

How do Hours Worked Vary with Income? Cross-Country Evidence and Implications

Alexander Bick, Arizona State University
Nicola Fuchs-Schündeln, Goethe University Frankfurt, CFS and CEPR
David Lagakos, UCSD and NBER

August 17, 2017

Abstract

This paper builds a new internationally comparable database of hours worked to measure how hours vary with income across and within countries. We document that average hours worked per adult are substantially higher in low-income countries than in high-income countries. The pattern of decreasing hours with aggregate income holds for both men and women, for adults of all ages and education levels, and along both the extensive and intensive margin. Within countries, hours worked per worker are also decreasing in the individual wage for most countries, though in the richest countries, hours worked are flat or increasing in the wage. One implication of our findings is that aggregate productivity and welfare differences across countries are larger than currently thought.

Bick: Arizona State University, Tempe, AZ 85287 (alexander.bick@asu.edu). Fuchs-Schündeln: Goethe University Frankfurt, Germany (fuchs@wiwi.uni-frankfurt.de). Lagakos: University of California, San Diego, La Jolla, CA 92093 (lagakos@ucsd.edu). We thank Mark Aguiar, Andy Atkeson, David Atkin, Angus Deaton, Maya Eden, Chad Jones, Pete Klenow, Aart Kraay, Norman Loayza, Alan Manning, Valerie Ramey, Richard Rogerson, Andres Santos, Matthias Schündeln, Jesse Shapiro, three anonymous referees, as well as numerous conference and seminar participants for helpful comments and suggestions. Paul Reimers provided truly outstanding research assistance. For further excellent research assistance we thank Caleb Johnson, Patrick Kiernan, Andre Ortseifen, Ang Xing Yi, and Shu Zhang. Fuchs-Schündeln gratefully acknowledges financial support from the European Research Council under Starting Grant No. 262116, and from the Cluster of Excellence "Formation of Normative Orders" at Goethe University Frankfurt. All potential errors are our own. We declare that we have no relevant or material financial interests that relate to the research described in this paper.

1 Introduction

One of the most basic facts in macroeconomics is that aggregate income per capita varies greatly across countries (Klenow and Rodríguez-Clare, 1997; Hall and Jones, 1999; Caselli, 2005). Much less is known about how aggregate hours worked vary across countries. Consider the basic question: are average hours worked higher for adults in high-income countries or for adults in low-income countries? Due to data limitations, the economics literature does not have an answer to this question. This is unfortunate, because if hours enter directly into preferences, then measures of average hours worked at the country level are a key input to understanding welfare differences across countries (Jones and Klenow, 2016).

In this paper, we create a new database of average hours worked using recent household survey data from 80 countries of all income levels. Unlike several other existing databases on hours worked, our data are internationally comparable and allow for rich disaggregate comparisons across and within countries. The surveys we employ are nationally representative and cover workers in all sectors, including the self-employed, which represent the majority of the workforce in low-income countries. We focus most of our analysis on a set of 49 *core countries*, for which international comparability of hours data is as high as possible. In particular, we require that the data from these core countries satisfy three basic criteria. First, the surveys cover the entire calendar year (rather than, say, one month of the year). This is necessary to prevent any bias induced by seasonality in labor demand. Second, hours worked are measured in a consistent way: actual (rather than usual) hours in all jobs (not just the primary job), and in the week prior to the interview. Finally, hours worked cover the production of goods or services counted in the National Income and Product Accounts (NIPA). Thus, our hours measures cover unpaid work in agricultural or non-agricultural businesses, as well as wage employment, but do not cover home-produced services, such as child care.¹

Our main finding is that average hours worked per adult are substantially higher in low-income countries (the bottom third of the world income distribution) than in high-income countries (the top third). In low-income countries, adults work 28.5 hours per week on average, compared to 19.0 hours per week in high-income countries. This difference is both statistically and economically significant, with cross-country differences in average hours per adult (9.5 hours per week) being twice as large as the decline in hours per adult in the United States over the twentieth century (4.7 hours per week) (Ramey and Francis, 2009). In percentage terms, adults in low-income countries work about fifty percent more hours per week than adults in high-income countries. As one simple

¹For a smaller set of countries, we document that hours spent on home production of services in low- and high-income countries follow the same pattern as hours spent on producing goods and services counted in NIPA; we present these findings in Section 4.4.

summary statistic, the slope coefficient from a regression of log hours per adult on log GDP per hour is -0.15. We also decompose average hours per adult into an extensive margin (employment rate) and intensive margin (average hours per worker). We find that cross-country differences in hours per adult are shaped by both margins. Employment rates are higher in the poorest third than in the middle third, and similar between the middle and top third of the world income distribution. Average hours per worker increase between the poorest and the middle third, and then decrease substantially for the richest third. Overall, employment rates account for about three quarters of the decline in hours per adult between low- and high-income countries, while hours per worker account for about one quarter.

After having documented the facts on the aggregate level, we compute average hours by sex, age group, and educational attainment. We find that the pattern of higher hours in low-income countries is quite broad-based, being present in each of the disaggregated categories. We also decompose average hours per worker into three broad sectoral aggregates: agriculture, manufacturing, and services. Hours per worker in agriculture are similar across the world income distribution, while manufacturing and services workers work 7.9 and 13.0 more hours per week in the low- than in the high-income countries. The results from the sectoral breakdowns are reassuring if one has the prior that hours worked are measured more accurately in manufacturing and services. We show that the differences in aggregate hours across countries are not driven by different compositions in age, education, or sectors across countries.

