Gender Barriers, Structural Transformation, and Economic Development*

Gaurav Chiplunkar
University of Virginia

Tatjana Kleineberg
World Bank

February 2025

Abstract

The representation and significance of women in the labor force have grown significantly over the past five decades around the globe. Using nationally representative data from more than 90 countries, this paper documents distinct *gender* patterns in employment transitions across both sectors and occupations during this period. Using a model of occupational and sectoral choice and focusing on six major economies, we find that declining gender barriers – defined as gender-specific distortions in employment and wages – were a key driver of the observed rise in female labor force participation, expansion of the service sector, and increases in real GDP per capita from 1970 to 2018, but with substantial variation across countries.

Keywords: Economic Growth, Structural Transformation, Gender, Misallocation

*We thank Claudia Olivetti for a very insightful discussion, as well as Julieta Caunedo, Cheryl Doss, Kevin Donovan, Doug Gollin, Charles Gottlieb, Michael Peters, Barbara Petrongolo, Markus Poschke, Tommaso Porzio, Todd Schoellman, and participants in conferences and seminars at the American University, Boston University, Columbia University, Cornell University, CREI, Dartmouth University, Northwestern University, the Richmond Federal Reserve, University of Virginia, the World Bank, the ASSA in San Antonio, the CEPR STEG Annual Workshop, the EFEG session of the NBER Summer Institute, NEUDC, SED, SEA Annual Conference, and WADES for helpful suggestions and comments. Facundo Ulivarri provided excellent research assistance. This project benefited from funding from the World Bank Research Support Budget (RSB). The findings, interpretations, and conclusions expressed in this paper are entirely those of the authors. They do not necessarily represent the views of the World Bank and its affiliated organizations, or those of the Executive Directors of the World Bank or the governments they represent. 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

The role of women in the labor force has undergone a profound transformation over the past century, evolving from a largely marginalized position (Boserup, 1975) to one of increasing representation, influence, and leadership in the global workforce (Goldin, 2024). In the 1970s, for instance, there were, on average, fewer than four women for every ten men in professional and managerial positions across a large cross-section of countries. By the 2010s, this number increased to more than eight women for every ten men. This growing representation of women in the labor market not only reflects progress toward gender equality, but also has significant implications for countries' aggregate employment structures, productivity, and economic growth.[1] Over the same period, economies around the world have undergone fundamental structural changes, with employment transitioning from agriculture to manufacturing and services (Kuznets, 1973).

In this paper, we study the evolution of gendered labor market opportunities, economic growth, and structural change empirically and theoretically. First, we use micro-data from 91 countries at different stages of development and over five decades to document new evidence on the link between economic development and gendered labor market outcomes. We then propose a model to study these patterns. Our framework integrates two broad explanations for the observed evolution of gendered labor market outcomes and economic growth: first, economic mechanisms such as sector- or occupation-specific technological change, income effects in consumption, and improved education; and second, gender-specific distortions arising from discrimination and social norms. We structurally estimate our model to six major economies, which offer high-quality gender-disaggregated data on *both* wages and employment over five decades. We then use the estimated model to study the effects of changing gender distortions on structural change and aggregate growth in these countries. We find that gender distortions are important and have declined in most countries. On average, the decline in gender distortions in our sample countries accounts for a substantial share of overall economic growth and reallocation of employment in the labor market.

We begin by providing new evidence on gendered employment patterns across occupations and sectors. Our analysis draws on censuses and nationally representative household and labor force surveys from 91 countries, spanning five decades (1970-2018) and covering over 300 country-years. Using these data, we first show that the canonical process of

[1] See Cuberes and Teignier (2014); Hsieh et al. (2019); Ranasinghe (2024a,b); Chiplunkar and Goldberg (2024).

structural transformation – 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, they exit agriculture and primarily transition out of the labor force. At higher levels of development, women re-enter the labor force, primarily in the service sector. Unlike men, women's representation in manufacturing remains low across all development levels. We further document gender differences in occupational choices *within* sectors, showing that men increasingly work in professional and industrial jobs in richer countries, while women's employment (relative to men) expands faster in clerical jobs but slower in professional and industrial jobs.

In addition to gender employment gaps, we also analyze the evolution of gender wage gaps for a select set of countries which provide high-quality micro-data on hourly earnings. We find that wage gaps have narrowed over time across all occupation-sector pairs. However, women in the 2010s still earned, on average, only 70 cents for every dollar earned by men. Furthermore, we observe a weak correlation between gender wage gaps and countries' development levels – in contrast to gender employment gaps that tend to decrease in richer countries. Additionally, gender employment and wage gaps do not necessarily correlate across occupation-sector pairs.

We then propose a model of occupational and sectoral choice that can speak to these empirical findings and that builds on Hsieh et al. (2019). The documented changes in gendered employment and wage patterns – both across countries and within countries over time – can be partly driven by previously studied economic channels, such as income effects in consumption, sector- or factor-specific technological change, or shifts in the supply of, or returns to human capital.[2] Specifically, technological change and income effects can increase labor demand in the "female-gendered" service sector, disproportionately benefiting women (Ngai and Petrongolo, 2017; Fan et al., 2023). Additionally, the narrowing of gender education gaps has expanded the supply of women's human capital (Evans et al., 2021; Feng et al., 2023). On the other hand, the documented changes in gendered labor market outcomes can also be driven by a decline in barriers that women face in certain occupations and sectors, which would increase women's *labor supply* in them.[3]

Our model enables us to disentangle the effects of economic mechanisms from those

[2]See for example Ngai and Pissarides (2007); Boppart (2014); Herrendorf et al. (2013, 2014); Comin et al. (2021); Porzio et al. (2022); Huneeus and Rogerson (2024).

[3]See Jayachandran (2021) for a review of social norms around women's labor force participation and Hyland et al. (2020) for a review of legislative changes over time and across countries that improve women's rights for fair remuneration and for accessing job opportunities.

of changing gender barriers. In the model, individuals differ by gender and ability, while occupation-sector pairs differ in wage rates per unit of human capital, returns to ability, and costs or barriers that men and women face when working in them. Individuals receive idiosyncratic preference shocks across occupation-sector pairs and then choose their occupation and sector. We model the option of not participating in the labor force as an additional choice, which we call the "home sector" (Hsieh et al., 2019; Ngai and Petrongolo, 2017). On the demand side, consumers have non-homothetic preferences, allowing them to shift consumption away from agriculture and toward services as their income rises – a key driver of structural transformation emphasized in the literature.[4] Wage rates per human capital unit in each occupation-sector and prices in each sector are endogenously determined in equilibrium.

Following Hsieh et al. (2019), we define two types of gender barriers: first, "gender norms" which capture the excess utility costs that women incur relative to men when working in a given occupation-sector pair; and second, female "wage discrimination", which captures differences in remuneration that women receive *per human capital unit* relative to men. In the absence of distortions, an efficient allocation of talent ensures that high-ability workers sort into sectors and occupations that offer the highest returns to their ability. However, gender barriers distort these choices, shifting workers away from the jobs where they have a comparative advantage. For instance, high-ability women may opt out of the labor market or work in clerical jobs with low returns to their ability rather than pursuing professional or managerial jobs with higher returns. Such distortions therefore lead to a misallocation of talent, reducing aggregate productivity and output.

We estimate the model to a core sample of six large economies – India, Indonesia, Brazil, Mexico, Canada and the United States – for which we have employment and hourly wage data spanning four to five decades. Gender barriers in each occupation-sector and country-year are quantified as "wedges" which ensure that our model equations precisely match the data on gender employment and wage gaps after accounting for economic fundamentals and individuals' employment and consumption choices in equilibrium. For example, the model expresses the observed gender wage gap in each occupation-sector as a function of a model-implied gender gap in human capital and the wage discrimination women face relative to men. To measure gender barriers, we assume that men and women have the same productivity within each occupation-sector, conditional on their model-implied human capital levels.[5]

[4] See Boppart (2014); Herrendorf et al. (2014); Comin et al. (2021).

[5] In robustness checks, we relax this assumption and allow for gender-specific comparative advantage

Our estimates of gender barriers reveal three main findings. First, gender norms have declined significantly over time, particularly in the service sector and in professional, managerial, and clerical jobs. Second, wage discrimination has also declined across all sectors and occupations, but relatively less in professional jobs. Third, trends in gender barriers vary considerably across countries: they worsened in India, showed little improvement in Indonesia, and declined substantially in Brazil, Mexico, Canada, and the United States. As a validation, we show that our estimated gender barriers strongly correlate across countries and over time with empirically measured gender norms and labor market constraints.

Using our estimated model, we then quantify the contribution of changing gender barriers to countries' structural transformation and economic growth over the past five decades, in the spirit of a growth accounting exercise. To do so, we simulate counterfactual employment and growth paths for each country while holding gender barriers in each occupation-sector fixed at their calibrated values from the 1970s but allowing all other variables, such as technologies and skill distributions, to evolve over time. We find that declining gender barriers had large effects on female labor force participation, structural transformation, and economic growth in all countries of our sample, except India. Averaging across countries, we find that changes in gender barriers between 1970 and the 2010s accounted for 40%-45% of employment growth and 20%-35% of output growth in the manufacturing and service sectors. Declining gender barriers particularly contributed to the rapid rise of the service sector, revealing a novel channel for explaining the premature deindustrialization observed in developing countries in recent decades (Rodrik, 2016; Huneeus and Rogerson, 2024). On average, changes in gender barriers accounted for 28% of observed real GDP per capita growth in our sample, with significant heterogeneity across countries – ranging from 10% in Indonesia and 28% in the United States to 40% in Canada and Mexico, and over 50% in Brazil. In contrast, worsening gender barriers in India imply that economic growth would have been higher if these barriers had not changed.

In a series of additional analyses, we examine the mechanisms and robustness of our findings. First, we show in our baseline model that declining gender norms had a larger effect on economic growth than reductions in wage discrimination. Second, we show that our aggregate results are similar in a model where workers sort into occupations and sectors based on idiosyncratic *talent* shocks rather than *preference* shocks, as in the baseline

across occupation-sector pairs.

model. Despite similar aggregate effects, the relative importance of wage discrimination and gender norms reverses in a model with talent shocks, which aligns with the findings from Hsieh et al. (2019) and is further discussed in Section 7.2. Third, we show that non-homothetic preferences *amplify* the effects of declining gender barriers on aggregate and sectoral outcomes because falling barriers make economies richer, shifting consumption and labor demand away from agriculture and toward manufacturing and services. Lastly, we perform a robustness check that accounts for gender-specific comparative advantage across occupation-sector pairs, which we quantify from a benchmark economy – the United States in 2015 – which we assume to have no wage discrimination, following Lee (2024). After adjusting for these estimated gender productivity differences in all country-years, changes in gender barriers remain relevant, accounting for 21% of observed growth in GDP per capita, compared to 28% in the baseline.

Related Literature. A large literature documents gender disparities in the labor market, but only a few recent studies provide cross-country comparisons and examine the macroeconomic implications of gender roles. Cuberes and Teignier (2014), Albanesi et al. (2023), and Olivetti et al. (2024) provide excellent reviews. Several papers study the link between gender roles and structural transformation in the United States, including Ngai and Petrongolo (2017), Ngai et al. (2024), Moro et al. (2017), and Rendall (2018). Across countries, Gottlieb et al. (2023), Bridgman et al. (2018) and Bick et al. (2022) use rich time-use surveys to examine the implications of gender differences in hours worked. Our quantitative approach – which measures gender barriers as wedges and quantifies their importance for aggregate outcomes – aligns most closely with Hsieh et al. (2019), Lee (2024), and Chiplunkar and Goldberg (2024). Hsieh et al. (2019) study the aggregate implications of talent misallocation due to gender and racial barriers in the United States. They find, similar to us, that a reduction in barriers explains 20-40% of US economic growth over the last decades. Lee (2024) documents that gender barriers in poor countries are higher in non-agriculture sectors, also aligning with our findings.

While the literature on structural transformation emphasizes sectoral employment transitions, we further document transitions across occupations within sectors, with a focus on gender. We find that gender differences across sectors are more important in poorer countries, while occupational differences matter more in richer countries, in line with recent work by Caunedo et al. (2023) and Bandiera et al. (2022). We therefore develop one consistent framework to *jointly* study the extensive margin of entering the labor force and the intensive margin of choosing a particular sector and

occupation. The model incorporates the drivers of employment choices that have been emphasized in the macro-labor and structural transformation literature, which include heterogeneity in workers' ability and hence their comparative advantage, gender-specific barriers, differential returns to ability across occupation-sectors (Hsieh et al., 2019; Cassan et al., 2024), as well as non-homothetic demand, sector-specific technological change, and changing human capital supply (Ngai and Pissarides, 2007; Herrendorf et al., 2014; Boppart, 2014; Ngai and Petrongolo, 2017; Comin et al., 2021; Porzio et al., 2022; Feng et al., 2023).

The paper is organized as follows: Section 2 describes the data and provides new evidence on the link between economic development and gendered labor market outcomes. Section 3 presents the theoretical model and Section 4 describes the identification and model quantification. Sections 5 and 6 present the estimation results and the counterfactual simulations. Section 7 examines mechanisms and tests the robustness of our analysis. Section 8 concludes.

2 Empirical Facts

2.1 Data

Data Sources and Sample Description. We use data from the Integrated Public Use Microdata Series (IPUMS International, 2020), which provides harmonized individual-level data on demographic and employment variables from nationally representative censuses, and household and labor force surveys for many countries and years. We extract employment information by occupation, sector, and gender for 91 countries and 305 country-years. The time coverage ranges from 1960 to 2020 and includes, on average, 3-4 rounds of data for each country. Online Appendix C shows that the data has good coverage over the time period and the entire development spectrum. We complement the IPUMS data with labor force surveys from the World Bank Global Labor Database (GLD) and the World Bank i2d2 database.[6]

For our quantitative exercise, we use a core sample of six large economies – India, Indonesia, Mexico, Brazil, Canada, and the United States – which additionally offer high-quality micro data on hours worked and compensation, enabling us to calculate workers' hourly wages. In addition, these countries provide data for more than four

[6] A description of the database can be found at https://github.com/worldbank/gld.

decades, allowing us to study the process of structural transformation and the evolution of gendered labor market outcomes over a long period of time. The six economies span a wide income spectrum (Figure OA.7 in the Online Appendix) and cover around 25-30% of the world population. For Indonesia, we complement the IPUMS census data with the SAKERNAS labor force survey to extend the time coverage of wage data to 2018. For India, we use survey data from the National Sample Surveys (NSS), the Employment-Unemployment Survey (EUS), and the Periodic Labor Force Survey (PLFS) up to 2018.

Sector and Occupation Classification. We define three market sectors – agriculture, manufacturing, and services – by assigning the more detailed industry codes from IPUMS and the labor force surveys to these three categories. The categorization follows the literature (Herrendorf et al. (2013); Herrendorf and Schoellman (2018)) and is outlined in Online Appendix C, Table OA.6. We define a fourth sector as the "home sector" to which we attribute unemployed and inactive individuals, following Hsieh et al. (2019).

At the occupation level, we use 1-digit ISCO codes. We aggregate the top three ISCO codes, which results in the following seven occupation categories: (1) professionals, (2) clerks, (3) service workers, (4) skilled agricultural workers, (5) crafts and trades workers, (6) plant and machine operators, and (7) elementary occupation workers. To minimize measurement errors arising from small sample sizes in specific occupation-sector-gender cells, we impose two restrictions. First, we consider only two occupations within the agriculture sector: skilled agricultural workers and elementary occupations.[7] Second, we model the occupation of "skilled agricultural workers" only in the agriculture sector and we assign all workers in that occupation to the agriculture sector. More information is provided in Online Appendix C.

Construction of the Data Set. First, we calculate employment shares of men and women in each occupation-sector and country-year. We allocate part-time workers proportionally between the reported occupation-sector and the home sector based on their hours worked. This adjustment is particularly important because women are more likely to work part time. Second, we compute hourly earnings for men and women in each occupation-sector for the limited sample of countries that provide data on hours worked and earnings. Further details on the data cleaning and dataset construction are available in Online Appendix C.

[7]We attribute the small percentage of workers who work in different occupations in the agriculture sector proportionally to the same occupation in the other two sectors.

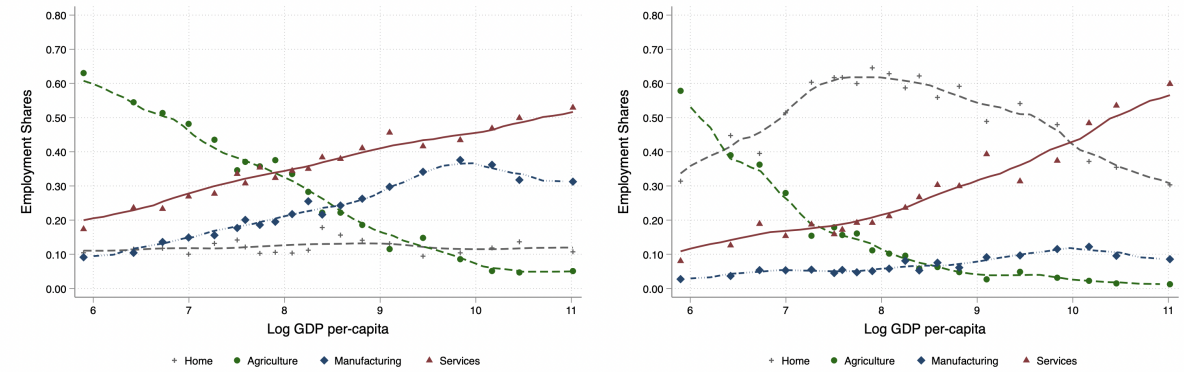
2.2 Sectoral Employment Transitions by Gender

We first document how sectoral employment choices differ between men and women and how these gendered employment choices change with economic development.

