Why are Average Hours Worked Lower in Richer Countries?

Alexander Bick, ASU and Federal Reserve Bank of Minneapolis
Nicola Fuchs-Schündeln, Goethe University Frankfurt and CEPR
David Lagakos, UCSD and NBER
Hitoshi Tsujiyama, Goethe University Frankfurt

November 30, 2019

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 prevalence of subsistence self-employment.

Email: alexander.bick@asu.edu, fuchs@wiwi.uni-frankfurt.de, lagakos@ucsd.edu, and hitoshi.tsujiyama@hof.uni-frankfurt.de. The views expressed in this paper do not represent the view of the Federal Reserve Bank of Minneapolis or the Federal Reserve System. We thank Nora Strecker and Peter Egger for providing us with percentile data on gross and net incomes for different countries. For helpful comments we thank Per Krusell, Johanna Wallenius and seminar audiences at BI Norwegian Business School, Bonn, the IFS, the IIES Stockholm, Melbourne, Monash, MOVE Barcelona, Queensland, the Reserve Bank of Australia, Sciences Po, Vienna, and Zurich as well as conference participants at the 2019 T2M conference and the 2019 SED meetings. Fuchs-Schündeln gratefully acknowledges financial support from the European Research Council under the ERC Consolidator Grant No. 815378. All potential errors are our own.

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 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 model of labor supply by heterogeneous households with preferences in which income and substitution effects do not necessarily offset each other. For maximum comparability with other studies we assume a standard form of preferences in which consumption and hours are additively separable, utility in consumption takes a CRRA form, and the disutility of hours worked is increasing and convex (MaCurdy, 1981; Keane, 2011). Households face marginal taxes 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 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 (Ohanian et al., 2008) and across countries (e.g. Jensen, 2019). Our approach then is to calibrate our model using cross-country evidence on statutory labor tax rates assembled by Egger et al. (2019) drawing on data from the IMF, the World Bank and other official government and private company sources. We assume flexible functional forms for the tax rates on labor and consumption that vary only with labor productivity. 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. 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 – around 71 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. Thus, the overall effect of a 10 percent increase in the wage rate is a decrease in labor supply of around 1 percent. The model's calibrated CRRA parameter on consumption is 1.23, which is higher than the standard value of 1 assumed in balanced growth preferences (King et al., 1988). 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 poor 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 and intensive margin and a "traditional sector" based on subsistence self-employment. The extended model has several new features that are included in order to match the patterns of the

[1]The 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).

two margins and self-employment across countries described above. There are heterogeneous families that make a labor supply decision for each of their members, who differ in costs of supplying labor, facing a fixed cost of supplying any labor. The "traditional sector" allows households to produce output using a decreasing-returns production function. Decreasing returns capture the reality that on a farm or in a family business, hours of work will become less and less productive as more hours are supplied. The "modern sector" allows household members to work in competitive labor markets for as many hours as they choose at the going wage. Subsistence self employment is rarely taxed in practice (Jensen, 2019), so we assume that taxes are only levied on the modern sector. We also allow for fixed cost of working à la Rogerson and Wallenius (2013), but that differ in the two sectors, capturing the fact that subsistence self-employment work is easily available, while accessing the modern sector is associated with higher costs. This is especially true in poor countries, where the modern sector is likely to be concentrated in cities but the majority of people still lives in rural areas. As before, we calibrate the extended model to match the cross-country data on average hours per adult. We calibrate fixed costs to match the patterns of employment rates in the richest and poorest countries, and the relative productivity of the modern sector using the shares of workers in subsistence self employment in the richest and poorest countries.

The extended model does well overall in matching both the convex shape of the extensive margin and the concave shape of the intensive margin, which we do not target directly. The intuition is as follows. First, the decreasing returns to hours in the traditional sector lead to relatively low hours per worker in this sector over the entire world income per capita range. As countries grow richer and the relative productivity of the market sector increases, individuals move rapidly out of the traditional sector into the market sector, which leads to initially stable *aggregate* hours per worker despite decreasing *sectoral* hours per worker in both sectors. Decreasing fixed cost of working in the market dampen the decrease in the employment rate with development, and thus generate the convex shape of the employment rate with development. Reassuringly, the model also matches the individual-level gradient of hours worked on income, which is not targeted at all. In both the model and data, these gradients are negative in poor countries and become flatter with GDP per capita (Bick et al., 2018; Costa, 2000; Aguiar and Hurst, 2007). Both of these successes in matching the data lend credence to the model's disaggregate predictions.

The extended model predicts that income effects explain the majority – around 80 percent – of the relationship between hours worked and income levels across countries. The reason income effects dominate here is the same as in the simple model: cross-country variation in tax rates are dwarfed by the variation in labor productivity levels. The extensive and intensive margins and self-employment patterns are shaped by modestly lower fixed costs of market work in richer countries and higher relative productivity of the modern sector. Yet these rich disaggregate features do not fundamentally change the cross-country drivers of aggregate labor supply. The CRRA parameter on consumption in now calibrated to be 1.22, which is somewhat lower but still quite similar with the estimate from the benchmark model. 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. Simple Model of Labor Supply

In this section we present a simple model of labor supply 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 level. 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.