We next ask how hours vary with income within countries, and to what extent the within-country patterns help account for the aggregate patterns we document thus far. Our main measure of individual income is the hourly wage from paid employment, which we compute for all workers who are employed as wage workers.² We find that when pooling individuals across all core countries, hours per worker fall with wages, just as they do in the aggregate: the slope of log hours on log wages is -0.10, compared to a slope of log hours on log GDP per hour of -0.12. When regressing log hours on both log wages and log country GDP per hour together, the slope on log wages falls only modestly, while the coefficient on GDP per hour becomes substantially smaller in absolute terms. Even with country fixed effects, the coefficient on log wages remains similar in magnitude. This suggests that residents of poor countries work more hours on average mainly because of their low wages, rather than because of aggregate factors prevalent in poor countries. We find that this effect is stronger for men than for women, for whom both country effects and individual wages play important roles.

Finally, we explore how the elasticity of hours to wages varies country by country. We find that individual slopes are negative for the majority of countries in our database. Interestingly, these

²We also look at wages plus self employment earnings for a broader set of workers, but with the caveat that self-employment earnings are less well measured than earnings from paid employment.

slopes are systematically lower in countries with lower GDP per capita than in richer countries like the United States, where hours are increasing in individual wages. This is consistent with historical evidence showing that hours-wage slopes for employed workers used to be negative or flat but are nowadays positive in the United States (Aguiar and Hurst, 2007; Costa, 2000), and used to be negative in other OECD countries as well (Huberman and Minns, 2007). The results for the extensive margin are in line with the intensive margin evidence. Using education as a proxy for the permanent income of non-working individuals we find that employment rates are flat by education in poor countries, but increasing by education in richer countries. Our findings suggest that the change from a decreasing hours-income slope to a positive one within countries may be a fundamental feature of the development process.

We conclude by discussing implications of our findings for cross-country differences in aggregate labor productivity and welfare. In the absence of data on hours worked, development accounting typically relies on GDP per worker as a measure of labor productivity, which implicitly treats hours per worker as identical across countries (see e.g. Hall and Jones, 1999; Caselli, 2005; Hsieh and Klenow, 2010). Across our core countries, GDP per worker is a factor 14.3 times larger in the richest third than in the poorest third of countries, while GDP per hour is 16.5 times as large. Thus, after taking hours into consideration, cross-country TFP differences are even larger than previously thought. Our findings also have implications for welfare differences across countries. Building on the work of Jones and Klenow (2016), we construct a flow measure of utility from consumption and disutility of work. Using our hours data, plus standard measures of consumption per capita, we calculate that welfare differs by a factor of 19 between high-income and low-income countries, compared to a factor of 12 when we ignore differences in hours worked. To put it succinctly, residents of low-income countries are not just consumption poor, but also leisure poor.

The rest of this paper is structured as follows. Section 2 highlights the paper's contribution relative to the literature. Section 3 describes our underlying data sources, and our efforts to construct internationally comparable data on hours worked. Section 4 documents that aggregate hours worked per adult are decreasing in the country income level, compares the data to U.S. time-series, and provides some evidence that hours on home services are also decreasing in income. Section 5 presents hours worked differences across countries by gender, education, age, and sector, and shows that differences in aggregate hours are not driven by differences in the demographic composition across countries. Section 6 then analyzes how hours vary with income at the individual level within our sample countries. Section 7 quantifies implications of the aggregate results for the measurement of labor productivity and welfare differences across poor and rich countries. Section 8 concludes.

2 Related Literature

Our study is the first to measure how hours worked vary with income across and within countries of all income levels. Prior studies of aggregate hours worked across countries have almost exclusively focused on rich countries, and in particular on the United States and European countries (e.g., Bick et al., 2017). Explanations of the U.S.-Europe gap in average hours have focused largely on differences in taxation (e.g., Prescott, 2004; Rogerson, 2006; McDaniel, 2011; Bick and Fuchs-Schündeln, 2017), institutions (Alesina et al., 2005) and social security systems (Erosa et al., 2012; Wallenius, 2013; Alonso-Ortiz, 2014). Other studies have focused on understanding changes in hours worked over time, though these have also concentrated on rich countries. For example, McGrattan and Rogerson (2004) and Ohanian et al. (2008) measure changes in hours among OECD countries over time, while Ramey and Francis (2009), and Francis and Ramey (2009) focus on the long-run decline in hours worked in the United States. Aguiar and Hurst (2007) and Costa (2000) study how hours vary with income within the United States historically, and Huberman and Minns (2007) focus on these patterns for a number of OECD countries.

The existing evidence on hours worked from developing countries is quite limited. The study by Lee et al. (2007) presents some evidence on hours from largely non-representative establishment surveys covering wage earners in manufacturing. Their data thus excludes the self-employed and those working in agriculture, which together form the vast majority of all workers in the developing world. Caselli (2005) considers hours worked data for 28 countries from the International Labor Organization (ILO), though just two of these countries are in the bottom half of the world income distribution. Gollin et al. (2014) compare average hours worked among workers in the agricultural and non-agricultural sectors of a large set of countries using nationally representative surveys. Their data are comparable across sectors within each country, though not necessarily comparable across countries. Bridgman et al. (2017) focus on the documentation of household production hours from time use surveys covering 43 countries of all income levels, and show that household production hours and total hours worked are decreasing in development. Similar to us, they find that male hours worked in the market decrease in GDP per capita. In contrast to our results, female hours worked in the market increase in GDP per capita in their data. However, it is unclear how market hours are measured in Bridgman et al. (2017) regarding the distinction between market and household hours in agriculture, and the treatment of unpaid family workers. Jones and Klenow (2016) consider hours worked in their study of welfare differences across countries. Their micro data cover however only 3 middle- and 3 low-income countries, and of these 6 only 2 qualify as core countries in our study. For their extended analysis, they rely on hours data from the Penn World Tables (PWT). None of these studies analyzes how hours worked vary with wages on the individual level.