Figure 1: Sectoral Employment Transitions by Gender

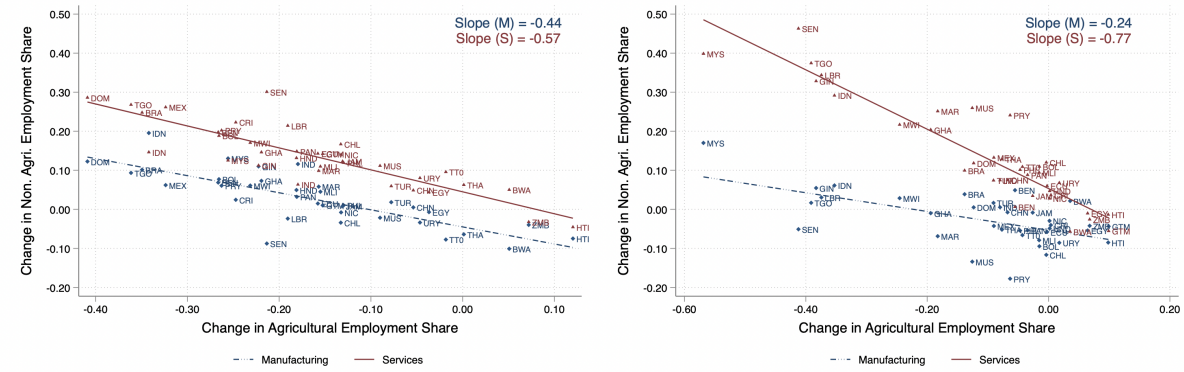
(a) Transitions Across Sectors: Men

(b) Transitions Across Sectors: Women



(c) (De)Industrialization: Men

(d) (De)Industrialization: Women



Notes: Figures (a) and (b) present a non-parametric fit of sectoral employment shares for men and women against log real GDP per-capita in 2010 US dollars. The shares are calculated as a fraction of the overall population, which includes the share of inactive individuals in the "home sector". The sample pools data from 91 countries and 305 country-years between 1960-2019. Figures (c) and (d) show how changes in male and female employment shares in agriculture (horizontal axis) correlate with employment changes in manufacturing (dash blue line) and services (solid red line). The sample is restricted to low- and middle-income countries. For each country, we calculate changes in sectoral employment shares over time, conditional on labor force participation, using the first (1970-1990) and last (2000-2019) years of available data. We report the coefficients from a linear regression model for manufacturing (Slope (M)) and services (Slope (S)) in the upper right corner of each graph.

Figure 1a plots the sectoral employment shares of men against countries' log real GDP per capita, pooling the data across all countries and years. The graph shows the well-documented pattern of structural transformation: poorer countries have high employment

shares in agriculture, while richer countries have lower agricultural employment shares and higher shares in manufacturing and market services. The share of men in the home sector (i.e., not in the labor force) is small and relatively stable across different levels of economic development.

Figure 1b shows different employment patterns for women. Similar to men, women's employment shares in agriculture are high in poor countries and low in richer ones. However, in low- and middle-income countries, lower agricultural employment shares among women coincide with higher shares in the home sector, i.e., with low female labor force participation. In richer countries, female labor force participation is higher again, with most women working in the service sector.[8] Female employment shares in manufacturing remain consistently low across all levels of economic development.

These gender differences in sectoral employment suggest that higher female labor force participation might increase employment in services relative to manufacturing. Hence, changes in gender roles and rising female labor force participation over recent decades may have contributed to the premature deindustrialization and service-led growth observed in many developing and emerging countries (Rodrik, 2016; Fan et al., 2023; Huneeus and Rogerson, 2024). To further examine this possibility, Figures 1c and 1d present how changes in agricultural employment shares correlate with corresponding changes in manufacturing and services between the 1970s and 2010s, separately for men and women, and restricting the sample to low- and middle-income countries. For each country, we calculate changes in sectoral employment shares over time, conditional on labor force participation, using the first (1970-1990) and last (2000-2019) years of available data. Figure 1c shows that a 10 percentage-point (p.p.) decline in men's agricultural employment share is, on average, associated with an increase of 4.4 p.p. in manufacturing and 5.7 p.p. in services. For women, the shift toward services is stronger with employment shares increasing by 2.4 p.p. in manufacturing and by 7.7 p.p. in services (Figure 1d).[9]

[8]Figure 1b confirms the U-shaped relationship between female labor force participation and economic development, which has been extensively studied in the literature (Goldin, 1995, 2024; Mammen and Paxson, 2000).

[9]Figure OA.1 in the Online Appendix shows that gender differences are more pronounced when including changes in labor force participation. For men, a 10 p.p. decline in the combined share of agricultural and home sector employment is, on average, associated with a 5 p.p. increase in manufacturing and services. For women, these numbers are 1.7 p.p. for manufacturing and 8.4 p.p. for services. Figure OA.2 in the Online Appendix additionally shows that the flattening of total manufacturing employment shares in Brazil and Mexico during the past decades is driven by women, while men's employment shares in manufacturing continue to rise.

2.3 Occupational Employment Transitions by Gender

We next examine gender differences in occupational employment transitions. While the literature on structural transformation primarily focuses on sectoral transitions, a recent literature has highlighted the importance of occupational choices for talent allocation, aggregate productivity, and growth (Hsieh et al., 2019; Cassan et al., 2024).

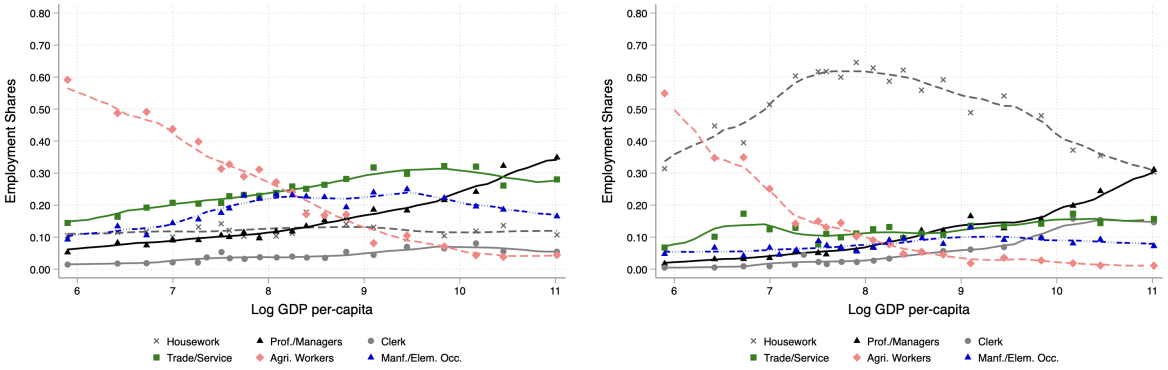
Figure 2 shows that employment in agricultural jobs is high in poor countries but lower in richer countries for both genders, reflecting the broader employment decline in the agriculture sector. In richer countries, men's employment shares increase across all other occupations. The increase is particularly large in professional jobs where men's employment shares rise from less than 8% in poor countries to over 35% in the richest ones (black triangles). In contrast, male employment shares in clerical jobs (gray dots) remain relatively low, ranging between 5% and 8% across all development levels.

For women, smaller agricultural employment shares in middle-income countries do not lead to employment growth in other occupations but instead to lower female labor force participation. In richer countries, female employment shares expand rapidly in professional occupations (from 5% to 35%) and in trade and service jobs (from 8% to 20%). Unlike men, women's employment shares rise substantially in clerical jobs (from 1% to 20%), while their employment remains low in machine-operating and elementary occupations.

Figure 2: Occupational Employment Shares by Gender and by Sector

(a) Transition Across Occupations: Men

(b) Transition Across Occupations: Women



Notes: This Figure presents a non-parametric fit of occupational employment shares against log real GDP per-capita in 2010 US dollars, separately for men and women. The shares are calculated as a fraction of the overall population, including inactive individuals in the home sector. The sample pools data from 91 countries and 305 country-years between 1960-2019.

More generally, sectors employ occupations in different proportions, creating a direct link between sectoral and occupational employment transitions. Figure OA.3 in the Online Appendix shows this link by plotting occupational employment shares separately *within* the manufacturing and service sectors against countries' GDP per capita. Overall, the service sector uses professional and clerical occupations more intensely at all development levels. The manufacturing sector is highly concentrated in trade and service occupations in poorer countries but diversifies across occupations in richer ones. Gender differences across occupations within these sectors are relatively small in poorer countries but increase in richer ones. The largest gender gap emerges in clerical occupations, where women account for the entire employment increase with development. Men and women both contribute to the rise in professional and managerial jobs in rich countries, but men remain 5-10 p.p. more likely to work in these occupations than women.

2.4 Gender Segregation across Sectors and Occupations

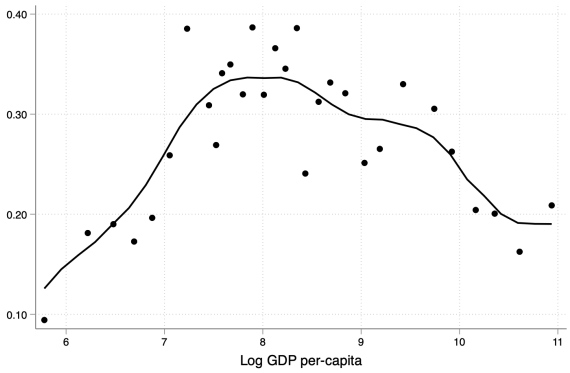
Having documented gender differences in employment choices across sectors and occupations, we now use the Theil Information H-Index to systematically measure gender segregation across occupation-sector pairs in each country-year. Similar to the Gini coefficient, the H-Index measures segregation on a normalized scale between 0 and 1 by using the entropic "distance" between the observed empirical distribution and an ideal gender-equal distribution. A higher value indicates greater segregation, while zero indicates perfect gender equality. The index allows decomposing total gender segregation across occupation-sector pairs into two additive components, which separately measure gender differences *across* sectors and across occupations *within* sectors. Appendix A.1 formally defines the index and decomposition.

Figure 3a plots the segregation across occupation-sector pairs against countries' real GDP per capita, pooling data across all countries and years. The figure shows an inverted U-shape relationship, indicating that gender segregation is lower in poor and rich countries but high in middle-income countries. Figure 3b plots the share of total gender segregation that is explained by differences across occupations *within* sectors. This component is significantly higher in richer countries, where it accounts for up to 50% of overall segregation, compared to 20% in poorer countries.[10] This finding is consistent

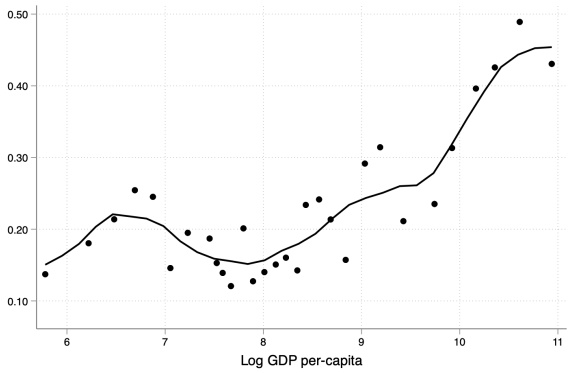
[10]Changes in labor force participation are attributed to across-sector variation, as inactive individuals are assigned to the home sector, which consists of a single occupation. To assess the importance of the extensive employment margin, we also compute the H-index conditional on labor force participation. The results remain similar and the importance of the within-sector variation in richer countries becomes

Figure 3: H-Index of Gender Segregation across Sectors and Occupations

(a) Aggregate Index



(b) Fraction Explained Within Sector



Notes: This figure plots the H-Index, which measures gender segregation across occupation-sector pairs, against the log of real GDP per capita in 2010 US dollars. The sample pools data from 91 countries and 305 country-years between 1960-2019. Figure (a) shows the level of the segregation index and Figure (b) shows the share of overall segregation that is explained by segregation across occupations *within* sectors.

with Bandiera et al. (2022), who document that the division of labor and the number of occupations increase with development.

2.5 Gender Gaps within Countries over Time

We next use a core sample of six large economies – India, Indonesia, Brazil, Mexico, Canada, and the United States – to examine the evolution of gender employment and wage gaps within countries over time and to implement our quantitative exercise. These countries provide individual-level data on employment outcomes and hourly wages spanning at least four decades. To document patterns over time, we compute female-to-male employment ratios by dividing the employment shares of women in each occupation-sector pair by those of men. Similarly, female-to-male wage ratios divide the average wages of women in each occupation-sector pair by those of men. Ratios of 1 indicate gender parity. Ratios below (above) 1 indicate that women have lower (higher) representation or wages in the occupation-sector relative to men.

Gender Employment Gaps over Time. Changes in employment ratios over time largely confirm the cross-sectional patterns documented above, particularly highlighting the transition of women from the home sector to the service sector. In the 1970s, women were, on average, 8 times more likely to stay at home than men in our sample

even more pronounced.

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	8.07	3.63		
Agriculture	0.27	0.27	0.56	0.73
Manufacturing	0.28	0.28	0.50	0.69
Services	0.65	0.93	0.54	0.71
Panel B. Occupations				
Professional	0.33	0.82	0.58	0.70
Clerks	1.08	1.41	0.80	0.86
Craft, Trade, Service	0.39	0.57	0.46	0.67
Agricultural	0.25	0.24	0.56	0.72
Machine Op., Elementary	0.39	0.50	0.51	0.73
Panel C. Countries				
India			0.56	0.60
Indonesia			0.54	0.73
Brazil			0.57	0.72
Mexico			0.65	0.75
Canada			0.44	0.71
United States			0.47	0.70

Notes: This table reports female-to-male employment ratios (Columns 1-2) and wage ratios (Columns 3-4) in the 1970s and 2010s. The ratios divide employment shares (average wages) of women in each occupation sector by those of men. Panels A-C report average ratios for sectors, occupations, and countries. Averages of employment (wage) ratios are weighted by the employment (income) of each occupation-sector.

countries, which decreased to 3.6 in the 2010s (Table 1, Panel A). Women remain most underrepresented in the agriculture and manufacturing sectors, which employ only 2.7 women for every 10 men, with little change over time. In contrast, women's employment expanded rapidly in the service sector, increasing from 6.5 to 9.3 women for every 10 men over the past decades. Turning to occupations, Panel B of Table 1 shows that women were underrepresented in all occupations except clerical jobs in the 1970s. The largest catch-up occurred in professional and managerial jobs, where women's employment increased from 3.3 to 8.2 for every 10 men. Women's overrepresentation in clerical jobs grew further over time, rising from 11 to 14 women for every 10 men. Employment gaps in other

occupations improved over time, but remain large in the 2010s.

Gender Wage Gaps over Time. Table 1 further documents changes in gender wage gaps over time (columns 3-4). In the 1970s, women earned, on average, 50-56 cents for every dollar earned by men across all sectors. Wage gaps exhibited more variation across occupations in the 1970s with women earning 46 cents per dollar earned by men in craft, trade, and service occupations but 80 cents in clerical jobs. Gender wage gaps have narrowed significantly over time, but remain persistent in the 2010s, with women earning, on average, around 70 cents for every dollar earned by men in most sectors and occupations. Comparing wage gaps across countries shows little correlation with economic development (Panel C). In the 1970s, gender wage gaps across countries range from 0.44 to 0.65, which improves to 0.6-0.75 in the 2010s. Moreover, we find that the levels and trends of wage and employment ratios are not necessarily correlated across occupations-sectors. For instance, in the 2010s, women were better represented in the service sector than in agriculture or manufacturing, yet wage ratios remained similar across sectors (Panel A, Columns 2 and 4). Likewise, gender employment gaps in professional and clerical occupations narrow rapidly over time and with economic development, but their wage gaps remain significant at all development levels.[11]

2.6 Discussion and Model Implications

Overall, our findings emphasize distinct gender patterns in labor force participation, occupational and sectoral choices, and wage gaps along the development spectrum and within countries over time. The documented narrowing of gender employment and wage gaps over time can be partly explained by well-studied economic drivers of structural transformation and economic growth, such as income effects in consumption, sector-specific technological change, or rising educational attainment. For instance, technological change or income effects have increased the *labor demand* in the predominantly female service sector, disproportionately benefiting women (Ngai and Petrongolo, 2017). However, gendered labor market outcomes also respond to changes in the barriers that women face in the labor market, such as gender norms or wage discrimination.

To disentangle the effects of economic forces from gender-specific barriers, and to assess their relative importance, we develop a quantitative model of occupational and

[11] Figure OA.4 in the Online Appendix plots each country's employment and wage ratios by sector and occupation against GDP per capita. These figures thus illustrate changes in each gender gap both *within* countries over time and *across* countries with economic development.

sectoral choice that incorporates the above-mentioned mechanisms and can be taken to the data. The fact that gender differences in occupational choices are particularly important at higher levels of economic development underscores the need for a unified framework that incorporates labor force participation as well as occupational and sectoral choices to study the effects of gender roles on macroeconomic outcomes over long time periods and across countries with varying levels of economic development. Our model framework allows us: (i) to measure gender barriers for each occupation-sector and country-year by matching model equations to observed gender employment and wage gaps; and (ii) to implement counterfactual simulations which quantify the effects of changes in gender barriers on countries' paths of structural transformation and economic development, in the spirit of a growth accounting exercise.