2.1. Environment

Output is produced using a constant-returns production function using labor as the sole input: $Y = AL$, where A represents aggregate labor productivity and L is aggregate 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:

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

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 percent 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 $T(y(z))$ to household income $y(z) = zwh$, and a linear consumption tax with a tax rate τ_c on consumption $c(z)$. The government spends G on "government consumption expenditure," which does not enter the households' utility function directly. The remaining government tax collections are redistributed as lump-sum transfers Υ to households such that the budget is balanced. These transfers can be best interpreted as the amount of direct transfers to households and the amount of public goods provided by the government that are a substitute for private consumption expenditures.

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

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

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

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

Two sets of variables differ exogenously by development: (i) the aggregate productivity level, A; (ii) the size of the tax-and-transfer system (τ_c, T, G, Υ). 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 simple model we need to discipline the size of tax and transfer systems across countries. To do so we draw on 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 Benabou (2002) with net income $\tilde{y}$ being given by

$$\tilde{y} = y - T(y) = y - \left(y - \lambda y^{1-\tau}\right) = \lambda y^{1-\tau}, \tag{4}$$

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, they shared for each country 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, again from the IMF government statistics we obtain the share of social benefits over GDP Υ/Y, and calculate the share of government consumption expenditure, G/Y as a residual, namely as the share of government expenditure excluding social benefits over GDP.

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 piece-wise linear interpolation of the averages for all variables over low-, middle-, and high-income countries.[3] 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 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 estimates 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). This indicates that 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 2.3 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 of worked in richer countries.

[2] We take the sample of 62 countries with information on the share of government revenues coming from labor income taxes also for the progressivity parameter, i.e. the sample of countries is consistent for the different fiscal inputs.

[3] For the government expenditure inputs in Figure 1c, we do the interpolation on total government revenues over GDP and government consumption expenditure over GDP, and then impose that social benefits over GDP cannot be negative.

Figure 1: Cross-Country Differences in the Tax-and-Transfer System

(a) Fiscal Progressivity

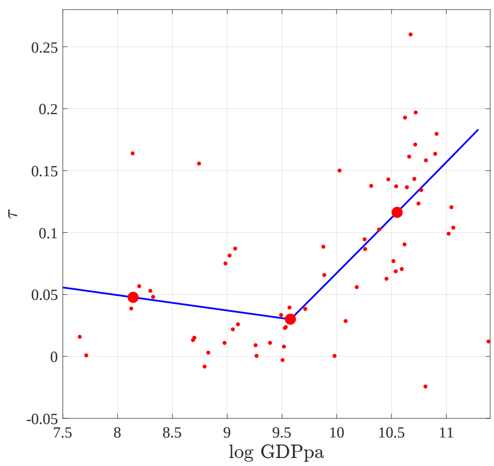

(b) Share of Government Revenues Coming from Labor Income Taxation

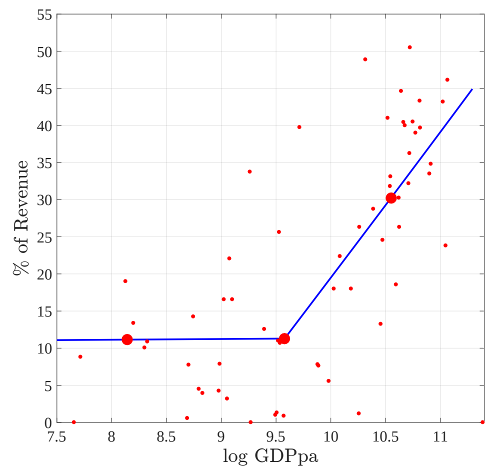

(c) Gov't Revenue Relative to GDP

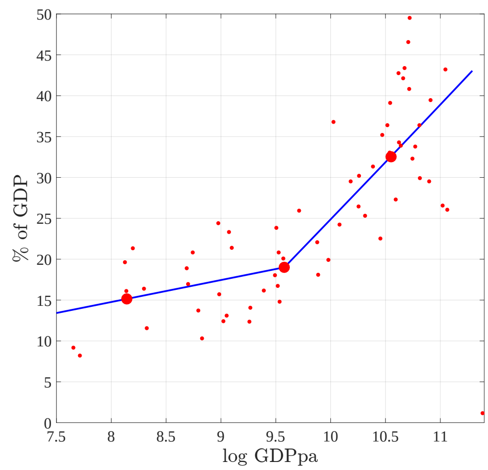

Note: The small red dots represent each country in our sample, and the large red dots the average by country-income group.

2.3. Calibration

With the tax data in hand, we calibrate the model so that it matches average hours worked in both the poorest tercile of the world income distribution as well as the richest. 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. Thus, by construction the model is calibrated to explain the entire decline in hours worked present in the data. To calibrate the model, 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 a Frisch elasticity in line with estimates in the literature (see Keane, 2011). We normalize $\alpha = 1$. We then jointly calibrate γ and the labor-productivity level in the poorest tercile of the income distribution, A^{poor}, to match average hours worked per adult and GDP per adult in the poorest tercile relative to the average rich country.[4] The parameter γ thus determines the size of income effects in preferences, and is identified from the hours differences between the richest and poorest countries not explained by differences in tax systems.

