From Rural Fields to Urban Kitchens: Structural Change and the Decline of Women's Work in India*

Michael Peters†, Pamela Torola‡, Lindsey Uniat§, and Fabrizio Zilibotti¶

December 2024

Abstract

India's GDP per capita grew threefold between 1987 and 2019, coinciding with rapid urbanization. During this period, female labor force participation (FLFP) declined significantly. Consistent with this observation, we document a pronounced urban-rural participation gap, where FLFP is higher in poorer, rural labor markets. Using time-use data, we show that this is primarily driven by an extensive margin: in rural districts, women often engage in part-time activities, typically related to agriculture and informal family businesses. These activities are less common in urban areas, where some women take formal jobs, but a larger share withdraws from the labor market to focus on home production. We propose and estimate a model of household labor supply that aligns with these trends. The main drivers of the urban-rural participation gap are higher spousal incomes in cities, which reduce the marginal utility of female labor, and labor market distortions that depress women's urban wages below their marginal product. Counterfactual simulations show that economic growth is unlikely to provide a sharp reversal of this trend in future decades unless it is accompanied by changes in gender norms and labor market institutions.

*We thank the Development Policy and Finance team at the Bill and Melinda Gates Foundation and the Tobin Center for Economic Policy at Yale University for funding. This material is based upon work supported by the National Science Foundation Graduate Research Fellowship under Grant No. (NSF grant number DGE-2139841). Youdan Zhang and Carmen Arbaizar Mazas provided outstanding research assistance.

†Yale University and NBER. m.peters@yale.edu

‡Yale University. pamela.torola@yale.edu

§Yale University. lindsey.uniat@yale.edu

¶Yale University and NBER. fabrizio.zilibotti@yale.edu

1 Introduction

In the past three decades, India underwent a remarkable transformation characterized by fast economic growth and rapid urbanization. In this context of modernization and economic development, one might expect to see a growing role for women in labor markets. The observation that India's growth was largely driven by services (Fan et al., 2023) should, in principle, reinforce this expectation, as historically the growth of the service sector has encouraged female labor participation in the West, due to the comparative advantage of women in providing services relative to manufacturing jobs. On the contrary, female labor force participation (FLFP) fell from a low 36% in 1987 to a dismal 27% in 2019. Although puzzling, this decline in FLFP is not unique to India, but rather a common pattern in today's developing economies (Agte et al., 2024; Boserup, 1970; Goldin, 1994).

The decline in FLFP as incomes rise, particularly in the early stages of development, can be explained by an income effect: As economic growth boosts spousal earnings, the incentives for women to engage in market work diminish. This trend is evident in India, where we observe a striking urban-rural participation gap. In cities, where households are richer, women are almost half as likely to participate in the labor force compared to rural areas. Our analysis of time-use data points to an additional factor: stark differences in the types of labor activities available to women in these settings. In rural areas, women typically engage in low-hour, often unpaid, informal activities. In urban areas, far fewer women engage in informal work. Instead, they either join the formal labor market or exit the workforce entirely, focusing on home production and leisure. Since the latter option prevails in India, urbanization contributes to the decline in female labor force participation.

We show that a simple neoclassical model of household labor supply can effectively capture these observed trends. Our theory highlights the importance of indivisibilities: formal jobs may offer high returns but come with a minimum hour requirement. This feature, for which we provide robust empirical evidence, introduces a nonconvexity in the household decision-making process. It encourages specialization, where one member engages in formal market work while the other allocates time to home production or flexible informal work. Moreover, rising male earnings in urban areas create a strong income effect, discouraging FLFP due to the high opportunity costs of home production.

The argument outlined above points at economic forces inducing specialization within the household: as the return to market activity grows, one partner works in the market while the other provides home production activities that the household values. But why, then, does specialization occur along gender lines? In our model, we account for three key differences between men and women. First, individuals may have different skills. Second, women may face labor market discrimination, meaning they are paid less than their marginal product. Third, women's extensive margin of market work might be distorted by traditional gender norms, which impose a utility loss on the household if the woman works. This supply distortion captures the social stigma that is associated with women working outside the home. Each of these mechanisms contributes to creating a comparative advantage for women in home production, leading to gender-based specialization. Importantly, we do not assume that women naturally possess an advantage in home production. Instead, we emphasize the importance of unequal labor market opportunities.

We calibrate our model using microdata on earnings and time use from both rural and urban India. Our calibration allows for two key differences. First, productivity and access to formal labor markets can be different across urban and rural districts. Second, we allow gendered distortions to differ between urban and rural areas, with women potentially facing greater wage discrimination in cities or being subject to distinct gender norms.

Despite its simplicity, our model successfully replicates salient moments related to relative employment and earnings, both between formal and informal work and across gender lines. In addition, it captures the allocation of hours between work in the market, home production, and leisure among household members. Notably, the model accurately reproduces the observed urban-rural participation gap, shedding light on the significant disparities in FLFP between urban and rural districts.

Our estimates suggest that urban areas exhibit higher productivity, particularly in formal jobs, and that access to formal employment is easier in cities. At the same time, urban women experience somewhat *weaker* social norms but face *stronger* wage discrimination in the workplace once employed. This inference is driven by three key empirical moments: the large gender wage gap in urban formal jobs, the stark urban-rural earnings gap in both formal and informal sectors, and the observed participation gap. The fact that men earn substantially more in urban areas points to higher productivity, while the relatively lower earnings of urban women suggest that they are paid below their marginal product. Both factors contribute to the low rates of female participation. Our model indicates that if the gender norms that influence labor supply were identical in urban and rural areas, the participation gap would be even greater than observed. Thus, the model requires relatively liberal norms in urban areas to align with empirical data on labor force participation.

We then use our calibrated model for two quantitative exercises. First, in the spirit of an"ex-post" analysis, we structurally decompose the observed urban-rural participation gap into the effects of productivity differences, labor market discrimination, and social norms. Our findings reveal that labor market discrimination accounts for the lion's share of the low female participation rate in Indian cities. Higher formal wages (and the associated income effects through increased spousal earnings) also contribute, but play a less significant role quantitatively. While traditional norms are an important determinant of the low FLFP in India, they are not responsible for the rural-urban gap—to the opposite, they reduce such gap.

Second, we use our model to forecast the future trajectory of FLFP in India. We find that productivity growth in formal jobs—a hallmark of future economic development—induces a U-shaped pattern of labor supply. Our estimates place India on the downward-sloping side of this "U," implying that, in the near term, economic growth could further reduce female participation. The theory predicts that continued productivity growth will eventually reverse this trend: as wages rise and access to formal labor markets improves, the opportunity cost of home production increases, causing the substitution effect to outweigh the income effect and driving FLFP upward. But how substantial is this mechanism in practice? Our analysis suggests that even after crossing to the upward-sloping side of the "U," the positive effect of growth on FLFP is relatively modest. While economic forces alone would eventually lead to higher participation, the progress would be painfully slow.

In contrast, if economic development is accompanied by reductions in labor market distortions and shifts in gender norms related to female labor supply, FLFP could grow relatively rapidly in the coming decades. Labor market discrimination, in particular, plays a crucial role by creating a wedge between male and female earnings. As male wages rise with economic development, while female wages remain low, the income effect leads to reduced female labor participation. However, raising women's wages increases the opportunity cost of home production, ultimately boosting FLFP. This explains why urbanization and the expansion of the service sector are associated with increased female labor supply in more advanced economies. Likewise, changes in social norms and the decline in gender stereotypes can significantly speed up the increase in FLFP over time.

Our quantitative analysis underscores that while economic forces—such as income effects and frictions in accessing formal labor markets—play a significant role in explaining recent trends, relying solely on economic growth is unlikely to drive substantial changes in the coming decades. Meaningful progress requires a broader process of economic development that includes institutional and cultural transformations.

Related Literature: Our paper contributes to a substantial body of literature on FLFP and economic development. On the empirical side, Goldin (1994) uses cross-country data to highlight a U-shaped pattern in FLFP driven by income and substitution effects, with related insights provided by Agte et al. (2024). These findings broadly align with the historical evidence for advanced countries. A recent paper by Ngai et al. (2024) documents a U-shaped trend in women's hours worked in the United States during the period 1870-2019. The study finds a mild decline in female labor hours until the mid-20th century, followed by a substantial and sustained increase thereafter, providing a comprehensive perspective on historical labor dynamics for women.

On the theoretical side, shifts in FLFP have been associated with structural change (Ngai and Petrongolo, 2017; Rendall, 2018; Ngai et al., 2024), evolving social norms (Fernández, 2013; Fogli and Veldkamp, 2011; Jayachandran, 2015; Olivetti et al., 2024), and technological advances in household appliances (Greenwood et al., 2005). Our work incorporates the insights of this literature underscoring the urban-rural dimension and the role of labor market frictions as key factors shaping female participation.

Our study specifically examines India, where FLFP remains strikingly low compared to other countries at similar stages of economic development—see Fletcher et al. (2019) for a comprehensive survey. Among the existing empirical studies of FLFP in India, Klasen and Pieters (2015) argue that income effects and changes in the sectoral structure of employment have contributed to the decline in aggregate FLFP. They also note the rural-urban gap. Deshpande and Kabeer (2024) emphasize the importance of accounting for unpaid work in family businesses, especially given its prevalence among women in West Bengal. Afridi et al. (2020) highlight the unique role of agriculture as a significant source of female employment. Related to our findings, Li (2023) also use time-use data to assess female employment across urban and rural areas. Unlike our focus on cross-sectional differences, however, their study focuses on time series trends within each location. Jayachandran (2020) and Munshi and Singh (2024) discuss the role of a variety of gender norms as a deterrent of FLFP in India.

From the standpoint of theory, our work builds on a substantial literature in household labor supply (e.g., Aguiar and Hurst (2007); Boerma and Karabarbounis (2021)), incorporating key market frictions such as discrimination—recognized as a major factor in female labor supply decisions (Hsieh et al., 2019; Chiplunkar and Kleineberg, 2022; Chiplunkar and Goldberg, 2021). The structure of our model is shares many features with Ngai and Petrongolo (2017) while integrating additional dimensions like social norms and paydiscrimination in urban labor markets that are especially relevant in the Indian context.

Our main contribution lies in the explicit use of microlevel time-use data, enabling us to directly measure labor inputs across household members and precisely estimate our model. In this regard, our study shares similarities with Gottlieb et al. (2024), who analyze cross-country time-use data but abstract from the urban-rural dimension.

Roadmap: Section 2 outlines the data sources. Section 3 presents some empirical analysis. Section 4 introduces the theoretical model. Section 5 discusses model estimation and the fit of targeted and nontargeted moments. Section 6 performs counterfactual simulations. Section 7 concludes. Additional empirical and theoretical results are provided in the Appendix.

2 Data

We use two data sources: the Periodic Labour Force Survey (PLFS) and the Time Use Survey (TUS), both conducted by the Ministry of Statistics and Program Implementation (MoSPI). In the following subsections, we provide a description of the data sources.

Figure 1 illustrates the geographic coverage of our datasets for 2019. We restrict our sample to districts with over 50 observations in a given year. Panel 1a displays the districts available in the PLFS, 1b in the Time Use Data. The geographical distribution of districts is quite similar between PLFS and TUS, covering 680 and 676 districts, respectively.

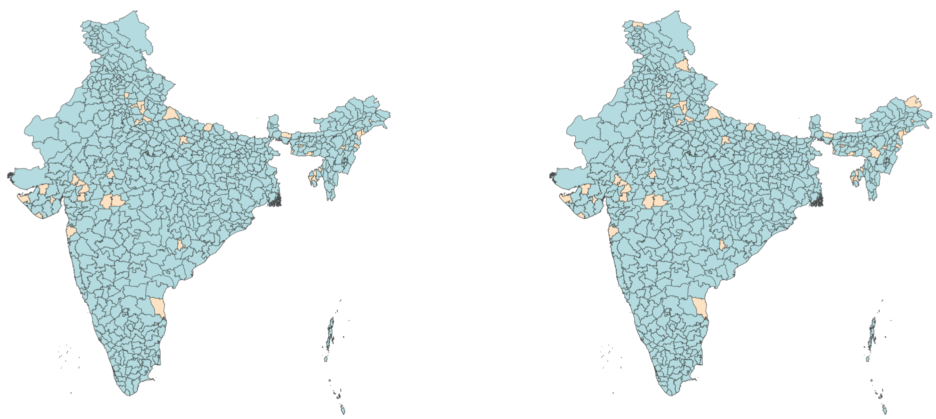

(a) PLFS (680 districts) **(b)** Time Use (676 districts)

Notes: Map of the availability of districts by dataset. Available districts are colored in blue. Panel 1a displays districts available in PLFS; Panel 1b, in TUS.

Figure 1: Data Availability

Table 1 presents summary statistics for each of these datasets, focusing on individuals aged 25-60 years old. Demographic characteristics, such as average age (around 36 years) and balanced gender ratios, are consistent across the datasets. However, there are slight differences in urbanization rates. Labor force participation rates also quite similar, approximately 60% in both PLFS and TUS. Educational attainment is categorized into five levels: Not literate, Primary and below, Secondary and below, Bachelor equivalent, and Postgraduate and above, and the shares are quite similar across datasets.

The two datasets also exhibit notable differences in sectoral employment shares. The service sector represents 37.9% in PLFS, and only 16.2% in TUS. The industrial sector shows an inverse pattern: it accounts for only 23% of employment in PLFS, but rises to 40.4% in TUS. The agricultural employment share is relatively consistent across datasets.

Table 1: Descriptive Statistics for PLFS and TUS

	PLFS	Time Use
Number of Observations	197395	254689
Number of Districts	675	676
Mean Age	40.32	39.91
Percentage Female	50.55%	50.42%
Percentage Urban	44.84%	40.18%
Labor Force Status		
Domestic Duties	33%	34.76%
In the Labor Force	63.01%	61.7%
Sector of Economic Activity		
Agriculture	33.02%	37.06%
Industry	23.42%	40.5%
Services	39.97%	17.35%
Highest Educational Attainment		
Not literate	21.35%	24.56%
Primary and below	39.67%	35.98%
Secondary equivalent	24.59%	24.65%
Bachelor equivalent	10.43%	11.08%
Postgraduate and above	3.96%	3.73%

2.1 Periodic Labor Force Survey

The Periodic Labour Force Survey (PLFS) is a nationally representative survey that collects income, consumption and employment information for over 100,000 urban and rural households. In total these data cover about 400,000 individuals.[1] The survey covers the whole of the Indian Union except the villages in Andaman and Nicobar Islands, due to difficulties in access. PLFS is conducted yearly and organized in quarterly waves. At the time of this writing, PLFS has released four years of data. PLFS 1 covers mid-2017 to mid-2018; PLFS 2, mid-2018 to mid-2019; PLFS 3, mid-2019 to mid-2020; and PLFS 4, mid 2020 to mid-2021.

PLFS divides the Indian population into two types of regions, rural and urban. PLFS is representative of rural and urban areas within each NSS Region (a subdivision of state). New urban and rural panels are added each quarter. However, rural households are only visited once, while urban households are revisited quarterly for a total of four quarters. After this, they drop out of the sample. In our analysis, we utilize all rural data and the first visit of each urban household, converting the data into a representative repeated cross-section.