The Total Economy Database, run by the Conference Board, recently released data on annual hours worked per worker, in addition to employment rates, for an unbalanced panel of countries, with the earliest data coming from 1950. These data are also included now in the PWT. Many of these data points are however interpolated or extrapolated, especially for the low- and middle-income countries. Of the data points that have an actual source available, the quality is highly questionable in terms of consistency of hours measurement, activity coverage, and potential biases from seasonality. We detail some of these data quality issues in Online Appendix Section A.4.³

3 Data

In this section, we describe the survey data underlying our analysis. We then introduce the criteria that we use to define the set of core countries, which are those that have the most scope for international comparability. Next, we describe how we measure hours per adult, employment rates, and hours per worker.

3.1 Data Sources

Our analysis draws on nationally representative household surveys. The key advantage of using household surveys, as opposed to firm surveys or administrative records, is that our measures of labor supply are not restricted to activities for which individuals receive a wage, but also include self-employed and unpaid family work. As is well known, the self-employed form an important fraction of the workforce in all countries, and particularly so in developing countries, see e.g. Gollin (2008).

All of the surveys we employ are publicly available for researchers, mostly via an application through national statistical agencies or similar institutions. We were able to collect nationally representative data for 80 countries with a population of at least one million. For 32 of our countries we can draw from harmonized data sets, for which efforts have already been made to standardize questions across countries. These comprise the European Union Labor Force Survey (EULFS; 26 countries) and the International Public-Use Microdata Project (IPUMS; 6 countries). For the remaining 48 countries, we draw on country-specific censuses, household or labor force surveys,

³To give some concrete examples, from the 304 country-year observations that come from country-years in which the country's GDP would place it into the lowest tercile of the world income distribution today, 89, i.e. almost one third, are inter- or extrapolated. Of the remaining data points, the vast majority, namely 196, come from data provided by the Asian Productivity Organization (APO). For 83 of the 196 APO data points, there is no information available regarding the original data source, and for a further 71 the original data source does not contain any information on hours per worker. Hence, only for 42 data points from 4 countries is there any information on hours per worker in the original data, but the quality is still unclear. Some of the other low-income data points in the Total Economy Database come from either Hofman (1998) or Maddison (2001), which in turn often extrapolate data or impute average hours per worker using only the number of statutory public holidays and vacation days.

including 19 surveys conducted as part of the World Bank’s Living Standards Measurement Studies (LSMS).⁴

When multiple years of appropriate data are available, we choose the year closest to 2005. Most of our data are within a few years of 2005; exact years, data sources, and sample sizes for all countries are given in Online Appendix Table C.1. Our sample sizes range from 5,000 to over 700,000 individuals. We focus on all individuals of at least age 15, whom we refer to as “adults”.⁵

3.2 Core Countries

The key measurement challenge we face is that not all of our surveys are conducted in the same way, and more specifically, not all surveys collect hours information in the same way. To ensure that international comparability is as high as possible, we focus our main analysis on a set of *core countries* which we define to be those that satisfy the following three criteria.

1. Activity Definition: We restrict attention to hours worked in the production of output that is counted in NIPA. These include hours worked in wage employment as well as hours in own-account agricultural or non-agricultural work, whether or not that output is sold or used for own consumption (see e.g. Gollin et al., 2014). This is important if we want to maintain a nationally representative sample of workers, particularly in the poorest countries, where agricultural work and self-employment are widespread. Not included in our definition of hours worked are hours spent on non-market services, such as cleaning or home-provided child care; we return to the issue of home-produced services in Section 4.4.

2. Hours Worked Information: We focus on actual hours worked, rather than usual hours worked, since individuals may work more or less than usual in a given week, due to e.g. over-time or sickness. We also focus on all jobs, rather than just the primary job, since many individuals have multiple jobs. Finally, we focus only on surveys that ask respondents about hours worked in the last week or in a recent reference week, since longer time periods may lead to recall bias.

3. Time Coverage: We restrict attention to surveys that cover the entire calendar year. While all of our surveys are nationally representative in terms of the covered population, some are conducted over the entire year, and others are conducted over only a few months or weeks. Using these partial-year surveys creates potentially biased estimates of hours worked unless the survey period happens to be representative of the entire year. This bias may be more pronounced in developing countries, which are largely agricultural and hence seasonal. Online Appendix Section A.1 provides a detailed explanation of how we determine the time coverage of each survey and which surveys qualify as covering the entire year according to our definition.

⁴Note that this does not imply that these standardized surveys are all in our sample of core countries. All ELFS countries are core countries, while none of the IPUMS countries and only 9 of the LSMS countries are core countries.

⁵The United States is an exception here as the youngest available age is 16.

Out of the 80 countries in our sample, 49 qualify as core countries. All the non-core countries satisfy the first criterion on the activity definition but have either non-standard hours worked information or cover less than the full calendar year (or both).

3.3 Measuring Employment and Hours Worked

Our measures of employment rates and hours worked rest on two key variables: the self-reported employment status and actual hours worked in all jobs in the last week. We measure the employment rate as the fraction of all adults that report being employed or have positive hours worked. We measure hours per worker as the average hours worked in all jobs in the reference week among all those who are employed. Both measures are calculated with the individual survey weights. We measure hours per adult as the product of the employment rate and average hours per worker. We provide more details on our calculations in Online Appendix Section A.2. Since the surveys of the core countries cover the entire year, vacations, sick leave, or other reasons for seasonality should in principle be covered, and multiplying the reported numbers by 52 gives an estimate of annual hours worked per adult (see our discussion of potential biases in Online Appendix Section A.5).