Figure 2: Key Facts: Simple Model vs. Data

(a) Average Hours per Adult

(b) Within-Country Hours-Wage Gradients

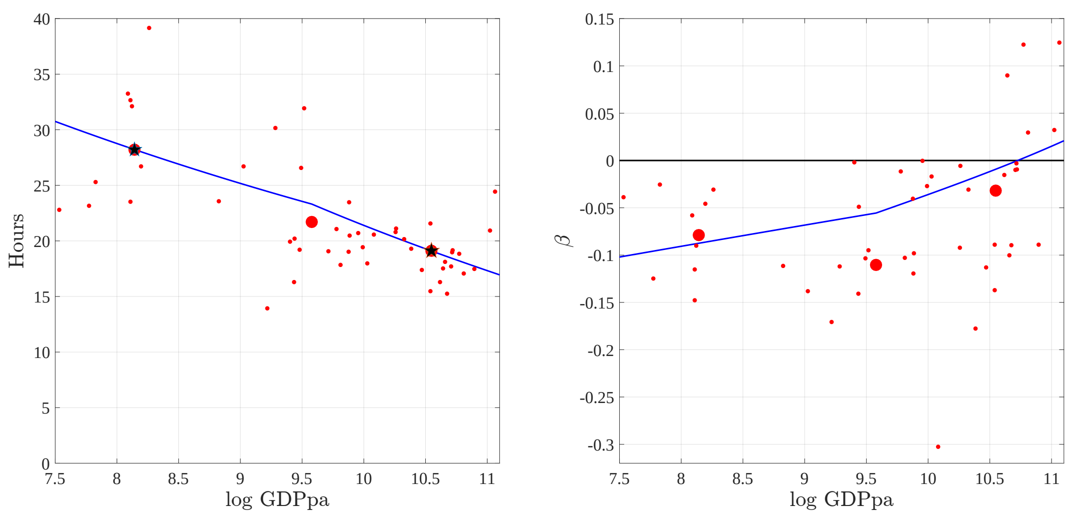

Note: The small red dots represent each country in our sample, and the large red dots the average by country-income group. The model predictions are displayed by the blue line. Explicitly and implicitly targeted moments are marked with a star.

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

[4] Average hours worked per adult in the poorest tercile are 29.3 hours, and in the richest tercile 19.1 hours, see Table 2)

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. Our estimate within the large range of the 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 we do. Hence, our estimate might be more useful in the context of large productivity differences.[5]

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 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, and even reverses in some advanced economies. This reversal has been documented in the U.S. history as well, by Costa (2000) and Aguiar and Hurst (2007).

Figure 2b plots the model's predictions for hours-wage gradients by income level (blue line) and the in the cross-country data (red dots). The model does quite well in matching the levels and cross-country variation in hours-wage gradients, even though it does not direct 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

[5] Another elasticity commonly reported in the literature is the Hicksian elasticity of labor supply, which summarizes how hours respond to a change in wages holding the level of income fixed. The Hicksian elasticity in our model is given by $1/(1/\phi+\gamma)$, and is 0.45 at our calibrated parameter values. This is somewhat higher than the average of 0.31 across 22 studies reported by Keane (2011), which have estimated Hicksian elasticities ranging from 0.02 to 1.32, and comparable to the average of 0.50 and 0.59 reported for micro studies and macro studies by Chetty et al. (2012).

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 as some validation of 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 from the rising and increasingly progressive tax systems that come with development (highlighted in Figures 1a, 1b, and 1c). The main reason is that the 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. Quantitative Counterfactuals

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 TFP to the level of the richest countries. This counterfactual simulates how hours worked would look across the world income distribution 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 capita. The purpose of the second counterfactual is to compute hypothetical hours worked by income level assuming that tax systems were the only source of variation across countries.

Table 1 shows the results of these counterfactuals. In the first two rows, the table states the full model predictions for the average poor and the average rich country, as well as the difference between both. 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 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. Income effects

Table 1: 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	21.7	6.5	71%
Taxes & Transfers	28.2	24.6	3.6	40%

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 in two counterfactual experiments. The first, 'Income Effects,' 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.

alone explain 71 percent of the difference between the poorest and richest tercile in the data. Taxes and transfers explain 40 percent by themselves. Not that because the income effect and taxation interact non-linearly 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 (=$A^{rich}/A^{poor} = 1/0.062$). The average labor tax rates varies from 12.5 percent in the poorest tercile of countries to 16.5 percent in the richest. The average consumption tax rises from 12.5 percent in the poorest tercile to 27 percent in the richest. As a simple summary metric of the size of the tax system, taxes relative to GDP are around 15 percent in the poorest tercile and 33 percent in the richest. Taking a simple ratio of these, the overall tax burden is about 2.2 times larger in the richest tercile than the poorest. At conventional elasticities of labor supply, which our model delivers, these tax differences are just large enough to match the substantially lower hours worked of the richest countries. Quantitatively, the enormous changes in income are therefore dominant in explaining the declines in hours across the world GDP per capita distribution.

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. As we review below, 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 them out of self-employment and into market wage work.

