

NBER WORKING PAPER SERIES

WHY ARE AVERAGE HOURS WORKED LOWER IN RICHER COUNTRIES?

Alexander Bick
Nicola Fuchs-Schündeln
David Lagakos
Hitoshi Tsujiyama

Working Paper 26554
<http://www.nber.org/papers/w26554>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2019

Fuchs-Schuendeln gratefully acknowledges financial support from the European Research Council under the ERC Consolidator Grant No. 815378. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2019 by Alexander Bick, Nicola Fuchs-Schündeln, David Lagakos, and Hitoshi Tsujiyama. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Why are Average Hours Worked Lower in Richer Countries?

Alexander Bick, Nicola Fuchs-Schündeln, David Lagakos, and Hitoshi Tsujiyama

NBER Working Paper No. 26554

December 2019

JEL No. E24,H2,O11

ABSTRACT

Why are average hours worked per adult lower in rich countries than in poor countries? Two natural candidates to consider are income effects in preferences, in which leisure becomes more valuable when income rises, and distortionary tax systems, which are more prevalent in richer countries. To assess the importance of these two forces, we build a simple model of labor supply by heterogeneous individuals and calibrate it to match international data on labor income taxation, government transfers relative to GDP, and hours worked per adult. The model predicts that income effects are the main driving force behind the decline of average hours worked with GDP per capita. We reach a similar conclusion in an extended model that matches cross-country patterns of labor supply along the extensive and intensive margins and of the prevalence of subsistence self-employment.

Alexander Bick
Department of Economics
W. P. Carey School of Business
Arizona State University
P.O. Box 879801
Tempe, AZ 85287-9801
alexander.bick@asu.edu

David Lagakos
Department of Economics, 0508
University of California, San Diego
9500 Gilman Drive
La Jolla, CA 92093
and NBER
lagakos@ucsd.edu

Nicola Fuchs-Schündeln
Goethe University Frankfurt
House of Finance
60323 Frankfurt
Germany
fuchs@wiwi.uni-frankfurt.de

Hitoshi Tsujiyama
Goethe University Frankfurt
hitoshi.tsujiyama@hof.uni-frankfurt.de

1. Introduction

Both time-series evidence and cross-country data on hours worked point to a pattern of lower average hours per adult in economies with higher income per capita. At the turn of the 20th century, U.S. adults worked an average of 28 hours per week, whereas nowadays, hours per adult have fallen to 24 hours per week (Ramey and Francis, 2009). Hours in European countries have also fallen systemically over the last century (Boppart and Krusell, 2018). For example, German adults worked around 28 hours in the 1950s compared to around 17 hours today (Ohanian et al., 2008). The cross-section of countries in the world today shows a remarkable similarity with these time series trends. In the bottom third of the world income distribution, adults work 29 hours per week on average, while in the top third adults average just 19 hours per week (Bick et al., 2018).

So why are average hours lower in richer countries? While there is no consensus in the literature, two hypotheses are natural candidates to consider. The first is income effects in preferences, which lead to an increase in the demand for leisure as individual wages rise. This view dates back at least to Keynes (1930), who argued that declining hours around the turn of the 20th century were due to higher income levels (see Ohanian, 2008, for a modern interpretation). Recently, the income-effects view has been embraced by Boppart and Krusell (2018), who reconcile the decrease in aggregate hours with standard balanced growth facts relying on a new specification of preferences in which income effects in labor supply dominate substitution effects. Similarly, Restuccia and Vandenbroucke (2013) adopt Stone-Geary preferences to capture income effects in labor supply and to explain the declining U.S. hours worked over the last century.

The second hypothesis is tax-and-transfer systems, which are much more extensive on average in richer economies (see e.g. Besley and Persson, 2014; Jensen, 2019). For example, Ohanian et al. (2008) argue that the declining pattern of hours in Europe since the 1950s is due almost entirely to rising tax rates. More generally, a large literature argues that Europe-U.S. differences in average hours worked can be traced to differences in taxation of labor income (e.g. Prescott, 2004; Rogerson, 2006, 2008; McDaniel, 2011; Bick and Fuchs-Schündeln, 2017, 2018; Bick et al., 2019). This literature has concluded that higher marginal tax rates discourage labor supply along the extensive margin and lower the hours of those employed. In the cross-country context, the idea is that the lower hours in richer countries may be explained at least in part by their more distortionary labor taxes and more generous transfers.

In this paper we quantitatively assess the importance of income effects relative to tax-and-transfer systems in driving the cross-country decline in average hours worked in GDP per capita. To do so, we build a simple model of labor supply by heterogeneous households with preferences in which income and substitution effects do not necessarily offset each other. Households face marginal tax rates on consumption and on labor income and receive lump-sum transfers from the government. Labor markets are competitive and labor is the only factor of production. Countries differ exogenously in two basic ways: first, in their levels of aggregate labor productivity, and second, in the size of their tax-and-transfer systems.

One key challenge in our quantitative analysis is that the literature does not have an agreed-upon estimate of the size of income effects relative to substitution effects in aggregate labor supply that we can take off the shelf to calibrate our preferences. Nor can one directly infer preferences from time-series or cross-country variation in aggregate hours worked, since features of tax-and-transfer systems also vary over time and across countries. Our approach then is to calibrate our model using cross-country evidence on statutory non-linear labor tax rates assembled by [Egger et al. \(2019\)](#). Given these labor and consumption tax rates across countries of different GDP per capita levels, we then calibrate the income effects in preferences and labor productivity differences to match average hours worked per adult across countries.

The model predicts that income effects explain the bulk – namely 77 percent – of the relationship between average hours worked and GDP per capita. Intuitively, the reason is that differences in labor productivity between poor and rich countries, and hence wage rates, are an order of magnitude larger than differences in tax systems. Between the top and bottom terciles of the world income distribution, labor productivity differs by a factor of over 16, whereas taxes as a fraction of GDP vary by a factor of only around 2. Germans face more taxes than Ghanaians, in other words, but these differences are dwarfed by the massive income gaps between Germany and Ghana. The model’s calibrated preferences feature income effects that dominate substitution effects only modestly, with an implied Marshallian elasticity of labor supply of -0.10. While simple, we show that the model is quantitatively consistent in matching another salient feature of the data, namely the within-country hours-wage gradients, which are negative for most countries but increase in income per capita.

Still, given the stylized nature of our model, it is natural to wonder how credible

its implications are for the importance of income effects and taxes in explaining aggregate hours worked. In particular, the model abstracts away from the extensive margin (employment rates) and intensive margin (hours per worker) of labor supply. [Bick et al. \(2018\)](#) show that the behavior of the two margins over the development spectrum is strikingly different. Employment rates fall rapidly between low- and middle-income countries, and then are flat or even slightly increasing towards the high-income countries. Hours per worker, on the other hand, are concave in income per capita, with a slight increase from poor to middle-income countries and then a marked decrease between middle- and high-income countries. The model also ignores the structural transformation from self-employment to market production that occurs over the development process ([Gollin, 2008](#)). Given that the self-employed may be constrained in how many productive hours they can supply ([Bandiera et al., 2017](#)), this structural transformation of how labor is supplied may be important for aggregate patterns of hours worked.¹

To address these concerns we extend the model to include an extensive margin (via a fixed cost of participation à la [Rogerson and Wallenius \(2013\)](#)) and an intensive margin, as well as a “traditional sector” based on subsistence self-employment characterized by a decreasing-returns production function. The “modern sector” allows household members to work in competitive labor markets for as many hours as they choose at the going wage. However, working in the modern sector comes at a higher fixed cost, capturing the fact that subsistence self-employment work is easily available, while accessing the modern sector is costly especially in poor countries, where it is often concentrated in cities. Subsistence self-employment is rarely taxed in practice ([Jensen, 2019](#)), so we assume that labor income taxes are only levied on the modern sector. As before, we calibrate the extended model to match the cross-country data on average hours per adult in countries belonging to the poorest and richest third of the world income distribution. We show that the model also replicates the convex shape of the extensive margin and the concave shape of the intensive margin, i.e. the facts for the middle-income countries, which we

¹The movement from self-employment to market work is one of the most salient features of the development process, though it has not been incorporated so far into the literature about the determinants of aggregate hours worked. [Ngai and Pissarides \(2008\)](#) and [Bridgman et al. \(2018\)](#) distinguish between market and non-market work, which is distinct but related, and most of the rest of the literature on structural change has focused on explaining how employment shares and value added move from agriculture to industry and then to services as countries grow richer (see e.g. [Herrendorf et al., 2014](#)). Dividing the economy into these three sectors is not important for our arguments, though, and in practice much of the work in each of these three sectors is subsistence self-employment in poor countries ([Gollin, 2008](#)).

do not target. Sectoral reallocation from the traditional into the modern sector is key to generate these shapes, as we explain in detail below. Reassuringly, similar to the simple model, the extended model also matches the individual-level gradient of hours worked on income, which is not targeted at all. Both of these successes in matching the data lend credence to the model's disaggregate predictions.

The extended model predicts that income effects explain more than the entire decrease in hours worked with development. The reason income effects dominate here is twofold. First, it is the same as in the simple model: cross-country variations in tax rates are dwarfed by the variation in labor productivity levels. Secondly, faster productivity growth in the modern than in the traditional sector leads families to switch into the modern sector, where employment rates are lower. This is an additional force for decreasing hours per adult which comes through sectoral reallocation. For taxes and transfers, sectoral reallocation works in the opposite direction, because increasing labor income taxes by development decrease the attractiveness of modern sector work. The implied Marshallian elasticity is now -0.18 : sectoral reallocation strengthens the effects of increasing labor productivity on hours worked. We conclude that while tax and transfer systems may be the primary factor explaining differences in hours worked among advanced economies, the decline in work hours that comes with development is mostly accounted for by income effects in preferences. Moreover, while structural transformation and fixed costs are clearly important for understanding the behavior of the extensive and intensive margins of labor supply across countries, these features do not fundamentally alter our conclusions about drivers of cross-country differences in aggregate labor supply.

The rest of this paper is structured as follows. Section 2 presents our model of aggregate labor supply, calibrates it to international data on taxes and hours worked, and explores the relative roles of income effects and taxes and transfer systems in explaining cross-country differences in labor supply. Section 3 reviews the cross-country facts about the extensive and intensive margins of labor supply and subsistence self-employment. Section 4 presents the extended model that is enriched to include two margins of labor supply and structural change, and Section 5 assesses the quantitative importance of income effects and taxes in this model. Last, Section 6 concludes.