3 Model

We now present the set-up of the occupational and sectoral choice model (Roy, 1951), solve individuals' employment and consumption choices and firms' production decisions, and define the equilibrium of the model.

3.1 Setup, Gender Barriers, and Preferences

The economy consists of a mass N_g of individuals of gender g, who are either male m or female f. Each individual has an ability z that is drawn from a gender specific ability distribution $F_g(z)$. There are O occupations and J sectors. Occupation-sector pairs oj differ in their returns to ability z, denoted by κ_{ojg}, which can potentially vary across gender. Workers receive idiosyncratic preference shocks ν^i_{oj} across all occupation-sector pairs oj.

Gender Barriers. Two gender-specific barriers can distort the efficient allocation of workers across occupation-sector pairs. First, each worker of gender g incurs a utility cost A_{ojg} when working in a given occupation-sector oj. This utility cost can capture general amenities that influence the desirability of working in the occupation-sector pair. However, systematic differences in utility costs between men and women likely indicate underlying gender norms or barriers. We therefore define the *excess* utility cost that women incur relative to men when working in the same occupation-sector as a gender-specific barrier. Throughout the paper, we refer to this distortion as "gender

norms" and we formally define it as: $\lambda_{oj} = 1 - A_{ojf}/A_{ojm}$.[12] The second gender distortion arises if men and women earn a different wage rate *per human capital unit* in the same occupation-sector pair. We refer to this wage differential as female "wage discrimination", which we formally define as the percentage difference between gender-specific wage rates w_{ojg} in each occupation-sector pair: $\tau_{oj} = 1 - w_{ojf}/w_{ojm}$.[13]

Consumer Preferences. Individuals have non-homothetic CES preferences across agriculture, manufacturing, and service goods, allowing sectoral expenditure shares to change with income levels (Comin et al. (2021)).[14] Preferences for a consumption bundle C over sectors $k = \{A, M, S\}$ are defined by the following constraint:

$$\sum_k \theta_k^{\frac{1}{\sigma}} \left(\frac{C_k}{C^{\varepsilon_k}} \right)^{\frac{\sigma-1}{\sigma}} = 1, \tag{1}$$

where $\sigma > 0$ is the standard CES elasticity of substitution across sectors, ε_k is the non-homothetic elasticity of substitution, and θ_k are sectoral preference shifters. Setting $\varepsilon_k = 1 \forall k$ corresponds to the standard homothetic CES preferences. Individuals maximize their utility subject to the budget constraint $\sum_k p_k C_k = I_{ojg}(z)$. The income earned by an individual of gender g and ability z who works in occupation-sector pair oj is equal to $I_{ojg}(z) = w_{ojg} \exp(z\kappa_{ojg})$, where w_{ojg} is the wage rate per human capital unit in oj, which can vary across genders due to gender wage discrimination, and κ_{ojg} are occupation-sector-specific returns to ability.

Services are a CES composite of home and market services $S = \{hs, ms\}$, following Ngai and Petrongolo (2017) and Ngai et al. (2024) so that:

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

[12] Our model and employment data cannot distinguish whether gender differences in utility costs, λ_{oj}, stem from intrinsic gender-specific preferences or from restrictive and coercive gender norms. For example, the underrepresentation of women relative to men in engineering jobs could reflect either a genuine preference differential (in the absence of constraints) or societal factors such as stigma, discouragement, or discrimination of women. Given this limitation, we restrict our analysis to a growth accounting exercise and we refrain from any normative welfare analysis, in line with the literature (Hsieh et al., 2019; Chiplunkar and Goldberg, 2024; Ranasinghe, 2024a; Cuberes and Teignier, 2014; Lee, 2024).

[13] In the baseline model, we assume that men and women with the same ability draw z have the same productivity in each occupation-sector pair. Hence, we do not allow for any gender-specific comparative advantage across occupation-sector pairs. We relax this assumption in Section 7.4.

[14] Such income effects have been emphasized in the literature as an important driver of structural transformation, e.g., Herrendorf et al. (2013); Alder et al. (2022); Comin et al. (2021); Conte (2022); Fan et al. (2023).

where $\sigma_s > 0$ is the elasticity of substitution between home and market services, and α_s are the preference weights with $\sum_s \alpha_s = 1$.

3.2 Consumption and Employment Choices

Consumption Choice across Sectoral Goods. With non-homothetic preferences, individuals of gender g and ability z who work in occupation o and sector j spend the following share of their income on goods from sector k:[15]

$$\varphi_k(I_{ojg}(z), p) = \theta_k \left(\frac{p_k}{P_{ojg}(z)} \right)^{1-\sigma} \left(\frac{I_{ojg}(z)}{P_{ojg}(z)} \right)^{(\varepsilon_k - 1)(1-\sigma)}, \tag{2}$$

where the price index $P_{ojg}(z)$ for an individual with income $I_{ojg}(z)$ is implicitly defined:

$$P_{ojg}(z) = \left[\sum_k \theta_k p_k^{1-\sigma} \times \left(\frac{I_{ojg}(z)}{P_{ojg}(z)} \right)^{(\varepsilon_k - 1)(1-\sigma)} \right]^{\frac{1}{1-\sigma}}. \tag{3}$$

Appendix A.2 provides the proof. From Equation (2), the ratio of expenditure shares across sectors k and m is given by:

$$\frac{\varphi_k}{\varphi_m} = \frac{\theta_k}{\theta_m} \left(\frac{p_k}{p_m} \right)^{1-\sigma} \left(\frac{I_{ojg}(z)}{P_{ojg}(z)} \right)^{(\epsilon_k - \epsilon_m)(1-\sigma)}. \tag{4}$$

Similar to homothetic CES preferences, the ratio of sectoral expenditure shares depends on preference shifters (θ_k/θ_m), and a substitution effect that arises from relative sectoral prices (p_k/p_m). With non-homothetic preferences, sectoral expenditure shares further depend on an income effect, whose importance is governed by the relative difference between the elasticities $\varepsilon_k - \varepsilon_m$. Within the service composite, the respective consumption shares of home and market services are given by the standard CES expression:

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

where φ_S is the expenditure share on the entire service composite and P_S is the CES price index of the service composite (cf. Online Appendix B.1).

[15] We denote the consumption sectors by k and the sector of individuals' employment by j. Workers' occupation-sector pair of employment oj matters for workers' income which affects their sectoral expenditure shares due to non-homothetic preferences.

Employment Choice Across Occupation-Sector Pairs. The utility of a worker i of gender g and ability z who works in an occupation-sector pair oj is given by:

$$U^i_{ojg}(z) = A_{ojg} \times V(I_{ojg}(z), P_{ojg}(z)) \times \nu^i_{oj},$$

where A_{ojg} are the gender-specific utility costs that an individual of gender g incurs when working in occupation-sector oj, ν^i_{oj} are idiosyncratic preference shocks across occupation-sector pairs, and $V(I_{ojg}(z), P_{ojg}(z))$ is the indirect utility from consumption, which equals real income $I_{ojg}(z)/P_{ojg}(z)$. We assume that preference shocks ν^i_{oj} are i.i.d. and follow a Fréchet distribution with dispersion parameter σ_ν. Under this assumption, the share of workers of gender g and ability type z working in occupation-sector pair oj is equal to:

$$\pi_{oj|gz} = \frac{\left[A_{ojg} \times V(I_{ojg}(z), P_{ojg}(z))\right]^{\sigma_\nu}}{\sum_{j'} \sum_{o'} \left[A_{o'j'g} \times V(I_{o'j'g}(z), P_{o'j'g}(z))\right]^{\sigma_\nu}}. \tag{6}$$

Equation (6) shows that workers' employment choices depend on the gender-specific utility cost A_{ojg} and the real income that they receive when working in a given occupation-sector.

3.3 Production and Equilibrium

Human Capital Supply: We then aggregate across individuals' employment choices to solve for the aggregate supply of human capital in each occupation-sector pair oj:

$$H_{oj} = \sum_g N_g \int_z \pi_{oj|gz} \exp(z\kappa_{ojg}) dF_g(z), \tag{7}$$

which integrates over the human capital units of each ability z and gender g type, weighted by the probability that such a type chooses a given occupation-sector oj.

Production: A representative firm in each sector j produces output Y_j with a Cobb Douglas production function that differs in productivity B_j and uses human capital from each occupation as inputs with expenditure shares γ_{oj}. Firms' profit maximization is then given by:

$$\Pi_j = \max_{H_{oj}} p_j B_j \prod_o H_{oj}^{\gamma_{oj}} - \sum_o w_{oj} H_{oj}, \tag{8}$$

where p_j are sectoral prices and w_{oj} are the wage rates per human capital unit that equal workers' marginal product. Following Hsieh et al. (2019), we assume that men are paid the undistorted wage rate w_{oj} while women face wage discrimination τ_{oj} that introduces a wedge between their effective wage rate and their marginal product so that $w_{ojf} = (1 - \tau_{oj}) w_{oj}$.[16]

Equilibrium: The exogenous parameters of the model characterize preferences over sectoral consumption $\{\sigma, \sigma_s, \alpha_s, \{\varepsilon_j, \theta_j\}_{\forall j}\}$, production functions $\{\{B_j\}_{\forall j}, \{\gamma_{oj}, \kappa_{ojg}\}_{\forall oj}\}$, the dispersion of preference shocks across occupation-sector pairs $\{\sigma_\nu\}$, the mean and standard deviation of gender-specific ability distributions $\{\mu_{gz}, \sigma_{gz}\}$, and gender barriers $\{\{\tau_{oj}\}_{\forall oj}, \{A_{ojg}\}_{\forall ojg}\}$. Given these parameters, the equilibrium is defined in each country-year by a vector of sectoral prices and occupation-and-sector-specific wage rates $\{\{p_j\}_{\forall j}, \{w_{oj}\}_{\forall oj}\}$ so that: (i) workers make optimal consumption and employment choices; (ii) firms in each sector hire human capital from each occupation to maximize profits; (iii) labor markets clear in each occupation-sector pair; and (iv) good markets clear in each sector.

4 Model Quantification

To quantify our model parameters, we determine a set of parameters externally, either by fitting them directly to observed data moments, or by taking estimates from the literature. The remaining parameters are then obtained by fitting the model's equilibrium conditions to selected data moments using a simulated method of moments approach in an iterative algorithm that is described in Appendix B. Table 2 defines all model parameters and summarizes how they are determined.

[16] This wage discrimination can be micro-founded, as in Hsieh et al. (2019), by assuming that entrepreneurs have a disutility δ_{oj} for hiring women. In equilibrium, this disutility is then exactly compensated by the profits that entrepreneurs earn from paying women below their marginal product, so that entrepreneurs are on the margin indifferent between hiring men or women.

Table 2: Model Parameters and Baseline Values

Sym.	Definition	Determination	Values
Panel A. Preference Parameters			
σ	Elasticity of substitution across j	Comin et al. (2021)	0.25
ε_j	Income effect non-homoth. CES	Comin et al. (2021)	$\{0.01, 1, 1.3\}$
θ_j	Sectoral preference shifters	Calibrated across countries	
σ_ν	Dispersion of pref. shocks over oj	Hsieh et al. (2019)	2
σ_s	Elasticity of substitution b/w home and market services	Ngai and Petrongolo (2017)	2.3
α_s	Preference weight on home services	Ngai and Petrongolo (2017)	0.2
Panel B. Ability Distribution and Returns to Ability			
μ_{gz}	Avg. ability z by gender	Avg. Years of Schooling	
σ_{gz}	S.D. of ability z by gender	S.D. Schooling Yrs	
κ_{oj}	Returns to ability in oj	Mincer Regressions	
Panel C. Production Technology			
γ_{oj}	Expenditure shares on H_{oj}	Wage Expenditure by oj	
B_j	Sectoral productivity of sector j	Zero profit conditions	
Panel D. Gender Barriers			
λ_{oj}	Gender norms/excess utility cost	Gender Emp. Ratios	
τ_{oj}	Female wage discrimination in oj	Gender Wage Ratios	

Notes: This table defines all model parameters and describes how we determine them in our quantification exercise. The preference parameters listed in Panel A are constant across countries and time, but sectoral preference shifters θ_j can vary across countries. All other parameters can vary across countries and years.

4.1 Externally Determined Parameters

Preference Parameters: We set the dispersion of workers' preference shocks across occupation-sector pairs to $\sigma_\nu = 2$, following Hsieh et al. (2019) and Cassan et al. (2024). For the CES preferences over home and market services, we set the elasticity of substitution to $\sigma_s = 2.3$, and the preference weight on home services to $\alpha_s = 0.2$, following Ngai and Petrongolo (2017). For the non-homothetic CES preferences, we use estimates from Comin et al. (2021), setting the elasticity of substitution to $\sigma = 0.25$ and the parameters that discipline income effects to $\varepsilon_j = \{0.01, 1, 1.3\}$ for $j = \{A, M, S\}$.[17] We assume preference parameters to be the same across countries and over time. Sectoral preference shifters θ_j

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

can vary across countries and are estimated jointly with the remaining parameters as described in Section 4.2.

Ability Distributions: We assume that ability z follows a normal distribution with a mean and variance that can vary across gender, countries, and over time, so that $z \sim N(\mu_{gz}, \sigma_{gz})$. We calibrate the mean and variance of each distribution to the observed distribution of years of schooling for each gender, country, and year.[18]

Returns to Ability: We estimate the returns to ability for each occupation-sector κ_{oj} by estimating a Mincer regression, similar to Fan et al. (2023).[19] Recall that a worker of gender g and ability z who works in an occupation-sector pair oj earns income equal to: $I_{ojg}(z) = w_{ojg} \times \exp(z\kappa_{oj})$. To estimate κ_{oj} for the six countries in our core sample, we take the log of this structural equation and use our individual-level data to estimate the following Mincerian wage regression separately for each country-year:

$$\ln(\text{Wage}_i) = \alpha_{ojg} + \kappa_{oj}(\text{YrsSchool}_i) + \gamma_1 \text{Exp}_i + \gamma_2 \text{Exp}_i^2 + u_i, \tag{9}$$

where Wage_i is individual i's hourly income, 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 wage rates in a given occupation-sector-gender cell. The coefficients $\hat{\kappa}_{oj}$ estimates the returns to schooling for each occupation-sector pair. We assume no returns to education in the home sector, setting $\kappa_{home} = 0$.

4.2 Parameters Estimated within the Model

The remaining parameters to be calibrated are sectoral preference shifters $\{\{\theta_j\}_{\forall j}\}$, sectoral productivity and occupational expenditure shares $\{\{B_j\}_{\forall j}, \{\gamma_{oj}\}_{\forall oj}\}$, and gender barriers $\{\{\tau_{oj}\}_{\forall oj}, \{A_{ojg}\}_{\forall ojg}\}$. We now describe the intuition of our quantification strategy, while Appendix B describes the numerical procedure and iterative algorithm.

[18] We assume for now that the distribution of ability (or education) is exogenously determined and constant in our counterfactuals. This assumption can be relaxed by modeling education choices endogenously as a function of education costs and education returns. Holding the distribution constant in our counterfactuals provides a lower-bound for the importance of changes in gender barriers because declining gender barriers over time increased women's return to education and therefore would have increased their incentives to invest in human capital.

[19] In the baseline calibration, we restrict the returns to education κ_{oj} to be the same for men and women. We experimented with a model version in which the returns can vary by gender, estimating κ_{ojg}. Overall, we find little variation across gender in these parameters and our qualitative and quantitative results remain mostly unchanged in both versions.

Gender Norms (λ_{oj}): We infer the gender-specific utility cost of working in each occupation-sector A_{ojg} from the occupational and sectoral choices shown in Equation (6):

$$\pi_{oj|g} \propto \underbrace{A_{ojg}}_{\text{Utility Cost}} \times \underbrace{I_{ojg}/P_{ojg}}_{\text{Real Income}}.$$

This equation shows that the share of individuals of gender g who chooses to work in a given occupation-sector oj depends on the real income (or indirect utility) that they receive and the utility cost they incur when working in oj. Using Equation (6) and computing the real income term within our model, we can then infer A_{ojg} as a residual that exactly matches the observed employment shares by occupation, sector, and gender in each country-year. A_{ojg} are only identified *relative* to other occupation-sector pairs for each gender g, so we normalize $A_{home,g} = 1$ for men and women. For each gender, working in any other occupation-sector can incur either a utility cost or premium relative to working at home. Gender norms are then given by $\lambda_{oj} = 1 - A_{ojf}/A_{ojm}$, capturing the *excess* costs that women incur relative to men when working in occupation-sector oj relative to staying at home. For example, a value of $\lambda_{oj} = 0.2$ indicates that women face an *additional* cost equivalent to 20% of their real income in occupation-sector pair oj compared to men. A value of $\lambda_{oj} = 0$ implies gender parity in utility costs, while larger values indicate higher barriers for women relative to men.

Wage Discrimination (τ_{oj}): For each oj, the model expresses the average gender wage ratio as a function of a model-implied human capital ratio and wage discrimination:

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

where we use $w_{ojf} = (1-\tau_{oj})w_{ojm}$ and where $\overline{H}_{ojg}$ is the average human capital of workers with gender g, who choose to work in the occupation-sector pair oj. $\overline{H}_{ojg}$ is endogenously determined in equilibrium and depends on workers' ability distributions $F_g(z)$, occupation-sector-specific returns to human capital κ_{oj}, and the extent to which workers sort across occupation-sector pairs based on their ability.[20] Using Equation (10) and computing the gender human capital gap within our model, we can then infer female

[20] As previously mentioned, we assume in the baseline calibration that men and women have the same productivity in each occupation-sector pair conditional on their ability z, abstracting from any gender-specific comparative advantage across occupation-sector pairs. We relax this assumption in a robustness check in Section 7.4.

wage discrimination τ_{oj} as a wedge to exactly match the observed female-to-male wage ratios in each occupation, sector, country, and year.

