Gender Barriers, Structural Transformation, and Economic Development*

Gaurav Chiplunkar
University of Virginia

Tatjana Kleineberg
World Bank

July 2024

PRELIMINARY. PLEASE DO NOT SHARE

Abstract

Using nationally representative data across many countries spanning five decades between 1970-2018, we document distinct *gender* patterns in the process of structural transformation – both across sectors, and within sectors across occupations. While gender gaps in occupational and sectoral employment have declined over time, they remain salient even today, and even in the most advanced economies. Gender segregation in employment is largely driven by variation *across* sectors in low- and middle-income countries, but across occupations *within* sectors in high-income countries. Gender wage gaps have declined more slowly, remain persistent, and have no clear correlation with economic development. Interpreted through the lens of a Roy model and a development accounting framework, the reductions of gender barriers during the past decades — defined as gender-specific wedges in employment and remuneration — account for around a third of the observed employment transitions towards manufacturing and services on average, with significant heterogeneity across countries.

Keywords: Gender, Structural Transformation, Economic Development

*We would like to thank Cheryl Doss, Kevin Donovan, Doug Gollin, Charles Gottlieb, Michael Peters, Barbara Petrongolo, Markus Poschke, Claudia Olivetti, and participants in conferences and seminars at the World Bank, Northwestern University, Cornell University, Dartmouth University, Boston University, American University, Richmond Federal Reserve, University of Virginia, CEPR STEG Annual Workshop, NEUDC, ASSA, SED, and WADES 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

As countries grow richer, the structure of their economies experience fundamental changes as workers transition from agricultural work into manufacturing and service employment. This process of structural transformation as well as its drivers and implications for aggregate productivity and growth have been extensively studied in the literature (e.g., Kuznets (1973); Maddison (1980)). Key drivers of structural transformation include income effects in consumption, sector- or factor-specific technological change, accumulation of human capital, etc.[1]

In this paper, we focus on the role of gender in understanding the process of structural transformation. We begin by providing new evidence on the gender dimension of sectoral and occupational employment transitions for individuals in over 90 countries, spanning multiple decades, and covering a wide spectrum of economic development. We document that men and women exhibit distinct patterns, both in moving across sectors (agriculture to non-agriculture), as well as across occupations (elementary to professional/managerial ones). We further document that while gender wage gaps in occupations and sectors across countries have been decreasing over time, they remain substantial and persistent even today, and even in the most advanced economies. Second, we develop a unified framework that allows us to parsimoniously quantify the role of: (i) economic/technological progress; and (ii) evolving gender norms over time, in driving the patterns that we observe in the data. We calibrate the model for a subset of countries where we observe rich data on education, hourly wages, and employment choices of men and women, and find that a reduction in the barriers faced by women in the labor market (both in employment choices as well as wages) have played a non-trivial role in explaining the economic growth in these countries. On average, changes in gender barriers explain around half the growth in services and manufacturing sectors, as well as around half the growth in professional and managerial jobs in these countries. Put together, they account for around a quarter of the the output growth of the service sector. However, these effects are heterogeneous across countries: for example, in India–which experienced rapid, service-led economic growth over the past decades–gender barriers account for only 6% of the output growth in the service sector. On the other hand, gender barriers account for 27% of the observed output growth in the service sector in the U.S. and 38% in Brazil.

To document new facts on the gender dimension of structural transformation, we collect a data set that relies on censuses, household, and labor force surveys that covers 91 countries and spans almost five decades from 1970 to 2018. In terms of empirical evidence, we first show that the canonical process of structural transformation, i.e., employment transitions from agriculture to manufacturing and services with economic development, is primarily driven by men. Women follow a different pattern: at low levels of development, women exit agriculture and first sort into the home sector. At higher development levels, women then enter the labor force again and mostly sort into the (market) service sector. Unlike their male counterparts, women's employment share in manufacturing remains small and relatively constant across development levels.

Importantly, we also find important gender differences in occupational choices across countries and over time. Changes in occupational employment have received less attention in the literature on structural trans-

[1] See Ngai and Pissarides (2007); Olivetti (2014); Ngai, Olivetti and Petrongolo (2024); Huneeus and Rogerson (2024); Lagakos and Shu (2023); Herrendorf, Rogerson and Valentinyi (2013, 2014); Comin, Lashkari and Mestieri (2021); Acemoglu and Guerrieri (2008); Caunedo and Keller (2023); Porzio, Rossi and Santangelo (2022).

formation, but the occupational and sectoral dimension is tightly linked as we show that sectors differ in their occupational employment structure which changes substantially over time. Across 1 digit ISCO occupation codes, we find that men are more likely to work in crafts and trade, machine operating, and professional occupations, while women are over-represented in clerical occupations. Women's employment in high-skilled occupations is much larger in high- compared to low-income countries: the share of women working in professional and clerical occupations is half that of men in low-income countries, while it is equal in professional occupations and up to 3 times higher in clerical occupations in high-income countries. A convergence in occupational choices across genders during the past decades has also been emphasized in the macro literature as an important determinant of talent allocation and aggregate productivity (Hsieh et al., 2019).

To measure gender segregation across occupations and sectors more systematically, we compute the Theil entropy index of segregation. Similar to the Gini index, this metric quantifies the degree to which an observed distribution of men and women across occupation-sectors deviates from an egalitarian allocation. We find an inverted-U shaped in gender segregation along countries' development levels. Gender segregation first increases between low- and middle-income countries and then decreases again in high-income countries. In poor countries, there is little division of labor and most men and women work in agricultural jobs, resulting in low gender segregation. With economic development, women transition out of the labor force, while men move into manufacturing and service jobs, thus increasing gender segregation in middle-income countries. At advanced levels of economic development, segregation declines again as more men and women find employment in the service sector. A key advantage of the entropy index is that it can be additively decomposed into the share of gender segregation that can be attributed to variation across sectors as opposed to within a sector across occupations. We find that only 15-25% of gender segregation is explained by differences across occupations within sectors in low- and middle-income countries, which increases to more than 60% in high-income countries. Together, the evidence shows that gender differences in occupational and sectoral employment choices are large and that each dimension matters differentially in low- and high-income countries, emphasizing the importance of studying occupational and sectoral employment choices simultaneously when studying a large cross-section of countries and long periods of time. To put it more bluntly: a scenario in which all women in the service sector work on lower rungs of the job ladder (e.g., as cleaners) or one in which women are equally represented in top positions (e.g., in managerial jobs) has very different implications for the effective supply of human capital in the sector and the general allocation of talent in the labor market.

To complement the cross-country evidence, we then narrow our focus to a "core sample" of six large countries (namely, India, Indonesia, Brazil, Mexico, Canada, and the U.S.) which cover a wide range of the development spectrum. A key advantage of these countries is that we can obtain high-quality individual-level data on employment, education, and hourly wages for five consecutive decades. We leverage these data to document how gender differences in employment and wages evolve over time *within* each of these countries.[2] Over the past five decades, female labor force participation increased substantially in all countries except India. However, female relative to male employment shares increased very unevenly across sectors and occupations with a substantial increase in services in all countries except India, but with little or no changes in agriculture and manufacturing. Focusing on occupations, we find that women's representation in professional occupations (relative to men) is higher in the richer countries of our sample and

[2]This data availability further allows us to estimate our structural model for this core sample of six countries, as we will describe in more detail below.

improvements over time are also steeper in the richer countries (US, Canada, but also Mexico) compared to the poorer countries of our sample (India and Indonesia). A similar pattern holds for clerical occupations, in which women are substantially over-represented compared to men in recent periods for the US, Canada, Mexico and Brazil.

We further document gender wage gaps across sectors and occupations and over time for the countries of our core sample. We find that significant gender wage gaps still persist even in the most recent period and in the richest countries of our sample, which is in line with the findings of a large literature on gender wage gaps (Fabrizio et al., 2018; OECD, 2023). Moreover, unlike the improvements we saw in gender employment gaps in richer countries, we find that there is little correlation between gender wage gaps and countries' economic development. For example, we find remarkably similar gender wage gaps for Indonesia, Brazil, Canada and the USA in the most recent period of our sample with women earning on average between 75 and 79 cents for each dollar earned by men. Similarly, gender employment gaps and gender wage gaps are not necessarily correlated across occupations or sectors: for example, women are most represented in the service sector, while their wage gap relative to men is similar or even larger in services (compared to other sectors). These findings show that similar trajectories of economic growth across countries can differentially affect men's and women's employment choices. For instance, the expansion of the service sector in the United States contributed to a reduction in gender employment gaps (as shown in Ngai and Petrongolo (2017)), while the service-led growth in India resulted in falling female labor force participation over the past decades.

Several economic channels–which are commonly-cited and studied in the literature on structural transformation–can help to partially explain these patterns in men and women's employment and wages over time. First, technological change could be biased towards specific sectors or occupations where women are initially more represented. Second, men and women's education choices might change over time which would impact their supply of effective human capital or returns to human capital might change across sectors or occupations due to skill-biased technological change. Third, as individuals and countries grow richer, income effects could tilt consumers' expenditure shares away from agriculture (a necessity) and towards the service sector (a luxury), where women are initially more represented. In addition, gender-specific mechanisms can also directly affect women's employment choices and wages relative to men and these forces can also change over time. Similar to Hsieh et al. (2019), we consider gender-specific barriers in the form of gender-specific "preferences"–that affect the utility cost of working in a given occupation-sector pair–and in the form of female wage discrimination which affects women's remuneration relative to an equivalent man when working in the same occupation-sector pair.

To distinguish between "general" economic and gender-specific mechanisms, we develop a theoretical framework that parsimoniously incorporates each of the above-mentioned channels, thus allowing us to study the role of gender in the economy and its aggregate implications. The aim of the theoretical exercise is twofold: first, to quantify the magnitudes of (and changes in) gender norms and wage discrimination across countries and over time, *after* taking into account differences in the economic fundamentals in them; and second, to simulate counterfactual growth trajectories for countries that can help assess the importance of gender barriers in impacting their process of structural transformation and economic growth. In a nutshell, individuals in our model differ by gender and ability (schooling). Conditional on the fundamentals of the economy, captured by the economic and non-economic channels described previously, they choose: (i)

whether to participate in the labor force or not; and (ii) conditional on working, their occupation and sector of work. On the demand side, we incorporate non-homothetic preferences that allow for sectoral expenditure shares to vary by income (Alder et al., 2022; Fan et al., 2021). Wages and prices are then determined in equilibrium. The non-economic channels therefore distort workers' employment choices, thus misallocating talent in the economy, and potentially lowering productivity and growth.

We then take the model to the data. Our calibration strategy imposes minimal restrictions on the fundamental economic and non-economic parameters i.e., we do not impose any dynamic restrictions on their evolution, their relationship with the level of economic development, or whether one gender systematically has an (dis-)advantage or faces more constraining barriers as compared to the other. In the spirit of an accounting exercise, gender barriers (norms and discrimination) within each occupation-sector (country-year) are then then calibrated as residual wedges that match the observed gender gaps in employment and wages, *after* accounting for economic factors. As we demonstrate in Section 7.2, despite being residuals, our estimates of gender barriers strongly correlate with empirically measured social norms and labor market constraints in the World Bank's "World, Business, and the Law" (WBL) database (World Bank, 2019, 2020; Hyland et al., 2020).

We find that gender norms, i.e, the excess utility costs faced by women for participating in occupation-sectors (relative to men), are on average lower in richer countries than poorer ones. While they have declined across all countries in our sample since the 1970s, the decline has been largest in middle-income (Brazil and Mexico) and high-income (Canada and USA) countries as opposed to low-income ones (India and Indonesia). Furthermore, gender norms in the service sector, and especially for clerical jobs, are lowest in middle- and high-income countries, while they are lowest in agricultural jobs in low-income countries. For wage discrimination, we find women in the 1970s earned around 53% of their marginal product in manufacturing and services, and 58% of their marginal product in agriculture. Wage discrimination was highest in trade and service occupations (50%) and professional/managerial jobs (45%), and lowest in clerical jobs (30%). While wage discrimination has decreased over time across most occupation-sectors and countries, it continues to persist even today and even in the most developed countries. 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.

Last, we use our estimated model to quantify the importance of gender barriers in explaining the observed paths of structural transformation and growth during the past decades in our core sample of six countries. To do so, we hold wage discrimination and gender norms *fixed* at the levels that we calibrated in the first year (in the 1970s) for each country, while allowing for all other (economic) parameters to evolve according to their calibrated values. To assess the importance of changes in gender barriers, we then compare the counterfactual path of structural transformation and economic growth for each country, to the trajectory that we observe in the data. We find that changes in gender barriers had large effects on labor force participation, structural transformation, and economic growth in all countries except India. Averaging across the six countries of our sample, we find that changes in gender barriers explain 44% of the observed increase in manufacturing employment shares, and 52% of the observed increase in the service sector employment share. Effects on agricultural employment are only very small. Across occupations, changes in gender barriers explain around half of the increase in professional, crafts and trade, and machine operating occupations.

Overall, improvements in gender barriers over time account for a substantial share of observed output (real value added) growth across sectors, explaining 4% of agricultural output growth, 10% of manufacturing output growth, and 23% of output growth in the service sector when averaging across the six countries of our core sample. However, these effects are heterogeneous across countries: for example, India had rapid, service-led growth during our sample period, but gender barriers account only for 6% of the output growth in the service sector. On the other hand, gender barriers account for 27% of the observed output growth in the service sector in the U.S. and 38% in Brazil.

Put together, we draw the following conclusions from our analysis: first, there are important differences in how men and women's employment choices transition across sectors and occupations along countries' development process with important heterogeneity across countries. Second, we find that significant gender barriers continue to persist even today, and even in the most advanced economies, despite improvements over time in most countries. Third, we find that changes in gender barriers had important effects on structural transformation and economic development. Reductions of gender barriers have substantially contributed to the observed process of structural transformation in most countries of our sample and they explain, on average, 23% of the observed output growth in the service sector. Last, we find that important differences exist across countries in the levels and trends of gender barriers, which warrants further comparative micro-economic studies to examine the sources of similarities and differences.

