Structural Change in Labor Supply and Cross-Country Differences in Hours Worked

Alexander Bick, Arizona State University
Nicola Fuchs-Schündeln, Goethe University Frankfurt and CEPR
David Lagakos, Boston University, CEPR and NBER
Hitoshi Tsujiyama, Goethe University Frankfurt

July 23, 2021

Abstract

This paper studies how structural change in labor supply along the development spectrum shapes cross-country differences in hours worked. We emphasize two main forces: sectoral reallocation from self-employment to wage work, and declining fixed costs of wage work. We show that these forces are crucial for understanding how the extensive margin (the employment rate) and intensive margin (hours per worker) of aggregate hours worked vary with income per capita. To do so we build and estimate a quantitative model of labor supply featuring a traditional self-employment sector and a modern wage-employment sector. When estimated to match cross-country data, the model predicts that sectoral reallocation explains more than half of the total hours decrease at lower levels of development. Declining fixed costs drive the rise in employment rates at higher levels of income per capita, and imply higher hours in the future, in contrast to the lower hours resulting from income effects and expansions in tax-and-transfer systems.

Email: alexander.bick@asu.edu, fuchs@wiwi.uni-frankfurt.de, lagakos@bu.edu, and hitoshi.tsujiyama@hof.uni-frankfurt.de. We thank Isaac Otoo and Paul Reimers for outstanding research assistance. For helpful comments and assistance we thank Timo Boppart, Peter Egger, Per Krusell, Kjetil Storesletten, Nora Strecker, Stephen Terry, Johanna Wallenius, and audiences at Athens, BI Norwegian Business School, Bocconi, Bonn, the CEPR Annual Macro Conference, Columbia, CREI, Florida, Goethe University Frankfurt, Hitotsubashi University, the IFS, the IIES Stockholm, Johns Hopkins SAIS, Melbourne, Monash, MOVE Barcelona, Oslo, Queensland, the Reserve Bank of Australia, the SED Annual Meetings, Sciences Po, Stockholm School of Economics, the T2M conference, University College Dublin, Vienna, Western Australia, and Zurich. Fuchs-Schündeln gratefully acknowledges financial support from the European Research Council under the ERC Consolidator Grant No. 815378 and from the DFG under the Leibniz-Preis. All potential errors are our own.

1. Introduction

Patterns of hours worked vary substantially with average income levels. Both time-series evidence and cross-country data point to a pattern of lower hours per adult in economies with higher income per capita (Ramey and Francis, 2009; Bick, Fuchs-Schündeln, and Lagakos, 2018; Boppart and Krusell, 2020). Underlying this decrease are stark differences in the behavior of the extensive margin (employment rates) and intensive margins (hours per worker) of aggregate hours worked over the development spectrum. Employment rates decrease sharply from low to middle-income levels, and then rise between middle- and high-income levels. Hours per worker increase between poor and middle-income countries, and then decline sharply between middle- and high-income countries. In other words, employment rates are convex and hours per worker are concave in income per capita.

The literature has focused on two broad explanations for why aggregate hours worked vary across countries. The first is preferences in which income effects dominate substitution effects in the long-run, leading desired hours to fall in response to permanent increases in wages (see Boppart and Krusell, 2020, and the references therein). The second is differences in tax-and-transfer systems, which are the leading explanation of why adults work fewer hours in Europe than in the United States (see e.g. Prescott, 2004; Rogerson, 2008; Ohanian, Raffo, and Rogerson, 2008; McDaniel, 2011; Bick and Fuchs-Schündeln, 2018). The fact that richer countries have more extensive tax-and-transfer systems than poorer countries on average (Besley and Persson, 2014) suggests that tax system differences may also help explain why adults in richer countries tend to work less. Yet, neither explanation seems well equipped to explain the convex behavior of employment rates and concave patterns of hours per worker with income per capita. If strong income effects and more distortionary tax systems reduce incentives to supply labor as income levels rise, why would employment rates fall strongly and then later rise with development, while hours per worker rise and then fall strongly?

In this paper we study a new explanation of why average hours worked vary across countries, which we refer to as *structural change in labor supply*. We emphasize two main forces. The first is the sectoral reallocation from self-employment to wage work that has long been studied in development economics (Lewis, 1954; Kuznets, 1966; Gollin, 2008). A key difference between self-employment and wage work is that wage work usually entails job search and travel costs, which are modeled as "fixed costs" by a

large literature (see e.g. Keane, 2011; Rogerson and Wallenius, 2013). Our insight is that movements into market wage work, with its higher fixed costs, must then reduce labor supply along the extensive margin. At the same time, this sectoral reallocation raises hours per worker, which are higher in wage work than in self-employment. The second force is fixed costs of wage work that fall with development, due to decreases in costs of job search, as emphasized by Poschke (2019), Martellini and Menzio (2020), and Donovan and Schoellman (2021), or declining transportation costs with development, as suggested by e.g. Adamopoulos (2011) and Bryan, Glaeser, and Tsivanidis (2020). To the extent that fixed costs fall as GDP per capita grows, this may increase labor supply along the extensive margin.

We formalize these features in a two-sector model with heterogeneous households and an intensive and extensive margin of labor supply. Sectoral reallocation from self-employment to market wage work is modeled as movements from a "traditional sector," with no fixed costs of labor supply but decreasing returns to labor supply, into a modern sector, with fixed costs but constant returns to labor supply. Decreasing returns to hours worked in the traditional sector are meant to capture seasonal variation in labor productivity in agriculture or the small scale of farms or family businesses (Paxson, 1993; Adamopoulos and Restuccia, 2014). This feature leads to lower average hours per worker in the traditional sector in equilibrium. Thus, sectoral reallocation from the traditional to the modern sector increases aggregate labor supply along the intensive margin, but decreases aggregate labor supply along the extensive margin due to the positive fixed costs of work in the modern sector. The model's declining fixed costs with development raise employment rates for those already in the modern sector. At the same time, decreasing fixed costs also drive sectoral reallocation by drawing more households into the modern sector, which serves to lower employment rates. In richer countries, where most workers are in the modern sector, the former force dominates, and the net effect of declining fixed costs is to raise employment rates.

In addition to these new forces, the model features the two standard channels studied in the literature: income effects in labor supply that dominate substitution effects, and distortionary tax-and-transfer systems. Preferences follow the MaCurdy (1981) specification, which is a special case of the preferences of Boppart and Krusell (2020), and used by other analyses of aggregate labor supply (e.g. Heathcote, Storesletten, and Violante, 2014). Tax-and-transfer systems follow the literature by featuring marginal taxation on

labor, with an arbitrary degree of progressivity, as well as consumption taxation and lump-sum transfers. We allow all of these features of the tax system to vary with an economy's income level.

We estimate the model to match patterns of labor supply in poor and rich countries drawing on the database of Bick et al. (2018). To parameterize tax-and-transfer systems, we rely on the database of gross income distributions and the associated net incomes based on statutory labor income tax codes assembled by Egger, Nigai, and Strecker (2019), which covers a wide range of countries of all income levels. We show as novel evidence that these data imply increasing progressivity and increasing labor income taxation between middle-income and rich countries. The model parameters are estimated to match differences between the average poor and rich country in the extensive and intensive margins of labor supply, as well as traditional sector shares and average incomes. We show that our model parameters are disciplined by patterns of structural change as well as more standard moments of the data.

The estimated model does well in replicating the convex pattern of employment rates in income per capita, in addition to the concave pattern of hours per worker. Sectoral reallocation in the lower half of the development spectrum leads to rapidly decreasing employment rates, but slightly increasing hours per worker. By contrast, the decreasing fixed costs of work in the modern sector serve to raise employment rates in the richer countries. In counterfactual simulations, we show that both having fixed costs of work apply only in the modern sector, as well as allowing them to change by development, are crucial features to match the two margins of labor supply.

We then use the estimated model to quantify the relative role of the driving forces in the decrease in hours per adult. The model predicts that sectoral reallocation explains more than half of the overall decrease in hours in the bottom half of the world income distribution, matching the importance of income effects. For higher income levels the role of sectoral reallocation is greatly reduced, but taxes and transfers gain in importance as a driver of the hours decrease as countries grow richer. Declining fixed costs of work become a strong countervailing force to the decrease in hours in the upper half of the development spectrum.

We conclude by illustrating our model's implications for the future of hours worked. Many since Keynes (1930) have argued that hours worked will continue to fall in the future as strong income effects make individuals reluctant to supply labor, or as tax-

and-transfer systems continue to expand. Our study suggests that if economies continue to witness decreases in the fixed costs of work—such as provided by internet platforms or through the recently observed increase in work from home—it is possible that hours will not fall as one might expect; they may even rise. As an illustration, continuing our observed trends out into the future would lead to roughly *constant* average hours worked for the next sixty years, meaning that decreasing fixed costs largely offset the decreases in hours implied by income effects and expanded tax systems. This exercise highlights the importance of distinguishing between structural change in labor supply and the latter forces as determinants of cross-country differences in hours worked.

Our paper builds on a growing literature that tries to understand the patterns of structural change that underlie the process of economic growth and development. Most of this literature has focused on the movement of workers from agriculture into industry and then services (see e.g. Buera and Kaboski, 2012; Herrendorf, Rogerson, and Valentinyi, 2014; Storesletten, Zhao, and Zilibotti, 2019; Alder, Boppart, and Müller, 2021). Dividing the economy into these three sectors is not important for our arguments, though, and in practice much of the work within these three sectors is subsistence self-employment in poor countries (Gollin, 2008). The movement of workers from subsistence self-employment to market work is one of the most salient features of development, though it has not been incorporated so far into the literature about the determinants of aggregate hours worked. Ngai and Pissarides (2008) and Bridgman, Duernecker, and Herrendorf (2018) distinguish between market and non-market work, though neither paper tries to account for extensive and intensive margins of aggregate labor supply.[1]

The rest of this paper is structured as follows. Section 2 presents the cross-country facts about hours worked and self-employment. Section 3 presents a model of structural change in labor supply. Section 4 estimates the model, and Section 5 illustrates how structural change in labor supply is important for matching the cross-country data. Section 6 assesses the quantitative importance of the different drivers of cross-country differences in hours worked per adult, and Section 7 turns to predictions for future hours as countries grow richer. Section 8 concludes.

[1]The same is true of Vandenbroucke (2009) and Kopytov, Roussanov, and Taschereau-Dumouchel (2021), who focus on the decrease in the relative price of leisure goods over time, and Restuccia and Vandenbroucke (2013) and Cervelatti and Sunde (2013), who focus more on the effect of the rise in life expectancy on hours worked.

2. Patterns of Labor Supply Across Countries

In this section, we review the main aggregate facts of labor supply across countries, which come from Bick et al. (2018), focusing on the extensive and intensive margins of labor supply and how they vary with income per capita. We then present some new evidence on average hours per worker by sector.

The data underlying the facts presented in the figures and tables of this section come from labor force surveys from 48 countries worldwide, as detailed in Bick et al. (2018). We use the core countries from that study, which have data that are best internationally comparable, but omit Laos, since it lacks data on education. The measure of hours used includes all activities producing output counted in NIPA. These include hours worked in wage employment as well as hours in own-account agricultural or non-agricultural work, whether or not that output is sold or used for own consumption, see Gollin (2008). Not included in our definition of hours worked are hours spent on non-market services, such as cleaning or home-provided child care.

2.1. Patterns of Labor Supply Across Countries: A Review

Figure 1 shows the key fact of decreasing hours worked per adult (15 years or older) over the development spectrum. The vertical lines divide countries into those belonging to the poorest, middle, or richest tercile of the world income distribution according to the Penn World Tables. Hours per adult decrease by 9.1 hours between the average poor and rich country, with most of the decrease taking place between poor and middle-income countries. In Bick et al. (2018), we show that this decrease is not driven by compositional effects.

Underlying the aggregate hours worked data are the extensive and intensive margins of labor supply. Panel A of Figure 2 plots the employment rates (the extensive margin) for men and women separately against GDP per adult. For both men and women, employment rates fall strongly between low- and middle-income countries, but then slightly increase towards the high-income countries. Taking averages across country-income groups and genders, employment rates fall by 22.1 percentage points from the low- to middle-income countries, and then rise by 2.2 percentage points between the middle- and high-income countries. Women have lower employment rates than men on average, but both genders exhibit remarkably similar patterns across the distribution of income per capita, with declines at low income levels and increases afterwards. Thus,

Figure 1: Average Hours Worked per Adult Across Countries

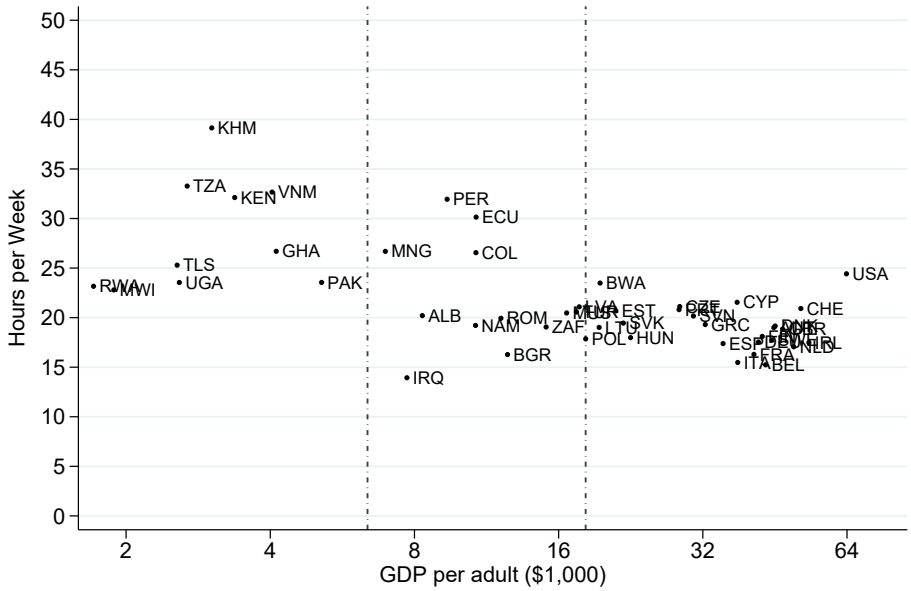

Note: This figure plots average weekly hours worked per adult by GDP per adult in thousands of international dollars. The vertical lines represent the division between low- and middle-income countries, and between middle- and high-income countries. Data source: Bick et al. (2018).

in the cross-country data, the uptick in employment rates is not particular to women, as has been posited by Goldin (1995) and others, but a more general feature of how labor supply evolves with development.

Panel B of Figure 2 plots hours per worker (the intensive margin) for men and for women across countries. Unlike employment rates, hours per worker show a slight *increase* between low- and middle-income countries, but then fall substantially between middle- and high-income countries. For both sexes, the decline is by 5.5 hours per week between the middle- and high-income countries. This concave pattern of hours per worker is evident for both men and women, with changes of remarkably similar magnitude across different levels of income per capita. Appendix Tables A.1 and A.2 show that the patterns of labor supply are not only the same for both genders, but also for different age groups.