2. A Simple Model of Labor Supply

In this section, we present a simple model of labor supply featuring households that we use to learn about the roles of income effects and tax-and-transfer systems in explaining cross-country differences in average hours worked per adult. The model features preferences in which income effects and substitution effects in labor supply do not necessarily cancel out when income rises, as in standard balanced-growth preferences. Households are heterogeneous in their individual productivity levels. Labor and consumption are taxed at the margin and a portion of tax collections are returned to the households as transfers. These features allow us to calibrate the model and quantify the sources of cross-country differences in labor supply. Due to the within-country heterogeneity, the model generates predictions not only for the relationship of hours and income on the aggregate level, but also for within-country hours-wage gradients. Moreover, heterogeneity gives a meaningful role to progressive taxation, which varies systematically by development.

2.1. Environment

Output is produced using a constant-returns production function with labor as the sole input: $Y = AL$, where A represents aggregate labor productivity and L is aggregate effective hours. Labor and output markets are perfectly competitive. Countries differ in the level of aggregate labor productivity, A , though we leave off country subscripts for convenience.

Each country is populated by a measure one of heterogeneous households that differ only in their productivity of labor. Formally, each household is endowed with one unit of time and has individual productivity z with $\log z \sim N(0, \sigma_z^2)$. Each household makes a labor-leisure choice given the following preferences (MaCurdy, 1981; Keane, 2011):

$$u(c, h) = \frac{c^{1-\gamma}}{1-\gamma} - \alpha \frac{h^{1+\frac{1}{\phi}}}{1+\frac{1}{\phi}}, \quad (1)$$

where c and h are individual consumption and hours worked, $\gamma \geq 0$, and $\phi \geq 0$. The parameter γ governs the strength of the income effects, and ϕ the curvature of disutility in hours worked. In a dynamic setting, ϕ represents the Frisch elasticity of labor supply, meaning the percentage change in hours that comes from a one-percent increase in the wage holding fixed the marginal utility of consumption. Boppart and Krusell

(2018) show that these preferences are consistent with balanced growth even if the income effect dominates, i.e. if $\gamma > 1$. As in [Vandenbroucke \(2009\)](#), the within-country heterogeneity in our model generates predictions for how hours vary with income for different income groups that can be tested in the data.

The government applies a progressive income tax $\mathcal{T}(y)$ to household income $y = zwh$, and a linear consumption tax τ_c on consumption, c . The government spends G on “government consumption expenditure,” which does not enter the households’ utility function. The remaining government tax collections are redistributed as lump-sum transfers Υ to households such that the budget is balanced. These transfers may represent direct transfers to households and/or public goods provided by the government that are a substitute for private consumption expenditures.

The government’s budget constraint in equilibrium is given by:

$$G + \Upsilon = \int [\mathcal{T}(y(z)) + \tau_c c(z)] dF_z, \quad (2)$$

where F_z is the cumulative distribution function of z . The household’s budget constraint in turn is:

$$(1 + \tau_c)c = zwh - \mathcal{T}(y) + \Upsilon. \quad (3)$$

Two sets of variables differ exogenously by development: (i) the aggregate productivity level, A ; and (ii) the size of the tax-and-transfer system $(\tau_c, \mathcal{T}, G, \Upsilon)$. We then characterize the equilibrium over a range of A values, representing countries at different levels of the development spectrum.

2.2. Tax-and-Transfer Systems

To parameterize the model, we need to discipline the size of tax-and-transfer systems across countries. 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 have already assembled a comprehensive database of statutory tax rates 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 put forth by [Bénabou \(2002\)](#) with net income \tilde{y} being given by

$$\tilde{y} = y - \mathcal{T}(y) = y - (y - \lambda y^{1-\tau}) = \lambda y^{1-\tau}, \quad (4)$$

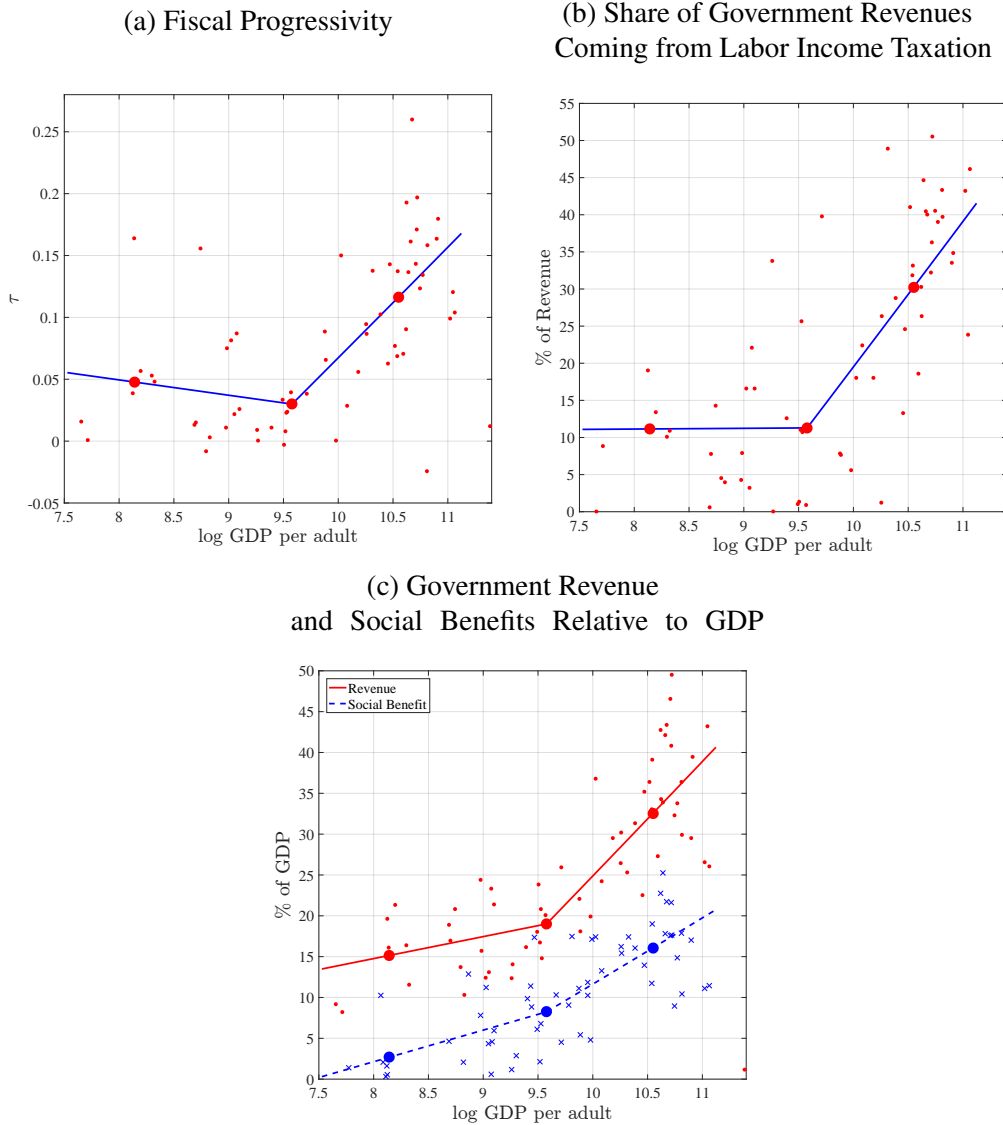
where λ is informative about the level of taxation and τ about the progressivity. For $\tau = 0$, $1 - \lambda$ represents a proportional tax on income, whereas for $\tau = 1$, net income is independent of gross income. We estimate τ for each country based on the data set compiled by [Egger et al. \(2019\)](#). Specifically, for each country they shared average gross incomes at each percentile of the income distribution and the implied net income, where the latter is calculate for a single individual without children using statutory tax codes excluding any transfers that are not incorporated directly into the tax system. Taking logs of Equation (4), 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. \(2019\)](#).²

We then 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 calibration 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/Y equal to the share of social benefits over GDP, which we obtain from the IMF government statistics.

Figures [1a](#), [1b](#), and [1c](#) plot the resulting components of the tax system that we use in the calibration. In the calibration, we do 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 suitable for studying how hours change with income across countries of different GDP per capita levels, which 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

²We 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. \(2019\)](#).

Figure 1: 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.

increase from middle- to high-income countries (Figure 1a). 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 1b). The estimate of government revenue relative to GDP increases somewhat from the poor to middle-income countries and then sharply from the middle to the richest countries (Figure 1c). Similarly, the size of government transfers also increases over the

development spectrum.

How does the overall burden of taxes vary by income on average across countries? Figure 1c provides one answer to this question. 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.

2.3. Calibration

With the tax data in hand, we proceed with the calibration of the model, using average targets from the poorest and richest terciles of countries in the world income distribution. We first normalize the level of labor productivity in the richest tercile of countries, A^{rich} , to be one. We set the value of $\phi = 1$ which gives us an elasticity in line with estimates for the extensive and intensive margin in the literature (see Keane, 2011). Last, we normalize $\alpha = 1$. The distribution of individual labor productivity is estimated using the panel component of the Current Population Survey (CPS), and is assumed to capture permanent differences across households. We estimate a panel fixed effect regression of log income per hour on individual fixed effects, following Lagakos and Waugh (2013), and take the variance of the individual effects to be the variance of the permanent productivity differences in our model.

We then jointly calibrate γ and the labor productivity level in the poorest tercile of the world income distribution, A^{poor} , to match average hours worked per adult and GDP per adult in the average poor country relative to the average rich country. Thus, by construction the model is calibrated to explain the entire decline in hours worked present in the data. The parameter γ determines the size of income effects in preferences, and is identified from the hours differences between the poorest and richest countries not explained by differences in tax systems.