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 contributes to several strands of the literature that study structural transformation, occupational choices, talent allocation, and the importance of gender roles in driving economic choices and outcomes.

A large 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 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.

Another literature studies workers' reallocation from agriculture to manufacturing and then to services. More recently, several papers have also focused on the role of gender in driving structural change (Cuberes and Teignier, 2014, 2016; Moro, Moslehi and Tanaka, 2017; Olivetti, 2014; Ngai and Petrongolo, 2017; Rendall, 2018; Ngai, Olivetti and Petrongolo, 2024), or a reversal of gender gaps in education as countries grow richer (Ying et al., 2023). Several papers study how men and women allocate their time between market

and home work within the household by collecting rich data from time-use surveys and using this data to analyze the implications of this time allocation for countries' structural transformation and economic development (cf. (Gottlieb, Doss, Gollin and Poschke, 2019) for a cross-country analysis and (Ngai, Olivetti and Petrongolo, 2024) for a historical analysis in the US). We contribute to this literature by documenting gender differences in employment choices and wages across industries and occupations and by quantifying the implications of these outcomes on aggregate productivity and output. Our model incorporates the occupational and sectoral dimension and allows us to quantify the effect of "non-economic" gender barriers (as opposed to economic driving forces) on countries' structural transformation and economic development.

In this spirit, our paper relates to Lee (2022), who 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% and GDP per-capita by 3.6%. 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. In that sense, it is similar to, and extends the analysis by Hsieh, Hurst, Jones and Klenow (2019), who show that a reduction in gender barriers explains 20-40% of overall growth in the US from 1960 to 2010. During the same period, the economic structure of the US–and of other countries–experienced significant transformations across sectors which can affect men and women's employment outcomes differentially due to initial gender differences in sectoral employment choices and due to differential improvements in gender barriers across occupations and sectors. Our paper therefore contributes to the literature by modeling structural transformation–and its commonly-studied drivers–jointly with occupational choices and a potentially resulting misallocation of talent.

3 Data

Data Sources and Sample Description

Our primary data source is 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 years. Our final sample contains 91 countries and 288 country-years which range from 1970 to 2018 and which include, on average, 3 rounds of data for each country. Data coverage increases in recent decades (1990 onwards) but previous decades are also well represented in our sample (as shown in Appendix Table D1).

The data requirements for our quantitative exercise are even higher as we need information on workers' hours worked and compensation in addition to their education and employment variables. In addition, we want to observed these employment and wage outcomes across occupations, sectors, and genders for a long time period (1970-2015) to study the process of structural transformation. We therefore focus on six "core countries" for our quantitative application which fulfill these data requirements. This sample includes India, Indonesia, Mexico, Brazil, Canada, and the United States; hence, covering a wide range of the income spectrum and accounting for 25-30% of the world population. 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 (around 2018). Appendix D provides more information on the data construction.

Classification of Sectors and Occupations

We aggregate the harmonized sector classifications from the IPUMS and labor force survey data into three sectors: (a) Agriculture; (b) Manufacturing and (c) Market Services as shown in Table D4 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 for structural transformation and for home production (Moro, Moslehi and Tanaka, 2017; Ngai and Petrongolo, 2017; Bridgman, Duernecker and Herrendorf, 2018). The underlying data classifies occupations at the 1 digit ISCO occupation codes as reported in Table D5. We aggregate the top three 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 aggregate the top three ISCO occupation codes (senior officials, directors, managers, technicians, etc) due to small sample sizes in some country-years; (2) Clerks, which include secretaries, librarians, cashiers, etc.; (3) Service Workers, which include travel, 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, and manufacturing jobs etc.; and (7) Elementary Occupation Workers such as street vendors, porters, 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.[3] The manufacturing and service sector then consist of six occupations as we attribute all "skilled agricultural workers" to the agricultural sectors. Home services are modeled as a separate sector.

4 Empirical Facts

4.1 Sectoral and Occupational Employment Transitions by Gender

We now document how occupational and sectoral employment choices differ between men and women and how their employment choices evolve along countries' development path.

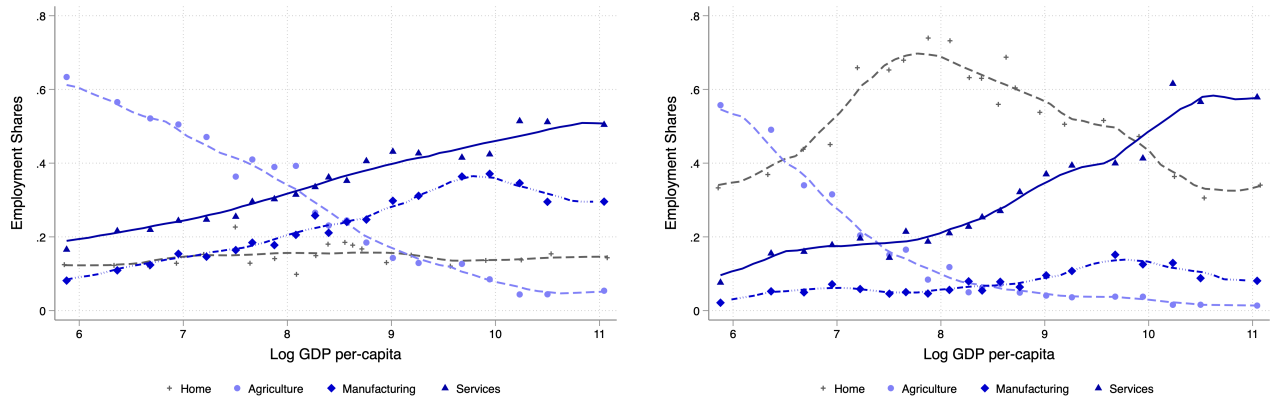
We first focus on men and women's employment choices across sectors, i.e., agriculture, manufacturing, market services, and home services. Figure 1(a) plots the employment shares of men in each sector against countries' log of GDP per capita, pooling all countries and all years of our sample. Men's employment transitions exhibit the standard patterns of structural change that have been extensively documented in the

[3] We attribute the small percentage of workers who report to work in different occupations in the agriculture sector to the other two sectors according to the share of workers in that occupation who work in either manufacturing or the service sector in the remaining data.

Figure 1: Sectoral and Occupational Employment Shares by Gender

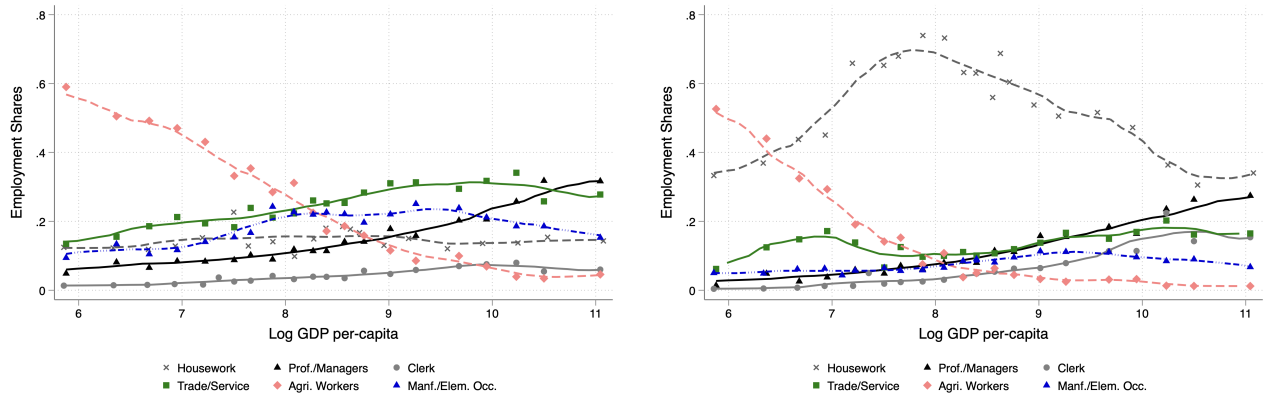
(a) Across Sectors: Male

(b) Across Sectors: Female



(c) Across Occupations: Male

(d) Across Occupations: Female



Notes: This figure is a non-parametric plot of the employment shares 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. Figures (a) and (c) report the employment shares for men, while (b) and (d) report them for women.

literature: the (male) employment share decreases in agriculture, is hump-shaped in manufacturing, and increases in market services, as countries grow richer. The share of men in the "home sector" (i.e., not in the labor force) is small and relatively constant across countries' economic development. Figure 1(b) shows a very different picture for women's sectoral employment shares. At low levels of economic development, women leave agriculture and mostly leave the labor force all together by sorting into the home sector. At higher levels of economic development, women then re-enter the labor force to work mostly in the service sector. Hence, female labor force participation (FLFP) follows a U-shaped pattern across countries' GDP per capita, which has been documented in the literature.

Next, we document gender-specific employment shares across occupations, which received less attention in the literature on structural transformation, but which have been emphasized in the literature as an important determinant of talent allocation and aggregate productivity (Hsieh et al., 2019). Figures 1(c) and 1(d) show that the employment shares of men and women decline in agricultural occupations with economic

development, in line with the decreasing agricultural sector which is heavily concentrated in agricultural occupations. Employment shares of men increase in all other occupations and the share of men in the home sector is roughly constant across development levels. Employment in professional occupations experiences the steepest increase from less than 8% in the poorest countries to over 35% in the richest countries. Male employment shares also increase in trade and service jobs and in machine-operating jobs, while increases are small for men in clerical jobs. Women's employment shares are again marked by the large increase and then decrease in "home service" employment. In the first stages of development, women who leave agricultural jobs sort mostly into home services, while the employment shares in other occupations see only very moderate changes. At higher development levels, female employment shares expand rapidly in clerical and in professional occupations, reaching in the richest countries close to 20% in clerical jobs and close to 30% in professional occupations. Contrary to men, women's employment shares in trade and service jobs or in machine operating jobs experience only a very modest increase and their levels remain substantially lower than the male counterparts.

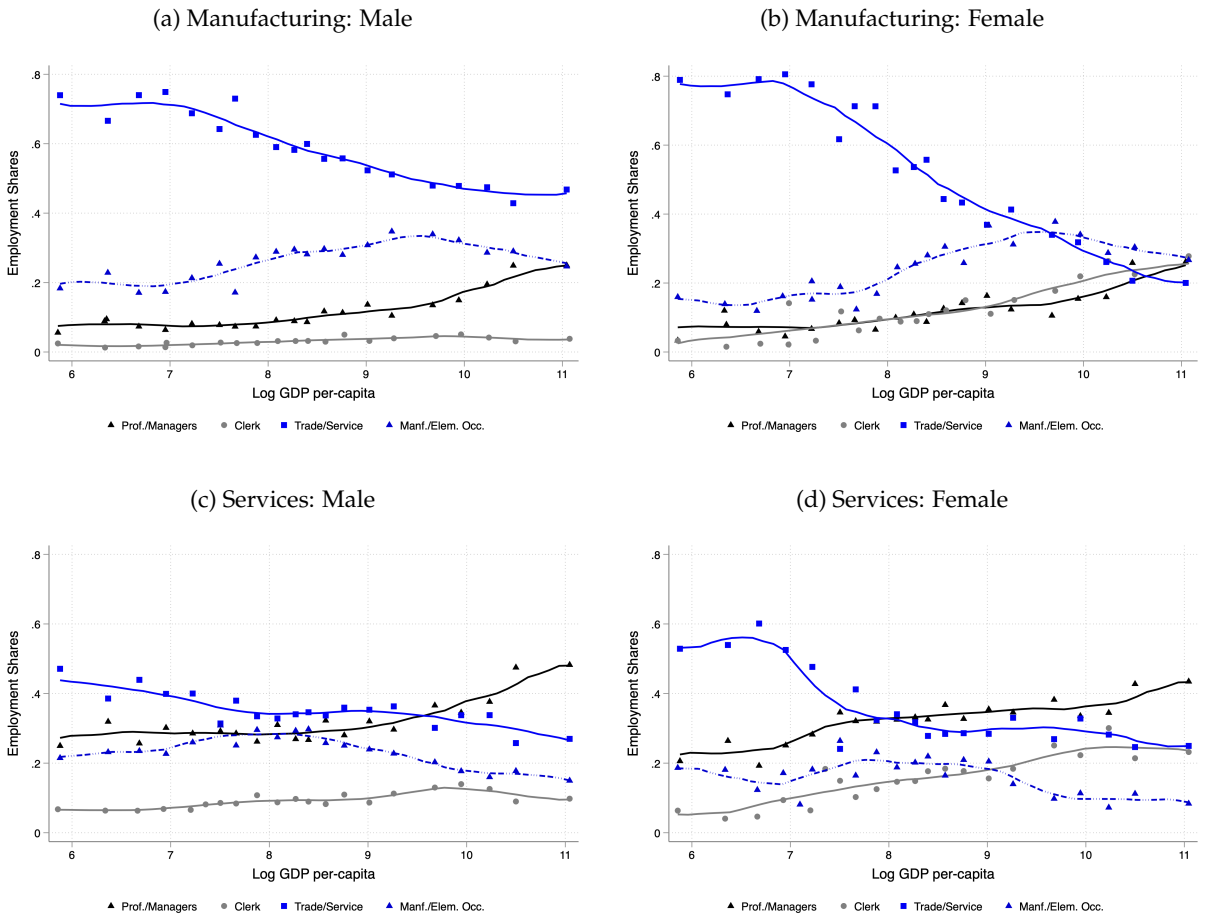
These gender differences in occupational choices and changes in women's employment choices over time can have important implications for structural transformation as sectors' production technologies differ in the intensity by which they rely on different occupations. Figure 2 plots occupational employment shares *within a given sector* separately for men and women against countries' development level. Conditional on working in the manufacturing sector, Figures 2(a) and 2(b) show that 75-80% of men and women work as craft and trade workers in poor countries which decreases in rich countries to less than 50% for men and to roughly 20% for women. Male employment shares in machine operating, elementary, and professional occupations increase within manufacturing, while the male employment share in clerical jobs remains constant. For women, the decrease in craft and trade jobs in manufacturing is even steeper and is compensated by larger increases in employment shares in clerical, professional, and elementary occupations. Overall, employment in the manufacturing sector is very concentrated in craft and trade occupations in poor countries, while manufacturing production in richer countries has more division of labor and spread workers across multiple occupations.

