔA_{oj} fixed at the values of each country's initial year, while allowing all other (economic) parameters to evolve according to our data and calibration. To quantify the importance of changes gender barriers, we then use the difference between each country's observed path in the data and the simulated paths from the counterfactual. We find that changes in gender barriers have large effects on sectoral output and employment growth and on the rise in female labor force participation. The magnitude of the effects varies across sectors with the smallest effects on agriculture and the largest ones on the service sector. Finally, a reduction in gender barriers accounts for around 25-30% of the growth in real value added (output) in middle-income countries–Brazil (29%) and Mexico (24%), as well as high-income ones–Canada (30%) and USA (25%). These barriers account for 17% of the growth in real value added in Indonesia, and only 4% in India.

The rest of the paper is organized as follows: Section 2 discusses our contribution to the literature, Section 3 describes the data and Section 4 presents the stylized facts. Section 5 builds the theoretical model and Section 6 describes the identification and calibration exercise. Sections 7 and 8 discuss the results from the calibration and counterfactual simulations respectively, and Section 9 offers a short conclusion.

2 Literature

Our paper links several branches of the literature that study structural transformation, occupational choices, human capital allocation, and the importance of gender roles in driving economic choices and outcomes. A large literature studies the reallocation of workers from agriculture to manufacturing and then to services. A recent literature, primarily in high-income countries, focuses on the role of gender in driving structural change. (Cuberes and Teignier, 2014, 2016; Moro, Moslehi and Tanaka, 2017; Olivetti, 2013; Ngai and Petrongolo, 2017; Rendall, 2018). We contribute to this literature by documenting gender-specific employment choices and wages across many countries and times periods, and by focusing on occupations and sectors.

Our paper is closest to Lee (2020) and Gottlieb, Gollin, Doss and Poschke (2021). Lee (2020) studies the effects of gender barriers on cross-country differences in agricultural productivity and finds that low-income countries have higher frictions for women in non-agricultural employment. Setting these frictions to US levels increases labor productivity by 21.3 percent and GDP per-capita by 3.6 percent. Gottlieb, Gollin, Doss and Poschke (2021) build a two-sector model of heterogeneous households where individuals differ by their marital status and gender. Our analysis on the other hand, studies the effects of gender barriers on occupational and sectoral choices, the (mis-)allocation of talent, and its macroeconomic implications. Another relevant paper is Hsieh, Hurst, Jones and Klenow (2019), who show that a reduction in gender barriers in the United States can explain 20-40 percent of the country's overall growth from 1960 to 2010.

Our analytical approach builds on Roy (1951) and assumes that workers have preferences for each occupation and they select into occupations based on their comparative advantage. Gender barriers in certain occupations or sectors (due to social norms or discrimination) distort the occupation-sector choices based on individuals' comparative advantage, which can affect the allocation of talent, as well as the structure and productivity of the aggregate economy. We model a home sector to which we attribute individuals who do not participate in the formal labor force. This is a particularly important for our focus on gender since female and male employment evolves differently as countries grow richer. Our exercise is similar in spirit to Hsieh, Hurst, Jones and Klenow (2019), who show that a reduction in gender barriers in the United States can explain 20-40 percent of the country's overall growth from 1960 to 2010.

Another strand of the literature measures and studies women's LFP which discusses the U-shaped pattern of female LFP over countries' development process. Goldin (1994) shows this pattern for the United States, and Heath and Jayachandran (2016); Fletcher et al. (2017); Mammen and Paxson (2000); Psacharopoulos and Tzannatos (1989) show it in other countries. We expand on these empirical facts by distinguishing between home and market sectors (agriculture, manufacturing, services), and several occupations within each market sector to relate these patterns directly to countries' process of structural transformation. Our model measures gender norms and wage discrimination for each occupation-sector which allows us to simulate counterfactuals that quantify the importance of gender on countries' process of structural transformation.

3 Data

Data Sources and Sample of Countries

Our primary data source is the The Integrated Public Use Microdata Series (IPUMS International, 2020) that harmonizes individual-level data on education and employment variables from nationally representative censuses, household and labor force surveys for many countries and over time. We use a sample of 91 countries and 273 country-years ranging from 1960-2018 to document how occupational and sectoral employment changes for men and women along countries' development spectrum. On average, we have 3 rounds of data for each country. While there is better coverage across countries in recent decades (1990 onwards), coverage across countries in older decades is non-trivial as well (See Table A.1.1).

A key requirement for our quantitative exercise is the availability of high-quality data on hourly wages in each occupation-sector-gender category over a long period of time (1970-2015). We therefore focus our attention on six countries (34 country-years) for which this is available, namely: India, Indonesia, Mexico, Brazil, Canada, and the United States. Put together, this sample covers a wide range of the income spectrum and accounts for 25-30% of the world population. We also complement the IPUMS data with data on sectoral value-added which we obtain from the Economic Transformation Database (ETD). For India and Indonesia, we complement the IPUMS data with data from labor force surveys (PLFS for India and SAKERNAS for Indonesia) to extend the time coverage to the most recent years (2018). Appendix E provides an elaborate discussion on the data construction.

Classification of Sectors and Occupations

We aggregate the harmonized sector classifications into (a) Agriculture; (b) Manufacturing and (c) (Market) Services as shown in Table A.1.3 and discussed in Herrendorf, Rogerson and Valentinyi (2013) and Herrendorf and Schoellman (2018). We create a category "Home Sector" to which we attribute unemployed and inactive individuals. This classification follows a recent literature that examines the role of gender and occupation choices (Moro, Moslehi and Tanaka, 2017; Ngai and Petrongolo, 2017), and is consistent across countries (Bridgman et al., 2018). Occupations are classified with the 1 digit ISCO-88 occupation codes as reported in Table A.1.4. We aggregate the top 3 occupation codes (managers, professionals, and technicians) due to small sample sizes for some country years.

Our main analysis focuses on seven occupation categories: (1) Professionals, which include managers, technicians, senior officials, legislators, directors, etc.; (2) Clerks, which include secretaries, librarians, cashiers, etc.; (3) Service Workers, which include travel, housekeeping and personalcare workers, along with those in shop, sales and service jobs; (4) Skilled Agricultural Workers, which include those in subsistence and market-oriented agricultural production; (5) Crafts and Trades Workers such as builders, painters, blacksmiths, electricians, etc.; (6) Plant and Machine Operators such as those workers in mining, metal, glass, wood, etc.; and (7) Elementary Occupation Workers such as street vendors, domestic helpers, prters, manual laborers, etc. Some occupation-sector-gender categories are very sparsely populated, so that we limit the agricultural sector to two occupations: skilled agricultural workers and elementary occupations. The manufacturing and the service sector consist of six occupations as we attribute all "skilled agricultural workers" to the agricultural sectors. Home is modeled as a separate sector.

4 Empirical Evidence on Employment Transitions by Gender

4.1 Sectoral and Occupational Employment by Gender

We first document employment transitions for men and women across sectors and occupations across countries at various stages of economic development. We being by examining how the gender-specific employment shares in the three sectors (agriculture, manufacturing and services) change with log GDP per-capita. For men, we see the standard pattern of structural change: as countries grow richer, male workers transition from agriculture to manufacturing and then to services (Figure 1(a)). The share of men in the home sector is low and relatively constant across at different stages of economic development. On the contrary, the sectoral employment pattern is very different for women (Figure 1(b)). At low levels of economic development, women first leave agriculture, but mostly leave the labor force all together (sorting into the home sector). However, at higher income levels, women then re-enter the labor force, to work mostly into the service sector. Hence, female labor force participation (FLFP) follows a U-shape across countries' GDP per capita, a feature that is well documented in the literature.

Figures 1(c) and 1(d) then plot the gender-specific employment shares across occupations across countries at various stages of economic development (measured by their GDP per capita). The employment share (for both men and women) in agricultural jobs declines with economic development (similar to the agriculture sector discussed previously). Consequently, employment shares in other occupations, and in particular, professionals, clerks, craft workers, and machine operators increase. Men tend to be over-represented in craft and trade jobs, machine-operators (mostly manufacturing), and professional occupations, while women are over-represented in jobs related to clerks and the home sector.

The above patterns highlight the importance for studying gender-specific employment transitions across either sectors or occupations. We now briefly show that the nature of occupational transitions follow gender-specific patterns even within sectors as well (Figure 2). More specifically, from Figures 2(a) and 2(b) craft, trade and service workers (such as blacksmiths, plumbers, electricians, etc.) dominate manufacturing employment for both men and women in poor countries (75-80%). However, they are less important in rich countries–around 50% for men and only 20% for women. Consequently, employment share in machine operator, and professional and elementary occupations increase for men, whereas women are over-represented in clerk occupations (such as secretaries, librarians, cashiers, etc.), along with elementary and professional occupations.

As compared to the manufacturing sector, occupations are less concentrated in the service sector for poor countries (Figures 2(c) and 2(d)). Approximately 50% of men and women work in trade and service workers, 20% in professionals, 20% in elementary occupations, and 10% clerks. At higher income levels, the shares of service and elementary workers decrease while clerks and professionals increase. Within the service sector, women are heavily over-represented among clerk occupations, while men are over-represented among machine-operators, elementary workers, and professionals.

4.2 Gender Gaps in Employment and Wages

To directly compare gender differences in employment choices, we compute the relative probability of a woman, as compared to a man, to be employed in an occupation-sector (within a country-year) i.e., we calculate the ratio of the share of working-age women in each occupation-sector to that of working-age men in the same occupation-sector. Similarly, to compare wage gaps, we compute the female-to-male wage ratios, which is the ratio of the average hourly wage in each occupation-sector (within a country-year) earned by a woman as compared to a man. Note that in both cases, a ratio of 1 implies gender parity (in employment or wages), while a ratio less (more) than 1 indicates that men (women) work/earn more in that occupation-sector than women (men). Given that wage data is not available across all countries, we restrict our sample to six countries that will form our core sample, namely: India, Indonesia, Brazil, Mexico, Canada and the United States. Figure 3 then plots the employment and wage ratios against GDP per capita for these six countries. For ease of interpretation, we only report (and connect) the ratios in the first and last year for each country. This allows us to analyze differences across countries and changes within-country over time. As discussed in Section 3, we cover several decades between 1970-75 and 2010-18 for thse countries. The shortest time span is India (1983-2018) while US and Mexico have the longest coverage (1970-2015).

Figure 3(a) shows the gender employment ratio for sectors.[1] In all countries (except India), the service sector has the highest female-to-male employment ratio which also improves most over time. In Canada and the US–where FLFP is high–the service sector employs more

[1] The graph excludes the home sector, where the share of women is multiple times that of men.

women than men (ratio > 1). Manufacturing and agriculture have much lower female-to-male employment ratios and see little improvement over time. In these sectors, all countries still employ more than two men for every woman in the most recent year (ratio < 0.5).[2] India stands out with low female employment, particularly in the manufacturing and service sectors where women are only a third as likely to be employed as compared to men (ratios 0.3), which does not improve over time (1983-2018).

Figure 3(b) shows the gender employment ratio across occupations. In all countries (except India), the female-to-male employment ratios increase over time in almost all occupations. A main driver is the increase in FLFP that is particularly strong during our sample period for Indonesia, Mexico, and Brazil. However, large gender differences remain across occupations (in levels and trends). In the US, Canada, Mexico and Brazil, clerks and professionals have the highest female-to-male employment ratios which increase over time. In the US, Canada, Mexico, and Brazil more women than men work as clerks (in the most recent period).[3] For Canada and the US this also applies to clerks and professionals. Indonesia and India start with low female-to-male employment ratios in most occupations due to very low FLFP. India shows almost no improvement over time, while Indonesia shows a sharp increase in female employment–particularly among professionals and clerks.

Turning to wage gaps in Figure 3(c), the female-to-male wage ratios improve over time across all countries, with the notable exception being India again, which has also seen a drop in FLFP over the years. Despite these improvements, wage ratios remain below 1 in all sectors and all countries. Levels are strikingly similar across the poorest and richest countries. In Brazil, Mexico, and the US, the service sector has the lowest female-to-male wage ratios despite of having the highest employment ratio (as we discussed above).

Figure 3(d) shows gender wage ratios for occupations. In most occupations, the female-to-male wage ratios improve over time.[4] Despite improvements over time, female-to-male wage ratios remain below 1 for all occupations (except clerks in Canada). Wage ratios vary substantially across occupations but they are similar in poor and rich countries. Consider for example the wage ratio in professional occupations in the most recent sample year: For each

[2]The only exception is agriculture in Indonesia where the ratio is slightly above 0.5 but is constant over time.

[3]For readability, we omit the employment ratio for clerks for the US and Canada as these countries have up to three times more female than male clerks.

[4]Wage ratios for clerks and professionals in India and Indonesia are an exception as they start with relatively large female-to-male wage ratios at the beginning of the sample period and then remain constant. Other countries then catch-up to these higher levels and the wage ratios of clerks and professionals converge to similar levels across all countries.

dollar earned by a men, a woman earns roughly 74 cent in Brazil and the US and on average 78-80 cents in India, Indonesia, Mexico and Canada. In many countries, professional occupations have the lowest female-to-male wage ratio despite of having relatively high employment ratios. This finding indicates that employment and wage gaps do not necessarily close simultaneously in given occupations or sectors.

4.3 Gender Segregation Across and Within Sectors

Segregation indices offer an additional way of presenting gender segregation across occupations and sectors. One example is the Theil Index. The Theil Index measures an entropic "distance" the population is away from the "ideal" egalitarian state which equals an index value of 0. Higher numbers for the index indicate more segregation in the population of interest. The equations of the Theil index are shown in Appendix A.2. Figure 4(a) plots the global Theil Index for our full sample of 91 countries against GDP per capita. We see an inverted U-shape where gender segregation across occupation-sector-pairs first increases and then decreases with GDP. This is perhaps unsurprising, given that most individuals (irrespective of gender) work in agriculture in poor countries. The transition of women out of the labor force and men into manufacturing and services worsens segregation for middle-income countries, which declines for high-income countries as women re-enter the labor force in services. A key advantage of the Theil index is that it further be decomposed into the share of segregation that is due to across-sector variation as compared to the part that is due to within-sector (across-occupation) variation. Figure 4(b) plots the share of gender segregation that is explained by segregation across occupations within each sector. The within-sector component explains around 30-40% of the segregation in low- and middle-income countries, which increases to around 60% for high-income ones.[5] Put together, the above discussion underscores the importance of examining the role of gender in occupational choices, along with sectoral ones.

4.4 Summary of Empirical Results and Model Implications

We can summarize our empirical findings in three steps.

[5]In this analysis, changes in labor force participation are attributed to across-sector variation since we attribute individuals who are not in the labor force to the home sector (which has no occupations). To abstract from the extensive margin, we re-compute the Theil index conditional on labor force participation. For this, we find the same inverted U-shape for the global Theil index and the within-sector variation now becomes even more important, explaining between 40-70% of the overall gender segregation.