Figure 2: Employment Rates and Hours per Worker Across Countries

Panel A: Employment Rates

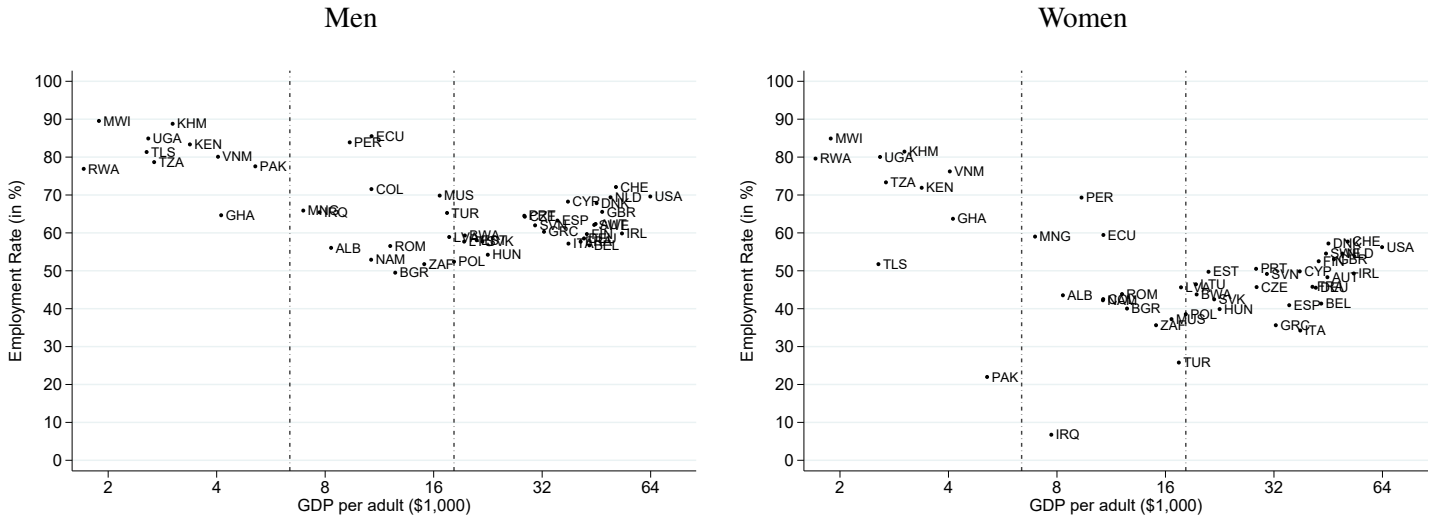

Panel B: Hours per Worker

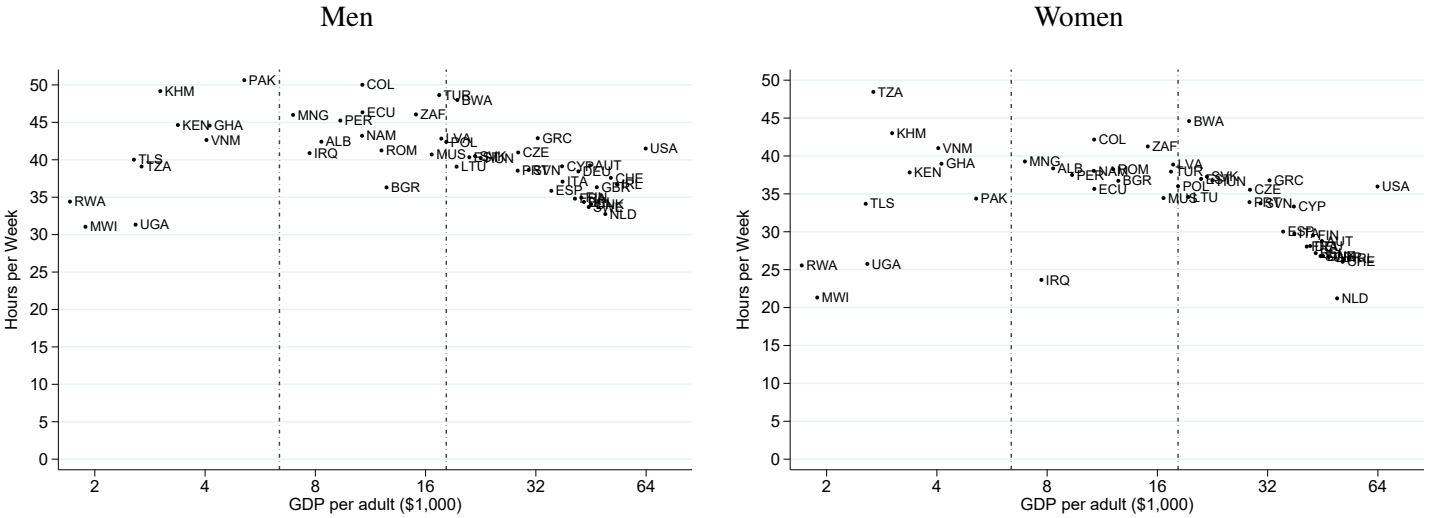

Note: This figure plots employment rates (in Panel A) and hours per worker (in Panel B) by GDP per adult in thousands of international dollars. The vertical lines represent the division between low- and middle-income countries, and between middle- and high-income countries. The left-hand panels are for men, and the right-hand panels are for women. Data source: Bick et al. (2018).

2.2. Patterns of Labor Supply Across Countries: Sectoral Facts

We turn next to sectoral facts, where we distinguish between a sector characterized by subsistence self-employment (traditional sector) and one characterized by wage employment (modern sector). Our empirical proxy for individuals working in the traditional sector is self-employed individuals with low education and unpaid family workers with low education. This is a close proxy for working in subsistence self-employment, though certainly not exact, but something we can measure in a comparable way across the countries in our data.[2]

Figure 3a shows the share of all workers working in the traditional sector: In the poor countries, almost two thirds of workers (64 percent) work in the traditional sector. This share rapidly decreases to 19 percent in the middle-income countries, and to only 6 percent in the high-income countries. Thus, over the development process, there is a marked structural change in labor supply, as workers move from subsistence self-employment to wage work. This pattern is well known; see e.g. Gollin (2008). What has not been studied systematically is hours worked for traditional-sector and modern-sector workers. Figures 3b and 3c show the corresponding cross-country patterns. Traditional-sector hours are slightly increasing by development, being 3.8 hours lower in poor than in rich countries. By contrast, modern sector hours are strongly decreasing by development: they are 11.3 hours higher in poor than in rich countries. As a result, for the poor and middle-income countries hours are markedly lower in the traditional than in the modern sector, namely by 10.9 and 5.8 weekly hours, respectively. Only for the rich countries are hours higher in the traditional sector, with a difference of 4.2 hours.

Taking the patterns of sectoral hours worked per worker and sectoral shares of workers together, it becomes clear that the modest increase in weekly hours worked per worker between low- and middle-income countries does not arise because of an increase in sectoral hours worked per worker, but is due to a compositional effect: hours are markedly lower in the traditional than in the modern sector in both low- and middle-income countries, and the substantial decrease in the share working in the traditional sector between low- and middle-income countries thus causes the small increase in average hours worked per worker. The small increase in hours per worker in the traditional

[2]For a subset of countries, we also have information on the number of employees. Classifying only self-employed individuals with low education and without any employees as being in the traditional sector has a negligible impact on the statistics presented in this section.

Figure 3: Sectoral Facts Across Countries

(a) Traditional Sector Share

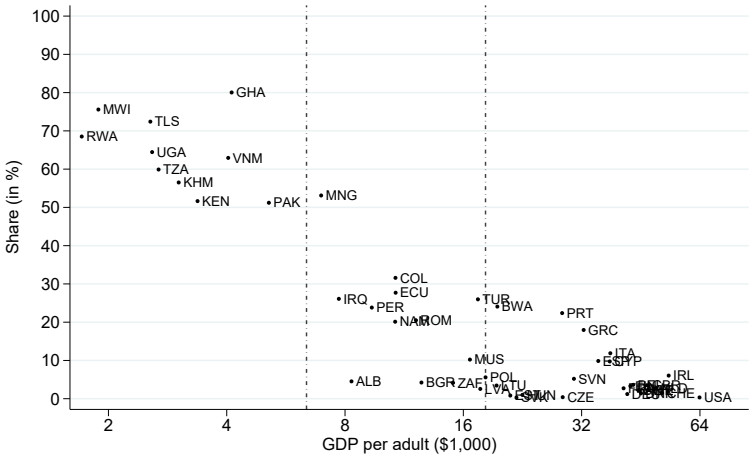

(b) Hours per Worker Traditional Sector

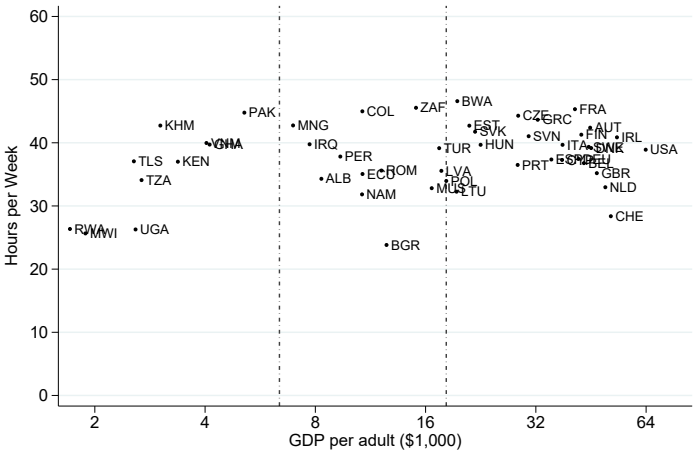

(c) Hours per Worker Modern Sector

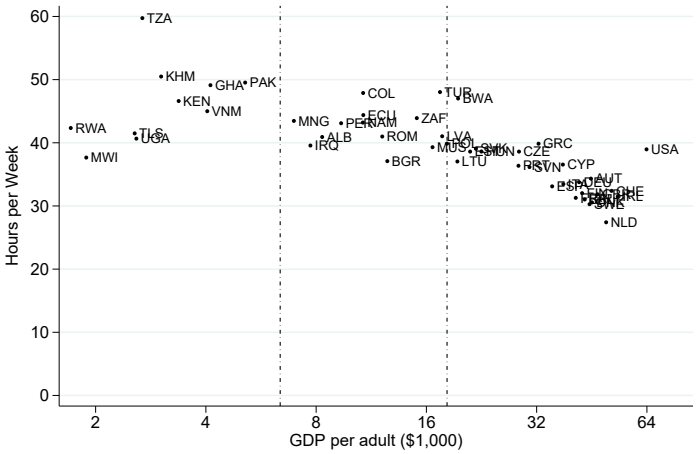

Note: This figure plots the share of employment in the traditional sector (Panel a), the average hours per worker in the traditional sector (Panel b), and the average hours per worker in the modern sector (Panel c). The vertical lines represent the division between low- and middle-income countries, and between middle- and high-income countries. Data source: Bick et al. (2018)

sector marginally adds to this increase. Thus, the initial fairly flat part in hours worked per worker over development is driven by this compositional effect. The decreasing part between middle- and high-income countries, by contrast, is driven by the strong decrease of 7.3 hours per worker in the modern sector between these two country income groups, with the large majority of individuals working in the modern sector in both country income groups. Note that these patterns are again robust for both genders and different age groups; see Appendix Tables A.3 to A.7.[3]

3. Model of Labor Supply and Structural Change

The objective of the paper is to understand the role of structural change in labor supply in explaining the facts presented in the previous section, in particular the convex pattern of employment rates in income per capita, and the concave pattern of hours per worker. To this end we build a model of labor supply along the extensive and intensive margins with two sectors: a traditional self-employment sector and a modern wage sector. Given that the empirical patterns in question are similar for both genders and across age groups, we abstract from modeling age and gender. The model also includes two more standard determinants of labor supply: preferences in which income effects potentially dominate substitution effects, and distortionary tax-and-transfer systems.

3.1. Environment

Households and Individuals There is a continuum of households of mass one in each country, and a continuum of individuals of mass one in each household. We allow for risk sharing within a household, and no insurance across households (as in Heathcote et al., 2014). Households in the model are meant to capture different kinds of informal insurance networks within a country, which might exist not only between members literally living in the same household, but also within villages or other groups (see e.g. Townsend, 1994, and Fafchamps and Lund, 2003).

[3] One potentially relevant factor for the decrease in hours per worker in the modern sector is the regulation of work hours. The World Bank's Doing Business Report provides country-level data on two main types of hours regulation used in practice, namely legal limits on the number of hours that can be worked per day, and legal limits on the number of days that can be worked per week. Neither series is correlated with GDP per capita, however, as Figure A.1 in the Appendix shows. Thus, while differential regulation of hours may well be important for explaining variation in hours across countries of similar income levels, it is not a prime candidate for explaining why individuals in richer countries work less on average than those in poorer countries.

Households differ in their modern sector productivity z with $\log z \sim N(0, \sigma_z^2)$. This is broadly consistent with the views of Caselli and Coleman (2001) and Porzio and Santangelo (2019), who posit that human capital does not matter in the agricultural sector, but only in non-agricultural work. Individuals within a household differ only in their individual fixed disutility of work $\eta \in \mathbb{R}_+$, where η is continuously distributed with the CDF $F(\eta)$ and the PDF $f(\eta)$. This individual fixed disutility of work is multiplied with a sector-specific disutility of work $\bar{u}_S$, where S represents the sector. We assume the standard utility function of an individual considered in Heathcote et al. (2014) but augment it by including the disutility of work:

$$u(c,h;S,\eta) = \frac{c^{1-\gamma}}{1-\gamma} - \alpha \frac{h^{1+\frac{1}{\phi}}}{1+\frac{1}{\phi}} - \bar{u}_S \eta I_{h>0}, \tag{1}$$

where c and h are individual consumption and hours worked, and $I_{h>0}$ is an indicator equal to 1 if the individual works. In what follows, variables c and h refer to the individual level, and C and H to the household level. Each household is headed by a household head who maximizes the sum of the utility of all household members with equal weight.

Two Sectors of Production The household can decide to work either in the traditional (T) or in the modern sector (M). The two sectors differ along five dimensions, summarized in Table 1, the first three relating to the production technology and productivity levels. The modern sector features a constant returns to scale technology with aggregate labor productivity A_M. Household income in the modern sector is thus equal to the hourly wage w times effective household hours worked (i.e. household hours multiplied by market productivity z).

The traditional sector, by contrast, features a decreasing-returns-to-scale technology, and household income in the traditional sector equals $Y_T = A_T H^\rho$ with $0 < \rho < 1$, where A_T is the traditional sector labor productivity. The decreasing returns to scale technology provides a parsimonious way of capturing the small scale of family farms or businesses in the developing world (Adamopoulos and Restuccia, 2014), the seasonality in agricultural production (Paxson, 1993), and the lack of well-functioning land or capital markets to increase scale (see e.g. Jayachandran, 2006, or Karlan et al., 2014). Thus, the marginal product of labor is decreasing in labor supply for those in the traditional sector, and the household is the residual claimant on all profits earned in the

Table 1: Differences Between Modern and Traditional Sector

	Modern, M	Traditional, T
Returns to scale in labor supply	1	ρ
Sector productivity level	A_M	A_T
Household productivity	z	1
Fixed cost of labor supply	$\bar{u}_M$	0
Taxation of labor income	yes	no

traditional sector. Fourth, working in the modern sector is associated with positive fixed costs $\bar{u}_M > 0$, while accessing the traditional sector is not associated with a fixed cost, i.e. $\bar{u}_T = 0$. Last, we assume that labor income in the traditional sector remains untaxed, due to limited tax enforceability in this sector (see Jensen, 2019).

3.2. Equilibrium Analysis

Household's Problem The household head faces a two-stage maximization problem. In the first stage, she chooses household hours H, consumption C, and the sector of employment S. In a second stage, given household hours and consumption, she chooses individual hours h and consumption c. We solve the maximization problem by backward induction.

Given (C,H,S), the second stage maximization problem amounts to

$$\begin{aligned} \max_{\{c(\cdot),h(\cdot)\}} \quad & \int u\left(c(\eta),h(\eta);S,\eta\right)dF(\eta) \\ \text{s.t.} \quad & \int c(\eta)dF(\eta) = C \\ & \int h(\eta)dF(\eta) = H. \end{aligned} \tag{2}$$

The first-order condition for consumption implies perfect consumption risk sharing within the household, i.e. $c(\eta) = C$ for all η. Also, due to the separability of disutility arising from working at the extensive and intensive margin, there is no variation within the household in optimal hours worked conditional on working. The optimal hours function thus can be expressed as

$$h(\eta) = \begin{cases} h^* > 0 & \text{for } \eta \leq \eta^* \\ 0 & \text{otherwise.} \end{cases}$$

The household head's problem therefore reduces to determining a threshold level η^*: all household members with a disutility of work below this threshold level work the same positive hours $h^* = \frac{H}{F(\eta^*)}$, and all household members with a disutility above this threshold level do not work. Given household hours H, hours worked per working household member are decreasing in the share of household members working, i.e., $\frac{dh^*}{d\eta^*} < 0$.

Substituting the optimal decisions into the objective function of the problem (2) gives the household utility:

$$U(C,H,S) \equiv \frac{C^{1-\gamma}}{1-\gamma} - \alpha \frac{H^{1+\frac{1}{\phi}}}{1+\frac{1}{\phi}} \left(F(\eta^*)\right)^{-\frac{1}{\phi}} - \bar{u}_S \int_0^{\eta^*} \eta dF, \tag{3}$$

which looks different from the individual utility (1), mirroring the result in Constantinides (1982).

After substituting the optimal consumption c and hours $h(\eta)$ described above, the first-order condition of the second stage problem (2) with respect to η^* is given by

$$\alpha \frac{h^{*1+\frac{1}{\phi}}}{1+\frac{1}{\phi}} f(\eta^*) + \bar{u}_S \eta^* f(\eta^*) = -\alpha h^{*\frac{1}{\phi}} F(\eta^*) \frac{dh^*}{d\eta^*}. \tag{4}$$

The first term on the left-hand side of Equation (4) is the disutility from working h^* hours for the new workers of mass $f(\eta^*)$ that start working when the optimal threshold level η^* is marginally changed. The second term adds to this the fixed utility costs incurred by these workers. These marginal utility losses of the new workers are equated at the optimum with the marginal utility gains the already existing workers of mass $F(\eta^*)$ enjoy because of the decrease in their hours worked, which is expressed on the right-hand side. Equation (4) thus implicitly defines the optimal threshold level η^* as a function of household hours $H = h^* F(\eta^*)$.