In this section we review these disaggregate 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 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).[6] Table 2 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, as measured by GDP per adult in the Penn World Tables.[7] 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 22.6 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 6.1 hours per week between middle- and high-income countries. Appendix Tables A.1 and A.2 shows that these patterns are similar among men and women separately, and across broad age groups. Hence, we abstract from age and sex in the model that follows.

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

[6] We omit Laos, since there are no data on education for Laos, which we need to pin down the sector of work in Section 3. Thus, all data on hours and sectors in this paper refer to 48 countries.

[7] In our data, we have 10 countries in the poorest tercile, 15 countries in the middle tercile, and 23 countries in the richest tercile.

Table 2: 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).

subsistence work, though certainly not exact, and something we can measure in a comparable way across the countries in our data (see Bick et al. (2018).) For comparison with our model later, we define these workers as the "traditional sector," and the balance to be the "modern sector."

Table 3 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 3.6 hours higher in rich than in poor countries in the traditional sector. By contrast, they are 11.5 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 11 and 5.3 weekly hours, respectively. Only for the rich countries are hours higher in the traditional sector, with a difference of 4.3 hours.

The last row of Table 3 shows the share of all workers working in the traditional sector: In the poor countries, almost two thirds of workers (64.6 percent) work in the traditional sector. This share rapidly decreases to 19.9 percent in the middle-income countries, and only 5.9 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 1.4 weekly hours worked per worker between low- and middle-income countries documented in Table 2 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-

Table 3: 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

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.[8] 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 7.2 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.

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 and show that it matches all the aggregate and disaggregate patterns described above, including a number of non-targeted moments. 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

[8] The small increase in hours per worker in the traditional sector marginally adds to this increase.

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 $\log z \sim N(0, \sigma_z^2)$. Individuals within a family differ only in their individual fixed disutility of work η. The instantaneous utility function of an individual is

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

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 depend on the sector the family is working, 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 traditional sector 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.

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.

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

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

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 η^*: 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:

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

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:

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

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

Equilibrium Wage The competitive market 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 $w = A$.

Government Budget The government budget is balanced in equilibrium:

$$\Upsilon = (1-g) \sum_{S=T,M} \left[\left(\int T_S(y(z)) + \tau_{c,S} c(z) \right) 1_{\{S(z)=S\}} dF_z \right]. \tag{9}$$

Equilibrium A stationary equilibrium consists of a set of decision rules $\{c(\cdot), h(\cdot), \eta^*\}$, a wage rates w, and the government policies $\{\tau_{c,S}, S, T_S(\cdot), \Upsilon, g\}$ such that

(i) given the price and policies, the decision rules solve families' problem (6) and (8),

(ii) the marginal profit condition is satisfied: $w = A$, 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, respectively; (ii) the size of the tax-and-transfer system (τ, T, g), 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 labor productivity differences between poor and rich countries are a standard force of structural change, different fixed cost of working in the modern sector are not. Why do we introduce this force of structural change? First, fixed cost 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 show in the Appendix and explain in Section 5, 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 take some parameters of the model from the literature, and calibrate the remaining 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.

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 9). The distribution of modern sector productivity is estimated using the panel component of the Current Population Survey (CPS), and is assumed to capture permanent differences across families. We do so by estimating a panel fixed effect regression of log income per hour on individual fixed effects, following Lagakos and Waugh (2013). We take the variance of the individual effects to be the variance of the permanent productivity differences in our model. The implementation of the tax-and-transfer system follows the description in Section 2.2.

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.[9] Last, the weight on the disutility of working α is just a normalization to get the level of hours per worker right.

Given these normalizations, all other parameters are jointly calibrated to replicate key moments from the data. As for the aggegrate 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.

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 not a proof of identification. We have in total 8

[9] Since 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 4: Calibrated Parameters and Calibration Targets

	Parameter	Value	Target	Data	Model
$\bar{u}_M^{poor}$	mean disutility of working in M sector in poor countries	5.51	avg. ER in poor countries	74.5	72.6
$\bar{u}_M^{rich}$	mean disutility of working in M sector in rich countries	2.44	avg. ER in rich countries	54.5	54.5
A_T^{poor}	T sector productivity in poor countries	0.02	avg. % of workers in T sector in poor countries	64.3	64.8
A_T^{rich}	T sector productivity in rich countries	0.12	avg. % of workers in T sector in rich countries	5.9	6.0
ρ	DRS parameter in T sector	0.86	avg. hours per worker in T sector in poor countries	35.4	35.1
ϕ	cuvature of disutility of working	0.55	avg. hours per worker in M sector in poor countries	46.3	46.0
γ	cuvature of consumption	1.20	avg. hours per worker in rich countries	35.0	35.0
A_M^{poor}	M sector productivity in poor countries	0.07	avg. output per adult in rich vs. poor countries	11.4	11.4

free parameters and choose 8 moments to be replicated. The mean disutilities 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 function, 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 rich relative to poor countries. The choice of these moments ensures that both labor input and output in poor and rich countries are consistent with the data.[10]

Table 4 states the targeted moments, which we all match perfectly, and lists the calibrated parameter values. The mean fixed cost of working in the modern sector is 2.3 times (5.51/2.44) higher in poor than in rich countries. The constant returns to scale parameter is with 0.86 very 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 about 3.5 times (0.07/0.02) larger than in the traditional sector. In rich countries, this difference is more than twice as large (1/0.12). Comparing poor and rich countries, the productivity gap amounts to a factor of over 14 in the modern sector, and slightly below 6 in the traditional sector. The calibrated curvature parameter for the disutility of working is with 0.55 consistent with the estimates surveyed in Blundell and MaCurdy (1999), Domeij and Flodén (2006), and Keane (2011). Finally, our calibrated value of the curvature of consumption is 1.20, and thus implies that income effects dominates, but at the same time is not too far from the log specification.