PLFS defines employment in four different ways: (i) usual principal activity status (UPAS), (ii) usual subsidiary activity status, (iii) current weekly activity status, and (iv) current daily activity status. An individual is considered to be employed as a principal activity (UPAS), if they

[1] The PLFS is the new version of the subject *Employment and unemployment* on the quinquennal National Sample Surveys (NSS), conducted by the National Sample Survey Office (NSSO).

report being employed for more than half of the preceding year. This usual principal activity is the measure of labor force participation that we use. In PLFS, income is reported conditional on employment status. Income is reported from three sources of employment: (i) regular wage or salary, (ii) casual/daily wage, and (iii) self-employment. Income from regular wage employment is asked for the last month; income from self-employment is asked for the last 30 days; and income from casual wage employment is asked for the past week. In order to make all the three sources of income comparable, we construct a weekly measure of income from regular wage or salary, and for income from self-employment. We then combine the three measures of weekly income for each individual to get a value of labour income per week.

2.2 Time Use Survey

The Time Use Survey (TUS) is also conducted as part of the National Sample Surveys (NSS) by the Ministry the Ministry of Statistics and Program Implementation (MoSPI). TUS collects data on time use over a 24-hour period for individuals over age 6 in a random sample of Indian households. Similar to PLFS, TUS covers the whole of the Indian Union except the villages in Andaman and Nicobar Islands. The duration of the survey is one year and, in this paper, we make use of the 2019 round.

For our analysis, we construct 7 activity groups based on ICATUS 2016: formal employment, household employment, subsistence employment, other employment, home production, leisure, and education.

1. Formal employment: Division 11 ("Employment in corporations, government and non-profit institutions")
2. Household employment: Divisions 12 and 13 ("Employment in household enterprises to produce goods" and "Employment in households and household enterprises to provide services")
3. Subsistence employment: Major division 2
4. Other employment: Divisions 14-18 (breaks, seeking employment, commuting, etc.)
5. Home production: Major divisions 3 and 4
6. Leisure: Major divisions 5 and 7-9
7. Education: Major division 6

3 Empirical Analysis

In this section, we combine the datasets outlined above to examine women's participation in the labor market over time and across regions. The left panel of Figure 2 presents the time series of FLFP in India over the past four decades. We show the aggregate FLFP for women aged 25-60 in green, while the participation rate for rural and urban women is represented in orange and blue, respectively.

Three notable trends emerge. First, FLFP in India is strikingly low: only about one-third of women are in the workforce. Second, between the late 1980s and 2019, the participation rate actually *declined*, falling from approximately 35% in 1987 to 27% by 2019.[2] In contrast, male labor force participation has remained consistently high, hovering around 96% throughout this period (see Figure A-3 in the Appendix). Finally, there is a significant urban-rural gap in FLFP, with women in urban areas much less likely to work, though urban participation has seen a slight increase over the past decade.

The right panel shows why the urban-rural FLFP-gap might play an important role for the secular fall of the aggregate FLFP rate: India, like all developing countries, has been urbanizing fast. Between 1987 and 2018, the urban share increased from 22 percent to 31 percent and cities have been major driver of the rise in living standards. Such reallocation, however, might come at the expense of women's attachment of the labor market given their substantially lower participation rate.

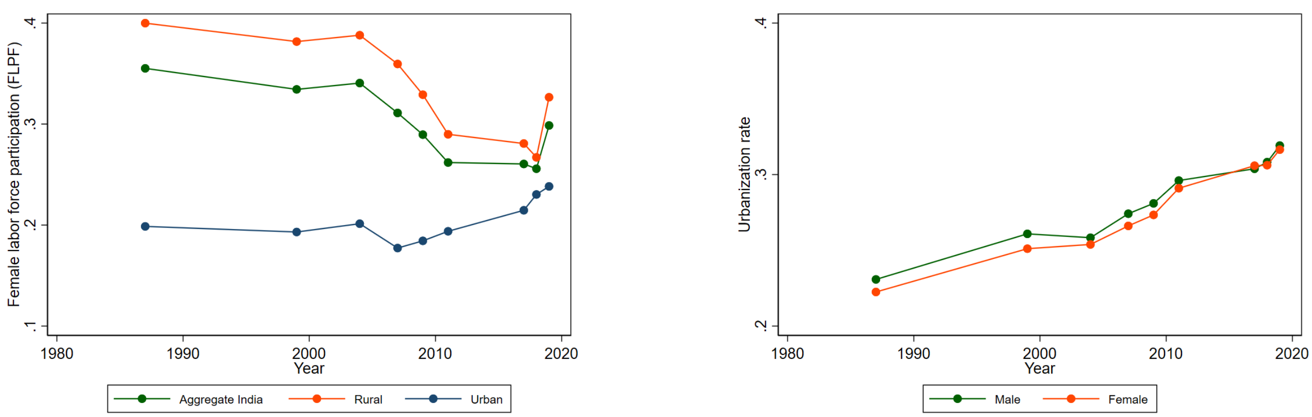

(a) Female LFP **(b)** Urban Share

Panel (a) shows the labor force participation rate for women aged 25–60 in India (green), with a breakdown for urban areas (blue) and rural areas (orange). Panel (b) shows the urbanization rate in India for men (green) and women (orange). The data are sourced from NSS and PLFS.

Figure 2: Female Labor Force Participation and Urbanization: 1987-2019

3.1 The Urban-Rural FLFP Gap in India

To provide more rigorous evidence on the urban-rural FLFP gap documented in Figure 2, we consider a regression of the form

[2] The sudden increase in rural FLFP after 2018 is unexpected and may be related to changes in the PLFS sampling frame.

$$LFP_{ir} = \delta_{s(r)} + \beta \text{Urban}_{ir} + \alpha_1 \Upsilon_r + \alpha_2 \Upsilon_r^2 + \gamma s_r^{SERV} + x_{ir}' \varrho + \epsilon_{ir}, \tag{1}$$

where LFP_i is a dummy that takes the unit value when a given individual i living in district r is employed, $\delta_{s(r)}$ is set of state fixed effects, $Urban_{ir}$ is an indicator whether the individual lives in a city, Υ_r is a measure of economic development at the district level (see below for details of the construction), s_r^{SERV} is the employment share of services within the non-agricultural sector, and x_{ir} is a rich set of individual controls including educational attainment, age, marital status, caste and religion. All standard errors are clustered at the district level.

Table 2: The Urban-Rural Gap in Female Labor Force Participation

	All Women					Married Women
	(1)	(2)	(3)	(4)	(5)	(6)
Urban (indiv)	-0.125***	-0.137***	-0.129***	-0.122***	-0.122***	-0.129***
	(0.007)	(0.007)	(0.006)	(0.006)	(0.006)	(0.007)
Development Index				-0.022***	-0.022***	-0.024***
				(0.003)	(0.003)	(0.004)
(Development Index)^2				0.006***	0.006***	0.006***
				(0.001)	(0.001)	(0.001)
Service share (nonag.)					-0.034*	-0.034*
					(0.018)	(0.019)
R^2	0.14	0.16	0.20	0.20	0.20	0.19
N	617,128	616,539	616,533	565,311	565,247	485,070
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes	Yes
District FE	Yes	Yes	Yes	Yes	Yes	Yes
Education	No	Yes	Yes	Yes	Yes	Yes
Age	No	No	Yes	Yes	Yes	Yes
Marital Status	No	No	Yes	Yes	Yes	Yes
Caste # Religion	No	No	Yes	Yes	Yes	Yes

Notes: Standard errors are clustered at the district level. "Urban" is an indicator variable whether the woman lives in an urban area. The "Development Index" is the first principal component of a principal component analysis using district level income per capita, the agricultural employment share and the urbanization rate. The "Service Share" is the share of the non-agricultural labor force that works in services. Columns 3-6 control for a full set of age fixed effects, for marital status, and a full set of interactions between religion and caste fixed effects. In column 6 we only focus on married women.

In Table 2, we report the estimates of equation (1). Column 1 shows that, within states, urban women are 12.5% less likely to participate in the labor force. Given an average participation rate of around 30%, this implies that the participation rate for urban women is roughly two-thirds that of rural women. In Columns 2 and 3, we show that this gap is not explained by observable heterogeneity. In fact, the gap slightly increases after controlling for education (Column 2), and additional controls such as age, marital status, and a full set of interacted caste and religion fixed effects (Column 3).

In Columns 4 and 5, we add additional controls for economic development at the district level. First, we control for the level of development using a district-level development index, which we construct as the first principal component from a principal component analysis (PCA) based on district-level per capita income, agricultural employment share, and urbanization rate.[3] Including this index does not significantly alter the urban-rural FLFP gap. Additionally, the relationship between the development index and FLFP follows the familiar U shape, where economic development initially reduces but later increases FLFP. We explore this U-shaped pattern in more detail in Section 3.2. Column 5 adds a control for the share of the labor force employed in services at the district level. As expected, we find that FLFP is higher in districts with a higher prevalence of service employment, although this relationship is not statistically significant.

In Columns 1-5, we focused on women aged 25–60. One might expect that the labor supply of married women differs qualitatively, given their access to other sources of family income, particularly their husband's earnings, which could encourage a specialization in home production over market work. In Column 6, we replicate the analysis from Column 5, restricting the sample to married women. The results remain both qualitatively and quantitatively similar, which is unsurprising since 85% of women aged 25-60 are married.

In Table 2, we focused solely on the PLFS data, which is arguably the most reliable source for measuring employment outcomes. However, the rural-urban FLFP gap is similarly pronounced in the Time Use data. This is demonstrated in Table 3, where we restrict the sample to married women. We observe consistent qualitative results across both datasets, although the significance levels of some coefficients vary.

First, in both datasets and specifications, we consistently observe a significantly negative relationship between urbanization and FLFP. Second, the effect of the development variable qualitatively follows a U-shaped pattern, though its significance varies across datasets.

The time series of FLFP presented in Figure 2 suggests a convergence between rural and urban labor markets in terms of women's participation over time. In Appendix Section A-0.3, we further investigate this trend by analyzing the dynamics through regressions based on equation (1). Specifically, we allow the coefficient β on the urban dummy to vary across different time periods in the full panel of regions. The results show that β_t follows a similar trajectory to the pattern displayed in Figure 2, even when controlling for individual characteristics. Therefore, changes in

[3] We provide a detailed explanation of this construction in Appendix Section A-0.2. There, we also show that the first principal component positively loads on urbanization and income per capita, and negatively on agricultural share.

Table 3: The Urban-Rural Gap in Female Labor Force Participation: Comparison of Data Sources

	PLFS		TUS	
	(1)	(2)	(3)	(4)
Urban	-0.109***	-0.117***	-0.087***	-0.087***
	(-14.44)	(-14.60)	(-29.38)	(-11.07)
Development Index	-0.019***	-0.021***	-0.020***	-0.022***
	(-3.83)	(-4.09)	(-14.72)	(-4.47)
(Development Index)^2	0.006***	0.006**	0.002***	0.002
	(2.63)	(2.56)	(3.87)	(0.94)
Service share (nonag.)	-0.012	-0.009		0.077*
	(-0.33)	(-0.23)		(1.89)
R^2	0.16	0.15	0.15	0.16
N	102,346	88,032	111,120	111,109
State FE	Yes	Yes	Yes	Yes
Individual Covariates	Yes	Yes	Yes	Yes

Notes: Standard errors are clustered at the district level. "Urban" is an indicator variable whether the woman lives in an urban area. The "Development Index" is the first principal component of a principal component analysis using district level income per capita, the agricultural employment share and the urbanization rate. The "Service Share" is the share of the non-agricultural labor force that works in services. All specifications control for a full set of age fixed effects and a full set of interactions between religion and caste fixed effects. We use data from the PLFS (columns 1 and 2) and the time use survey (columns 3 and 4). For all specifications we focus on married women.

the composition of urban locations do not explain the convergence. Instead, the findings suggest that urban and rural labor markets in India have followed distinct but converging paths in terms of women's labor market participation.

3.2 FLFP and Economic Development: The Cross-Sectional U Shape

To analyze the U-shaped relationship between FLFP and economic development without imposing a specific quadratic functional form, we estimate a non-parametric regression using decile indicators of the district-level development index. This approach allows us to capture the shape of the relationship between FLFP and development more flexibly:

$$LFP_{ir} = \delta_{s(r)} + \sum_{d=1}^{10} \alpha_d 1\{\Upsilon_r \in d\} + \beta \text{Urban}_{ir} + \gamma s_r^{SERV} + x'_{ir}\varrho + \epsilon_{ir}. \tag{2}$$

Here, $1\{\Upsilon_r \in d\}$ is an indicator for Υ_r being the dth decile of the cross-sectional distribution Υ_r. The 1st (10th) decile consists of the districts with the lowest (highest) level of development. We use the 10th deciles as the baseline category and normalize its coefficient to 0 and focus on all married women aged 25-60. The results for all women are identical.

In Figure 3, we plot the coefficients of α_d across all deciles for both the PLFS (green) and the TUS (orange) datasets. Remarkably, the trends are quite consistent across the two datasets. First,

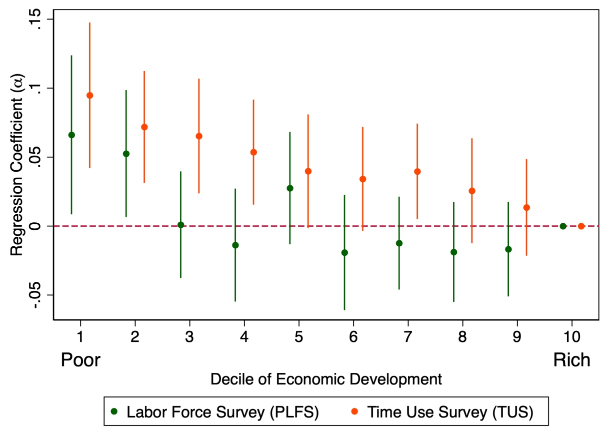

Source: 2019, PLFS and TUS. Calculated by the authors. Notes: We depict the estimates of α_d based on equation (2). We focus on married females ages 25–60. The vertical lines represent the 95% confidence interval based on district-level clustered standard errors.

Figure 3: Economic Development and FLFP: The U Shape Across Indian Districts

we observe a clear decline in FLFP as local development increases. Second, particularly in the PLFS, this relationship appears to plateau at higher levels of economic development, consistent with the U-shaped pattern when restricting the relationship to be quadratic, as shown in Tables 2 and 3.

3.3 Income Effects: FLFP and Spousal Earnings

A key aspect of our theory is the presence of income effects, where female labor supply decreases as household income rises. In Table 4, we use cross-sectional data from the PLFS to explore the relationship between FLFP and spousal income. We adopt the same specification as in equation (1), adding controls for the husband's total earnings, education, and sector of employment, while accounting for women's observable characteristics like age and education.

Columns 1-3 focus on the full sample and show a significantly negative relationship between spousal earnings and female LFP. Specifically, a 10% increase in spousal earnings leads to nearly a 1% reduction in FLFP. Notably, even after controlling for spousal income, the negative urban coefficient remains robust, indicating that income effects alone do not fully explain the lower FLFP in cities. In our structural model we attribute this to both higher productivity in urban areas (driving up incomes) and institutional factors specific to urban labor markets that independently affect female labor supply.

Columns 2 and 3 introduce further controls for spousal characteristics that may be correlated with household income and female labor supply. Interestingly, the husband's education is found to lower FLFP (controlling for the woman's education), suggesting that permanent income plays a larger role than current income. Additionally, accounting for the husband's employment sector does not significantly alter these results.

In the final columns of Table 4, we apply the same specification separately to rural and urban samples.[4] These results reveal that income effects are equally strong in both rural and urban settings, with comparable quantitative magnitudes.