Table 1 shows the calibration targets and the calibrated parameters. We take the data on hours worked from Bick et al. (2018), in which we carefully construct these data for 49 countries, and also report averages by terciles of the world income distribution.³ Hours worked include any hours spent on producing output counted in NIPA;

³There are two slight differences in the data used in this paper and in Bick et al. (2018). First, we drop Laos, because it lacks information on self-employment, which we need for the calibration of the extended

Table 1: Calibrated Parameters and Calibration Targets

	Parameter	Value	Target	Data	Model
γ	curvature of consumption	1.23	avg. hours per adult in poor vs. rich countries	1.48	1.48
A^{poor}	productivity in poor countries	0.06	avg. output per adult in poor vs. rich countries	0.09	0.09

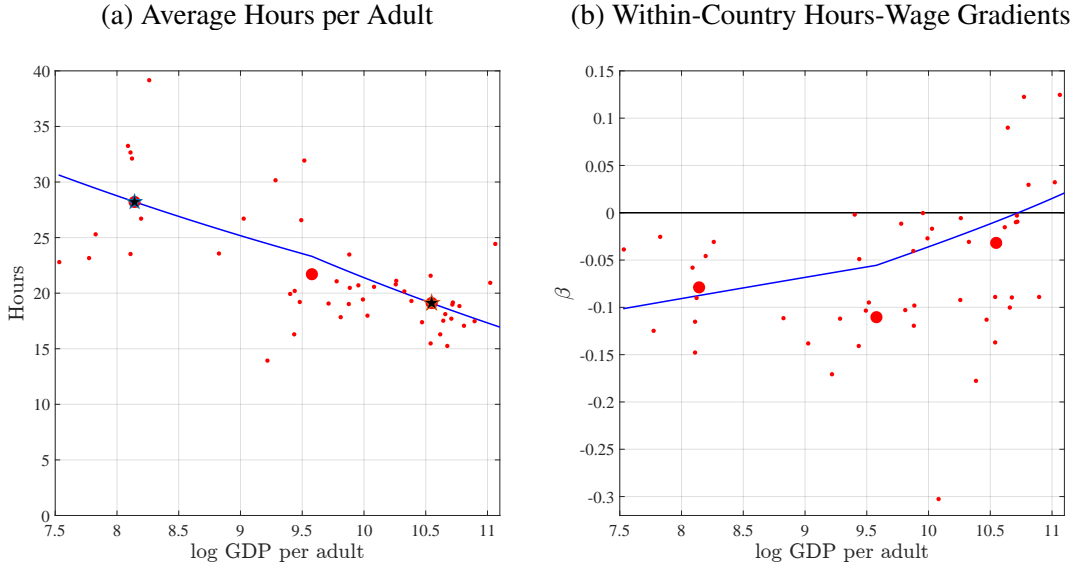
especially, hours spent in informal work, self-employment, or production of goods for self-consumption are also included in this measure. Data on output per adult come from the Penn World Tables. Average hours worked per adult in the poorest tercile are 28.2 hours, and in the richest tercile 19.1 hours (see Table 3), such that the ratio of hours in poor to rich countries amounts to 1.48. Average output per adult in the poorest tercile is only 0.09 times of average output per adult in richest tercile. The resulting calibrated parameter values are $\gamma = 1.23$ and $A^{poor} = 0.062$. The model’s value of γ is higher than the standard value of 1 used to generate balanced growth with constant hours (King et al., 1988). This points to an important role for income effects in generating the decline in hours present in the data.

As shown by Keane (2011), an advantage of our simple preference specification is that it allows for closed-form solutions for several common elasticities of labor supply studied by the literature. Our analysis is most informative about the Marshallian elasticity of labor supply, which summarizes how hours respond to a one-percent permanent increase in wages. The Marshallian elasticity in our model is $(1 - \gamma)/(1/\phi + \gamma)$, which is -0.10 at our calibrated parameter values.⁴ Thus, the overall effect of a 10 percent increase in the wage rate is a decrease in labor supply of around 1 percent. Our estimate lies within the large range of estimates reported in 22 studies studied by Keane (2011), which range from -0.47 to 0.51, and below his average of 0.06. None of these studies use cross-country aggregate evidence to measure a Marshallian elasticity, though, as

model. Second, we define terciles of the world income distribution based on output per adult rather than output per capita. We have data on 10 countries in the poorest tercile, 15 countries in the middle tercile, and 23 countries in the richest tercile.

⁴We get exactly the same estimate for the Marshallian elasticity when regressing average predicted hours from the “income effect” experiment in Table 2 on average wages on the country level. In the extended model, we calculate Marshallian elasticities based on such an estimation, see Section 5.

Figure 2: Key Facts: Simple Model vs. Data



Note: The small red dots represent each country in our sample, and the large red dots the average by country income group. The model predictions are displayed by the blue line. Targeted moments are marked with a star. The red dots in the right panel are the beta-coefficients from the following regression, run separately for each country on the sample of workers in paid employment: $\log(h_i) = \alpha + \beta \log(w_i) + \delta_1 \text{age}_i + \delta_2 \text{age}_i^2 + \varepsilon_i$.

we do. Hence, our estimate might be more useful in the context of large productivity differences.

Figure 2a plots the model's predictions for average hours by income level against the data (small red dots).⁵ The large red dots are average hours by income tercile, and the stars are the calibration targets. Overall, the model matches the decline of hours per adult in log GDP per adult well.

One way to cross-check the model's income effect on labor supply is to compare its predictions for the within-country gradients of hours worked to wages to their counterparts in the data. Bick et al. (2018) estimate these hours-wage gradients for 46 countries and document that they are negative in poorer countries and increase with GDP per capita. Only some of the world's highest income countries, like the United States and the United Kingdom, have positive gradients. In other words, within most countries in the world, individuals with lower wages work more hours than those with higher wages. As countries get richer, the gap in hours of those with lower and higher wages closes,

⁵We match predictions from the model to actual countries from the data by output per adult.

and even reverses in some advanced economies. This reversal has been documented in the U.S. history as well, see e.g. [Costa \(2000\)](#), [Aguiar and Hurst \(2007\)](#), and [Heathcote et al. \(2014\)](#).

Figure 2b plots the model’s predictions for hours-wage gradients by income level (blue line) and in the cross-country data (red dots).⁶ The model does quite well in matching the levels and cross-country variation in hours-wage gradients, even though we did not target them in any way. The model’s predicted gradient of -0.07 for the average of the poorest tercile of countries (the thick red dot on the left) is quite close to the empirical value. For the middle tercile, the gradient is negative as in the data, but somewhat higher than in the data. The richest tercile gradient is close to zero both in the model and data, and the model accurately predicts the modestly positive gradients for the very richest countries.

The fact that the model reproduces the negative hours-wage gradients within most countries provides some validation to the model’s quantitative predictions for income effects on labor supply. Perhaps more subtly, the increasing gradients with GDP per capita help corroborate the model’s calibrated tax and transfer systems. In the model, the rising gradients with development come entirely from the rising and increasingly progressive tax systems that come with development (highlighted in Figures 1a, 1b, and 1c). The main reason is that less productive households respond to the larger government transfers by lowering their hours more than those with higher productivity. In the richest countries, transfers are so large relative to potential wage income for the least productive workers that they work even less than those with higher productivity, in spite of strong income effects in preferences.

2.4. Decomposition

The purpose of the calibrated model is that it can help decompose the importance of income effects and tax-and-transfer systems in explaining the decrease in aggregate hours. We decompose these two forces using two counterfactual exercises. The first takes the tax-and-transfer system of the poorest countries and keeps that fixed while raising aggregate labor productivity to the level of the richest countries. This counterfactual simulates how hours worked would look across the world income distribution

⁶In [Bick et al. \(2018\)](#), we report the coefficients separately for men and women, show that they line up well quantitatively with time-series estimates from the US provided in [Costa \(2000\)](#), and show that they are similar when self-employed workers are included in the sample.

Table 2: Counterfactual Experiments in the Simple Model

	Poor	Rich	Diff.	% Expl.
Data	28.2	19.1	9.1	
Model	28.2	19.1	9.1	
Income Effect	28.2	21.2	7.0	77%
Taxes & Transfers	28.2	25.2	3.0	33%

Note: This table reports average hours worked per adult in the poorest and richest terciles of the world income distribution in the data, the calibrated simple model, and two counterfactual experiments. The first, ‘Income Effect,’ varies only labor productivity, A , across countries but holds fixed the size of tax-and-transfer systems. The second, ‘Taxes & Transfers,’ holds fixed A and varies the size of the tax-and-transfer systems. The last column reports the percent of the difference in the data explained by the model under each counterfactual.

assuming only an income effect. The second takes the income level of the poorest countries as fixed, but changes the size of the tax-and-transfer system with actual GDP per adult. The purpose of the second counterfactual is to compute hypothetical hours worked by income level assuming that tax systems are the only source of variation across countries.⁷

Table 2 shows the results of these counterfactuals. In the first two rows, the table reports hours worked in the data and model for the average poor and the average rich country, as well as their simple difference, which amounts to 9.1 hours. The following two rows then provide the results from the decomposition exercise, with the third row indicating the decrease in hours between the average poor and the average rich country that can be attributed to income effects. The fourth row indicates this decrease for the fiscal inputs. The columns state the predictions for the average low- and high-income country from each decomposition, the predicted change between the average poor and rich country, and the percent the respective decomposition contributes to the total predicted change in the model. By construction, average hours in the poor country are always matched. Income effects alone explain 77 percent of the difference in hours between the poorest and richest tercile in the data. Taxes and transfers explain 33 percent by themselves. Note that because the income effect and taxation interact non-linearly,

⁷Our experiments take into account that progressive tax codes are always defined relative to the average income of a country. Appendix B.1 provides the details of how we account for this feature.

the two numbers do not add up to 100 percent.

Why do income effects explain more than taxes? The intuition comes from comparing the magnitude of the productivity (and hence income) differences across countries to the tax differences. Labor productivity differences between the richest tercile of countries and poorest tercile are a factor of 16.2 ($= A^{rich}/A^{poor} = 1/0.062$). As discussed in Section 2.2, the overall tax burden is about 2.2 times larger in the richest tercile than the poorest. The implied average consumption tax rate rises from 15.4 percent in the poorest tercile to 27.2 percent in the richest, and the average labor income tax rate, measured as labor income tax divided by income, from 1 percent to 7.6 percent. Nevertheless, even these tax differences could in principle explain all of the hours difference leaving no role for income effects if labor supply would be much more elastic than implied by our calibration. However, for our parameter estimates, which are in line with conventional elasticities of labor supply, the enormous changes in income are dominant in explaining the decline in hours across the world GDP per adult distribution.⁸ As a robustness check, we also calibrate an alternative version of the model in which we assume full redistribution of government revenues. This maximizes the negative effect of taxes on hours by maximizing the outside income provided by transfers. In this robustness check, taxes and transfers explain 40 percent of the decrease in hours, and income effects 71 percent. Thus, the model robustly attributes the majority of the decrease in hours to income effects.