4.4.3. Model Fit

The first question our paper asks is whether the three fundamental forces of income effects, tax and transfer systems, and structural change can generate the patterns of decreases in hours worked per adult, employment rates, and hours worked per worker with development observed in the data. While we calibrate the model to aggregate moments from the average poor and high income countries, the shapes of different aggregate 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 aggregate modern sector productivity A_M increases linearly in log from A_M^{poor} to A_M^{rich}, and extend the process to productivity levels below A_M^{poor} and above A_M^{rich}. We then interpolate/extrapolate the aggregate traditional sector productivity A_T and disutility of labor supply $\bar{u}_M$ along the process for A_M. Last, we also interpolate/extrapolate the fiscal inputs (i.e., tax progressivity, share of government revenue coming from labor income taxes, ratio of government revenue to GDP, share of govern-

[10]Note 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.

ment consumption). We do this in a piece-wise linear fashion, separately for low- to middle-income countries and for middle- to high-income countries.[11]

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.[12] 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 explaining the behavior of aggregate hours per worker stems from structural change. Figure 3d shows the decrease in of the employment share in the traditional sector, which is replicated well over the full range of development. Besides this 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. The intuition for the latter result is the production technology in the traditional sector. With decreasing returns in the traditional sector, workers do not want to reduce their hours too much because of the 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 fairly flat hours per worker in the aggregate between poor and middle-income countries, and the steeper fall off between the middle- and high-income countries.

Finally, we can analyze how well the model fits a completely untargeted moment, namely the within-country hours-wage elasticity. Aggregate hours worked are the result of individual hours worked decisions. How do hours worked vary with income on the

[11]The social benefits over GDP are calculated as a residual from the shares of government consumption and revenue. For countries in which the residual is negative, we assume there are no social benefits.

[12]We also mark average hours per adult and hours per worker in poor countries as targets, since they are calculated exclusively from targeted moments.