Assuming that the individual fixed utility cost of working is uniformly distributed with $\eta \sim U(0,1)$ and thus $F(\eta) = \eta$, we can solve η^* in closed form:

$$\eta^* = \min\left\{ \left(\frac{1}{\bar{u}_S} \frac{\alpha}{1+\phi} H^{1+\frac{1}{\phi}} \right)^{\frac{\phi}{1+2\phi}}, 1 \right\}. \tag{5}$$

This equation illustrates the trade-off that the household head faces between intensive and extensive margins of labor supply. If the household head chooses the traditional sector, she will make all household members work, i.e., $\eta^* = 1$, as the fixed cost of working in that sector is zero, $\bar{u}_T = 0$. If she instead chooses the modern sector where $\bar{u}_M > 0$, she will make only the members with low individual fixed cost η work, i.e., $\eta^* < 1$, unless household hours H are sufficiently high (see Appendix B for the upper bound of H that gives an interior solution). As H increases, the household head sends more workers by choosing a higher η^* in order to keep individual hours ($h^* = H/\eta^*$) relatively low. This force interacts with the curvature on labor supply ϕ in Equation (5).

Equation (5) is a novel way to separate out and solve for intensive and extensive margin labor supply decisions at the individual level. Heathcote et al. (2014) and Heathcote and Tsujiyama (2021) analytically solve for intensive margin labor supply decisions when there are labor productivity shocks that are privately insurable within a risk-sharing unit. Our model is complementary to their framework in that we provide an analytical characterization of extensive margin labor supply decisions. Equation (5) also enables us to solve the household utility (3) in closed form (see Appendix B). This result helps us handle two-dimensional heterogeneity with respect to z and η and thus drastically reduces the computational burden, which allows us to conduct a full estimation of the model.

In the first stage, the household head solves the following maximization problem of the household:

$$\begin{aligned} \max_{C,H,S\in\{T,M\}} \quad & U(C,H,S) \\ \text{s.t.} \quad & (1+\tau_c)C = Y_S - \mathscr{T}_S(Y_S) + \Upsilon, \\ \text{where} \quad & Y_M = wzH \text{ and } Y_T = A_T H^{\rho} \end{aligned} \tag{6}$$

where τ_c is a linear consumption tax, $\mathscr{T}_S(Y_S)$ are non-linear income taxes, and Υ are lump-sum transfers. We denote the solution to the household's problem by $\{C(z), H(z), S(z)\}_{z\in\mathbb{R}_{++}}$.

Equilibrium Wage The competitive modern sector clears such that

$$L = \int zH(z) \cdot 1_{\{S(z)=M\}} dF_z,$$

where F_z is the CDF of z. In equilibrium, the market-clearing wage is given by $w = A_M$.

Government Budget The government budget is balanced in equilibrium:

$$G + \Upsilon = \sum_{S=T,M} \left[\int \left(\mathscr{T}_S(Y(z)) + \tau_c C(z) \right) 1_{\{S(z)=S\}} dF_z \right], \tag{7}$$

where G is government spending that does not affect the utility of households.

Equilibrium A stationary equilibrium consists of a set of decision rules $\{c(\eta), h(\eta)\}$ and $\{C(z), H(z), S(z), \eta^*(z)\}$, a wage rate w, and the government policies $\{\tau_c, \mathscr{T}_T(\cdot), \mathscr{T}_M(\cdot), G, \Upsilon\}$ such that

(i) given the price and policies, the decision rules solve households' problems (2) and (6),

(ii) the wage clears the labor market: $w = A_M$, and

(iii) the government budget constraint (7) is satisfied.

3.3. The Process of Development

We assume that three sets of variables differ exogenously by development: (i) the aggregate productivity levels in the modern and traditional sector A_M and A_T; (ii) the size of the tax-and-transfer system $(\tau_c, \mathscr{T}_T(\cdot), \mathscr{T}_M(\cdot), G, \Upsilon)$, and (iii) the fixed cost of working in the modern sector $\bar{u}_M$. We solve steady-states for each level of development.

The differential fixed costs of market work across countries could reflect decreasing search frictions that accompany development (Poschke, 2019; Martellini and Menzio, 2020; Donovan and Schoellman, 2021), or declining transportation costs with development (Adamopoulos, 2011; Bryan et al., 2020), as we posited earlier. High fixed cost of market work in poor countries could also be related to the costs of urban-rural migration, as posited by a large literature in development economics. The development literature since Harris and Todaro (1970) has viewed the unemployment risk associated with migration to cities as a key cost of accessing wage employment. Fields (1975) emphasized how migration to cities would be concentrated among the more educated, and more productive workers, as in our model. This sorting pattern is consistent with a number of subsequent empirical studies of rural-urban migration (see Lagakos, 2020, and the references therein).

4. Model Estimation

This section describes how we estimate the model to minimize the distance between eight moments of the data from the high- and low-income countries and their model counterparts. We show how the model's preference parameters, fixed costs of modern work and production functions are identified by the sectoral employment shares and patterns of hours worked in the data. To match the data, the estimated model requires income effects that dominate substitution effects, modest decreasing returns to hours worked in the traditional sector, and fixed costs of work that are higher in poor countries than rich ones. Overall, we find that the model does well in matching the convex pattern of employment rates in income per capita, and the concave pattern of hours per worker, over the full development spectrum.

4.1. Tax-and-Transfer Systems

To parameterize the model, we need to discipline the size of tax-and-transfer systems across countries. For this model input, we generate novel empirical evidence for how the degree of tax progressivity, the importance of labor income taxes, and the share of social benefits change over the full development spectrum. To do this, we draw on two different data sets, and use available information for as many countries as possible for each input.

Our main data source is Egger et al. (2019), who assemble a comprehensive database of statutory tax codes across countries. To do so, they draw on official data from the IMF, the World Bank, the OECD, other government sources from individual countries, and data on taxation by private companies. To operationalize these data for use in our quantitative analysis, we assume the functional form for a progressive tax system used by Bénabou (2002) and Heathcote et al. (2014), with net income $\tilde{y}$ being given by

$$\tilde{y} = y - \mathscr{T}(y) = y - \left(y - \lambda y^{1-\tau}\right) = \lambda y^{1-\tau}, \tag{8}$$

where λ is informative about the level of taxation and τ about the progressivity. For $\tau = 0$, $1-\lambda$ represents a proportional tax on income, whereas for $\tau = 1$, net income is independent of gross income. We estimate τ for each country based on the data set compiled by Egger et al. (2019). Specifically, for each country we combine data on the average gross income at each percentile of the income distribution and the implied net

income, where the latter is calculated for a single individual without children, and using statutory tax codes excluding any transfers that are not incorporated directly into the tax system. Taking logs of Equation (8), we estimate τ for each country from a regression of log net earnings on log gross earnings. We then set λ such that the equilibrium share of government revenue coming from labor income taxes corresponds to the one in the data, which we also obtain from Egger et al. (2019).[4] Chang, Chang, and Kim (2018) and Holter, Krueger, and Stepanchuk (2019) also estimate the degree of tax progressivity for the OECD countries using Equation (8), but our estimation covers the whole development spectrum.[5]

We set the consumption tax rate such that the equilibrium government revenue to GDP ratio equals its data counterpart in the Egger et al. (2019) data, assuming a balanced budget. Thus, consumption taxes in our parameterization implicitly also contain revenues coming from tariffs or corporate taxes, assuming that all these revenues are raised as linear taxes on households. Finally, we redistribute only a fraction of government revenues to households. Specifically, we set Υ relative to the aggregate output equal to the share of social benefits over GDP, which we obtain from the IMF government statistics.[6]

Figure 4 plots the resulting components of the tax-and-transfer system that we use as an input into the model. We apply a piecewise linear interpolation of the averages for all variables over countries belonging to the poorest, middle, and richest terciles of the world income distribution, as measured by GDP per adult in the Penn World Tables. This simplification leaves the model ill-suited to explain differences in hours worked *within* groups of countries of similar income levels, but makes it useful for understanding why hours change with income across countries of different GDP per capita levels, which

[4]We take the sample of 62 countries with information on the share of government revenues coming from labor income taxes also for the estimation of the progressivity parameter, i.e. the sample of countries is consistent for the different fiscal inputs coming from Egger et al. (2019).

[5]Appendix Figure A.2 shows that our measure of tax progressivity for the OECD countries lines up more closely with Holter, Krueger, and Stepanchuk (2019) than with Chang, Chang, and Kim (2018). As in Holter, Krueger, and Stepanchuk (2019) the calculation of net incomes in Egger et al. (2019) which we use for our estimates of progressivity relies on statutory tax codes, whereas Chang, Chang, and Kim (2018) rely on reported gross and net incomes from household surveys and also include transfers not incorporated directly into the tax system.

[6]These social benefits include the provision of medical services, unemployment compensation, and social security pensions. We think of these as direct transfers to households, and treat all other government expenditure as either not affecting household utility, or alternatively as an added utility shifter that does not affect household choices.

Figure 4: Cross-Country Differences in the Tax-and-Transfer System

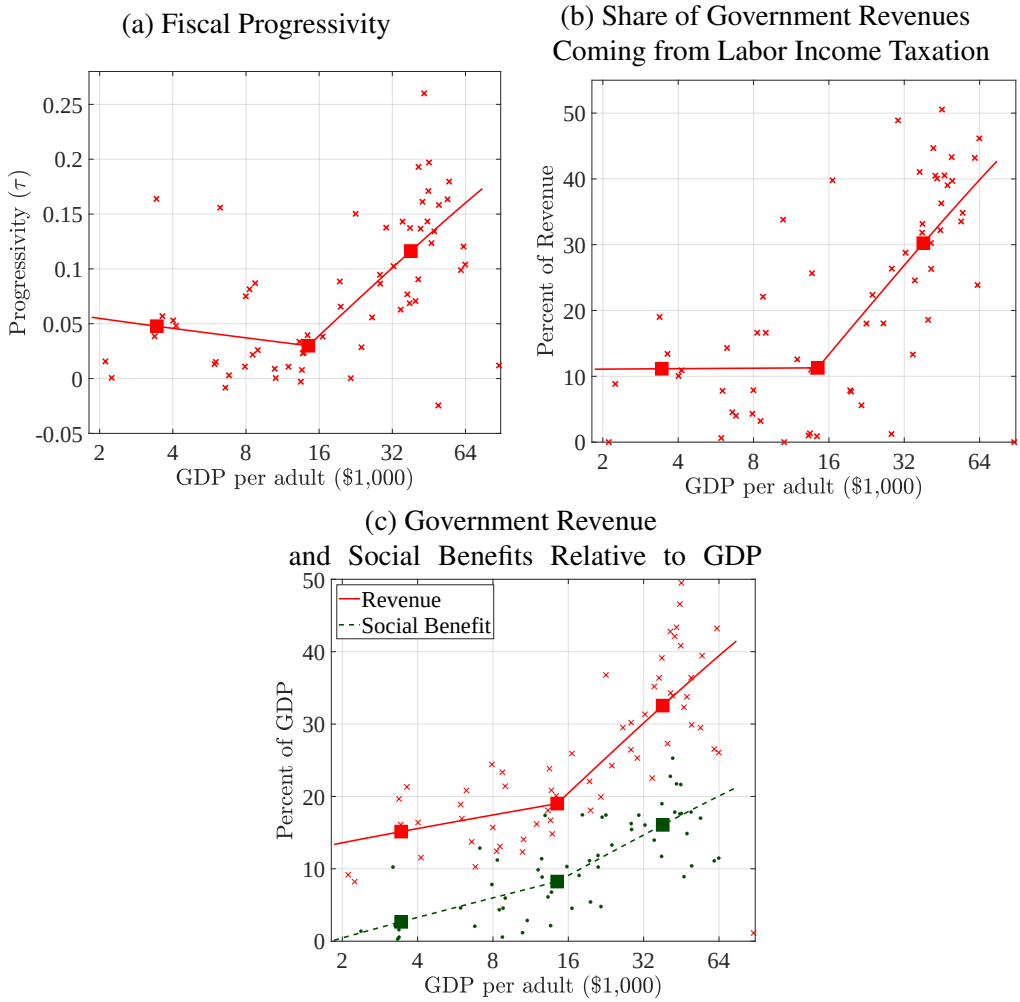

Note: The small dots represent each country in our sample, and the large dots the averages by country-income group. The lines show the piecewise linear interpolation.

is the goal of the paper. The estimate of progressivity is slightly U-shaped, exhibiting a slight decrease from low- to middle-income countries and a substantial increase from middle- to high-income countries (Figure 4a). The share of government revenue coming from labor income taxes is small and almost flat from low- to middle-income countries, but sharply increases from middle- to high-income countries (Figure 4b). The size of government transfers relative to GDP increases from the poor to middle-income countries and then somewhat stronger from the middle to the richest countries (Figure 4c). Similarly, the estimate of government revenue increases over the development spectrum,

with a sharp increase from the middle to the richest countries.

These government revenues relative to GDP can be interpreted as the overall burden of taxes. Taking this perspective, in the poorest tercile countries in our data, taxes are on average around 15 percent of GDP. In the richest tercile, in contrast, taxes are about 33 percent of GDP. Thus, by this metric, the tax burden is about 2.2 times as high in the richest countries as in the poorest. Since these taxes are distortionary, and because redistribution of taxes is perceived as outside income by the households, this will translate qualitatively into lower hours worked in richer countries in our model.

4.2. Other Exogenous Model Inputs

We set the level of modern sector labor productivity in the average country in the richest tercile, A_M^{rich}, to be 100, and the weight on the disutility from labor, α, to be one. Individual labor productivities, capturing permanent differences across households, are assumed to be log-normally distributed with mean 1. To estimate the variance of the distribution, we exploit the panel component of the Current Population Survey (CPS). We estimate a panel fixed effects regression of log income per hour, following Lagakos and Waugh (2013), and take the variance of the individual fixed effects (0.224) to be the variance of the (log) permanent productivity differences in our model.

4.3. Estimation

We jointly estimate the remaining eight model parameters, to replicate eight key moments from the data. Our estimation procedure is therefore a minimum distance estimation, which is a special case of the generalized method of moments (GMM). We construct all our estimation targets for the "average poor country" and "average rich country," leaving the rest of the world income distribution un-targeted. The moments we target are: (1 and 2) the average employment rates in the poor and rich countries; (3 and 4) average hours per worker in the traditional and modern sectors of the poor countries; (5) average hours per worker in the rich counties; (6) the ratio of average GDP per adult in the poor to rich countries; and (7 and 8) the average traditional sector employment shares in the poor and rich countries.

The parameters to estimate are $\Theta = \{\gamma, \phi, \bar{u}_M^P, \bar{u}_M^R, \rho, A_T^P, A_T^R, A_M^P\}$. These represent the curvature parameters on consumption and hours worked in preferences, the fixed costs of market work in the poor and rich countries, the returns-to-scale parameter in the traditional sector, the traditional-sector productivity terms in the poor and rich country,

and the modern-sector productivity term in the poor country (with A_M^P set to 100).

Since the model is highly nonlinear, and to ensure that our choice of estimated parameters leads to a global minimum, we apply a global optimization algorithm using quasi-random numbers generated by the Sobol sequence (see Arnoud, Guvenen, and Kleineberg, 2019). We estimate Θ by minimizing the percentage point differences between model and data moments for the employment rates and the traditional sector shares, and the percent differences for the remaining four moments. We find similar results using the inverse of the variance-covariance matrix of the moments as a weighting matrix.

Table 2 reports the targeted moments in the model and data, including bootstrapped 95 percent confidence intervals for the data moments. Overall, the estimated model does quite well in matching the data. On average, the absolute difference between model and data moments is less than one percent. The only moment with a difference larger than one percent is the employment rate in the poor countries, which has a relatively large confidence interval in the data. The other moment with a relatively large confidence interval is the traditional-sector share in the poor country. In general, we get larger confidence intervals for moments in the poor countries because we just have ten poor countries in our database, and there is some variability in values within the poor countries. Moreover, as Figures 2 and 3a show, the employment rate and the traditional sector share both show a clear decreasing pattern within the poor country group, increasing the bootstrapped standard errors around the average of the poor countries.