3. Disaggregate Patterns of Labor Supply Across Countries

While the model above provides a simple way to account for aggregate hours differences across countries, it necessarily abstracts from the rich disaggregate patterns of labor supply in the data. To start with, the extensive and intensive margins behave quite differently along the development process (Bick et al., 2018). Moreover, labor supply in developing countries is largely directed toward subsistence self-employment activities. As countries become richer, they undergo structural change which takes workers out of self-employment and into market wage work. In this section we review these disag-

⁸For the evaluation of tax changes the Hicksian elasticity, which summarizes how hours respond to a change in wages holding the level of income fixed, is particularly relevant, see Keane (2011). The Hicksian elasticity in our model is given by $1/(1/\phi + \gamma)$, and is 0.45 at our calibrated parameter values. This is line with the estimates from 22 studies reported by Keane (2011), which range from 0.02 to 1.32 with an average of 0.31, and comparable to the averages of 0.50 and 0.59 reported for micro studies and macro studies by Chetty et al. (2012).

Table 3: Extensive and Intensive Margins by Income Group

	Low	Middle	High
Hours per Adult	28.2	21.7	19.1
Employment Rate (Extensive)	74.5	52.4	54.6
Hours per Worker (Intensive)	38.4	41.3	35.1

Note: This table reports average weekly hours worked per adult, average employment rates, and average weekly hours worked per worker by country income group. Source: [Bick et al. \(2018\)](#).

aggregate patterns of labor supply and discuss their potential roles in shaping aggregate labor supply across countries. We then turn in the following section to a model that can potentially match the disaggregate facts and can be used to re-assess the importance of income effects and taxes in explaining aggregate hours differences.

We begin with the cross-country patterns of labor supply along the extensive and intensive margins, drawing on the findings from [Bick et al. \(2018\)](#). Table 3 reports average employment rates (the extensive margin) and hours per worker (the intensive margin), as well as average hours per adult, for three country income groups: those belonging to the poorest, middle, and richest terciles of the world income distribution. While both margins of labor supply show a decrease between low- and high-income countries, they behave very differently over the entire development spectrum. Employment rates fall strongly between low- and middle-income countries, namely by 22.1 percentage points, but then slightly increase towards the high-income countries. Hours per worker, on the other hand, show a slight increase between low- and middle-income countries, but then fall by 6.2 hours per week between middle- and high-income countries.⁹ Last, Table 3 shows that the decrease in the employment rate between poor and rich countries is quantitatively more important than the decrease in hours per worker for the total decrease in hours per adult: employment rates generate three quarters of the total decrease (see [Bick et al. \(2018\)](#)).

We turn next to subsistence self-employment rates across countries. Empirically, we focus on self-employed individuals with low education, which is a close proxy for subsistence work, though certainly not exact, and something we can measure in a com-

⁹In [Bick et al. \(2018\)](#), we perform permutation tests and show that none of the two increases in the two margins of labor supply is statistically significant, while both of the decreases are.

Table 4: Sectoral Hours Worked and Sectoral Shares

	Country Income Group		
	Low	Middle	High
Traditional Sec. Hours	35.4	36.6	39.2
Market Sec. Hours	46.3	42.3	35.0
Traditional Sec. Share	64.3	18.6	5.7

parable way across the countries in our data. For comparison with our model later, we define these workers as the “traditional sector,” and the balance to be the “modern sector.”

Table 4 shows in the first two rows the average hours per worker in the traditional and the modern sector, respectively, separately for the three country income groups. Looking across columns, we find that hours worked per worker are 3.8 hours higher in rich than in poor countries in the traditional sector. By contrast, they are 11.3 hours lower in rich than in poor countries in the modern sector. Thus, hours per worker are strongly decreasing in development in the modern sector, and slightly increasing in the traditional sector. As a result, looking across rows, for the poor and middle-income countries hours are markedly lower in the traditional than in the modern sector, namely by 10.9 and 5.7 weekly hours, respectively. Only for the rich countries are hours higher in the traditional sector, with a difference of 4.2 hours. The last row of Table 4 shows the share of all workers working in the traditional sector: In the poor countries, almost two thirds of workers (64.3 percent) work in the traditional sector. This share rapidly decreases to 18.6 percent in the middle-income countries, and only 5.7 percent in the high-income countries. Thus, over the development process, workers shift quickly from the traditional into the market sector.

Taking the patterns of sectoral hours worked per worker and sectoral shares of workers together, it becomes clear that the modest increase of 2.9 weekly hours worked per worker between low- and middle-income countries documented in Table 3 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 market 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 in-

crease in average hours worked per worker.¹⁰ 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.

Appendix Tables A.1 to A.7 show that the patterns presented in Tables 3 and 4 are similar among men and women, and across broad age groups. Hence, we abstract from age and sex in the model that follows.

4. Extended Model of Aggregate and Disaggregate Labor Supply

We now extend the aggregate model to feature the rich disaggregate patterns of labor supply discussed in the section above. In particular, we add a household labor supply decision that includes an extensive and intensive margin, and a traditional sector featuring self-employment, in addition to the modern sector with labor supplied in competitive markets. We model families as an informal insurance mechanism to generate the employment rates well below 1 also in the poorest countries. We calibrate the model to match the aggregate and disaggregate patterns described above for the average poor and average rich country, and show that it also replicates the facts for the middle-income countries. We then use the model to re-assess the roles of income effects and taxes in driving aggregate labor supply.

4.1. Environment

There is a continuum of families of mass one in each country, and a continuum of individuals of mass one in each family. We assume perfect insurance within a family, and no insurance across families (see [Heathcote et al., 2014](#)). Families are meant to capture different kinds of informal insurance networks within a country, which might exist not only within families, but also within villages or other groups (see e.g. [Townsend, 1994](#), and [Fafchamps and Lund, 2003](#)).

Families differ in their modern sector productivity z with $\log z \sim N(0, \sigma_z^2)$.¹¹ Indi-

¹⁰The small increase in hours per worker in the traditional sector marginally adds to this increase.

¹¹Similarly, [Porzio and Santangelo \(2017\)](#) assume that human capital does not matter in the agricultural sector, but only in the non-agricultural one. [Caselli and Coleman \(2006\)](#) introduce a menu of different production technologies that differ in the use they make of skilled and unskilled labor.

viduals within a family differ only in their individual fixed disutility of work η . The instantaneous utility function of an individual is

$$u(\tilde{c}, \tilde{h}; \eta) = \frac{\tilde{c}^{1-\gamma}}{1-\gamma} - \alpha \frac{\tilde{h}^{1+\frac{1}{\phi}}}{1+\frac{1}{\phi}} - \bar{u}_S \eta I_{\tilde{h}>0}, \quad (5)$$

where \tilde{c} and \tilde{h} are individual consumption and hours worked, \bar{u}_S is the utility cost of working a positive number of hours, which depends on the sector the family is working in, and $I_{\tilde{h}>0}$ is an indicator equal to 1 if the individual works. In what follows, variables c and h with a tilde refer to the individual level, and without a tilde to the family level. With the exception of the fixed cost of working, individual preferences thus take the same form as in the model described in Section 2.

Each family is headed by a family head who maximizes the sum of the utility of all family members with equal weight. The family can decide to work either in the traditional or in the modern sector. In the modern sector, family income is equal to the hourly wage w times effective family hours worked (i.e. family hours multiplied with market productivity z). The modern sector features a constant returns to scale technology. The traditional sector by contrast features a decreasing returns to scale technology, and family income in the traditional sector equals $y_T = A_T h^\rho$ with $\rho < 1$, where A_T is the traditional sector labor productivity. The decreasing returns to scale technology captures the partial absence of land, labor, and capital markets in developing countries (see e.g. [Jayachandran, 2006](#), or [Karlan et al., 2014](#)).¹²

4.2. Equilibrium Analysis

Family's Problem The family head faces a two-stage maximization problem. In a first stage, she chooses family hours h , consumption c , and the sector of employment S . In a second stage, given family hours and consumption, she chooses individual hours \tilde{h} and consumption \tilde{c} . We solve the maximization problem by backward induction.

¹²[Storesletten et al. \(2017\)](#) present a model of structural transformation in which only the agricultural sector is split into a traditional and a modern sector. [Boppart et al. \(2019\)](#) analyze differential effects of productivity growth in agriculture - which is characterized by a nested CES production function - and non-agriculture in a general equilibrium set-up.

Given (c, h, S) , the second stage maximization problem amounts to

$$\begin{aligned} \max_{\{\tilde{c}(\cdot), \tilde{h}(\cdot)\}} \quad & \int u(\tilde{c}(\eta), \tilde{h}(\eta); S, \eta) dF(\eta) \\ \text{s.t.} \quad & \int \tilde{c}(\eta) dF(\eta) = c \\ & \int \tilde{h}(\eta) dF(\eta) = h, \end{aligned} \tag{6}$$

where F is the CDF of η .

The first order condition for consumption implies perfect consumption risk sharing within the family, i.e. $\tilde{c}(\eta) = c$ 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 family in optimal hours worked conditional on working. The optimal hours function thus can be expressed as

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

The family head's problem therefore reduces to determining a threshold level $\eta^*(z)$: all family members with a disutility of work below this threshold level work the same positive hours $\tilde{h}^*(\eta^*) = \frac{h}{F(\eta^*)}$, and all family members with a disutility above this threshold level do not work. Given family hours h , individual hours worked are decreasing in the threshold level, $\frac{d\tilde{h}^*}{d\eta^*} < 0$.

Substituting the optimal decisions into the objective function of the problem (6) gives the family utility:

$$U(c, h) \equiv \frac{c^{1-\gamma}}{1-\gamma} - \alpha \frac{h^{1+\frac{1}{\phi}}}{1+\frac{1}{\phi}} (F(\eta^*))^{-\frac{1}{\phi}} - \bar{u}_S \int_0^{\eta^*} \eta dF. \tag{7}$$

Note that the family utility looks different from the individual utility (5), see [Constantinides \(1982\)](#). In the first stage, the family head solves the following maximization problem of the family:

$$\begin{aligned} \max_{c, h, S \in \{T, M\}} \quad & U(c, h) \\ \text{s.t.} \quad & (1 + \tau_{c, S})c = y_S - \mathcal{T}_S(y_S) + Y, \\ \text{where} \quad & y_M = wz h \text{ and } y_T = A_T h^\rho \end{aligned} \tag{8}$$

where the taxes paid may depend on the sector. We denote the solution to the family's problem by $\{c(z), h(z), S(z)\}_{z \in \mathbb{R}_{++}}$.