Figures 2(c) and 2(d) show that employment shares in the service sector are less concentrated in poor countries (compared to manufacturing). Approximately 50% of men and women work in craft, trade and service jobs, 20% in professional, 20% in elementary, and 10% in clerical jobs. At higher development levels, the share of craft, trade and service workers decreases to less than 30% while the share of machine operating and elementary jobs decreases somewhat for men and more for women. Employment in professional occupations increases for both genders while it increases only for women in clerical jobs. Even in the richest countries, men remain over-represented in professional and machine operating jobs, while women are over-represented in clerical jobs.

4.2 Gender Segregation Across and Within Sectors

As an alternative way of measuring gender differences in employment choices, we use the Theil Index of segregation to measure overall gender segregation across all occupation and sector pairs. The index measures the entropic "distance" between the current distribution and an "ideal" egalitarian state, where the most unequal distribution takes the value of 1 and the egalitarian state corresponds to an index value of

Figure 2: Occupation Structure within Manufacturing and Services by Gender



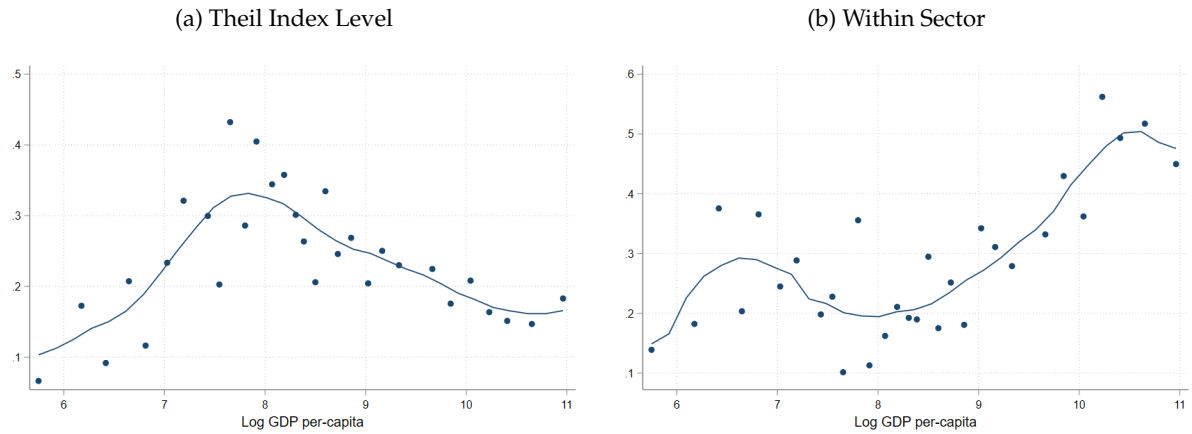
Notes: This figures is a non-parametric plot of the share of men and women in each occupation against the log of real GDP per-capita. Figures (a) and (b) show the occupational distribution among workers in the manufacturing sector. Figures (c) and (d) show the distribution for workers in the service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

0 (see Appendix B.1 for details). A key advantage of the Theil index is that the overall gender segregation across occupations and sectors can be decomposed into the share of segregation that is due to differences across sectors compared to the share that is due to differences across occupations within sectors.

Figure 3(a) shows that the aggregate Theil Index exhibits an inverted U-shape along countries' development levels. In poor countries, gender segregation across occupations and sectors is low, which then first increases and then decreases. The increase is driven by the decline of agricultural employment for both genders, which leads many women to exit the labor force while men sort into manufacturing and market services. The decrease in segregation at higher development levels then reflects the re-entry of women into the labor force and a convergence in occupational and sectoral choices across men and women. Figure 3(b) plots the share of total gender segregation that is explained by segregation across occupations within sectors against countries' GDP per capita. The within-sector component explains only 15-25% of segregation in low- and middle-income countries, but its importance increases to 60% in high-income countries.[4] This finding is

[4]In this analysis, changes in labor force participation are attributed to across-sector variation since we classify individuals who are

Figure 3: Gender Segregation across Sectors and Occupations



Notes: This figure plots the Theil Index that measures gender segregation across occupation-sector pairs against the log of 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.

in line with Bandiera, Elsayed, Heil and Smurra (2022) who document changes in the organization of labor across jobs along countries' economic development. Put together, the evidence presented above shows that gender differences in occupational and sectoral employment choices are large and that each dimension matters differentially in low- and high-income countries, emphasizing the importance of studying occupational and sectoral employment choices simultaneously.

4.3 Gender Gaps in Employment and Wages

To complement the cross-country evidence presented in the previous sections, we now document how gender differences in employment choices evolve *within* countries over time. To do so, we focus on our core sample of six countries which includes India, Indonesia, Brazil, Mexico, Canada and the United States. These countries cover different development levels and offer high-quality data for five or more decades, including individual-level data on hours worked and wages.

To document how gender differences in sectoral and occupational choices evolve over time, we compute "employment ratios" which divide the share of women in a sector (occupation) by the share of men in the same sector (occupation). To document the evolution of gender differences in wages, we then compute "wage ratios" which divide the average hourly wage of women by the average hourly wage of men in the same sector (occupation). 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 a given occupation or sector than women (men).

not in the labor force as working in home services (which we classify as a sector which has no occupational dimension). Re-computing the Theil index conditional on labor force participation shows similar results for the global Theil index and attributes even more weight to the within-sector variation.

Gender Employment Gaps Over Time

Table 1 documents these employment and wage ratios in the beginning and end of our sample (i.e., for the 1970ies and 2010s). Panel A shows income-weighted average ratios for each sector, averaging across occupations and across all six countries. Panel B shows similar income-weighted averages ratios for occupations, averaging across occupations and countries. Lastly, Panel C shows income-weighted average ratios for each of the six countries, averaging across sectors and occupations. Appendix Figure A.1.3 provides more detailed information by plotting the employment and wage ratios against each countries' GDP per capita for sectors and occupations, showing the full evolution of gaps over time for each country-sector pair or each country-occupation pair.

From Column (1) of Panel A in Table 1 shows that women in the 1970s were on average almost 11 times more likely to stay at home than men. Employment ratios were smallest in agriculture and manufacturing, which employed around 3 women for every 10 men, with a somewhat better ratio in the service sector (6.2 women for every 10 men). From Panel B, women were also under-represented in professional, craft, trade and service, and machine operating and elementary occupations, which employed around 4 women for every 10 men, while clerical jobs had on average an over-representation of female workers compared to male workers, employing 11.6 women for every 10 men.

In the 2010s, male and female employment choices converged to some extent, which was accelerated by a large entry of women into the labor force: Women in the 2010s were on average 6 times more likely to stay at home than men, compared to 11 times in the 1970s. Female-to-male employment ratios increased over time very unevenly across sectors and occupations: the average ratios in manufacturing and agriculture remained almost constant while the employment ratio increased in the service sector from 0.62 to 0.85. For occupations, the ratios increased very little in agricultural, machine operating, and elementary jobs, slightly more in crafts, trade, and service jobs, and most noticeably in clerical and professional jobs. In professional occupations, the average female-to-male employment ratio doubled from 0.4 to 0.8. In clerical jobs, women's initial over-representation relative to men further increased from 1.16 to 1.36.

Appendix Figure A.1.2(a) shows the entire distribution of gender employment ratios by pooling across all occupations, sectors, and countries. As we can see, women were under-represented in over 90% of jobs in the 1970s. Despite a clear rightward shift in the entire distribution, women continue to be under-represented in around 80% of occupation-sector-countries. Appendix Figure A.1.3 shows the heterogeneity across countries by plotting the employment ratios against each countries' GDP per capita for sectors and occupations. This allows us to further examine the change in gender employment gaps with the process of economic development. For example, from Figure A.1.3(a), we see that India is a stark outlier — it starts with the lowest levels of female LFP and shows no noticeable improvement in it over time. Indonesia exhibits a large increase in female relative to male employment in the service sector, a small increase in agriculture but a decrease in manufacturing. Brazil and Mexico see large increases in all sectors with the steepest slope in services. Canada and US see little or no increases in manufacturing and agriculture but sharp increases in the service sector. Overall, consistent with the previous patterns, the employment gaps in services is lowest in richer compared to poorer countries. In fact, for the U.S. and Canada, the ratio exceeds 1 in recent years, indicating that women are over-represented in the service sector compared to men.

Table 1: Gender Employment and Wage Ratios

	Employment Ratio		Wage Ratio	
	1970-75	2010-18	1970-75	2010-18
	(1)	(2)	(3)	(4)
Panel A. Sectors				
Home	10.92	6.01		
Agriculture	0.29	0.32	0.72	0.80
Manufacturing	0.27	0.26	0.56	0.79
Services	0.62	0.85	0.56	0.80
Panel B. Occupations				
Professional	0.40	0.81	0.64	0.77
Clerk	1.16	1.36	0.84	0.95
Craft, Trade, Service	0.41	0.55	0.5	0.69
Agricultural	0.28	0.31	0.74	0.81
Machine Op., Elementary	0.37	0.41	0.52	0.74
Panel C. Countries				
India			0.55	0.67
Indonesia			0.65	0.76
Brazil			0.69	0.79
Mexico			0.79	0.90
Canada			0.63	0.75
USA			0.63	0.77

Notes: Columns (1)-(2) and (3)-(4) report the employment and wage ratios in the first and last decade that we observe countries. 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. Panel A reports an income-weighted average of employment and wage ratios across occupations and countries, while Panel B reports an income-weighted average across sectors and countries, and Panel C reports an income-weighted average across sectors and occupations. Columns (1) and (2) are not reported in Panel C since the fraction of men and women always add up to 1 within a country.

Large gender differences also remain across occupations — both in levels and trends (Appendix Figure A.1.3(b)). There are two key patterns that are worth highlighting: first, despite an improvement over time, women (as compared to men) are significantly less likely to be employed in professional and managerial jobs in poor countries (India and Indonesia). This ratio does improve significantly with economic development, completely closing in rich countries (U.S. and Canada). Second, except for India and Indonesia, women are over-represented in clerical jobs, and increasingly so with economic development (U.S. and Canada vs Brazil and Mexico), as well as over time. Put together, our patterns highlight the role of examining gender differences not only across sectors, but also across occupations.

Gender Wage Gaps Over Time

Columns (3) and (4) in Panel A of Table 1 show the female-to-male wage ratios for the beginning and end of our sample. Average wage ratios are particularly low in manufacturing and services in the 1970s where

women earned on average only 56 cents for each dollar earned by men. In the 2010s, the average wage ratios improve close to 0.8 in all sectors. More heterogeneity remains across occupations (Panel B), where women's earnings in the 1970s were particularly low in craft and trade jobs (0.50), machine operating jobs (0.60) and professional/managerial jobs (0.64). These ratios improved in the 2010s respectively to 0.70, 0.74, and 0.77. More generally, there have been a significant closing of gender gaps across all occupations (except skilled agricultural jobs). Panel C further shows the wage ratio for each country by taking the income-weighted average across all occupation and sector pairs. Lastly, from Panel C, we see that the patterns across sectors and occupations are not a country-specific story. In the 1970s, women earned around 60-70 cents for every $1 earned by a man, which improved to around 70-80 cents by 2010s. This highlights two patterns of interest: first, even today, and even in the most developed countries, significant gender gaps persist. Comparing across all columns of Table 1 shows that employment and wage ratios are not necessarily correlated: For example, average employment ratios are substantially larger in the service sector than in other sectors, but wage ratios are similar or even smaller in the service sector. Moreover, unlike the distinct gender patterns we document in employment transitions and economic development (Figure 1), there appears to be little correlation between gender wage gaps between rich and poor countries.

Lastly, pooling across occupations, sectors, and countries, Appendix Figure A.1.2(b) shows the change in distribution of gender wage ratios over time. There are two points to highlight: first, women earned less than men in over 95% of all sector-occupation-country jobs in the 1970s compared to around 85% in the 2010s. Appendix Figures A.1.3(c) and A.1.3(d)) show the heterogeneity in wage ratios across countries in more detail. The key insight from these graphs suggest an improvement in wage gaps in sectors and occupations across all countries. The only notable exception appears to be a stagnation in clerical and professional jobs in India.

4.4 Key Takeaways and Model Implications

Our empirical patterns provide three key takeaways.

First, employment transitions across sectors and occupations along countries' development process have salient gender patterns. For sectors, men follow the standard patterns of structural transformation. Women, on the other hand, follow a different pattern — at low levels of economic development, women exit agriculture and first sort into the home sector. At higher development levels, women then enter the labor force again and mostly sort into the (market) service sector. Female employment in manufacturing remains small and relatively constant across development levels. Across occupations, men are more likely to work in trade and service, machine operating, and professional jobs. Women are over-represented, most notably, in clerical occupations. Representation of women in professional and managerial jobs have approached gender parity in recent years middle-income countries (Brazil and Mexico), while being are over-represented in recent years in U.S. and Canada.

Second, we find that gender segregation across occupation-sectors follows an inverted U-shape pattern across countries' development process; and that the share of this segregation that is explained by gender differences across occupations within a sector matters more with economic development. This finding emphasizes that the occupational and sectoral dimension are both important when studying the effects of

gender roles on employment choices. Given that the relative importance of these dimensions changes over countries' development stages, it is important that long-run studies which cover a wide cross-section of countries incorporate both dimensions into the analysis.

Third, 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. Moreover, a good gender balance in employment can co-exist with large gender pay gaps, which is most salient in the service sector, suggesting that non-pecuniary factors are important in explaining men and women's employment choices.

Do these empirical patterns reflect a change in underlying barriers faced by women (such as social norms, discrimination in hiring, etc.) in productively contributing to the economy? Or do they reflect economic channels, such as changes in technology, (gender-specific) human capital accumulation, or the returns to it, etc.? Disentangling these channels are important from the perspective of understanding what drives the inclusion of women in the economy. In the light, we develop a theoretical framework that incorporates these forces and then allows us to transparently decompose the observed changes (in employment and wage gaps, for example), through these two channels: economic and non-economic. Furthermore, we can flexibly calibrate the model and use it to simulate counterfactuals that quantify the importance of gender-specific mechanisms on countries' paths of structural transformation and economic development.