Our definition of whether a country in our sample is a low- (bottom third of the world income distribution), middle- or high-income (top third of the world income distribution) country is based on GDP per capita in the year 2005 for all countries in the Penn World Tables version 9.0 (see Feenstra et al. (2015) for a detailed description of the PWT). Specifically, we use expenditure-side real GDP at chained PPPs in 2011 US \$ (*rgdpe*). We find similar levels of average GDP per capita when comparing each of these terciles in our core and full set of countries to all countries in the PWT; see Online Appendix Table C.2. When plotting aggregate hours worked against GDP per capita, we use GDP per capita for each country from the same year for which we have the hours data.

4 Aggregate Hours by Income Across Countries

In this section, we document that average hours worked per adult are substantially higher in low-income countries than in high-income countries, and that the same holds true for both the extensive and the intensive margins of labor supply. Last, we provide suggestive evidence that also hours spent on the production of home services are decreasing by development.

4.1 Average Hours Worked Per Adult

Figure 1 plots average weekly hours per adult against log GDP per capita for our core countries. Vertical lines separate the three terciles of the world income distribution; 11 countries fall in the

bottom tercile, 15 in the middle tercile, and 23 in the top tercile. The figure shows that average hours per adult are downward sloping in income per capita. The poorest countries in the world range from a low of around 24 hours per week in Malawi, Rwanda, and Uganda to a high of almost 40 hours per week in Cambodia. The richest countries average between a low of around 16 hours in Italy, Spain, Belgium, and France and a high of 24.4 hours in the United States. Iraq has the lowest hours per adult in our sample, which is driven entirely by women, as discussed in Section 5.1.

Panel A of Table 1 reports in the first row the average hours per adult by income tercile in our core countries. In these countries, average hours per adult are 28.5 hours per week in low-income countries, compared to 21.7 hours in middle- and 19.0 hours in high-income countries. In terms of economic significance, the 9.5 higher weekly hours in the low-income group correspond to 50 percent higher hours than in the high-income group.⁶ Regressing the logarithm of hours on the logarithm of GDP per hour worked yields a slope coefficient of -0.15.⁷

Given that the number of core countries is relatively small, particularly in the lower end of the income distribution, we conduct statistical tests of the hypothesis that average hours worked in all countries are drawn from the same distribution. We do so using permutation tests, which have more favorable small-sample properties than other commonly used tests, such as *t*-tests (Lehmann and Romano, 2005). Panel B of Table 1 reports the results of these permutation tests. For the core countries, shown in the first row, the observed difference in mean hours between the low- and high-income groups is 9.5 hours per week, between the low- and middle-income group 6.8 hours per week, and between the middle- and high-income group 2.7 hours per week. All *p*-values are well under one percent. We conclude that the decreasing average hours over the income terciles are unlikely to be a coincidence.

Rows 2 and 3 of Table 1, Panel A, report average hours per adult by income tercile in two broader sets of countries: (i) the core plus those countries having a survey that covers only part of the year, and (ii) all countries in our data set, regardless of how hours are measured. Covering more countries comes however at the cost of lower international data comparability. Across the 73 core plus partial-year survey countries, average hours worked are 26.7 in the low-, 22.5 in the middle-, and 19.4 in the high-income countries. Thus, within the low-income countries average hours worked are slightly lower in this group than in the core, while hours worked in the middle- and high-income groups are similar to the core. Going to the full set of 80 countries, average hours

⁶In the main analysis, we take unweighted averages across countries. When weighting by population, hours differences between the bottom and top thirds of the world income distribution are similar: averages in the low-, middle- and high-income groups are 27.8, 21.8, and 20.3 hours per week.

⁷We regress on the logarithm of GDP per hour worked rather than GDP per capita, because GDP per hour worked is an aggregate productivity measure analogous to the individual wage, which we use as a regressor in a similar regression in Section 6. Figure C.2 in the Online Appendix shows the corresponding scatter plot of the logarithm of average hours worked per adult and the logarithm of GDP per hour worked.

per adult rise by 0.2 to 0.3 weekly hours in all income groups (see also Figure C.1 in the Online Appendix.) As Panel B of Table 1 shows, all differences across the income groups are significant at the one percent level in the broader set of countries. We conclude that our finding of higher hours per adult in poor countries than in rich countries holds in a broader set of countries as well as in our core countries. From here on we focus on the core countries due to their higher degree of international comparability.⁸

4.2 Comparison to Time-Series Data from the United States

Our research is motivated by the question of how hours worked vary with income in the cross-section of countries. Complementary evidence on the relationship between income and hours worked comes from time-series data. With time-series data, one can ask how hours looked like in the currently rich countries back when they were poor. Comprehensive and reliable historical data on hours worked are unfortunately hard to obtain, and the discussion on the complications of constructing reliable hours worked data in Section 3 makes clear why this is the case. However, for the United States, data spanning over 100 years are available from Ramey and Francis (2009). Yet, even 100 years ago, the United States was as rich as current middle-income countries, so these time-series data do not span as large an income range as our data does.

Figure 2 plots Ramey and Francis (2009)'s U.S. time-series of average hours worked per adult (individuals aged 14+) from 1900 to 2005 (gray line), and average hours per adult from our data (black dots). Average adult hours per week in the United States decreased from 27.7 hours per week in 1900 to 23.0 hours in 2005, corresponding to a decline of 4.7 hours per week. Interestingly, the patterns for hours in the cross-section of countries is quantitatively similar over the range of GDP per capita spanned by the United States in the last century, only with slightly higher overall hours in the United States.