Equilibrium Wage The competitive market sector clears such that

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

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:

$$G + \Upsilon = \sum_{S=T,M} \left[\int (\mathcal{T}_S(y(z)) + \tau_{c,Sc}(z)) 1_{\{S(z)=S\}} dF_z \right]. \quad (9)$$

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

- (i) given the price and policies, the decision rules solve families' problems (6) and (8),
- (ii) the marginal profit condition is satisfied: $w = A_M$, and
- (iii) the government budget constraint (9) is satisfied.

4.3. 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 $(\tau_{c,T}, \mathcal{T}_T(\cdot), \tau_{c,M}, \mathcal{T}_M(\cdot), G, \Upsilon)$, and (iii) the utility cost of working in the modern sector \bar{u}_M , which captures the variation of accessibility of modern sector jobs. We solve steady-states for each level of development.

While differential labor productivity growth between sectors is a standard force of structural change, different fixed costs of working in the modern sector are not. Why do we introduce this force of structural change? First, fixed costs of working in the modern sector are a plausible source of cross-country heterogeneity. In poor countries, the modern sector is often allocated exclusively in cities, making it very costly for a large part of the population who live in rural areas to access it. Moreover, getting the paperwork done

to register as a formal worker is likely also more complicated. Secondly, as we explain in Section 4.4.3, without this extra degree of freedom, it is still possible to replicate the decrease in hours worked per adult by development. However, we fail to replicate the differential behavior of the two margins of labor supply between the middle- and high-income countries.

4.4. Taking the Model to the Data

We calibrate the model parameters to data coming from the countries belonging to the poorest and richest third of the world income distribution. We then ask whether the model can explain the patterns we see over the full development spectrum, both across and within countries. In the baseline calibration, we assume that labor income taxes are only raised on modern sector incomes, due to enforceability problems in the traditional sector. At the end of Section 5, we show that our main results are robust to different assumptions on sectoral taxation.

4.4.1. Exogenous Model Inputs

We assume that the individual fixed utility cost of working is uniformly distributed with $\eta \sim U(0, 1)$, which allows us to solve the second stage of the family head maximization problem in closed form (see Appendix B.2). The implementation of the tax-and-transfer system and the estimation of the distribution of the family-level modern sector productivity z follow the descriptions in Sections 2.2 and 2.3.

4.4.2. Calibration

We introduce a few normalizations. First, we normalize $A_M^{rich} = 1$, i.e. the average modern sector productivity in countries belonging to the richest third of the world income distribution is set to 1. Secondly, we normalize $\bar{u}_T = 0$, i.e. there are no fixed cost associated with working in the traditional sector.¹³ Last, we again normalize $\alpha = 1$.

Given these normalizations, all other parameters are jointly calibrated to replicate key moments from the data. As in the simple model in Section 2, we construct our calibration targets for the “average” poor and “average” rich country, not using any targets from the middle-income countries.

¹³Since there is no empirical counterpart of the employment rate of a family in a given sector, we identify the fixed cost of working from the overall employment rates. As a consequence, all that matters are the fixed cost of working in the modern sector compared to the fixed cost of working in the traditional sector.

Table 5: Calibrated Parameters and Calibration Targets

Parameter		Value	Target	Data	Model
\bar{u}_M^{poor}	fixed cost of working in M sector in poor countries	5.55	avg. ER in poor countries	74.5	72.6
\bar{u}_M^{rich}	fixed cost of working in M sector in rich countries	2.23	avg. ER in rich countries	54.6	54.6
A_T^{poor}	T sector productivity in poor countries	0.03	avg. % of workers in T sector in poor countries	64.3	64.8
A_T^{rich}	T sector productivity in rich countries	0.14	avg. % of workers in T sector in rich countries	5.7	5.8
ρ	DRS parameter in T sector	0.72	avg. hours per worker in T sector in poor countries	35.4	35.1
ϕ	curvature of disutility of working	0.45	avg. hours per worker in M sector in poor countries	46.3	46.0
γ	curvature of consumption	1.23	avg. hours per worker in rich countries	35.1	35.1
A_M^{poor}	M sector productivity in poor countries	0.08	avg. output per adult in poor vs. rich countries	0.09	0.09

While all parameters are jointly calibrated, some moments are more informative for some parameters than others. In the following, we provide some informal discussion, acknowledging that these arguments are of course not a formal proof of identification. We have in total 8 free parameters and choose 8 moments to be replicated. The fixed costs of working in the modern sector in poor and rich countries are crucial for matching the employment rates in poor and rich countries, respectively. Productivities in the traditional sector in poor and rich countries are (among other things) informative about the fraction of workers in the traditional sector in poor and rich countries, respectively. The decreasing returns to scale parameter in the production technology in the traditional sector is crucial for determining optimal hours in this sector, and we choose the corresponding hours per worker in the poor countries as a target. Obviously, the curvature of the disutility of hours worked affects any hours choice, and thus also the hours per worker in the modern sector in poor countries. The curvature of the consumption func-

tion, i.e. the strength of the income effect, is chosen such that hours per worker in rich countries, or respectively the decrease relative to poor countries, are replicated. Finally, we set productivity in the modern sector in poor countries such that we match output per adult in poor relative to rich countries. The choice of these moments ensures that both labor input and output in poor and rich countries are consistent with the data.¹⁴

Table 5 states the targeted moments, which the model matches near perfectly, and lists the calibrated parameter values. The fixed cost of working in the modern sector is 2.5 times ($5.55/2.23$) higher in poor than in rich countries. The returns to scale parameter is with 0.72 close to the estimate of [Guner et al. \(2008\)](#), although the setups are not fully comparable. In poor countries, productivity in the modern sector is around 3 times ($0.08/0.03$) larger than in the traditional sector. In rich countries, this difference is more than three times as large ($1/0.14$). Comparing poor and rich countries, the productivity gap amounts to a factor of over 12 in the modern sector, and less than 5 in the traditional sector. The calibrated curvature parameter for the disutility of working is with 0.45 consistent with the intensive-margin estimates surveyed in [Blundell and MaCurdy \(1999\)](#), [Domeij and Flodén \(2006\)](#), and [Keane \(2011\)](#). Finally, our calibrated value of the curvature of consumption is 1.23, exactly as in the simple model. It thus implies that income effects dominate, but at the same time is not too far from the log specification.

4.4.3. Model Fit

Before conducting the decomposition exercise, we analyze the model fit of the extended model. While we calibrate the model to aggregate moments from the average low- and high-income countries, the shapes of different variables over the entire development spectrum, especially the differential shapes of the two margins of labor supply, are non-targeted moments.

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 increase linearly in $\log(A_M)$. Secondly, as in the simple model, we assume all fiscal inputs (i.e., tax progressivity, share of government revenue coming from labor income taxes, ratio of government rev-

¹⁴Note that the only non-targeted statistics in that regard are hours per worker separately by sector in rich countries. We refrain from targeting them separately because aggregate hours per worker in rich countries are effectively only determined by hours per worker in the modern sector.

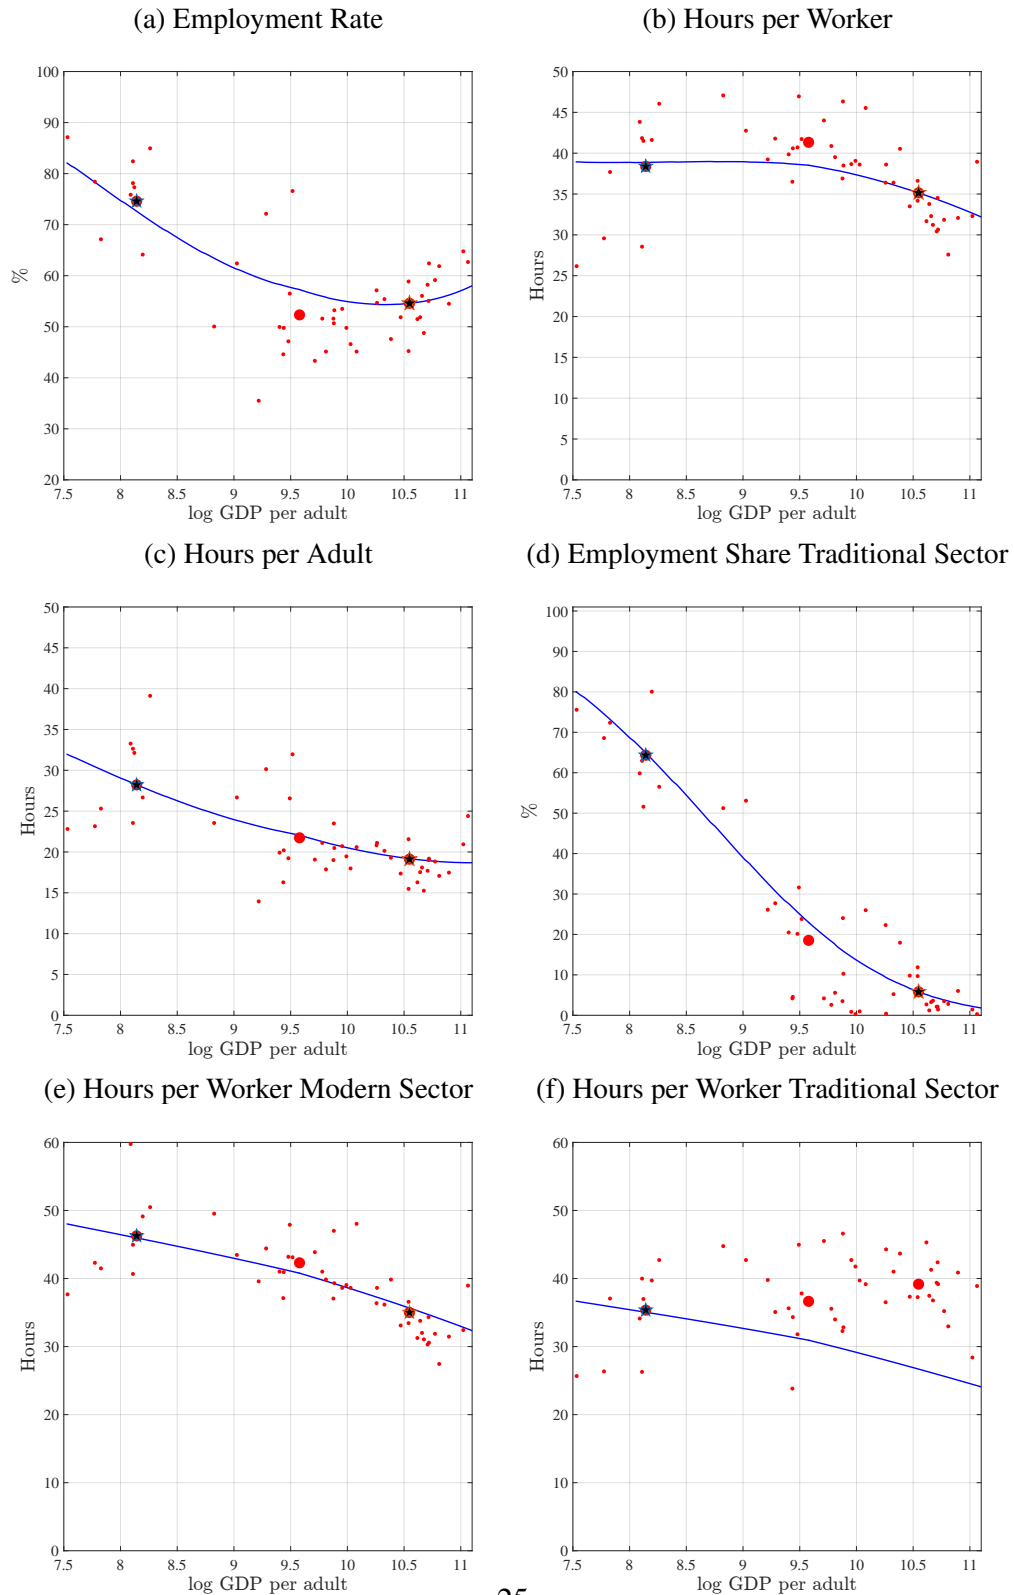
enue to GDP, share of government consumption) change piecewise linearly in $\log(A_M)$. For each level of A_M , we then solve for optimal hours.