Figure 3: Key Facts: Model vs. Data

(a) Employment Rate

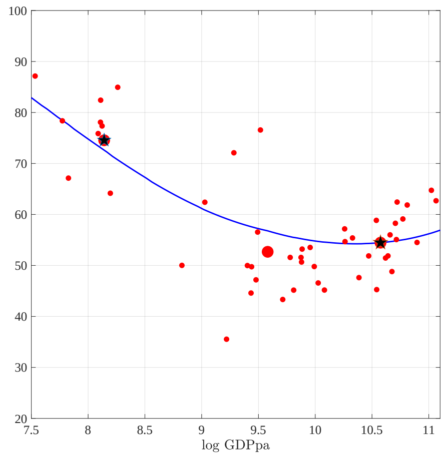

(b) Hours per Worker

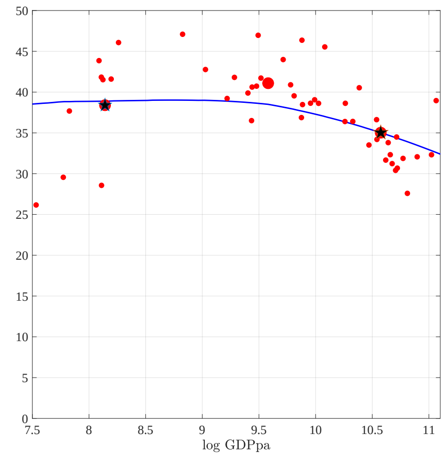

(c) Hours per Adult

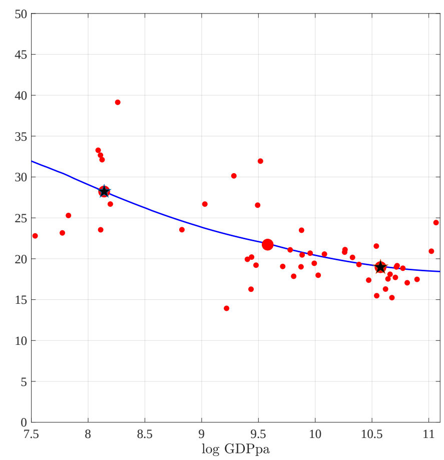

(d) Employment Share Traditional Sector

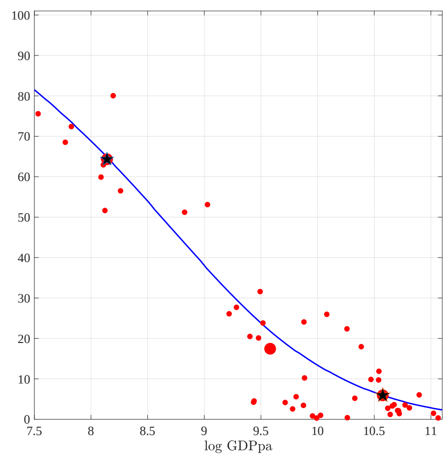

(e) Hours per Worker Modern Sector

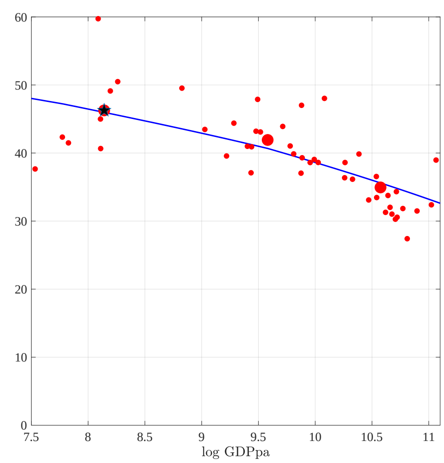

(f) Hours per Worker Traditional Sector

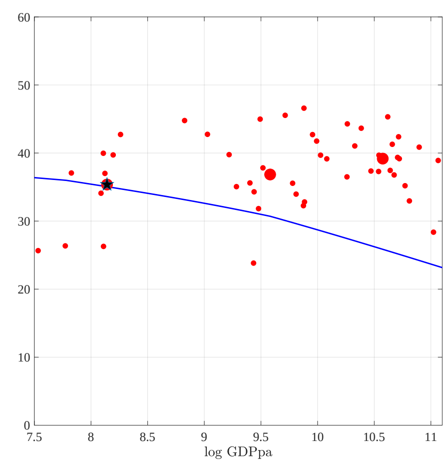

Note: The small red dots represent each country in our sample, and the large red dots the average by country-income group. The model predictions are displayed by the blue line. Explicitly and implicitly targeted moments are marked with a star.

individual level? The hours-wage elasticity on the individual level is estimated in Bick et al. (2018) relying on wages in paid employment. While workers in paid employment are clearly a selected group, their wages are relatively straightforward to measure by dividing monthly earnings by monthly hours. Moreover, paid employment closely resembles the modern sector definition, and we compare the hours-wage elasticities in paid employment to the predicted ones in the modern sector. In Bick et al. (2018), we run the following regression on the country level:

$$log\left(h_i\right) = \alpha + \beta log\left(w_i\right) + \delta_1 age_i + \delta_2 {age_i}^2 + \varepsilon_i. \quad (10)$$

The resulting β-coefficients are shown in Figure 4.[13] In the majority of countries, low-wage individuals work more hours than high-wage individuals. For the low-income countries, the average hours-wage elasticity is negative with -0.08, and it stays negative and becomes even slightly larger in absolute terms for the middle-income countries with -0.09. However, in the richest countries in the sample the relationship turns around and high-wage individuals work more hours than low-wage individuals. On average, the hours-wage elasticity is -0.04 for the high-income countries.

Figure 4 compares the estimated hours-wage elasticity in each country from the data with the model predictions. The model endogenously generates the turning of the hours-wage elasticity from negative in poor countries to positive in rich countries. Quantitatively, the model is somewhat off, starting the increase in the elasticity too early in the development process.

5. Decomposing Aggregate Hours Worked in the Extended Model

In this section, we provide a quantitative evaluation of the relative importance of the three driving forces for the patterns in the data, namely income effects, taxes and transfers, and structural change. At the same time, this section provides intuition for the model predictions presented in the previous subsection.

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

[13] In 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.

Figure 4: Hours-Wage Elasticity in Data and Model

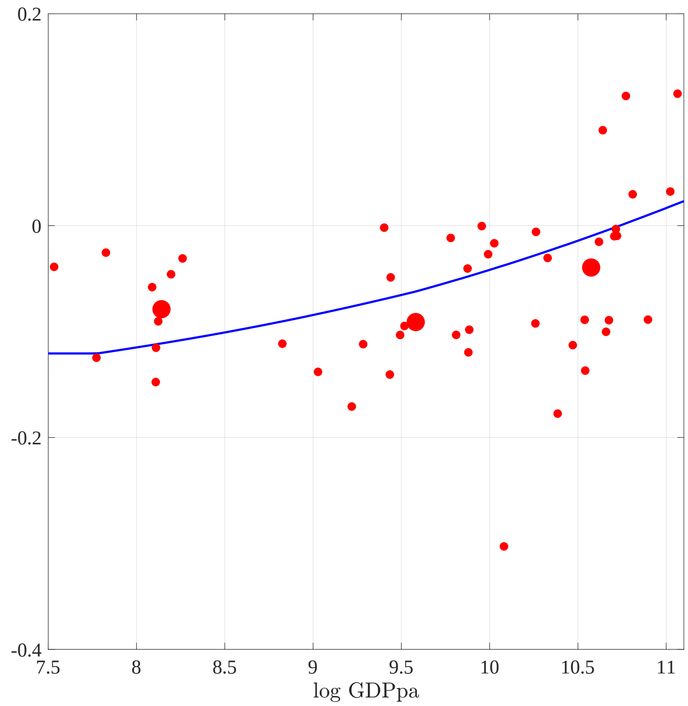

from the average poor to the average rich country from this exercise only, as a percentage of the total change in the model. Table 5 shows the results for these exercise and is constructed in the same way as Table 1. In the first two rows, the table states the full model predictions for the average poor and the average rich country, as well as the difference between both. The following rows then provide the results from the decomposition exercise. The first three columns focus on hours per adult, and state the predictions for the average 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. The next three column do the same for the employment rate, and the final three for hours per worker.