Table 4: Income Effects: FLFP and Spousal Earnings

	Full Population			Rural		Urban	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Log Wage Husband	-0.095***	-0.086***	-0.072***	-0.094***	-0.066***	-0.092***	-0.076***
	(0.003)	(0.003)	(0.003)	(0.004)	(0.004)	(0.003)	(0.003)
Education Husband		-0.009***	-0.007***		-0.005***		-0.010***
		(0.001)	(0.000)		(0.001)		(0.001)
R^2	0.23	0.23	0.24	0.29	0.30	0.12	0.12
N	310,013	310,013	309,706	180,448	180,270	129,564	129,435
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
District FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Individual Covariates	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Husband Sector Controls	No	No	Yes	No	Yes	No	Yes

Notes: Standard errors are clustered at the district level. Husband's education is a continuous variable with 13 distinct values, ranging from "Not literate" to "Postgraduate". Individual covariates include age, educational attainment, and an interaction term between religion and caste category. We include dummy variables to control for the husband's employment in the formal or informal sector, as well as for his sector of employment (agriculture, manufacturing, or services).

3.4 Female Labor Force Participation and Time Use

So far, we have analyzed female labor supply using the traditional, binary measure of labor force participation (FLFP). However, this measure may be imprecise in a developing economy like India, especially in rural areas where many work arrangements are informal or involve subsistence activities. Such nuances could contribute to the observed urban-rural FLFP gap, as rural women may report time spent on farms as work, while urban women classify similar activities as home production. On the other hand, this gap might also *underestimate* differences in participation if rural women underreport informal work, which could be more prevalent in rural areas. To better understand this gap, we now turn to time use data.

Our sample is the set of household heads and spouses in the 2019 TUS. We restrict our analysis to couples with one person of each gender, and ages 25-60. Using the 6 activity groups outlined Table 5, we classify each individual into one of three mutually exclusive groups:

1. Formal workers: Anyone who performs formal work > 0 hours in the survey day

2. Informal workers: Anyone who performs 0 hours of formal work, but performs at least some household employment, subsistence, or other employment

[4]For brevity, we report only the baseline specification (1) and the full specification (3).

Table 5: Work Status Division

	Formal	**Informal**
Female	3,076	28,091
of who are also doing informal work (%)	14.01	
Male	11,173	53,002
of who are also doing informal work (%)	8.58	

Notes: The table presents the count of formal and informal workers by gender in our 2019 TUS sample. Informal workers report only informal work hours, while formal workers report more than 0 formal work hours. Rows 2 and 4 show the percentage of formal female and male workers, respectively, who also report some informal work.

3. Neither: Anyone who documents 0 hours spent on formal or informal work

Under this classification, a worker is either classified as formal or informal, but not both. These mutually exclusive categories align with the data on individual time use, as the vast majority of those engaged in formal work spend zero hours on subsistence or household employment—two primary forms of informal work. As shown in Table 5, only 14% of women and 9% of men who report any formal work also engage in informal work within these categories. This distinction likely captures two separate modes of work that individuals must choose between: formal and informal sectors.

In the left panel of Figure 4, we illustrate the distribution of these three activity categories among rural and urban women. The figure reveals a clear pattern: consistent with the substantial rural-urban FLFP gap, many more women in cities do not report any hours spent on either formal or informal activities. However, there is a significant shift away from informal work in urban areas. Among urban women, the share engaged in formal rather than informal work is about 6%. In contrast, formal employment is nearly absent in rural India.

This stark contrast between formal and informal work raises the question of whether rural women engaged in informal sectors actually identify as workers. The answer is found in the right panel of Figure 4, where we document the share of women who report being in the labor force versus performing domestic duties. In both rural and urban areas, formal employment is strongly associated with reporting work as the primary activity. This differs sharply in informal sectors, where approximately 20% of women engaged in informal work do not report themselves as being in the labor force. Given the higher number of women in informal work in rural areas, this implies that the rural-urban FLFP gap may underestimate the true decline in FLFP in cities.[5]

Figure 4 focuses solely on the extensive margin of work. In contrast, Figure 5 delves into the intensive margin, which is crucial for measuring overall labor inputs—especially if many women only work a few hours in informal jobs. This distinction proves important, as it aligns with what we observe in the data. In the left panel, we display the distribution of formal and informal hours (conditional on reporting positive hours in either category) for rural women. The right panel shows

[5] Among women who report neither formal nor informal work, almost none identify as being in the labor force.

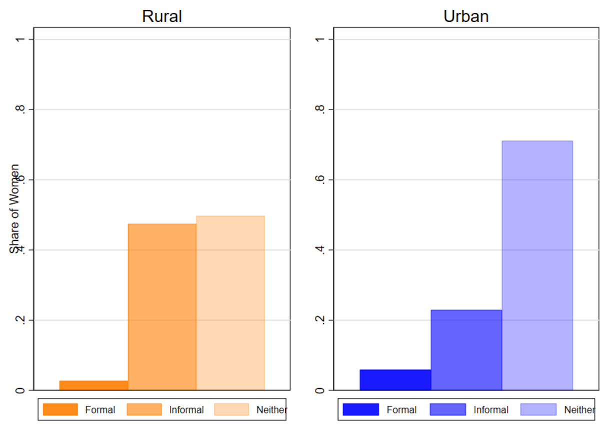

(a) The Nature of Work

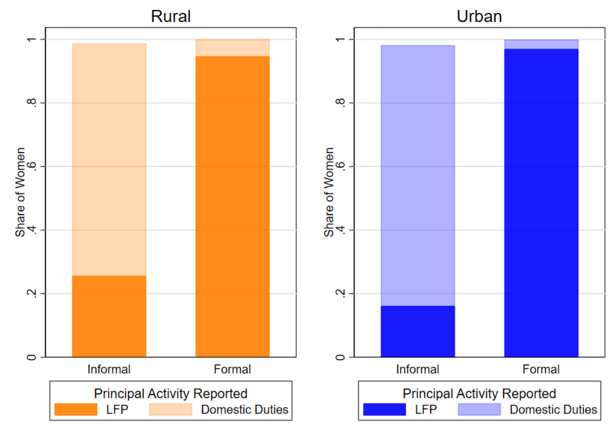

(b) Work vs. Labor Force Participation

Notes: In the left panel we report the share of women doing formal, informal, or no employment in rural and urban India. In the right panel we report the women reporting to be in the labor force (dark shade) or performaing domestic duties (light shade) as a function of their employment status in rural and urban India.

Figure 4: Time Use and Female Labor Force Participation in Rural and Urban India

the same for urban women.

The results in Figure 5 reveal substantial variation in the distribution of hours across activities. Formal jobs tend to be more rigid, with most women working at least four hours in rural areas and five hours in cities, conditional on working in the formal sector. In contrast, informal work offers greater flexibility, allowing women to work significantly fewer hours. The distribution is more dispersed, with a considerable number of women working just one or two hours. Interestingly, despite the formal/informal divide, the distribution of hours remains relatively similar between rural and urban areas. This suggests that formal jobs in cities do not have markedly different hour-requirements than their rural counterparts.

Taken together, Figures 4 and 5 suggest an important difference between rural and urban labor markets in India that could explain the differences in women's labor market choices. Formal jobs seem to be rigid in nature and offer less flexibility in the number of hours spent. Informal jobs, by contrast, provide such flexibility but the opportunities to engage in such activities in cities might be harder to come by. As a consequence, urban women might decide to drop out of the labor force entirely if only given the choice of formal, full-time employment opportunities. To analyze whether this mechanism can in fact explain observed labor market choices and time-use patterns, we require a structural model of household labor supply. This is where we turn now.

4 Theory

To explain these empirical findings, we propose a parsimonious model of household labor supply. Cities and rural areas differ in three dimensions: overall labor market opportunities (such as productivity or the ease of finding formal jobs), gender specific labor market frictions (such as

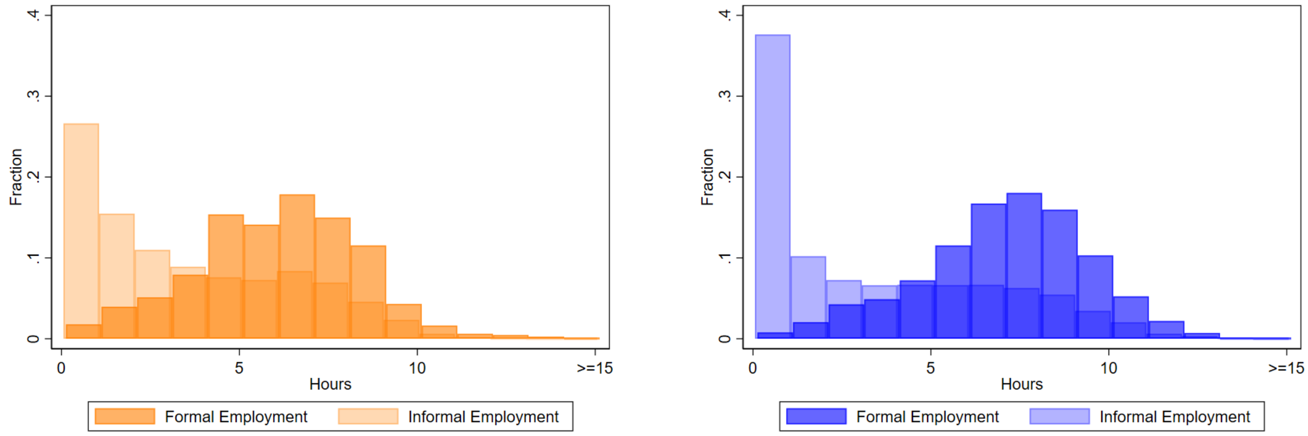

(a) Rural Areas **(b)** Urban Areas

Notes: The figure shows the distribution of hours among women in informal (light shade) and formal (dark shade) employment. In the left (right) panel we focus on rural (urban) areas.

Figure 5: The Intensive Margin of Formal and Informal Work

distortions for women to access market work), and differences in gender norms that might result in women staying out of the labor force. We calibrate our model separately for urban and rural areas, and then use our theory to quantitatively decompose the observed urban-rural participation gap into these different components.

4.1 Environment

We consider a neoclassical model of household labor supply, where each household consists of a husband (m) and a wife (w). The household model is unitary, meaning labor supply decisions are made to maximize household utility. Each spouse is endowed with one unit of time, which can be allocated to formal work f, informal work i, home production h, or leisure l. Both informal and formal work yield income that the household spends to purchase market consumption goods C.

Preferences: We parameterize household utility as

$$\mathcal{U} = \ln \mathcal{C} + \beta \ln L - U_{FLFP} \mathbb{1}\{f_w + i_w > 0\} \tag{3}$$

where

$$\mathcal{C} = \left(\psi C^{\frac{\sigma-1}{\sigma}} + \mathcal{H}^{\frac{\sigma-1}{\sigma}}\right)^{\frac{\sigma}{\sigma-1}} \tag{4}$$

$$\mathcal{H} = \left(h_w^{\frac{\eta-1}{\eta}} + h_m^{\frac{\eta-1}{\eta}}\right)^{\frac{\eta}{\eta-1}} \tag{5}$$

$$L = \left(l_w^{\frac{\gamma-1}{\gamma}} + l_m^{\frac{\gamma-1}{\gamma}}\right)^{\frac{\gamma}{\gamma-1}}. \tag{6}$$

Household utility $\mathcal{U}$ depends on household consumption $\mathcal{C}$ and household leisure L. Additionally, we allow for female labor to be a disamenity at the household level: if female working hours are positive, i.e. $f_w + i_w > 0$, household utility is reduced by U_{FLFP}. This term broadly captures female labor supply distortions driven by gender norms. Mathematically, it operates like a fixed cost of female labor supply. This gendered utility loss is the only difference between men and women in terms of preferences.

Equation (4) shows that household consumption $\mathcal{C}$ is a CES aggregate of market goods C and home production $\mathcal{H}$. Home production $\mathcal{H}$, in turn, is a CES aggregate of the home production hours contributed by each partner (see Equation (5)). Lastly, household leisure L is a CES aggregate of the leisure hours of both partners.

In terms of the respective elasticities σ, η, and γ, we assume that

$$\gamma \ < \ 1 \ < \ \sigma \ < \ \eta. \tag{7}$$

Thus, spousal leisure hours l_w and l_m are complements ($\gamma < 1$), capturing the idea that leisure hours are typically spent together. In contrast, market consumption C and home production $\mathcal{H}$ are imperfect substitutes—for instance, a home-cooked meal versus a meal in a restaurant ($\sigma > 1$). Finally, spousal home-production hours h_w and h_m are substitutes in producing $\mathcal{H}$ and we assume the input of both partners to be more substitutable than home production and market consumption ($\eta > \sigma > 1$). Note that we treat men and women symmetrically in the provision of home-produced goods, meaning there is no gender-specific comparative advantage in home production tasks.

Formal and Informal Work: To purchase market goods C, both household members can choose to earn income by working in the market. There are two types of work: formal f and informal i, which differ in three main dimensions.

First, they vary in their returns per unit of time spent. We assume individuals are heterogeneous in their formal skills but provide the same number of efficiency units in the informal sector. Specifically, any worker engaged in informal work earns $A^I i$, where i is the time allocated to informal work. In contrast, a worker who spends f hours in the formal sector earns $A^f \epsilon f$, where A^f is the formal wage, ϵ represents the realization of an individual's idiosyncratic efficiency units, and f is the time spent on formal work.

Second, building on our empirical finding in Figure 5 that formal work is time-intensive, we assume formal work is relatively inflexible, requiring a minimum of $\bar{f}$ hours to generate income, such that $f \geq \bar{f}$. Informal work, however, is more flexible and can be performed even for small amounts of time. This idea aligns with the model of Erosa et al. (2022), which differentiates occupations based on how wages convexly depend on hours worked.

Third, not all households have access to formal jobs. We capture this restriction through the random variable $\mathcal{A} \in \{0, 1\}$, assumed to be independent and identically Bernoulli distributed, where $\mathcal{A} = 1$ indicates that the household has access to formal employment. This setup reflects labor market frictions, where the absence of formal employment is not a result of low labor supply but rather limited labor demand.[6]

The Household Budget Constraint: The household's financial budget constraint is given by

$$C = A^i(i_m + (1-\tau)i_w) + A^f \left(\epsilon_m f_m \mathbb{1}\{f_m \geq \bar{f}\} + \epsilon_w (1-\tau) f_w \mathbb{1}\{f_w \geq \bar{f}\}\right) \times \mathcal{A}, \tag{8}$$

where $\tau \in (0, 1)$ is a wedge that captures gender discrimination or other institutional factors that reduce the pay of women.[7] If the household does not have formal labor market access, i.e., $\mathcal{A} = 0$, market consumption is solely determined by informal hours, $i_m + (1-\tau)i_w$. If the household has access to the formal labor market, $\mathcal{A} = 1$, there are both extensive and intensive labor supply choices: which market to supply hours to, and how many hours to supply. For informal work, the choice of hours is unconstrained, while in the formal market, hours must exceed the minimum threshold $\bar{f}$.