First, employment transitions across sectors and across occupations along countries' development process have salient gender patterns. For sectors, men follow the standard patterns of structural transformation. Women follow a different pattern. When they exit agriculture they first sort into the home sector and then enter into the service sector at higher income levels. Female employment in manufacturing remains small and relatively constant across income levels. For occupations, men enter more into crafts, trade, and machine-operating occupations along with the managerial and professional occupations. Clerk occupations grow fast and are very female-dominated.

Second, we find that gender gaps in employment and wages vary substantially across sectors and occupations. Employment gaps tend to be smaller in rich countries but wage gaps are strikingly similar in poor and rich countries. Both gaps improve over time for most countries, but sizable gender inequality still persists in the most recent period in all countries. A good gender balance in employment can co-exist with large gender pay gaps, which is most salient for professional occupations and to a lesser degree for the service sector. This empirical finding can be driven by wage discrimination or by sorting of workers based on their comparative advantage. Disentangling these channels is key to understanding the drivers behind these pattern–a feature that we will build into our theoretical framework. For example, we allow workers to differ in their productivity across sectors and occupations. Workers sort into occupation-sector-pairs based on their comparative advantage, so that small female employment shares imply that only the most productive woman sort into the occupation. As more woman enter the occupation, the employment share increases but average productivity and hence average wages can decrease.

Third, the global and within-sector Theil Index of segregation shows that variation across sectors and across-occupations within-sectors are both important in explaining the overall gender segregation in the economy. In particular, the within-sector segregation across occupations starts to dominate as countries move along the path of economic development. This is an important dimension that we incoroporate in our model.

The three observations discussed above point to the fact that several economic and "non-economic" mechanisms can affect employment choices and wage gaps between men and women during the process of economic development. Economic drivers include technological change, which for example, can be biased towards specific sectors or occupations. In addition, workers' effective human capital in sectors or occupations can change due to higher education or due to skill-biased technological change. As countries grow richer, income effects can also change the consumption basket, for example, by shifting expenditure away from agriculture and towards services (where women might have a comparative advantage). The occupational mix of workers in sectoral production might also change with economic development. On the other hand, there could also be "non-economic" drivers such as gender-specific amenities/barriers/wedges that can make it more or less attractive for women (relative to men) to work in certain occupations and/or sectors. In particular, we consider two such barriers that vary across occupation-sector-pairs: “wage discrimination” that allows employers to pay women only a fraction of their marginal product, and “gender norms”, which can be thought of as gender-specific preferences or gender-specific amenities/barriers to working in certain occupation-sectors (or in the labor force all together).

To decompose the observed empirical changes and quantify the importance of these "non-economic" channels, we develop (and estimate) a Roy model of occupational and sectoral choices, which incorporates each of the above channels. In particular, we allow gender barriers, productivity/skills, and returns to human capital to differ across occupations, sectors, countries, and over time. Gender barriers distort workers’ occupational and sectoral sorting based on their comparative advantage which generates a misallocation of talent and lowers productivity and growth. We flexibly estimate the model to fit several key data moments separately for each country-year. We identify gender barriers for each occupation-sector as residuals/wedges to match the observed employment and wage gaps after accounting for the above-mentioned economic factors. In the spirit of a growth accounting exercise, we then evaluate counterfactuals that quantify how much changes in gender barriers contributed to each country’s employment transition, output growth, and welfare.

5 Model

We now describe the model set up, solve for individuals’ employment choices and firms’ production decisions, and define the equilibrium.

5.1 Setup and Preferences

Model Setup: The economy consists of occupations o, sectors $j = (A, M, ms, hs)$ (agriculture, manufacturing, market services, and home services), and a mass N_g of male and female individuals which we denote by $g = \{f, m\}$. Each individual of gender g has an ability z and chooses to work in an occupation-sector-pair oj with the “home sector" being one possible choice.

Non-homothetic Preferences: Individuals have non-homothetic preferences in the PIGL class over agriculture, manufacturing and services $j = \{A, M, S\}$, where services are a composite of home and market services. Non-homothetic preferences imply that sectoral expenditure shares change with income, which can be an important driver of structural transformation that has been extensively studied in the literature (Herrendorf, Rogerson and Valentinyi, 2013; Alder, Boppart and Müller, 2022; Comin, Lashkari and Mestieri, 2021). We can represent PIGL preferences with an indirect utility function:

$$V(I_{ojgz}, p) = \frac{1}{\eta}\Big[\frac{I_{ojgz}}{P}\Big]^{\eta} - D(p), \tag{1}$$

where $I_{ojg}(z)$ is the income earned by an individual of gender g and ability z who works in an occupation-sector pair oj and p is a vector of sectoral prices. P and $D(p)$ are homogeneous of degree zero and one respectively and are parameterized as:

$$P = \prod_{j=\{A,M,S\}} p_j^{\omega_j} \quad \text{and} \quad D(p) = \sum_{j=\{A,M,S\}} \nu_j \ln p_j$$
$$\text{where:} \sum_j \omega_j = 1 \quad \text{and} \quad \sum_j \nu_j = 0$$

Roy's Identity implies that individuals' sectoral expenditure shares are given by:

$$\varphi_j(I_{ojg}(z), p) = \omega_j + \nu_j \left[\frac{I_{ojgz}}{P}\right]^{-\eta}, \tag{2}$$

as we prove in Appendix C.1. Equation (2) highlights how income and prices affect sectoral expenditure shares. ω_j denotes the sectoral expenditure shares in the limit when real income approaches infinity. ν_j captures the income effect: an expenditure share increases with income if $\nu_j < 0$ (as is the case for services) and decreases with income if $\nu_j > 0$ (as is the case for agriculture). Preferences are homothetic if $\nu_j = 0$.

Preferences for Home and Market Services: Similar to Ngai and Petrongolo (2017), we assume that services are a CES composite of home and market services $s = \{hs, ms\}$ given by:

$$C_S = \Big[\sum_{s' \in \{hs, ms\}} \alpha_{s'} C_{s'}^{\frac{\eta_s - 1}{\eta_s}}\Big]^{\frac{\eta_s}{\eta_s - 1}},$$

where η_s is the elasticity of substitution between home and market services and α_s are the preference weights across home and market services with $\sum_{s'} \alpha_{s'} = 1$. Expenditure shares for home and market services are given by the standard CES formulation:

$$\varphi_{s'} = \alpha_{s'} \left[\frac{p_{s'}}{P_S} \right]^{(1-\eta_s)} \times \varphi_S, \tag{3}$$

where φ_S is the expenditure share on services and $P_S = \left[\sum_{s'} \alpha_{s'} p_{s'}^{1-\eta_s} \right]^{\frac{1}{1-\eta_s}}$ is the CES price index. The proof is provided in Appendix C.2.

5.2 Income and Occupational Choice

We now define and solve workers' occupational and sectoral choice problem. Utility of a worker i of gender g and ability z who works in an occupation-sector pair oj is given by:

$$U^i_{ojg}(z) = V(I^i_{ojg}(z), p) - A_{ojg} + \varepsilon^i_{oj},$$

where $V(I^i_{ojg}(z), p)$ is workers' indirect utility as defined in Equation (1). A_{ojg} are gender-specific utility costs of working in an occupation-sector pair which capture a wide range of factors including amenities, preferences, norms or entry costs that can vary across occupation-sector pairs oj and across genders g. To disentangle the gender-specific component, we compute the difference in utility costs between men and women: $\triangle A_{oj} = A_{ojf} - A_{ojm}$. We refer to this difference as "gender norms" as it captures the additional cost that women incur when working in an occupation-sector oj relative to men.

ε^i_{oj} are idiosyncratic preference shocks for working in each occupation-sector pair.

Assumption 1: We assume that preference shocks ε^i_{oj} are extreme value Gumbel distributed across occupation-sector-pairs with a dispersion parameter σ_ε.

With this assumption, we can express the share of workers of gender g and ability z who chooses to work in occupation-sector pair oj as:

$$\Pr(oj|g,z) = \frac{\exp\left[\frac{1}{\sigma_\varepsilon} V(I_{ojg}(z), p) - \frac{1}{\sigma_\varepsilon} A_{ojg}\right]}{\sum_{j'} \sum_{o'} \exp\left[\frac{1}{\sigma_\varepsilon} V(I_{o'j'g}(z), p) - \frac{1}{\sigma_\varepsilon} A_{o'j'g}\right]}. \tag{4}$$

This Equation shows that workers select into occupation-sectors based on their comparative advantage (i.e., real income earned), the gender-specific amenity or utility costs A_{ojg}, and their idiosyncratic preference shocks.

Occupation-sector pairs differ in their returns to ability κ_{oj} so that $z^{\kappa_{oj}}$ are the effective units of human capital that a worker of gender g and ability z can supply to occupation-sector pair oj. For each unit of human capital, a occupation-sector-pair pays a wage rate w_{ojg} which can vary for men and women. We model female wage discrimination as a "wedge" τ_{oj} between women's wage rate and their marginal product so that: $\{w_{ojm}, w_{ojf}\} = \{w_{oj}, (1-\tau_{oj})w_{oj}\}$. This assumes that men receive the effective wage rate w_{oj}, while women are paid only a fraction $(1-\tau_{oj})$ of their marginal product. To micro-found this wedge, we follow Hsieh et al. (2019) and assume that entrepreneurs have a disutility δ_{oj} of hiring women, which is compensated with the profits that arise from paying women below their marginal product. The total income of an individual of gender g and ability z, who works in an occupation-sector oj is therefore given by: $I_{ojg} = w_{ojg} \times z^{\kappa_{oj}}$.

5.3 Production, Aggregation and Equilibrium

Aggregate Supply of Human Capital: To derive the total human capital that is supplied to each occupation-sector pair, we can sum across individuals' employment choices:

$$H_{oj} = \sum_g N_g \int_z \Pr(oj|g,z) z^{\kappa_{oj}} dF(z),. \tag{5}$$

where κ_{oj} represents the occupation-sector-specific returns to ability.

Production: A representative firm in each sector produces output Y_j with a Cobb Douglas production function, using as input the human capital from each occupation, so that:

$$Y_j = B_j \prod_o H_{oj}^{\gamma_{oj}}, \tag{6}$$

where γ_{oj} are the Cobb Douglas expenditure shares for human capital from each occupation which sum to 1 across occupations ($\sum_o \gamma_{oj} = 1$). B_j is sector-specific productivity, and H_{oj} is the total human capital that is supplied to oj as shown in Equation (5). Firms' profits are therefore given by:

$$\pi_j = p_j Y_j - \sum_o w_{oj} H_{oj},$$

where p_j are sectoral prices.

Aggregate Expenditure Shares: PIGL preferences allow us to aggregate sectoral expenditure shares across individuals in a tractable, closed-form solution despite its non-homothetic nature. Aggregate expenditure shares in sector j are given by:

$$\Phi_j = \omega_j + \nu_j \times \frac{P^\eta}{I} \times \sum_o \sum_j \sum_g \int_z I_{ojg}^{(1-\eta)} \times \Pr(g) \times \Pr(oj|gz) dz, \tag{7}$$

which we derive in Appendix C.3.

Equilibrium: Exogenous parameters of the model characterize preferences $\{\alpha_s, \eta_s, \omega_j, \nu_j, \eta\}$, the dispersion of preference shocks across occupation-sector-pairs (σ_ε), the ability distribution $z \sim F(z)$, the production side $\{\gamma_{oj}, B_j, \kappa_{oj}\}$, and gender barriers $\{\tau_{oj}, A_{ojg}\}$. Given these parameters, the equilibrium in each country-year is defined by a vector of sectoral prices and occupation-sector-specific wage rates $\left\{\{p_j\}_{\forall j}, \{w_{oj}\}_{\forall oj}\right\}$ which ensure that:

1. Workers make optimal consumption and employment choices.
2. Firms in each sector hire human capital from each occupation to maximize profits.
3. Labor markets clear in each occupation-sector pair equalizing human capital supply and demand.
4. Good markets clear in each sector.

6 Model Calibration

We calibrate a set of parameters to the literature or to data moments outside of our model. The remaining parameters are then calibrated by fitting our model's equilibrium conditions to the data in an iterative algorithm which is described in Appendix D.

6.1 Parameters Calibrated Outside of the Model

Preferences: Preference parameters do not vary across countries or over time. For the CES preferences over home and market services, we follow Ngai and Petrongolo (2017) and set the elasticity of substitution to $\eta_s = 2.3$, and the share of home services to $\alpha_{hs} = 0.3$. For the preference shocks across occupation-sector-pairs, we follow Hsieh et al. (2019) and set the dispersion parameter to $\sigma_\epsilon = 2$. For the PIGL preferences, we follow Fan, Peters and Zilibotti (2021) and set the elasticity of substitution to $\eta = 0.395$. We estimate the remaining PIGL parameters $\{\omega_j, \nu_j\}$ by indirect inference, as described below.

Ability Distribution and Returns to Human Capital: Ability distributions vary across countries and over time and follow a log-normal distribution, so that $z \sim \ln N(\mu, \sigma_z)$. We calibrate $\{\mu, \sigma_z\}$ for each country-year from the moments of the observed schooling distribution. Occupation-sector specific returns to human capital κ_{oj} determine the income of an individual of gender g and ability z who works in that occupation-sector in the following way: $I_{ojg}(z) = w_{ojg} \times z^{\kappa_{oj}}$. This provides us with a structural equation that we can take to the data to estimate κ_{oj}, similar to Fan, Peters and Zilibotti (2021). Specifically, we use individual-level data for each country-year to estimate the following Mincerian wage regression:

$$\ln I_i = \alpha_{oj} + \kappa_{oj} \ln(\text{YrsSchool}_i) + \varepsilon_i, \tag{8}$$

where I_i is income/earnings and YrsSchool_i are years of schooling of an individual i. α_{oj} are occupation-sector fixed effects that capture average wages and other unobserved factor that can affect individuals' wages in each occupation-sector. The coefficient $\hat{\kappa}_{oj}$ is estimated for each occupation-sector-pair and corresponds through the lens of our model to the occupation-sector-specific returns to human capital. We estimate the above equation separately for each country-year. We set $\kappa_{home} = 0$.

Occupational Cobb Douglas Expenditure Shares: For every country-year and each sector, we further fix the Cobb Douglas occupational expenditure shares γ_{oj} to the wage expenditure shares that we observe in the data.

6.2 Parameters Calibrated using Model's Equilibrium Conditions

We estimate the remaining parameters by fitting our model's equilibrium conditions to key data moments. The remaining parameters are PIGL preferences $\mathscr{U} = \{\omega_j, \nu_j\}_{\forall j}$, sectoral productivity $\mathscr{P} = \left\{\{B_j\}_{\forall j}\right.$, and gender barriers $\mathscr{B} = \left\{\{\tau_{oj}\}_{\forall oj}, \{A_{ojg}\}_{\forall ojg}\right\}$, which consist of wage discrimination and gender norms. Here we provide the intuition of the estimation strategy while we describe the numerical procedure and the iterative algorithm in Appendix D.