4.3 Extensive and Intensive Margins

Differences in hours worked per adult stem from differences in employment rates, which represent the extensive margin, and average hours per worker, which represent the intensive margin. Figure 3 plots employment rates (top panel) and average hours per worker (bottom panel) for our core countries. The figure shows that employment rates are decreasing for much of the income distribution, with a modest increase for the richest countries, while hours per worker follow more of a hump-shaped pattern.

⁸We conclude that potential biases from seasonality, measurement of vacation days and public holidays, and child labor would all likely increase measured hours worked differences between rich and poor countries; see Online Appendix Section A.5.

Table 2 reports the average employment rates and hours per worker in each country income group. In the low-income countries, the average employment rate is 75.3 percent. In middle- and high-income countries, employment rates are 52.7 and 54.5 percent, respectively. Along the intensive margin, workers in low-income countries average 38.4 hours per week, compared to 41.1 hours and 35.0 hours in the middle- and high-income countries; thus, on average there is a mild hump-shape in hours per worker among the three income groups. For employment rates, the large difference between the low- and middle-income countries (22.6 percentage points) is statistically significant at the one-percent level, while the (negative) difference between the middle- and high-income countries is not significant. For hours per worker, the opposite is true. The small low-middle difference is statistically insignificant, while the large middle-high difference (6.1 hours per week) is significant at the one-percent level. Thus, average hours per adult are shaped by the two margins differently, with employment rates accounting for the decline in hours per adult between low- and middle-income countries, and hours per worker accounting for the decline between middle- and high-income countries. Overall, we calculate that employment rates account for around three quarters of the cross-country differences in hours per adult, while hours per worker account for around one quarter; see Online Appendix Section A.3 for details.

4.4 Time Spent on Production of Home Services

Until now, we have focused attention entirely on hours worked in the production of output counted in NIPA. A large literature has emphasized broader notions of work, however, including hours spent on home production of services (Parente et al., 2000; Aguiar and Hurst, 2007; Ngai and Pissarides, 2008; Ramey, 2009; Ngai and Pissarides, 2011; Aguiar et al., 2012, 2013; Rendall, 2015; Duernecker and Herrendorf, 2016). In this section, we consider hours of home production using a smaller set of countries for which we have data.

Hours spent producing home services are notoriously hard to measure. The two most important reasons are the difficult differentiation between leisure and home production of services in some categories, and the possibility of multi-tasking. Both difficulties apply especially when it comes to child care, but can also arise in other categories like cooking (see Aguiar and Hurst (2007) and Ramey (2009) for excellent discussions of the difficulties of measuring leisure and home production hours in general). Questions covering time spent on home production of services are therefore not usually included in labor force surveys or censuses. However, a few of the surveys we use do in fact ask about time spent on some categories of home production of services. We complement these surveys with data from the Multinational Time Use Study (MTUS).⁹

⁹For each country, we use the year closest to 2005. Online Appendix Table C.3 provides an overview of the countries with data on time use by income terciles. All data from the bottom and middle terciles, except South Africa, and data from Russia come from our main data source for the respective country. All other data come from the MTUS.

We provide evidence on average weekly hours spent in five aggregated service categories, namely cooking (including preparing food and washing dishes), cleaning, child care, shopping, and collecting water and firewood. These data should be considered suggestive evidence: we do not apply the same standards to ensure comparability across countries that we apply when calculating hours worked in the market or in the production of home produced goods. The MTUS covers all five categories except collecting water and firewood. The other individual country surveys often cover only a subset of the categories. For each category and each income tercile, we have data from at least four countries, with the exception of hours spent on collecting water and firewood, which has minimal data outside the bottom tercile.

Table 3 presents the average hours spent on each of the time use categories by income group, with the number of countries for each category and group in parentheses. Average hours are lowest for the high-income countries in every single category except shopping. The totals amount to 25.3 weekly hours in the bottom tercile, 25.8 hours in the middle tercile, and 18.2 hours in the top tercile. We conclude that our main finding of higher hours worked in low-income than in high-income countries is still present once we consider time spent on broader categories of work, at least using these data. Our findings here are consistent with those of Bridgman et al. (2017). Using richer data on broad time use categories from 43 countries, they find that home production hours decrease in GDP per capita.

5 Hours Worked by Income for Different Demographic Groups

So far, we document that aggregate hours worked are decreasing in GDP per capita. We now turn our attention to potential heterogeneity related to this fact. We document hours worked per adult by sex, education, and over the life cycle, and hours worked per worker by sector across our core countries. We show that the finding that average hours worked per adult are substantially higher in low-income countries than in high-income countries is quite broad-based. The positive hours difference between low- and high-income countries is statistically significant for both sexes, all education groups, and all sectors except agriculture. Moreover, the decline of aggregate hours by income is not driven by different compositions of the population across countries.

5.1 Average Hours per Adult by Sex

We start by analyzing average hours worked for both sexes across our set of core countries. As Table 4 shows, for both men and women hours per adult are decreasing by development. For the low-income countries, men average 32.7 hours per week, while in the middle- and high-income countries they average 27.3 and 23.8 hours per week, for a difference of 8.9 hours per week be-

tween low and high. Women work fewer hours than men in all income groups, but show the same pattern of higher hours in poorer countries, with a very similar difference of 10.0 hours between low- and high-income countries.¹⁰

5.2 Average Hours by Education Group

Patterns of hours worked have been shown to differ systematically by education group within the United States, see e.g. Aguiar and Hurst (2007). Do similar patterns arise in countries in other parts of the world income distribution? Our data allow us to consistently define three broad education groups: those with (1) less than secondary school, (2) secondary school completed (but not more), and (3) more than secondary school. We restrict attention in this exercise to adults aged 25 and above, so as to focus on those who have most likely completed schooling.