5 Model

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

5.1 Overview of the Model

We develop a gender-specific model of occupational and sectoral choice model (Roy, 1951) that incorporates gender-specific barriers and a rich set of economic forces that have been documented to be important drivers of structural transformation in the literature. They include: (i) technological change biased towards specific sectors or occupations; (ii) changes in gender-specific effective human capital supply due to schooling changes across the world; (iii) skill-biased technological change; (iv) changes in the sorting of workers based on their comparative advantage reflected by gender-specific occupation-sectoral returns to human capital; (v) inclusion of non-homothetic preferences, which allow consumers to change their consumption basket across sectors as they become richer.[5]

In addition to the above economic forces, we also incorporate two gender-specific (non-economic) barriers that can also directly affect women's employment choices and wages relative to men (Hsieh, Hurst, Jones

[5]Income effects have been shown to lead consumers to increase their expenditure share in services (a luxury good) and to decrease it in agriculture (a necessity), as they grow richer (Alder, Boppart and Müller, 2022; Herrendorf, Rogerson and Valentinyi, 2013; Comin, Lashkari and Mestieri, 2021).

and Klenow, 2019). First, we model a gender-specific dis-utility of participating in the labor force and working a specific occupation-sector. We call these "gender norms" that reflect underlying amenities, barriers, entry costs, social norms, or any other factor that can make it differentially attractive for women and men to work in a given occupation-sector (relative to staying at home). Second, we incorporate "wage discrimination" faced by women, which captures differences in the wage rate per human capital unit between men and women.

Put together (as we will show later), our model therefore allows us to parsimoniously decompose the observed gender gaps in employment and wages into those shaped by economic changes, as opposed to changes in gender barriers. To illustrate this with a specific example, we have documented a significant closing of gender gaps in services in our earlier discussion. They could be driven by technological growth in services, or a change in the mix of occupations used in the services sector, or income effects from countries getting richer over time, etc. — all economic forces. On the other hand, they could also reflect a reduction in gender norms, amenities, and other gender barriers. Our model allows a parsimonious decomposition between these two.

5.2 Setup and Preferences

Model Setup: The economy consists of O occupations and J sectors, namely agriculture A, manufacturing M, market services S^m, and home services S^h. There is a mass N_g of individuals of each gender g, which is either male m or female f. 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. Workers incur a utility cost A_{ojg} and receive a wage rate per human capital unit w_{ojg} when working in occupation o and sector j. Wage rates and utility costs can differ across men and women. Each sector produces the final good competitively.

Preferences: Individuals have non-homothetic preferences over agriculture, manufacturing and services $j = \{A, M, S\}$, where services are a composite of home and market services (Comin et al., 2021; Conte, 2022; Cravino and Sotelo, 2019; Lewis et al., 2022; Sposi et al., 2021). Non-homothetic preferences allow consumers to change their expenditure shares across sectors as they become richer, which has been documented in the literature as an important driver of structural transformation and determinant for the allocation of workers across sectors.[6] In particular, we use non-homothetic CES preferences — developed in Comin, Lashkari and Mestieri (2021) — which define preferences for a consumption bundle C over the three sectors j (agriculture, manufacturing, and services) implicitly through the following constraint:

$$\sum_j \theta_j^{\frac{1}{\sigma}} \left(\frac{C_j}{C^{\varepsilon_j}} \right)^{\frac{\sigma-1}{\sigma}} = 1, \tag{1}$$

where $\sigma > 0$ is the standard CES elasticity of substitution across sectors, ε_j is the non-homothetic elasticity of substitution, and θ_j are sectoral preference-shifters. As we will see below, $\varepsilon_j = 1$ would reduce these preferences to the standard homothetic CES preferences. Individuals maximize their utility subject to the

[6] See Herrendorf, Rogerson and Valentinyi (2013); Alder, Boppart and Müller (2022); Comin, Lashkari and Mestieri (2021); Fan, Peters and Zilibotti (2021); Lewis, Monarch, Sposi and Zhang (2022); Sposi, Yi and Zhang (2021).

standard budget constraint $\sum_j p_j C_j = I_{ojg}(z)$, where $I_{ojg}(z)$ is the income earned by an individual of gender g and ability z who works in the occupation-sector pair oj.

Home and Market Services: We assume a preference structure similar to Ngai and Petrongolo (2017) and Ngai, Olivetti and Petrongolo (2024), where services are a CES composite of home and market services $S = \{S^h, S^m\}$ given by:

$$C_S = \left[\sum_{s' \in \{S^h, S^m\}} \alpha_{s'} C_{s'}^{\frac{\sigma_s - 1}{\sigma_s}} \right]^{\frac{\sigma_s}{\sigma_s - 1}},$$

where $\sigma_s > 0$ is the elasticity of substitution between home and market services, and α_s are the preference weights across home and market services, with $\sum_{s'} \alpha_{s'} = 1$.

Consumption Choice across Sectoral Goods: An individual of gender g and ability z who works in occupation o and sector j spends the following shares of their expenditure on agriculture, manufacturing, and the service composite:

$$\varphi_j(I_{ojg}(z), p) = \theta_j \left(\frac{p_j}{P} \right)^{1-\sigma} \left(\frac{I}{P} \right)^{(\varepsilon_j - 1)(1-\sigma)}, \tag{2}$$

where the price index P is implicitly defined by the following constraint:

$$P = \left[\sum_j \theta_j p_j^{1-\sigma} \times \left(\frac{I}{P} \right)^{(\varepsilon_j - 1)(1-\sigma)} \right]^{\frac{1}{1-\sigma}}. \tag{3}$$

Appendix B.2 provides the proof. Note from Equation (2) that we can express the ratio of individuals' expenditure share in sector j relative to their expenditure share in manufacturing in the following way:

$$\frac{\varphi_j}{\varphi_m} = \frac{\theta_j}{\theta_m} \left(\frac{p_j}{p_m} \right)^{1-\sigma} \left(\frac{I}{P} \right)^{(\epsilon_j - \epsilon_m)(1-\sigma)} \tag{4}$$

This equation highlights how sectoral preference parameters, income levels, and substitution effects through prices impact individuals' sectoral expenditure shares. Specifically, like in homothetic preferences, the ratio of sectoral expenditure shares depend on preference-shifters θ_j/θ_m and on the substitution effect that arises from relative sectoral prices (p_j/p_m). In addition, the non-homothetic preferences now have an income effect, whose magnitude is governed by $\varepsilon_j - \varepsilon_m$.

Lastly, the respective consumption shares across home and market services within the service composite are given by:

$$\varphi_{s'} = \alpha_{s'} \left(\frac{p_{s'}}{P_S} \right)^{1-\sigma_s} \varphi_S \tag{5}$$

where φ_S is the expenditure share on the entire service composite and $P_S = \left[\sum_{s'} \alpha_{s'} p_{s'}^{1-\sigma_s} \right]^{\frac{1}{1-\sigma_s}}$ is the CES price index of the service composite. The proof is provided in Appendix B.3.

5.3 Occupational Choice

The utility of a worker i of gender g and ability z who works in an occupation-sector pair oj is given by:

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

where $V(I_{ojg}(z), p)$ is workers' indirect utility from consumption, i.e., $C_{ojg}(z) = I_{ojg}(z)/P(I_{ojg}(z), p)$. Note that individuals' income and the corresponding price index vary across gender-ability types and across the chosen occupation-sector pairs, so that we index these variables by $ojg(z)$. 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. Employment choices of men and women identify this utility cost only in relative terms across all options, so that we normalize the cost of working in home services to 0 for men and women in every country-years. In the rest of the paper, we refer to "gender norms" as the excess utility costs that women face, relative to men, when working in an occupation-sector oj relative to working in home services, i.e., $\triangle A_{oj} = A_{ojf} - A_{ojm}$.

ν^i_{oj} are idiosyncratic preference shocks for working in each occupation-sector pair, which we assume are Extreme Value Gumbel distributed across all occupation-sector pairs with a dispersion parameter σ_ν. Given this assumption, the share of workers of gender g and ability type z who chooses to work in occupation-sector pair oj is equal to:

$$\Pr(oj|g,z) = \frac{\exp\left[\frac{1}{\sigma_\varepsilon}V(I_{ojg}(z),p) - \frac{1}{\sigma_\varepsilon}A_{ojg}\right]}{\sum_{j'}\sum_{o'}\exp\left[\frac{1}{\sigma_\varepsilon}V(I_{o'j'g}(z),p) - \frac{1}{\sigma_\varepsilon}A_{o'j'g}\right]}. \tag{6}$$

This equation shows that employment choices of workers across occupation-sectors depend on the real income they earn in an occupation-sector and the gender-specific utility cost A_{ojg} that they incur when working in the occupation-sector.

We allow occupation-sector pairs to differ in their returns to ability κ_{ojg}. The effective human capital units of a worker of gender g and ability type z who works in an occupation-sector pair oj is then given by: $z^{\kappa_{ojg}}$. For each unit of human capital, an occupation-sector pays a wage rate w_{ojg}, which can differ across men and women. Hence, an individual of gender g and ability z who works in an occupation-sector oj earns income equal to: $I_{ojg}(z) = w_{ojg}z^{\kappa_{ojg}}$. We model female wage discrimination as a "wedge" τ_{oj} between women's wage rate and their marginal product so that men are paid by their marginal product $w_{ojm} = w_{oj}$, while women are paid a fraction $(1-\tau_{oj})$ of their marginal product, i.e., $w_{ojf} = (1-\tau_{oj})w_{oj}$.[7] In the estimation, we allow the gender-specific utility costs of working A_{ojg} and the wage discrimination τ_{oj} to vary across all occupation-sector pairs and across all country-years.

[7] This wedge can be micro-founded for example, as in Hsieh, Hurst, Jones and Klenow (2019), by assuming that entrepreneurs have a disutility δ_{oj} of hiring women. In equilibrium, this disutility is then exactly compensated by the profits that entrepreneurs earn from paying women below their marginal product.

5.4 Production and Equilibrium

Aggregate Supply of Human Capital: To solve for the aggregate supply of human capital in each occupation-sector oj, we assume that ability is log-normally distributed with mean μ_z and standard deviation σ_z. We then integrate over the human capital units of each worker conditional on having chosen a given occupation-sector oj as follows:

$$H_{oj} = \sum_g N_g \int_z \Pr(oj|g,z) z^{\kappa_{oj}} dF(z), \tag{7}$$

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:

$$Y_j = T_j \prod_o H_{oj}^{\gamma_{oj}}, \tag{8}$$

where γ_{oj} are sector-specific expenditure shares on workers from each occupation which sum to 1 within each sector: $\sum_o \gamma_{oj} = 1$. T_j is a sector-specific productivity term and H_{oj} is the total human capital that is supplied to occupation-sector oj (see, Equation 7). For a vector of sectoral prices (p_j), firms' profits (π_j) are given by $p_j Y_j - \sum_o w_{oj} H_{oj}$.

Equilibrium: The exogenous parameters of the model characterize preferences $\{\sigma, \sigma_s, \alpha_s, \{\varepsilon_j, \theta_j\}_{\forall j}\}$, the dispersion of preference shocks across occupation-sector pairs (σ_ν), the mean and variance of the ability distribution $z \sim F(z)$, the production function $\{\{T_j\}_{\forall j}, \{\gamma_{oj}, \kappa_{ojg}\}_{\forall oj}\}$, and gender barriers $\{\{\tau_{oj}\}_{\forall oj}, \{A_{ojg}\}_{\forall ojg}\}$. Given these parameters, the equilibrium in each country-year is defined by a vector of sectoral prices and occupation-and-sector-specific wage rates $\{\{p_j\}_{\forall j}, \{w_{oj}\}_{\forall oj}\}$ so 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 supply and demand for human capital.
4. Good markets clear in each sector.

6 Model Calibration

6.1 Overview

Our calibration exercise attempts to flexibly allow for gender barriers, production technologies, education, and returns to human capital to differ across occupations, sectors, countries, and over time. Gender barriers distort workers' occupational and sectoral sorting away from their comparative advantage, which

can generate a misallocation of talent and can lower aggregate productivity and growth. We flexibly estimate the model to fit several key data moments separately for each country-year. Gender barriers for each occupation-sector are then residual wedges that match the observed gaps in employment and wages, *after* accounting for the economic structure of the model. As we will show later, even though they are calibrated as residuals, our estimates for gender barriers correlate strongly with indicators of norms in the households, and equality at the workplace in these countries. Lastly, in the spirit of an accounting exercise, we then evaluate counterfactuals that quantify the contribution of changes in gender barriers to each country's trajectories in employment choices (structural transformation), sectoral real value added, and welfare over the last four decades.

Turning to the calibration exercise, we first discuss parameters that we take from the literature, or from data moments outside of our model. The remaining parameters are then calibrated by fitting the model's equilibrium conditions to the data in an iterative algorithm which is described in Appendix C.

6.2 Parameters Calibrated Outside of the Model

Preferences: We first calibrate the preference parameters which 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 $\sigma_s = 2.3$. For the preference shocks across occupation-sector pairs, we follow Hsieh et al. (2019) and set the dispersion parameter to $\sigma_\nu = 2$. For the elasticities in the non-homothetic CES preferences, we follow Comin et al. (2021) and set $\sigma = 0.53$, and $\varepsilon_j = \{0.4, 1, 1.2\}$.[8] The sectoral preference shifters θ_j are estimated by indirect inference across all country-years.