Table 3 reports the estimated parameters and bootstrapped 95 percent confidence intervals. The estimated curvature on consumption, γ, is 1.21, and its confidence interval lies entirely above one, implying income effects that dominate substitution effects. This is consistent with the literature on income effects driving hours worked decreases, though that literature has not agreed on a specific range for this parameter value. Our estimate is also broadly in line with that of Heathcote et al. (2014), who estimate dominant income effects in labor supply using entirely different data, on a panel of U.S. households covering the last several decades. The estimated curvature on labor supply, ϕ, is 0.49, and its confidence interval of 0.26 to 0.79 lies within a range of commonly used values at the individual level (see the intensive-margin estimates surveyed in Blundell and MaCurdy, 1999, and Keane, 2011).

The fixed costs of labor supply in the modern sector, $\bar{u}_M^P$ and $\bar{u}_M^R$, are estimated to be

Table 2: Moments Targeted in the Estimation

Moments	Data	Model
Employment Rate, Poor	74.5 (68.0, 81.1)	72.6
Employment Rate, Rich	54.6 (52.5, 56.7)	54.6
Hours per Worker, Traditional Sector, Poor	35.4 (35.2, 35.5)	35.1
Hours per Worker, Modern Sector, Poor	46.3 (46.1, 46.5)	46.0
Hours per Worker, Rich	35.1 (35.1, 35.2)	35.1
GDP per Adult ratio, Poor as % of Rich	9.0 (6.6, 11.4)	9.0
Traditional Sector Employment Share, Poor	64.3 (58.5, 70.1)	64.8
Traditional Sector Employment Share, Rich	5.7 (3.1, 8.4)	5.8

Note: This table reports the moments targeted in the estimation and their values in the data and in the model. The range reported below each data moment is its bootstrapped 95 percent confidence interval, computed by re-sampling with replacement from the cross-country database of Bick et al (2018).

2.04 and 0.86, respectively, in poor and rich countries, meaning substantially larger fixed costs in the poorer economy. The confidence interval for $\bar{u}_M^P$ is fairly wide, reflecting (as we argue further below) the relatively larger variance in traditional sector shares in the poorer countries. The confidence interval for $\bar{u}_M^R$ is somewhat narrower, and the p-value for the hypothesis that the two fixed costs are the same is 0.06.

The magnitudes of the model's estimated fixed costs are consistent with previous estimates in the literature. For example, expressing the fixed costs in units of time, workers in the modern sector of rich countries would be willing to work 26 percent more hours to avoid these fixed costs of work. This translates into 464 hours per year. The equivalent number for poor counties is 598 hours per year. The value for rich countries is well within the wide range of estimates of reported fixed costs by French (2019) in his study of labor supply and health shocks in the United States. In monetary

Table 3: Parameter Estimates and Confidence Intervals

Parameter	Description	Estimate (Confidence Interval)
γ	Curvature of consumption in preferences	1.21 (1.12, 1.41)
ϕ	Curvature of labor supply in preferences	0.49 (0.26, 0.79)
$\bar{u}_M^P$	Fixed cost of labor supply, poor countries	2.04 (1.12, 2.70)
$\bar{u}_M^R$	Fixed cost of labor supply, rich countries	0.86 (0.57, 1.19)
ρ	Returns to scale in traditional sector	0.77 (0.61, 0.89)
A_T^P	Traditional sector productivity, poor countries	2.36 (1.94, 2.72)
A_T^R	Traditional sector productivity, rich countries	13.33 (11.69, 15.36)
A_M^P	Modern sector productivity, poor countries	7.69 (5.71, 11.69)

Note: This table reports the estimated parameters. The confidence interval is the 2.5th and 97.5th percentiles of 500 bootstrapped parameter estimates. Levels of modern and traditional sector productivity are not directly comparable unless $\rho = 1$. Fixed costs of labor supply apply only to the modern sector. A_M^R is set to a value of 100.

terms, workers in our model's poor country would be willing to give up 26 percent of their consumption to avoid the fixed costs of work; the corresponding number for rich countries is 21 percent.

The returns to scale parameter in the traditional sector, ρ, is estimated to be 0.77, consistent with substantial decreasing returns in household labor supply there. The confidence interval for ρ is fairly wide, but we can easily reject constant returns, as in the modern sector. Our estimated returns to scale parameter is close to the estimate of Guner, Ventura, and Xu (2008), although the setups are not fully comparable.

Productivity in the traditional sectors in poor and rich countries, A_T^P and A_T^R, are estimated to be 2.36 and 13.33, meaning that the poor countries are around 18 percent as productive as rich countries in their traditional sectors. The estimated modern-sector

productivity in the poor country, A_M^P, is 7.69, compared to a value A_M^R set to 100, i.e. the poor countries are less than 8 percent as productive as rich countries in their modern sectors. The larger productivity differences in the modern sector with development are one driver of sectoral reallocation in the model, and are consistent with the literature on skill-biased technology differences across countries (Caselli and Coleman, 2006; Malmberg, 2020). This literature has argued that productivity differences across countries are larger for skill-intensive tasks than for unskilled tasks.

4.4. Identification

While all parameters are jointly identified from the data, in the point-identification sense of Lewbel (2019), some moments are more informative for some parameters than others. To help systematically illustrate how the model's parameters are identified from the data moments, we compute the derivative of each moment with respect to each parameter. Kaboski and Townsend (2011) report a similar matrix in their structural estimation, for example. In practice this means re-solving the model eight additional times, each time raising one parameter by one percent above its estimated value, keeping other parameters, tax rates and transfers constant. We then express these derivates in percentage point changes for the variables originally measured in percentage points (employment rates, traditional sector shares and the GDP ratio), and in percent changes for variables measured in hours per week. Table 4 reports these derivatives. For expositional purposes, we report the largest derivative in each column in bold face, to illustrate which moments respond most to each parameter, and we underline the largest derivative in each row to show which parameter has the biggest impact on each moment.

As Table 4 shows, the curvature parameter on consumption, γ, affects most of the moments and has the largest impact on the employment rate in the rich country. This moment also is most affected by γ. The largest impacts on hours per worker in the rich countries, and in the modern sector of the poor countries, also come from γ, showing that this parameter is broadly important in determining hours worked, and sectoral shares, in the model. A similar conclusion holds for the curvature on hours worked, ϕ, though with generally smaller impacts on the model moments.

Perhaps more interestingly, the fixed cost parameters, $\bar{u}_M^P$ and $\bar{u}_M^R$, have somewhat different impacts in the rich and poor countries. An increase in $\bar{u}_M^R$ has the biggest effect on the employment rate in rich countries, consistent with standard intuition about how fixed costs work in models of labor supply (Rogerson and Wallenius, 2013). In contrast,

Table 4: Derivative of Moments to Parameters

Moment	γ	ϕ	$\bar{u}_M^P$	$\bar{u}_M^R$	ρ	A_T^P	A_T^R	A_M^P
Employment Rate, Poor	0.0	-0.2	0.0	0.0	-0.2	0.5	0	-0.5
Employment Rate, Rich	**-1.0**	-0.1	0.0	**-0.2**	0.0	0.0	0.1	0.0
Hours / Worker, Traditional, Poor	-0.3	-0.1	0.0	0.0	**0.3**	-0.1	0.0	0.0
Hours / Worker, Modern, Poor	-0.2	0.2	0.2	0.0	0.0	0.0	0.0	0.0
Hours / Worker, Rich	-0.8	0.0	0.0	0.1	0.0	0.0	-0.1	0.0
GDP p.a., Poor/Rich	0.2	0.0	0.0	0.0	0.0	0.0	0.0	0.1
Traditional Share, Poor	0.5	**-0.3**	**0.4**	0.0	-0.3	**0.9**	0.0	**-0.9**
Traditional Share, Rich	0.4	0.0	0.0	0.1	-0.1	0.0	**0.3**	0.0

Note: This table reports the derivatives of each moment with respect to each parameter. The numbers in the table are the percentage point changes in all variables expressed in percentage points (the employment rates and traditional in the rich and poor countries plus the GDP per adult ratio), and the percent changes in the variables measured in hours per week (hours per worker in the rich country and poor countries by sector. The largest value in each column is bold faced, and the largest value in each row is underlined.

$\bar{u}_M^P$ has the largest impact on the size of the traditional sector, which it expands. This result highlights how households in the model can respond to higher fixed costs both by cutting employment rates or by entering the traditional sector and avoiding fixed costs altogether. In poorer countries, where the modern sector is relatively less attractive, households tend to choose the latter.

The returns-to-scale parameter in the traditional sector, ρ, has the largest impact on average hours in the traditional sector. This moment also responds the most to ρ, consistent with the intuition that hours are held down in developing economies by diminishing returns to hours of work. The ρ value less than one is also consistent with recent experimental evidence from the developing world showing that subsistence workers raise their hours worked in response to exogenous increases in assets (see in particular Bandiera, Burgess, Das, Gulesci, Rasul, and Sulaiman, 2017).

The sectoral productivity terms—A_T^P, A_T^R and A_M^P—have the expected effects on model moments. An increase in A_T^P mostly expands the traditional sector in the poor country, and an increase in A_T^R mostly expands the traditional sector in the rich country. An increase in A_M^P mostly increases the size of the modern sector in poor countries. The

moment that is the most responsive to A_M^P is the employment rate in poor countries. The reason is that when productivity in the modern sector rises, households move out of the traditional sector and therefore face fixed costs of supplying labor. They respond by reducing the number of household members who are employed.

Altogether, the derivatives in Table 4 help illustrate the central role that structural change in labor supply plays in disciplining our model quantitatively. The standard intuition that fixed costs of labor supply inform employment rates is only partially correct. Households, particularly in poorer economies, also respond to the presence of fixed costs of market work by switching to self-employment. The lower average hours in the traditional sector are informed largely by the extent of decreasing returns there, meaning that decreasing returns are an important factor keeping hours per worker low in poor countries. The curvature parameters in preferences inform labor supply directly, and also indirectly through the choice of sector. The lesson is that when trying to match moments across countries with vastly different income levels, it is not just the income levels that inform the model, but also differences in the nature of work.

As an additional exercise to illustrate how the model's parameters are identified, we re-estimate the model eight additional times, each with a different counterfactual assumption about the value of one of the moments targeted in the estimation. This exercise builds on the idea of the "sensitivity matrix" proposed by Andrews, Gentzkow, and Shapiro (2017), but computes the actual change in each estimated parameter rather than taking a first-order approximation. The results are presented in Appendix Table A.8. In general they are less transparent than the derivatives in Table 4 since they involve a full re-estimation of the model rather than a local change in one parameter. Yet they paint a similar picture, with fixed costs parameters being sensitive to structural change moments, productivity terms being sensitive to the size of the traditional sector, and preference parameters, particularly ϕ, responding to all moments.

4.5. Model Fit

In this section, we discuss the model fit relative to two sets of non-targeted moments. The first and most important ones are the shapes of the different variables over the entire development spectrum. While we target labor supply facts in the average poor and rich country, the facts for all other countries—in particular the entire middle of the world income distribution—are not targeted. Secondly, we analyze the model's success in replicating sectoral patterns in consumption and earnings in poor countries, where the

traditional sector plays a large role.

To construct country-specific model values for all variables, we proceed as follows. First, we assume the logarithm of aggregate traditional sector productivity $\log(A_T)$ and the fixed cost of working in the modern sector $\bar{u}_M$ both change linearly in $\log(A_M)$. Second, we assume all fiscal inputs (i.e., tax progressivity, the share of government revenue coming from labor income taxes, the ratio of government revenue to GDP, and the share of government consumption) change piecewise linearly in $\log(A_M)$. For each level of A_M, we then solve for optimal allocations.

Figure 5 compares the model predictions (blue line) against the data (small red dots). The large red dots denote the averages by country groups in the data, and the stars mark the subset of targeted moments. Panel (a) in the upper left corner shows that hours per adult decrease at a similar rate in the model as in the data. In fact, average hours per adult in the middle income countries are replicated exactly by the model. Panel (b) then focuses on the share of workers in the traditional sector: as in the data, this share decreases rapidly between poor and middle income countries, and approaches zero in the richest countries. The lower two panels (c) and (d) show that the model replicates the different behavior of the two margins of hours per adult. Employment rates are decreasing strongly between low- and middle-income countries, with a modest increase for the richest countries, while hours per worker are similar between low- and middle-income countries on average, and substantially lower in the richer countries. Thus, while the model does not perfectly match the shapes, it generates both the convex decrease in the employment rates, including the uptick in the richest countries, and the concave decrease in hours per worker over the development spectrum.

The success behind generating the different shapes of the two margins of labor supply stems from structural change, i.e. households moving from the traditional to the modern sector over the development spectrum, and decreasing fixed costs of work in the modern sector. Because of the absence of fixed costs in the traditional sector, the employment rate in the traditional sector is always 1 in the model. By contrast, the employment rate in the modern sector is significantly below 1, but increasing in development due to the decrease in the fixed cost of working there, $\bar{u}_M$. The strong decrease in the traditional sector share from 64 percent to less than 20 percent between low- and middle-income countries generates the decrease in the employment rate between these two country groups. In both middle- and high-income countries, by contrast, the large

Figure 5: Estimated Model

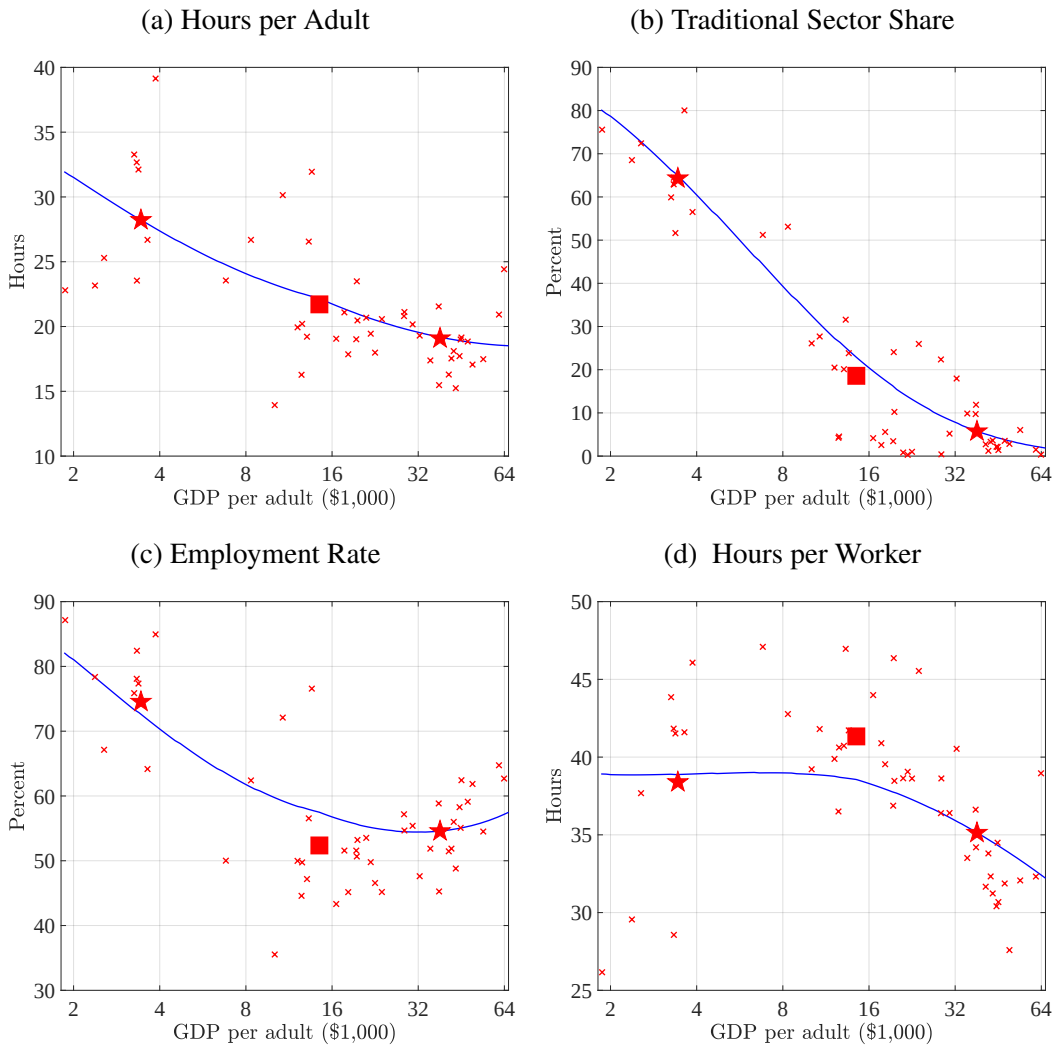

Note: This figure plots average hours per adult, employment rates, hours per worker and the traditional sector share of employment in the data and in the estimated model. Stars represent calibration targets, and squares represent un-targeted means of the low-, middle- or high-income country groups.

majority of households works in the modern sector, and the decrease in the fixed cost in this sector generates the mild increase in the employment rate.