Table 5 reports average hours per adult by education group; see Figure C.4 of the Online Appendix for the plots. All three education groups feature higher hours in the poorer countries. Among individuals with less than secondary school, average hours are 31.8 in the low-income countries compared to 18.3 in the middle- and just 12.5 in the high-income countries. Thus, for the lowest education group the difference between hours in low- and high-income countries amounts to 19.3 hours.¹¹ Individuals in low-income countries with secondary school completed work 13.7 more hours per week on average than their counterparts in high-income countries, and for individuals with more than secondary education the difference between low- and high-income countries falls to 12.5 hours. Within each country group, average hours are higher for more educated individuals than for the less educated, though less so for the low-income countries. We return to this evidence in Section 6.5 below.

5.3 Hours Worked over the Life Cycle

In this subsection, we document hours worked over the life cycle for the three country income groups. Figure 4 plots average hours for five year age groups, starting at age 15-19 and ending at age 85-89. Since we do not have panel data, we cannot distinguish between cohort and age effects. We interpret the data as age effects, but one should keep in mind the caveat that we could be capturing cohort effects at least to some extent. The well-known hump-shape of hours over the life cycle is present for all three country income groups. Most importantly, the pattern of decreasing

¹⁰Online Appendix Figure C.3 plots average hours per adult for men (top panel) and women (bottom panel) for each country. A notable feature of the graphs by sex is that female hours are substantially lower for countries with large Muslim populations, such as Iraq, Pakistan, and Turkey.

¹¹This dramatic difference is partly caused by an age-composition effect, because individuals with less than secondary education are most prevalent among the old in high-income countries. Focusing on individuals aged 25 to 54 (instead of all individuals aged over 25) reduces the hours difference between high- and low-income countries of the lowest education group from 19.3 hours to 10.0 hours.

hours by country income is present at each single age group. The largest differences arise for older individuals. Starting at the age group 55-59, the differences in hours worked between low- and high-income countries are increasing up to age 65-69, at which point they start decreasing again. This points to the absence or existence of social security programs as an important driver of hours worked differences around the retirement age. On average, individuals aged 55+ work 21.8 hours per week in the low-income countries, compared to 12.1 and 7.6 hours in the middle- and high-income countries. Among individuals younger than 55, the average hours difference between low- and middle-income countries amounts to 4.9 hours, and between middle- and high-income countries to 0.4 hours.¹²

5.4 Hours Per Worker by Sector

Which sectors contribute most to the patterns of hours per worker that we document? To answer this question, we compute average hours per worker by three broad sectoral aggregates, which we can define consistently across countries. These are agriculture (including forestry and fisheries), manufacturing (including mining and utilities), and services. We assign each worker to one of these sectors based on their primary sector of employment (though their hours cover all jobs). We focus on hours per worker (rather than hours per adult), since industry is only well defined for those currently working.

Table 6 presents the average hours per worker by industry. Among agricultural workers, differences in average hours are statistically insignificant across the three country groups, with the low-, middle- and high-income countries working 36.0, 38.3 and 39.7 hours per week, respectively. Manufacturing workers work longer hours in the low-income countries, at 44.9 hours per week, compared to 42.5 in the middle-income and 37.0 in the high-income countries. The differences are even more substantial for services, where workers in the low-income countries average 47.7 hours, compared to 41.8 and 34.7 hours per week in the middle- and high-income countries. It is reassuring that the hours differences are so pronounced in manufacturing and services, if one has the prior that hours worked in these two sectors are measured more accurately since the prevalence of paid employment is higher. In contrast to age and education, the country income groups do not share a common pattern of hours worked differences across sectors.

5.5 Compositional Effects on Aggregate Hours

Given that we document substantial differences in hours worked by age, education, and sector in all three country income groups, the question arises whether the differences in aggregate hours by income arise due to substantial different compositions of the population by development. To

¹²De Magalhães et al. (2017) find similar patterns for several African countries.

gauge how important the composition of the population by country is for determining the aggregate hours differences, we conduct several counterfactual exercises. Starting with the age composition, we first compute in each country average hours per person for 5 year age groups. We then calculate hypothetical average weekly hours per adult in each country by multiplying U.S.-population weights for the age groups with average hours of the corresponding age group in each country, and then summing up over all age groups. We do a similar counterfactual exercise for the educational composition relying on the three education groups, and another one combining the age-education composition. Last, we conduct the same exercise for hours per worker, imposing the U.S. sectoral composition.

The results are shown in Table 7. The first two rows relate to the age structure. Average weekly hours per adult essentially do not change when the U.S. age structure is imposed. Imposing the U.S. age structure on poor countries shifts weight from the older population to the younger one; since both work similarly fewer hours than prime-aged individuals, the shift has on net a minimal effect. Rows 3 and 4 focus on the educational composition (and thus on the population 25+). The hypothetical hours difference between low- and high-income countries imposing the U.S. education composition is somewhat larger than the actual hours difference. By imposing the U.S. education composition, the share of higher educated individuals in the low-income countries rises, increasing the hypothetical hours there. However, since the hours gradient in education is smallest in the low-income countries, the difference in hypothetical hours is still relatively close to the difference in actual hours. Row 5 then combines the age and education structure, and finds a very similar difference in hypothetical hours as in actual hours between high- and low-income countries. The hypothetical values for hours per worker if the U.S. sectoral composition is imposed in all countries are presented in the last row of Table 7. What we find, not surprisingly, is that the hypothetical hours per worker would be even higher in low-income countries if those countries had the U.S. industry composition, given the small share of agriculture in the United States and the low average hours in the agricultural sector in low-income countries.