Heterogeneity: We assume that households differ along three dimensions: each spouse's formal efficiency units (ϵ_m, ϵ_w), the disamenity of female labor supply U_{FLFP}, and the access to formal jobs $\mathcal{A}$. Formally, we assume that the 4-tuple $\varphi \equiv (\epsilon_m, \epsilon_w, U_{FLFP}, \mathcal{A})$ is drawn from a joint distribution F such that

$$F(\varphi) = P\left[\epsilon_m \leq e_m, \, \epsilon_w \leq e_w, U_{FLFP} \leq u, \mathcal{A} = a\right]. \tag{9}$$

Equation (9) allows us to capture aspects of spatial and marital sorting in a reduced form. For instance, a positive correlation between ϵ_m and ϵ_w could represent assortative matching in the marriage market, where human capital is positively correlated across spouses. A correlation between skills ϵ and U_{FLFP} may reflect systematic differences in gender norms that are linked to formal skills, such as education. Similarly, formal labor market access $\mathcal{A}$ could vary across households within a given location or differ systematically between urban and rural labor markets. When we estimate our model below, we will impose more parametric restrictions on the joint distribution F in (9).

The Household Decision Problem: A household is fully parameterized by the 4-tuple $\varphi \equiv (\epsilon_m, \epsilon_w, U_{FLFP}, \mathcal{A})$. Given the realization of the spouses' skill draws, gender norms, and labor

[6] Note that we assume the shock is realized at the household level, that is, either both the husband and wife have access to formal job opportunities, or neither of them does. We consider this positive correlation realistic, as access to formal jobs typically depends on factors such as location or social networks, which are determined at the household rather than the individual level. The assumption of perfect correlation is made for simplicity.

[7] For simplicity, we assume that this discrimination applies equally to formal and informal activities.

market access, the household's decision problem is as follows:

$$V(\varphi) \;=\; \max_{\{f_g, i_g, h_g, l_g\}_{g=w,m}} \mathcal{U}, \tag{10}$$

subject to the financial budget constraint (8) and the time budget constraint

$$1 = f_g + i_g + h_g + l_g \text{ with } f_g, i_g, h_g, l_g \in [0,1] \text{ for } g \in \{w, m\}. \tag{11}$$

Note that the indivisibility constraint $f_g \geq \bar{f}$ is already imposed in the budget constraint (8).

4.2 Optimal Time Allocation and Labor Force Participation

We now characterize the optimal time allocation, which is the solution to the utility maximization problem (10), subject to the financial and time budget constraints (8) and (11). The inflexibility of formal work introduces a discrete choice in the household's decision-making process. To solve for the optimal time allocation of the household, we first consider how time would be optimally allocated under each of the possible cases, and then choose the case that provides the highest overall utility. The (seven) cases are:[8]

1. Both spouses works informally;
2. The husband works informally and the wife does not work;
3. One spouse works formally and the other does not work (two cases);
4. One spouse works formally and the other works informally (two cases);
5. Both spouses work formally.

Consider, first, the case where both spouses work in the informal sector (Case 1). In this situation, the household's decision problem simplifies to allocating time between informal work, home production, and leisure. Given the linearity of the budget constraint in hours, the household optimally allocates time across informal work, leisure, and home production, balancing the trade-off between leisure time, informal income, and the value of home production. More formally, the household solves the following problem:

[8] Note that the budget constraint (8) is linear in both f_g and i_g. This implies that, generically, the same spouse would never engage in both formal and informal work. Our assumptions—specifically, the norm and discrimination against female work together—rule out the possibility that the wife works informally and the husband does not work.

$$\max_{\{i_w,h_w,i_m,h_m\}} \ln\left(\psi A^i(i_m+(1-\tau)i_w)^{\frac{\sigma-1}{\sigma}} + \left(h_m^{\frac{\eta-1}{\eta}} + h_w^{\frac{\eta-1}{\eta}}\right)^{\frac{\eta}{\eta-1}\frac{\sigma-1}{\sigma}}\right)^{\frac{\sigma}{\sigma-1}}$$
$$+\beta \ln\left((1-i_m-h_m)^{\frac{\gamma-1}{\gamma}} + (1-i_w-h_w)^{\frac{\gamma-1}{\gamma}}\right)^{\frac{\gamma}{\gamma-1}} - U_{FLFP}\mathbb{1}\{f_w+i_w>0\}$$

When $\tau = 0$ and $U_{FLFP} = 0$, the labor supply of men and women enters the program symmetrically. In this case, both the husband and wife supply the same number of market hours, contribute equally to home production, and take the same amount of leisure. However, the wedge τ reduces the returns to market work for women and hence induces a specialization in home production. If we continue to assume $U_{FLFP} = 0$, both spouses participate in informal labor market activities, but the wife works fewer hours. When there is a norm against women's work ($U_{FLFP} > 0$), the household may decide that the wife does not participate in the labor market at all (Case 2).

Next, suppose the husband works in the formal sector (Cases 3, 4, and 5).[9] In our model, consistent with the data, formal workers earn higher wages than informal workers. Interestingly, *ceteris paribus*, when the husband works in the formal sector and earns a higher salary, it reduces the likelihood of the wife working at all. This can be explained by an income effect: in households where one spouse earns a high wage and works long hours (given the inflexibility of formal employment in terms of hours), market consumption C at the household level is high and the marginal value of the other spouse's contribution to home production increases. The presence of gender norms against women's work, coupled with labor market discrimination, further biases this outcome, leading to the typical situation where some husbands work in the formal sector while their wives do not work. Note that this is not always the case. When both the husband and wife have very high productivity in the formal sector, they may both choose to work and substitute home production with private consumption. However, as we will see, this pattern is relatively uncommon in the calibrated model.

Figure 7 illustrates the discussion above with the aid of an example. The figure depicts the decision of a specific household where the husband has potential earnings 50% higher if he secures a job in the formal sector (i.e. $\epsilon_m A_f = 1.5A_i$). In the left panel, the wife's productivity in the formal sector, ϵ_w, is assumed to be low, implying that the relevant choice for her is to either work in the informal sector or stay out of the labor force. In the right panel, the wife's productivity in the formal sector is assumed to be high, implying that she would rather work in the formal than in the informal sector. The parameters used in the figure align with the calibration discussed in the next section.

[9] It is theoretically possible for the wife to work in the formal sector while the husband works in the informal sector or does not work at all. However, we do not emphasize these possibilities in our discussion because they are empirically rare occurrences and are not significant in our quantitative analysis.

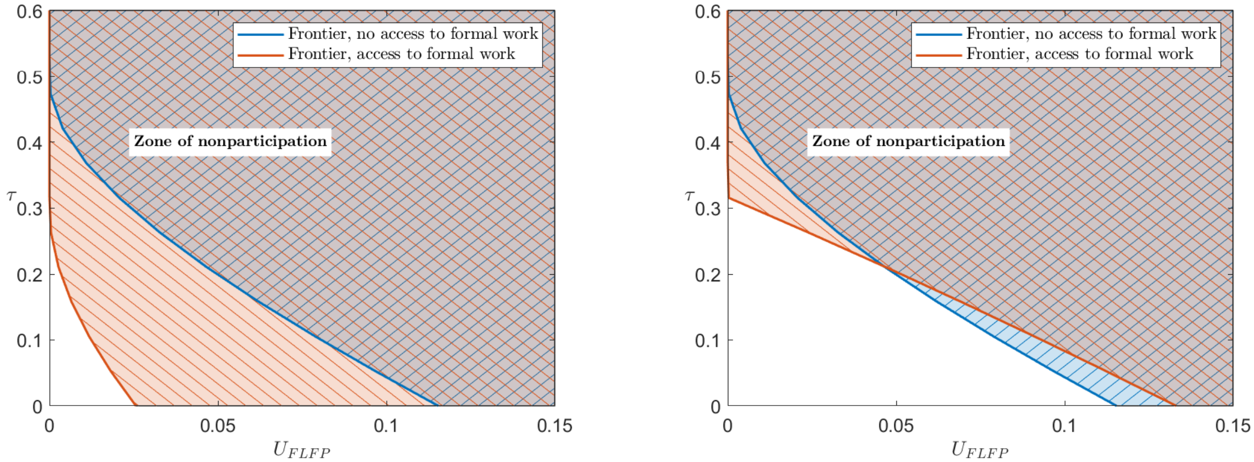

(a) Advantage in informal sector: $\epsilon_w A_f < A_i$ **(b)** Advantage in formal sector: $\epsilon_w A_f > A_i$

Notes: The figures display FLFP as a function of labor market discrimination τ and gender norms U_{FLFP}. We assume that the husband has a productivity advantage in formal work, i.e. $\epsilon_m A_f > A_i$. In the left panel we assume that the wife has low formal skills (i.e. $\epsilon_w A_f < A_i$), in the right panel she has a high formal skills (i.e. $\epsilon_w A_f > A_i$). The orange shaded area refers to the case where the household has access to formal jobs ($\mathcal{A} = 1$), the blue shaded area to the case where the household has no access for formal jobs ($\mathcal{A} = 0$).

Figure 7: Income Effects on Female Labor Force Participation: An Example

The figure shows combinations of values for the norm U_{FLFP} and labor market discrimination τ at which the household is indifferent to the wife participating in the labor market or staying out. Each panel illustrates two scenarios: the orange curve represents the case where the *household* has access to formal employment ($\mathcal{A} = 1$), and the blue curve represents the case where the household has no access to formal employment ($\mathcal{A} = 0$). Combinations of U_{FLFP} and τ to the left of each curve indicate that the wife participates in the labor market, while combinations to the right lead her to remain at home, focusing on home production.

Consider, first, the left panel. In this case, when the husband works in the formal sector, it reduces the range of parameters under which the wife also participates in the labor market. In other words, improved job opportunities for husbands tend to decrease women's labor participation, conditional on labor market institutions and gender norms. This reduction is driven by an income effect. Similarly, higher labor market discrimination τ and more adverse gender norms U_{FLFP} reduce FLFP.

Next, consider the right panel. The blue curve (no access to the formal labor market) is the same as in the left panel. However, the red curve shows that when the household has access to the formal labor market and the wife is highly productive, the range of U_{FLFP} and τ values where the wife participates in the labor market increases due to the higher opportunity cost of staying out. Additionally, the interplay between income and opportunity cost effects causes the orange and blue curves to intersect in this example. When labor market discrimination is high (large τ), the income effect dominates, making wives less likely to work when the household has access to better-paid formal jobs. On the other hand, when labor market discrimination is moderate (small τ), the opportunity cost effect dominates, and access to formal jobs expands the range of parameters where the wife participates in the labor market.

Taking Stock: Our model shows that women, due to social norms or labor market discrimination, tend to choose fewer hours of market work and focus more on home production. When good jobs are both scarce and rigid, requiring a minimum number of hours, women face a tough choice: either sharply reduce home production or stay out of the labor market. With urbanization and rising incomes, many women opt for the latter, moving from rural fields to urban kitchens as their husbands' formal income increases.

5 Quantification

To evaluate if our theory can explain the observed patterns in men's and women's work across urban and rural areas, we first calibrate the model to rural India. We then assess whether higher labor market earnings, along with shifts in norms and institutions, can account for the decline in female labor participation and changes in time allocation as we move from rural to urban areas.

5.1 Calibration

We assume that the fixed cost of market work, U_{FLFP}, the access to formal jobs, $\mathcal{A}$, and formal efficiency ϵ are i.i.d. Specifically, we assume that $U_{FLFP} \in [\bar{u}, \infty)$ follows a Pareto distribution with p.d.f. $f(U_{FLFP}) = \theta \frac{\bar{u}^{\theta}}{U_{FLFP}^{\theta+1}}$ and we denote by p the (Bernoulli) probability that $\mathcal{A} = 1$. Furthermore, we parameterize the distribution $\tilde{F}(e_m, e_w)$ as a multivariate log-normal distribution given by

$$\begin{pmatrix} \ln \epsilon_m \\ \ln \epsilon_w \end{pmatrix} \sim \mathcal{N}\left(\begin{pmatrix} \mu_m \\ \mu_w \end{pmatrix}, \begin{pmatrix} \sigma_\epsilon^2 & \rho\sigma_\epsilon^2 \\ \rho\sigma_\epsilon^2 & \sigma_\epsilon^2 \end{pmatrix} \right). \tag{12}$$

Hence, we allow the skills of men and women to be correlated within the household and assume that the dispersion of men's and women's skills, σ_ϵ^2 to be the same.

The model is thus fully parameterized by the parameters of the Pareto distribution for U_{FLFP} θ and $\bar{u}$, the Bernoulli probability p, the parameters of the distribution in (12), the preference parameters β and ψ, the elasticities σ, η, and γ, the indivisibility in hours for formal jobs $\bar{f}$, and the productivities A_f and A_i. Note that A_f is not separately identified from μ_m and μ_w. In a cross-sectional model, μ_m and μ_w pin down the relative productivity of formal jobs for men and women, respectively. However, changes in A_f and A_i can capture productivity growth over time or productivity differences across space.

This leaves thirteen parameters plus two productivities, A_i and A_f, to calibrate. For now, we set exogenously the tail parameter of the Pareto distribution of norms θ, the elasticities σ, η, and

γ, and the indivisibility parameter $\bar{f}$. Specifically, in line with our empirical results reported in Figure 5 we fix $\bar{f}$ to 50% of non-leisure hours, which translates to seven hours. This leaves eight parameters that we need to calibrate.

Estimation for Rural India: Without loss of generality, we set $A_i = 1$ in rural India by a choice of units. Since A_f is not separately identified, for the reason explained above, we also set it to unity. We will allow the productivities A_i and A_f to change when we transition to urban India. In rural India, we abstract from labor market discrimination against women and set $\tau = 0$. This assumption is reasonable, since in rural areas, women typically work in the fields where income directly benefits the household.[10]

We then estimate the remaining eight parameters $\{p, \mu_m, \mu_w, \sigma_\epsilon, \rho, \bar{u}, \psi, \beta\}$ to match eleven empirical moments for rural India. Specifically, we target three moments on the extensive margin of labor supply, one moment on the intensive margin of labor supply, four moments pertaining to individuals' time use, and three moments related to earnings differentials, both across sectors and men and women. The eleven moments are:

- Extensive margin: (1) the share of men working in the formal sector, (2) the share of women working in the formal sector, (3) the share of women performing neither formal nor informal work (i.e. FLFP);

- Intensive margin: (4) the average hours men spend on informal work;

- Time use: (5) the average hours men spend on home production, (6) the average hours women spend on home production, conditional on not participating in the market, (7) the average hours women spend on home production, conditional on participating in the market, (8) the average hours men spend on leisure;

- Earnings: (9) the average earnings premium of formal relative to informal work for men, (10) the average earnings ratio between women and men working in the informal sector, (11) the average earnings ratio between women in formal work and men in informal work.

Even though all parameters are jointly estimated, there is an intuitive mapping between the moments and the structural parameters. In particular, the scarcity of formal jobs, captured by a low p (with $p = 1$ representing a frictionless environment), is crucial to explain the low share of both men and women working in formal labor markets (moments 1 and 2), given the high formal earnings premium. The parameters of the multivariate log-normal distribution of skills ($\mu_m, \mu_w, \sigma_\epsilon$,

[10] Allowing for labor market discrimination in the formal rural sector would not change any of the results. As we will see, the model predicts that women do not work in the formal sector even in the absence of any labor market discrimination.