Besides the sectoral share, sectoral hours are important for aggregate hours per worker. Due to the decreasing returns to scale in the traditional sector, model-predicted hours per worker in the modern sector are higher than in the traditional sector over the full development spectrum (see Appendix Figure A.3). In both sectors, hours are

decreasing with development, only at a slightly faster rate in the modern sector. The sectoral reallocation from the traditional to the modern sector generates the flat hours per worker in the aggregate between poor and middle-income countries. The fall-off between the middle- and high-income countries then largely mimics the decrease in hours per worker in the modern sector.

Two further non-targeted moments of importance are the ratios of earnings and consumption for households working in the modern vs. traditional sector. This is most informative in the poor countries, where the majority of households work in the traditional sector. In the model, households in the average poor country who work in the modern sector have 2.1 times higher earnings and 2.0 times higher consumption than households in the traditional sector. This is in line with evidence that households working in non-agriculture have 2.1 times higher earnings and 1.7 times higher consumption than households working in agriculture in poor countries (see Gollin, Lagakos, and Waugh, 2014).

5. Investigating the Mechanisms of Structural Change

Structural change in labor supply has two channels in our model: sectoral reallocation from the traditional to the modern sector, and the declining fixed costs of work with development. In this section, we investigate the importance of these two channels in matching the cross-country data. To do so, we carry out two exercises. First, we estimate a model in which we assume that the same fixed costs of work apply in the traditional sector as in the modern sector. Secondly, we estimate a model in which fixed costs only apply in the modern sector, but are not allowed to change by development.

The first exercise, in which fixed costs of work apply in both sectors, leaves the number of parameters and moments unchanged from the baseline exercise, but re-estimates all the parameters. Figure 6 shows the resulting model predictions; the estimated parameters are in Appendix Table A.9. This model with fixed costs of work in the traditional sector can match all the targets except the traditional sector share in the rich countries. However, it gets the predictions for the rest of the world income distribution more or less entirely wrong. Employment rates and consequently hours per adult are now hump-shaped, with a strongly counterfactual increase in both margins when moving from low- to middle-income economies. These counterfactual predictions are driven by fixed costs of working that are low and slightly increasing over the development spectrum, rather

Figure 6: Model with Fixed Costs of Working in Both Sectors

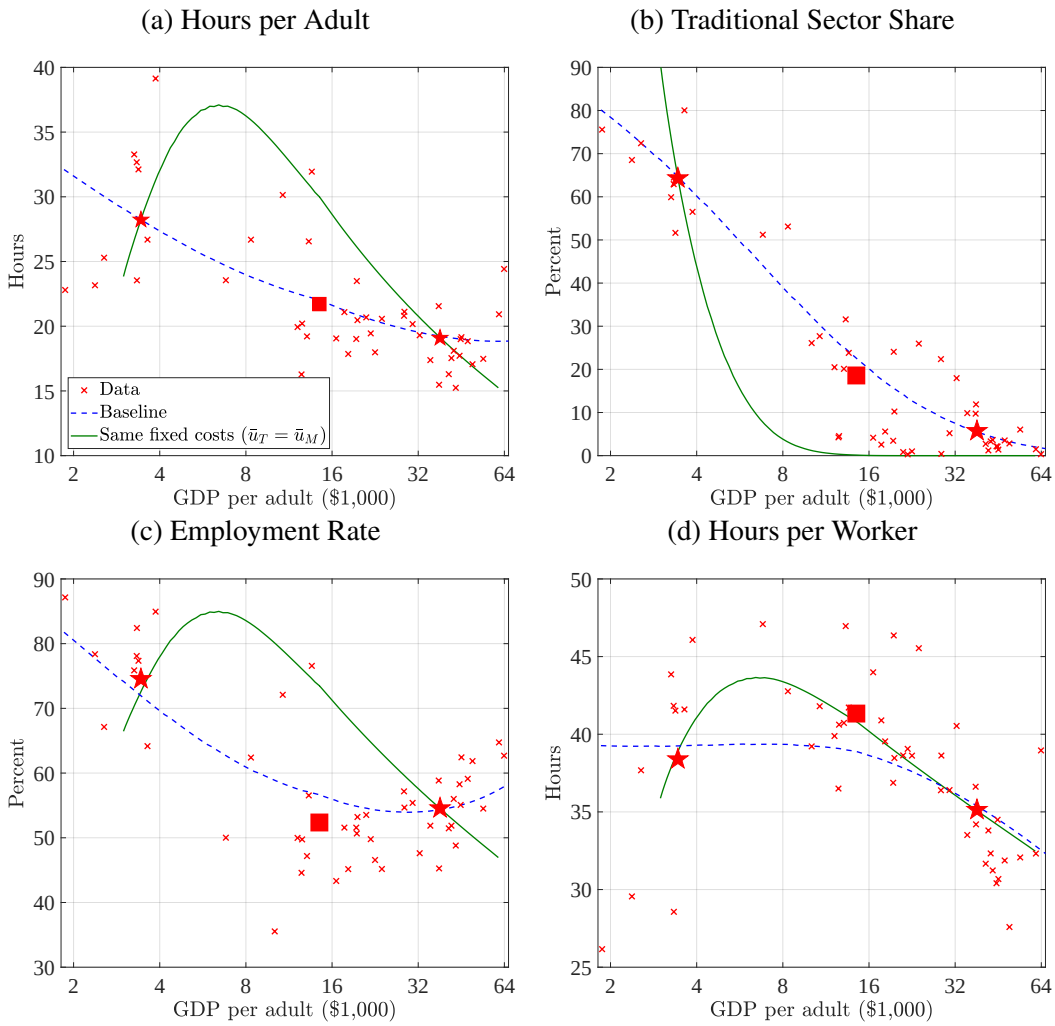

Note: This figure plots average hours per adult, employment rates, hours per worker and the traditional sector share of employment in the data and in a re-estimation of the model that restricts the fixed costs of working to be the same in both sectors. Stars represent calibration targets, and squares represent un-targeted means of the low-, middle- or high-income country groups. The dashed line reproduces the predictions from the main estimation of the model.

than decreasing. The resulting employment rate is lower in the traditional than in the modern sector. Both employment rates then fall over the development spectrum, and as workers move into the modern sector, employment rates rise counterfactually. The hump in the overall employment rate is thus a result of the rapid shift from the traditional sector into the modern sector, with the decreasing part driven in part by the increasing

fixed costs.

In the second re-estimation exercise, we require the fixed costs of work to apply only in the modern sector, as in the baseline estimation, but we impose them to stay constant by development. This implies that we have one less parameter to be estimated, namely $\bar{u}_M^R$. Consequently, we also use one less estimation target, and target average hours per adult in the rich countries, rather than the average employment rates and hours per worker in rich countries separately.

The results are presented in Figure 7, and estimated parameter values are reported in Appendix Table A.10. The model matches the decreasing traditional sector share and also hours worked per adult well. However, it badly overestimates the decrease in the employment rate, and fails to replicate the decrease in hours per worker from middle- to high-income countries. With constant fixed costs of working over the development spectrum, the model misses the countervailing force that induces an increase in the employment rate. Consequently, the income effects have to be smaller now than in the baseline model, and in fact the estimation of γ implies that income and substitution effects basically cancel each other out, resulting in almost flat hours per worker in the upper end of the development spectrum, when almost all workers are in the modern sector. The sectoral reallocation into the modern sector, together with increasing taxes and transfers, are strong enough to generate the decrease in hours per adult, given the absence of decreasing fixed costs of work as a counteracting force.

These two exercises make it clear that fixed costs applying only to the modern sector and changing with development are two crucial features to match the data over the full development spectrum. In the appendix, we show that a model without any traditional sector can also not replicate the differential shapes of the two margins of labor supply over the development spectrum (see Appendix Figure A.4). We also find that assuming constant returns to scale in the traditional sector exaggerates the convex pattern of employment rates and concave pattern of hours per worker with income per capita (Appendix Figure A.5) . On the other hand, assuming full taxation in the traditional sector makes very little difference in matching the data (Appendix Figure A.6), suggesting this is not an important feature in understanding the patterns in question.

Figure 7: Model with Constant Fixed Costs Across Countries

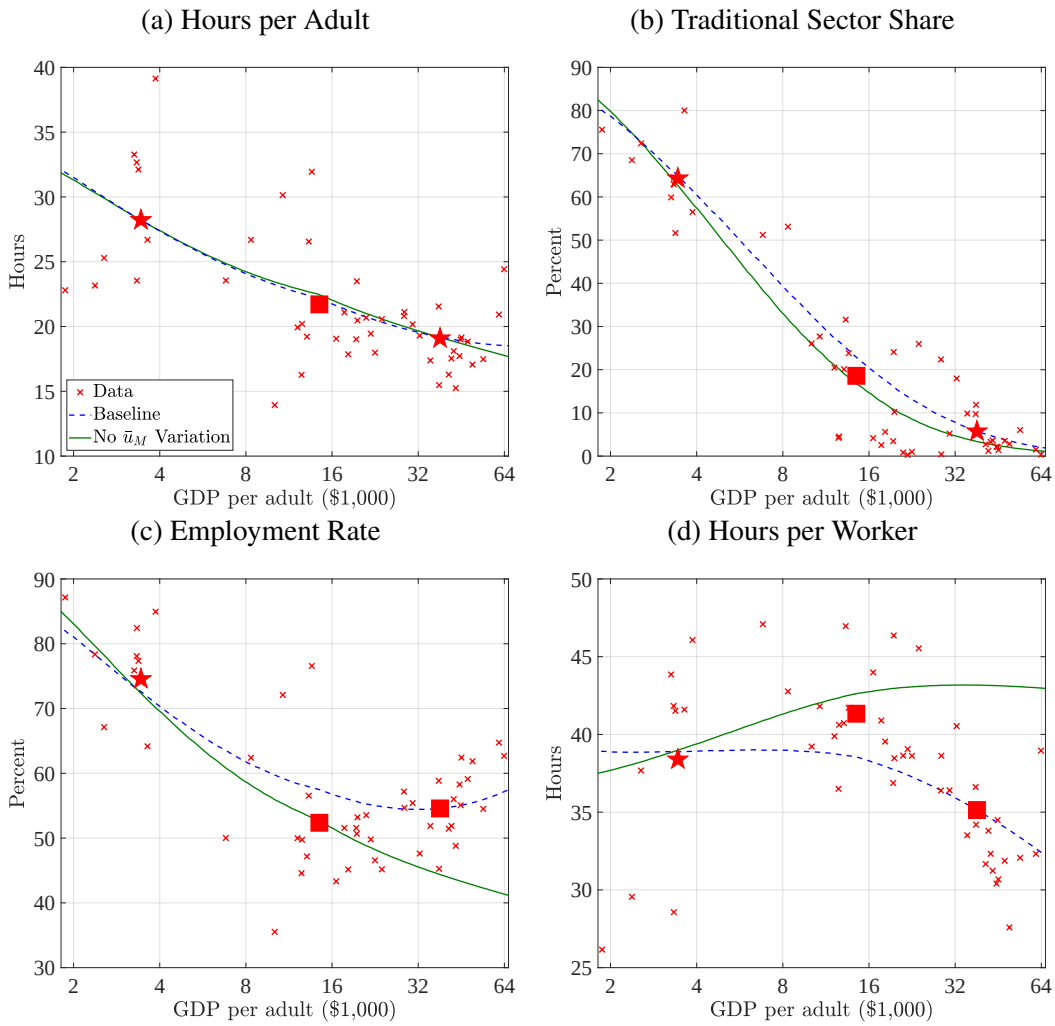

Note: This figure plots average hours per adult, employment rates, hours per worker and the traditional sector share of employment in the data and in a re-estimation of the model that restricts the fixed costs of working to be identical across countries, and present only in the modern sector. Stars represent calibration targets, and squares represent un-targeted means of the low-, middle- or high-income country groups. The dashed line reproduces the predictions from the main estimation of the model.

6. Decomposing the Driving Forces of the Decrease in Hours Worked

After establishing the importance of modeling structural change in labor supply, we now analyze the quantitative importance of the different driving forces in explaining the decrease in hours worked per adult between poor and rich countries. To do this, we start

Table 5: Decomposing Poor-Rich Differences in Average Hours Worked

	Hours	% Explained
Model	9.0	100.0
Higher Productivity	5.5	61.1
Higher Taxes & Transfers	2.3	25.6
Structural Change in Labor Supply		
Lower Fixed Costs	−2.7	−30.0
Sectoral Reallocation	3.8	42.2

Note: This table shows the marginal change in hours when, starting from the estimated baseline model for the average poor country, one model component after the other is set to the level of the average rich country. The first column shows the change in hours, the second column the percent explained of the total hours change. The table shows the mean of six possible decomposition exercises that change the ordering of the steps. Appendix Table A.14 also reports the minimum and maximum effects among the different orderings.

with the estimated baseline model for the average poor country, then set one component after the other to the level of the average rich country in a cumulative fashion, and calculate the marginal change in hours by adding that specific component. To give a concrete example, we first change the productivity level in both sectors to the level of the rich countries, afterwards we additionally set the tax-and-transfer system to the level of the rich country, and finally let the fixed costs in the modern sector fall to the level of the average rich country. In all these three decomposition steps, we do not allow for sectoral reallocation: households are required to continue working in the same sector they optimally chose given the poor countries' environment, but they can optimally adjust their hours. The role of sectoral reallocation is assessed in the last step by allowing households to choose their sector. We conduct the decomposition exercise in a cumulative way, since the effect of sectoral reallocation between poor and rich countries on hours worked is ultimately the residual. However, this approach implies that the marginal effects of the first three steps, i.e. by how much hours per adult change when turning on a specific feature, depend on the order in which they are turned on, and there exist six possible orderings. Rather than reporting the results from each ordering, Table 5 reports the average effect across all orderings (Appendix Table A.14 reports also the minimum and maximum effects).

Row 1 of Table 5 shows the hours difference between the average poor and rich country predicted by the full model (9 hours), and the subsequent rows show the effect attributed to each model feature separately. The higher productivities in the modern and traditional sectors in rich than in poor countries predict on average a decrease of 5.5 hours, accounting for 61.1 percent of the overall predicted hours gap. The model thus predicts that on average almost two thirds of the observed decrease in hours worked per adult between poor and rich countries can be attributed to an income effect, which is the most important driving factor of the decrease. The second most important driver of the decrease in hours per adult between poor and rich countries is one of our two features of structural change in labor supply, namely sectoral reallocation, i.e. the move from the traditional into the modern sector of production over the course of the development process. Sectoral reallocation accounts for 42.2 percent of the total hours decrease. Higher taxes and transfers have a more modest effect and predict on average an hours decrease of only 2.3 hours, accounting for a quarter (25.6 percent) of the total hours gap. Finally, the second feature of structural change in labor supply, lower fixed costs of working in the modern sector in rich countries, alone predicts *higher* hours in rich countries, resulting on average in a negative difference of 2.7 hours between poor and rich countries. Thus, while the decreasing fixed costs of work are an important model feature to explain the convex shape of the employment rate and the concave shape of hours per worker, they do not contribute to explaining the overall decrease in hours. On the contrary, based on the decreasing fixed costs alone, hours in rich countries would be higher than in poor countries.

Table 6 shows the results of the same decomposition exercise, but focusing on the hours decrease between poor and middle-income countries (columns 1 and 3) or between middle-income and rich countries (columns 2 and 4), respectively. Hours per adult decrease by 6.0 hours at the lower end and 3.0 hours at the upper end of the development spectrum, which affects the explanatory power of the model components in terms of percentages. The role of income effects is increasing between the lower and the upper half of the development spectrum, which stems from the fact that aggregate labor productivities change more between middle-income and rich countries than between poor and middle-income countries.[7] As Figure 4 shows, taxes and transfers do not

[7] Appendix Table A.16 shows the ratio and differences of the different inputs in the decompositions between low- and middle- and middle- and high-income countries. These ratios and differences are not necessarily the same across both halves of the development spectrum.