Ability Distribution and Returns to Human Capital: Ability distributions vary across countries and over time and follow a log-normal distribution $z \sim \ln N(\mu_z, \sigma_z)$. We calibrate the mean and variance of this distribution $\{\mu_z, \sigma_z\}$ for each country-gender-year from the moments of the observed distribution of years of schooling.[9] We estimate the returns to human capital for each occupation-sector and gender κ_{ojg} by estimating a Mincer regression, similar to Fan, Peters and Zilibotti (2021). Recall that an individual of gender g and ability z who works in an occupation-sector pair oj earns income equal to: $I_{ojg}(z) = w_{ojg} \times z^{\kappa_{ojg}}$. To estimate κ_{ojg}, we take logs of this structural equation and use our individual-level data to estimate the following Mincerian wage regression separately for each country-year:

$$\ln(\text{wage}_i) = \alpha_{ojg} + \kappa_{ojg} \ln \text{YrsSchool}_i + \gamma_1 Exp_i + \gamma_2 Exp_i^2 + u_i \tag{9}$$

where wage_i are individuals' hourly wages, YrsSchool_i are years of schooling, Exp_i is experience, and α_{ojg} are occupation-sector-gender fixed effects that capture, among other things, the average wages in a given

[8] As in Comin et al. (2021), we can only identify ϵ_j relative to a base sector, for which we choose manufacturing and which we normalize to 1.

[9] In the current model, we assume that the distribution of ability/schooling is exogenously given, which implies that the distribution is constant in our counterfactual simulations. This assumption can be relaxed in the future by modeling education choices endogenously as a function of education costs and returns. Holding the ability distribution constant provides a lower-bound of our counterfactual results as the observed decrease in gender norms over time increased women's return to education and therefore increased their human capital accumulation.

occupation-sector-gender pair. The coefficient $\hat{\kappa}_{ojg}$ captures the returns to schooling for each occupation-sector-gender pair.[10]

6.3 Parameters Calibrated Using the Model

We estimate the remaining parameters by fitting our model's equilibrium conditions to key data moments, which include sectoral productivity and preference shifters $\mathcal{U} = \{\{T_j\}_{\forall j}, \{\{\gamma_{oj}\}_{\forall oj}, \{\theta_j\}_{\forall j}\}$, and gender barriers $\mathcal{B} = \left\{\{\tau_{oj}\}_{\forall oj}, \{A_{ojg}\}_{\forall ojg}\right\}$, which consist of wage discrimination and gender norms. Sectoral preference shifters θ_j do not vary across country-years, but we allow the other parameters listed here to differ flexibly across country-years. In the following sections, we describe the intuition of the estimation strategy, while we explain the numerical procedure and iterative algorithm in Appendix C.

Gender Norms: Individuals choose an occupation-sector pair oj based on the real income that they earn in the occupation-sector and based on the gender-specific utility cost that they incur when working in the occupation-sector pairs. From Equation (6), we can see that they share of men and women who chooses to work in a given occupation-sector pair depends on:

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

where $V(I,p)$ is workers' indirect utility that depends on their real income and sectoral prices, and where A_{ojg} represents the gender-specific utility cost of working in oj. From observed employment choices of men and women, we can only identify their utility costs A_{ojg} relative to all other employment options. We therefore normalize $A_{home,g} = 0$ for men and women, so that A_{ojg} captures the additional disutility that workers incur when working in occupation-sector oj relative to working in home services. To calibrate this utility cost A_{ojg}, we use our model to compute worker's indirect utility $V(I,p)$ when working in each occupation-sector, which is an equilibrium outcomes as it depends on wages and prices that have to ensure that labor and goods market clear. For each country year, we then use Equation (6), the model-implied indirect utility measure $V(I,p)$, and data on gender-specific employment shares $\Pr(oj|g)$ to infer the utility costs A_{ojg} in each occupation-sector which ensure that our model exactly matches the observed employment shares of men and women.

Wage Discrimination: To infer female wage discrimination, we assume that men are paid their marginal product w_{oj}. We can then express women's wage per human capital unit as: $w_{ojf} = (1-\tau_{oj}) \times w_{oj}$, where τ_{oj} represents wage discrimination, which we allow to be positive or negative. Given this assumption, we can express the female-to-male wage ratio as a function of female wage discrimination and the (model-implied)

[10] We experimented with a model version in which we restricted κ_{oj} to be the same for both genders. As reported in Figure B3, there is little variation across gender in these parameters and our qualitative and quantitative results remain mostly unchanged in both versions. We set $\kappa_{home} = 0$.

female-to-male human capital ratio, so that:

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

where $\overline{H}_{ojg}$ is the average human capital of workers with gender g who choose to work in the occupation-sector pair oj. This average human capital by gender is endogenously determined in the model, as it depends on workers' ability distributions, occupation-sector-gender-specific returns to human capital, and workers' occupational and sectoral choices in equilibrium which are a function of endogenously determined wages and prices. To calibrate tau_{oj}, we therefore solve for the model-implied female-to-male human capital ratio and then infer wage discrimination to match the observed female-to-male wage ratio perfectly. When calibrating the gender barriers $\mathcal{B}$, we do not impose any direction or any restrictions on the values across countries or over time, which allows for the possibility that women may have an advantage or disadvantage of working in a particular occupation-sector.

Sectoral Productivity and Preferences, Prices and Wages: To infer sectoral productivity T_j for each country-year, we target the ratio of value added across sectors. To pin down the levels of value added over time, we target countries' real GDP per capita growth rates (in international USD). Since labor is the only production factor in our model, we use observed wage bill in each sector to measure value added. We then solve for sectoral prices p_j and for occupation-sector-specific wage rates w_{oj} to ensure that goods markets clear in each sector and to ensure that labor markets clear in each occupation-sector pair. Preference shifters θ_j are identified by indirect inference from Equation (2) as follows from Equation (1). We normalize θ_m in manufacturing to 1 without loss of generality. We then use the average ratio of aggregate expenditure shares, relative sectoral prices, and aggregate real income across countries (similar to Equation 4) to estimate θ_j. We find values of 0.6 for agriculture and 1.6 for services.

7 Results on Gender Barriers and Model Validation

We now describe our estimates of wage discrimination (τ_{oj}) and gender norms ($\triangle A_{oj} \equiv A_{ojf} - A_{ojm}$) for each occupation-sector-country-year. As a validation exercise, we further show that our model-implied estimates correlate with non-targeted data moments that measure gender norms and women's legal rights across countries and over time.

7.1 Gender Barriers

Table 2 summarizes our estimates of gender barriers. Columns (1)-(2) report gender norms and columns (3)-(4) report wage discrimination, respectively in the first and last year of our sample. The panels show how gender barriers differ and evolved across sectors, occupations, and countries. Panel A shows gender barriers for each sector, averaging across occupations and countries. Panels B documents average gender

barriers for each occupation. Panel C averages gender barriers separately for each country. Appendix Figure A.1.4 shows the changes in these barriers separately for each country and over time.

Table 2: Gender Norms and Female Wage Discrimination

	Gender Norms (ΔA)		Wage Discrimination (τ)	
	1970-75	2010-18	1970-75	2010-18
	(1)	(2)	(3)	(4)
Panel A. Sectors				
Agriculture	7.71	5.52	0.42	0.33
Manufacturing	7.46	5.71	0.47	0.33
Services	5.74	3.49	0.48	0.32
Panel B. Occupations				
Professional	6.46	3.34	0.45	0.3
Clerk	5.59	2.62	0.3	0.29
Craft, Trade, Service	6.55	4.38	0.50	0.37
Agricultural	7.87	5.67	0.41	0.33
Machine Op., Elementary	7.45	5.11	0.4	0.31
Panel C. Countries				
India	7.40	7.62	0.44	0.42
Indonesia	5.08	6.20	0.45	0.3
Brazil	9.59	3.57	0.40	0.32
Mexico	8.41	4.54	0.39	0.24
Canada	4.95	1.33	0.55	0.36
USA	4.68	0.92	0.54	0.36

Notes: Columns (1)-(2) and (3)-(4) report gender norms ($\triangle A$) and wage discrimination (τ), respectively in the first and last time period that we observe in our sample. Panel A reports the income-weighted averages of gender norms and wage discrimination for each sector, while Panel B reports them for each occupation. Panel C reports average gender norms and wage discrimination separately for each country of our sample, averaging across sectors and occupations.

Gender Norms ($\triangle A$): columns (1) and (2) of Table 2 show that gender norms decreased in all sectors and occupations over time. The largest decrease has been in the service sector and in professional and clerical occupations. Looking across countries in Panel C, we see that gender norms have changed little (or in fact worsened) in India and Indonesia, while all other countries experienced substantial and comparable declines during the last five decades. Appendix Figure A.1.4 shows the heterogeneity in levels and changes of gender barriers separately for each country. The Figure shows that gender norms declines most for the service sector and for clerical and professional occupations in Mexico, Brazil, US, and Canada with little improvements in India and Indonesia.

Wage Discrimination (τ): Columns (3) and (4) of Table 2 show that wage discrimination across sectors has uniformly decreased by around 10-15 percentage points (pp) between the 1970ies and the 2010s (Panel A). The averages mask a substantial heterogeneity in wage discrimination across occupations as shown in Panel

B. For example, in the 1970s, wage discrimination was the highest in craft, trade, and service occupations (50%), followed by professional occupations (45%). In the 2010s, these occupations have seen the large declines in wage discrimination (around 15 pp). On the other hand, wage discrimination was the lowest (though non-trivial in magnitude with 30%) in clerical jobs, with saw no change over time. Looking across countries shows that wage discrimination declined substantially in all countries except in India, where the decline is very low (0.02 pp). Indonesia sees a substantial decline in wage discrimination (15 pp) despite the small deterioration in gender norms. In contrast to gender norms, there appears to be no systematic correlation between the levels of wage discrimination and countries' economic development, as we observe relatively large levels of wage discrimination in India, but also in the US and Canada. The decline over time is smallest in the low-income countries of our sample (India and Indonesia), followed by Brazil (8 pp) and Mexico (15 pp) and the decline is steepest in the high-income countries of Canada and the US (around 20 pp).

7.2 Model Fit and Validation

In the model, we estimate gender norms and wage discrimination in the spirit of an accounting exercise by inferring them as "structural residuals" which ensure that our model's equilibrium conditions perfectly fit the observed data on gender employment and wage ratios. In this section, we show that these estimates relate closely to targeted and non-targeted data moments.

Model Fit with Targeted Moments: Appendix Figure B1 pools the data across all occupations, sectors, countries, and years and shows a strong negative correlation between the gender norms ($\triangle A_{ojct}$) and the observed gender employment gaps in the data. Figure B2 documents a strong correlation between wage discrimination (τ_{ojct}) and gender wage gaps. As we target the employment choices and gender wage ratios perfectly, any difference in the correlation between these targets and the estimates is explained by the remaining moments in the estimating equations, i.e., by real income earned in each occupation-sector in the case of gender norms and the gender human capital gap in the case of wage discrimination.

Validating the Model Estimates with Non-Targeted Measures of Gender Norms and Women's Rights: We now examine the extent to which our estimates of gender norms and wage discrimination relate to measurable changes in gender norms and women's labor market constraints, which we do not target in our estimation. To do so, we take advantage of the World Bank's "World, Business, and the Law" (WBL) database (World Bank, 2019, 2020; Hyland, Simeon and Goldberg, 2020) which 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 normative questions such as "Can a woman get a job in the same way as a man?" or "Is there a legislation on sexual harassment at the workplace?" as well legal indicators such as "Can a woman sign a contract" or "Does the legal system explicitly prohibit gender discrimination in the workplace". To evaluate the correlation between our estimates and these measures, we regress our estimated values of gender norms and wage discrimination on the WBL indicators across the countries and decades of our sample as follows:

$$y_{ojct} = \alpha + \beta \text{WBL}_{ct} + \gamma \ln \text{GDP p.c.}_{ct} + u_{ct}, \tag{11}$$

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 report robust standard errors in parentheses. Panel A of Table B1 shows a strong negative relationship between our estimated gender norms ($\triangle A_{ojct}$) and 4 WBL indicators which measure gender equality equality in the household and in society (e.g., gender parity in heading a household, laws against domestic violence, and rights to divorce and re-marriage). Panel B of Table B1 shows a negative correlation between our estimated wage discrimination and measures of gender equality at the workplace (e.g., whether women can get a job in the same way as a man, and laws against gender discrimination and sexual harassment at the workplace). Our estimates of both gender barriers also correlate negatively and significantly with the total WBL index across countries and years which captures overall gender equality in countries' legal system. Put together, the results show that our estimated gender barriers 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.

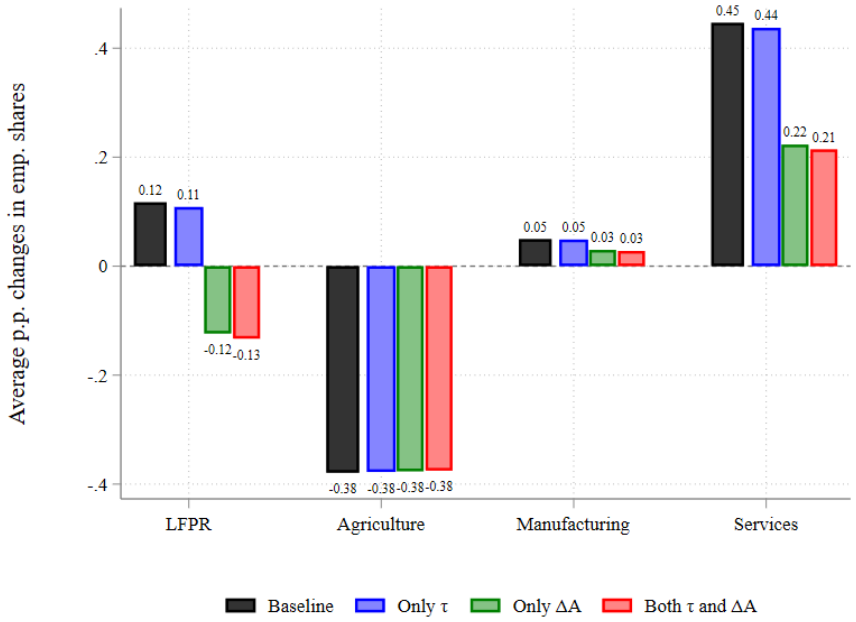
8 Importance of Gender Barriers in Explaining Economic Growth

We now use our estimated model to answer the following question: for each country, how much of the structural transformation and sectoral output growth that we observed over the last five decades can be explained by changes in gender barriers? To assess the importance of changes in gender barriers for these outcomes, we simulate counterfactual scenarios that hold these gender barriers constant at the values that we calibrated in the first year of each country, while we allow all other parameters to evolve over time according to the values that we calibrate for each year from the data. More specifically, we allow sector- and occupation-specific technological change (γ_{oj}, T_j), gender-specific education outcomes (μ_{zg}, σ_{zg}), and returns to human capital (κ_{ojg}) to evolve over time according to our estimated values, while keeping only gender barriers ($\tau_{oj}, \triangle A_{ojg}$) fixed at the estimated values in the first period of each country. We consider three sets of counterfactuals which first, fix only female wage discrimination (τ_{oj}); second, fix only gender norms ($\triangle A_{oj}$); and third fix gender norms and female wage discrimination simultaneously.