Table 6: Structural Parameters and Model Fit: Rural India

Parameter	Estimate	Moment	Data	Model
p	0.33	Male formal employment share	0.09	0.08
		Female formal employment share	0.03	0.01
μ_m	-0.34	Avg earnings ratio, formal men/inf. men	1.70	1.72
μ_w	-1.06	Avg earnings ratio, formal women/inf. men	1.05	1.06
σ_ϵ	0.66	Avg earnings ratio, inf. women/inf. men	0.48	0.80
ρ	0.00	Avg informal hours, men	0.52	0.48
$\bar{u}$	0.06	Female share not working	0.50	0.57
ψ	1.25	Avg home prod. hours, men	0.05	0.10
		Avg home prod. hours, women out of LF	0.51	0.55
		Avg home prod. hours, women in the LF	0.36	0.20
β	0.71	Avg hours leisure, men	0.43	0.40

Notes: The table reports the estimated structural parameters (column 2), the target moments (column 3), and the moments in the data (column 4), and the calibrated model (column 5). The following parameters are set exogenously: $\theta = 2$, $\sigma = 1.5$, $\eta = 5$, $\gamma = 0.21$, and $\bar{f} = 0.5$. Furthermore, $A_f = A_i = 1$ as discussed in the text.

and ρ) together with the preference parameter ψ are then responsible for matching the average earnings ratios between formal and informal workers, and between men and women (moments 9—11) and the allocation of time (moments 4—7). Note that we separately target the average home production hours of women who are in and out of the labor force. Since almost all men work, applying the same distinction to men is unnecessary.

Conditional on the other parameters, the scale parameter of the distribution of norms, $\bar{u}$, ensures that the model matches the low participation rate of women (moment 3). The preference parameter β is determined by the time people spend on leisure (moment 8).

Genetic Algorithm: We estimate the vector $\{p, \mu_m, \mu_w, \sigma_\epsilon, \rho, \bar{u}, \psi, \beta\}$ using a genetic algorithm to minimize the Euclidean (weighted) norm of the difference between the target data moments and the model predictions. Moments pertaining to shares in the formal sector, share of women not working, and men's hours in home production are up-weighted.[11] In Table 6 we report the estimated structural parameters (column 2) and the target moments, both in the data (column 4) and the calibrated model (column 5).

Our model accurately matches the empirical target moments. In the estimated model, 57% of women stay out of the labor force, which is close to 50% in the data. 8% of men and 1% of women in rural areas engage in formal activities, closely aligning with the target moments of 9% and 3%, respectively. Women not in the labor force devote 55% of their time to home production, very near the empirical observation of 50%. However, the model is less successful in replicating the

[11]In practice, we assign a weight of 3 to the share of women and men in the formal sector; we assign a weight of 2 to the share of women out of the labor force; lastly, we assign a weight of 1.5 to men's mean home production hours.

hours spent on home production by *working* women. In the data, these women spend 36% of the non-sleeping time on home production. In the model, they only spend 20%. The reason for this discrepancy is that the model struggles to capture the large number of women working very few hours in the data and hence spending a substantial amount of time on home production. Many of these women work in subsistence agriculture, where the norm against women's doing any work is likely less relevant or fails to exhibit a less sharp discontinuity at zero hours. Note however, that even empirically, labor force participation reduces home production hours by almost 30% (0.36/0.51).

The model accurately predicts men's leisure time and closely reproduces the average earnings ratios. In particular, our model is consistent with the fact that, among men, average earnings are 70% higher in the formal sector (relative to the informal sector) and that women in formal jobs earn the same amount as men in informal occupation. The least accurate fit is the earnings ratio for women working in the informal sector relative to men. Empirically, women earn about half as much as men, while our model predicts that women earn 80% as much. Again, this is due to the difficulty of replicating the significant proportion of women in the data who work only a few hours and therefore have low earnings. Overall, the model achieves a very good fit with the target moments.

5.2 From Rural Fields to Urban Kitchens

Next, we use the model with preference and skill parameters estimated from rural India data to predict household behavior in urban India, allowing technology (A_f, A_i) and institutions (τ, $\bar{u}$, and p) to vary between the two environments. This exercise has two objectives. First, we test whether the model can capture key changes in labor supply behavior, particularly the decline in FLFP. Second, we aim to identify the relative importance of the driving forces of this transformation: stronger norms against female work in urban areas, labor market discrimination, the availability of formal jobs, and income effects. Finally, in Section 6, we assess whether further economic development will likely lead to a continued decline in FLFP or a reversal of this trend.

We fix all preference parameters to the values estimated in Table 6. Next, we separately estimate A_f, A_i, p, $\bar{u}$, and τ for urban India.[12] In simpler terms, we allow urban districts to differ from rural ones in the average productivity of formal and informal jobs, the availability of formal jobs, the intensity of norms against women's work, and labor market discrimination.

We anchor A_i to reflect the relative earnings differences between rural and urban India. Since earnings in informal activities are 71% higher in urban areas compared to rural ones, we set $A_i = 1.71$.[13] The remaining parameters are calibrated to match the same eleven target moments

[12] The parameter $\bar{u}$ reflects the stringency of norms, which could be considered part of the vector of preferences. Since norms may, in principle, be extrinsic, we treat them as part of the institutional framework.

[13] Because hours worked are endogenous, the actual earnings ratios will differ slightly. However, the difference

Table 7: Structural Parameters and Model Fit: Urban India

Parameter	Estimate	Moment	Data	Model
p	0.63	Male formal employment share	0.31	0.31
		Female formal employment share	0.06	0.02
μ_m	-0.34	Avg earnings ratio, formal men/inf. men	1.31	1.58
μ_w	-1.06	Avg earnings ratio, formal women/inf. men	1.09	0.98
σ_ϵ	0.66	Avg earnings ratio, inf. women/inf. men	0.43	0.46
ρ	0.00	Avg informal hours, men	0.59	0.55
$\bar{u}$	0.04	Female share not working	0.71	0.72
ψ	1.25	Avg home prod. hours, men	0.03	0.04
		Avg home prod. hours, women out of LF	0.48	0.54
		Avg home prod. hours, women in the LF	0.32	0.25
β	0.71	Avg hours leisure, men	0.38	0.40

Notes: The table reports the estimated structural parameters (column 2), the target moments (column 3), and the moments in the data (column 4), and the calibrated model (column 5). The following parameters are set exogenously: $\theta = 2$, $\sigma = 1.5$, $\eta = 5$, $\gamma = 0.21$, and $\bar{f} = 0.5$. Furthermore, we set $A_i = 1.71$ to match the earning premium of urban-to-rural in the informal sector and set $A_f = 2.21$ as discussed in the text.

as in Table 6, but for urban India. The results are:

1. $A_f = 2.21$ (compared to $A_f = 1$ in rural districts);

2. $p = 0.63$ (compared to $p = 0.33$ in rural districts);

3. $\bar{u} = 0.04$ (compared to $\bar{u} = 0.06$ in rural districts);

4. $\tau = 0.19$ (compared to $\tau = 0$ in rural districts).

These results are highly intuitive. In urban areas, formal jobs are more than twice as productive as in rural India. They are also much more abundant in urban districts, as shown by the sharp increase in p. The third significant change is the introduction of gender pay discrimination, represented by τ. Our estimate indicates that women are paid 19% less than observationally equivalent men, a disparity that, as we show below, plays a key role in explaining the decline in female labor participation. Conversely, our estimation reveals no significant increase in the social norm discouraging women from working. If anything, the parameter $\bar{u}$ slightly declines, suggesting somewhat more permissive social norms in urban areas.

Table 7 shows that the estimated model accurately captures the targeted moments for urban India. Most notably, the model predicts that the share of non-working women will be 74%, which is very close to the data moment of 71%, up from 57% (respectively, 50%) in rural India. The share of men employed in formal jobs rises to 31% (both in the model and data). The share of women

is small: in our calibrated model, the earnings ratio is 1.76. We use A_i as the anchor because selection effects are more pronounced in the formal sector, where individuals have heterogeneous productivities.

taking formal jobs is lower in the model than in the data. However, both the model and the data agree that the share of women in formal jobs has doubled compared to rural areas. The model continues to accurately capture men's leisure time and the home production hours of non-working women, though it slightly underpredicts the hours spent on home production by working women.

The model also captures the three average earnings ratios being targeted—even more accurately than in rural areas. It captures the gender gaps especially well: on average, women earn 46% of what men earn in informal jobs. In formal jobs, women earn roughly the same as men working informal jobs. However, the model predicts a higher premium for men working in formal jobs than the data suggest. Interestingly, both in the data and in the model, the earnings premium for working in the formal sector relative to the informal sector is *lower* in cities, despite the fact that productivity is estimated to increase more in formal than informal jobs. This result reflects changes in selection, whereby the expansion of the formal work force reduces the efficiency of the marginal formal worker.

The results in Table 7 are also instructive to understand how the model distinguishes changes in social norms ($\bar{u}$) from labor market discrimination (τ). If norms against women working were strengthening, we would expect to see many urban women rejecting well-paid jobs due to social constraints, leading to relatively higher earnings for women in urban areas compared to rural regions. However, the data show the opposite trend. Urban women are not working primarily due to a lack of sufficiently well-paid job opportunities combined with higher household incomes, rather than strong norms discouraging them from accepting well-paid positions.

In summary, the model successfully captures and explains the decline in FLFP as one moves from rural to urban areas. As in the data, the model predicts that this decline is largely driven by reduced informal employment, while the share of women in formal jobs doubles, though it remains very low overall. The model attributes these changes mainly to a combination of income effects and labor market discrimination. Although women in informal jobs are better paid in urban areas than in rural India, the increase is not sufficient to offset the income effect of higher wages earned by their husbands. The widening gender gap drives further specialization, ultimately increasing the overall supply of home production by women. In other words, it shifts many women from rural fields to urban kitchens.

To summarize the model's success to capture labor supply, Figure 8 presents a quantitative comparison between the model's predictions and the data, broken down by gender. Panel (a) highlights the differences in female labor supply between rural and urban India. While few women work in formal sectors in both areas, urban districts show higher formal participation. The main shift, however, is from informal employment to nonparticipation. In contrast, panel (b) reveals that men experience a significant shift from informal to formal employment in urban areas. Given that earnings in the formal sector are higher, this reallocation of men into the formal jobs contributes to womens' low participation through higher household earnings.

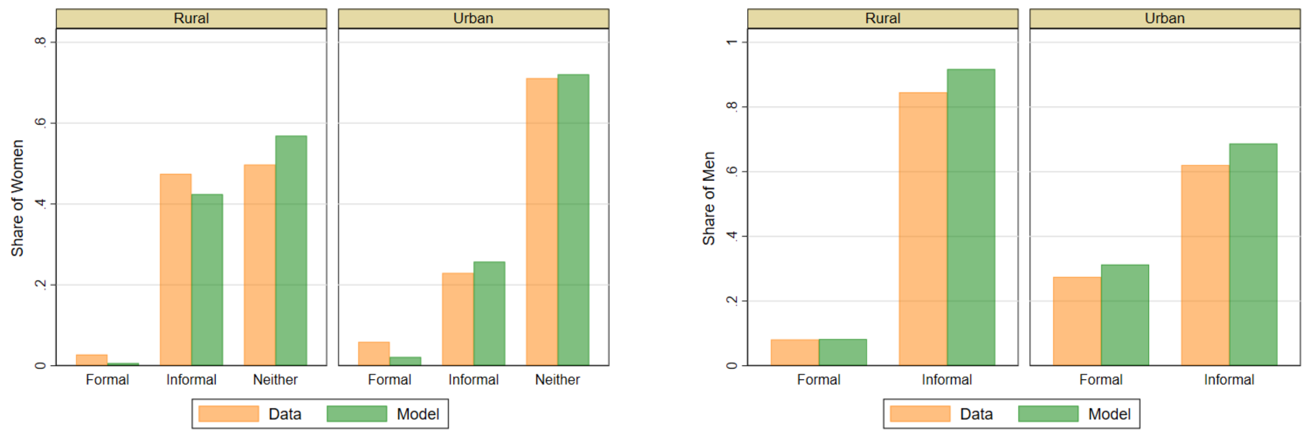

(a) Female Labor Supply

(b) Male Labor Supply

Notes: Panel (a) shows the breakdown of female labor supply across formal labor markets, informal labor markets, and nonparticipation in rural and urban India. Panel (b) shows the breakdown of male labor supply between formal and informal labor markets in both rural and urban India (we omit nonparticipation, as it accounts for a very small share of men's labor supply).

Figure 8: Female and Male Labor Supply in India: Data vs. Model

5.3 Nontargeted Moments: Income Effects

A key determinant of female labor supply in our model is the presence of income effects. Even though our estimation did not explicitly target the strength of such income effects, our model captures their importance quite well. Figure 9 reports the joint distribution of labor across spouses within households. Specifically, it shows the share of women working formally, informally, or not at all, as a function of the employment status of their husband — a proxy of spousal income. Panel (a) presents results for rural India, panel (b) focuses on urban India. Two key patterns emerge. First, a strong income effect: when the husband works in the formal sector (and earns more), the wife is much more likely to stay out of the labor market and focus on home production. Second, despite not targeting these conditional employment distributions within couples, our model successfully captures the main features of the relationship between women's labor supply and their husband's earnings. For urban labor markets, our model is quantitatively very close to what we see in the data. The same is true for rural locations, conditional on the husband working informally. For rural women with husbands working in formal jobs, our model predicts too little female participation. Note, however, that only 9% of all men work in formal jobs in rural areas and that, empirically, many women in informal jobs work few hours.

(a) Rural India

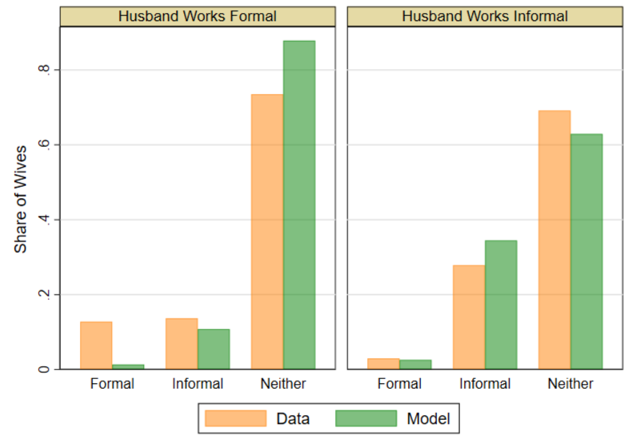

(b) Urban India

Notes: The figure shows the share of women doing informal, formal or no market work activity as a function of their husband's labor force status. In the left panel we report the distribution for rural India, while in the right panel we report the distribution for urban India. We always depict the data in light orange and the model in green.

Figure 9: Nontargeted Moments

5.4 Accounting for the Rural-Urban Female Labor Force Participation Gap

Tables 6 and 7 illustrate how our model accounts for the urban-rural participation gap. This gap is driven by a combination of three factors: higher productivity (A_i and A_f), better access to labor markets (p), and variations in norms and discrimination ($\bar{u}$ and τ). In the following section, we quantify the contribution of each of these forces using our model.

The results of this exercise are illustrated in Figure 10. In the left panel, we adjust the rural calibration by setting the respective variables—A_i, A_f, p τ, $\bar{u}$— *individually* to their urban counterparts. Hence, each bar represents the counterfactual share of women participating in market work if only that variable were to change. In the right panel, we change these variables *cumulatively*, thereby accounting for the full differences between urban and rural India.

Consider first the individual effects of the different variables, shown in the left panel. In the leftmost bar, we report the FLFP rate in rural India, which, in our model, is 43%. The remaining bars reflect changes in economic aspects (productivity and labor market frictions) shown in blue and changes in institutions (demand and supply distortions) shown in red. Two factors push female participation upward: higher informal wages ($A_i^{URBAN} > A_i^{RURAL}$) and more liberal gender norms ($\bar{u}^{URBAN} < \bar{u}^{RURAL}$), each increasing participation by about 30 percentage points. Conversely, higher formal wages ($A_f^{URBAN} > A_f^{RURAL}$), lower labor market frictions ($p^{URBAN} > p^{RURAL}$), and greater distortions for women ($\tau^{URBAN} > \tau^{RURAL}$) *decrease* participation. That higher distortions reduce women's incentives to participate in the labor market is expected. That this is also true for improvements in formal productivity and in the access to formal jobs is more surprising and a consequence of income effects: both raise spousal incomes and hence lower the marginal value of household consumption and female participation. Quantitatively, we find that higher formal productivity and better labor market access would reduce FLFP in rural India by around 3%-5% and that a 19% tax on women's marginal product would cause most rural women to exit the labor force entirely.