Inferring Gender Barriers: Since average human capital by gender, occupation, and sector $\overline{H}_{ojg}$ and workers' real income I_{ojg}/P_{ojg} are jointly determined in equilibrium, we solve for them in an iterative algorithm, which simultaneously solves for equilibrium wage rates in each occupation-sector and sectoral prices. We impose no restrictions on the values of gender barriers across countries, time, or genders, allowing women, relative to men, to receive either a penalty or a premium in wages and utility when working in a given occupation-sector.

Sectoral Preference Shifters, Prices, and Productivity: In the joint estimation, we further quantify sectoral preference shifters θ_j, productivities B_j, and prices p_j to ensure that zero-profit conditions and goods market clearing conditions hold in each sector. These conditions require that consumers' expenditure and wage bills equal firms' revenue in each sector, so that $p_j Y_j = p_j C_j = WageBill_j = \sum_o w_{oj} H_{oj}$.

With non-homothetic CES preferences, sectoral expenditure shares for $j = \{A, M, S\}$ depend on the product between sectoral preference shifters and sectoral prices, as shown in Equation (2). We therefore first solve for the joint term, which we denote as $\psi_j \equiv \theta_j p_j^{1-\sigma}$, by matching the model-implied sectoral expenditure to the empirically observed sectoral wage bill.[21] In a second step, we then separate preference shifters and prices by making the identifying assumption that preference shifters can vary across countries but not over time. Under this assumption, *changes* in ψ_j over time pin down the growth in sectoral prices. Without loss of generality, we normalize prices for sectors $j = \{A, M, S\}$ to 1 in the first period of each country, which infers the values of preference shifters θ_j for each country. To solve for the separate prices of home and market services within the service composite, we use Equation (5) which is the standard CES expression for sectoral expenditure shares.

In the estimation procedure, we further match the data on real GDP per capita growth rates for each country. To compute real income or GDP growth in our model, we compute a Fisher price index between the first and last year for each country and use it to deflate the model-implied income measure. In a final step, we infer the production parameters B_j and γ_{oj} from firms' profit functions and first order conditions (cf. Equation 8). Appendix B provides more details on the numerical procedure and iterative algorithm.

[21] The calibration exactly matches sectoral wage bills in the data, because it exactly matches empirical employment shares and average wages in each occupation-sector oj when inferring A_{ojg} and w_{ojg}.

5 Gender Barrier Estimates and Model Validation

We now present our estimates of gender barriers, and show that the estimates correlate with empirical measures of gender norms and women's legal rights across countries and over time.

Gender Norms (λ_{oj}). Table 3 shows that gender norms have, on average, decreased across all sectors and occupations during the past decades (columns 1-2), especially in the service sector and in professional and clerical occupations. A comparison across countries in Panel C shows that gender norms have worsened in India, changed little in Indonesia, but improved substantially in the other countries in our sample. Brazil and Mexico exhibited very high gender norms in the 1970s, but have seen the largest improvements over time – consistent with rising female labor force participation in these countries. In the 2010s, gender norms show a clear negative correlation with economic development: they are highest in India, lower in Mexico, Brazil, and Indonesia, and significantly lower in Canada and the United States.

Wage Discrimination (τ_{oj}). Table 3 shows that average wage discrimination has declined across all sectors, from around 45% in the 1970s to around 30% in the 2010s (Columns 3-4, Panel A). Levels and changes of wage discrimination vary substantially across occupations but are relatively similar across sectors (Panel A-B), highlighting the importance of occupations to study gender barriers in the labor market. In the 1970s, wage discrimination ranged from 52% in crafts and trade occupations, to 36% in professional jobs, and 15% in clerical jobs. Unlike gender norms, wage discrimination has shown only modest improvements in professional and clerical occupations. A comparison across countries in Panel C shows that wage discrimination increased slightly in India, but declined substantially in all other countries, reaching average levels of 25-30% in the 2010s.[22] Unlike gender norms, wage discrimination does not exhibit a clear correlation with economic development. Additionally, improvements in one gender barrier do not necessarily imply progress in the other one. For example, while Indonesia reduced wage discrimination by 14 p.p., gender norms remained largely unchanged.[23]

[22]The worsening of gender barriers in India is consistent with a large literature that studies India's low and declining female labor force participation, along with the continued persistence of restrictive social norms, despite strong and rapid economic growth over the last decades. See recent work by Fletcher et al. (2017); Deshpande and Singh (2021); Deshpande and Kabeer (2024); Deshpande and Singh (2024).

[23]Figure OA.5 in the Online Appendix shows the evolution of gender barriers across countries and over time by plotting each country's estimated gender norms and wage discrimination by sector and

Table 3: Measures of Gender Norms and Female Wage Discrimination

	Gender Norms		Wage Discrimination	
	1970-75	2010-18	1970-75	2010-18
	(1)	(2)	(3)	(4)
Panel A. Sectors				
Agriculture	0.73	0.67	0.42	0.29
Manufacturing	0.73	0.66	0.48	0.32
Services	0.63	0.38	0.43	0.29
Panel B. Occupations				
Professional	0.73	0.40	0.36	0.31
Clerks	0.59	0.28	0.15	0.14
Craft, Trade, Service	0.66	0.54	0.52	0.34
Agricultural	0.73	0.69	0.42	0.29
Machine Op., Elementary	0.73	0.61	0.48	0.28
Panel C. Countries				
India	0.70	0.77	0.34	0.37
Indonesia	0.50	0.48	0.41	0.27
Brazil	0.87	0.47	0.46	0.33
Mexico	0.81	0.57	0.34	0.25
Canada	0.64	0.26	0.56	0.30
United States	0.50	0.16	0.53	0.32

Notes: This table reports the gender norm λ_{oj} and wage discrimination estimates in the 1970s and 2010s for the countries in our core sample. Panels A-C present average barriers by sector, occupation, and country. Averages are weighted by the total income in each occupation-sector and country-year.

Model Validation. As a validation, we examine whether our estimates of gender barriers – which are inferred as structural residuals – relate to measurable changes in social norms and women rights. To do so, we use the World Bank's "Women, Business, and the Law" (WBL) database, which evaluates 35 aspects of countries' legal code to construct indicators that measure gender equality in the labor market, at the workplace, and in the legal system over the past five decades and across 190 countries (World Bank, 2019; Hyland et al., 2020).

We merge the data sets at the country-year level and regress our estimates of wage discrimination $\hat{\tau}_{ojct}$ or gender norms $\hat{\lambda}_{ojct}$ on the different indicators from the WBL data for each country-year, controlling for the log of GDP per capita.

occupation against GDP per capita.

Table OA.1 in the Online Appendix shows a negative correlation between our wage discrimination estimates and measures of gender equality at the workplace, and between our gender norm estimates and WBL indicators related to household norms, which measure, for example, gender parity in heading a household and laws against domestic violence. Estimates of both gender barriers correlate negatively and significantly with the overall WBL index. These results indicate that our estimates capture meaningful information about measurable changes in underlying social and gender norms across countries and over time.

In terms of model fit, we find a strong correlation between our estimated gender barriers and their empirical targets in the data across occupations, sectors, countries and years, i.e., gender norms λ_{ojct} correlate with observed gender employment ratios; and wage discrimination τ_{ojct} with observed gender wage gaps (cf. Figure OA.6 in the Online Appendix).[24]

6 Effects of Changes in Gender Barriers

We now use our estimated model to quantify the share of countries' overall growth and structural transformation that can be attributed to changes in gender barriers over the past five decades. To do so, we simulate counterfactuals in which gender barriers $\{\tau_{oj}, \lambda_{oj}\}$ are held fixed at their calibrated values from each country's initial year, while other parameters – such as sector- and occupation-specific technological change $\{\gamma_{oj}$, $B_j\}$, gender-specific education distributions $\{\mu_{zg}$, $\sigma_{zg}\}$, and returns to human capital κ_{oj} – evolve according to their calibrated trajectories over time. We first fix both gender barriers at their initial levels and then gender norms and wage discrimination separately. We simulate counterfactual equilibrium paths for each country's employment structure, sectoral output, wages, and real GDP. To quantify the share of observed growth explained by changes in gender barriers, we compute $1 - \hat{g}_x/g_x$, where g_x is the observed growth rate in the data, and $\hat{g}_x$ is the counterfactual growth rate that we obtain when fixing gender barriers at their 1970s levels.

[24]Since our quantification approach directly targets gender employment and wage gaps, any difference between these targets and our estimates stems from additional terms in the estimating equations, namely: real income in the case of gender norms, and human capital gaps in the case of wage discrimination.

6.1 Effects on Sectoral Employment and Output

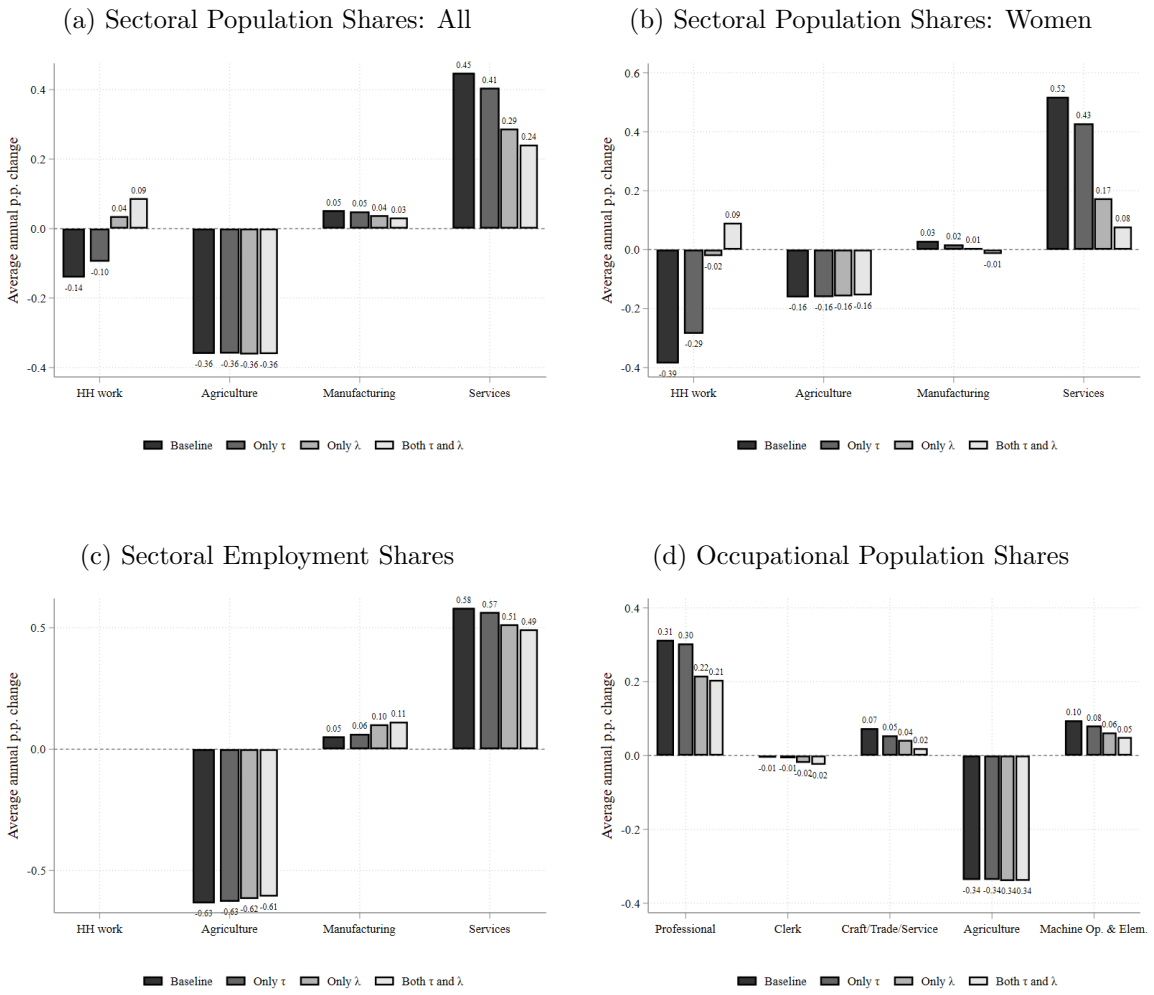
Employment Transitions across Sectors. Figure 4 shows the average effects of changes in gender barriers on sectoral employment shares across the six countries of our core sample. The black bars show the average annual percentage point (p.p.) change in sectoral employment shares that we observe between the 1970s and 2010s in the data. Population share in the home and agriculture sectors declined respectively by an average of 0.14 p.p. and 0.36 p.p. per year (Figure 4a). In contrast, population shares in the manufacturing and service sectors increased respectively by 0.05 p.p. and 0.45 p.p. per year. Figure 4b shows that the reallocation from the home sector to market services was primarily driven by women, while changes in the agriculture and manufacturing sectors were smaller for women relative to men.

These transitions would have looked significantly different if gender barriers had not changed since the 1970s. Without changes in wage discrimination and gender norms (lightest gray bar on the right), the population share in the home sector would have *increased* by 0.09 p.p. per year, which would have resulted in slower employment growth in manufacturing (0.03 p.p.) and services (0.24 p.p.), with minimal effects on the agriculture sector. It follows that changes in both gender barriers together account, on average, for approximately 40% of the observed changes in employment in manufacturing (1-0.03/0.05) and 45% in services. Hence, declining gender barriers significantly accelerated structural transformation in our sample countries by enabling women to transition from the home sector primarily into market services, and, to a lesser extent, into manufacturing, while leaving agricultural employment trends largely unchanged. These effects are primarily driven by changes in gender norms rather than changes in wage discrimination, as shown by the two darker gray bars in the middle.

The measures in Figure 4a represent the shares of the total working-age *population* in each sector, including the share that is inactive and works in the home sector. Instead, the literature typically focuses on sectoral employment shares *conditional* on labor force participation, which does not account for changes in labor force participation. Using this measure, Figure 4c shows that keeping gender barriers at their 1970s levels more than *doubles* the growth in manufacturing employment shares, while slowing the growth in services by 16%. These findings imply that changes in gender barriers and women's entry into the labor market over the past decades increased the absolute *number* of workers in manufacturing but reduced its *employment share* among active workers because the women who joined the labor market disproportionally entered the service sector. Hence,

our results indicate that falling gender barriers over the last decades contributed to the observed deindustrialization and rise of the service sector – a novel channel in addition to the ones documented in a recent literature (Rodrik, 2016; Fan et al., 2023; Huneeus and Rogerson, 2024).

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



Notes: This figure shows average annual percentage point changes in sectoral and occupational employment shares between the 1970s and 2010s for the six countries in our core sample. The top panel shows sectoral employment shares changes in the aggregate and for women, including the home sector share. Figure (c) instead measures sectoral employment shares *conditional* on labor force participation, excluding the home sector share. Figure (d) shows changes in occupational employment shares, whose measures include the home sector share (however, we omit the home sector bars from Figure (d) because they are identical to Figure (a)). The four bars in each graph represent: (i) the data; (ii) counterfactuals that fix wage discrimination at the 1970s level; (iii) counterfactuals that fix gender norms at the 1970s level; and (iv) counterfactuals that fix both gender barrier at the 1970s levels.

Employment Transitions across Occupations. Figure 4d documents how changes in gender barriers affected occupational employment shares over the past decades. In the data (black bars), rising female labor force participation and declining agricultural employment have coincided with employment growth in professional and managerial occupations (0.31 p.p. per year), machine-operating and elementary occupations (0.10 p.p.), and crafts, trade, and service occupations (0.07 p.p.). Overall, declining gender barriers account for a third of the observed growth in professional occupations, half for elementary occupations, and around 70% for craft, trade, and service occupations (lightest gray bars on the left). Clerical occupations – where women are overrepresented – would have experienced a steeper decrease if barriers had not changed. Effects are again primarily driven by changes in gender norms rather than changes in wage discrimination, as shown by the two darker gray bars in the middle.

Table 4: Change in Sectoral Output Explained by Gender Barriers

	Agriculture	Manufacturing	Services
	(1)	(2)	(3)
India	-5.90	-3.59	-2.94
Indonesia	1.81	4.45	14.98
Brazil	34.45	44.80	61.64
Mexico	12.59	23.59	53.46
Canada	11.41	32.52	44.72
USA	16.50	16.85	32.75
Average	11.81	19.77	34.10

Notes: This table reports the share of sectoral output growth, which is explained by changes in gender barriers between the 1970s and 2010s. To calculate this share, we compute $(1 - \hat{g}/g)$, where g represents the sectoral output growth from our baseline calibration, and $\hat{g}$ is the counterfactual output growth obtained when both gender barriers (τ_{ojct} and λ_{ojct}) are fixed at their calibrated values from the 1970s. The last row presents a simple average across the six countries in our core sample.