Table 6: Decomposing Differences in Average Hours Worked by Income Level

	Hours		% Explained	
	Poor-Middle	Middle-Rich	Poor-Middle	Middle-Rich
Model	6.0	3.0	100.0	100.0
Higher Productivity	3.3	2.2	55.0	73.3
Higher Taxes & Transfers	0.7	1.6	11.7	53.3
Structural Change in Labor Supply				
Lower Fixed Costs	−1.4	−1.3	−23.3	−43.3
Sectoral Reallocation	3.4	0.4	56.7	13.3

Note: This table shows the marginal change in hours when, starting from the estimated baseline model for the average poor country, one model component after the other is set to the level of the average middle-income country (first and third columns), or when, starting from the estimated baseline model for the average middle-income country, one model component after the other is set to the level of the average rich country (second and fourth columns). The first and second columns show the change in hours, the third and fourth columns the percent explained of the total hours change. The table shows the mean of six possible decomposition exercises that change the ordering of the steps. Appendix Table A.15 also reports the minimum and maximum effects among the different orderings.

change much between poor and middle-income countries, and consequently also play a minor role in explaining the decrease in hours at this development stage. In contrast, they are an important driver of the hours decrease in the upper half of the development spectrum: there, they explain 53.3 percent of the decrease, reflecting the large increase in taxes and transfers at this stage of development. The findings of the model are thus also in line with the literature that establishes taxes as an important driver of hours differences between different rich countries and in the development of hours in the rich countries over the last half century (see e.g Prescott, 2004 and Ohanian et al., 2008).

Turning to the two forces of structural change in labor supply, sectoral reallocation is a very important driver of the hours decrease at the lower end of the development spectrum, where households move rapidly from the traditional into the modern sector. Its importance mimics the one of income effects at this stage of the development process: both explain slightly more than half of the total hours decrease. However, sectoral reallocation loses importance as a driver of the hours decrease in the upper half of the development spectrum, because already in the middle-income countries more than 80 percent of the workers work in the modern sector. The opposite holds true for the coun-

teracting role of decreasing fixed costs: this role is somewhat muted at the lower end of the development spectrum, but then becomes very strong at the upper end, since the fixed costs only apply when households work in the modern sector. Decreasing fixed costs alone would on average predict 1.3 higher hours in rich than in middle-income countries, rather than 3.0 lower ones.

7. Hours Worked in the Future

Will people continue to work less in the future, as economies continue to grow? Prognosticators since Keynes (1930) have posited that, indeed, strong income effects will lead households to demand more leisure, and work less, as labor productivity continues to rise (see Ohanian, 2008, and the references therein). The elegant preference specification of Boppart and Krusell (2020) captures the secular decline in hours with productivity growth through just this force, and predicts continued reductions in labor supply going forward. Observers such as Oppenheimer (2002) have rung the alarm over automation eliminating jobs in the future, and academic evidence does suggest a negative link between recent adoption of robots and employment, at least in the short run (see e.g. Acemoglu and Restrepo, 2020). Future increases in tax rates, and more generous transfer systems, provide additional incentives to reduce labor supply, and cross-country evidence suggests that tax-and-transfer systems will continue to expand as countries grow richer.

Our analysis suggests that declining fixed costs of work may be an important channel moving in the opposite direction, potentially stabilizing or even raising average hours worked in the future. To illustrate this point, Figure 8 computes the model's predictions for economies with income levels up to $128,000 (in 2011 International Dollars in PPP). Assuming a growth rate of GDP per adult of 2 percent per year, it will take the average rich country in our data 61 years to reach this level. For the United States, this would correspond to roughly a doubling of GDP per adult.

The model predictions in Figure 8 extrapolate from our cross-country data, meaning that they assume the same relationship between GDP per adult, the size of tax-and-transfer systems, and fixed costs of work as in our main analysis. The model's predictions for hours worked, and the two margins of aggregate labor supply, are driven therefore by higher productivity levels, higher tax rates and larger transfers, and lower fixed costs of work. The figure also includes a breakdown akin to our decomposition

Figure 8: Declining Fixed Costs and the Future of Hours Worked

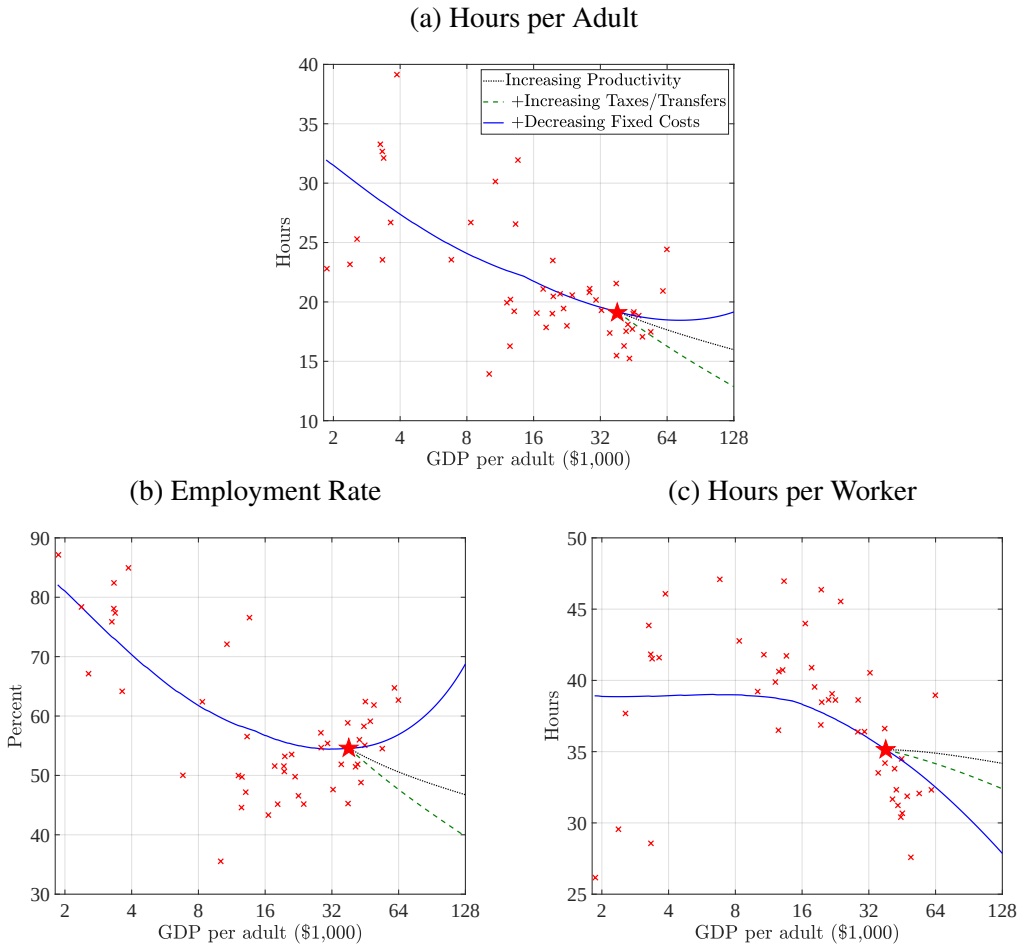

Note: This figure plots the model's predicted average hours per adult, employment rates and hours per worker for higher income levels. The star represents the average rich country in the data.

exercise, namely asking what would happen if only productivity continued to increase, if both productivity and taxes and transfers increased, or if additionally fixed costs continued to decrease. Since sectoral reallocation hardly plays any role anymore at this stage of development, we do not separately consider the role of sectoral reallocation.

The black dotted line in Figure 8a shows that if only productivity would continue to increase, average hours per adult would fall by 3 hours per work relative to the average rich country today. This is similar to what Keynes (1930) and others had in mind, with income effects dominating substitution effects in preferences. Adding the impact of an expanding tax-and-transfer system would lead to a reduction of another 3.2 hours

per week (the green dashed line). Thus, the model without structural change in labor supply would predict sizeable decreases in hours worked in the future, consistent with the standard views.

The solid blue line in Figure 8a shows that once declining fixed costs are added back into the model, average hours are basically unchanged compared to the average rich country today. The decreasing fixed costs are thus such a strong countervailing force that they completely offset the two other forces leading hours to fall. Figures 8b and 8c show the breakdown for the two margins. Increasing productivity and higher taxes both reduce employment and hours per worker, whereas the decreasing fixed costs operate in opposite directions on the two margins. As more people are drawn into employment, average hours per worker are reduced. In other words, decreasing fixed costs point to a world where more people work in the future, but where each supplies fewer hours on average.

Obviously, it is unclear what the future will hold for productivity growth, tax rates, and fixed costs of market work. The goal of this section is to show that a future with less work is not a foregone conclusion. Recent advancements in gig-work like Uber or Lyft probably have reduced—and will continue to reduce—the cost of working, as did the surge in working from home amid the COVID-19 pandemic (see, e.g. Bick et al., 2021). To the extent that the fixed costs of working continue to fall, average hours worked could remain similar to current levels or even grow. This point underscores the importance of distinguishing between the drivers of hours worked across countries, and in particular in separating income effects in labor supply from the changing nature of labor supply as countries develop, as we have emphasized.

8. Conclusion

This paper explores a new source of variation in hours worked across countries, which we call *structural change in labor supply*. We emphasize two forces that accompany development: sectoral reallocation from self-employment to market wage work, and declining fixed costs of market work. The first leads employment rates to fall, and hours per worker to rise, as households transition into work arrangements characterized by high fixed costs but constant, rather than decreasing, marginal products of labor. The second leads employment rates to rise as more and more individuals enter the labor market to take advantage of jobs with lower fixed costs.

We estimate the model to match patterns of labor supply in low- and high-income countries using rich data on labor supply along the intensive and extensive margins, rates of labor taxation, and the size of government transfers. The estimated model is largely successful in matching the observed data, including the concave pattern of hours per worker in GDP per capita, and the convex pattern of employment rates. Variants of our model that ignore the transition from self-employment to market work, and the decline in fixed costs of work with development, make counterfactual predictions for these intensive and extensive margins of aggregate labor supply.

Our study offers a new perspective on the future of hours worked relative to existing studies, which focus on income effects in labor supply, automation, and expanding tax-and-transfer systems. These channels point to a future of ever lower labor supply as productivity rises, new technologies are developed, and the "welfare state" continues to expand. Our model shows that declining fixed costs are an offsetting force that can increase employment rates in the coming decades. Extrapolating from recent experience—including the rise in work from home resulting from the pandemic—suggests that structural change in labor supply is likely to be a quantitatively important factor shaping the future of work.

References

Acemoglu, D. and P. Restrepo (2020). Robots and Jobs: Evidence from US Labor Markets. *Journal of Political Economy 128*(6), 2540–69.

Adamopoulos, T. (2011). Transportation Costs, Agricultural Productivity, and Cross-Country Income Differences. *International Economic Review 52*(2), 489–521.

Adamopoulos, T. and D. Restuccia (2014). The Size Distribution of Farms and International Productivity Differences. *American Economic Review 104*(6), 1667–1697.

Alder, S., T. Boppart, and A. Müller (2021). A Theory of Structural Change That Can Fit the Data. *American Economic Journal: Macroeconomics*. Forthcoming.

Andrews, I., M. Gentzkow, and J. M. Shapiro (2017). Measuring the Sensitivity of Parameter Estimates to Estimation Moments. *Quarterly Journal of Economics 132*(4), 1553–1592.

Arnoud, A., F. Guvenen, and T. Kleineberg (2019). Benchmarking Global Optimizers. Technical report. University of Minnesota.

Bandiera, O., R. Burgess, N. Das, S. Gulesci, I. Rasul, and M. Sulaiman (2017). Labor Markets and Poverty in Village Economies. *Quarterly Journal of Economics 132*(2), 811–870.

Bénabou, R. (2002). Tax and Education Policy in a Heterogeneous-Agent Economy: What Levels of Redistribution Maximize Growth and Efficiency? *Econometrica 70*(2), 481–517.

Besley, T. and T. Persson (2014). Why Do Developing Countries Tax So Little? *Journal of Economic Perspectives 28*(4), 99–120.

Bick, A., A. Blandin, and K. Mertens (2021, February). Work from Home Before and After the COVID-19 Outbreak. Working Paper 2017, Federal Reserve Bank of Dallas.

Bick, A. and N. Fuchs-Schündeln (2018). Taxation and Labor Supply of Married Couples Across Countries: A Macroeconomic Analysis. *Review of Economic Studies 85*(3), 1543–1576.

Bick, A., N. Fuchs-Schündeln, and D. Lagakos (2018). How Do Hours Worked Vary with Income? Cross-Country Evidence and Implications. *American Economic Review 108*(1), 170–199.

Blundell, R. and T. MaCurdy (1999). Labor Supply: A Review of Alternative Approaches. In O. Ashenfelter and D. Card (Eds.), *Handbook of Labor Economics*, Chapter 27, pp. 1559–1695. Elsevier Science B.V.

Boppart, T. and P. Krusell (2020). Labor Supply in the Past, Present, and Future: A Balanced-Growth Perspective. *Journal of Political Economy 128*(1), 118–157.

Bridgman, B., G. Duernecker, and B. Herrendorf (2018). Structural Transformation, Marketization and Household Production around the World. *Journal of Development Economics 133*, 102–126.

Bryan, G., E. Glaeser, and N. Tsivanidis (2020). Cities in the Developing World. *Annual Review of Economics 12*(1), 273–297.

Buera, F. J. and J. P. Kaboski (2012). The Rise of the Service Economy. *American Economic Review 102*(6), 2540–69.

Caselli, F. and W. J. Coleman (2001). The U.S. Structural Transformation and Regional Convergence: A Reinterpretation. *Journal of Political Economy 109*(3), 584–616.

Caselli, F. and W. J. Coleman (2006). The World Technology Frontier. *American Economic Review 96*(3), 499–522.

Cervelatti, M. and U. Sunde (2013). Life Expectancy, Schooling, and Lifetime Labor Supply: Theory and Evidence Revisited. *Econometrica 81*(5), 2055–2086.

Chang, B. H., Y. Chang, and S.-B. Kim (2018). Pareto Weights in Practice: A Quantitative Analysis Across 32 OECD Countries. *Review of Economic Dynamics 28*, 181–204.

Constantinides, G. M. (1982). Intertemporal Asset Pricing with Heterogeneous Consumers and without Demand Aggregation. *Journal of Business 55*, 253–267.

Donovan, K. and T. Schoellman (2021). The Role of Labour Market Frictions in Structural Transformation. Working paper, Yale University.

Egger, P. H., S. Nigai, and N. M. Strecker (2019). The Taxing Deed of Globalization. *American Economic Review 109*(2), 353–390.

Fafchamps, M. and S. Lund (2003). Risk Sharing Networks in Rural Philippines. *Journal of Development Economics 71*, 261–287.

Fields, G. S. (1975). Rural-Urban Migration, Urban Unemployment and Underemployment, and Job-Search Activity in LDCs. *Journal of Development Economics 2*(2), 165–187.

French, E. (2019). The Effects of Health, Wealth, and Wages on Labour Supply and Retirement Behaviour. *Review of Economic Studies 72*(2), 395–427.

Goldin, C. (1995). The U-Shaped Female Labor Force Function in Economic Development and Economic History. In T. P. Schultz (Ed.), *Investment In Women's Human Capital and Economic Development*, pp. 61 – 90. University of Chicago Press.

Gollin, D. (2008). Nobody's Business But My Own: Self-Employment and Small Enterprise in Economic Development. *Journal of Monetary Economics 55*(2), 219–233.

Gollin, D., D. Lagakos, and M. E. Waugh (2014). The Agricultural Productivity Gap. *Quarterly Journal of Economics 129*(2), 939–993.

Guner, N., G. Ventura, and Y. Xu (2008). Macroeconomic Implications of Size-Dependent Policies. *Review of Economic Dynamics 11*, 721–744.

Harris, J. R. and M. P. Todaro (1970). Migration, Unemployment and Development: A Two-Sector Analysis. *American Economic Review 60*(1).

Heathcote, J., K. Storesletten, and G. Violante (2014). Consumption and Labor Supply with Partial Insurance: An Analytical Framework. *American Economic Review 104*(7), 2075–2126.

Heathcote, J. and H. Tsujiyama (2021). Optimal Income Taxation: Mirrlees Meets Ramsey . *Journal of Political Economy* (forthcoming).

Herrendorf, B., R. Rogerson, and A. Valentinyi (2014). Growth and Structural Transformation. In P. Aghion and S. N. Durlauf (Eds.), *Handbook of Economic Growth*, pp. 855–941. Elsevier.

Holter, H., D. Krueger, and S. Stepanchuk (2019). How Do Tax Progressivity and Household Heterogeneity Affect Laffer Curves? *Quantitative Economics 10*(4), 1317–1356.

Jayachandran, S. (2006). Selling Labor Low: Wage Responses to Productivity Shocks in Developing Countries. *Journal of Political Economy 114*(3), 538–575.

Jensen, A. (2019). Employment Structure and the Rise of the Modern Tax System. Unpublished Working Paper, Harvard University.