Wage Discrimination: We estimate wage rates w_{oj} and women's wage discrimination τ_{oj} in each occupation-sector and each country-year by matching men's and women's average wage data. In our model, individuals of gender g who work in occupation-sector oj earn average wages:

$$\overline{\text{wage}}_{ojg} = w_{ojg}\overline{H}_{ojg}, \tag{9}$$

where $\overline{H}_{ojg}$ is the average human capital that men or women supply to an occupation-sector oj, which is defined as: $\overline{H}_{ojg} = \int_z \Pr(oj|g,z) z^{\kappa_{oj}} dF(z)$.

This measure of average human capital is not observed in the data as it depends on the extent to which workers sort into occupation-sectors due to their comparative advantage ($z^{\kappa_{oj}}$) or due to other factors such as gender barriers or amenities. We therefore use our model and an iterative algorithm to compute the average human capital measure. We then use this model-implied measure of average human capital and data on average wages for each occupation-sector and gender to infer men's and women's wage rates w_{ojg} in each occupation-sector according to Equation 9. We assume men's wages to be undistorted, so that: $\{w_{ojm}, w_{ojf}\} = \{w_{oj}, (1-\tau_{oj})w_{oj}\}$.

To provide intuition, we can write the observed female-to-male wage ratio in each occupation sector as a function of female wage discrimination (τ_{oj}) and the female-to-male ratio of (model-implied) human capital:

$$\underbrace{\frac{\overline{\text{wage}}_{ojf}}{\overline{\text{wage}}_{ojm}}}_{\text{Obs. Wage Gap}} = \underbrace{(1-\tau_{oj})}_{\text{Female Wage Discr}} \times \underbrace{\frac{\overline{H}_{ojf}}{\overline{H}_{ojm}}}_{\text{Avg. HC Gap}} . \tag{10}$$

Gender Norms: Individuals choose an occupation-sector pair oj based on real income and based on a gender-specific utility cost which captures amenities and gender norms (as shown in Equation 4 and presented in a simplified form here):

$$\Pr(oj|g) \propto \left[\underbrace{V(I,p)}_{\text{Real Income}} - \underbrace{A_{ojg}}_{\text{Utility Cost}} \right],$$

where $V(I,p)$ is workers' indirect utility from consumption that they can obtain from the real income earned in oj (Equation 1), and A_{ojg} represents the gender-specific utility cost of working in oj.

For each gender, we can only identify the relative utility costs A_{ojg}. We therefore normalize $A_{home,g} = 0$ without loss of generality so that A_{ojg} expresses the additional utility cost of working in each occupation-sector relative to the home sector for each gender. To estimate A_{ojg}, we use our model to compute worker's indirect utility $V(I,p)$ in each occupation-sector. We then use this model-implied measure and data on men's and women's employment shares $\Pr(oj|g)$ to infer A_{ojg} for each occupation-sector according to Equation 4.

Sectoral Prices and Sectoral Productivity Growth: We normalize productivity in the baseline period to 1 for every sector and every country, so that $B_{j0} = 1$. We then infer the productivity levels in subsequent years to target real value added growth in each sector (which we obtain from the ETD database). Sectoral prices p_j are solved in the model to ensure that the good's market clears in each sector.

PIGL Parameters: The PIGL parameters ω_j and ν_j are identified by indirect inference from Equation 7. For each sector, we regress the observed expenditure shares from the data on model-implied real income across the country-years of our sample. The constant of this regression identifies the PIGL expenditure shares ω_j and the coefficient on the real income measure identifies ν_j.

7 Estimation Results and Model Validation

7.1 Estimation Results for Gender Barriers

Our estimates of wage discrimination and gender norms for each occupation-sector-country-year quantify the part of gender differences in wages and employment that cannot be explained by differences in human capital or skill sorting.

Figure 5 shows how gender barriers evolve over time for the six countries in our sample. Figures 5(a) and 5(b) focus on gender norms which decrease over time for most occupations and sectors. India and Indonesia have higher gender norms in most occupations and sectors and saw less improvements over time. Level differences in gender norms across sectors also varies across countries. Relative gender norms across sectors or occupations also vary across countries. Brazil, Mexico, Canada, and the US have smaller gender norms in the service sector and in clerk and professional occupations relative to other jobs, while the opposite is true for India and Indonesia.

Figures 5(c) and 5(d) show the results for female wage discrimination. At the beginning of the sample period, most countries had high wage discrimination, paying women only between 40 and 60% of their marginal product. Wage discrimination decreased over time in most occupations and sectors but remains large in the most recent year in many occupations and sectors. In Brazil, Mexico, Canada and the US, women face more wage discrimination in professional occupations (where they are paid between 65-75% of their marginal value). Levels in female wage discrimination is strikingly similar across the countries of our sample. In contrast to the gender norms discussed previously, we see that there is no discernable correlation in wage discrimination and economic development. Today's developed countries do not perform better in fair remuneration than today's emerging countries.

Appendix Figures A.3.2(a) and A.3.2(b) compare how the distribution of gender barriers across all occupation-sector-pairs changes over time by plotting the CDFs for the first and last year across the six countries of our sample. There has been a substantial reduction in barriers over time and both distributions move closer to 0 (Kolmogorov-Smirnov p-value=0.00). Despite the improvements, female wage discrimination and gender norms remain persistent in around 80% of occupation-sectors in the most recent sample year of our sample countries.[6]

7.2 Model Fit and Validation

Model Fit

Figure B1 pools the data across all occupations, sectors, countries, and years and shows a strong negative correlation between the calibrated gender norms $\triangle A_{ojct}$ and the observed gender employment gaps in the data. Figure B2 documents the same for the correlation between wage discrimination τ_{ojct} and gender wage gaps. Figure B3 shows that our gender barriers also closely track the gender wage and employment gaps for each country over time. As an additional model validation, Figure B4 shows that our model closely replicates the share of nominal value-added across sectors.

[6] The lower panel of Figure A.3.2 shows that these patterns are not specific to the choice of our sample period as results are similar when we use data for all countries for 1980 to 2010.

Validating the Model Estimates with Measures of Social Norms

To examine the extent to which our estimates of gender barriers capture measurable changes in women's underlying social norms and labor market constraints, we use the World Bank's "World, Business, and the Law" (WBL) database (World Bank, 2019, 2020; Hyland, Simeon and Goldberg, 2020). The WBL data set evaluates 35 aspects of countries' legal code to create 8 indicators, which measure gender equality in the labor market, at the workplace, and in the legal code across 190 countries and over five decades (1970-2020). The indicators include answers to gender normative questions such as "Can a woman get a job in the same way as a man?" or "Can a woman work at night in the same way as a man"?). We then regress our estimated gender barriers on the WBL indicators in the following way:

$$y_{ojct} = \alpha_c + \alpha_t + \mathrm{WBL}_{ct} + \ln \mathrm{GDP\ p.c.}_{ct} + \varepsilon_{ct}, \tag{11}$$

where the dependent variable y_{ojct} is either female wage discrimination τ_{ojct} or gender norms $\triangle A_{ojct}$ and where WBL_{ct} is a specific indicator from the WBL data for a country c in year t. We control for real GDP per-capita and for country and year fixed effects to control for unobserved time-invariant country-specific social norms, or other unobserved (economic or social) changes over time. We cluster standard errors at the country-level.

Panel A of Table B1 shows a strong negative relationship between our estimated gender norms $\triangle A_{ojct}$ and 5 WBL indicators which measure gender equality in mobility and at the workplace (e.g., equality in getting a job, working at night, working in industrial and 'dangerous' jobs). Panel B of Table B1 examines the correlation between our estimated gender norms and 5 measures of gender equality in the household (e.g., equality as household head, rights to remarry and ownership, and legislations against domestic violence). We find a negative correlation in 4 out of the 5 indicators which are statistically significant for 3 out of the 5 indicators.

Panel A of Table B2 shows a negative correlation between estimated wage discrimination τ_{ojct} and the 5 WBL indicators on gender equality in mobility and at the workplace (as described above). In Panel B, we regress our estimated wage discrimination on WBL indicators that measure gender equality at the workplace including whether it offers paid maternity leave. We find a negative and significant correlation only for the provision of paid maternity leave.

Put together, the results show that the gender barriers (estimated as residuals in our model) contain important information about and correlate with changes in underlying social and gender norms that are measured in the data across countries and over time.

8 Quantitative Results: Importance of Gender Barriers

Given our estimates of gender barriers, we can now answer the following question: how much of structural transformation can be attributed to the changes in gender barriers in each country of our sample?

In our model, the gendered employment transitions from agriculture and homework to manufacturing and services can be driven by multiple channels: changes in gender barriers (wage discrimination or gender norms: τ_{oj}, A_{ojg}), changes in production technologies (occupation-, sector-, or skill-biased technological change: γ_{oj}, B_j, κ_{oj}), changes in human capital (μ_{zg}, σ_{zg}), or by income effects which affect consumers' expenditure shares (ω_j, ν_j).

To assess how much changes in gender barriers contribute to each country's process of structural transformation and productivity, we simulate counterfactuals which hold gender barriers constant at the values that we observe in the first year of each country. First, we fix only female wage discrimination (τ_{oj}), second only gender norms ($\triangle A_{oj}$), and third, gender norms and female wage discrimination simultaneously. We allow all other parameters to evolve according to the data and our calibration. For each counterfactual, we solve for workers' employment choices, wages, and sectoral prices which are consistent with the general equilibrium of the model. To quantify the importance of changes in gender barriers, we then compare the counterfactual path to the actual path in the data (to which we estimate the model).

Effects of Gender Barriers on Employment Transitions

We first examine the importance of changes in gender barriers on sectoral and occupational employment shares. In Figure 6, we take the average across the six countries of our sample and show how sectoral and occupational employment shares change in the baseline and in each counterfactual during the time period of our sample (1970-2015). In the baseline data (black bars), labor force participation increased on average by 0.19 p.p. each year. For our median sample period of approximately 40 years, this corresponds to a total change of 7.6 p.p. A counterfactual simulation that fixes wage discrimination (τ_{oj}) to the values in

1970 would have generated a 0.15 p.p. annualized increase in the LFPR which implies that the reduction in wage discrimination can explain around 21% (1-0.15/0.19) of the observed increase in LFP. While this magnitude is non-trivial, changes in gender norms had much larger effects. If gender norms had not changed since their initial values in the 1970s, LFP would have actually decreased by 0.21 p.p. each year or by 0.25 p.p. if neither wage discrimination nor gender norms had changed. Changes in gender norms also had large effects on sectoral employment. Employment share in agriculture decreased at 1.4 p.p. each year on average across countries in our sample. Without changes in wage discrimination and gender norms, the employment share in agriculture would have decreased even faster–around 1.5 p.p. a year. These changes in agriculture and LFPR would have consequently slowed employment growth in manufacturing and services. For example, the employment share in services would have grown by only 0.93 p.p. instead of 1.77 p.p. per year, implying that changes in gender norms explain around half, or 47.5% (1-0.93/1.72) of the observed growth in service employment. Similarly, changes in gender norms explain just over a third, or 36.8% of the observed growth in the employment share in manufacturing. Figure 6(b) shows the effects for employment by occupation. Changes in gender barriers explain around half the changes in professional occupations (46.8%), around 40% and 60% for service workers and clerks respectively. and about a third for machine operators. Most of these effects are again driven by changes in gender norms.

Effect of Gender Barriers on Sectoral Output

Similar to the above exercise, we now turn to calculating how much of the change in sectoral output between 1970-75 and 2010-18 can be explained by changes in gender barriers. We report the results in Table 2. Columns (1)-(3) report the fraction of output growth in agriculture, manufacturing and services that are explained by gender barriers, while Column (4) report the the same for aggregate output. Aggregate output is calculated as an expenditure-share weighted average of the sectoral output. We report the results for each country, along with a sample average.

Turning to the results, from Column (1), changes in gender barriers account for 15% of growth in agricultural output on average. This ranges from 9% in Indonesia and around a quarter in Brazil, to around 12-14% for the other countries. From Column (2), changes in gender barriers account for 16% of growth in manufacturing output on average. For low-income countries (India and Indoensia), this is only 5-7%, while around 20% for the middle- and high-income countries. From Column (3), a quarter of the growth in services output can be attributed to changes in gender barriers on average. However, this is lowest in India (4%), around a third in Brazil, and around a quarter for the other countries.

Overall from Column (4), changes in gender barriers from 1970-2015 explain around 20% of changes in aggregate output. However, there is a large variation across countries, ranging from 4% in India, 17% in Indonesia, and round 25-30% in the other countries. Put together, we conclude that they were an important driver of reallocation of workers across sectors as well as had large effects on output growth.

9 Conclusion

The paper documents stylized facts about the gender dimension of structural transformation across multiple countries over the last five decades. We find substantial gender gaps in employment and wages across occupations-sector pairs, which narrow over time, but still persist today even in the most developed countries. To quantify the effects of gender barriers on economic outcomes, we develop a general equilibrium Roy model that incorporates standard economic drivers of occupational and sectoral choices as well as gender barriers through the form of wage discrimination and gender norms. We estimate the model for six countries across five decades, and use our estimated model for a counterfactual analysis. We find that the reduction in gender barriers over the last five decades had large effects on sectoral employment changes and output growth. The importance of changes in gender barriers varies across sectors and countries with larger effects for the service sector and small effects for agriculture.

Our analysis (intentionally) does not propose specific policies that could bolster gender parity in the labor market, but we view our quantitative model as a useful framework that allows decomposing observable changes in empirical data patterns into a part that is due to standard economic channels and another part that is due to changes in gender barriers. In addition, our general equilibrium framework is useful to aggregate changes in individual choices to quantify the macroeconomic and sectoral effects of changing gender barriers. Future research should explore the underlying factors that led to larger declines in gender barriers in some countries (like Brazil) than in others (like India).

References

Alder, Simon, Timo Boppart, and Andreas Müller, "A theory of structural change that can fit the data," *American Economic Journal: Macroeconomics*, 2022, *14* (2), 160–206.

Bridgman, Benjamin, Georg Duernecker, and Berthold Herrendorf, "Structural transformation, marketization, and household production around the world," *Journal of Development Economics*, 2018, *133*, 102–126.

Comin, Diego, Danial Lashkari, and Martí Mestieri, "Structural change with long-run income and price effects," *Econometrica*, 2021, *89* (1), 311–374.

Cuberes, David and Marc Teignier, "Gender inequality and economic growth: A critical review," *Journal of International Development*, 2014, *26* (2), 260–276.

— **and** — , "Aggregate effects of gender gaps in the labor market: A quantitative estimate," *Journal of Human Capital*, 2016, *10*, 1–32.

Fan, Tianyu, Michael Peters, and Fabrizio Zilibotti, "Service-Led or Service-Biased Growth? Equilibrium Development Accounting across Indian Districts," Technical Report, National Bureau of Economic Research 2021.

Fletcher, Erin, Rohini Pande, and Charity Maria Troyer Moore, "Women and work in India: Descriptive evidence and a review of potential policies," 2017.