While the left panel shows the isolated impact of each factor, the right panel emphasizes potential complementarities by displaying these effects sequentially, starting from the rural calibration and gradually adjusting A_i, A_f, p, etc., to their urban levels. This sequential analysis confirms the earlier findings: informal productivity boosts participation, but more productive formal labor markets, through both higher A_f and lower frictions p, reduce it. Discrimination significantly curbs participation (causing female participation to fall below 5%), while more liberal gender norms ($\bar{u}$) help offset these forces. Overall, the urban-rural participation gap reflects the combination of these forces.

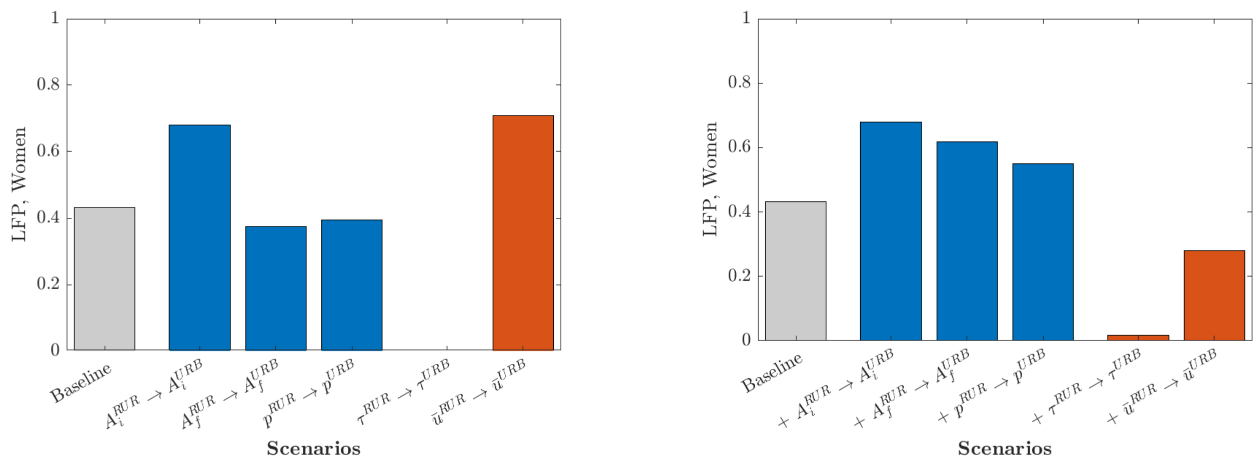

(a) Individual variables

(b) Sequential

Notes: In the left panel, we present the counterfactual female participation rate when each parameter is individually adjusted from its rural to urban value. In the right panel, parameters are adjusted sequentially: starting with A_i, followed by A_f, and so on. The last bar in the right panel thus reflects the female participation rate in our calibrated model.

Figure 10: Accounting for the Urban-Rural Participation Gap

6 The Future of FLFP in India

Our study was motivated by the observation that economic development in India has been accompanied by a decline in female labor force participation (FLFP). This trend, combined with the pronounced urban-rural participation gap, raises questions on whether future economic growth can lead to a reversal of this decline.

In this section, we use our model to run counterfactual experiments to forecast the future of FLFP in India. We are particularly interested in understanding whether economic growth per se (specifically, higher productivity and improvements in access to formal labor markets) can have significant effects on female participation in the near future or whether institutional and cultural changes are needed at the same time.

6.1 Economic Growth vs. Economic Development

To answer these questions we subject our calibrated model to different counterfactual scenarios about the future path of India's economy. We begin by considering counterfactual experiments that focus on purely economic variables. Future economic growth is likely to entail rapid wage growth and an increasing role for formal labor markets. We capture this evolution through two scenarios:

GROWTH1: Productivity in the formal sector (A_f) grows at an annual rate of 5% from 2020 to 2100.

GROWTH2: In addition to the growth in A_f, access to formal labor markets, parameterized by p, increases linearly from its estimated level ($p = 63\%$) in 2020 to $p = 100\%$ in 2100.

In both scenarios, the productivity of informal jobs is held constant.

Next, we consider scenarios where economic growth is accompanied by institutional and cultural changes:

DEVO1: Building on GROWTH2, pay discrimination against women declines. The discrimination parameter τ decreases linearly from 0.19 in 2020 to zero in 2100, allowing women to get paid their marginal product and hence receive equal pay for equal productivity.

DEVO2: In addition to pay discrimination, social norms against women joining the labor force also decline. We capture this gradual transformation by assuming that the scale parameter of the Pareto distribution $\bar{u}$ decreases linearly from 0.04 in 2020 to zero by 2100, progressively eliminating cultural barriers to female labor force participation.

We think of these scenarios as capturing the differences between economic growth and broad-based economic development.

6.2 Results

We summarize the results of these different scenarios in Figure 11. The upper panels show, respectively, the evolution of FLFP and that of working hours for participating women. The lower panels show the participation of women in formal and informal labor markets.

Economic Growth: We first consider the effects of economic growth, i.e. GROWTH1 and GROWTH2. The upper left panel shows that FLFP follows a U-shaped trajectory. Initially, increasing households' access to formal labor markets (GROWTH2) has a *negative* effect on the extensive margin of FLFP. As more husbands gain access to formal jobs and earn higher wages, their spouses drop out of *informal* labor markets (see lower right panel) and specialize in home production. FLFP continues to decline until 2036 due to income effects before beginning a slow recovery as substitution effects dominate. By 2100, even under GROWTH2, FLFP reaches just over 40%, which is lower than today's participation rate in rural India.[14] Hence, according to our model, economic growth alone is unlikely to substantially raise women's participation rate in the near future. Eventually, FLFP would rise to the level of today's advanced economies but the progress would be painfully slow.

The two lower panels decompose the overall participation rate into formal (panel c) and informal (panel d) activities. The discrepancy between the two growth scenarios reflects the presence of labor market frictions (p). This is particularly evident in women's participation in the informal sector, where nearly 15% of women continue to work informally despite the rise in formal productivity (recall that we assume A_i to be constant). This outcome is driven by the fact that we model labor market frictions at the household level. Thus, as long as $p < 1$, some households lack access to formal employment, and both partners work informally. Consequently, the decline in informal employment in the GROWTH2 scenario, particularly in the early decades, is primarily due to falling labor market frictions for the husband.

These patterns are also visible on the intensive margin. As shown in panel b, economic growth significantly increases the number of hours worked by participating women. This reflects the decline in informal employment, where most women work few hours. Once the informal sector disappears, all women, conditional on working, spend at least the minimum hours, $\bar{f}$, at work.

Economic Development: The results change significantly when economic growth is accompanied by institutional and cultural changes. Under DEVO1 (shown in dashed red), declining pay discrimination leads to a steady rise in FLFP, surpassing 40% by 2060 and 50% by 2090. Intuitively, lower discrimination encourages more women to enter the labor market. This is evident in the lower panels, which show an increase in participation in both formal and informal work relative to both growth scenarios. Since lower discrimination raises the marginal product of women's hours, it counteracts the growth-induced income effects. This is also why falling discrimination increases hours, conditional on working.

Finally, under DEVO2, we also allow for a reduction of cultural norms against women working.

[14]Appendix Figure B-1 shows that under scenario GROWTH2 FLFP would take until year 2180 to reach 70% if A_f continues to grow at 5% annually.

In this scenario, FLFP increases significantly faster than in all previous cases. By 2100, over 70% of women are participating in the labor force. Notably, there is an initial surge in informal employment before a gradual transition to formal jobs, indicating that social norms are an important barrier for women to access even informal jobs. This highlights the critical role of cultural shifts in shaping the future of FLFP in India.

Taking Stock: Our quantitative theory predicts that economic growth alone is unlikely to reverse the decline in FLFP in India without significant institutional and cultural changes. In the short horizon, productivity growth and improved access to formal jobs could even lead to a further decline in FLFP. The subsequent reversal is likely to be painfully slow. In contrast, addressing distortions and norms is expected to have a more significant impact, aligning with cross-country evidence from Agte et al. (2024). Policies targeting pay equity and reducing cultural barriers can significantly enhance female labor market participation, fostering a more inclusive workforce and equitable growth trajectory.

7 Conclusion

Structural change and urbanization have transformed the developing world in recent decades. In this paper, we examine the implications of these developments for women's roles in the workplace. Our analysis focuses on India, a country that has experienced fast income growth, a rapid shift from agriculture to services with limited industrialization, and substantial urbanization. However, these changes have not translated into an increase in FLFP: overall participation declined from 36% in 1987 to 27% in 2019. We also document a significant urban-rural participation gap, with FLFP being notably higher in rural labor markets than in urban areas.

To explain these findings and project the future of women's participation in labor markets, we propose a parsimonious model of household labor supply, calibrated using detailed time-use data from both rural and urban India. Our theory emphasizes three key factors that shape labor division within households. First, income pooling within households results in an income effect: higher spousal earnings reduce the marginal utility of income and, consequently, the incentive to work. Second, formal jobs often come with minimum hour requirements, pushing one household member toward specialization in home production. Third, women face potential distortions: they may either be underpaid relative to their marginal productivity ("demand distortions") or face barriers to market participation due to social norms ("supply distortions").

We calibrate our model to microdata on time use, capturing how households allocate time between formal and informal work, home production, and leisure. The model indicates that the low female participation rate in Indian cities is driven primarily by higher spousal incomes and demand distortions. Supply distortions, if anything, play a countervailing role, reflecting more

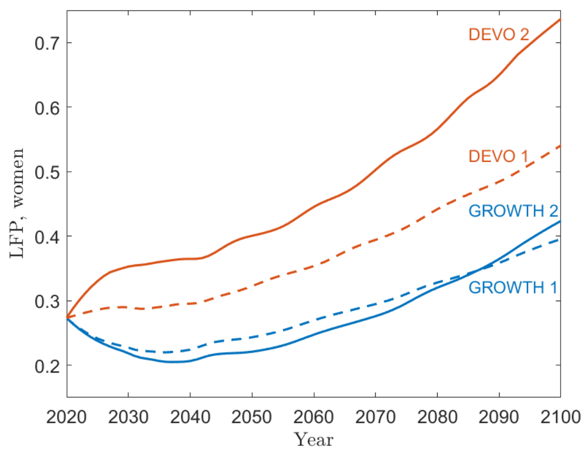

(a) Labor force participation

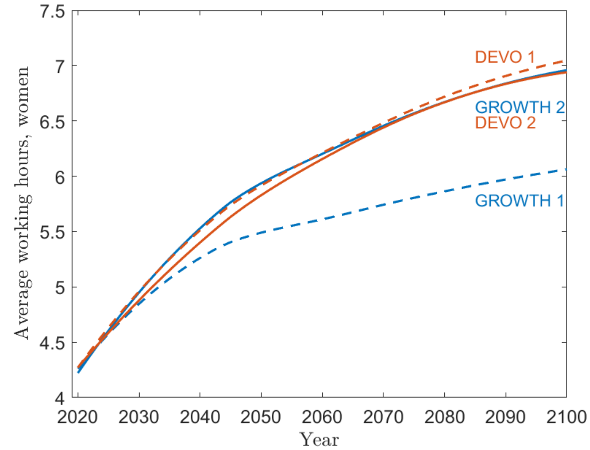

(b) Average hours worked | working

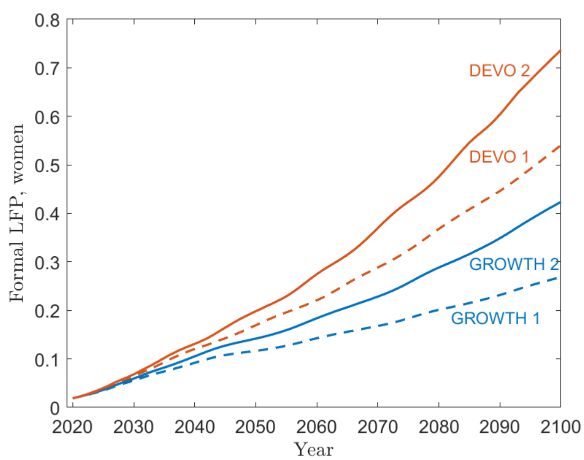

(c) Formal participation

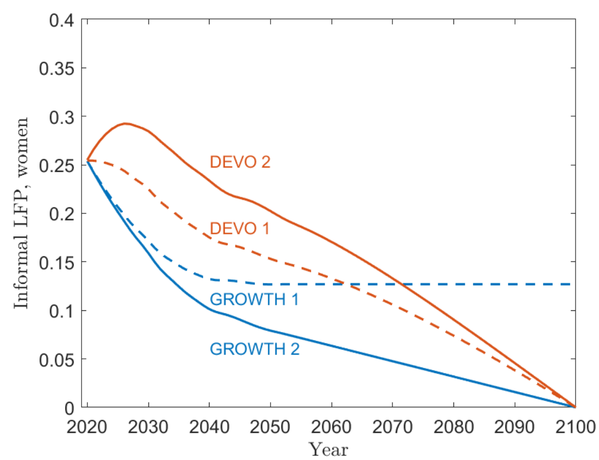

(d) Informal participation

Notes: Each panel shows the evolution of womens' labor supply in urban India under four scenarios: GROWTH 1 (blue - dashed), where A_f (formal hourly wage) grows by 5% yearly; GROWTH 2 (blue - solid), where in addition p (the probability that the household has access to formal jobs) increases linearly from 0.63 to 1; DEVO 1 (red - dashed), where growth is accompanied by a decline in the discrimination parameter τ from 0.19 to zero; DEVO 2 (red - solid), where additionally $\bar{u}$ (the scale parameter of the Pareto distribution for norms against women's work) declines from 0.04 to zero. Panel (a) shows female labor force participation (extensive margin). Panel (b) shows women's average hours worked conditional on working (intensive margin). Panels (c) and (d) show the proportion of women working in the formal and informal sector, respectively.

Figure 11: The Future of Female Labor Force Participation in India.

liberal attitudes in urban areas. Looking ahead, the model suggests that further productivity growth in formal jobs will have ambiguous effects on female participation and is unlikely to lead to substantial increases. Instead, reducing labor market discrimination and more liberal gender norms appear to be the most critical factor for boosting FLFP in India.

References

Afridi, F., M. Bishnu, and K. Mahajan: 2020, 'Gendering technological change: Evidence from agricultural mechanization'. *IZA Discussion Paper*.

Agte, P., O. Attanasio, G. Pinelopi, R. Aishwarya Lakshmi, R. Pande, M. Peters, C. T. Moore, and F. Zilibotti: 2024, 'Gender Gaps and Economic Growth: Why Haven't Women Won Globally (Yet)'. *Yale Working Paper*.

Aguiar, M. and E. Hurst: 2007, 'Life-cycle prices and production'. *American Economic Review* **97**(5), 1533–1559.

Boerma, J. and L. Karabarbounis: 2021, 'Inferring inequality with home production'. *Econometrica* **89**(5), 2517–2556.