For taxes and transfers (row 4), the exercise is relatively straightforward: in this decomposition exercise, we let only the fiscal inputs vary as in the data. To completely abstract from structural change, we additionally impose that all families remain working in the sector they optimally choose in the average low-income country. To capture income effects, we conceptually would like to let only labor productivity increase and calculate predicted hours for the average rich country from this exercise. The complication arises that the model features labor productivity in two sectors, and labor productivity in the modern sector rises faster than in the traditional sector, which is a driving force of structural change. To disentangle income effects from structural change, we let in a first step labor productivity in the traditional sector increase proportionally with labor productivity in the modern sector, i.e. we keep the ratio $\frac{A_M}{A_T}$ constant across countries, and again additionally shut down sectoral choice by imposing that every family continues working in the same sector it optimally chooses in the average low-income country. We thereby shut down the effects of structural change and attribute the change in hours from this counterfactual exercise to the income effect. The results for this exercise are shown in row 3. In a second step, we let TFP in both sectors rise as in the benchmark and allow families to optimally choose the sector of work. We attribute the resulting change in hours from this exercise relative to the previous exercise to structural change. This is the part of the effect of structural change on hours that works via relative labor productivity changes in both sectors, and we call it "sectoral reallocation" (row 6). An additional effect of structural change on hours comes through the change in the fixed cost of working in the modern sector with development. To analyze the importance of this component of structural change, we again impose no sectoral reallocation of families and let the fixed cost of working vary as in the full model (we call this $\bar{u}_M$, see row 7). Put differently, in this exercise all effects come through the households working in the modern sector changing their decision of how many family members work and how many hours each working member works (row 5). We add up both components of structural change to get the total effect of structural change on hours worked.

Hours per adult are 9.1 hours lower in the average rich than in the average poor country in the model. The model attributes 67% of this overall decrease of 9.1 hours to income effects. Income effects are thus still the most important driving force of the decrease in hours by development. Taxes and transfers explain 21% of the decrease, and structural change 21%; both thus play a similar, and significantly smaller, role in explaining the decrease in hours. For structural change, this total effect is the sum of two opposing effects: sectoral reallocation explains 54% of the total decrease in hours between rich and poor. However, the decrease in the fixed cost of working in the modern sector predict an increase in hours between rich and poor of 3 hours, thus going in the

Table 5: Counterfactual Experiments

	Hours per Adult			Empl. Rate (in pp.)			Hours per Worker		
	Level	$\Delta_{\text{P-R}}$	Expl.	Level	$\Delta_{\text{P-R}}$	Expl.	Level	$\Delta_{\text{P-R}}$	Expl.
Full Model									
Poor	28.2	–	–	72.6	–	–	38.9	–	–
Rich	19.1	9.1	–	54.5	18.1	–	35.0	3.9	–
Predictions for Rich									
Income Effect	22.1	6.1	67%	67.0	5.6	31%	33.0	5.9	152%
Tax. & Trans.	26.3	1.9	21%	71.4	1.2	7%	36.9	2.0	52%
Struct. Change	26.3	1.9	21%	59.4	13.2	73%	45.0	-6.1	-156%
Sec. Realloc.	23.3	5.0	54%	47.4	25.2	139%	46.9	-8.0	-207%
$\bar{u}_M$	31.3	-3.0	-33%	84.6	-12.0	-66%	37.0	1.9	50%

Note: $\Delta_{\text{P-R}}$ represents the difference between the model prediction for the average poor country and rich country under different settings. The predictions for the poor country are the same for all specifications. The fraction explained (Expl.) corresponds to dividing $\Delta_{\text{P-R}}$ for a given specification with the $\Delta_{\text{P-R}}$ for the full model.

opposite direction. Overall, the extended model reaches the same conclusion as the aggregate model in Section 2, namely that income effects are the most important driver of the decrease in aggregate hours over the development spectrum.

To understand the effects of the different driving forces better, we separately analyze their effects on the two margins, starting with the employment rate in columns 4 to 6. The model predicts an 18.1 percentage point lower employment rate in the average rich than in the average poor country. The model attributes only small parts of this decrease to income effects and taxes and transfers, namely 31% and 7%, respectively. Structural change is the most important driver of the decrease in the employment rate, explaining 73% of the decrease. This comes through sectoral reallocation, which alone predicts an even larger decrease of the employment rate than in the data, namely by 25 percentage points between the average rich and the average poor country. In the model, a family working in the traditional sector sends every member to work, given that there are no fixed cost of working in this sector. Thus, the traditional sector employment rate is always 1. By contrast, the fixed costs of working in the modern sector make it optimal not to send the family members with the highest disutility λ to work, generating a modern sector employment rate below 1. Note that, while in the data we cannot observe sectoral employment rates, in the model this concept exists, given the allocation of families to sectors. The higher TFP growth in the modern than in the traditional sector, and the resulting shift of families from the traditional to the modern sector depicted in Figure 3d, thus leads to a strong decrease in the employment rate with development. The decrease in the fixed cost of working in the modern sector counterweighs this effect, as it predicts an increase in the employment rate within the modern sector. Taken together, structural change also generates the shape of the decrease of employment rates over the development spectrum, predicting the entire decrease to happen between the low- and middle-income countries. In this range of the development spectrum, sectoral reallocation is rapid, leading to a decrease in the employment rate. Between middle- and high-income countries, sectoral reallocation plays a minor role, and the reduction in the fixed cost of working in the modern sector rather generates an increase in the employment rate.