Goldin, Claudia, "The u-shaped female labor force function in economic development and economic history," *Working Paper 4707, National Bureau of Economic Research*, 1994.

Gottlieb, Charles, Douglas Gollin, Cheryl Doss, and Markus Poschke, "Gender, Work and Structural Transformation," 2021.

Heath, Rachel and Seema Jayachandran, "The causes and consequences of increased female education and labor force participation in developing countries," Technical Report, National Bureau of Economic Research 2016.

Herrendorf, Berthold and Todd Schoellman, "Wages, human capital, and barriers to structural transformation," *American Economic Journal: Macroeconomics*, 2018, *10* (2), 1–23.

— **, Richard Rogerson, and Akos Valentinyi**, "Two perspectives on preferences and structural transformation," *American Economic Review*, 2013, *103* (7), 2752–89.

Hsieh, Chang-Tai, Erik Hurst, Charles Jones, and Peter Klenow, "The Allocation of Talent and U.S. Economic Growth," *Econometrica*, 2019, *87* (5), 1439–1474.

Hyland, Marie, Djankov Simeon, and Pinelopi Koujianou Goldberg, "Gendered laws and women in the workforce," *Unpublished Manuscript*, 2020.

IPUMS International, "Minnesota Population Center, Integrated Public Use Microdata Series, International," *Version 7.3 [dataset]. Minneapolis, MN: IPUMS, 2020. https://doi.org/10.18128/D020.V7.3*, 2020.

Lee, Munseob, "Allocation of Female Talent and Cross-Country Productivity Differences," *Unpublished Manuscript*, 2020.

Mammen, Kristin and Christina Paxson, "Women's work and economic development," *Journal of economic perspectives*, 2000, *14* (4), 141–164.

Moro, Alessio, Solmaz Moslehi, and Satoshi Tanaka, "Does home production drive structural transformation?," *American Economic Journal Macroeconomics*, 2017, pp. 116–146.

Ngai, L. Rachel and Barbara Petrongolo, "Gender Gaps and the Rise of the Service Economy," *American Economic Journal: Macroeconomics*, 2017, *9* (4), 1–44.

Olivetti, Claudia, "The female labor force and long-run development: the american experience in comparative perspective," *Technical report, National Bureau of Economic Research*, 2013.

Psacharopoulos, George and Zafiris Tzannatos, "Female labor force participation: An international perspective," *The World Bank Research Observer*, 1989, *4* (2), 187–201.

Rendall, Michelle, "Female market work, tax regimes, and the rise of the service sector," *Review of Economic Dynamics*, 2018, *28*, 269–289.

Roy, Andrew Donald, "Some thoughts on the distribution of earnings," *Oxford Economic Papers*, 1951, *3* (2), 135–146.

World Bank, "Women, Business and the Law 2019: A Decade of Reform," *Technical Report*, 2019.

_ , "Women, Business and the Law 2020," *Technical Report*, 2020.

Figures

Figure 1: Sectoral and Occuaptional Employment by Gender

(a) Sectoral: Male

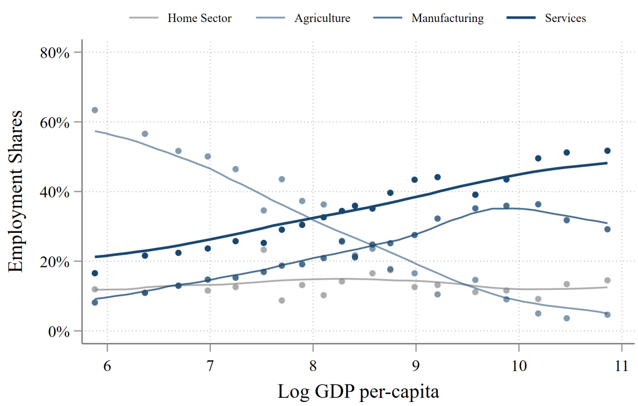

(b) Sectoral: Female

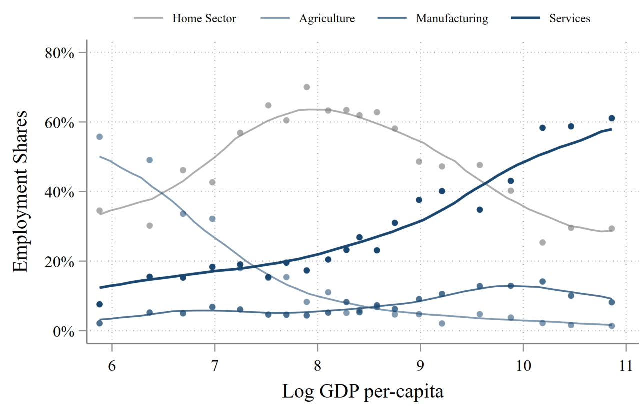

(c) Occupational: Male

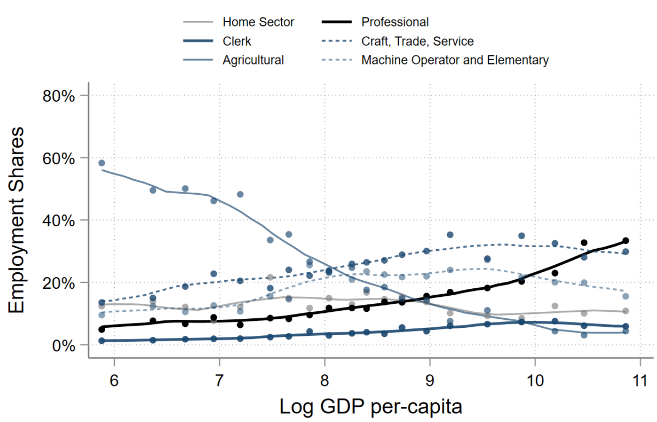

(d) Occupational: Female

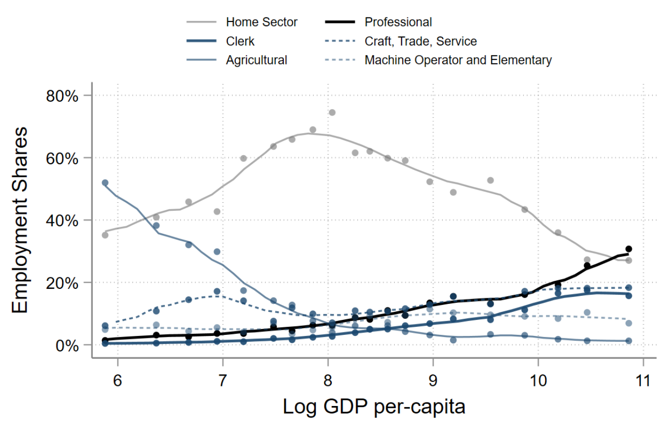

Notes: This figure is a non-parametric plot of the share of men and women in each sector and occupation against the log of real GDP per-capita in 2010 US dollars. The sample pools all available country years from the IPUMS data.

Figure 2: Occupation Structure within Manufacturing and Services by Gender

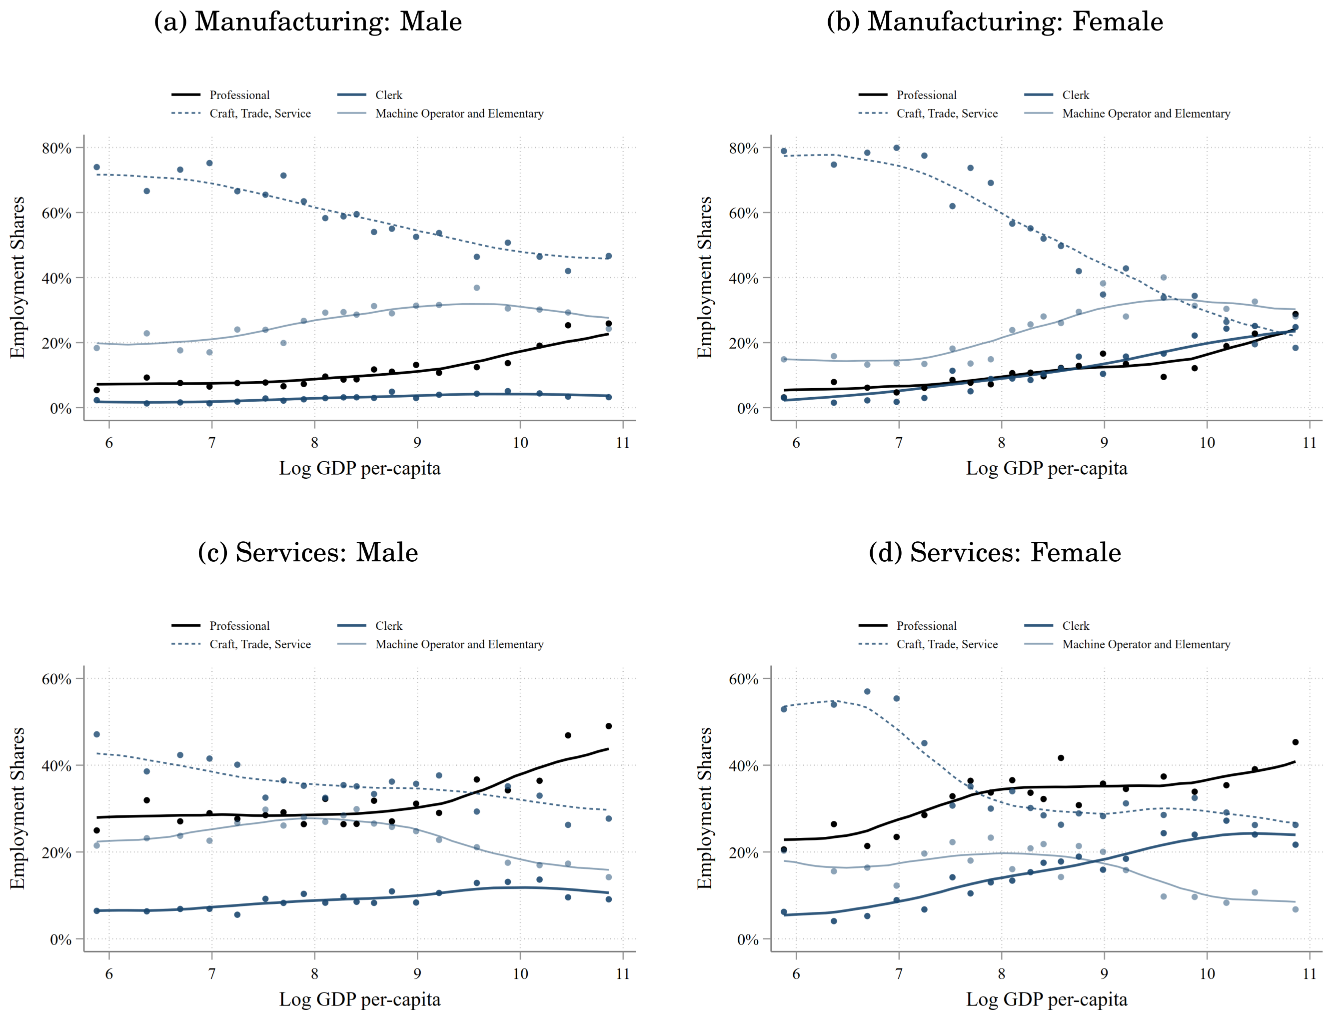

Notes: This figures is a non-parametric plot of the share of men and women in each occupation against the log of real GDP per-capita. Figures (a) and (b) show the occupational distribution among workers in the manufacturing sector. Figures (c) and (d) show the distribution for workers in the service sector. The sample pools all available country years from the IPUMS data.

Figure 3: Employment and Wage Ratios over Time

(a) Employment Ratio by Sector

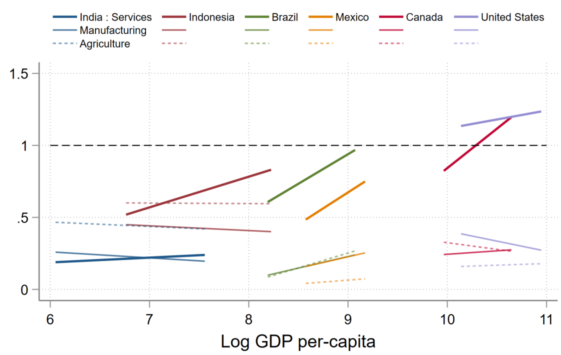

(b) Employment Ratio by Occupation

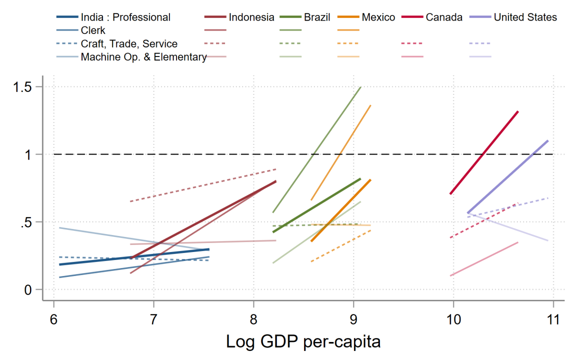

(c) Wage Ratio by Sector

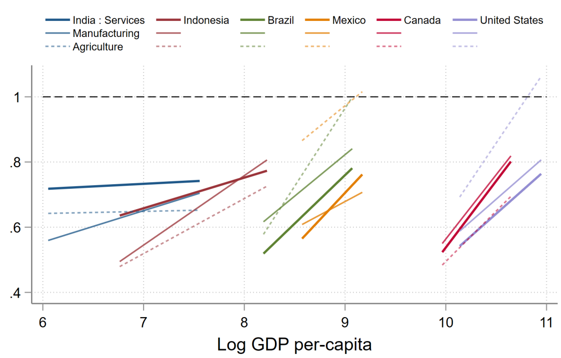

(d) Wage Ratio by Occupation

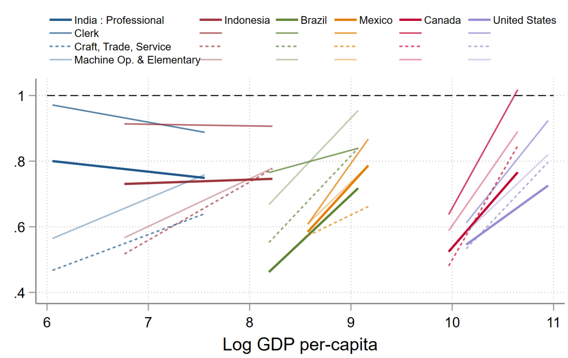

Notes: This figure plots the employment and wage ratios for selected countries over time against the log of real GDP per-capita in constant 2010 US dollars. The time period covers between 1970 and 2018 depending on data availability and the horizontal dimension of the graph shows how fast countries grew during the sample period. Employment ratios divide the share of women working in an occupation-sector by the share of men. Wage ratios divide the average wage of women in an occupation-sector by the average wage of men. Figures (a) and (c) show these ratios by sector while figures (b) and (d) show them by occupation. Figure (b) excludes the clerk occupation for the US and Canada as their employment ratios exceed 2 which makes the graph hard to read.