Boserup, E.: 1970, *Woman's Role in Economic Development*. George Allen and Unwin, London.

Chiplunkar, G. and P. K. Goldberg: 2021, 'Aggregate Implications of Barriers to Female Entrepreneurship'. Working Paper 28486, National Bureau of Economic Research.

Chiplunkar, G. and T. Kleineberg: 2022, 'Gender Barriers, Structural Transformation, and Economic Development'. Working paper.

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.

Erosa, A., L. Fuster, G. Kambourov, and R. Rogerson: 2022, 'Hours, occupations, and gender differences in labor market outcomes'. *American Economic Journal: Macroeconomics* **14**(3), 543–590.

Fan, T., M. Peters, and F. Zilibotti: 2023, 'Growing Like India - the Unequal Effects of Service-Led Growth'. *Econometrica* **91**(4), 1457–1494.

Fernández, R.: 2013, 'Cultural Change as Learning: The Evolution of Female Labor Force Participation over a Century'. *American Economic Review* **103**(1), 472–500.

Fletcher, E. K., R. Pande, and C. T. Moore: 2019, 'Women and Work in India: Descriptive Evidence and a Review of Potential Policies'. In: *India Policy Forum*, Vol. 15. pp. 149–216.

Fogli, A. and L. Veldkamp: 2011, 'Nature or Nurture? Learning and the Geography of Female Labor Force Participation'. *Econometrica* **79**(4), 1103–1138.

Goldin, C.: 1994, 'The U-shaped female labor force function in economic development and economic history'. *National Bureau of Economic Research Working Paper* (4707).

Gottlieb, C., C. Doss, D. Gollin, and M. Poschke: 2024, 'The gender division of work across countries'. *IZA Discussion Paper*.

Greenwood, J., A. Seshadri, and M. Yorukoglu: 2005, 'Engines of Liberation'. *Review of Economic Studies* **72**(1), 109–133.

Hsieh, C.-T., E. Hurst, C. I. Jones, and P. J. Klenow: 2019, 'The Allocation of Talent and U.S. Economic Growth'. *Econometrica* **87**(5), 1439–1474.

Jayachandran, S.: 2015, 'The Roots of Gender Inequality in Developing Countries'. *Annual Review of Economics* **7**(1), 63–88.

Jayachandran, S.: 2020, 'Social Norms as a Barrier to Women's Employment in Developing Countries'. Working Paper 27449, National Bureau of Economic Research.

Klasen, S. and J. Pieters: 2015, 'What Explains the Stagnation of Female Labor Force Participation in Urban India?'. *The World Bank Economic Review* **29**(3), 449–478.

Li, N.: 2023, 'Women's work in India: Evidence from changes in time use between 1998 and 2019'. *World Development* **161**, 106107.

Munshi, K. and S. Singh: 2024, 'Social Status, Economic Development and Female Labor Force (Non) Participation'. Working Paper 32946, National Bureau of Economic Research.

Ngai, L. R., C. Olivetti, and B. Petrongolo: 2024, 'Gendered Change: 150 Years of Transformation in US Hours'. Working Paper 32475, 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.

Olivetti, C., J. Pan, and B. Petrongolo: 2024, 'The Evolution of Gender in the Labor Market'. Working Paper 33153, National Bureau of Economic Research.

Rendall, M.: 2018, 'Female Market Work, Tax Regimes, and the Rise of the Service Sector'. *Review of Economic Dynamics* **28**, 269–89.

APPENDIX A: EMPIRICAL RESULTS

In this section, we discuss our empirical analysis and the construction of the data in more detail.

A-0.1 ICATUS Activity Definition

Table A-I: International Classification of Activities for Time-Use Statistics 2016

Major division	*Division*	*Activity title*
1		**Employment and related activities**
	11	Employment in corporations, government and non-profit institutions
	12	Employment in household enterprises to produce goods
	13	Employment in households and household enterprises to provide services
	14	Ancillary activities and breaks related to employment
	15	Training and studies in relation to employment
	16	Seeking employment
	17	Setting up a business
	18	Travelling and commuting for employment
2		**Production of goods for own final use**
	21	Agriculture, forestry, fishing and mining for own final use
	22	Making and processing goods for own final use
	23	Construction activities for own final use
	24	Supplying water and fuel for own household or for own final use
	25	Travelling, moving, transporting or accompanying goods or persons related to own-use production of goods
3		**Unpaid domestic services for household and family members**
	31	Food and meals management and preparation
	32	Cleaning and maintaining of own dwelling and surroundings
	33	Do-it-yourself decoration, maintenance and repair
	34	Care and maintenance of textiles and footwear
	35	Household management for own final use
	36	Pet care
	37	Shopping for own household and family members
	38	Travelling, moving, transporting or accompanying goods or persons related to unpaid domestic services for household and family members
	39	Other unpaid domestic services for household and family members
4		**Unpaid caregiving services for household and family members**
	41	Childcare and instruction
	42	Care for dependent adults
	43	Help to non-dependent adult household and family members
	44	Travelling and accompanying goods or persons related to unpaid caregiving services for household and family members
	49	Other activities related to unpaid caregiving services for household and family members

5		**Unpaid volunteer, trainee and other unpaid work**
	51	Unpaid direct volunteering for other households
	52	Unpaid community- and organization-based volunteering
	53	Unpaid trainee work and related activities
	54	Travelling time related to unpaid volunteer, trainee and other unpaid work
	59	Other unpaid work activities
6		**Learning**
	61	Formal education
	62	Homework, being tutored, course review, research and activities related to formal education
	63	Additional study, non-formal education and courses
	64	Travelling time related to learning
	69	Other activities related to learning
7		**Socializing and communication, community participation and religious practice**
	71	Socializing and communication
	72	Participating in community cultural/social events
	73	Involvement in civic and related responsibilities
	74	Religious practices
	75	Travelling time related to socializing and communication, community participation and religious practice
	79	Other activities related to socializing and communication, community participation and religious practice
8		**Culture, leisure, mass media and sports practices**
	81	Attending/visiting cultural, entertainment and sports events/venues
	82	Cultural participation, hobbies, games and other pastime activities
	83	Sports participation and exercise, and related activities
	84	Mass media use
	85	Activities associated with reflecting, resting, relaxing
	86	Travelling time related to culture, leisure, mass media and sports practices
	89	Other activities related to culture, leisure, mass media and sports practices
9		**Self-care and maintenance**
	91	Sleep and related activities
	92	Eating and drinking
	93	Personal hygiene and care
	94	Receiving personal and health/medical care from others
	95	Travelling time related to self-care and maintenance activities
	99	Other self-care and maintenance activities

A-0.2 Constructing a Spatial Development Index

In order to examine how labor force participation varies with development, we need a measure of the latter across space in India. We construct a development index by running principal component analysis (PCA) on three standardized district-level variables: (i) urbanization rate, (ii) share of people working in agriculture, and (iii) average household consumption. We use the first principal component as our development index. We perform this exercise separately in each of our three data sets; Table A-II shows the factor loadings and variance explained for PC1, which are very similar across data sets. Reassuringly, PC1 captures urbanization, movement out of agriculture, and higher consumption – all indicators of development.

Table A-II: PCA1 Loadings and Variance Explained

	PLFS	Time Use
Std. Urban Share	0.59	0.59
Std. Agriculture Share	-0.59	-0.58
Std. HH Consumption	0.55	0.57
Variance Explained by PC1	71.51	70.71

Figure A-1 displays the correlation of the PC1 with standardized urban rates, agricultural shares, and log household consumption across districts. We observe a highly positive correlation of the First Principal Component with urban rates and mean household consumption (0.86 and 0.80, respectively). The correlation with the share of the labor force working in agriculture is -0.87. A similar picture can be depicted in TUS.

A-0.3 The Urban-Rural FLFP Gap Over Time

In this section, we extend our empirical analysis by using data on Labor Force Participation in India between 1987 and 2022. We combine data on five different waves from NSS, corresponding to the years 1987, 1999, 2004, 2007, and 2011, with yearly data from 2017 to 2022 from PLFS. First of all, we use a model identical to Equation (1), and run regressions for these 12 years to analyze the evolution of the urban coefficient, β. In this model, the coefficient β represents the average difference in employment between urban and rural individuals, controlling for state-specific effects, economic development at the district level, the employment share of services, and individual characteristics.

Figure A-2 plots the coefficients for each year. For each year, the urban coefficient is significantly negative. Around 2009, we observe a decrease in the magnitude of the coefficient, suggesting a decrease in the gap in FLFP between urban and rural areas. However, the size of the coefficient becomes larger again after 2018, almost returning to its original value by 2021. Overall, we observe that the gap between rural and urban areas has been rather stable over time.

A-0.4 Urban-Rural Differences in Labor Supply

In our main analysis we highlight the urban-rural differences in female labor supply. In Table A-III we show that this participation gap is indeed a phenomenon of women and not for the entire population. In columns 1 and 2 of Table A-III we replicate our main analysis for men. While there is indeed a negative effect of living in an urban area, the effect is one order of magnitude smaller than for women. In columns 5 and 6 we run our regression on the full sample and allow for an interaction between Urbanization and a female dummy. This coefficient is statistically negative and economically large.

This is also consistent with Figure A-3, where we depict the participation rate among men. This rate is remarkably stable and almost identical between rural and urban areas.

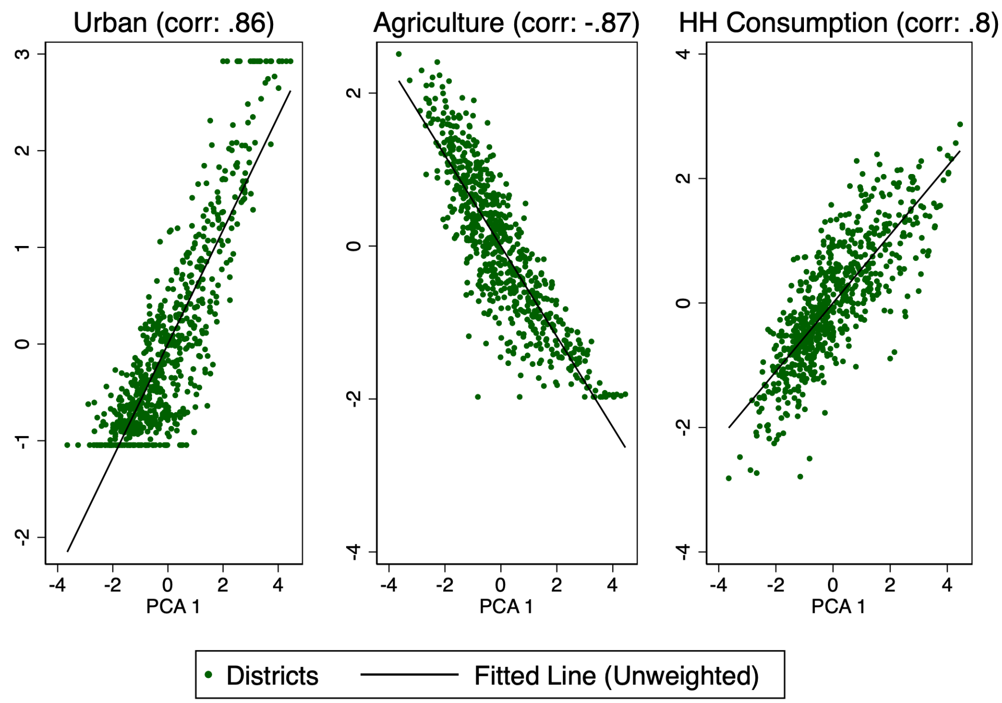

Source: 2019, PLFS. Calculated by the authors.
Notes: The figure illustrates the correlation between three standardized development variables: urban share, agriculture share, the logarithm of household consumption, and the first component of Principal Component Analysis (PCA). Each scatter point represents a different district, and the solid line represents the linear fitted line.

Figure A-1: Correlation of Development Variables in PLFS and PCA 1

Table A-III: The Urban-Rural Gap in Male and Female Labor Force Participation

	Men		Women		All	
	2019	2017-2022	2019	2017-2022	2019	2017-2022
	(1)	(2)	(3)	(4)	(5)	(6)
Urban (indiv)=1	-0.0180***	-0.0167***	-0.102***	-0.110***	-0.0187***	-0.0183***
	(0.00247)	(0.00126)	(0.0216)	(0.0206)	(0.00502)	(0.00508)
Development Index	-0.00513**	-0.00348***	-0.0180**	-0.0190***	-0.00695	-0.00672
	(0.00226)	(0.000471)	(0.00665)	(0.00626)	(0.00648)	(0.00645)
(Development Index)^2	0.000899**	0.000813***	0.00620*	0.00545**	0.00129	0.00127
	(0.000382)	(0.000174)	(0.00321)	(0.00263)	(0.00178)	(0.00179)
Service share (nonag.)	0.00134	-0.00406	-0.0152	-0.102**	-0.110***	-0.111***
	(0.00724)	(0.00397)	(0.0580)	(0.0436)	(0.0390)	(0.0389)
female=1					-0.663***	-0.663***
					(0.0693)	(0.0694)
Urban (indiv)=1 × female=1					-0.0890***	-0.0893***
					(0.0226)	(0.0226)
female=1 × Development Index					-0.00859	-0.00869
					(0.0143)	(0.0143)
female=1 × (Development Index)^2					0.00379	0.00381
					(0.00431)	(0.00430)
female=1 × Service share (nonag.)					0.110	0.110
					(0.0913)	(0.0914)
Constant	0.959***	0.965***	0.360***	0.431***	1.030***	1.030***
	(0.00491)	(0.00248)	(0.0398)	(0.0289)	(0.0270)	(0.0269)
State FE	Yes	Yes	Yes	Yes	Yes	Yes
Indiv. Covs.	Yes	Yes	Yes	Yes	Yes	Yes
Household FE	No	No	No	No	No	No
Year FE	No	Yes	No	Yes	No	Yes
Month FE	No	Yes	No	Yes	No	Yes
R^2	0.0570	0.0574	0.139	0.137	0.471	0.471
N	100036	555198	102351	565253	1122824	1120451

Source: 2017-2022 PLFS. Calculated by the authors.
Notes: The age range is limited to males between 25 and 60 years old. District-level clustered standard errors are in parentheses.