Figure 3 compares the model predictions against the data. The large dots denote the averages by country groups in the data, and the stars mark the subset of targeted moments. Figures 3a and 3b 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, the model generates both the convex decrease in the employment rates, and the concave decrease in hours per worker over the development spectrum. As a result, hours per adult decrease at a similar rate as in the data, see Figure 3c.

The success behind generating the different shapes of the two margins of labor supply stems from structural change. Figure 3d shows the strong decrease of the employment share in the traditional sector, which is replicated well over the full range of development. Due to no fixed cost of working 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 \bar{u}_M . In the model, each family chooses a sector, and we can thus calculate sector-specific employment rates. These are however not defined in the data, in which we cannot assign non-working individuals to a sector. 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 Figure 3a. In both middle- and high-income countries, by contrast, the large majority of families works in the modern sector, and the decrease in the fixed cost of working in this sector generates the mild increase in the employment rate. Without changing fixed cost of working, employment rates would fall over the full development spectrum, see Appendix Figure A.1.


Besides the sectoral share, sectoral hours are important for aggregate hours per worker. Hours per worker in the modern sector are higher than in the traditional sector in the low- and middle-income countries, and decrease at a slightly faster pace than in the traditional sector, see Figures 3e and 3f, respectively. With decreasing returns in the traditional sector, workers do not want to reduce their hours too much because of the

Figure 3: Key Facts: Model vs. Data



Note: The small red dots represent each country in our sample, and the large red dots the averages by country income group. The model predictions are displayed by the blue line. Explicitly and implicitly targeted moments are marked with a star.

Figure 4: Hours-Wage Elasticity in Data and Model



Note: The small red dots represent each country in our sample, and the large red dots the averages by country income group. The model predictions are displayed by the blue line. The red dots are the beta-coefficients from the following regression, run separately for each country on the sample of workers in paid employment: $\log(h_i) = \alpha + \beta \log(w_i) + \delta_1 \text{age}_i + \delta_2 \text{age}_i^2 + \varepsilon_i$.

relatively high marginal product of an extra hour at low hours of work. Note that this pattern is still quantitatively at odds with the data, where hours per worker in the traditional sector even display a modest increase. 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.

Figure 4 compares the estimated hours-wage elasticity in each country from the data with the predictions from the modern sector in the extended model. The modern sector is the model equivalent of paid employment measured in the data. Similar to the simple model, the extended model endogenously generates the turning of the hours-wage elasticity from negative in poor countries to positive in the richest countries. Quantitatively, the model is still somewhat off, but provides an even better match than the simple model.

5. Decomposing Aggregate Hours Worked in the Extended Model

In contrast to the simple model, we now have three fundamental driving forces for the patterns in the data, namely income effects (we let both A_M and A_T vary), changes

Table 6: Counterfactual Experiments

	Poor	Rich	Diff.	% Expl.
Data	28.2	19.1	9.1	
Model	28.2	19.3	8.9	
<i>Income Effect</i>				
No Sect. Realloc.	28.4	23.0	5.4	61%
Sect. Realloc.	28.4	17.2	11.2	126%
<i>Taxes & Transfers</i>				
No Sect. Realloc.	28.4	26.4	2.0	22%
Sect. Realloc.	28.4	28.3	0.1	1%
\bar{u}_M				
No Sect. Realloc.	28.4	32.1	-3.7	-42%
Sect. Realloc.	28.4	31.1	-2.7	-30%

Note: By construction, the predictions for the average poor country are the same for all specifications. The fraction explained (% Expl.) corresponds to dividing Diff. for a given specification with the Diff. for the Model.

in taxes and transfers, and changes in the fixed cost of working in the modern sector.

We proceed as in Section 2: starting with the model inputs for the average low-income country, we turn on different driving forces one by one (without recalibrating the model). In each of these exercises, we then compute the predicted change in hours from the average poor to the average rich country from this exercise only, as a percentage of the total change in the model. Table 6 shows the results for these exercises and is constructed in the same way as Table 2. In the first two rows, the table states hours per worker in the average poor and the average rich country, as well as the difference between both, in data and model. The following rows then provide the results from the decomposition exercise. For each of the three driving forces, we first show results without sectoral reallocation: in this exercise, we force families to remain working in the same sector they optimally choose in the average low-income country. In a second step, we allow families to optimally choose the sector of work. That way, we can directly analyze which role sectoral reallocation plays for each of the driving forces of hours.

Without allowing for sectoral reallocation, income effects explain 61 percent of the decrease in hours between poor and rich countries, and taxes and transfers 22 percent. These results are similar to the decomposition results of the simple model, which at-

tributed 77 percent of the decrease to income effects, and 33 percent to taxes and transfers. They are slightly smaller because in the extended model the traditional sector employment rate does not adjust and always stays at 1. Without sectoral reallocation, 64 percent of the families work in the traditional sector also in the rich country predictions, and thus the absence of any adjustment in this margin matters quantitatively.

Allowing for sectoral reallocation however changes these percentages significantly: with sectoral reallocation, the model attributes more than the full decrease, namely 126 percent, to income effects, and essentially nothing, namely 1 percent, to taxes and transfers. To understand the role of sectoral reallocation, it is key to remember that a movement from the traditional to the modern sector implies a decrease in employment rates and an increase in hours per worker, and that employment rates are quantitatively more important for hours per adult. The faster increase in labor productivity in the modern than in the traditional sector by development induces such a reallocation of families from the traditional to the modern sector in the income effect experiment. This implies a strong decrease in the employment rate and thus hours per adult. This effect is even larger than the income effect on hours without sectoral reallocation, and the combination of both imply that the model predicts a decrease of 11.2 hours between poor and rich countries solely caused by increasing labor productivities in both sectors, and the associated income effects and sectoral reallocation. We can estimate the implied Marshallian elasticity of the extended model by regressing the logarithm of average predicted hours on the country level from the income effect experiment on the logarithm of average after-tax wages. The implied Marshallian elasticity is with -0.18 somewhat more negative than in the simple model, since sectoral reallocation exacerbates the effect of rising aggregate labor productivities on hours.

For taxes and transfers, sectoral reallocation works in the opposite direction: since labor income is not taxed in the traditional sector, and level and progressivity of labor income taxes both increase with development, the increase in taxes and transfers induces an increase in the traditional sector share with development.¹⁵ This implies an increase in the employment rate and consequently in hours per adult. This increase almost exactly counteracts the negative effect of taxes and transfers on hours without sectoral reallocation. Thus, in net the increase in taxes and transfers predicts a decrease of only

¹⁵In fact, we find the same, though somewhat muted, increase in the traditional sector share with development if we assume full taxation of labor income also in the traditional sector: the increase in progressivity then still induces more families to work in the traditional sector.

0.1 hours per week between poor and rich countries, and thus explains only 1 percent of the total decrease.

The new factor of decreasing fixed cost of working in the modern sector predicts an increase in hours by development, namely by 3.7 hours per week without sectoral reallocation, and 2.7 hours per week with sectoral reallocation. It thus negatively contributes to the total decrease in hours. Quite obviously, the decrease in the fixed cost of working in the modern sector induces an increase in the modern sector employment rate and thus total hours. Allowing for sectoral reallocation shifts families into the modern sector, somewhat counterbalancing this effect. Thus, while the decreasing fixed costs of working are crucial to generate the differential shapes of the two margins of labor supply, they do not positively contribute to explaining the decrease in hours worked by development.