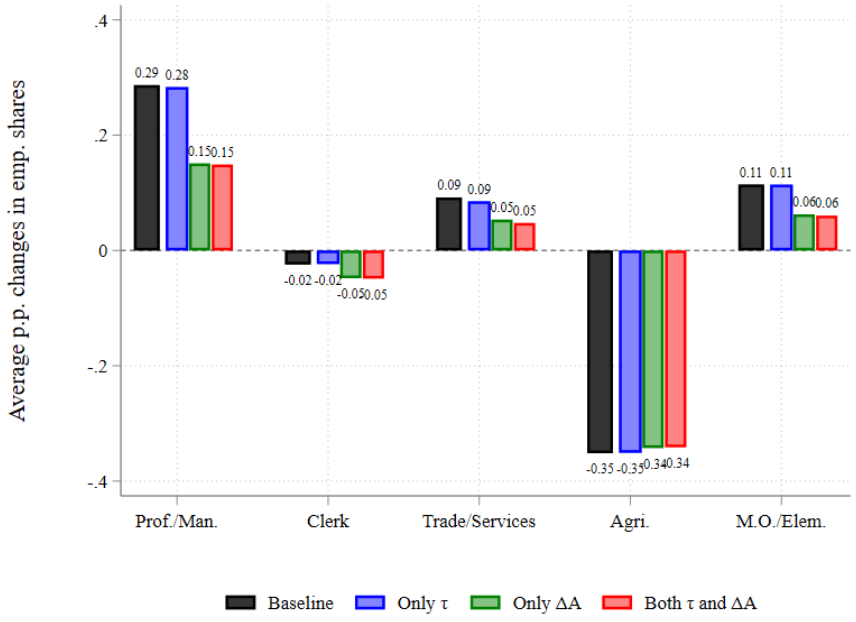
Conditional on these parameters, we then solve in each counterfactual for workers' employment choices and for the occupation-sector-specific wages and sectoral prices that ensure that goods markets and labor markets clear in equilibrium. Given non-homothetic preferences, consumption baskets can endogenously adjust in counterfactuals due to income effects. To quantify the importance of changes in gender barriers for each country, we then compare the counterfactual path of structural transformation and real income growth to the path that we actually observe in the data. More specifically, we compute the fraction of growth of some variable of interest (employment share, sectoral output, etc.) that can be attributed to changes in gender barriers over time as $1 - \hat{g}/g$, where g is the observed annualized growth rate of a variable of interest and $\hat{g}$ is the growth rate of the same variable in the counterfactual with fixed gender barriers.

Figure 4: Effect of Gender Barriers on Sectoral and Occupational Employment Shares

(a) Employment Shares Across Sectors



(b) Employment Shares Across Occupations



Notes: This figure reports the average percentage point changes (between the first and last year) in sectoral and occupational employment shares, averaged across the countries of our core sample. Figure (a) shows changes across sectors and Figure (b) across occupations. The labor force participation rate (LFPR) is defined as 1- the share of individuals in the home sector. The four bars for each sector/occupation correspond to (1) the changes observed in the data, (2) the changes in the first counterfactual which fixes wage discrimination τ_{ojct} to the values of the first period for each country, (3) the changes in the second counterfactual which fixes gender norms $\triangle A_{ojct}$, and (4) the changes in the third counterfactual which fixes both τ and $\triangle A$ to the first period. For each country, average percentage point changes add to zero across all sectors/occupations (when changing the sign of changes in labor force participation). Changes in labor force participation are identical for Figures (a) and (b) so that we omit it from Figure (b).

Effect of Changes in Gender Barriers on Structural Transformation

We first examine the importance of changes in gender barriers on structural transformation through changes in sectoral and occupational employment shares. In Figure 4, we compare average annual percentage point changes in sectoral and occupational employment shares in the data and in each counterfactual over the last decades, averaging across the six countries of our core sample.

Employment Transitions across Sectors: Figure 4(a) shows that the labor force participation rate (LFPR) increased on average by 0.12 pp per year in our sample. The employment share decreased in the agricultural sector on average by 0.38 pp per year, while it increased by 0.05 pp in manufacturing and by 0.45 pp in the service sector.

These employment transitions would have been very different if gender barriers had not changed over time, i.e., if both wage discrimination and gender norms had remained fixed at the 1970s levels for each country (red bars). The employment share in the service sector would have increased by only 0.21 pp per year on average. This finding implies that changes in gender barriers explain on average 53% (1-0.21/0.45) of the observed increase in the service sector employment share in our sample. Changes in gender barriers further explain 40% of the increase in manufacturing (1-0.03/0.05), while agricultural employment shares are less affected by gender norms and decrease similarly in all counterfactuals. Changes in gender norms were particularly important for explaining the observed increase in labor force participation. Without changes in gender norms, labor force participation would have decreased by 0.13 pp per year instead of the observed 0.12 pp increase. Comparing the effects of changes in gender barriers across counterfactuals shows that changes in wage discrimination alone have very small effects (blue bars) while changes in gender norms are an important driver of observed employment transitions (green and red bars).

Employment Transitions across Occupations: Figure 4(b) shows the changes in employment shares across occupations. We omit changes in labor force participation as these are identical to the numbers reported in Figure 4(a). Overall, we observe in our sample a large reallocation of workers away from agricultural occupations (-0.35 pp per year on average) towards professional occupations (+0.29 pp per year on average). Craft, trade, and service jobs as well as machine operating and elementary occupations increase by approximately 0.1 pp per year on average, while clerical jobs decrease slightly on average. If gender barriers had remained fixed at their 1970s values in each country, labor force participation would have decreased and average pp changes in employment shares would have been approximately 50% lower in professional, crafts and trade, as well as machine operating and elementary occupations. Changes in gender norms (as opposed to changes in wage discrimination) drive again the largest share of these effects.

Impact of Changing Gender Barriers on Sectoral Output

We now turn to calculating how much of the growth in real sectoral output (measured by real value added) over the last decades is explained by changes in gender barriers. We report these results in Columns (1)-(3) of Table 3 for agriculture, manufacturing and services respectively. We report the results for each country,

Table 3: Change in Sectoral Output Explained by Gender Barriers

	Sectoral Output		
	Agri.	Manf.	Services
	(1)	(2)	(3)
IND	-0.05	0.01	0.06
IDN	-0.07	-0.11	0.08
BRA	0.22	0.23	0.38
MEX	0.07	0.17	0.29
CAN	-0.01	0.16	0.29
USA	0.10	0.13	0.27
AVG	0.04	0.10	0.23

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.

along with a sample average at the bottom. From Column (1), changes in gender barriers account on average for 4% of growth in agricultural output. This ranges from negative 5-7% in India and Indonesia, to over a fifth in Brazil. From Column (2), changes in gender barriers account for 10% of growth in manufacturing output on average, which ranges from 1% in India, -11% in Indonesia, and around 15-20% in the other countries. Last, from Column (3), we see that 23% of the growth in service sector output can be attributed to changes in gender barriers on average, ranging from 6-8% in India and Indonesia, to around 30% in Mexico, Canada, and the U.S., and to 38% in Brazil. Put together, this implies that a decline in gender barriers was an important driver of structural transformation. Changes in gender barriers are particularly important for the transition towards the service sector in middle- and high-income countries, but changes in gender barriers also explain a non-trivial fraction of the output growth in manufacturing.

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

Acemoglu, Daron and Veronica Guerrieri, "Capital Deepening and Nonbalanced Economic Growth," *Journal of Political Economy*, 2008, *116* (3), 467–498.

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.

Bandiera, Oriana, Ahmed Elsayed, Anton Heil, and Andrea Smurra, "Economic Development and the Organisation of Labour: Evidence from the Jobs of the World Project," *Journal of the European Economic Association*, 2022, *20* (6), 2226–2270.

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

Caunedo, Julieta and Elisa Keller, "Capital-Embodied Structural Change," Technical Report, Unpublished Manuscript 2023.

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

Conte, Bruno, "Climate change and migration: the case of africa," 2022.

Cravino, Javier and Sebastian Sotelo, "Trade-induced structural change and the skill premium," *American Economic Journal: Macroeconomics*, 2019, *11* (3), 289–326.

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.

Fabrizio, Stefania, Lisa Kolovich, and Monique Newiak, "Pursuing Women's Economic Empowerment," *International Monetary Fund*, 2018.

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, Charyl Doss, Douglas Gollin, and Markus Poschke, "The gender division of work across countries," *Technical Report*, 2019.

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.

_ , _ , and _ , "Growth and Structural Transformation," *Handbook of Economic Growth*, 2014, *2*, 855–941.

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.

Huneeus, Federico and Richard Rogerson, "Heterogeneous paths of industrialization," *Review of Economic Studies*, 2024, *91* (3), 1746–1774.

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.

Kuznets, Simon, "Modern economic growth: findings and reflections," *The American economic review*, 1973, *63* (3), 247–258.

Lagakos, David and Martin Shu, "The role of micro data in understanding structural transformation," *Oxford Development Studies*, 2023, *51* (4), 436–454.

Lee, Munseob, "Allocation of Female Talent and Cross-Country Productivity Differences," *Unpublished Manuscript*, 2022.

Lewis, Logan T, Ryan Monarch, Michael Sposi, and Jing Zhang, "Structural change and global trade," *Journal of the European Economic Association*, 2022, *20* (1), 476–512.

Maddison, Angus, "Economic growth and structural change in the advanced countries," *Western Economies in Transition, eds.: I. Leveson and W. Wheeler. London: Croom Helm*, 1980.

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.

_ and Christopher A. Pissarides, "Structural Change in a Multisector Model of Growth," *American Economic Review*, March 2007, *97* (1), 429–443.

Ngai, L Rachel, Claudia Olivetti, and Barbara Petrongolo, "Gendered change: 150 years of transformation in US hours," Technical Report, National Bureau of Economic Research 2024.

OECD, "Gender Wage Gap Indicator," *doi: 10.1787/7cee77aa-en (Accessed on 04 November 2023)*, 2023.

Olivetti, Claudia, "The female labor force and long-run development: the american experience in comparative perspective," *Human Capital in History: The American Record, Ed. L. Platt Boustan, C. Frydman R.A. Margo University of Chicago Press*, 2014.

Porzio, Tommaso, Federico Rossi, and Gabriella Santangelo, "The Human Side of Structural Transformation," *American Economic Review*, August 2022, *112* (8), 2774–2814.

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.

Sposi, Michael, Kei-Mu Yi, and Jing Zhang, "Deindustrialization and industry polarization," Technical Report, National Bureau of Economic Research 2021.

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

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

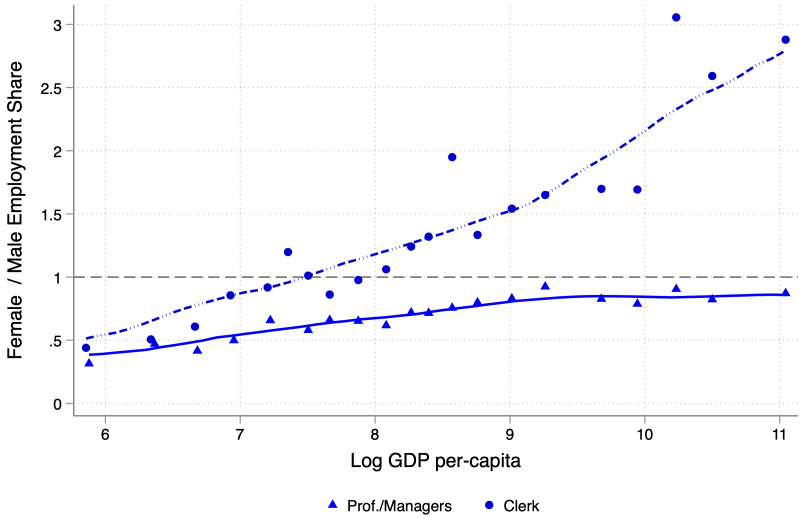
Ying, Feng, Jia Ren, and Michelle Rendall, "The reversal of the gender education gap with economic development," *Unpublished Manuscript*, 2023.

ONLINE APPENDIX: NOT FOR PUBLICATION

A Appendix Tables and Figures

A.1 Additional Empirical Evidence

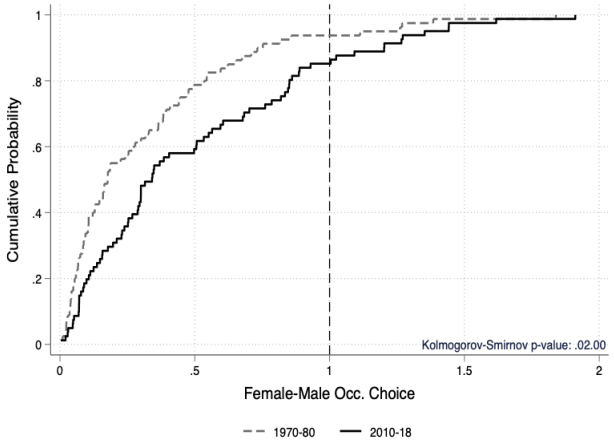
Figure A.1.1: Employment Gap in Professional and Clerical Occupations



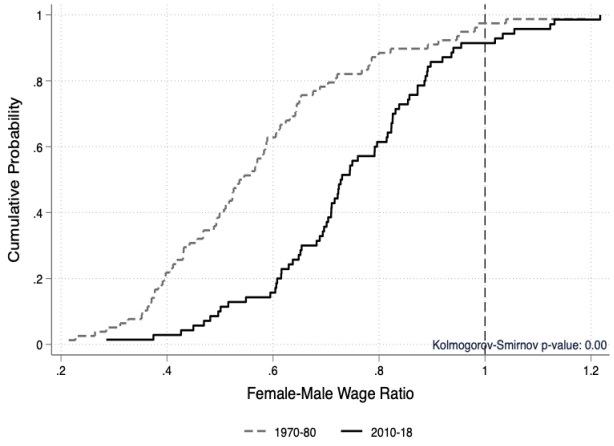
Notes: The above figures reports a non-parammetric fit between the ratio of the female to male employment share in professional and managerial occupations (blue triangles) and clerks (blue rounds) on the vertical axis, and log of real GDP per-capita (in 2010 USD) is on the vertical axis.