Figure 4: Gender Segregation across Sectors and Occupations

(a) Theil Index Level

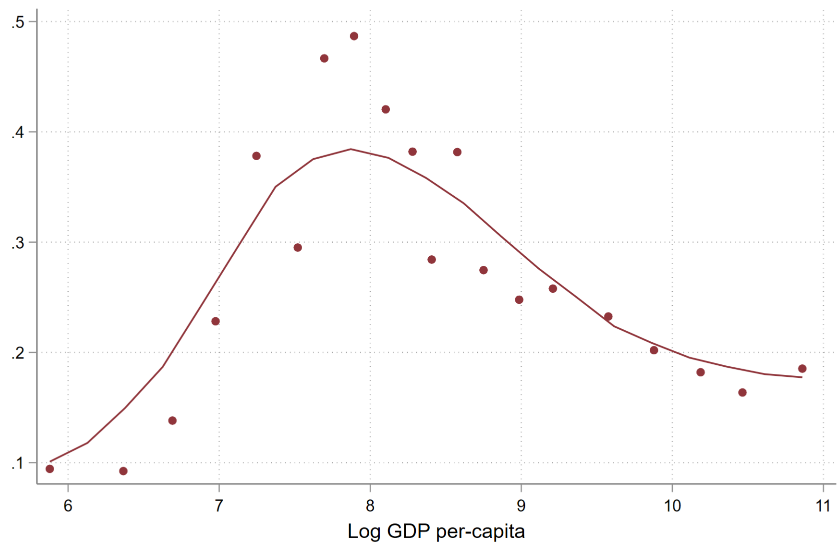

(b) Within Sector

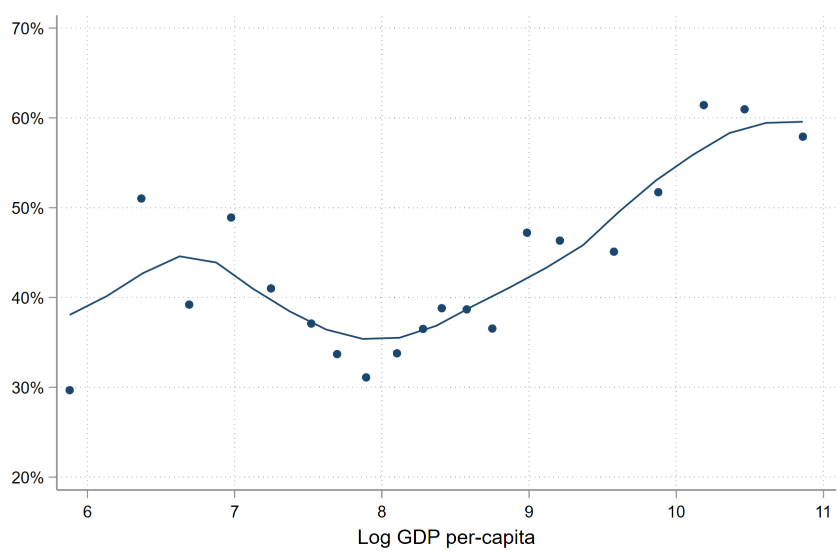

Notes: This figure plots the Theil Index that measures gender segregation across occupation-sector pairs against the lof og real GDP per capita. Figure (a) plots the level of the segregation Index and figure (b) plots the share of segregation that is explained by segregation across-occupations within-sectors.

Figure 5: Gender Norms and Wage Discrimination Over Time

(a) Gender Norms by Sector

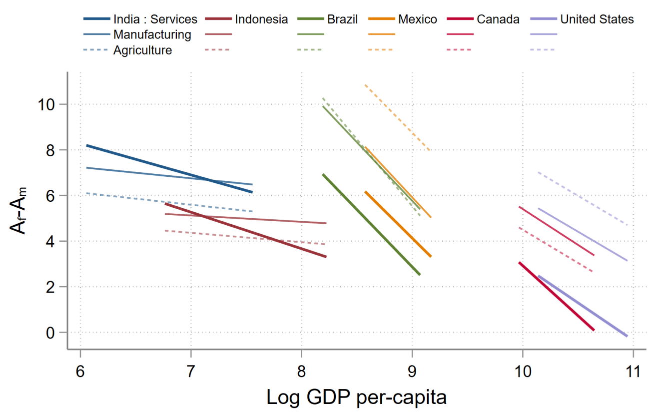

(b) Gender Norms by Occupation

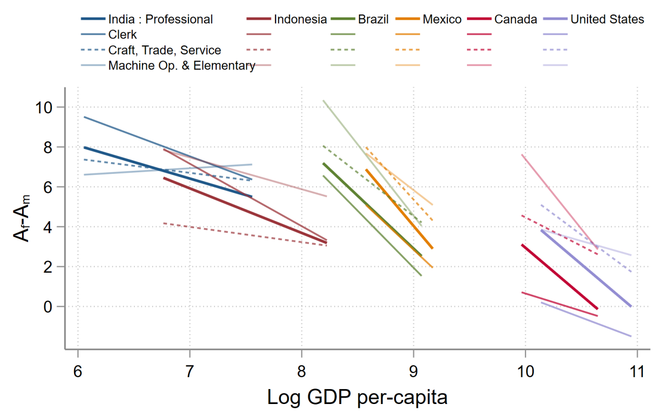

(c) Wage Discrimination by Sector

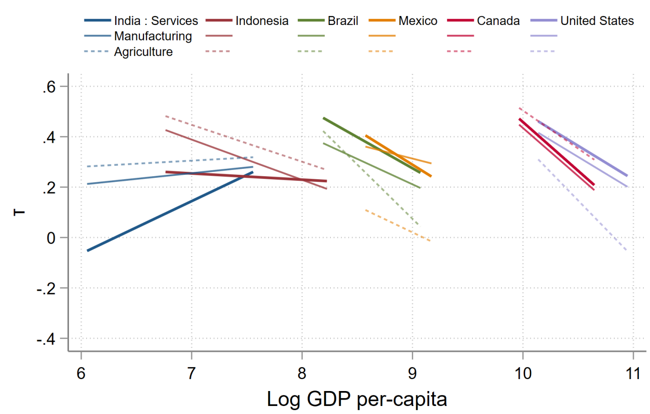

(d) Wage Discrimination by Occupation

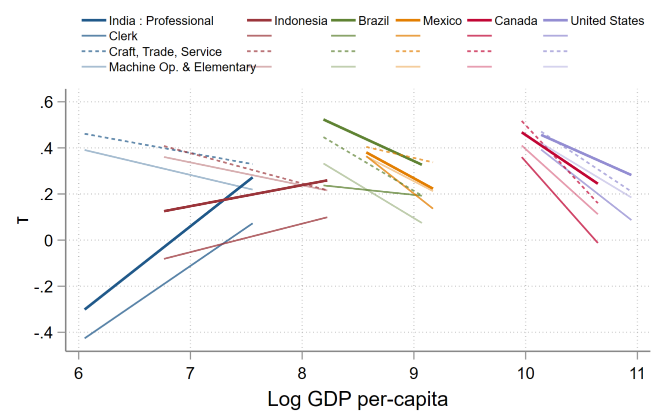

Notes: This figure plots the estimated gender norms ($\triangle A_{oj}$) and wage penalties (τ) for selected countries over time against the log of real GDP per-capita in constant 2010 US dollars. The time period covers between 1970 and 2018 depending on data availability and the horizontal dimension of the graph shows how fast countries grew during the sample period. Figures (a) and (c) show gender norms and wage discrimination by sector while figures (b) and (d) show them by occupation.

Figure 6: Gender Barriers and Change in Employment Shares Across Sectors and Occupations

(a) Employment Shares Across Sectors

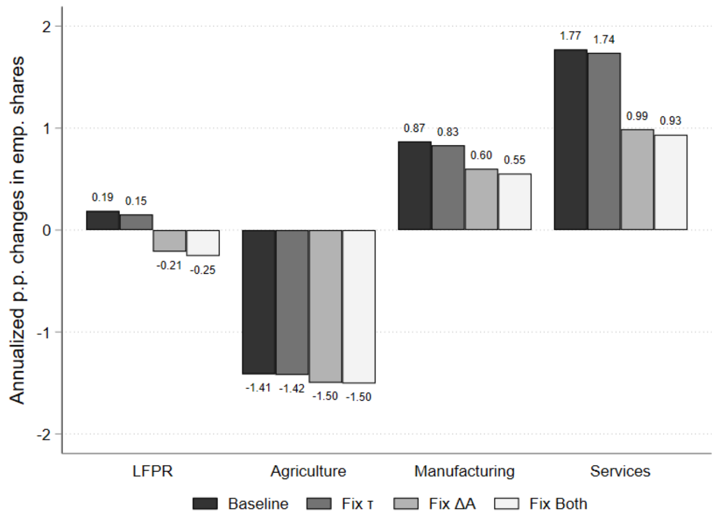

(b) Employment Shares Across Occupations

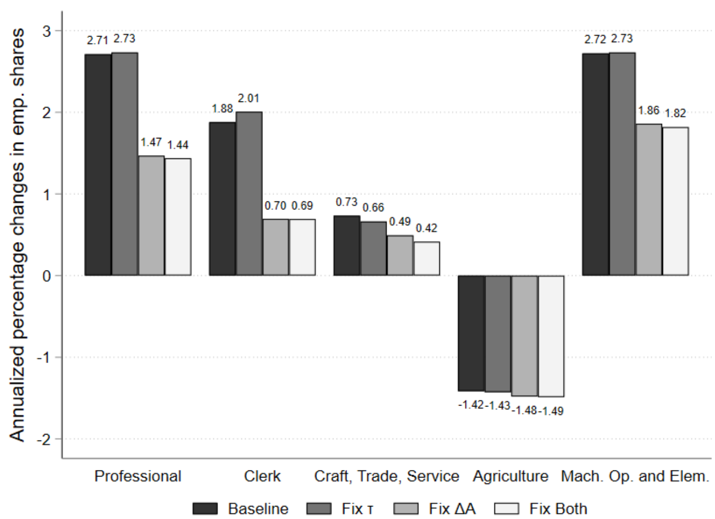

Notes: This figure reports the average annualized percentage point changes (between the first and last year) in sectoral and occupational employment shares across countries. Figure (a) shows changes across sectors, while Figure (b) shows changes across occupations. The labor force participation rate (LFPR) is defined as 1- share of individuals in the home sector.

Tables

Table 1: Gender Employment and Wage Ratios

	Employment Ratio			Wage Ratio		
	1970-75	2010-18	Change	1970-75	2010-18	Change
	(1)	(2)	(3)	(4)	(5)	(6)
Home	8.46	3.57	-11.98	0.04	0.07	0.07
Agriculture	0.28	0.30	0.04	0.17	0.07	-0.25
Manufacturing	0.27	0.27	0.02	0.11	0.12	0.02
Services	0.63	0.87	0.59	0.16	0.23	0.16
Professional	0.41	0.86	1.07	0.06	0.10	0.09
Clerk	1.19	1.41	0.54	0.02	0.02	-0.01
Craft, Trade, Service	0.41	0.56	0.33	0.12	0.14	0.03
Agricultural	0.26	0.28	0.03	0.16	0.06	-0.25
Machine Operator and Elementary	0.36	0.41	0.15	0.08	0.11	0.07

Notes: Columns (1)-(2) and (4)-(5) report the employment and wage ratios in the first and last year, averaged across countries. Columns (3) and (5) report the average annual percentage point change in these ratios. The employment ratio divides the share of women working in an occupation-sector by the share of men. The wage ratio divides average wage of women in an occupation-sector by the average wage of men. A ratio below 1 implies lower employment (or lower wages) for women relative to men.

Table 2: Change in Sectoral Output Explained by Gender Barriers

	Sectoral Output			Aggregate
	Agri.	Manf.	Services	Output
	(1)	(2)	(3)	(4)
IND	0.14	0.05	0.04	0.04
IDN	0.09	0.07	0.20	0.17
BRA	0.27	0.23	0.34	0.29
MEX	0.13	0.19	0.28	0.24
CAN	0.12	0.23	0.28	0.30
USA	0.12	0.19	0.25	0.25
AVG	0.15	0.16	0.23	0.21

Notes: This table reports the share of sectoral output growth that is explained by changes in gender barriers. To calculate this share, we compute $(1-\hat{g}/g)$ where g is the output growth observed in the data and $\hat{g}$ is the counterfactual output growth when gender barriers are fixed at their initial values.

ONLINE APPENDIX: NOT FOR PUBLICATION

A Tables and Figures

A.1 Sample Details, Sector and Occupation Classification

Table A.1.1: Coverage of All Countries Across Decades

Decade	Country-Years	Percentage
1960-69	13	4.76
1970-79	30	10.99
1980-89	44	16.12
1990-99	57	20.88
2000-10	81	29.67
2010-18	48	17.58
Total	273	100

Notes: The above table reports the coverage of country-years in our data across decades.

Table A.1.2: Coverage of Six Countries

Country	Years	GDP p.c. in 2010
India	1983 to 2018	$1,357
Indonesia	1976 to 2018	$866
Mexico	1960 to 2015	$9,271
Brazil	1970 to 2010	$11,286
Canada	1971 to 2011	$48,464
USA	1960 to 2015	$48,467

Notes: The above table reports the coverage of our data in the final sample of countries. Column 2 reports the years while Column 3 reports the GDP per-capita in 2010 US dollars from the World Bank data.

Table A.1.3: Classification of Sectors

Sector	IPUMS classification
Agriculture	Agriculture, Fishing and Forestry
Manufacturing	Mining, Construction, Electricity, Gas, Water
Market Services	Retail, Wholesale, Transport, Hotels, Education, Health
"Home Work"	Unemployed, Inactive or in Household Services

Notes: The above table shows the classification of sectors reported in the IPUMS data into Agriculture, Manufacturing, Market and Home services.

Figure A.1.1: Coverage of Countries

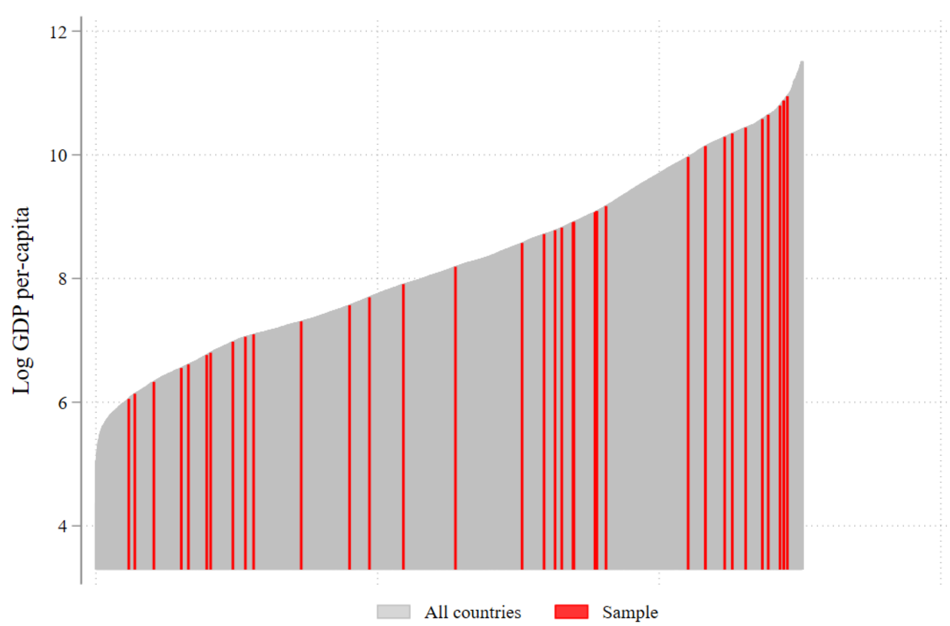

Notes: The above figure sorts all country-years by their GDP per-capita (USD 2010) and shows the coverage of countries in our final sample in red.