While structural change thus is the most important driver of the decrease in employment rates, income effects and fiscal inputs generate the decrease in hours per worker in the model. The model predicts a decrease of 3.9 hours per worker between poor and rich. Income effects play by far the most important role, predicting 152% of this decrease, namely a decrease of 5.9 hours, followed by fiscal inputs with 52% of the decrease. The income effects generated by the (proportional) increase in labor productivity work stronger on the intensive margin than on the extensive margin, given the fixed cost of working. The same holds true for the effect of taxes and transfers. Structural change, by contrast, predicts an increase in of 6.1 hours per worker between poor and rich countries, contrary to the decrease in the data. This is the mirror-image of its effect on the employment rate: the sectoral reallocation from the traditional into the modern sector predicts an increase in hours per worker, given the higher hours in the modern sector than in the traditional sector at each level of development (see Figure 3e and Figure 3e). At the same time, the decrease in the fixed costs of working in the modern sector induces families working in the modern sector to send more individuals to work, and rather decrease the hours per working family member. While structural change thus overall does not contribute to explaining the decrease in hours per worker between poor and rich countries, it is again crucial for explaining its shape: sectoral reallocation is the only driving force in the model predicting an increase in the hours per worker between low- and middle-income countries.

Summarizing, income effects are the most important driver for the decrease of hours per adult by development. Income effects and taxes and transfers work most strongly though the intensive margin, and structural change through the extensive one. Moreover, structural change is crucial in generating the different shapes of the two margins of labor supply over the development spectrum.

Last, we can also provide a decomposition of the within-country hours-wage elasticity. Since this elasticity is defined within the modern sector, structural change does by definition not play a role for it.[14] In the model, the coefficient of a regression of hours on wages is -0.11 in the poor countries, and -0.01 in the rich countries, thus increasing by 0.1. This increase is entirely explained by fiscal policies: since government transfers are increasing in development, as is tax progressivity, the hours-wage elasticity increases with development. The income effect does not matter at all for the change in the elasticity by development. However, it is the driver of the negative elasticity that prevails in most countries, since with income effects dominating substitution effects, an increase in the individual wage leads to a decrease in hours.

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 largely absent 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 from all income levels. We then use these data to discipline a simple aggregate 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 aggregate model, while transparent, 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,

[14] The only role could come through selection effects of working into the modern sector. We find that these effects are however minimal when we do a formal decomposition.

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 fine 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. We conclude that while tax and transfer systems may be the main feature explaining differences in hours worked across rich countries, 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.

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.

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). Intertermporal 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)

7. 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).

8. Data Appendix

8.1. Converting TFP from Penn World Tables into Sectoral Productivities

Let $\hat{A}$ be the smoothed measure of TFP from the Penn World Tables. Smoothing is done by regressing TFP against the logarithm of GDP per adult; $\hat{A}$ is the predicted value from this regression.

Let $\hat{S}_M$ be the smoothed market hours share from our data. The market hours share is calculated by multiplying the employment share E_M in the market sector (i.e. share of workers in the market sector) with hours per worker in the market sector H_M, and dividing this over total hours, i.e.

$$S_M = \frac{E_M H_M}{E_M H_M + E_T H_T}. \tag{8.1}$$

Smoothing is equally done by (linearly) regressing against the logarithm of GDP per adult (and capping at 1).

We make the assumption that the TFP measure from the PWT is an hours share-weighted average of the two sectoral productivities:

$$log(\hat{A}) = \hat{S}_M log(A_M) + (1 - \hat{S}_M) log(A_T) \tag{8.2}$$

Replacing A_T with equation **??** gives:

$$\begin{aligned} log(\hat{A}) &= \hat{S}_M log(A_M) + (1 - \hat{S}_M)(\phi_0 + \phi_1 log(A_M)) \\ &= (1 - \hat{S}_M)\phi_0 + (\hat{S}_M + (1 - \hat{S}_M)\phi_1) log(A_M) \end{aligned} \tag{8.3}$$

Thus, we can derive the country-specific A_M from the TFP measure of the PWT as

$$log(A_M) = \left[(1 - \hat{S}_M)\phi_0 + (\hat{S}_M + (1 - \hat{S}_M)\phi_1)\right]^{-1} \hat{A}. \tag{8.4}$$

Based on equation **??**, we can then also calculate the country-specific A_T.

9. Model Appendix

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

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

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

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

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

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