Kaboski, J. P. and R. M. Townsend (2011). A Structural Evaluation of a Large-Scale Quasi-Experimental Microfinance Initiative. *Econometrica 79*(5), 1357–1406.

Karlan, D., R. Osei, I. Osei-Akoto, and C. Udry (2014). Agricultural Decisions After Relaxing Credit and Risk Constraints. *Quarterly Journal of Economics 129*(2), 597–652.

Keane, M. P. (2011). Labor Supply and Taxes: A Survey. *Journal of Economic Literature 49*, 961–1075.

Keynes, J. (1930). *Economic Possibilities for our Grandchildren*. New York: W. W. Norton. Reprinted in Essays in Persuasion.

Kopytov, A., N. Roussanov, and M. Taschereau-Dumouchel (2021, May). Cheap Thrills: the Price of Leisure and the Global Decline in Work Hours. NBER Working Paper No. 27744.

Kuznets, S. (1966). *Modern Economic Growth: Rate, Structure, and Spread*. New Haven, CT: Yale University Press.

Lagakos, D. (2020). Urban-Rural Gaps in the Developing World: Does Internal Migration Offer Opportunities? *Journal of Economic Perspectives 34*(3), 174–92.

Lagakos, D. and M. E. Waugh (2013). Selection, Agriculture, and Cross-Country Productivity Differences. *American Economic Review 103*(2), 948–980.

Lewbel, A. (2019). The Identification Zoo: Meanings of Identification in Econometrics. *Journal of Economic Literature 57*(4), 835–903.

Lewis, A. W. (1954). Economic Development with Unlimited Supplies of Labor. *The Manchester School 22*(2), 139–91.

MaCurdy, T. E. (1981). An Empirical Model of Labor Supply in a Life-Cycle Setting. *Journal of Political Economy 89*(6), 1059–1085.

Malmberg, H. (2020). Skill Bias and Development – Evidence from Trade and Industry Data. Unpublished Working Paper, University of Minnesota.

Martellini, P. and G. Menzio (2020). Declining Search Frictions, Unemployment, and Growth. *Journal of Political Economy 128*(12), 4387–4437.

McDaniel, C. (2011). Forces Shaping Hours Worked in the OECD, 1960-2004. *American Economic Journal: Macroeconomics 3*, 27–52.

Ngai, L. R. and C. A. Pissarides (2008). Trends in Hours and Economic Growth. *Review of Economic Dynamics 3*(11), 239–256.

Ohanian, L., A. Raffo, and R. Rogerson (2008). Long-Term Changes in Labor Supply and Taxes: Evidence from OECD Countries, 1956-2004. *Journal of Monetary Economics 55*, 1353–1362.

Ohanian, L. E. (2008). Back to the Future with Keynes. *Federal Reserve Bank of Minneapolis Quarterly Review 32*(1), 10–16.

Oppenheimer, A. (2002). *The Robots Are Coming: The Future of Jobs in the Age of Automation*. New York, NY: Vintage Books.

Paxson, C. H. (1993). Consumption and Income Seasonality in Thailand. *Journal of Political Economy 101*(21), 39–72.

Porzio, T. and G. Santangelo (2019). Does Schooling Cause Structural Transformation? Unpublished Working Paper, Columbia University.

Poschke, M. (2019). Wage Employment, Unemployment and Self-Employment Across Countries. Unpublished Working Paper, McGill University.

Prescott, E. C. (2004). Why Do Americans Work So Much More Than Europeans? *Federal Reserve Bank of Minneapolis Quarterly Review 28*, 2–13.

Ramey, V. A. and N. Francis (2009). A Century of Work and Leisure. *American Economic Journal: Macroeconomics 1*(2), 189–224.

Restuccia, D. and G. Vandenbroucke (2013). A Century of Human Capital and Hours. *Economic Inquiry 51*(3), 1849–1866.

Rogerson, R. (2008). Structural Transformation and the Deterioration of European Labor Market Outcomes. *Journal of Political Economy 116*, 235–259.

Rogerson, R. and J. Wallenius (2013). Nonconvexities, Retirement and the Elasticity of Labor Supply. *American Economic Review 103*(4), 1445–1462.

Storesletten, K., B. Zhao, and F. Zilibotti (2019). Business Cycle during Structural Change: Arthur Lewis' Theory from a Neoclassical Perspective. NBER Working Paper No. 26181.

Townsend, R. M. (1994). Risk and Insurance in Village India. *Econometrica 62*(3), 539–591.

Vandenbroucke, G. (2009). Trends in Hours: The U.S. from 1900 to 1950. *Journal of Economic Dynamics and Control 33*, 237–249.

Appendix (For Online Publication Only)

A. Appendix Figures and Tables

Table A.1: Extensive Margin by Gender and Age

	Country Income Group		
	Low	Middle	High
All	74.5	52.4	54.6
Men	80.6	63.2	62.0
Women	68.5	42.1	47.7
Young (15-24)	57.4	32.4	37.9
Prime (25-54)	86.2	70.6	78.9
Old (55+)	69.8	30.5	24.0

Table A.2: Intensive Margin by Gender and Age

	Country Income Group		
	Low	Middle	High
All	38.4	41.3	35.1
Men	40.8	43.7	38.2
Women	35.0	37.0	31.5
Young (15-24)	36.1	39.8	32.6
Prime (25-54)	40.6	42.3	35.9
Old (55+)	32.6	37.5	33.6

Table A.3: Sectoral Hours Worked and Sectoral Shares: Men Only

	Country Income Group		
	Low	Middle	High
Traditional Sec. Hours	37.2	39.1	42.0
Modern Sec. Hours	47.4	44.4	38.0
Traditional Sec. Share	54.9	17.3	6.4

Table A.4: Sectoral Hours Worked and Sectoral Shares: Women Only

	Country Income Group		
	Low	Middle	High
Traditional Sec. Hours	33.1	33.0	33.8
Modern Sec. Hours	43.0	38.2	31.4
Traditional Sec. Share	75.7	20.8	4.8

Table A.5: Sectoral Hours Worked and Sectoral Shares: Young Only (15-24)

	Country Income Group		
	Low	Middle	High
Traditional Sec. Hours	32.8	33.0	31.1
Modern Sec. Hours	44.9	40.8	32.9
Traditional Sec. Share	59.9	15.2	2.3

Table A.6: Sectoral Hours Worked and Sectoral Shares: Prime-Aged Only (25-54)

	Country Income Group		
	Low	Middle	High
Traditional Sec. Hours	37.4	38.8	41.9
Modern Sec. Hours	47.1	43.0	35.6
Traditional Sec. Share	62.0	15.8	4.4

Table A.7: Sectoral Hours Worked and Sectoral Shares: Old Only (55+)

	Country Income Group		
	Low	Middle	High
Traditional Sec. Hours	31.5	34.2	36.2
Modern Sec. Hours	43.4	40.3	33.3
Traditional Sec. Share	82.4	40.5	16.2

Figure A.1: Cross-Country Differences in Hours Regulation

(a) Legal Limits on Hours per Day (b) Legal Limits on Days per Week

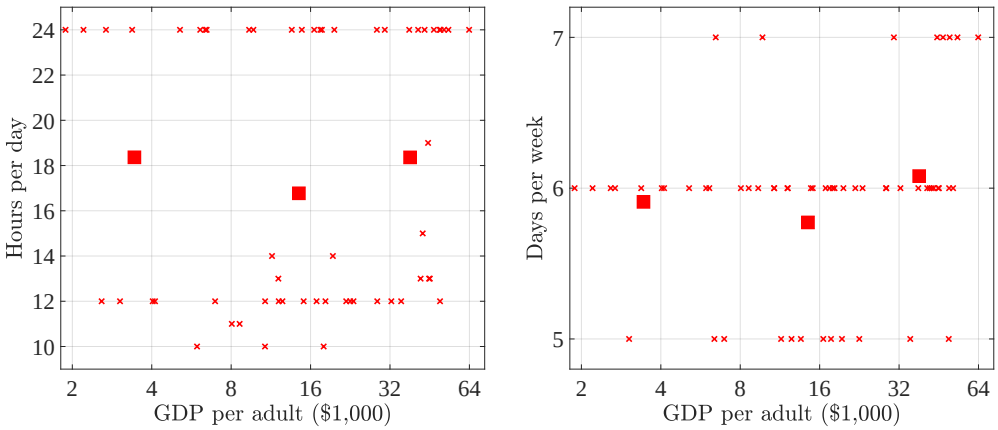

Source: Doing Business 2005, World Bank. Note: The small dots represent each country with available data, and the large dots the averages by country-income group. Legal limits on maximum hours worked per day, depicted in panel (a), may be defined either as mandatory maximum regular and overtime working hours per day or by mandatory minimum rest hours per day. If nothing is specified in the law, 24 hours are used. Legal limits on the maximum number of work days per week, depicted in panel (b), may be defined either by a number of mandatory rest days per week or by a mandatory minimum of consecutive hours of rest. If nothing is specified, seven days are used.

Figure A.2: Comparison of Progressivity Estimates

(a) Comparison with Holter et al. (2019) (b) Comparison with Chang et al. (2018)

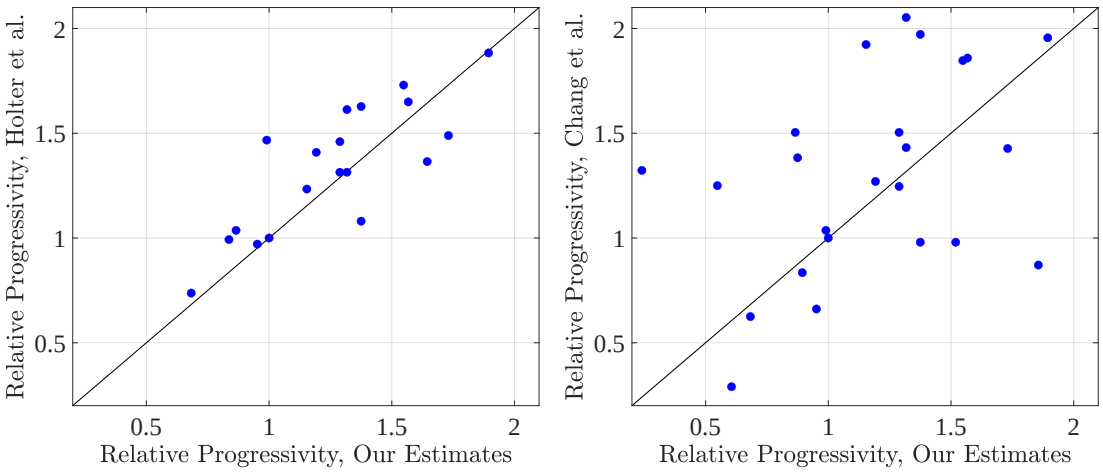

Note: This figure plots our estimates of country-specific progressivity relative to the one for the United States against those in Holter et al. (2019) (Panel a) and in Chang et al. (2018) (Panel b). The estimates shown in panel a are directly taken from table 1 in Holter et al. (2019), and those shown in panel b are constructed using table 3 in Chang et al. (2018).

Table A.8: Sensitivity of Parameter Estimates to Moments

	γ	ϕ	$\bar{u}_M^P$	$\bar{u}_M^R$	ρ	A_T^P	A_T^R	A_M^P
95% CI (% Deviation from Estimate)								
Lower Bound	-7.4	-46.9	-45.1	-33.7	-20.8	-17.8	-12.3	-25.7
Upper Bound	16.5	61.2	32.4	38.4	15.6	15.3	15.2	52.0
Moment								
Employment Rate, Poor	-1.7	14.8	-7.1	2.0	6.4	3.3	5.0	-3.0
Employment Rate, Rich	-1.9	6.2	0.5	1.5	5.0	3.6	4.4	1.7
Hours / Worker, Traditional, Poor	0.7	-2.6	0.0	-1.8	7.1	-0.4	1.6	-0.7
Hours / Worker, Modern, Poor	-1.4	20.9	-4.3	-1.0	5.5	2.7	5.7	-2.6
Hours / Worker, Rich	-1.6	-8.7	5.0	9.5	-4.2	1.0	-3.2	3.6
GDP p.a., Poor/Rich	1.6	-3.1	-3.5	-5.1	-3.1	9.7	-1.7	11.1
Traditional Share, Poor	-3.2	14.8	2.8	8.2	12.0	4.9	6.8	0.7
Traditional Share, Rich	-0.2	0.4	0.2	1.7	0.4	-0.1	3.1	-0.3

Note: This table reports the sensitivity of each parameter estimate to each targeted moment. Each row represents one re-estimation of the model, where we assume that one targeted moment is counterfactually higher than it actually is. The numbers in the table represent the percent change in each parameter estimate relative to their value in the main estimation. For moments measured in percentage points we consider a counterfactual increase of one percentage point in the moment. For moments measured in hours per week we consider a counterfactual increase of one hour per week. To put these estimates into perspective, the two top rows report the 95% confidence interval of the parameter estimates expressed as % deviations of the respective point estimate.

Figure A.3: Baseline Model Fit: Hours per Worker by Sector

(a) Modern Sector

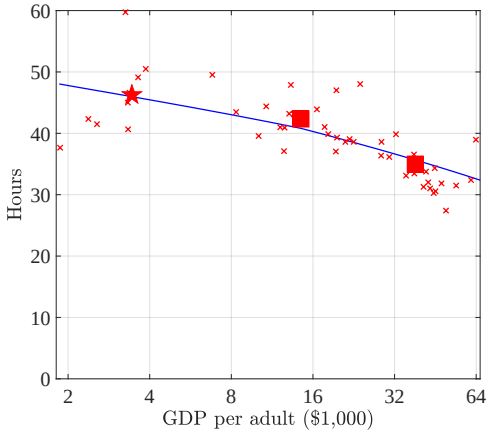

(b) Traditional Sector

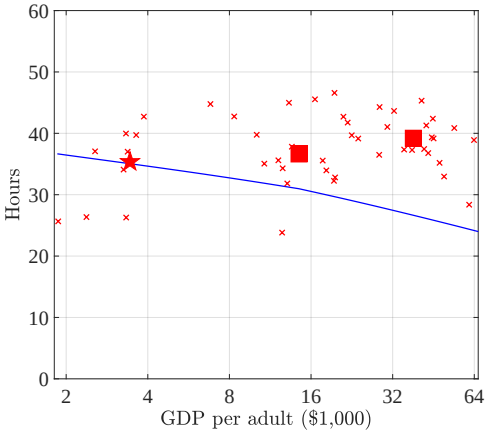

Table A.9: Estimated Parameters: Fixed Costs in Both Sectors

Specification	γ	ϕ	$\bar{u}_M^P$	$\bar{u}_M^R$	ρ	A_T^P	A_T^R	A_M^P
Baseline	1.21	0.49	2.04	0.86	0.77	2.36	13.33	7.69
$\bar{u}_M = \bar{u}_T$	1.43	1.75	0.23	0.27	0.32	2.52	3.19	3.24

Table A.10: Estimated Parameters: Constant Fixed Costs of Working in Modern Sector

Specification	γ	ϕ	$\bar{u}_M^P$	$\bar{u}_M^R$	ρ	A_T^P	A_T^R	A_M^P
Baseline	1.21	0.49	2.04	0.86	0.77	2.36	13.33	7.69
$\bar{u}_M^P = \bar{u}_M^R$	1.00	0.21	3.14	*	0.54	2.05	7.78	8.32