Summarizing, we unequivocally find that differences in the population composition across countries are not the main driver of the aggregate hours worked differences by development.

6 Individual Hours Worked and Wages

So far, we have provided evidence that hours are decreasing in income at the country level. In this section, we look at how hours vary with income within countries. We first pool all our individual data for our core countries and compute an hours-wage elasticity using individual-level wage data, and ask whether the aggregate hours-income relationship we document thus far is accounted for mostly by the individual elasticity, or whether there are aggregate features particular to poor coun-

tries that lead to higher hours there, such as a lack of social security programs. Finally, we compute the hours-wage slopes country by country, and ask how these slopes vary across countries.

6.1 Constructing Individual Wages

There are two major challenges in constructing wages on the individual level. First, we observe individual earnings only for individuals who are currently working. Thus, we focus on the intensive margin of labor supply, i.e. hours per worker. Second, while constructing hourly wages is relatively straightforward for individuals in paid employment, it is a challenging task for self-employed individuals, who make up the majority of the work force in developing countries. We thus construct two wage rates: the first is an individual hourly wage from paid employment, which we only construct for workers in paid employment; the second is an hourly wage including earnings from self-employment, which we construct for all individuals working positive hours. We exclude Namibia, Laos, and South Africa because of missing earnings data. We briefly summarize the construction of the two measures here, and more details as well as a validity check can be found in Online Appendix Section A.6.

Wage from Paid Employment For individuals in paid employment, we construct an hourly wage (expressed in 2011 PPP-adjusted US \$) by dividing monthly earnings from paid employment by actual weekly hours worked from paid employment in the reference week multiplied by 4.33. We only include individuals who are paid employees in the main job. If they also report being paid employees in any additional jobs, we sum up earnings and hours over all relevant jobs, if available. The European Union Labor Force Survey, which is our main data source for European countries, does not include individual earnings, and thus in this analysis we use the European Union Statistics on Income and Living Conditions (EU-SILC) for the European countries. Moreover, we have to recur to usual rather than actual hours for the European countries and the United States. The main caveat concerning evidence from the group of wage earners in paid employment is that they are a selected set of all employed, particularly in lower income countries (see Table A.1 in the Online Appendix).

Figure 5 presents a simple way of summarizing our individual-level data across all sample countries for which we can construct an individual wage. We sort individuals into wage deciles within each country, and calculate for each decile the average wage and average hours worked. We then plot these against each other for all 46 core countries, and connect all deciles within a given country with a line. Notably, there is substantial overlap of individual wages between countries of different income levels. While there appear to be significant country fixed effects, especially among the low-income countries, the decrease of hours by wages seems to hold both within and across countries.

Wage from Paid Employment Plus Self-Employment We also construct an alternative hourly wage which includes earnings from self-employment for a broader set of individuals, encompassing all those who are working and have non-missing earnings data. In most of our surveys, self-employment income is reported at the individual level, and so we construct an individual wage by summing all wage and self-employed earnings and dividing by total hours worked at all jobs. An important caveat is that although self-employment income is reported at the individual level, it is still possible that other family members supply hours of work in order to help earn the reported income. Furthermore, self-employed income is reported in different ways across countries, and there are limits to how well we can standardize self-employment income. Fortunately, as we show below, our overall findings are similar with and without self-employment earnings.

Online Appendix Figure C.6 plots average hours against average wages for the deciles of the wage distribution in each country. As in Figure 5, there is substantial overlap in wages across countries. The overall pattern is still a negative relationship between hours and wages, though substantially flatter than the pattern for wage earners, particularly in the poorer countries.

6.2 Individual vs. Country Income

Our findings thus far of a negative relationship of hours and GDP per capita at the country level raise a natural question: do individuals in poor countries work more hours because they themselves have low income, or because they live in a poor country? If it is the former, this points to a strong role for preferences in which income effects outweigh substitution effects (see e.g. Boppart and Krusell, 2016). If country effects dominate, this points to institutional features of poor countries that raise labor supply relative to rich countries.

To answer this question, we build a world data set containing all individual observations from all 46 core countries, and regress hours worked of individual i living in country c on the individual wage (w) and country GDP per hour worked ($GDPph$), including as controls age and age squared, and clustering standard errors at the country level:

$$\log(h_{ic}) = \alpha + \beta_w \log(w_{ic}) + \beta_{GDP} \log(GDPph_c) + \delta_1 age_{ic} + \delta_2 age_{ic}^2 + \varepsilon_{ic}. \quad (1)$$

The results are shown in Table 8, using the sample of wage earners from paid employment. Each column is a different regression specification. Panel A shows our findings for both sexes. In the first two columns, we show results if either only GDP per hour or only the individual wage are included as regressors. In this case, both turn out to have negative and highly significant coefficients of -0.12 and -0.10, respectively. When both regressors are included at the same time (third column), the coefficient on the individual wage drops slightly to -0.07, but the coefficient on country GDP per hour drops more substantially to -0.04. Both coefficients become insignificant.

In the last column, we replace GDP per hour with country fixed effects, and the coefficient on the individual wage remains again largely unchanged at -0.09. We thus conclude that a low individual wage correlates significantly with high hours worked per worker across the entire world income distribution. On top of this effect, individuals living in poorer countries tend to work more hours, though that effect is more modest.