Effects on Sectoral Output. In addition to the employment transitions, changing gender barriers also affected sectoral output growth.[25] Table 4 shows that changes in gender norms and wage discrimination account, on average, for 12% of real output growth in agriculture, 20% in manufacturing, and 34% in services, with substantial variation

[25]Unlike employment shares, our calibration does not exactly match sectoral output in each country-year. Instead, the calibration targets nominal wage shares within each sector and adjusts sectoral prices to ensure market clearing in each sector. The reported effects on sectoral output therefore compare the values between our baseline calibration and counterfactuals.

across our sample countries. A notable outlier is India, where sectoral output would have been 3-6% higher had gender barriers not changed – which is consistent with the worsening of gender barriers shown in Table 3. Among the other countries, we find the smallest impact on sectoral output growth in Indonesia and the largest in Brazil. Across sectors, reductions in gender barriers have contributed the most to output growth in the service sector (15-60% across countries), followed by manufacturing (5-45%), and the least in agriculture (2-34%).

6.2 Effects on Aggregate Income and Productivity

Given the importance of changes in gender barriers for employment transitions and sectoral output, we now examine their contribution to aggregate earnings and GDP.

Effects on Output and Productivity. Table 5 shows that changes in gender barriers over the past decades explain, on average, 28% of the observed real GDP per capita growth in our sample countries. The effects vary considerably across countries, ranging from -5% in India to over 50% in Brazil.

Falling gender barriers increase GDP per capita through two channels: first, by increasing women's labor supply (extensive margin), and second, by improving the allocation of talent within the labor market as workers improve their sorting into occupations and sectors based on their comparative advantage. The combined effect of these channels on workers' productivity – i.e., output or GDP *per market worker* – is theoretically ambiguous because an improved talent allocation raises average human capital per worker, but higher labor force participation lowers it because the marginal worker who enters the labor market is negatively selected on ability. Quantitatively, we find that changes in gender barriers account for 3% of real GDP growth *per worker* (Column (2) of Table 5), indicating that the productivity gains from improved sorting marginally outweigh the negative selection effect from the extensive margin.

The effects are primarily driven by changes in gender norms, which account for an average of 24% of the observed real GDP per capita growth, while reductions in wage discrimination explain only 4% (see Table OA.2 in the Online Appendix). Changes in gender norms are more important because they declined most in occupations that offer high returns to ability, such as professional and clerical jobs, while wage discrimination has decreased more in occupations with lower returns to ability, as shown in Table 3.

Table 5: Share of GDP and Earnings Growth Due to Changes in Gender Barriers

	Real GDP		Real Earnings	
	Per Capita	Per Worker	Men	Women
	(1)	(2)	(3)	(4)
India	-4.93	-0.35	-2.91	-27.51
Indonesia	9.93	1.00	-5.07	56.47
Brazil	55.98	5.68	-15.36	96.81
Mexico	39.43	2.50	-20.96	85.14
Canada	41.45	2.39	-14.01	83.68
USA	27.81	6.00	-11.76	79.34
Average	28.28	2.87	-11.68	62.32

Notes: This table reports the share of real GDP and real earnings growth, which is explained by changes in gender barriers between the 1970s and 2010s. The shares are calculated as described in the notes of Table 4. The last row presents a simple average across the six countries of our core sample.

Effects on Earnings for Men and Women. Table 5 further shows that changes in gender barriers over the past decades reduced real wage growth by an average of 12% for men while increasing it by 62% for women (Columns 3-4). These results are driven by changes in the workforce composition and equilibrium effects: declining gender barriers increase women's labor force participation and enable them to enter the occupations and sectors where they were most productive, including high-return jobs. In equilibrium, these changes increase the competition for men in high-return jobs, causing some lower-ability men to shift to lower-return occupation-sectors or to exit the labor force altogether. Table OA.2 in the Online Appendix shows that changes in gender norms alone account for nearly the entire effect on real earnings growth for men. For women, both gender barriers matter because declining wage discrimination directly raises women's wage rates per human capital unit, in addition to affecting their employment choices.[26]

Summary. Our counterfactual analysis shows that changes in gender barriers have played an important role in driving structural transformation and GDP growth over the past decades in our sample countries, primarily by increasing female labor force participation and expanding employment in the service sector and in professional, trade, and service occupations. Comparing countries at different development levels, we find that changes in gender barriers had the smallest – or even negative – effects in the poorer

[26]The direct effects of lower wage discrimination matter for earnings but not for aggregate GDP because we attribute all discriminatory taxes on women to firms' profits. Hence, effects on aggregate GDP are only driven by changes in employment choices.

countries of our sample (India and Indonesia), and the largest effects in the middle-income countries (Brazil and Mexico). The richest countries (Canada and the United States) fall in between but still exhibit large effects.

7 Mechanisms and Robustness of Results

We next perform a series of analyses and robustness checks to study specific mechanisms of our findings and to examine the sensitivity of our quantitative results.

7.1 Non-Homothetic Preferences and Income Effects

We first assess the importance of non-homothetic preferences for our findings by simplifying our baseline model to homothetic CES preferences, setting $\varepsilon_j = 1 \forall j$ in Equation (1). With homothetic CES preferences, the expenditure shares in sectors $j = \{A, M, S\}$ simplify to $\varphi_j = \theta_j \times (p_j/P)^{1-\sigma}$, where P is the standard CES price index and sectoral expenditure shares no longer depend on individuals' income levels. We re-estimate all parameters and gender barriers in the homothetic model and implement the same counterfactuals.

Table 6 compares the results between the non-homothetic benchmark (column 1) and the homothetic model (column 2). We find that non-homothetic preferences amplify the effects of changing gender barriers, accounting for 28% of observed real GDP per capita growth, compared to 23% with homothetic preferences. In the non-homothetic model, changing gender barriers also explain a larger share of output growth in manufacturing (20% vs 16%) and services (34% vs 28%), while accounting for a similar share of agricultural output growth. Intuitively, declining gender barriers increase individuals' income levels. Under non-homothetic preferences, these income gains shift expenditure shares away from agriculture and toward manufacturing and services, which amplifies the employment and output gains in these sectors.

7.2 Preference versus Talent Shocks

We next analyze how our quantitative results change when workers receive idiosyncratic *talent* shocks ν_{oj}^{ti} across occupation-sector pairs (Roy, 1951), instead of the *preference* shocks ν_{oj}^{i} used in the baseline model. With preference shocks an inefficient allocation of talent can emerge because workers differ in absolute ability z and occupation-sector pairs differ in their returns to ability κ_{oj}. This model setup captures a *vertical* dimension of

Table 6: Importance of Non-Homothetic Preferences

	Benchmark Model	Homothetic Demand
	(1)	(2)
Real GDP per capita	28.28	22.73
Real GDP per worker	2.87	1.54
Agricultural Output	11.81	10.77
Manufacturing Output	19.77	15.92
Service Output	34.10	27.75

Notes: This table reports the share of real GDP and sectoral output growth, which is explained by changes in gender barriers between the 1970s and 2010s for the benchmark model with non-homothetic preferences (Column (1)) and a model with homothetic CES preferences (Colum (2)). The shares are calculated as described in the notes of Table 4. We average across the six countries in our core sample.

comparative advantage: high-ability z types have an absolute advantage in all occupation-sectors but a comparative advantage in occupation-sectors with higher returns to ability. Gender barriers can then distort the sorting of high-ability z workers into high-return occupation-sectors. I.i.d. talent shocks across occupation-sector pairs instead represents a *horizontal* dimension of comparative advantage. When combining absolute ability z and i.i.d. talent shocks, workers' comparative advantage in a given occupation-sector pair depends on the triplet $\{\nu_{oj}^{ti}, z^i, \kappa_{oj}\}$.

Using talent shocks, instead of preference shocks, adds complexity to the model with non-homothetic preferences because the talent shocks directly affect workers' income in each occupation-sector pair. Under non-homothetic preferences, income levels further affect individuals' sectoral expenditure shares, and therefore the relative weights that they assign to nominal earnings, sectoral prices, and occupational disutility A_{ojg}. These additional equilibrium interactions complicate the employment choice problem, which no longer admits the tractable, closed-form solution proposed in Equation (6). To avoid this complexity, we use a model with homothetic preferences to examine the effects of gender barriers when shocks across occupation-sector pairs are specified either as preference or talent shocks. For both models, we re-estimate all parameters and implement all counterfactuals. Section B.2 in the Online Appendix solves the model with talent shocks.

Table 7 shows that changes in gender barriers account for a similar share of the observed GDP per-capita growth with preference or talent shocks (Columns 1-2). The talent shock model, however, explains a larger share of GDP growth per worker (7.2% vs. 1.5%) because it provides two levers to improve the talent allocation of workers: first, by

Table 7: Impact of Preference vs Talent Shocks

	Both Barriers		Gender Norms		Wage Discrimination	
	Preference	Talent	Preference	Talent	Preference	Talent
	(1)	(2)	(3)	(4)	(5)	(6)
Real GDP per capita	22.73	21.78	14.10	-6.89	7.48	29.13
Real GDP per worker	1.54	7.17	2.82	-2.93	-1.20	9.39

Notes: This table reports the share of real GDP growth, which is explained by changes in gender barriers between the 1970s and 2010s, averaging across the six countries in our core sample. The shares are calculated as described in the notes of Table 4. The columns presents the growth contributions in models where workers receive either idiosyncratic *preference* or *talent* shocks across occupation-sector pairs. Columns (1) and (2) report contributions for both gender barriers, columns (3) and (4) for gender norms only, and columns (5) and (6) for wage discrimination only. Both models use homothetic CES preferences on the demand side.

improving the sorting of high-ability z types into high-return occupation-sectors; and second, by improving the sorting of workers into occupation-sectors based on their talent shocks ν_{oj}^{ti}. The second channel is absent in the model with preference shocks.

Table 7 further shows that changes in *gender norms* have larger effects on aggregate growth in the model with preference shocks (Columns 3-4), whereas changes in *wage discrimination* matter more in the model with talent shocks (Columns 5-6). To understand this result, it is useful to examine how each model links workers' occupation-sector choices to groups' average human capital across occupation-sector pairs. In the model with preference shocks, a group's average ability $\bar{z}_{oj}^{g}$ is decreasing in the group's employment share $\pi_{oj|gz}$ only in *high-return* occupation-sectors. This relation arises because the first entrants of a group into the high-return occupation have the highest ability z, leading to a decline in the group's average ability as additional workers enter. The opposite holds for low-return occupations. In the model with talent shocks, a group's average occupation-sector-specific talent $\bar{\nu}_{oj}^{tg}$ decreases in the group's employment share $\pi_{oj|g}$ which holds for all occupation-sectors.

As a result, a model with talent shocks infers a large decline in women's average human capital in *high-return* occupations where women's employment shares have increased over time (e.g., professional occupations), as it predicts reductions in both average ability $\bar{z}_{oj}^{g}$ and average occupation-sector-specific talent $\bar{\nu}_{oj}^{tg}$. A model with preference shocks instead captures only the decline in average ability $\bar{z}_{oj}^{g}$. To match the observed changes in gender employment and wage gaps, the model with talent shocks therefore infers a larger decline in female wage discrimination τ_{oj} in these occupations, while the model with preference

Table 8: Robustness to the Dispersion of Preference Shocks σ_ν

	$\sigma_\nu = 1.5$	$\sigma_\nu = 2$ (Benchmark)	$\sigma_\nu = 4$
	(1)	(2)	(3)
All Gender Barriers	30.20	28.28	22.53
Gender Norms Only	26.67	23.68	14.25
Wage Discrimination Only	3.49	4.39	7.12

Notes: This table reports the share of real GDP growth, which is explained by changes in gender barriers between the 1970s and 2010s, for different levels of dispersion in workers' occupation-sector-specific preference shocks σ_ϵ. The shares are calculated as described in the notes of Table 4. We average across the six countries in our core sample. A smaller value of σ_ϵ indicates higher dispersion.

shocks infers a larger decline in gender norms λ_{oj} over time.

7.3 Dispersion of Preference Shocks

We next examine how our results change with the dispersion of workers' preference shocks across occupation-sector pairs σ_ν. Smaller values of σ_ν indicate higher dispersion. Our baseline model uses $\sigma_\nu = 2$, following Hsieh et al. (2019) and Cassan et al. (2024). We now consider alternative values, ranging from 1.5 (high dispersion) to 4 (low dispersion).

Table 8 shows that a higher dispersion of preference shocks increases the share of economic growth that is explained by changes in gender barriers. Intuitively, the model infers higher levels of gender norms to explain observed occupational choices when idiosyncratic preferences are more dispersed across occupation-sectors, and vice-versa for less dispersion. The estimates of wage discrimination are more stable because they are calibrated to observed wage gaps, which are less affected by the dispersion of preference shocks than employment choices. Hence, Table 8 shows that the importance of changes in gender norms decreases when preference shocks are less dispersed, while the importance of wage discrimination increases.

7.4 Gender-Specific Comparative Advantage

In our benchmark model, we assume that men and women with the same ability z are equally productive in each occupation-sector. Under this assumption, we infer female wage discrimination τ_{oj} to match observed gender wage gaps in each occupation-sector using Equation 10. However, if men or women had a systematic comparative advantage in specific occupation-sector pairs, then our estimates of τ_{oj} would also reflect these

Table 9: Robustness to Adjustments for Gender-Specific Comparative Advantage

Model:	Benchmark	With Comparative Adv.
	(1)	(2)
All Gender Barriers	28.28	20.77
Gender Norms Only	23.68	17.28
Wage Discrimination Only	4.39	3.12

Notes: This table reports the share of real GDP growth, which is explained by changes in gender barriers between the 1970s and 2010s. The shares are calculated as described in the notes of Table 4. We average across the six countries in our core sample. The columns compare the results from our benchmark model (Column (1)) and from a model that adjusts for a gender-specific comparative advantage across occupations-sector pairs (Column (2)). We measure such comparative advantage from a benchmark economy – the United States in 2015 – which we assume to have no wage discrimination.

productivity effects. Measuring true productivity differences across genders in each occupation-sector pair and disentangling them from gender barriers is challenging (if not impossible), as it would require detailed data on labor input and output, separately for men and women. Such granular data does not exist for a large set of occupations, sectors, years, and countries. Instead, we implement two analyses to examine the potential role of gender-specific comparative advantage and its implications for our results.

First, we examine whether our estimates of τ_{ojct} show greater variation across occupation-sector pairs, or across country-years. This information provides insights into the underlying factors that τ_{ojct} reflects, as gender-specific productivity differences (unlike gender barriers) are expected to vary more across occupation-sector pairs and less across country-years.[27] Regressing our estimates of τ_{ojct} on occupation, sector, or occupation-sector fixed effects explains at most 20% of the variation in τ_{ojct} with R^2 values ranging from 0.01 to 0.2. This finding indicates that a substantial part of the variation in our τ_{ojct} estimates is driven by variation across country-years, making them more likely to reflect gender barriers than systematic gender productivity differences within occupation-sector pairs.

In a second exercise, we repeat our counterfactual analysis while explicitly accounting for a gender-specific comparative advantage in each occupation-sector pair. To do so, we select a benchmark economy, the United States in 2015, where we assume no wage discrimination (Lee, 2024). Under this assumption, any observed gender wage

[27]More specifically, it is reasonable to assume that the relative ranking of occupation-sectors in terms of a potential male (or female) comparative advantage is consistent across countries and over time. For example, agricultural jobs are likely to be consistently more brawn-intensive, which can benefit men relative to women, than office jobs.

gap in a given occupation-sector pair – after adjusting for differences in model-implied human capital – reflects actual differences in productivity between men and women. Hence, we estimate these gender-specific comparative advantages for each occupation-sector pair for our benchmark economy, $\delta_{oj,US2015}$, and apply the estimates to all other countries and years. Under this adjustment, the effective human capital units of a woman who works in occupation-sector oj in any country-year is then given by: $h^i_{ojf}(z) = (1 - \delta_{oj,US2015}) \times \exp(z^i \kappa_{oj})$. Table 9 shows that changes in gender barriers account for 21% of observed real GDP per capita growth under this adjustment, compared to 28% in the benchmark model. The moderate reduction also holds when we fix either gender norms (17% vs 24%) or wage discrimination (3% vs 4%). Hence, changes in gender barriers over the past decades remain an important channel for explaining economic growth in our sample, even when adjusting for a gender-specific comparative advantage.

8 Conclusion

This paper provides new evidence on the gender dimension of structural transformation across a large cross-section of countries over the past five decades. We find significant gender gaps in employment and wages that have narrowed on average over time, but continue to persist even today, and even in the most developed countries. We develop a general equilibrium occupational and sectoral choice model to quantify the effects of changes in gender barriers on macroeconomic outcomes, after accounting for the economic channels emphasized in the literature. We find that declining gender barriers have made substantial contributions to the observed increase in female labor force participation, the rise of the service sector, and real GDP per capita growth, but with substantial variation across countries.

By developing a unified framework to study *both* occupational and sectoral choices, our model not only extends the approaches in the literature (Hsieh et al., 2019), but is particularly useful to study employment transitions over long time horizons and across countries that span a wide spectrum of economic development, where employment structures vary significantly across sectors and occupations. Our analysis cannot identify the factors that led to a rapid decline in gender barriers in some countries but not in others. Future comparative research is therefore needed to examine the underlying drivers of these changes and to propose specific policies that can reduce gender barriers and promote gender parity in the labor market.

References

Albanesi, S., C. Olivetti, and B. Petrongolo (2023). Families, labor markets and policy. *Handbook of Economics of the Family*.

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

Bandiera, O., A. Elsayed, A. Heil, and A. Smurra (2022). Economic development and the organisation of labour: Evidence from the jobs of the world project. *Journal of the European Economic Association 20*(6), 2226–2270.

Bick, A., N. Fuchs-Schündeln, D. Lagakos, and H. Tsujiyama (2022). Structural change in labor supply and cross-country differences in hours worked. *Journal of Monetary Economics 130*, 68–85.