Figure A.4: Model Fit without Traditional Sector

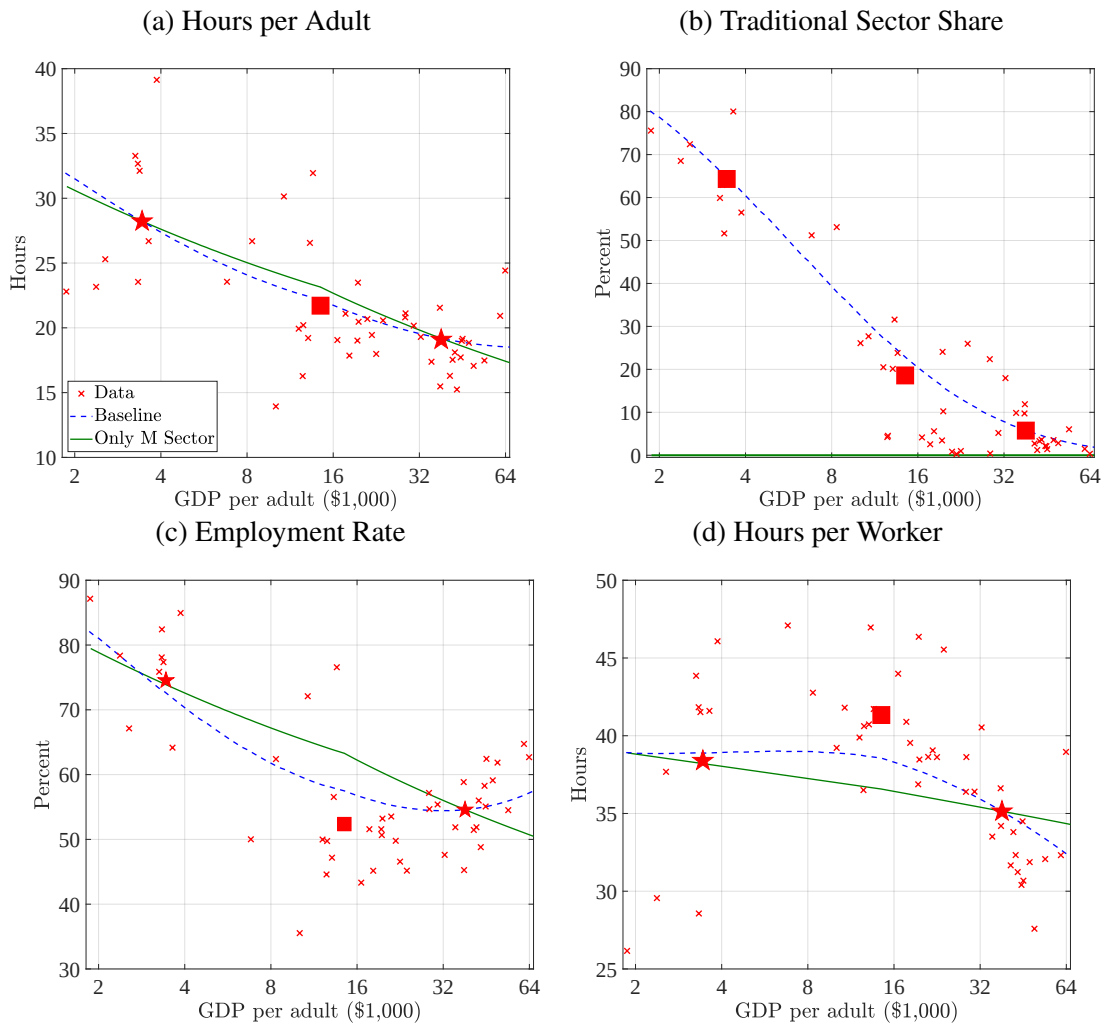

Table A.11: Calibrated Parameters: No Traditional Sector

Specification	γ	ϕ	$\bar{u}_M^P$	$\bar{u}_M^R$	ρ	A_T^P	A_T^R	A_M^P
Baseline	1.21	0.49	2.04	0.86	0.77	2.36	13.33	7.69
No T Sector	1.34	0.13	1.05	0.67	*	*	*	6.24

* A_T^{rich}, A_T^{poor} and ρ do not apply in the model without traditional sector.

Figure A.5: Constant Returns to Scale in Traditional Sector ($\rho = 1$)

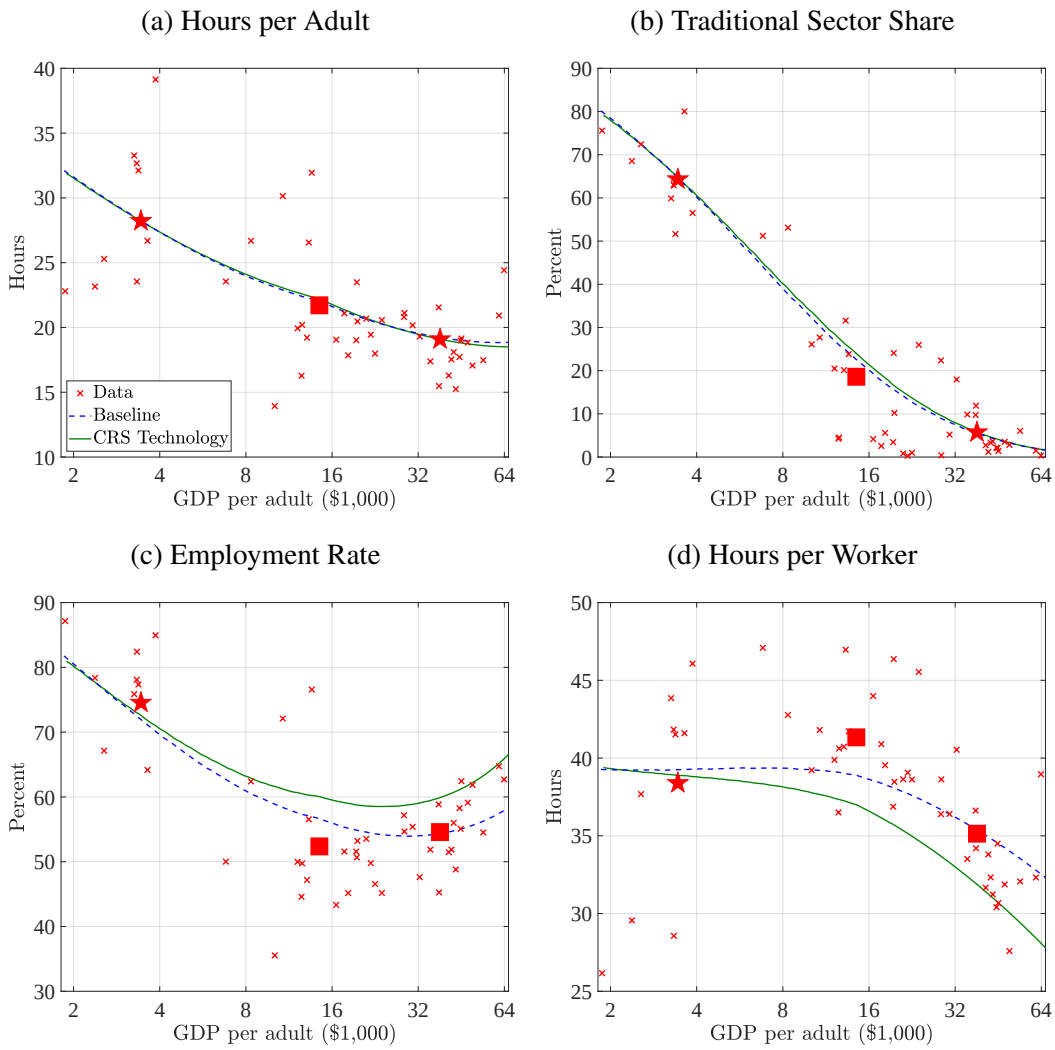

Table A.12: Calibrated Parameters: Constant Returns to Scale in Traditional Sector

Specification	γ	ϕ	$\bar{u}_M^P$	$\bar{u}_M^R$	ρ	A_T^P	A_T^R	A_M^P
Baseline	1.21	0.49	2.04	0.86	0.77	2.36	13.33	7.69
	1.17	0.81	1.76	0.66	1.00	2.66	17.67	7.41

Figure A.6: Taxation of Labor Income in Traditional Sector

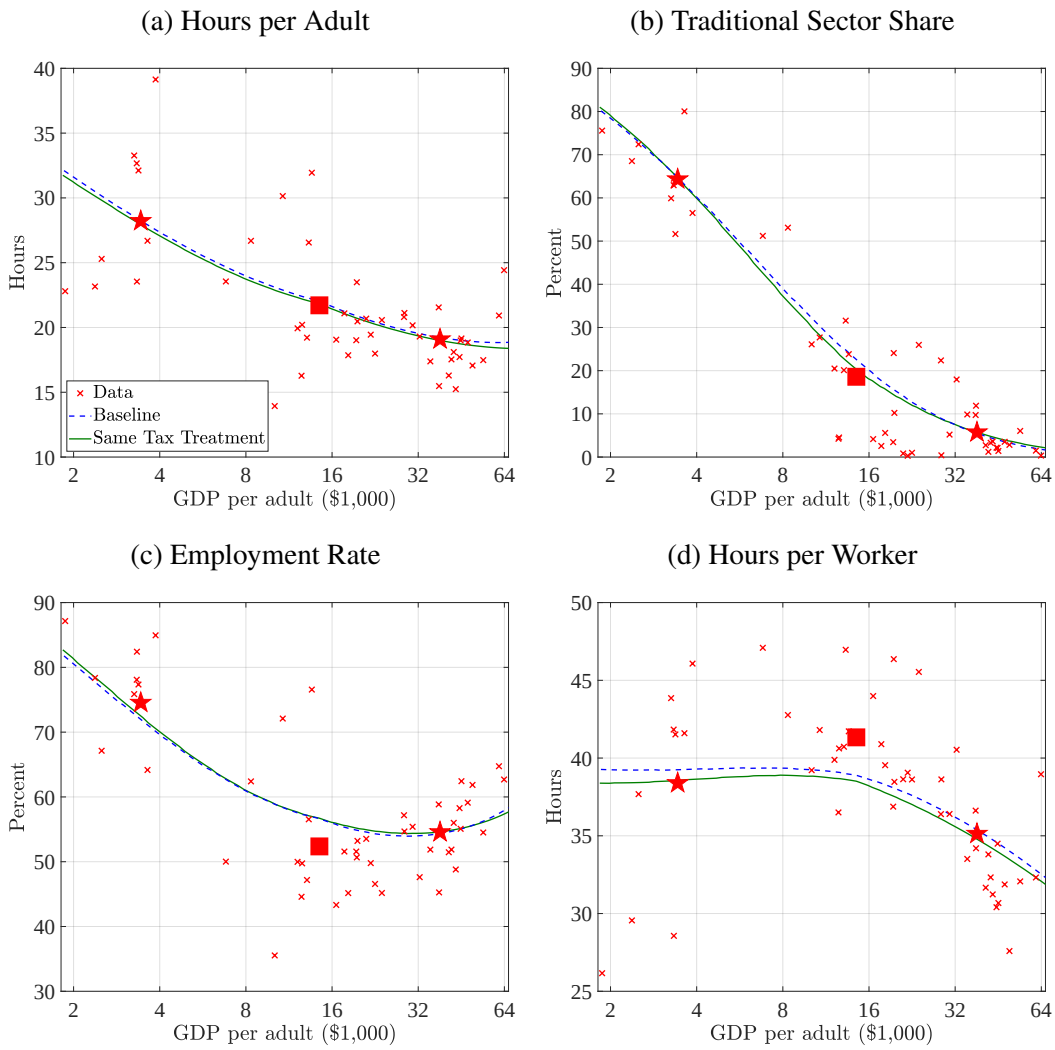

Table A.13: Calibrated Parameters: Taxation of Labor Income in Traditional Sector

Specification	γ	ϕ	$\bar{u}_M^P$	$\bar{u}_M^R$	ρ	A_T^P	A_T^R	A_M^P
Baseline	1.21	0.49	2.04	0.86	0.77	2.36	13.33	7.69
	1.21	0.49	2.05	0.86	0.80	2.35	12.52	7.72

Table A.14: Decomposing Poor-Rich Differences in Average Hours Worked

	Hours			% Explained		
	Mean	Min	Max	Mean	Min	Max
Model	9.0			100.0		
Higher Productivity	5.5	4.9	6.1	61.1	54.5	67.8
Higher Taxes & Transfers	2.3	2.0	2.6	25.6	22.2	28.9
Lower Fixed Costs	-2.7	-3.4	-1.9	-30.0	-37.8	-21.1
Sectoral Reallocation	3.8	3.8	3.8	42.2	42.2	42.2

Note: This table shows the mean, minimum, and maximum values of decomposition exercises that change the ordering of the steps. There are six different possible orderings since sectoral reallocation is by definition always the last component. Therefore the mean, min, and max are by construction identical for sectoral reallocation.

Table A.15: Decomposing Differences in Average Hours Worked by Income Level

Panel A: Poor-Middle

	Hours			% Explained		
	Mean	Min	Max	Mean	Min	Max
Model	6.0			100.0		
Higher Productivity	3.3	3.1	3.5	55.0	51.6	58.4
Higher Taxes & Transfers	0.7	0.6	0.8	11.7	10.0	13.4
Lower Fixed Costs	-1.4	-1.6	-1.2	-23.3	-26.7	-20.0
Sectoral Reallocation	3.4	3.4	3.4	56.7	56.7	56.7

Panel B: Middle-Rich

	Hours			% Explained		
	Mean	Min	Max	Mean	Min	Max
Model	3.0			100.0		
Higher Productivity	2.2	1.8	2.6	73.3	60.0	86.7
Higher Taxes & Transfers	1.6	1.4	1.8	53.3	46.7	60.0
Lower Fixed Costs	-1.3	-1.8	-0.7	-43.3	-60.0	-23.3
Sectoral Reallocation	0.4	0.4	0.4	13.3	13.3	13.3

Note: This table shows the mean, minimum, and maximum values of decomposition exercises that change the ordering of the steps. There are six different possible orderings since sectoral reallocation is by definition always the last component. Therefore the mean, min, and max are by construction identical for sectoral reallocation.

Table A.16: Differences in Estimated Model Inputs over the Development Spectrum

	Poor vs. Middle		Middle vs. Rich	
	Diff.	Ratio	Diff.	Ratio
$\bar{u}_M$	-0.70	0.66	-0.49	0.64
A_M	27.20	4.54	65.11	2.87
A_T	4.19	2.77	6.78	2.03

B. Solution to the Household Problem (2)

We solve the second stage household problem (2) as follows. Plugging the optimal consumption c and hours $h^*(\eta^*)$ into the objective function, the household head's problem becomes an unconstrained problem:

$$\max_{\eta^*} \quad -\left[\alpha \frac{h^*(\eta^*)^{1+\frac{1}{\phi}}}{1+\frac{1}{\phi}} F(\eta^*) + \bar{u}_S \int_0^{\eta^*} \eta dF\right].$$

Taking the first order condition and applying the chain rule and the Leibniz rule leads to Equation (4). Since $h^*(\eta^*) = \frac{H}{F(\eta^*)}$, we have

$$\alpha \frac{\left(\frac{H}{F(\eta^*)}\right)^{1+\frac{1}{\phi}}}{1+\frac{1}{\phi}} f(\eta^*) + \bar{u}_S \eta^* f(\eta^*) = -\alpha \left(\frac{H}{F(\eta^*)}\right)^{\frac{1}{\phi}} \frac{-H}{F(\eta^*)} f(\eta^*).$$

After straightforward algebra, we get

$$\eta^* F(\eta^*)^{1+\frac{1}{\phi}} = \frac{1}{\bar{u}_S} \frac{\alpha}{1+\phi} H^{1+\frac{1}{\phi}}.$$

Assuming $\eta \sim U(0,1)$ and thus $F(\eta) = \eta$, we obtain a closed form solution (5) for the optimal cutoff η^*. Since η^* must be bounded by one from above, there is a maximum H for an interior solution that is given by

$$H = \left(\bar{u}_S \frac{1+\phi}{\alpha}\right)^{\frac{\phi}{1+\phi}}.$$

We thus have two cases. First, if H is smaller than this threshold, then η^* is given by an interior solution in Equation (5) and the household utility (3) becomes

$$\begin{aligned} U(C,H,S) &= \frac{C^{1-\gamma}}{1-\gamma} - \alpha \frac{H^{1+\frac{1}{\phi}}}{1+\frac{1}{\phi}} \left(\frac{1}{\bar{u}_S} \frac{\alpha}{1+\phi} H^{1+\frac{1}{\phi}}\right)^{\frac{-1}{1+2\phi}} - \frac{\bar{u}_S}{2} \left(\frac{1}{\bar{u}_S} \frac{\alpha}{1+\phi} H^{1+\frac{1}{\phi}}\right)^{\frac{2\phi}{1+2\phi}} \\ &= \frac{C^{1-\gamma}}{1-\gamma} - \left[\alpha \frac{\phi}{1+\phi} \left(\frac{1}{\bar{u}_S} \frac{\alpha}{1+\phi}\right)^{\frac{-1}{1+2\phi}} + \frac{\bar{u}_S}{2} \left(\frac{1}{\bar{u}_S} \frac{\alpha}{1+\phi}\right)^{\frac{2\phi}{1+2\phi}}\right] H^{\frac{1+\phi}{\phi} \frac{2\phi}{1+2\phi}}. \end{aligned}$$

Second, if H is larger than the threshold, then $\eta^* = 1$ and the household utility (3)

is simply given by

$$U(C,H,S) = \frac{C^{1-\gamma}}{1-\gamma} - \alpha \frac{H^{1+\frac{1}{\phi}}}{1+\frac{1}{\phi}} - \frac{\bar{u}_S}{2}.$$