Figure A-2: The Urban-Rural Gap in Female Labor Force Participation, Independent Regressions

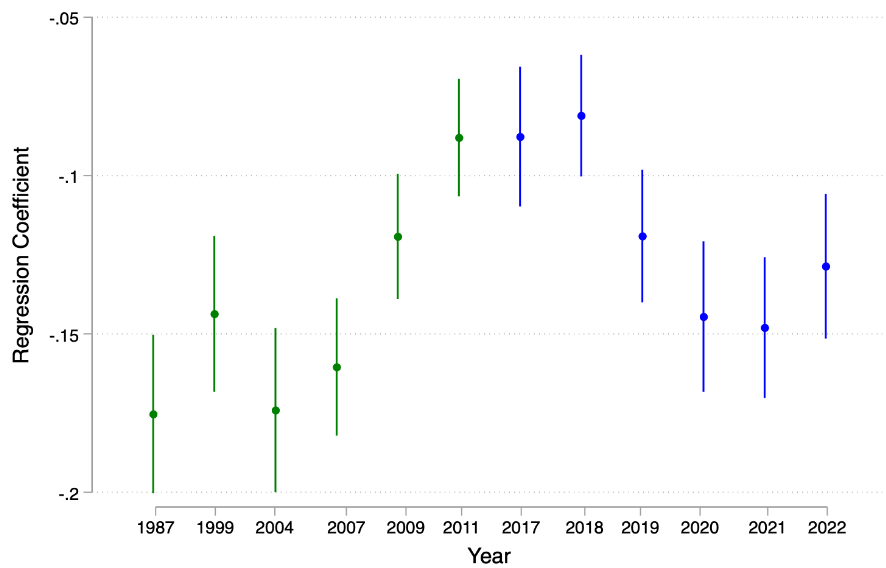

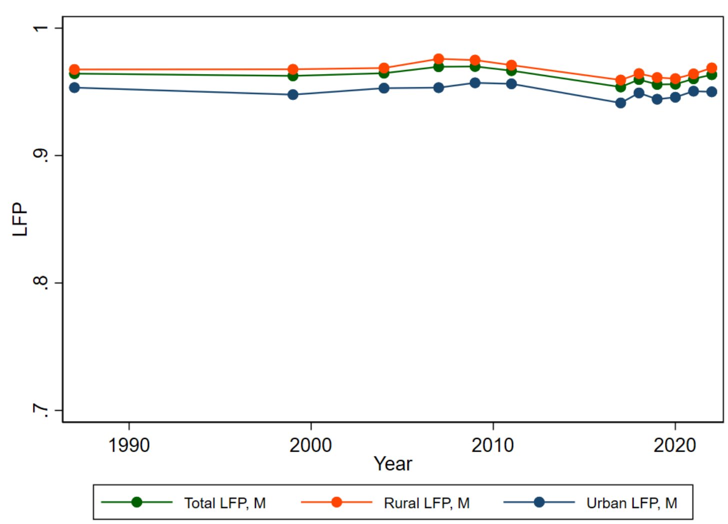

Notes: This figure depicts the participation rate among men. The age range is limited to individuals between 25 and 60 years old.

Figure A-3: Trends in India

A-0.5 Time Use

In this section we present some results based on time use data. In Figure A-4 we report the share of people that report either being in the labor force (FLP) or performing domestic duties as their principal activity as a function of whether they work in the formal sector, in the informal sector, or neither. Almost every man claims he is in the labor force no matter what job he does, while more than half of women do not report they are in the labor force even if they do some informal work. This gender difference is more clear in the group that does neither formal nor informal work.

In Figure A-5 we depict the intensive margin, that is the number of hours individuals spend

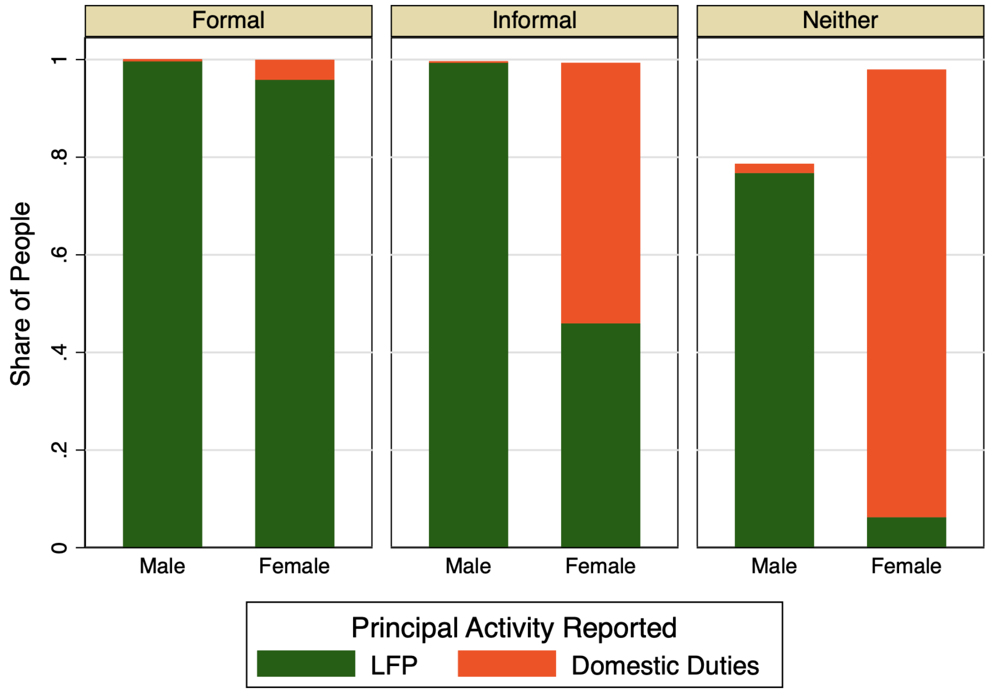

Notes: The figure illustrates how different genders report their principal activity based on their employment status.

Figure A-4: Gender Difference in Principle Activity Reported by Employment Status

on either work activities, leisure, or home production. Sleep and education are excluded from individuals' total hours in a day. Women spend substantially more time on home production irrespective of their work status and enjoy, on average, less leisure.

Finally, in Figure A-6, we depict the distribution of hours for men and women as a box plot. For men, the distribution of hours is quite similar irrespective of their sector of employment. For women, there is a pronounced difference in that the number of hours is consierably smaller in the informal sector.

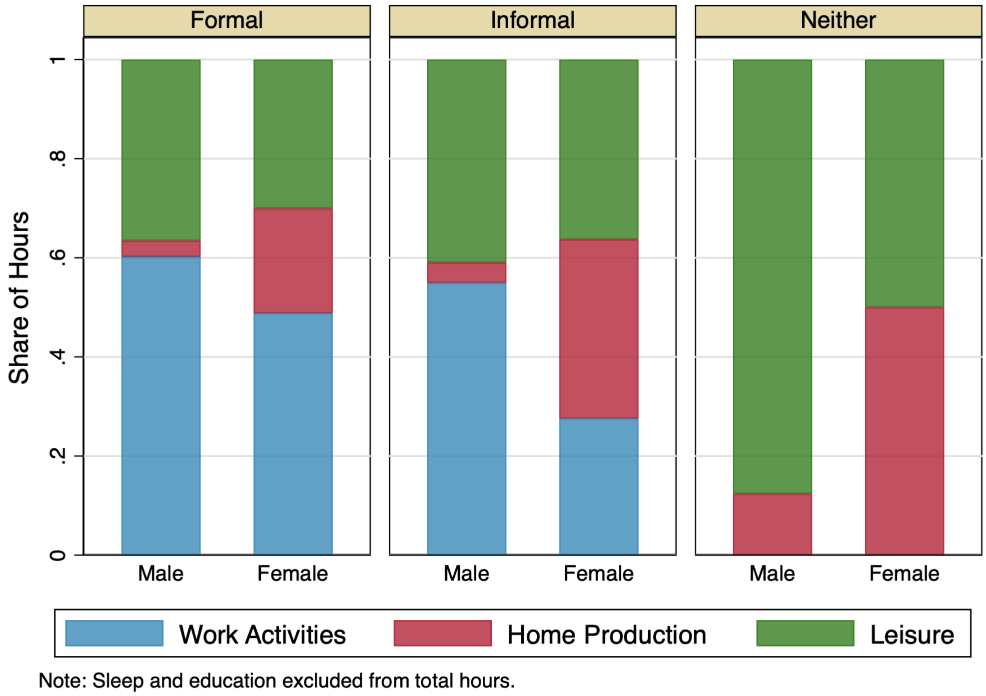

Notes: Sleep and education are excluded from individuals' total hours in a day. The figure illustrates the time allocation patterns of different genders based on their employment status.

Figure A-5: Gender Difference in Hours Spent w.r.t. Employment Status

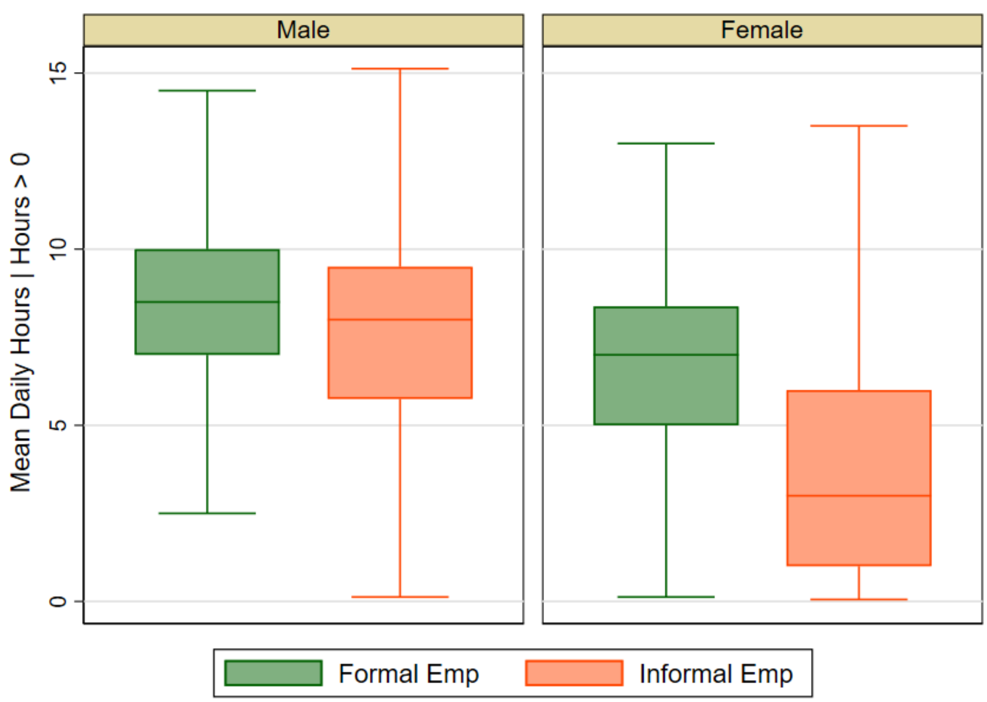

Source: 2019, TUS. Calculated by the authors.

Figure A-6: Distribution of Work Hours Spent w.r.t. Employment Status

APPENDIX B: ADDITIONAL MODEL RESULTS

B-1 The Future of FLFP in India: Extended

This section extends the exercise in Section , where we use our model to run counterfactual experiments to forecast the future of FLFP in India. The four scenarios are the same as before, but we allow A_f to grow beyond 2100. The aim of this extended exercise is to display how much longer growth alone takes to achieve a large FLFP. Figure B-1 shows the results.

GROWTH1: Productivity in the formal sector (A_f) grows at an annual rate of 5% from 2020 to 2200.

GROWTH2: In addition to the growth in A_f, access to formal labor markets, parameterized by p, increases linearly from its estimated level ($p = 63\%$) in 2020 to $p = 100\%$ in 2100. p remains at 100% thereafter.

In both scenarios, the productivity of informal jobs is held constant.

DEVO1: Building on GROWTH2, pay discrimination against women declines. The discrimination parameter τ decreases linearly from 0.19 in 2020 to zero in 2100, allowing women to receive equal wages for equal productivity. τ remains at zero thereafter.

DEVO2: In addition to declining pay discrimination, social norms against women working ($\bar{u}$) diminish linearly from 0.04 in 2020 to zero by 2100, eliminating cultural barriers to female labor force participation. $\bar{u}$ remains at zero thereafter.

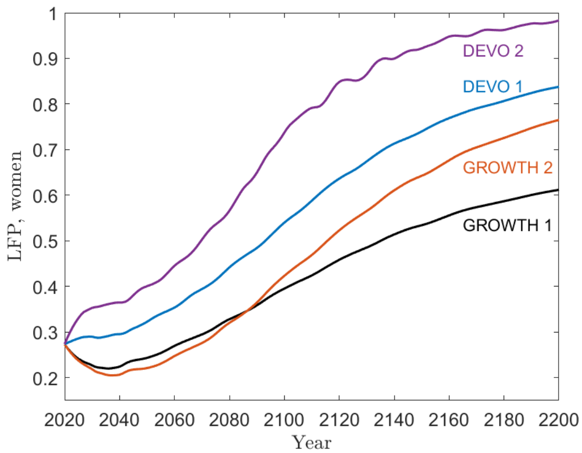

(a) Labor force participation

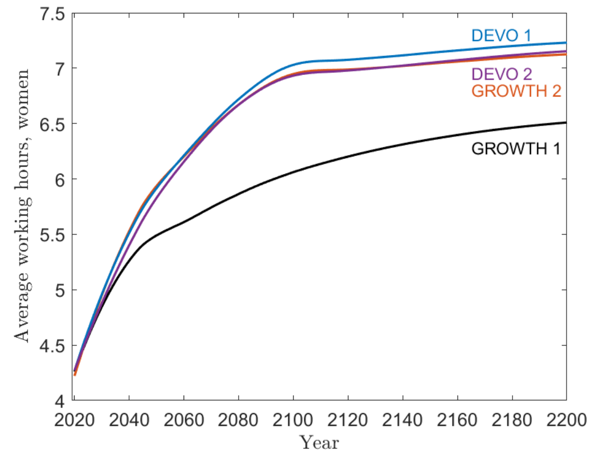

(b) Average hours worked | working

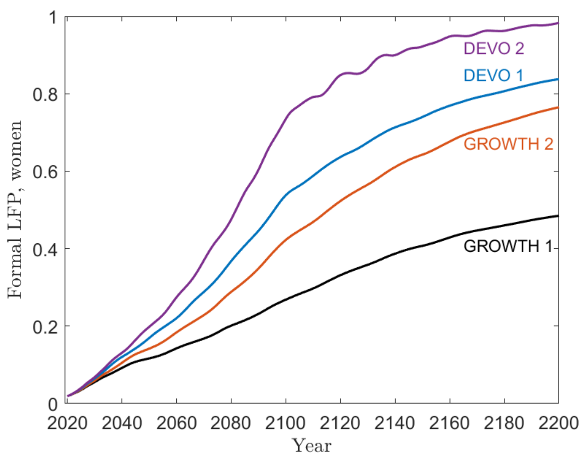

(c) Formal participation

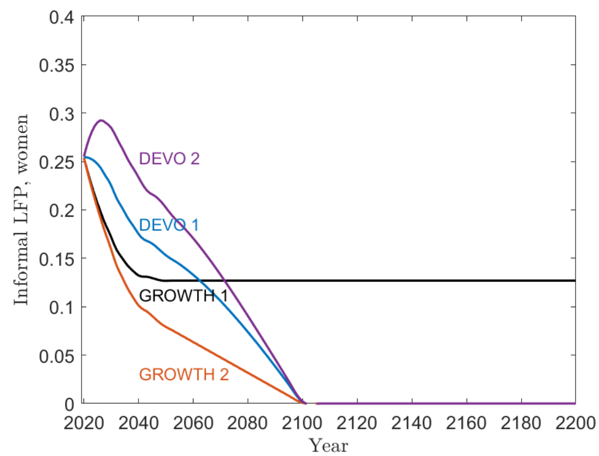

(d) Informal participation

Figure B-1

Notes: Each panel shows the evolution of womens' labor force involvement in urban India under four scenarios: GROWTH 1 (black), where A_f (formal hourly wage) grows by 5% yearly; GROWTH 2 (red), where in addition p | access to formal work | increases linearly from 0.63 to 1 in 2100 (and 1 thereafter); DEVO 1 (blue), where growth is accompanied by a decline in the discrimination parameter τ from 0.19 to zero in 2100; DEVO 2 (purple), where additionally $\bar{u}$ declines from 0.04 to zero in 2100. Panel (a) shows female labor force participation. Panel (b) shows women's average hours worked conditional on working. Panels (c) and (d) show the proportion of women working in the formal and informal sector, respectively.