In Appendix Tables A.8 to A.10, we show decomposition results from three further model variants with different assumptions about taxation. First, we assume that not only labor income, but also consumption remains untaxed in the traditional sector (Bachas et al., 2019). Second, we make the opposite assumption that both consumption and labor income are taxed in the traditional sector the same way as in the modern sector. Third, we keep the taxation assumptions as in the baseline model (no taxation of labor income, but taxation of consumption in the traditional sector), but assume full redistribution of government revenues to households, i.e. we set $G = 0$. For each specification, we recalibrate the model. The importance of the income effect for explaining the decrease in hours worked is almost completely unaffected by these different model assumptions: income effects always explain 121 percent to 126 percent of the decrease in hours by development. The implied Marshallian elasticities vary between -0.18 and -0.19. Not surprisingly, the quantitative importance of taxes and transfers varies more between these different set-ups. Assuming that the traditional sector is taxed the same way as the modern sector leads to the largest importance of taxes and transfers, both by maximizing their direct effect and by minimizing the counterweighting effect of sectoral reallocation. However, even in that case, taxes and transfers explain only 28 percent of the total decrease in hours by development. The decreasing fixed costs of working always predict an increase in hours by development, negatively contributing -24 to -35 percent to the decrease.

Thus, the main result of the decomposition exercise – income effects are the key

driver for the total decrease in hours worked by development – is very robust to different assumptions of taxation in the traditional sector and redistribution of government revenues. Not surprisingly, the importance of taxes and transfers for the decrease varies somewhat between these specifications, but never exceeds 28 percent. However, as in the simple model, taxes and transfers remain the sole driver of the turning of the within-country hours-wage gradient from negative to positive, which the extended model matches very well.

6. Conclusion

This paper asks why average hours worked are lower in rich countries than in poor countries. We consider two natural candidates: the more distortionary tax-and-transfer systems in rich countries, which are much smaller in poor countries, and income effects in preferences, which lead households to supply less labor when their income rises. We draw on detailed data on labor and non-labor taxation from a large set of countries of all development levels. We then use these data to discipline a simple model of labor supply, and we calibrate the model’s income effects to match the average differences in hours per adult across countries. The calibrated model predicts that income effects are the dominant force in lowering hours across the income spectrum. The reason is that cross-country differences in labor tax rates are modest in comparison with cross-country differences in labor productivity, and in turn average wage levels.

The simple model, while being straightforward to interpret, sweeps aside several salient disaggregate features of the cross-country data on labor supply. In particular, it does not address the facts that (i) employment rates decrease in a convex way over the development spectrum, but hours per worker in a concave way, and (ii) workers in poorer countries work primarily in self-employment, while those in richer countries work mostly in market wage work. To match these disaggregate patterns, we extend the simple model to include a household labor supply decision with an extensive and intensive margin of work, plus a subsistence self-employment sector, which has decreasing marginal product of labor. We show that this model does a good job at matching (i) and (ii) and also matches non-targeted moments, in particular the within-country hours-wage gradient.

We find that the extended model also predicts income effects to be the dominant force behind the overall decline in hours per adult across countries, even more so than

the simple model. The reason for this is that sectoral reallocation exacerbates the income effects, but attenuates the effects of increasing tax-and-transfer systems by development. We conclude that while tax-and-transfer systems may be the main feature explaining differences in hours worked across rich countries, and while they are the sole driver of the turning of the within-country hours-wage gradient from negative to positive over the development spectrum, they play at best a secondary role in explaining hours differences across the full development spectrum. Similarly, while structural transformation is crucial in explaining the extensive and intensive margins of labor supply across countries, matching these disaggregate features do not overturn the conclusion that income effects are the dominant force in explaining cross-country hours differences. In both the simple and the extended model, and independent of the exact specification of taxation and transfers, we always find Marshallian elasticities in the range of -0.10 and -0.19.

References

- Aguiar, M. and E. Hurst (2007). Measuring Trends in Leisure: The Allocation of Time Over Five Decades. *Quarterly Journal of Economics* 122(3), 969–1006.
- Bachas, P., L. Gadenne, and A. Jensen (2019). Consumption Taxes with Informality: Theory and Cross-Country Evidence from Expenditure Diaries. mimeo.
- 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.
- Benabou, 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., B. Brüggemann, N. Fuchs-Schündeln, and H. Paule-Paludkiewicz (2019, May). Long-term Changes in Married Couples’ Labor Supply and Taxes: Evidence from the U.S. and Europe Since the 1980s. *Journal of International Economics* 118, 44–62.
- Bick, A. and N. Fuchs-Schündeln (2017, May). Quantifying the Disincentive Effects of Joint Taxation on Married Women’s Labor Supply. *American Economic Review Papers & Proceedings* 107(5), 100–104.
- Bick, A. and N. Fuchs-Schündeln (2018, July). 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 (2018). Labor Supply in the Past, Present, and Future: A Balanced-Growth Perspective. *Journal of Political Economy* (forthcoming).
- Bridgman, B., G. Duernecker, and B. Herrendorf (2018). Structural Transformation, Marketization and Household Production around the World. *Journal of Development Economics* 133, 102–126.
- Chetty, R., D. Manoli, and A. Weber (2012). Are Micro and Macro Labor Supply Elasticities Consistent? A Meta-Analysis of Extensive Margin Elasticities. *NBER Macroeconomics Annual*, 1–56.
- Constantinides, G. M. (1982). Intertemporal Asset Pricing with Heterogeneous Consumers and without Demand Aggregation. *Journal of Business* 55, 253–267.
- Costa, D. L. (2000). The Wage and the Length of the Work Day: From the 1890s to 1991. *Journal of Labor Economics* 18(1), 156–181.
- Domeij, D. and M. Flodén (2006). The Labor-Supply Elasticity and Borrowing Constraints: Why Estimates are Biased. *Review of Economic Dynamics* 9, 242–262.
- Egger, P. H., S. Nigai, and N. M. Stecker (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.
- 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.
- Guner, N., G. Ventura, and Y. Xu (2008). Macroeconomic implications of size-dependent policies. *Review of Economic Dynamics* 11, 721–744.
- 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.
- Herrendorf, B., R. Rogerson, and A. Valentinyi (2014). *Handbook of Economic Growth*, Chapter Growth and Structural Transformation, pp. 855–941. Elsevier Science B.V.

- Jensen, A. (2019). Employment structure and the rise of the modern tax system. Unpublished Working Paper, Harvard University.
- 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*.
- King, R. G., C. I. Plosser, and S. T. Rebelo (1988). Production, Growth and Business Cycles I. The Basic Neoclassical Model. *Journal of Monetary Economics* 21(2-3), 195–232.
- Lagakos, D. and M. E. Waugh (2013, April). Selection, Agriculture, and Cross-Country Productivity Differences. *American Economic Review* 103(2), 948–980.
- MaCurdy, T. E. (1981). An Empirical Model of Labor Supply in a Life-Cycle Setting. *Journal of Political Economy* 89(6), 1059–1085.
- 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(8), 1353–1362.
- Ohanian, L. E. (2008). Back to the Future with Keynes. *Federal Reserve Bank of Minneapolis Quarterly Review* 32(1), 10–16.
- 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. (2006). Understanding Differences in Hours Worked. *Review of Economic Dynamics* 9, 365–409.
- 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.
- 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: Employment Rates by Country Income Group

	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	70.5	78.9
Old (55+)	69.8	30.5	24.0

Note: This table reports the percent of adults employed by country income group in total, by sex and by age group. Source: [Bick et al. \(2018\)](#).

Table A.2: Average Hours per Worker by Country Income Group

	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)	40.5	42.3	35.9
Old (55+)	32.6	37.5	33.6

Note: This table reports the average hours worked per employed adult by income group in total, by sex and by age group. Source: [Bick et al. \(2018\)](#).

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
Market 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
Market 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
Market 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	37.3	38.7	41.9
Market Sec. Hours	47.1	42.9	35.6
Traditional Sec. Share	62.2	15.9	4.4

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

	Country Income Group		
	Low	Middle	High
Traditional Sec. Hours	31.2	34.1	36.2
Market Sec. Hours	43.0	40.2	33.3
Traditional Sec. Share	83.0	41.4	16.3

Table A.8: Robustness: No Taxation of Traditional Sector

	Poor	Rich	Diff.	% Expl.
Data	28.2	19.1	9.1	
Model	28.2	19.1	9.1	
<i>Income Effect</i>				
No Sect. Realloc.	28.2	23.3	4.9	54%
Sect. Realloc.	28.2	16.9	11.3	124%
<i>Taxes & Transfers</i>				
No Sect. Realloc.	28.2	26.4	1.8	20%
Sect. Realloc.	28.1	29.0	-0.9	-10%
\bar{u}_M				
No Sect. Realloc.	28.2	31.5	-3.3	-36%
Sect. Realloc.	28.2	30.4	-2.2	-24%

Note: This table shows decomposition results for the model setup with no taxation of neither labor income nor consumption in the traditional sector. By construction, the predictions for the average poor country are the same for all specifications. The fraction explained (% Expl.) corresponds to dividing Diff. for a given specification with the Diff. for the Model.

Table A.9: Robustness: Full Taxation of Traditional Sector

	Poor	Rich	Diff.	% Expl.
Data	28.2	19.1	9.1	
Model	28.2	19.3	8.9	
<i>Income Effect</i>				
No Sect. Realloc.	28.4	23.1	5.3	60%
Sect. Realloc.	27.8	17.0	10.8	121%
<i>Taxes & Transfers</i>				
No Sect. Realloc.	28.4	25.7	2.7	30%
Sect. Realloc.	27.8	25.3	2.5	28%
\bar{u}_M				
No Sect. Realloc.	28.4	32.0	-3.6	-40%
Sect. Realloc.	27.8	30.8	-3.0	-34%