Boppart, T. (2014). Structural change and the kaldor facts in a growth model with relative price effects and non-gorman preferences. *Econometrica 82*(6), 2167–2196.

Boserup, E. (1975). The changing role of women in developing countries. *India International Centre Quarterly 2*(3), 199–203.

Bridgman, B., G. Duernecker, and B. Herrendorf (2018). Structural transformation, marketization, and household production around the world. *Journal of Development Economics 133*, 102–126.

Cassan, G., D. Keniston, and T. Kleineberg (2024). A division of laborers: Identity and efficiency in india. *Technical Manuscript*.

Caunedo, J., D. Jaume, and E. Keller (2023). Occupational exposure to capital-embodied technical change. *American Economic Review 113*(6), 1642–1685.

Chiplunkar, G. and P. K. Goldberg (2024). Aggregate implications of barriers to female entrepreneurship. *Econometrica 92*(6), 1801–1835.

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

Conte, B. (2022). Climate change and migration: The case of africa.

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

Deshpande, A. and N. Kabeer (2024). Norms that matter: Exploring the distribution of women's work between income generation, expenditure-saving and unpaid domestic responsibilities in india. *World Development 174*, 106435.

Deshpande, A. and J. Singh (2021). Dropping out, being pushed out or can't get in?

decoding declining labour force participation of indian women.

Deshpande, A. and J. Singh (2024). The demand-side story: Structural change and the decline in female labour force participation in india.

Evans, D. K., M. Akmal, and P. Jakiela (2021). Gender gaps in education: The long view. *IZA Journal of Development and Migration 12*(1).

Fan, T., M. Peters, and F. Zilibotti (2023). Growing like indiaâthe unequal effects of service-led growth. *Econometrica 91*(4), 1457–1494.

Feng, Y., J. Ren, and M. Rendall (2023). The reversal of the gender education gap with economic development. *Center for Economic Policy Research*.

Fletcher, E., R. Pande, and C. M. T. Moore (2017). Women and work in india: Descriptive evidence and a review of potential policies.

Goldin, C. (1995). The u-shaped female labor force function in economic development and economic history. *Investment in Womenâs Human Capital and Economic Development, ed. by T. P. Schultz, University of Chicago*, 61–90.

Goldin, C. (2024). Nobel lecture: An evolving economic force. *American Economic Review 114*(6), 1515–1539.

Gottlieb, C., C. Doss, D. Gollin, and M. Poschke (2023). Understanding the gender division of work across countries. Technical report.

Herrendorf, B., R. Rogerson, and A. Valentinyi (2013). Two perspectives on preferences and structural transformation. *American Economic Review 103*(7), 2752–89.

Herrendorf, B., R. Rogerson, and A. Valentinyi (2014). Growth and structural transformation. *Handbook of Economic Growth 2*, 855–941.

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

Hsieh, C.-T., E. Hurst, C. Jones, and P. Klenow (2019). The allocation of talent and U.S. economic growth. *Econometrica 87*(5), 1439–1474.

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

Hyland, M., S. Djankov, and P. K. Goldberg (2020). Gendered laws and women in the workforce. *American Economic Review: Insights 2*(4), 475–490.

IPUMS International (2020). 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*.

Jayachandran, S. (2021). Social norms as a barrier to womenâs employment in developing

countries. Technical report.

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

Lee, M. (2024). Allocation of female talent and cross-country productivity differences. *The Economic Journal 134*(664), 3333–3359.

Mammen, K. and C. Paxson (2000). Women's work and economic development. *Journal of Economic Perspectives 14*(4), 141–164.

Moro, A., S. Moslehi, and S. Tanaka (2017). Does home production drive structural transformation? *American Economic Journal Macroeconomics*, 116–146.

Ngai, L. R., C. Olivetti, and B. Petrongolo (2024). Gendered change: 150 years of transformation in us hours. Technical report, National Bureau of Economic Research.

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

Ngai, L. R. and C. A. Pissarides (2007). Structural change in a multisector model of growth. *American Economic Review 97*, 429–443.

Olivetti, C., J. Pan, and B. Petrongolo (2024). The evolution of gender in the labor market. *Handbook of Labor Economics*.

Porzio, T., F. Rossi, and G. Santangelo (2022, August). The human side of structural transformation. *American Economic Review 112*(8), 2774–2814.

Ranasinghe, A. (2024a). Gender specific distortions, entrepreneurship and misallocation. *Journal of Economic Dynamics and Control 162*, 104858.

Ranasinghe, A. (2024b). Misallocation across establishment gender. *Journal of Comparative Economics 52*(1), 183–206.

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

Rodrik, D. (2016). Premature deindustrialization. *Journal of Economic Growth 21*.

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

World Bank (2019). Women, business and the law 2019: A decade of reform. *Technical Report*.

APPENDIX

A THEORY

A.1 Theil Information H-Index and Decomposition

The Theil Information H-Index of segregation is based on Entropy measures which are defined as:

$$E = \sum_{g} \frac{N_g}{N} \ln\left(\frac{N}{N_g}\right),$$

where E measures the "diversity" of the population across gender groups g. N_g is the total number of workers of gender g and N is the total population. The entropy measure is equal to zero if all workers have the same gender, while it is maximized if workers are evenly distributed among genders. The entropy measure can also be calculated conditional on a given sector or occupation. Total gender segregation across all occupation-sector pairs is given by:

$$H_{oj} = \sum_{g} \sum_{o} \sum_{j} \frac{N_{oj}}{N \times E} \frac{N_{ojg}}{N_{oj}} \ln\left(\frac{N_{ojg}}{N_{oj}} \frac{N}{N_g}\right),$$

where N_{ojg} is the number of workers of gender g in occupation o and sector j. A larger H index implies that men and women are more segregated across occupations and sectors. The H-index is additively decomposable into gender segregation across sectors and segregation across occupations within sectors in the following way:

$$H_{oj} = H_j + \sum_{j} \sum_{o} \frac{N_j}{N} \frac{E_j}{E} H_o^j,$$

where H_j captures gender segregation across sectors and the second term captures gender segregation across occupations within sectors. Within-sector segregation depends on the relative size of each sector $(\frac{N_j}{N})$ and the relative entropy measures of each sector $\frac{E_j}{E}$. H_o^j captures the segregation across occupations within each sector which is equal to:

$$H_o^j = \sum_{g} \sum_{o} \frac{N_{oj}}{N_j \times E_j} \frac{N_{ojg}}{N_{oj}} \ln\left(\frac{N_{ojg}}{N_{oj}} \frac{N_j}{N_{gj}}\right),$$

where we simply treat the population in each sector as the universe. Similarly, gender segregation across sectors abstracts from any occupational differences and is given by:

$$H_j = \sum_g \sum_j \frac{N_j}{N \times E} \frac{N_{jg}}{N_j} \ln \left(\frac{N_{jg}}{N_j} \frac{N}{N_g} \right).$$

A.2 Deriving Sectoral Expenditure Shares

Individuals have non-homothetic preferences over sectors $k = \{A, M, S\}$ which are implicitly defined by the constraint:

$$\sum_k \theta_k^{\frac{1}{\sigma}} \left(\frac{C_k}{C^{\varepsilon_k}} \right)^{\frac{\sigma-1}{\sigma}} = 1, \tag{A.1}$$

which consumers maximize subject to the standard budget constraint: $\sum_k p_k C_k = I_{ojg}(z)$, where $I_{ojg}(z)$ is the income earned by a worker of gender g and ability z who works in oj. Taking the first-order conditions with a Lagrangian multiplier λ, we get:

$$\begin{aligned}
\left(\frac{\sigma-1}{\sigma} \right) \theta_k^{\frac{1}{\sigma}} C_k^{\frac{-1}{\sigma}} C^{\frac{-\varepsilon_k(\sigma-1)}{\sigma}} &= \lambda p_k \\
\Rightarrow \left(\frac{\sigma-1}{\sigma} \right) \underbrace{\sum_k \theta_k^{\frac{1}{\sigma}} C_k^{\frac{\sigma-1}{\sigma}} C^{\frac{-\varepsilon_k(\sigma-1)}{\sigma}}}_{=1} &= \lambda \underbrace{\sum_k p_k C_k}_{=I_{ojg}(z)} \\
\Rightarrow \lambda &= \left(\frac{\sigma-1}{\sigma} \right) \times \frac{1}{I_{ojg}(z)}.
\end{aligned}$$

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

$$\begin{aligned}
C_k &= \theta_k \left(\frac{p_k}{I_{ojg}(z)} \right)^{-\sigma} \left(\frac{I_{ojg}(z)}{P_{ojg}(z)} \right)^{\varepsilon_k(1-\sigma)} \\
&= \theta_k \left(\frac{p_k}{P_{ojg}(z)} \right)^{-\sigma} \left(\frac{I_{ojg}(z)}{P_{ojg}(z)} \right)^{\sigma+\varepsilon_k(1-\sigma)}.
\end{aligned}$$

Lastly, defining expenditure shares $\varphi_k \equiv p_k C_k / I_{ojg}(z)$, we get:

$$\varphi_k = \theta_k \left(\frac{p_k}{P_{ojg}(z)} \right)^{1-\sigma} \left(\frac{I_{ojg}(z)}{P_{ojg}(z)} \right)^{(\varepsilon_k-1)(1-\sigma)}.$$

B ESTIMATION ALGORITHM

We now describe the numerical procedure that estimates our model parameters. Specifically, we estimate sectoral preference shifters $\{\theta_j\}_{\forall j}$, productivity parameters $\{\{B_j\}_{\forall j}$ and $\{\gamma_{oj}\}_{\forall oj}\}$, and gender barriers $\{\tau_{oj}\}_{\forall oj}$ and $\{A_{ojg}\}_{\forall ojg}$. The estimation can be implemented separately for each country and consists of three nested loops.

Loop I: Guess overall income levels (i.e., nominal wages) in each country-year.

Loop II: Equation (2) shows that sectoral expenditure shares for $j = \{A, M, S\}$ depend on the joint term of sectoral preference weights and prices, so that we first solve for the combined term of $\psi_j := \theta_j p_j^{1-\sigma}$ for each country-year.

Loop III: Guess wage rates w_{ojg} and utility costs A_{ojg} for each ojg.

Step III.1: Compute income $I_{ojgz} = w_{ojg} \exp(z\kappa_{ojg})$, the non-homothetic-CES price index P_{ojgz} and hence indirect utility $V(I_{ojgz}, P_{ojgz})$ for each gender-ability type gz and each oj. Note that the price index P_{ojgz} of each income-type depends on the joint term $\psi_j := \theta_j p_j^{1-\sigma}$. Using the model-implied indirect utility values and the current guess of A_{ojg}, we then solve the occupational-and-sectoral choices for each gender-and-ability-type using Equation (6). Integrating the employment choice probabilities across ability-types, we can then solve for new values of utility costs A_{ojg}^{new}, which ensure that the model exactly matches the data on gender employment shares $\pi_{oj|g}$.

Step III.2: Next, we compute average human capital in each occupation-sector for each gender – which takes into account how workers select into occupation-sector pairs based on their comparative advantage, wage rates, and gender barriers:

$$\overline{H}_{ojg} \equiv \frac{H_{ojg}}{N_{ojg}} = \frac{\int_z \pi_{oj|gz} \exp(z\kappa_{ojg}) dF_g(z)}{\int_z \pi_{oj|gz} dF_g(z)}. \tag{B.1}$$

Using this measure of $\overline{H}_{ojg}$, we then solve for new values of wage rates at the occupation-sector-gender level w_{ojg}^{new}, which ensure that the model exactly matches the data on men and women's average wages in each occupation-sector pair by computing:

$$w_{ojg}^{new} = \frac{\overline{\text{wage}}_{ojg}}{\overline{H}_{ojg}}. \tag{B.2}$$

Step III.3: We iterate on the inner loop until the values of w_{ojg} and A_{ojg} converge. After convergence, female wage discrimination τ_{oj} is given by: $\tau_{oj} = 1 - w_{ojf}/w_{ojm}$.

Loop II: We then turn to the second loop to solve for the values of $\psi_j := \theta_j p_j^{(1-\sigma)}$.

Step II.1: Using Equation (2), we solve for the sectoral expenditure shares in $j = \{A, M, S\}$ for each income-type $ojgz$. We then aggregate these expenditure shares to the sector level, weighting them by the distribution of workers across $ojgz$ types and their respective income levels I_{ojgz}.[28] We then use the market clearing condition to solve for new values of ψ_j, which ensure that consumers' expenditure equals firms' revenue in each sector $j = \{A, M, S\}$ so that: $\psi_j^{new} = R_j/(E_j/\psi_j^{old})$, where $R_j = \sum_o w_{oj} H_{oj}$.

Step II.2: We iterate on this loop until the values of ψ_j converge.

Loop I: We next turn to the outer loop which solves for nominal income levels and sectoral prices by targeting each country's observed real GDP per capita growth rate.

Step I.1: To compute *real* income growth in the model, we first disentangle the joint term ψ_j into sectoral preference shifters θ_j and sectoral prices p_j. To do so, we assume that sectoral preference shifters can vary across countries but not over time, so that the *change* in ψ_j over time pins down the growth in sectoral prices for each country. Without loss of generality, we normalize prices for sectors $j = \{A, M, S\}$ to 1 in the first period of each country, pinning down the level of sectoral preference shifters θ_j for each country.

Step I.2: Given sectoral prices and consumers' expenditure for sectors $j = \{A, M, S\}$, we then use Equation (5) to decompose the price of the aggregate service composite into the respective prices of home and market services.[29]

Step I.3: Using the sectoral prices and quantities, we then compute the Fisher price index across the three market sectors $j = \{A, M, ms\}$ for each country. To compute real income growth in our model, we then deflate total market-based income (i.e., workers' wage bill or firms' revenue in the sectors $j = \{A, M, ms\}$) with the Fisher price index. We then adjust the nominal wage level in the last year of each country until the model-implied real income growth rate matches the real GDP per capita growth rates that we observe in the data.

Step I.4: We iterate until the estimates of all three loops converge.

[28] We assume that the profits – which are earned by paying women below their marginal product – are spent across sectors $j = \{A, M, S\}$ in the same proportions as the economy-wide wage bill.

[29] This step further uses the market clearing condition for home and market services which requires that total expenditure equals total revenues, so that $\varphi_k \times I = R_k = \sum_k w_{ok} H_{ok}$ in each sector $k = \{hs, ms\}$.

Online Appendix Material for

Gender Barriers, Structural Transformation, and Economic Development

by Gaurav Chiplunkar and Tatjana Kleineberg

2025

FOR ONLINE PUBLICATION

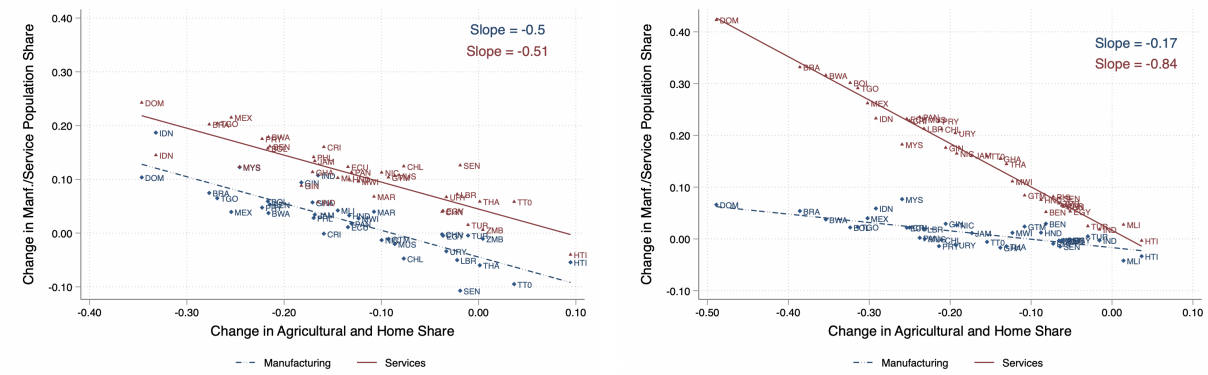
A TABLES AND FIGURES

A.1 Descriptive Patterns

Figure OA.1: Manufacturing and Service Growth in Developing Countries

(a) Men

(b) Women

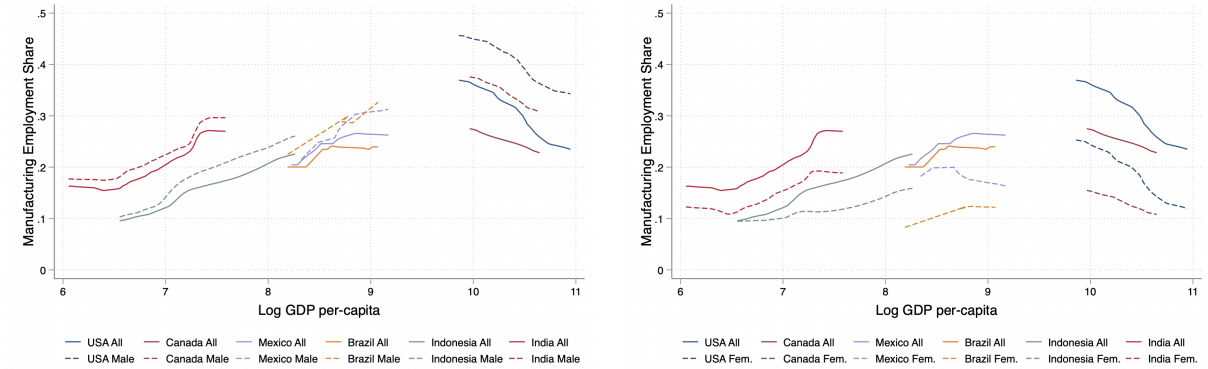


Notes: This figure shows how changes in the combined employment shares in agriculture and the home sector (horizontal axis) correlate with employment changes in manufacturing (dash blue line) and services (solid red line), separately for men and women. The sample is restricted to low- and middle-income countries. For each country, we calculate changes in sectoral employment shares over time, conditional on labor force participation, using the first (1970-1990) and last (2000-2019) years of available data. We report the coefficients from a linear regression model for manufacturing (Slope (M)) and services (Slope (S)) in the upper right corner of each graph.