Table A.1.4: Classification of Occupations

Code	Occupation	Classification	Sector	Details
1	Legislators, Senior Officials and Managers	Professional	M, S	Legislators and Senior Officials, General and
2	Professionals	Professional	M, S	Technical Managers,
3	Technicians and Associate Professionals	Professional	M, S	Professionals and Technicians
4	Clerks	Clerks	M, S	Secretaries, Librarians, Cashiers, Clerks
5	Service Workers and Shop and Market Sales	Services Workers	M, S	Travel, Housekeeping, Personalcare Workers, Shop and Market Sales and Service Workers
6	Skilled Agricultural and Fishery Workers	Skilled Agri.	A	Subsistence and Market-oriented Workers, Crop Growers and Animal Producers, Forestry and Fishery Workers
7	Crafts and Related Trades Workers	Craft/Trade Wrkrs	M, S	Builders, Painters, Blacksmiths, Electricians, Potters, Printers, Textile, Leather Workers
8	Plant and Machine Operators	Plant & Machine	M, S	Plant and Machine Operators in Mining, Metal, Glass, Wood, Chemical, Rubber, Transportation
9	Elementary Occupations	Elementary	M,S	Street Vendors, Domestic Helpers, Porters, Doorkeepers, Garbage Collectors, Manual and Transportation Laborers
10	Armed forces	Drop		
11	Other occupations, unspecified or n.e.c.	Drop		
97	Response suppressed	Drop		
98	Unknown	Drop		
99	NIU (not in universe)	Drop		

Notes: The above table shows the classification of occupations as reported in the IPUMS data. For our analysis in the paper, we aggregate them based on the ISCO 88 classification (Column 1) as well as report the sectors covered for each occupation (Column 4). Details on the classification and occupations can be found here.

A.2 Theil Index and Decomposition

The Theil Index of segregation is defined by:

$$T_{oj} = \sum_j \sum_o \frac{N_{oj}}{N} \log\left(\frac{N^f/N}{N^f_{oj}/N_{oj}}\right),$$

where N_{oj} is the number of workers in occupation o and sector j, N is the total population, and N^f is the total number of women in the population. A larger number implies more gender segregation across occupations and sectors. In the case of complete gender equality in employment choices, the ratio in the bracket is equal to 1 so that the whole index becomes equal to 0. The Theil index is additively decomposable into segregation across-sector and within-sector (across-occupation) in the following way:

$$T_{oj} = T_j + \sum_j \frac{N_j}{N} T^j_o,$$

where T_j is the Theil index for gender segregation across sectors and T^j_o is the Theil index for gender segregation across occupations in each sector j, which are defined as:

$$T_j = \sum_j \frac{N_j}{N} \log\left(\frac{N^f/N}{N^f_j/N_j}\right) \quad \text{and} \quad T^j_o = \sum_o \frac{N_{oj}}{N_j} \log\left(\frac{N^f_j/N_j}{N^f_{oj}/N_{oj}}\right).$$

A.3 Occupation and Wage Gaps Circa 1980 and 2000

Figure A.3.1: Gender Employment and Wage Ratios Circa 1980 and 2010

(a) Employment Ratio

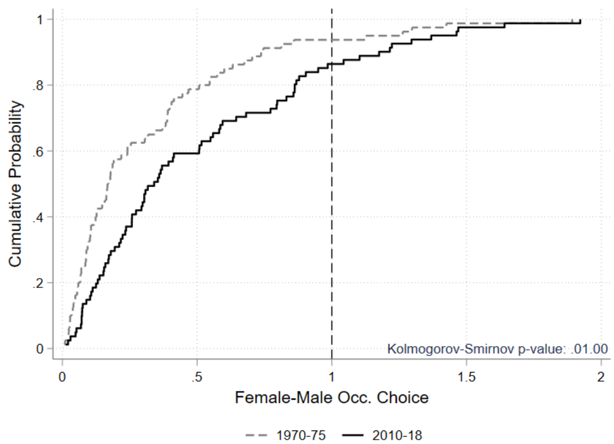

(b) Wage Ratio

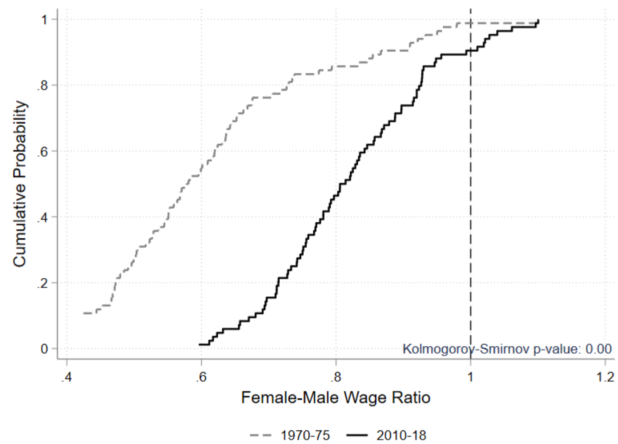

(c) Employment Ratio (Circa 1980-2010)

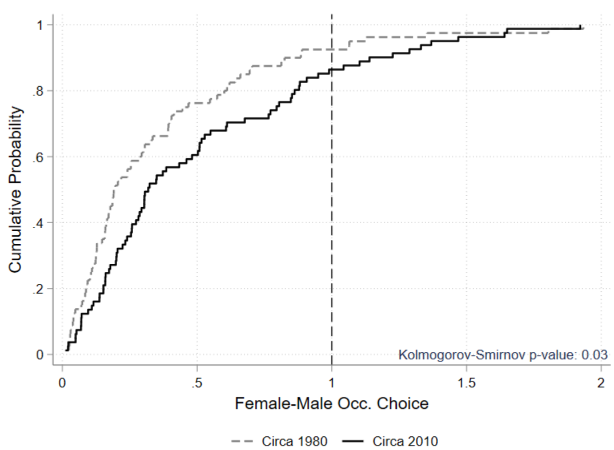

(d) Wage Ratio (Circa 1980-2000)

Notes: The above figure plots the CDF of the employment and wage gaps across all sectors and occupations. Figures (a) and (b) plot the employment and wage gaps using the first (dotted line) and last year (solid line) for each country. Figures (c) and (d) show the same distributions, but now using survey years across countries closest to 1980 and 2010. The employment ratio divides the share of women working in an occupation-sector by the share of men. The wage ratio divides the average wage of women in an occupation-sector by the average wage of men.

Figure A.3.2: Gender Norms and Wage Discrimination Over Time

(a) Gender Norms ($\triangle A_{ojct}$)

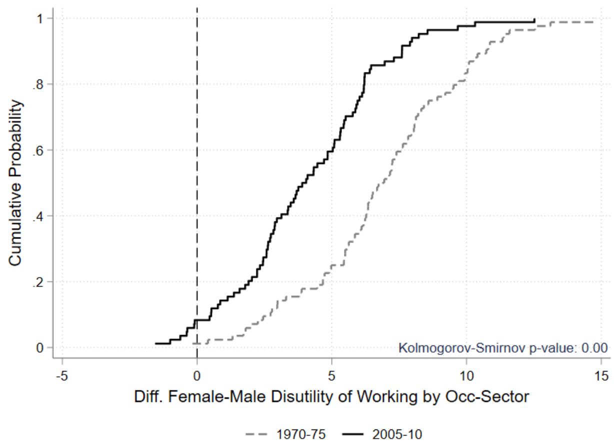

(b) Wage Discrimination (τ_{ojct})

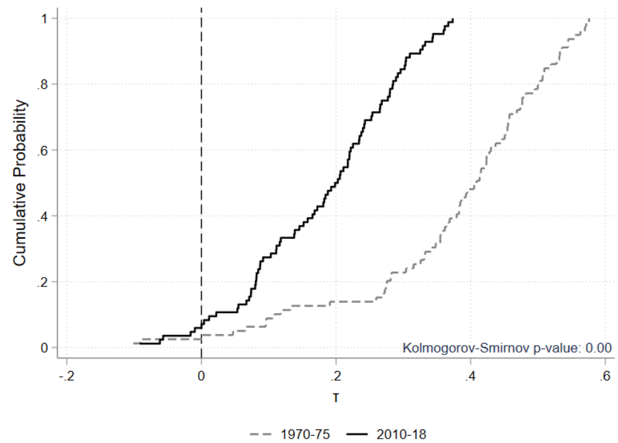

(c) Gender Norms (Circa 1980 and 2010)

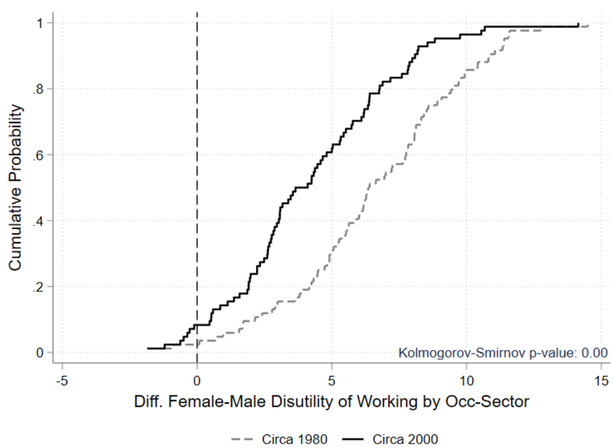

(d) Wage Discrimination (Circa 1980 and 2010)

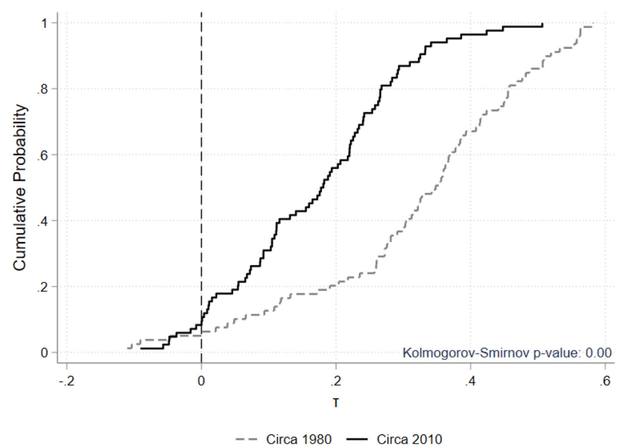

Notes: The above figure plots the CDF of gender norms ($\triangle A_{ojct}$) and female wage discrimination (τ_{ojct}) across sectors and occupations. Figures (a) and (b) plot gender norms and wage discrimination using the first (dotted line) and last year (solid line) for each country. Figures (c) and (d) show the same distributions, but using survey years closest to 1980 and 2010.

B Model Fit and Correlations with Social Norms

Figure B1: Correlations of Excess Gender Norms ($\triangle A$) and Gender Employment Gaps

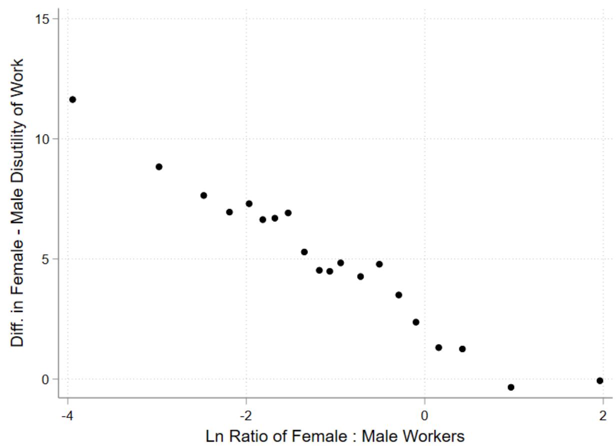

Notes: The above figure shows a binned scatter plot of the correlation between our estimated gender norms between men and women $\Delta A_{oj} = A_{ojf,ct} - A_{ojm,ct}$ and the observed log male to female workers in an occupation-sector across all country-years.

Figure B2: Correlations Wage Penalties (τ) and Gender Wage Gaps

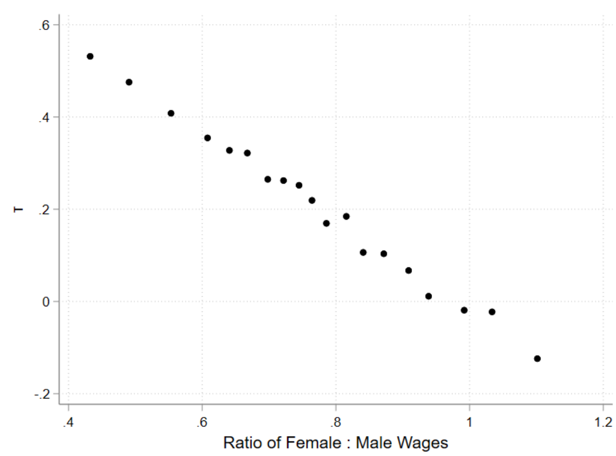

Notes: The above figure shows a binned scatter plot of the correlation between our estimated female wage penalty τ and observed wage gaps in all occupation-sectors and across all country-years.