Figure A.1.2: Distribution of Female-to-Male Employment and Wage Ratios

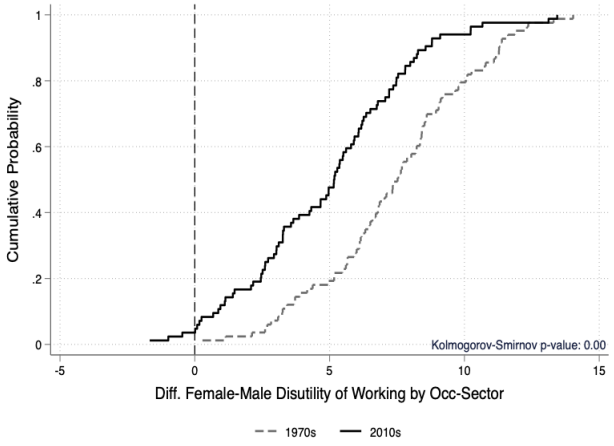
(a) Employment Ratio



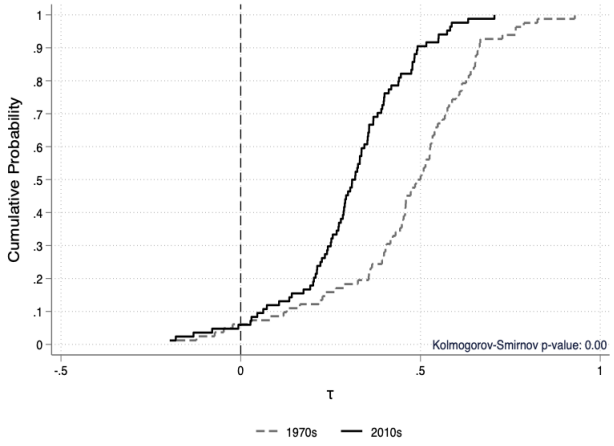
(b) Wage Ratio



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



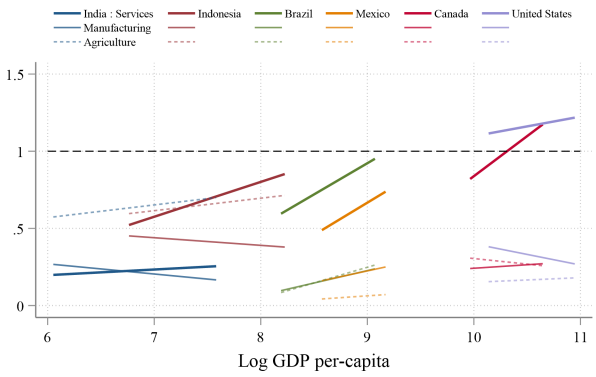
(d) Wage Discrimination (τ_{ojct})



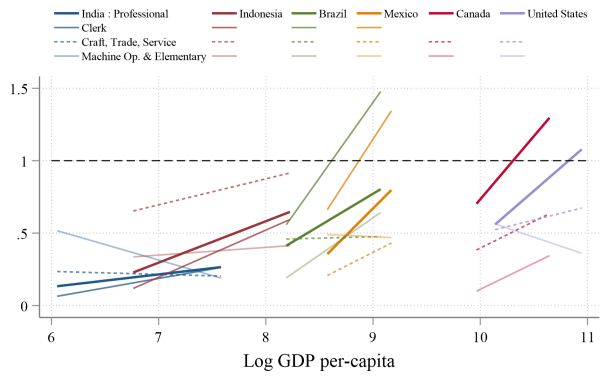
Notes: The above figure plots the CDF of the employment and wage gaps across all sectors and occupations. Figures (a) and (b) plot the employment and wage gaps using the first (dotted line) and last year (solid line) for each country. Figures (c) and (d) show the same distributions for gender norms ($\triangle A_{ojct}$) and female wage discrimination (τ_{ojct}). 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.1.3: Employment and Wage Ratios over Time

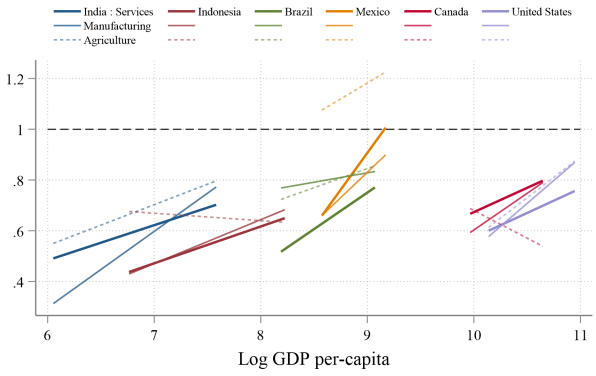
(a) Employment Ratio by Sector



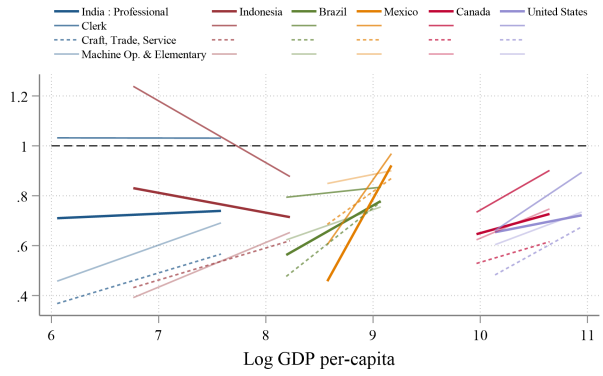
(b) Employment Ratio by Occupation



(c) Wage Ratio by Sector



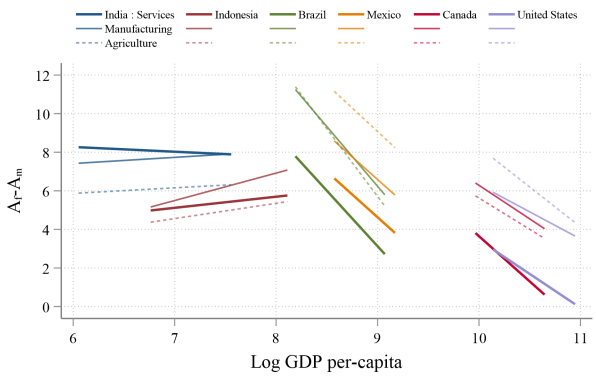
(d) Wage Ratio by Occupation



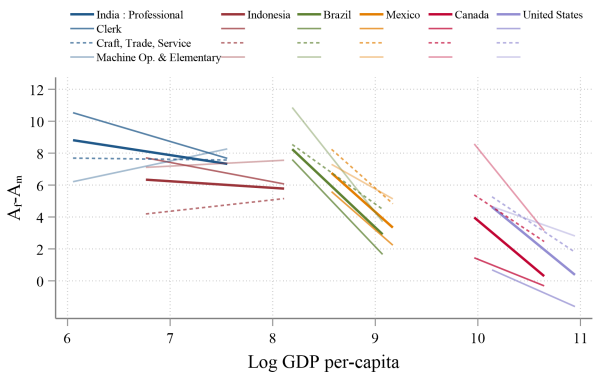
Notes: This figure plots the female-to-male employment and wage ratios for selected countries over time against the log of real GDP per-capita in constant 2010 US dollars. The time period covers between 1970 and 2018 depending on data availability and the horizontal dimension of the graph shows how fast countries grew during the sample period. Employment ratios divide the share of women working in an occupation or sector by the share of men. Wage ratios divide the average wage of women in an occupation or 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 A.1.4: Gender Norms and Wage Discrimination Across Countries Over Time

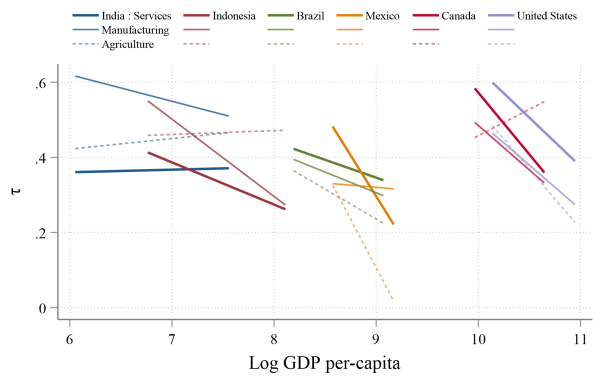
(a) Gender Norms ($\triangle A_{oj}$) by Sector



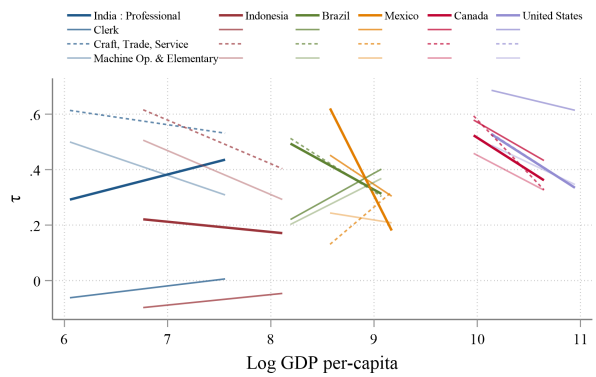
(b) Gender Norms ($\triangle A_{oj}$) by Occupation



(c) Wage Discrimination (τ_{oj}) by Sector



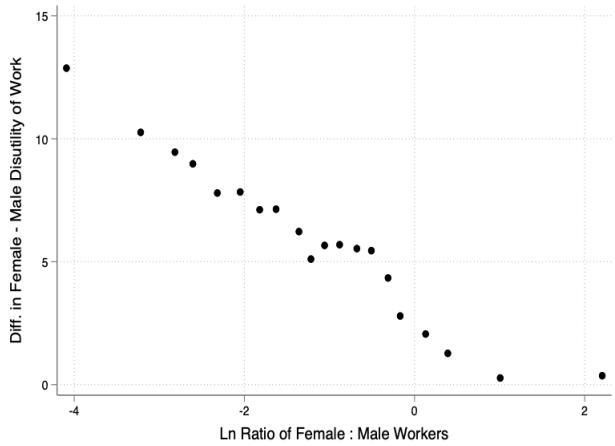
(d) Wage Discrimination (τ_{oj}) by Occupation



Notes: This figure plots the estimated gender norms ($\triangle A_{oj}$) and wage penalties (τ) for countries in our core sample for the first and last year in our sample, against the log of real GDP per-capita in constant 2010 US dollars. 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.

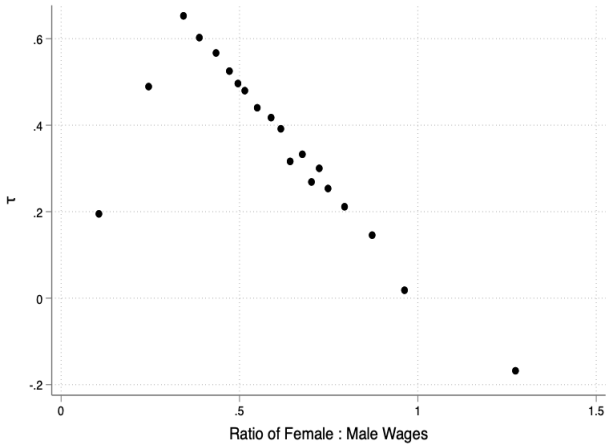
A.2 Model Fit and Correlations with Measured Gender Norms/Legal Indicators

Figure B1: Correlations of 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



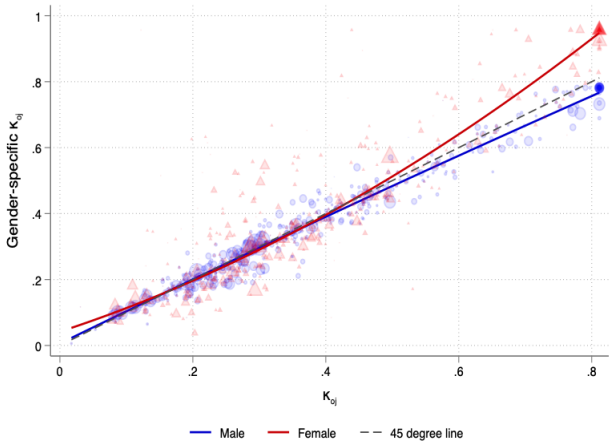
Notes: The above figure shows a binned scatter plot of the correlation between our estimated female wage penalty τ and observed wage gaps in all occupation-sectors and across all country-years.

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

	Coefficient	S.E.	p-value
	(1)	(2)	(3)
Panel A. Correlation with $\triangle A$			
Overall WBL Index	-0.33	(0.09)	0.00***
Index of Household Norms	-0.20	(0.08)	0.01**
Can a woman be head of household in the same way as a man?	-0.16	(0.19)	0.38
Is there legislation specifically addressing domestic violence?	-0.43	(0.14)	0.00***
Can a woman obtain a judgment of divorce in the same way as a man?	-0.32	(0.16)	0.05**
Does a woman have the same rights to remarry as a man?	-0.32	(0.16)	0.05**
Panel A. Correlation with τ			
Overall WBL Index	-0.05	(0.02)	0.03**
Index of Workplace Equality	-0.05	(0.02)	0.02**
Can a woman get a job in the same way as a man?	-0.03	(0.06)	0.65
Does the law prohibit discrimination in employment based on gender?	-0.09	(0.06)	0.12
Is there legislation on sexual harassment in employment?	-0.11	(0.04)	0.00***
Are there criminal penalties for harrassment at workplace?	-0.11	(0.04)	0.00***
Observations	525		

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.