Figure OA.2: Manufacturing Employment Shares: Aggregate and by Gender

(a) Aggregate and Male Shares

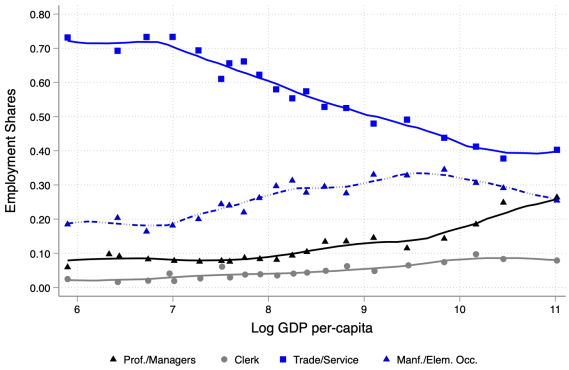
(b) Aggregate and Female Shares



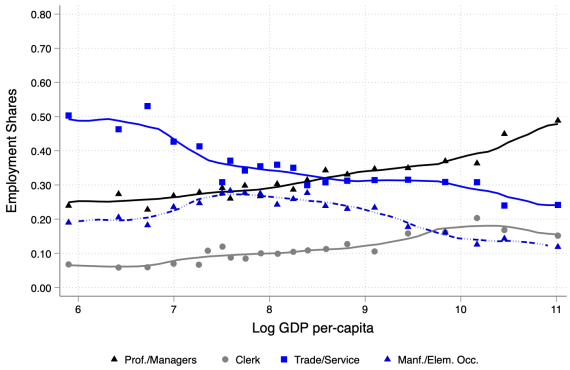
Notes: This figure plots manufacturing employment shares against log real GDP per-capita in 2010 US dollars over time for the six countries of our core sample: India, Indonesia, Brazil, Mexico, Canada, and the United States. The solid lines show changes in aggregate employment shares. The dashed lines show changes in male (left figure) and female (right figure) employment shares.

Figure OA.3: Occupational Transitions Within Manufacturing and Services

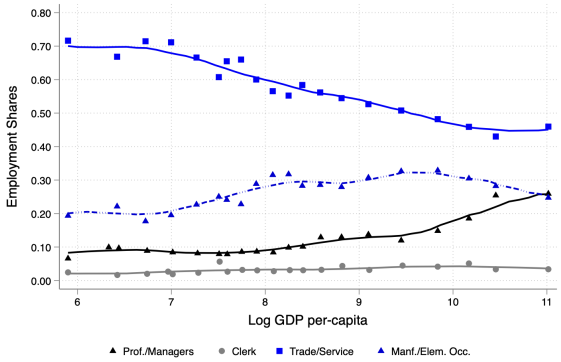
(a) Manufacturing: All



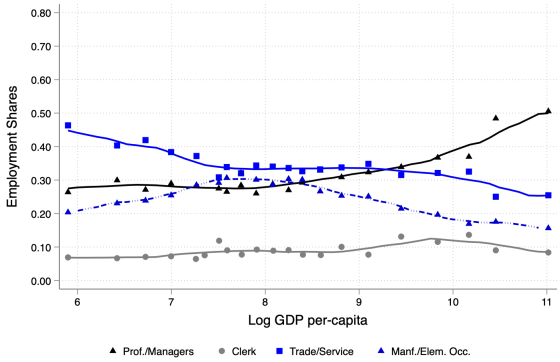
(b) Services: All



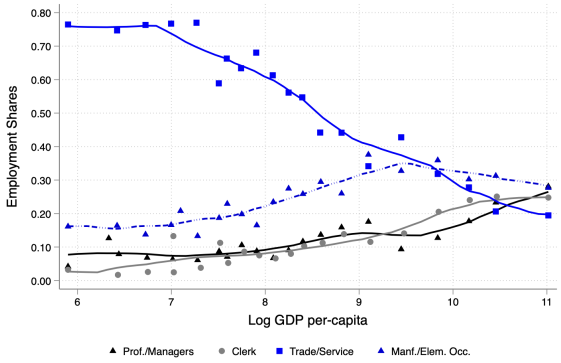
(c) Manufacturing: Men



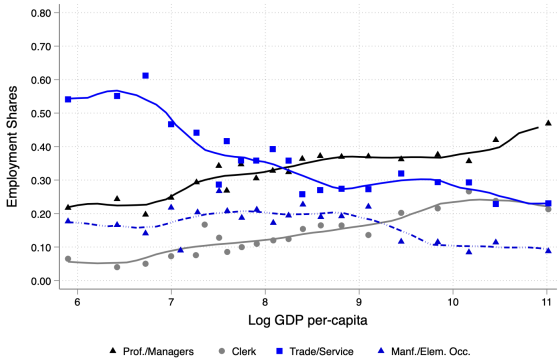
(d) Services: Men



(e) Manufacturing: Women



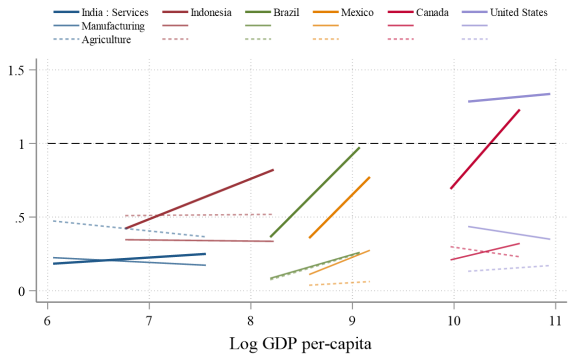
(f) Services: Women



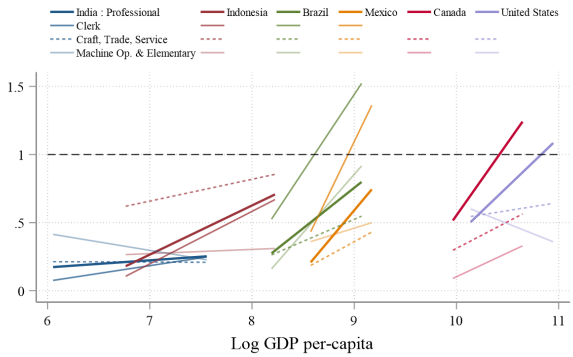
Notes: This figure presents non-parametric fits of occupational employment shares within manufacturing (left panel) and within the service sector (right panel) against log real GDP per-capita in 2010 US dollars – aggregate, and separately for men and women. The sample pools data from 91 countries and 305 country-years between 1960-2019.

Figure OA.4: Gender Employment and Wage Ratios over Time

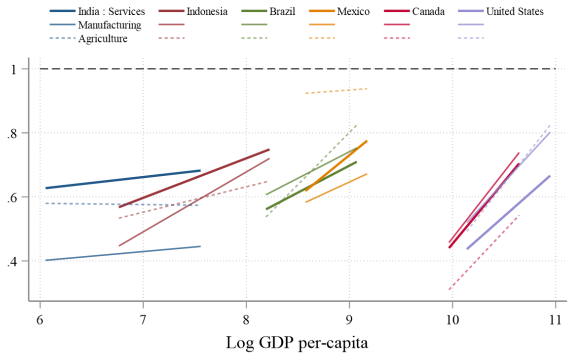
(a) Employment Ratio by Sector



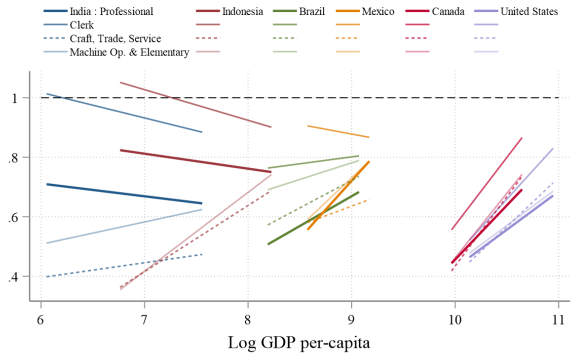
(b) Employment Ratio by Occupation



(c) Wage Ratio by Sector



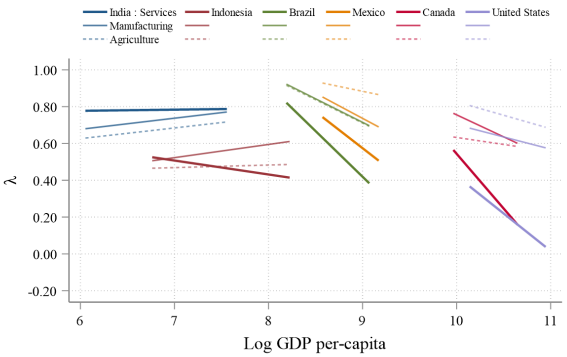
(d) Wage Ratio by Occupation



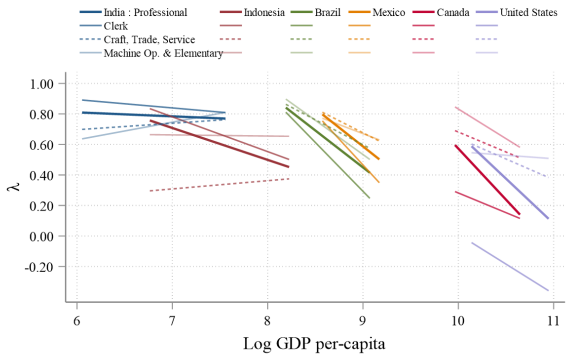
Notes: This figure plots the female-to-male employment and wage ratios for countries in our core sample against the log of real GDP per-capita in 2010 US dollars. The time period covers the first (1970s) and last (2010s) survey year of each country. Hence, the length of the graphs on the x-axis reflects how fast countries grew during this time period. Employment ratios divide the share of women working in an occupation-sector by the share of men. Wage ratios divide the average wage of women in an occupation-sector by the average wage of men. Figures (a) and (c) show these ratios by sector, while Figures (b) and (d) show them by occupation. Figure (b) excludes the clerk occupation for the United States and Canada as their employment ratios exceed 2 which makes the graph hard to read.

Figure OA.5: Gender Norms and Wage Discrimination Across Countries Over Time

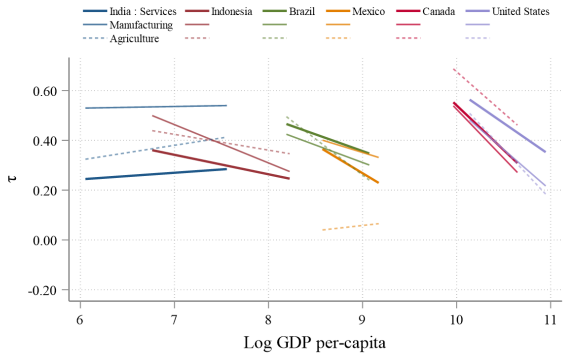
(a) Gender Norms (λ_{oj}) by Sector



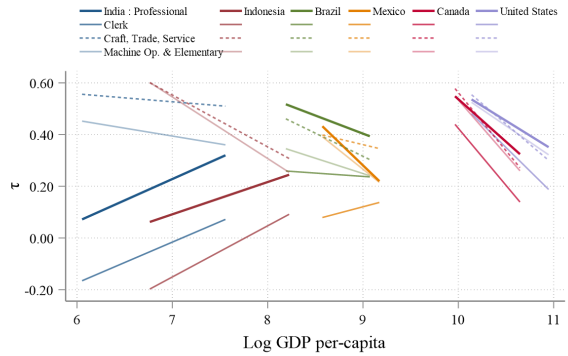
(b) Gender Norms (λ_{oj}) by Occupation



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



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



Notes: This figure plots the estimated gender norms λ_{oj} and wage discrimination τ_{oj} for our core-sample countries against the log of real GDP per-capita in 2010 US dollars. The time period covers the first (1970s) and last (2010s) survey year of each country. Hence, the length of the graphs on the x-axis reflects how fast countries grew during this time period. Figures (a) and (b) show gender norms by sector and occupation, and Figures (c) and (d) show wage discrimination by sector and occupation.

A.2 Model Fit and Model Validation

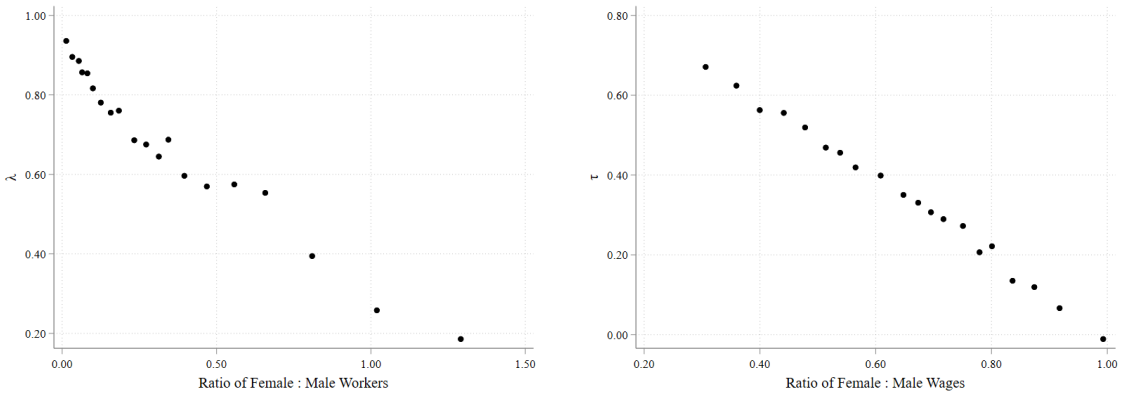
Table OA.1: Gender barriers and World, Business, and the Law Indicators

	Coef.	S.E.	p-value
Panel A. Correlation with τ			
Overall WBL Index	-0.08	(0.02)	0.00***
Index of Workplace Equality	-0.07	(0.02)	0.00***
Can a woman get a job in the same way as a man?	0.03	(0.04)	0.53
Does the law prohibit discrimination in empl. based on gender?	-0.12	(0.04)	0.01***
Is there legislation on sexual harassment in employment?	-0.17	(0.03)	0.00***
Are there criminal penalties for harrassment at workplace?	-0.17	(0.03)	0.00***
Panel B. Correlation with λ			
Overall WBL Index	-0.16	(0.02)	0.00***
Index of Household Norms	-0.11	(0.02)	0.00***
Can a woman be head of household in the same way as a man?	-0.04	(0.05)	0.45
Is there legislation specifically addressing domestic violence?	-0.17	(0.05)	0.00***
Can a woman obtain a divorce in the same way as a man?	-0.12	(0.05)	0.01**
Does a woman have the same rights to remarry as a man?	-0.12	(0.05)	0.01**

Notes: This table shows the OLS correlation of wage discrimination τ_{ojct} in Panel A and gender norms λ_{ojct} in Panel B with indicators of the "World, Business, and the Law" data as described in Section 5. Standard errors are clustered at the country level. *** is p<0.01, ** is p<0.05 and * is p<0.1.

Figure OA.6: Correlations of Gender Barriers and their Empirical Targets

(a) Gender Norms and Gender Employment Gaps (b) Wage Discrimination nd Gender Wage Gaps



Notes: This figure shows a binned scatter plot of the correlation between our estimates of gender barriers and their empirical targets in the data. Figure (a) shows the correlation between gender norms λ_{ojct} and observed employment gaps, Figure (b) between wage discrimination τ_{ojct} and observed wage gaps. The sample pools all occupation-sectors and country-years.

A.3 Importance of Gender Barriers and Aggregate Growth

Table OA.2: Change in Aggregate Outcomes Explained by Gender Barriers

	Real GDP		Real Earnings	
	Per Capita	Per Worker	Men	Women
	(1)	(2)	(3)	(4)
Panel A. Gender Norms				
India	-3.78	0.37	-3.11	-1.28
Indonesia	3.84	1.24	-6.96	21.49
Brazil	53.34	5.38	-14.39	87.41
Mexico	34.51	3.75	-20.07	66.01
Canada	31.97	2.95	-12.13	46.05
United States	22.21	8.37	-12.29	42.78
Average	23.68	3.68	-11.49	43.74
Panel B. Wage Discrimination				
India	-2.11	-0.72	0.44	-43.57
Indonesia	3.93	-0.27	-0.66	19.90
Brazil	5.09	-0.87	-0.80	29.64
Mexico	6.76	0.14	-3.80	33.71
Canada	8.01	-0.88	-2.89	52.47
USA	4.64	-1.58	-1.28	49.00
Average	4.39	-0.70	-1.50	23.52

Notes: This table reports the share of real GDP and real earnings growth, which is explained by changes in gender barriers between the 1970s and 2010s. To calculate this share, we compute $(1 - \hat{g}/g)$, where g represents the growth rate observed in the data and $\hat{g}$ is the counterfactual growth rate obtained when gender barriers are fixed at their calibrate values from the 1970s. The panels report the separate effects for changes in either gender norms (Panel A) or wage discrimination (Panel B). The last row reports a simple average across the six countries in our core sample.