Figure B3: Correlation of Calibrated Parameters and Targeted Data Moments for Specific Countries Over Time

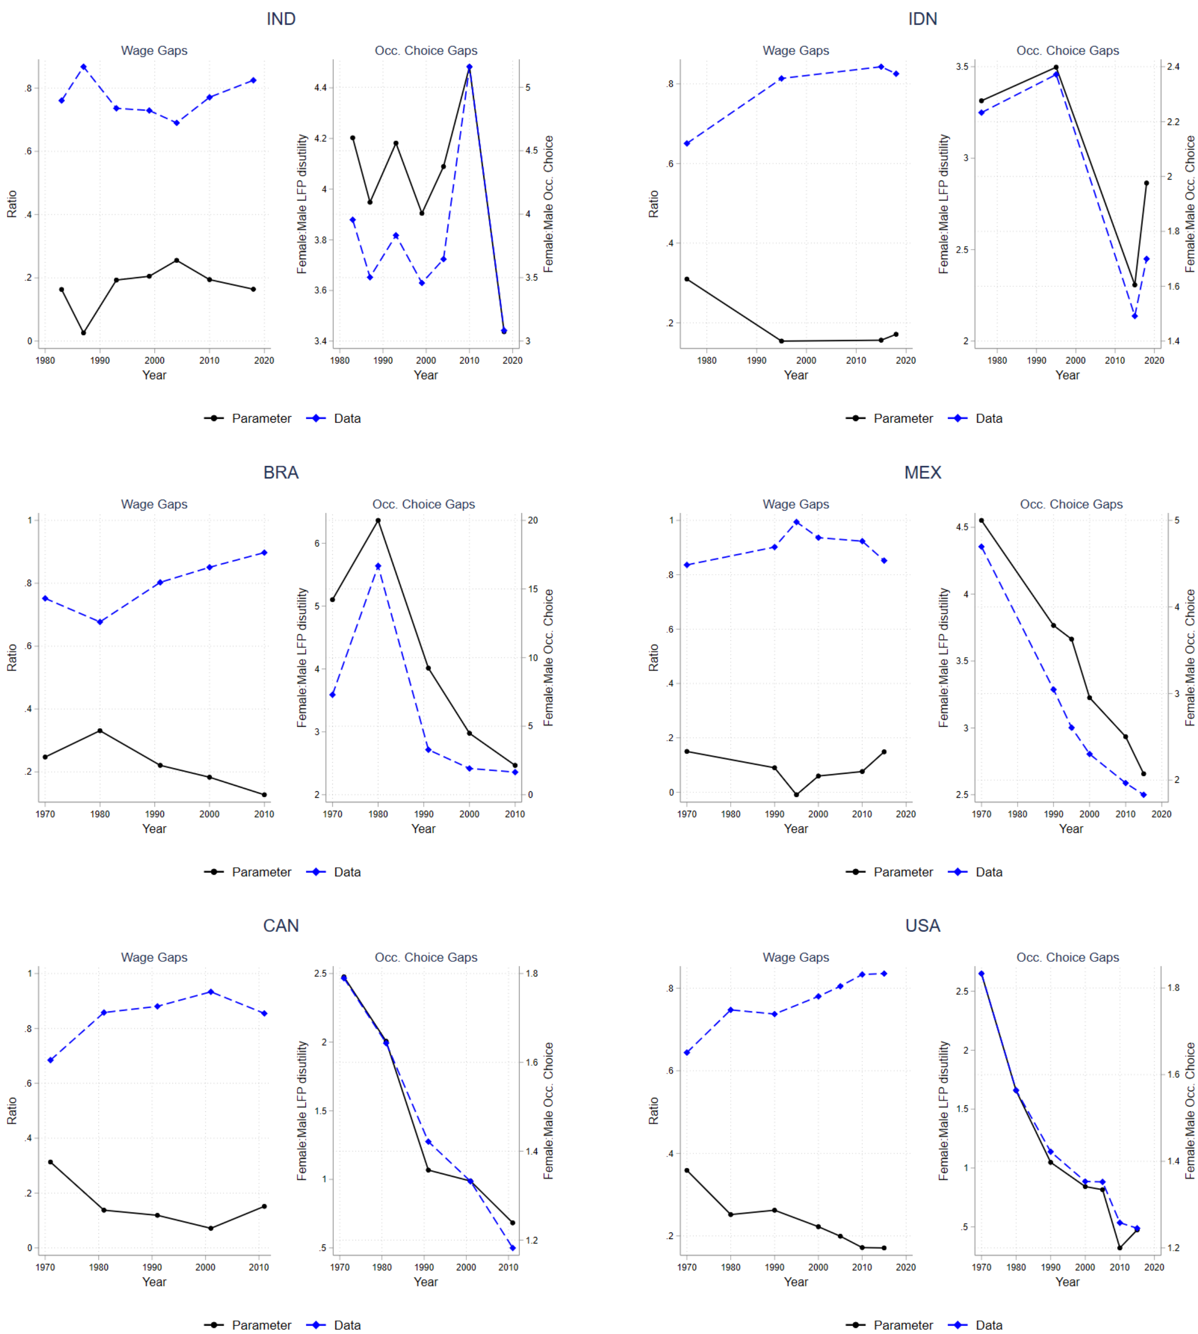

Notes: The above figures plot the average wage and employment ratios (dash blue line) and the estimated values of wage discrimination (τ_{oj}) and gender norms ($\Delta A_{oj} = A_{ojf} - A_{ojm}$) (solid black lines) over time for each country.

Figure B4: Correlations in Value-Added Shares in the Model and Data

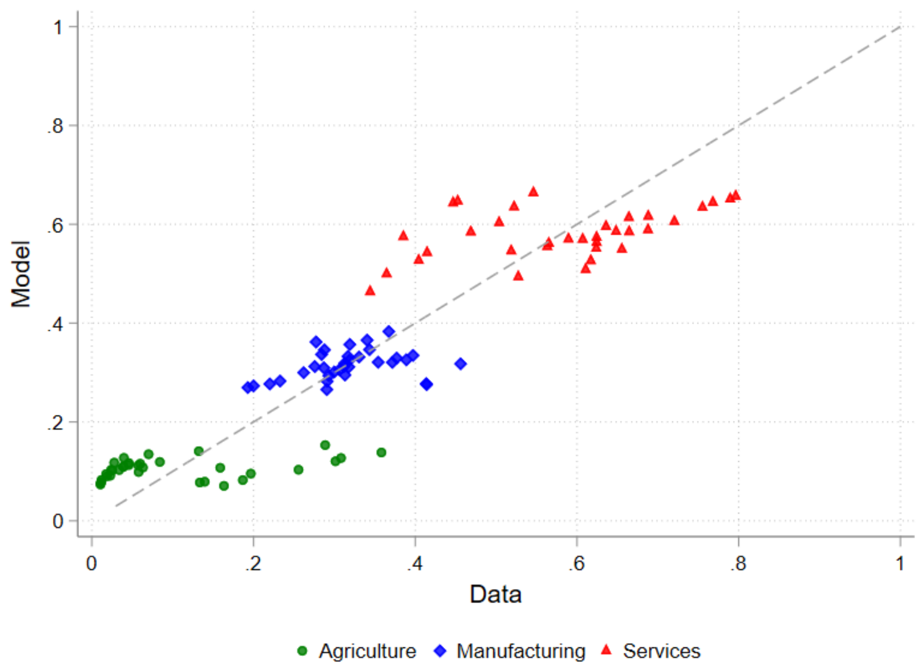

Notes: The above figure shows the correlation between the share of value-added in agriculture (green dot), manufacturing (blue diamond) and services (red triangle) in the data (horizontal axis) and the model (vertical axis) across all country-years in our sample.

Table B1: Correlation between $\triangle A$ and "World, Business, and the Law" Indicators

	Coefficient	S.E.	p-value
	(1)	(2)	(3)
Panel A. Gender Equality in Mobility and LFP			
Index of Mobility/LFP	-0.80	(0.05)	0.00***
Can a woman get a job in the same way as a man?	-0.54	(0.19)	0.10
Can a woman work at night in the same way as a man?	-0.78	(0.08)	0.01**
Can a woman work in a job deemed dangerous in the same way as a man?	-0.71	(0.02)	0.00***
Can a woman work in an industrial job in the same way as a man?	-0.59	(0.11)	0.03**
Panel B. Household Norms			
Index of Household Norms	-0.61	(0.03)	0.00***
Can a woman be head of household in the same way as a man?	-0.76	(0.04)	0.00***
Is there legislation specifically addressing domestic violence?	0.22	(0.11)	0.18
Does a woman have the same rights to remarry as a man?	-0.20	(0.09)	0.16
Do men and women have equal ownership rights to immovable property?	-0.76	(0.04)	0.00***
Observations	510		

Notes: This table shows the OLS correlation between gender norms $\triangle A_{ojct}$ (which we standardize to mean 0 and std dev 1) and indicators of the "World, Business, and the Law" database as described in Equation (11). Standard errors are clustered at the country level. *** is p<0.01, ** is p<0.05 and * is p<0.1.

Table B2: Correlation between τ and "World, Business, and the Law" Indicators

	Coefficient	S.E.	p-value
	(1)	(2)	(3)
Panel A. Gender Equality in Mobility and LFP			
Index of Mobility/LFP	-0.07	(0.02)	0.08*
Can a woman get a job in the same way as a man?	-0.04	(0.04)	0.43
Can a woman work at night in the same way as a man?	-0.04	(0.02)	0.19
Can a woman work in a job deemed dangerous in the same way as a man?	-0.07	(0.01)	0.04**
Can a woman work in an industrial job in the same way as a man?	-0.06	(0.01)	0.03**
Panel B. Gender Equality at the Workplace			
Index of Workplace Equality	0.02	(0.03)	0.51
Does the law prohibit discrimination in employment based on gender?	0.01	(0.01)	0.51
Ln(1+Paid Maternity Days Leave)	-0.03	(0.01)	0.04**
Observations	510		

Notes: This table shows the OLS correlation between female wage discrimination (τ_{oj}) and indicators of the "World, Business, and the Law" database as described in Equation (11). Standard errors are clustered at the country level. *** is p<0.01, ** is p<0.05 and * is p<0.1.

C Model Appendix

C.1 Deriving Aggregate Sectoral Expenditure Shares

Individual Sectoral Expenditure Shares

From Roy's Identity, the Marshallian demand is:

$$\begin{aligned} x_j &= -\frac{\partial V/\partial p_j}{\partial V/\partial I} \\ \Rightarrow \varphi_j(I,p) &= -\frac{\partial V/\partial p_j}{\partial V/\partial I} \times \frac{p_j}{I} \end{aligned}$$

From Equation (1), we get:

$$\begin{aligned} \frac{\partial V}{\partial I} &= \frac{1}{I} \times \left(\frac{I}{P}\right)^{\eta} \\ \frac{\partial V}{\partial p_j} &= -\frac{\omega_j}{p_j}\left(\frac{I}{P}\right)^{\eta} + \frac{\nu_j^h}{p_j} \end{aligned}$$

Defining $P = \prod_j p_j^{\omega_j}$, and substituting in the Roy's identity above, we get that individuals' expenditure share for a sector j is given by:

$$\varphi_j(I_{ojg}(z),p) = \omega_j + \nu_j^h \left(\frac{I_{ojgz}}{P}\right)^{-\eta}$$

C.2 CES Preferences over Home and Market Services

For a sector k where $m \in hs, ms$ i.e., home and market services, the individual's optimization problem can be given by:

$$\begin{aligned} &\min \sum_m p_k C_k \\ &\text{s.t. } C_s = \Big[\sum_k \alpha_k^{\frac{1}{\eta_s}} C_k^{\frac{\eta_s-1}{\eta_s}}\Big]^{\frac{\eta_s}{\eta_s-1}} \end{aligned}$$

Let $P_S = \left[\sum_k \alpha_k p_k^{1-\eta_s}\right]^{\frac{1}{1-\eta_s}}$ and λ be the Lagrange multiplier. Taking the first-order condition and solving we have:

$$\begin{aligned}
\lambda p_k &= \alpha_k^{\frac{1}{\eta_s}} \times \left(\frac{C_k}{C_S}\right)^{-\frac{1}{\eta_s}} \\
\Rightarrow C_k &= \alpha_k (\lambda p_k)^{-\eta_s} C_S \\
\therefore \frac{C_{hs}}{C_{ms}} &= \frac{\alpha_{hs}}{\alpha_{ms}} \times \left(\frac{p_{hs}}{p_{ms}}\right)^{-\eta_s} \\
\Rightarrow \frac{\varphi_{hs}}{\varphi_{ms}} &\equiv \frac{P_{hs}C_{hs}/I}{P_{ms}C_{ms}/I} = \frac{\alpha_{hs}}{\alpha_{ms}} \times \left(\frac{p_{hs}}{p_{ms}}\right)^{1-\eta_s}
\end{aligned}$$

Lastly, substituting back in the constraint, we have:

$$\begin{aligned}
C_S^{\frac{\eta_s-1}{\eta_s}} &= \lambda^{1-\eta_s}\left\{\sum_k \alpha_k p_k^{1-\eta_s}\right\} C_S^{\frac{\eta_s-1}{\eta_s}} \\
\lambda &= \left[\sum_k \alpha_k p_k^{1-\eta_s}\right]^{\frac{-1}{1-\eta_s}} = 1/P_S \\
\Rightarrow C_k &= \alpha_k \left(\frac{p_k}{P_s}\right)^{-\eta_s} C_s \\
\Rightarrow \varphi_k \equiv \frac{p_k C_k}{I} &= \alpha_k \left(\frac{p_k}{P_s}\right)^{1-\eta_s} \varphi_S
\end{aligned}$$

C.3 Aggregation of Sectoral Expenditure Shares

Given PIGL preferences, the expenditure share of an individual of gender g, ability z, working in occupation-sector oj is given by (see Equation (2)):

$$\varphi_j(I_{ojgz}, p) = \omega_j + \nu_j^h \left(\frac{I_{ojgz}}{P}\right)^{-\eta}$$

Therefore, the total expenditure on a sector j and the total income in the economy (across all occupations and gender) can be given by:

$$\begin{aligned}
E &= \sum_j \sum_o \sum_g N_g \int_z I_{ojg}(z) \Pr(oj|z,g) dF(z) \\
E_j &= \sum_o \sum_g N_g \int_z \varphi_j \times I_{ojg}(z) \Pr(oj|z,g) dF(z) \\
&= \sum_g N_g \left\{\omega_j \sum_o \int_z I_{ojg}(z) \Pr(oj|z,g) dF(z) + \nu_j P^\eta \sum_o \int_z I_{ojg}(z)^{1-\eta} \Pr(oj|z,g) dF(z)\right\}
\end{aligned}$$

Define:

Aggregating over returns to ability: $E(z^x) = \int_z z^x \Pr(oj|z,g)dF(z)$

Total supply of human capital: $H_{ojg} = E(z^{\kappa_{ojg}}$ (from Equation (**??**)

Total income in a sector j: $I_j^{TOT} = \sum_o \sum_g w_{ojg} H_{ojg}$

Total income in the economy: $I^{TOT} = \sum_j I_j^{TOT}$

Income share of a sector j: $\iota_j = I_j^{TOT}/I^{TOT}$

Therefore:

$$
\begin{aligned}
\Phi_{jg} = \frac{E_j}{E} &= \omega_j \times \frac{\sum_o \int_z I_{ojg}(z)\Pr(oj|z,g)dF(z)}{\sum_j \sum_o \int_z I_{ojg}(z)\Pr(oj|z,g)dF(z)} + \nu_j P^\eta \times \frac{\sum_o \int_z I_{ojg}(z)^{1-\eta}\Pr(oj|z,g)dF(z)}{\sum_j \sum_o \int_z I_{ojg}(z)\Pr(oj|z,g)dF(z)} \\
&= \omega_j \times \frac{\sum_o \sum_g N_g w_{ojg} E(z^{\kappa_{ojg}})}{\sum_j \sum_o \sum_g N_g w_{ojg} E(z^{\kappa_{ojg}})} + \nu_j P^\eta \times \frac{\sum_o \sum_g N_g w_{ojg} E(z^{\kappa_{ojg}(1-\eta)})}{\sum_j \sum_o \sum_g N_g w_{ojg} E(z^{\kappa_{ojg}})} \\
&= \omega_j \times \frac{\sum_o \sum_g N_g w_{ojg} E(z^{\kappa_{ojg}})}{\sum_j \sum_o \sum_g N_g w_{ojg} E(z^{\kappa_{ojg}})} \\
&+ \nu_j \times \underbrace{\frac{\sum_o \sum_g N_g w_{ojg} E(z^{\kappa_{ojg}(1-\eta)})}{\left[\sum_j \sum_o \sum_g N_g w_{ojg} E(z^{\kappa_{ojg}})\right]^{1-\eta}}}_{=g(z^{\kappa_{ojg}})} \times \left[\frac{\sum_j \sum_o \sum_g N_g w_{ojg} E(z^{\kappa_{ojg}})}{P}\right]^{-\eta} \\
&= \omega_j \iota_j + \nu_j \times g(z^{\kappa_{ojg}}) \times \left[\frac{I^{TOT}}{P}\right]^{-\eta}
\end{aligned}
$$

D Estimation Algorithm

Here we describe the numerical procedure that we use to estimate key model parameters by fitting our model's equilibrium conditions to data moments. With this procedure, we estimate the following parameters: PIGL preferences $\mathscr{U} = \{\omega_j, \nu_j\}_{\forall j}$, sectoral productivity $\{B_j\}_{\forall j}$, and gender barriers $\mathscr{B} = \left\{\{\tau_{oj}\}_{\forall oj}, \{A_{ojg}\}_{\forall ojg}\right\}$. To calibrate these parameters, we exactly match data on men's and women's occupation-sector choices $\Pr(oj|g)$, gender gaps in average hourly wages ($\overline{\text{wage}}_{ojf}/\overline{\text{wage}}_{ojm}$) and growth in sectoral real value added ΔY_j. We use data on sectoral value added shares to calibrate the PIGL preference parameters.