Figure B3: Comparison of κ_{oj} and Gender-Specific κ_{ojg}



Notes: The above figure reports the correlation between the returns to ability that is the same across gender (κ_{oj}) on the vertical axis with gender-specific returns (κ_{ojg}) on the vertical axis for men (blue circles) and women (red triangles). Each point is an estimate of κ obtained from Mincerian wage regressions described in Section 6. The dash black line is the 45 degree line.

B Mathematical Appendix

B.1 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:

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

B.2 Deriving Sectoral Expenditure Shares

Individuals of gender g, ability z, working in an occupation-sector oj have non-homothetic preferences over sectors j (agriculture, manufacturing, and services) given by the implicit constraint:

$$\sum_j \theta_j^{\frac{1}{\sigma}} \left(\frac{C_j}{C^{\varepsilon_j}}\right)^{\frac{\sigma-1}{\sigma}} = 1 \tag{12}$$

which they maximize subject to the standard budget constraint: $\sum_j p_j C_j = I_{ojg}(z)$. Taking the first-order conditions with a Lagrangean multiplier λ, we get:

$$\begin{aligned}
\left(\frac{\sigma-1}{\sigma}\right)\theta_j^{\frac{1}{\sigma}} C_j^{\frac{-1}{\sigma}} C^{\frac{-\varepsilon_j(\sigma-1)}{\sigma}} &= \lambda p_j \\
\Rightarrow \left(\frac{\sigma-1}{\sigma}\right) \underbrace{\sum_j \theta_j^{\frac{1}{\sigma}} C_j^{\frac{\sigma-1}{\sigma}} C^{\frac{-\varepsilon_j(\sigma-1)}{\sigma}}}_{=1} &= \lambda \underbrace{\sum_j p_j C_j}_{=I} \\
\Rightarrow \lambda &= \left(\frac{\sigma-1}{\sigma}\right) \times \frac{1}{I}
\end{aligned}$$

Substituting for λ from the above, using $C = I/P$ in the FOC, and rearranging the terms, we get:

$$\begin{aligned} C_j &= \theta_j \left(\frac{p_j}{I}\right)^{-\sigma} \left(\frac{I}{P}\right)^{\varepsilon_j(1-\sigma)} \\ &= \theta_j \left(\frac{p_j}{P}\right)^{-\sigma} \left(\frac{I}{P}\right)^{\sigma+\varepsilon_j(1-\sigma)} \end{aligned} \tag{13}$$

Lastly, defining expenditure shares $\varphi_j \equiv p_j C_j / I$, we get:

$$\varphi_j = \theta_j \left(\frac{p_j}{P}\right)^{1-\sigma} \left(\frac{I}{P}\right)^{(\varepsilon_j-1)(1-\sigma)}$$

B.3 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:

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

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 Calibration 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: preference shifters $\mathcal{U} = \{\theta_j\}_{\forall j}$, occupation-sector productivity parameters $\{\{T_j\}_{\forall j}, \{\gamma_{oj}\}_{\forall oj}\}$, and gender barriers $\mathcal{B} = \left\{\{\tau_{oj}\}_{\forall oj}, \{A_{ojg}\}_{\forall ojg}\right\}$. To calibrate these parameters, we match data on: (i) men's and women's occupation-sector choices $\Pr(oj|g)$; (ii) gender gaps in average hourly wages ($\overline{\text{wage}}_{ojf}/\overline{\text{wage}}_{ojm}$); (iii) share of wage-bill in sector j allocated to occupation o; (iv) sectoral wage-bills. The calibration proceeds in two loops:

Outer loop: Guess sectoral prices (p_j) and preference shifters (θ_j).

Inner loop: For a guess of the occupation-sector wage rates w_{oj} and amenities A_{ojg}, solve for the occupational choices of individuals. To do that:

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

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

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 I.2: 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:

$$\overline{H}_{ojg} \equiv \frac{H_{ojg}}{N_{ojg}} = \frac{\int_z \Pr(oj|g,z) z^{\kappa_{oj}} dF(z)}{\int_z \Pr(oj|g,z) dF(z)} \tag{C.2}$$

Solve for a new guess of wages w_{ojg}^{new} to perfectly fit observed gender-specific average wage in occupation-sector oj so that:

$$w_{ojg} = \frac{\overline{\text{wage}}_{ojg}}{\overline{H}_{ojg}} \tag{C.3}$$

Iterate on the inner loop until convergence.
τ_{oj} can then be calculated as: $\tau_{oj} = 1 - w_{ojf}/w_{ojm}$ and $\gamma_{oj} = \frac{\sum_g w_{ojm} H_{ojg}}{\sum_o \sum_j \sum_g w_{ojm} H_{ojg}}$.

Turning to the outer loop, we then find prices and preference shifters as follows:
Step O.1: Given the convergence in the inner loop, we can calculate aggregate expenditure in a sector j by using Equation (2) and summing across individuals in the economy (where we drop ojg for notational ease):

$$p_j C_j = \theta_j p_j \underbrace{\int_z \Pr(oj|g,z) \left(\frac{p_j}{P(p, I(z))}\right)^{-\sigma} \left(\frac{I(z)}{P(p, I(z))}\right)^{\varepsilon_j + \sigma} dF(z)}_{=\omega_j} \tag{C.4}$$

From the first-order condition of the firm and market clearing in the goods market, we know that $p_j C_j = w_j H_j$. To update θ_j, we fix a "base-sector" m (manufacturing in our case) and normalize $\theta_m = 1$. Then from Equation C.4, we regress $\frac{w_j H_j}{w_m H_m}$ on $\frac{\omega_j}{\omega_m}$ across countries to estimate θ_j^{new} (for agriculture and services).

Step O.2: We then use market clearing and the firms' optimality conditions, along with Equation C.4 to update prices (after normalizing $p_M = 1$):

$$p_j^{new} = \left[\frac{w_j H_j}{\int_z \Pr(oj|g,z) P(p, I(z))^{-\sigma} \Big(I(z)/P(p, I(z)) \Big)^{\varepsilon_j + \sigma} dF(z)} \right]^{\frac{1}{1-\sigma}} \qquad \text{(for } j = \{A, M, S\} \text{ sectors)}$$

$$p_k^{new} = \left[\frac{1}{\alpha_k} \times \frac{w_k H_k}{w_S H_S} \right]^{\frac{1}{1-\sigma_s}} \times P_S \qquad \text{(for } k = \{hs, ms\} \text{ sectors)}$$

Iterate on the outer loop until convergence. Lastly, using the production function and firms' optimality condition (Wage bill$_j = p_j Y_j = p_j T_j H_j$) we calculate $T_j =$ Wage bill$_j / p_j H_j$.

D Data Appendix

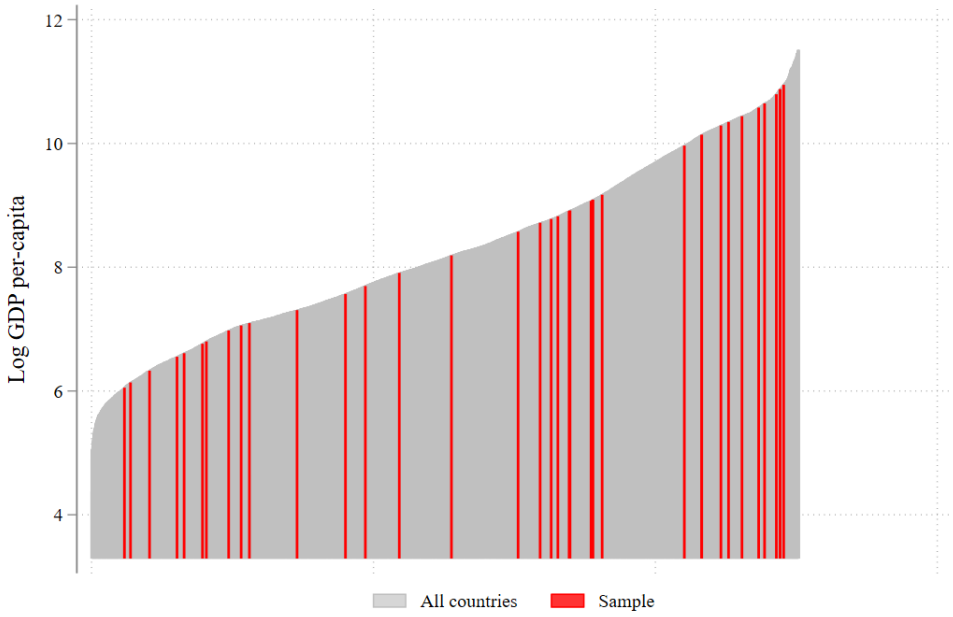
D.1 Sample Coverage

Table D1: 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.

Figure D1: 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 D2: Coverage of the Six Countries of the Core Sample

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.

D.2 Sample Definition, Industry-Occupation Classifications, and Data Cleaning

D.2.1 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* $\leq$ *3*
   drop if empstatd$\geq$*320 & empstatd*$<$*370* .
4. In case a specific survey does not have such detailed information on the employment status, we adjust the number of inactive people using the share by gender of inactive individuals in school, prision, disabled, etc. from the closest survey with data available.
5. We classify all individuals who are unemployed or out of the workforce in the "home sector" i.e., *empstat* $==$ *2* — *empstat* $==$ *3*
6. Table D3 provides the classification of education categories into years of education.

D.2.2 Industry and Occupation Classifications

1. Tables D4 and D5 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$\geq$*1 occ* $\leq$ *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*$\geq$ *5 occ* $\leq$ *6*

Table D3: 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

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* $\leq$ *3*

D.3 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.

D.3.1 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).

Table D4: 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

D.3.2 Measurement of Hours Worked

There are three types of weekly worked hours variables available in IPUMS:

1. **HRSWORK1:** reports the person's total worked hours per week at all jobs. It combines information of the following two variables.
2. **HRSUSUAL1:** reports the usual number of hours the person works in a typical week across all jobs.
3. **HRSACTUAL1:** reports the person's actual number of hours worked per week at all jobs.

Each variable has an corresponding categorized variable with the suffix 2 (ex.: HRSWORK2). Whenever we do not have the uncategorized variable but we do have the categorized one, we use the midpoint of the categories.

1. We use the variable HRSWORK1 or the midpoints of HRSWORK2 categories as much as possible. If none of those variables are present, we use either usual or actual hours depending on the case.
2. We trim the sample at 100 hours.
3. 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.
4. In case no information is available on hours worked for a specific survey, we replace it by the gender-industry-occupation average of the closest year.
5. 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.

Table D5: Classification of Occupations

Code	Occupation	Classification	Occ. Nb	Sector
1	Legislators, senior officials and managers	Professional	(1)	M,S
2	Professionals	Professional	(1)	
3	Technicians and associate professionals	Professional	(1)	M,S
4	Clerks	Clerks	(2)	M,S
5	Service workers and shop and market sales	Services Workers	(3)	M,S
6	Skilled agricultural and fishery workers	Skilled Agri.	(4)	A
7	Crafts and related trades workers	Craft/Trade Wrkrs	(5)	M,S
8	Plant and machine operators and assemblers	Plant & Machine	(6)	M,S
9	Elementary occupations	Elementary	(7)	A,M,S
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 the first three ISCO-88 codes as shown in column 3 and 4. The last column shows which occupations are represented in the respective sectors. Within the agricultural sector, we model only skilled agricultural and elementary occupations as the other occupation-sector cells are very sparse in many gender-country-years. More information on the ISCO-88 classification of occupations can be found here.

D.3.3 Countrywise Availability of Income and Hours

Brazil: 1970, 1980, 1991, 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: HRSWORK1 available for 1991, 2000 and 2010. HRSWORK2 availabe for 1980. For 1970 we use the gender-industry-occupation averages of the 1980 survey.
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: HRSWORK1 available 1981 onwards, HRSUSUAL2 available for 1971. 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, 2009, 2011, 2018

1. Since there is no information on hours worked for any year within IPUMS, we replace the source of data with the Indian Employment-Unemployment Survey (EUS) until 2009 and the Periodic Labor Force Survey (PLFS) for 2018. Both surveys have been harmonized the World Bnak's Global Labor Database project. (GLD)
2. Hours worked GLD: The variable records the hours of work last week for the individual's main job.
3. Wage GLD: Wage income is defined as "Last wage payment, primary job, excl. bonuses, etc (7-day ref period)". An additional variable records the frequency of payment, allowing us to convert it to weekly wages and then to hourly wages.

Indonesia: 1976, 1995, 2015, 2018

1. Income: INCWAGE from the previous month is available for both the years.
2. Hours worked: HRSWORKED1 available for 1995 and HRSACTUAL1 available for 1976.
3. Wage = Income/(4*Hrs Worked)
4. For 2015 and 2018, we complement the IPUMS data for Indonesia with data from the SAKERNAS survey. We pool data from 2013-2016 (2014 missing) and 2017-2019 respectively to increase the sample size.
5. Wages and hours worked in SAKERNAS: Wage income is available for 2013, 2015, 2016, 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: 1960, 1970, 1990, 1995, 2000, 2005, 2010, 2015

1. Income: INCTOT and INCEARN in the previous month. Not available for 2005.
2. Hours worked: HRSWORKED1 available for 1990, 1995, 2000, 2010. For 1960 and 1970 we use the gender-industry-occupation averages of the 1990 survey. For 2015 we use the gender-industry-occupation average of the 2010 survey.
3. Wage = Income/(4*Hrs Worked)

USA: 1960, 1970, 1980, 1990, 2000, 2005, 2010, 2015

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 don't need these two variables separately.
2. Hours worked: HRSWORK1 Available 1980 onwards. HRSWORK2 available for 1960 and 1970. MONTHSWRK is also available for all years
3. Wage = Income/(4*Hrs Worked*Months Worked)

D.3.4 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.