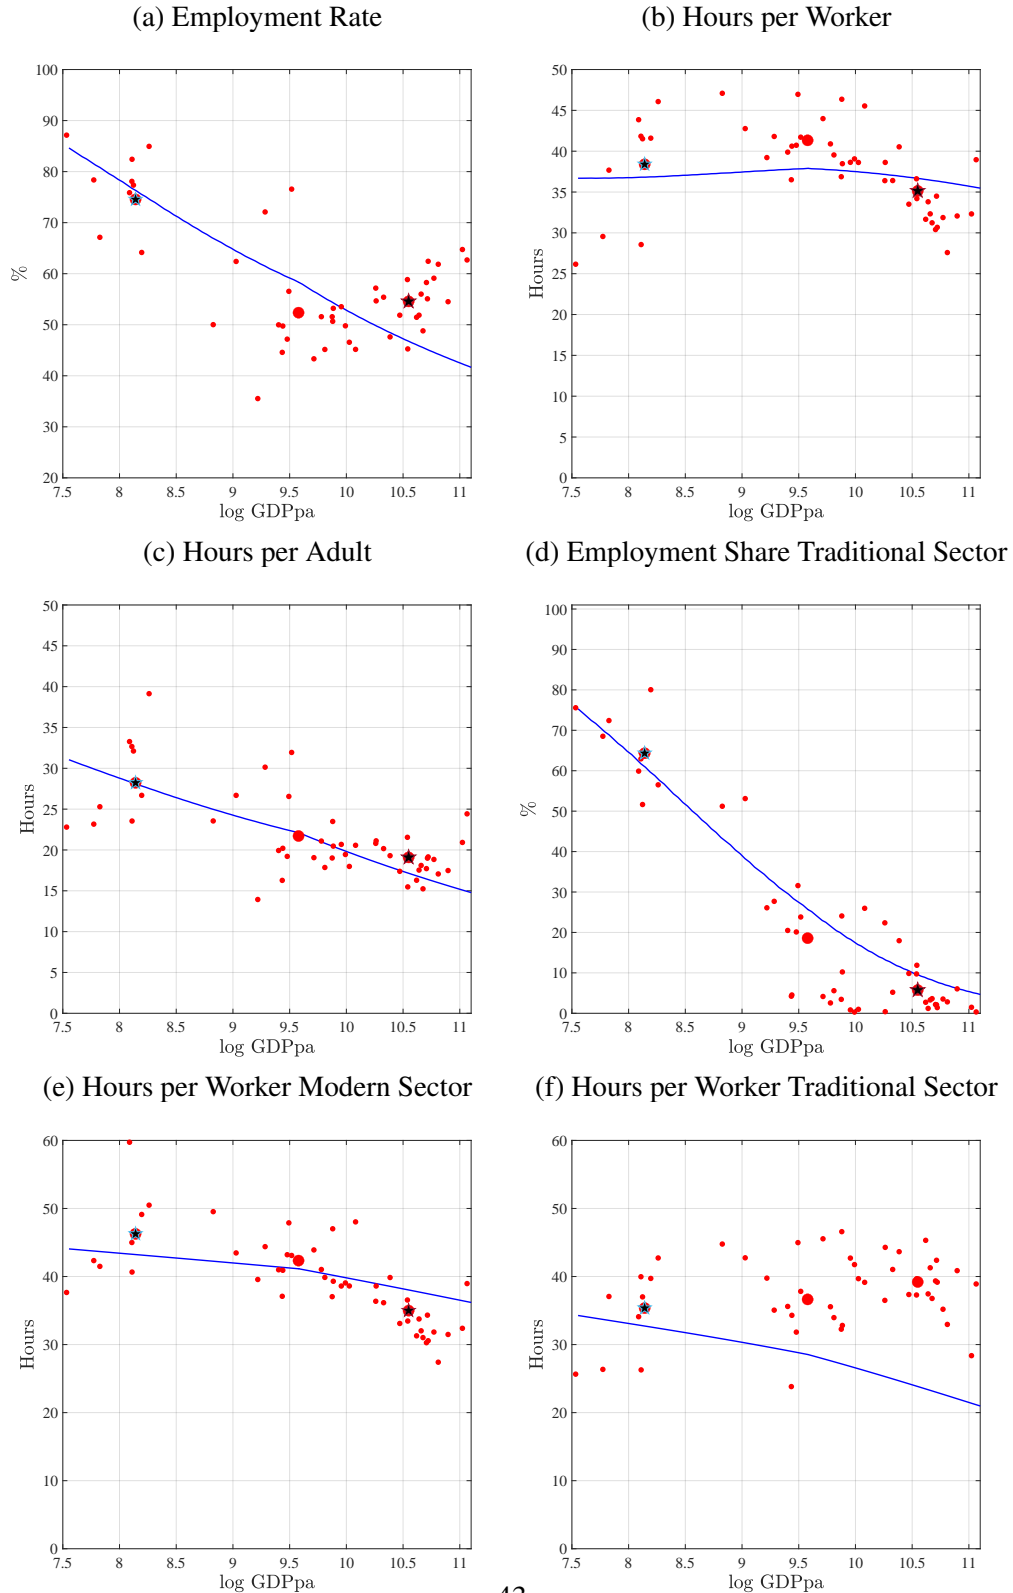
Note: This table shows decomposition results for the model setup with taxation of both labor income and consumption in the traditional sector. By construction, the predictions for the average poor country are the same for all specifications. The fraction explained (% Expl.) corresponds to dividing Diff. for a given specification with the Diff. for the Model.

Table A.10: Robustness: Full Redistribution

	Poor	Rich	Diff.	% Expl.
Data	28.2	19.1	9.1	
Model	28.2	19.1	9.1	
<i>Income Effect</i>				
No Sect. Realloc.	28.1	22.7	5.4	59%
Sect. Realloc.	28.1	16.8	11.3	124%
<i>Taxes & Transfers</i>				
No Sect. Realloc.	28.2	26.1	2.1	23%
Sect. Realloc.	28.1	28.1	0.0	0%
\bar{u}_M				
No Sect. Realloc.	28.2	32.2	-4.0	-44%
Sect. Realloc.	28.1	31.3	-3.2	-35%

Note: This table shows decomposition results for the model setup with only taxation of labor income, but not of consumption in the traditional sector. In contrast to the baseline model, it assumes full redistribution, i.e. $G = 0$. By construction, the predictions for the average poor country are the same for all specifications. The fraction explained (% Expl.) corresponds to dividing Diff. for a given specification with the Diff. for the Model.

Figure A.1: Model Fit: \bar{u}_M Not Changing with Development



Note: These are the model results from a version of the model in which the fixed cost of working in the modern sector, \bar{u}_M , are not allowed to change by development. We thus need to calibrate one parameter less, and the model is overidentified. The small red dots represent each country in our sample, and the large red dots the averages by country income group. The model predictions are displayed by the blue line. Explicitly and implicitly targeted moments are marked with a star.

B. Model Appendix

B.1. Counterfactual Experiments

We use the model to decompose the importance of income effects and tax-and-transfer systems in explaining the decrease in aggregate hours via two counterfactual exercises. When doing so, we take into account that progressive tax systems are in some way defined relative to the income level in a country. To best understand the issue, we restate the expression for after-tax labor income in our simple model:

$$A_c z h - \mathcal{T}_c(A_c z h) = A_c z h \left(1 - \frac{\mathcal{T}_c(A_c z h)}{A_c z h} \right) = A_c z h (1 - \delta_c(A_c z h)) \quad (\text{B.1})$$

where $\delta_c(A_c z h)$ is the average tax rate the household pays at any hours choice in country c . Using our specific functional form assumption for progressive taxation (4) yields

$$\delta_c(A_c z h) = \frac{\mathcal{T}_c(A_c z h)}{A_c z h} = \frac{A_c z h - \lambda_c(A_c z h)^{1-\tau_c}}{A_c z h} = 1 - \lambda_c(A_c z h)^{-\tau_c}. \quad (\text{B.2})$$

In our first experiment, we keep the tax-and-transfer system of the poorest countries fixed, while raising TFP to the level of the richest countries. Without any further adjustment for any hours choice h , a household with productivity level z in this counterfactual world would pay an average tax rate $1 - \lambda_{poor}(A_{rich} z h)^{-\tau_{poor}}$. Since in our calibration $A_{rich}/A_{poor} = 16.7$, this would imply that the household would face a much higher marginal tax rate than the household with the same individual productivity level and same hours given the poor country's TFP. This logic counteracts the idea of holding the tax system fixed. We address this by calculating net labor income as follows:

$$A_{rich} z h (1 - \delta_{poor}(A_{poor} z h)). \quad (\text{B.3})$$

In our second experiment, we take the income level of the poorest countries as fixed, but change the size of the tax-and-transfer system with actual GDP per adult. Without any further adjustment for any hours choice h , a household with productivity level z in this counterfactual world would pay an average tax rate $1 - \lambda_{rich}(A_{poor} z h)^{-\tau_{rich}}$, would thus face a much lower marginal tax rate than the household with the same individual productivity level and same hours given the rich country's TFP. This logic counteracts the idea of varying the tax system. We address this by calculating net labor income as

follows:

$$A_{poor}zh(1 - \delta_{rich}(A_{rich}zh)). \quad (\text{B.4})$$

We use the same approach for the counterfactual exercises in the extended model.

B.2. Second-Stage Solution of Family Problem

We solve the second stage family problem (6) as follows. Plugging the optimal consumption c and hours $\tilde{h}^*(\eta^*)$ into the objective function, the family head's problem becomes an unconstrained problem:

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

Taking the first order condition and applying the chain rule and the Leibniz rule leads to

$$\alpha \frac{\tilde{h}^*(\eta^*)^{1+\frac{1}{\phi}}}{1+\frac{1}{\phi}} f(\eta^*) + \bar{u}_S \eta^* f(\eta^*) = -\alpha \tilde{h}^*(\eta^*)^{\frac{1}{\phi}} F(\eta^*) \frac{d\tilde{h}^*(\eta^*)}{d\eta^*},$$

where f is the PDF of η . The first term on the LHS of this equation equals the marginal disutility from working \tilde{h}^* hours for the new workers of mass $f(\eta^*)$ that start working if the optimal threshold level η^* is marginally changed. The second term of the LHS adds to this the fixed utility cost incurred by these workers. These marginal utility losses of the new workers are equated with the marginal utility gain the already existing workers of mass $F(\eta^*)$ enjoy because of their decrease in hours worked, which is expressed on the RHS. The equation thus implicitly defines the optimal threshold level as a function of family hours, $\eta^* = \eta(h)$.

Since $\tilde{h}^*(\eta^*) = \frac{h}{F(\eta^*)}$, 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

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

Case of Uniform Distribution To make further progress, we assume $\eta \sim U(0, 1)$ and thus $F(\eta) = \eta$. In this case, we can solve for the optimal cutoff $\eta(h)$ in closed form:

$$\eta(h) = \left(\frac{1}{\bar{u}_S} \frac{\alpha}{1+\phi} h^{1+\frac{1}{\phi}} \right)^{\frac{\phi}{1+2\phi}}.$$

Notice that η^* must be bounded by one from above, so the maximum h for an interior solution is

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

We thus have two cases. First, if h is larger than this threshold, then $\eta^* = 1$ and the family utility is simply given by

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

Second, if h is smaller than the threshold, then $\eta^* = \eta(h)$ and equation (7) becomes

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