B THEORY

B.1 CES Preferences over Home and Market Services

For a sector $k \in hs, ms$, i.e., home and market services, consumers minimize:

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

For a CES price index of the aggregate service composite $P_S = \left[\sum_k \alpha_k p_k^{1-\sigma_s} \right]^{\frac{1}{1-\sigma_s}}$ and a Lagrange multiplier λ, the first-order condition is given by:

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

Lastly, substituting back in the constraint, we have:

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

B.2 Model with Homothetic CES Preferences

B.2.1 Indirect utility, expenditure, and occupation-sector choices with homothetic CES preferences

The non-homothetic CES preferences simplify to homothetic ones when setting $\epsilon_k = 1 \forall k$. In this case, the price index P no longer depends on individuals' income and Equation (3) simplifies to the standard CES price index:

$$P = \left[\sum_k \theta_k p_k^{1-\sigma} \right]^{\frac{1}{1-\sigma}}. \tag{OA.1}$$

We can then write the indirect utility from Section 3.2 by taking the exponential as follows:

$$U_{ojg}(z) = A_{ojg} \times \underbrace{\frac{w_{ojg} e^{z\kappa_{ojg}} \times \nu_{oj}^i}{P}}_{=V(I(z),P)}. \tag{OA.2}$$

Assuming that ν_{oj}^i is Frechet distributed, the solution to the occupational choice problem is given by:

$$\pi_{oj|gz} = \frac{\left[A_{ojg} w_{ojg} e^{z\kappa_{ojg}} \right]^{\sigma_\nu}}{\sum_{j'} \sum_{o'} \left[A_{o'j'g} w_{o'j'g} e^{z\kappa_{o'j'g}} \right]^{\sigma_\nu}}, \tag{OA.3}$$

where the CES price index P cancels because – in contrast to the non-homothetic CES – it no longer depends on individual income and therefore no longer varies across occupational and sectoral employment choices.

B.2.2 Homothetic Preferences and Talent Shocks across Occupation-Sector Pairs

With homothetic preferences, the model remains tractable when workers receive either preference ν_{oj}^i or talent shocks ν_{oj}^{ti} across occupation-sector pairs. With preference shocks, individuals' income varies only across $ojgz$ cells and $I_{ojg}(z) = w_{ojg} \exp(z\kappa_{ojg})$. With talent shocks ν_{oj}^{ti}, each worker's human capital and income in a given occupation-sector pair further depends on their idiosyncratic talent draw. Solving for the human capital supply in each occupation-sector pair in the talent-shock model therefore requires several steps. First, we use the properties of the Frechet distribution to solve for the average

talent levels of workers with gender g and ability z conditional on having chosen an occupation-sector pair oj, which we denote by $\bar{\nu}_{oj}^{tgz}$, and which is equal to:

$$\bar{\nu}_{oj}^{tgz} = \sigma_\nu \left(\frac{1}{\pi_{oj|gz}} \right)^{\frac{1}{\sigma_\nu}} \Gamma\left(1 - \frac{1}{\sigma_\nu}\right), \tag{OA.4}$$

where $\Gamma(.)$ is the gamma function. This equation illustrates the negative selection result that is standard in Roy models with job-specific talent shocks: for each gender-ability type, the average occupation-sector-specific talent is decreasing in the group's employment share $\pi_{oj|gz}$ because additional entrants into the occupation-sector are negatively selected based on their talent draw ν_{oj}^{ti}.

The average human capital units of workers of gender g and ability z who work in occupation-sector oj then combines the average occupation-sector-specific talent and the ability z so that:

$$\overline{H}_{ojgz} = e^{z\kappa_{ojg}} \times \bar{\nu}_{oj}^{tgz}. \tag{OA.5}$$

Lastly, to solve for the aggregate human capital supply in each occupation-sector, we integrate over ability-z-types and sum across genders so that:

$$H_{oj} = \sum_g N_g \int_z \pi_{oj|gz} \overline{H}_{ojgz} dF(z). \tag{OA.6}$$

C DATA AND MEASUREMENT

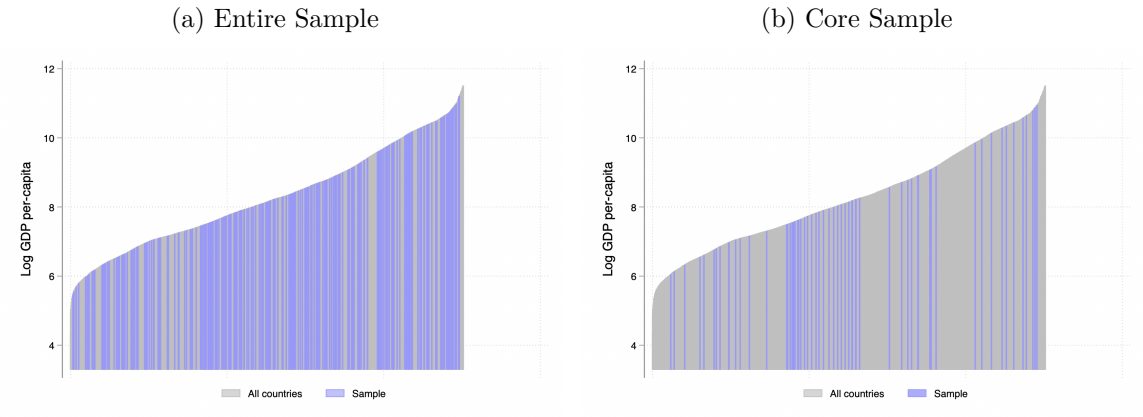
C.1 Sample Coverage

Table OA.3: Coverage of All Countries across Decades

Decade	# Country-Years	Percentage
1960-69	13	4.26
1970-79	31	10.16
1980-89	46	15.08
1990-99	62	20.33
2000-09	93	30.49
2010-19	60	19.67
Total	305	100

Notes: This table reports the coverage of country-years across all 91 countries and 305 country-years in our data by decade.

Figure OA.7: Sample Coverage across Countries



Notes: This figure sorts all country-years by their real GDP per-capita in 2010 US dollars and highlights the data coverage across all countries in our sample (Figure a) and across the countries in our core sample (Figure b) in blue.

Table OA.4: Coverage of Countries in the Core Sample

Country	Time Coverage	Real GDP p.c. in 2010
	(1)	(2)
India	1983 to 2018	$1,357
Indonesia	1976 to 2018	$866
Mexico	1970 to 2015	$9,271
Brazil	1970 to 2010	$11,286
Canada	1971 to 2011	$48,464
United States	1970 to 2015	$48,467

Notes: This table reports the coverage of our data for the six countries in our core sample. Column (1) reports the time coverage while Column (2) reports the real GDP per-capita in 2010 US dollars from the World Bank data.

C.2 Sample Definition, Sector-Occupation Classifications, and Data Cleaning

C.2.1 Sample Definition

1. We restrict the sample to individuals of ages 25 to 60 years.

2. For a very small share of observations, we observe a sector and occupation of workers' current job, but their employment status is missing. We count these individuals as "employed".

3. We drop individuals who are in the armed forces. Among the inactive population, we drop individuals who report that they do not work because they are in school, in prison, disabled, retired, or living on rents. For years where this information is not available, we assume that the fraction of inactive individuals for which this criterion applies is the same as in the closest survey-year.

4. We classify all individuals who are unemployed or out of the workforce as working in the "home sector."

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

C.2.2 Industry and Occupation Classifications

1. Tables OA.6 and OA.7 provide the classification of sectors and occupations.

Table OA.5: 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

2. For the small share of employed workers for which we observe only either an occupation code or a sector code, we make the following adjustments:

    - We set the occupation code to "Skilled Agricultural Workers" if the reported sector is agriculture, but the occupation code is missing, unknown, or unspecified.
    - We set the sector to "Agriculture" if the sector code is missing and the reported occupation is "Skilled Agricultural Workers".
    - We set the sector to "Manufacturing" if the sector code is missing and the reported occupation is either "Craft and Trades Workers" or "Plant and Machine Operators".
    - We set the sector to "Services" if the sector code is missing and the reported occupation is "Professionals," "Clerks," "Service workers," or "Elementary workers".
    - We drop all other employed individuals for which the occupation code is missing, unknown, or unspecified.

3. We group some occupations and sectors to avoid noisy measurements in occupation-

sector-gender cells that have only a very small sample size in some country years.

- We group the first three ISCO occupation codes as shown in Table OA.7 as one category of "Professionals".
- We re-assign workers that work as "Professionals," "Clerks," or "Service Workers" in the "Agriculture" sector to the "Services" sector.
- We re-assign workers that work as "Crafts and Trade workers" or "Plant and Machine Operators" in the "Agriculture" sector to the "Manufacturing" sector.
- We re-assign workers that work as "Skilled Agricultural workers" in the "Manufacturing" or "Service" sector to the "Agriculture" sector.

Table OA.6: Classification of Industry Codes in IPUMS

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

Table OA.7: Classification of Occupations

ISCO	Occupation	Classification	Sector
1	Legislators, senior officials and managers	Professional (1)	M,S
2	Professionals	Professional (1)	M,S
3	Technicians and associate professionals	Professional (1)	M,S
4	Clerks, secretaries, librarians, cashiers	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 trades workers (builders, textile)	Craft/Trade (5)	M,S
8	Plant and machine operators and assemblers	Plant & Machine (6)	M,S
9	Elementary occ (street vendors, manual labor)	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:*This table shows the classification of occupations as reported in the IPUMS data. We aggregate the first three ISCO codes as shown in Column (3). The last column shows which occupations are represented in the respective sectors. More information on the ISCO classification of occupations can be found here.

C.3 Measures of Earnings, Hours Worked, and Effective Employment

In this section, we first discuss the available information on hours worked and how we compute measures of "effective employment" in each occupation-sector-gender cell. We then describe how income is measured in IPUMS and the labor force survey data and how we compute our measures of hourly earnings.

C.3.1 Measurement of Hours Worked

Availability and measurement of hours worked in IPUMS vary across country-years. The most commonly available measures are:

1. **HRSWORK1:** reports the person's total hours worked per week at all jobs.
2. **HRSUSUAL1:** reports the person's usual number of hours worked in a typical week across all jobs.

3. **HRSACTUAL1:** reports the person's actual number of hours worked per week at all jobs.

For some country-years, only a categorized equivalent of the variable is available (marked with the suffix 2, e.g.: HRSWORK2). When the numerical variable is not available, we use the categorical variable by setting hours to the midpoint of each category.

1. We use the variable HRSWORK1 or the midpoints of HRSWORK2 categories when possible. If these variables are not available, we use either usual or actual hours worked per week.

2. We set observations to missing if reported hours are negative or exceed 100 hours per week.

3. When hours worked are missing for an employed individual, but are available in general for the country-year sample, then we replace the missing observations by the average hours worked in the same gender-sector-occupation cell within that country-year.

4. If no information on hours worked is available for a specific country year, we use the average hours worked of each gender-sector-occupation cell from the closest year of the country for which this variable is available. This adjustment applies to the samples of Brazil (1970), Mexico (1970,2005,2015), and Indonesia (2000, 2005, 2010).

5. To convert weekly hours worked into monthly or annual measures, we assume that individuals work for 4.3 weeks/month and for 12 months/year. In some cases (United States and Canada) we observe the number of months worked, which we then use to calculate annual hours worked.

C.3.2 Measurement of Effective Employment

Our quantitative exercise focuses on the extensive margin of employment choices. To capture differences in the intensive margin between part- or full-time workers, we make the following adjustments in the individual-level data to measure effective employment in each occupation-sector-gender cell in each country year:

- Workers who work less than 15 hours are assigned two-thirds to the "Home sector" and one-third to the reported occupation-sector.

- Workers who work between 15 and 25 hours are assigned one-half to the "Home sector" and one-half to the reported occupation-sector.
- Workers who work between 25 and 35 hours are assigned one-third to the "Home sector" and two-thirds to the reported occupation-sector.
- Workers who work more than 35 hours are fully assigned to the reported occupation-sector and unemployed and inactive individuals are fully assigned to the "Home sector".
- For country-years in which hours worked is not available in the individual-level data, we assume that the ratios between workers and "effective workers" in each occupation-sector-gender cell are the same as in the closest survey year in which this information is available. The relevant country-years for which we need to make this adjustment are Brazil 1970 (for which we use the ratios from Brazil 1980), Mexico 1970 (for which we use the ratios from Mexico 1990) and Mexico 2015 (for which we use the ratios from Mexico 2010).

C.3.3 Measuring Income

IPUMS provides three main income variables:

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 (wages, business, or farm) in the previous month or year.
3. INCWAGE: reports the person's weekly, monthly or annual wage or salary income.

The labor force surveys provide the equivalent of INCWAGE. Each country-year differs in the availability of these variables. When possible, we try to use the same income measure for a given country over time. More specifically, countries have the following availability:

1. Brazil: 1970, 1980, 1991, 2000, 2010
    - INCTOT for the previous month is available and used for all sample years.
    - Hours Worked per Week: HRSWORK1 available for 1991, 2000 and 2010. HRSWORK2 available for 1980. For 1970 we use the gender-sector-occupation average from the 1980 data.

- Earnings per hour = INCTOT/(4.33*Hrs Worked per Week).

2. Canada: 1971, 1981, 1991, 2001, 2011

   - INCTOT for the previous year is available and used for all sample years.
   - Hours Worked per Week: HRSWORK1 available 1981 onwards, HRSUSUAL2 available for 1971.
   - Months Worked per Year: MONTHSWRK is available and used to convert annual income to hourly income.
   - Earnings per hour = INCTOT/(4.33*Hrs Worked per Week*Months Worked).

3. India: 1983, 1987, 1993, 1999, 2004, 2009, 2011, 2018

   - There is no information on hours worked provided in IPUMS, so we instead rely on data from the labor force surveys, in particular the Indian Employment-Unemployment Survey (EUS) from 1983-2009 and the Periodic Labor Force Survey (PLFS) for 2011-2018. Both surveys have been harmonized in the World Bank's Global Labor Database (GLD).
   - Income is measured as "Last wage payment, primary job, excl. bonuses, etc. (7-day ref period)". An additional variable records the frequency of payment for each survey.
   - Hours Worked per Week: The variable records the hours of work last week for the individual's main job.
   - Wages per hour: We convert weekly wages to hourly wages using hours worked per week.

4. Indonesia: 1976, 1995, 2015, 2018

   - Income: In IPUMS, the only available income measure is INCWAGE of the previous month, which is available only for the years 1976 and 1995.
   - Hours worked: HRSWORKED1 is available for 1995 and HRSACTUAL1 is available for 1976.
   - For 1976 and 1995, we compute wages per hour = INCWAGE/(4.333*Hrs Worked per Week).
   - For the time period from 1994 to 2019, we complement the IPUMS data for Indonesia with data from the SAKERNAS labor force survey.

- Wage income is available in the SAKERNAS survey for most years and is defined as "Last wage payment, primary job, excl. bonuses, etc. (7-day ref period)".
- Hours worked in the last week is available in the SAKERNAS survey.
- For 1994-2018, we compute hourly wages by dividing weekly wage income by weekly hours worked.

5. Mexico: 1970, 1990, 1995, 2000, 2005, 2010, 2015

   - INCTOT of the previous month is available in 1970, 1995, and 2000.
   - INCEARN of the previous month is available in 1990, 1995, 2000, and 2015.
   - We use the variable INCTOT in 1970 and INCEARN in 2015.
   - Hours worked: HRSWORKED1 is available for 1990, 1995, 2000, and 2010. For 1970 we use the gender-sector-occupation averages of the 1990 survey. For 2015 we use the gender-sector-occupation average of the 2010 survey.
   - Earnings per hour = INCTOT or INCEARN / (4.33*Hrs Worked per Week)

6. United States: 1970, 1980, 1990, 2000, 2005, 2010, 2015

   - INCTOT in the previous year is available and used in all sample years.
   - Hours worked: HRSWORK1 is available from 1980 onwards. HRSWORK2 is available for 1970. MONTHSWRK is available for all years.
   - Earnings per hour = INCTOT/(4.33*Hrs Worked per Week*Months Worked).

C.3.4 Measuring Income: Data Cleaning and Aggregation

We first clean the data on hourly wages and then compute men and women's aggregate and average wages for all occupation-sector pairs. We do the cleaning separately for each country year.

We trim hourly wages at the 99th percentile for each occupation-sector-gender cell. If workers' income is missing, we impute it by estimating the following Mincer regression:

$$\ln(\text{wage}_i) = \alpha * \text{YrsSchool}_i + \beta * \text{Exp}_i + \gamma * \text{Exp}_i^2 + u_{ojg} + \epsilon_i, \tag{OA.7}$$

which regresses the log of hourly wages of individuals i on their years of schooling, experience, experience squared, and gender, sector, and occupation fixed effects. We

then use the predicted values from the Mincer regression to impute hourly income for all employed individuals for which income data is missing.

We compute annual total earnings for each individual by multiplying reported monthly income times 12 (Brazil, Mexico, Indonesia IPUMS), reported weekly income times 52 (India and Indonesia labor force surveys) or by simply using reported annual income (Canada and US IPUMS). We then aggregate income across individuals to the occupation-sector-and-gender level after adjusting individuals' survey weights for part-time employment as described above. Average earnings in each occupation-sector-gender cell is then equal to annual aggregate earnings divided by the number of "effective workers" | which again adjusts for part-time workers as described above.

Earnings in Home Sector: For each country-year, we set earnings in the Home Sector equal to the measured earnings of women who work in "elementary occupations" in the service sector.