Outer loop: Guess sectoral productivity B_j and PIGL parameters (ω_j, ν_j).

Inner loop: Guess occupation-sector wage rates w_{oj}, sectoral prices p_j, amenities A_{ojg}, and female wage discrimination τ_{oj}.

Step 1: Compute income $I_{ojgz} = w_{ojg} z^{\kappa_{oj}}$ and indirect utility $V(I_{ojgz}, p)$ for each gender-ability type gz using:

$$V(I_{ojgz}, p) = \frac{1}{\eta}\Big[\frac{I_{ojgz}}{\prod_j p_j^{\omega_j}}\Big]^{\eta} - D(p), \tag{12}$$

and compute occupational choices for each gender-ability-type using:

$$\Pr(oj|g,z) = \frac{\exp\left[\frac{1}{\sigma_\varepsilon} V(I_{ojgz}, p) - \frac{1}{\sigma_\varepsilon} A_{ojg}\right]}{\sum_{j'}\sum_{o'} \exp\left[\frac{1}{\sigma_\varepsilon} V(I_{o'j'gz}, p) - \frac{1}{\sigma_\varepsilon} A_{o'j'g}\right]}. \tag{13}$$

Step 2: Integrate these choice probabilities across z-types and solve for a new guess of amenities A_{ojg}^{new} to perfectly fit men's and women's observed employment shares in each occupation-sector $\Pr(oj|g)$.

Step 3: Compute average human capital in each occupation-sector and for each gender (taking into account how workers' selection into occupation-sectors is driven by their comparative advantage $z^{\kappa_{oj}}$ and other factors), using:

$$\overline{H}_{ojg} = \int_z \Pr(oj|g,z) z^{\kappa_{oj}} dF(z). \tag{14}$$

Step 4: Solve for a new guess of female wage discrimination τ_{oj}^{new} to perfectly fit observed gender wage gaps in each occupation-sector using:

$$\frac{\overline{\text{wage}}_{ojf}}{\overline{\text{wage}}_{ojm}} = (1-\tau_{oj}) \times \frac{\overline{H}_{ojf}}{\overline{H}_{ojm}}. \tag{15}$$

Step 5: Solve for a new guess of occupation-sector-specific wage rates w_{oj}^{new} from firms' first order condition:

$$w_{oj} = \frac{\gamma_{oj} p_j B_j \prod_o H_{oj}^{\gamma_{oj}}}{H_{oj}} \tag{16}$$

Step 6: Compute aggregate sectoral expenditure shares:

$$\Phi_j = \omega_j + \nu_j \frac{P^\eta}{I} \sum_o \sum_j \sum_g \int_z I_{ojgz}^{(1-\eta)} \frac{N_g}{N} \Pr(oj|g,z), \tag{17}$$

where $P = \prod_j p_j^{\omega_j}$ and $I_{ojgz} = w_{ojg} z^{\kappa_{oj}}$.

Step 7: Solve for a new guess of sectoral prices p_j^{new} to ensure that good markets clear in each sector:

$$p_j^{new} = \frac{\Phi_j \times I}{B_j Y_j} \tag{18}$$

Iterate on the inner loop until convergence.

For each convergence of the inner loop, we proceed with the outer loop which solves for PIGL preference parameters and sectoral productivity.

Step 8: Regress observed shares of sectoral value-added (ETD data) on a constant and a model-implied measure of real income according to Equation 17 using data across multiple countries and over time. The regression constant provides a new guess for ω_j^{new} and the coefficient provides a new guess for ν_j^{new}.

Step 9: Solve for a new guess of sectoral productivity B_j^{new} to match observed growth in secotal real value added ΔY_j in each country over time using the sector's production function (and normalizing $B_{j0} = 1$ in the first year of each country):

$$Y_j = B_j \prod_o H_{oj}^{\gamma_{oj}}. \tag{19}$$

E Data Appendix

E.1 Sample Definiton and Industry-Occupation Classifications

E.1.1 Sample Definition

1. We restrict the sample between the age of 18 to 65 years old.

2. For a small share of observations, we do observe the industry and occupation of their current/most recent job, but they are coded as "unemployed". This is mostly due to the recall periods differing on the survey. We therefore set them to "employed".

3. We drop those individuals in school, prision, disabled, ill, etc.
   keep if empstat$\geq$*1 & empstat* ≤ 3.

4. We classify all individuals who are unemployed or out of the workforce in the "home sector" i.e., *empstat == 2 | empstat == 3*

5. Table D1 provides the classification of education categories into years of education.

Table D1: Classification of Education

Code	Education	Years
0	NIU (not in universe)	NA
100	Less than primary completed	2
110	No schooling	1
120	Some primary	3
130	Primary (4 years)	4
211	Primary (5 years)	5
212	Primary (6 years)	6
221	General and unspecified track	9
222	Technical track	9
311	General track completed	12
312	Some college/university	14
320	Technical track	14
321	Secondary technical degree	12
322	Post-secondary technical education	14
400	University Completed	16
999	Unknown/Missing	NA

E.1.2 Industry and Occupation Classifications

1. Tables D2 and D3 provide the classification of industries and occupations

2. We re-classify some occupations since they are sparsely represented in industries (for example, professionals in agriculture). Currently, we are using the following:

   - Professionals and Clerks in Agriculture are assigned to Services:
     ind = 4 if ind == 2 occ≥1 occ ≤ 3

   - Agri Fisheries in Manufacturing reassigned to Agriculture
     ind = 2 if ind == 3 occ==4

   - Crafts/Trade Workers & Plant & Machine Operators in Agriculture re-assigned to Manufacturing i.e., *ind = 3 if ind == 2 occ≥ 5 occ ≤ 6*

3. Some individuals report an occupation but not the industry. Where a clear mapping exists, we classify them in the correct industry.
   Agriculture: *ind = 2 if missing(indgen) occisco == 4*
   Manufacturing: *ind = 3 if missing(indgen) & (occisco == 5 | occisco == 6)*
   Services: *ind = 4 if missing(indgen) & occisco ≤ 3*

Table D2: Classification of Industry Codes in IPUMS

Code	Industry	Classification
10	Agriculture, fishing, and forestry	Agriculture (1)
20	Mining and extraction	Manufacturing (2)
30	Manufacturing	Manufacturing (2)
40	Electricity, gas, water and waste management	Manufacturing (2)
50	Construction	Manufacturing (2)
60	Wholesale and retail trade	Services (3)
70	Hotels and restaurants	Services (3)
80	Transportation, storage, and communications	Services (3)
90	Financial services and insurance	Services (3)
100	Public administration and defense	Services (3)
110	Services, not specified	Services (3)
111	Business services and real estate	Services (3)
112	Education	Services (3)
113	Health and social work	Services (3)
114	Other services	Services (3)
120	Private household services	Services (3)
130	Other industry, n.e.c.	Services (3)
998	Response suppressed	NA
999	Unknown	NA

Table D3: Classification of Occupations

Code	Occupation	Classification
1	Legislators, senior officials and managers	Professional (1)
2	Professionals	Professional (1)
3	Technicians and associate professionals	Professional (1)
4	Clerks	Clerks (2)
5	Service workers and shop and market sales	Services Workers (3)
6	Skilled agricultural and fishery workers	Skilled Agri. (4)
7	Crafts and related trades workers	Craft/Trade Wrkrs (5)
8	Plant and machine operators and assemblers	Plant & Machine (6)
9	Elementary occupations	Elementary (7)
10	Armed forces	Drop
11	Other occupations, unspecified or n.e.c.	Drop
97	Response suppressed	Drop
98	Unknown	Drop
99	NIU (not in universe)	Drop

E.2 Calculating Hourly Wages

We discuss the calculation of hourly in three steps. First, we define the different ways in which income is measured in the IPUMS data. We then discuss the measurement of hours worked. Lastly, we discuss the available of the income and hours measurement specific to the countries in our sample. Usually, within a country, the measurement does not change over time.

E.2.1 Measurement of Income Earned

There are three types of income variables available in the IPUMS:

1. **INCTOT:** reports the person's total personal income from all sources in the previous month or year.
2. **INCEARN:** reports the person's total income from their labor (from wages, a business, or a farm) in the previous month or year.
3. **INCWAGE:** reports the respondent's weekly, monthly or annual wage and salary income.

INCTOT is most commonly available across almost all country-years. Therefore, to maintain consistency across the definition of income in our sample, we use INCTOT where available, even if others are available. All income variables are reported in local currency units, with varying frequency (as we will discuss below). We remove extreme outliers in the income distribution (above 9 million LCU).

E.2.2 Measurement of Hours Worked

1. Weekly hours are provided in most country-year surveys for all individuals in the workforce (HRSWORK1 variable in IPUMS). We trim the sample at 100 hours.
2. When hours worked are missing for an individual, but available for the country-year sample, we replace it by the gender-industry-occupation average within that country-year.

3. In case no information is available on hours worked, we set it to 40 hrs/week. From other country-years, the average varies between 40-45 hrs/week, so it is not a bad approximation.

4. In cases where income is available at the monthly or annual frequency, we assume an individual works for 4 weeks/month and 52 months/year. In some cases (USA and Canada) we do observe the number of months worked, which we use to calculate the wages.

E.2.3 Countrywise Availability of Income and Hours

Brazil: 1970, 1980, 1990, 2000, 2010

1. Income: INCTOT and INCEARN from all sources in the previous month are available. As discussed earlier, to be consistent across countries, we use INCTOT whenever reported by the individual. If not, we replace it by INCEARN.

2. Hours worked: Available for 1991, 2000 and 2010.

3. Wage = Income/(4*Hrs Worked)

Canada: 1971, 1981, 1991, 2001, 2011

1. Income: INCTOT, INCEARN, INCWAGE and INCSELF are available. INCEARN = INCWAGE + INCSELF. Like previously, we use INCTOT when reported and INCEARN in case it is missing.

2. Hours worked: Available 1981 onwards. Moreover MONTHSWRK is also available, which we use to construct the number of months worked by the individual

3. Wage = Income/(4*Hrs Worked*Months Worked)

India: 1983, 1987, 1993, 1999, 2004, 2018

1. Income IPUMS: INCWAGE is reported for the reference week.

2. Hours worked IPUMS: Not available for any of the rounds. Days worked data (DAYSWRK) available for 1993.

3. Wage IPUMS: We assume an individual works for 8 hrs a day × number of days (when reported). Otherwise, we set it to 40 hrs/week and define: Wage = Income/Hrs Worked.

4. For 2018, we complement the IPUMS data for India with data from the Indian Periodic Labor Force Survey (PLFS). We pool data from 2017-2019 to increase the sample size.

5. Wages and hours worked in PLFS 2018: Wage income is available in the PLFS for 2017, 2018, 2019 and is defined as "Last wage payment, primary job, excl. bonuses, etc (7-day ref period)". The PLFS further reports hours worked in the last week, which allows us to compute hourly wages.

Indonesia: 1976, 1995, 2018

1. Income: INCWAGE from the previous month is available for both the years.

2. Hours worked: Hours worked are available for 1995.

3. Wage = Income/(4*Hrs Worked)

4. For 2018, we complement the IPUMS data for Indonesia with data from the SAKERNAS survey. We pool data from 2017-2019 to increase the sample size.

5. Wages and hours worked in SAKERNAS 2018: Wage income is available for 2017, 2018, 2019 and is defined as "Last wage payment, primary job, excl. bonuses, etc (7-day ref period)". The SAKERNAS further reports hours worked in the last week, which allows us to compute hourly wages.

Mexico: 1970, 1990, 1995, 2000, 2005, 2010

1. Income: INCTOT and INCEARN in the previous month. Not available for 2005.

2. Hours worked: Available for 1990, 1995, 2000, 2010

3. Wage = Income/(4*Hrs Worked)

USA: 1970, 1980, 1990, 2000, 2005, 2010

1. Income: INCTOT and INCWAGE in the previous year available for all years. INCWAGE and INCSELF are also available, but I have checked that INCEARN = INCWAGE + INCSELF, so we donât need these two variables separately.

2. Hours worked: Available 1970 onwards. MONTHSWRK is also available for all years

3. Wage = Income/(4*Hrs Worked*Months Worked)

E.2.4 Home Sector & Real Wages

1. We trim the wage earnings within each country-year-industry-occupation-gender at the 1st and 95th percentile.

2. We impute the gender-specific wages for the "home sector" using the average wages in elementary occupations in the services sector within each country-year.

3. We set the returns to ability (κ) at home to be equal to 1 across both men and women.

4. We use the exchange rates (LCU/USD) and real GDP at current and constant prices from the World Bank data to convert all wages in LCU to real 2010 USD as follows: $w_{USD} = \frac{w_{LCU}}{ExchangeRate} \times \frac{GDP_{Constant\ 2010}}{GDP_{Current}}$.

5. Lastly, while aggregating, we use the person weights provided by the sample surveys to make the estimates representative of the population in that country-year.