The next two panels of Table 8 show the estimation results separately for men and women. The main difference between both is that, once both GDP per hour or country fixed effects and individual wages are included, the coefficient on the individual wage is more negative for men than for women, while the coefficient on GDP per hour worked even turns slightly positive for men, but remains negative for women. The positive coefficient on country productivity for men is driven by variation within the group of low-income countries (see Figure C.7 in the Online Appendix).

We find similar regression results for the sample including earnings from self-employment (see Table C.4 in the Online Appendix). The only substantive difference is that the negative coefficient on GDP per hour is smaller in an absolute sense in the specification without individual wages as regressors, and turns positive for both men and women once individual wages are added. Thus, even more so than in the sample of wage workers, high hours seem to be driven by low individual wages rather than low country income. We want to stress that all these estimates, as well as those relating to Equation 2 below, only capture partial correlations, and should not be interpreted as structural elasticities.

6.3 Individual Hours-Wage Elasticities by Country

How does the hours-income elasticity vary from country to country? To answer this question, we follow Costa (2000) and regress within each country the logarithm of individual hours worked on the log wage and age and age squared, separately for men and women:

$$\log(h_i) = \alpha + \beta_w \log(w_i) + \delta_1 \text{age}_i + \delta_2 \text{age}_i^2 + \varepsilon_i. \quad (2)$$

Figure 6 plots the country-specific coefficients β_w against log GDP per capita, for men in Panel (a) and for women in Panel (b). It shows a negative coefficient for low- and middle-income countries, which increases towards zero starting with the richer end of the middle-income countries, and turns positive for the richest countries. Thus, in the majority of countries hours are decreasing with the individual wage, while only in the richest countries this relationship is reversed. This cross-country evidence is in line with the historical evidence by Huberman and Minns (2007) for a subset of OECD countries, and the time-series evidence for the United States. Costa (2000) runs the same regression on U.S. data from different time periods, and finds a negative coefficient in the

1890s, which increases over time and turns positive in the female sample by 1973 and in the male sample by 1991. Her estimates are included in Figure 6 and are in line with the values we find in the respective stages of development in the cross-section.¹³

Thus, in line with the cross-country evidence of a decline in hours worked per worker with income, in the majority of countries hours per worker are also declining in the individual wage. Only for the richest countries does the relationship between hours and individual wages turn positive. The same holds true in the sample including earnings from self-employment; see Figure C.8 in the Online Appendix.

6.4 The Role of Division Bias

Since individual hours worked are the dependent variable in our regressions, but also feature as a denominator on the right hand side of the regression equation in individual wages, any classical measurement error in hours leads to a spurious negative correlation between hours and wages, see e.g. Borjas (1980). To gain some insights into the potential size of this division bias, we conduct the main robustness check recommended by Borjas (1980). The strategy is to run regressions in which the measure of hours on the left hand side is replaced with an alternative measure (e.g. usual hours worked rather than actual hours worked, which are used to construct wages). Table A.4 in the Online Appendix shows results for the five countries for which such an alternative hours measure is available. For Turkey, Mongolia, and Uganda, replacing hours with an alternative measure increases the coefficient β_w , providing evidence for division bias. However, the coefficient change is relatively moderate: the coefficient increases by on average 40%, but still remains negative. Only in the Peruvian sample, using an alternative hours measure as dependent variable turns the coefficient positive. By contrast, in the US sample the use of an alternative hours measure leaves the coefficient virtually unchanged. Overall, this leads us to conclude that the division bias might be present, but is unlikely to substantially alter our findings.

6.5 Variation of Employment Rates by Education

For non-working individuals, wage information is not present, so we use education as a proxy for permanent income. Table 9 shows employment rates by education for the age group 25 to 54 in the upper panel, and for the age group 55+ in the lower panel, and separately for men and women. Employment rates are essentially flat by education in the low-income countries for all four age-gender groups. By contrast, in the middle-income countries, and even more so in the high-income countries, employment rates are strongly increasing in education for all age-gender groups.

¹³Costa restricts the sample to individuals aged 25 to 64. If we do the same, our coefficient for the United States falls from 0.09 to 0.07 for men, and from 0.12 to 0.10 for women.

Of course, this evidence is only suggestive, as education is only a rough proxy for permanent income. However, the pattern of a positive relationship between employment rates and education in rich countries, and a flat relationship in poor countries, is in line with the patterns we find above between hours worked per worker and wages. This relationship is also positive in rich countries, but negative in poor countries.

One potential explanation for the diverging evidence from rich and poor countries could be the size of the welfare state: more generous welfare states in rich countries might create disincentives for low-productivity workers to participate in the labor market or work long hours. Another common pattern that we find is that both the gradient of employment rates in education, and the gradient of hours worked per worker in wages are more positive in high-income countries for women than for men. This points to within-household specialization and insurance in richer countries. In the low-income countries, by contrast, both gradients are similar for men and women.

7 Implications

Summarizing our findings thus far, we establish a robust negative relationship between hours worked and income on the aggregate level, for different demographic subgroups, and also on the individual level for the majority of countries. In this section we discuss two aggregate implications of these findings. The first is about development accounting, which tries to account for cross-country income differences using observable factors of production like physical and human capital. Our aggregate findings for average hours imply that labor productivity differences, and hence total factor productivity (TFP) differences, are even larger across countries than previously thought. The second implication is about welfare differences across countries. Our finding of higher hours worked for poorer countries, and poor individuals more generally, imply that welfare differences across countries are even larger than previously thought.

7.1 Development Accounting

What are the implications of the fact that hours per adult are higher in low-income countries than in high-income countries for the measurement of labor productivity differences? The literature on development accounting attempts to explain cross-country differences in output per worker using aggregate stocks of human and physical capital per worker. One basic piece of information missing from this literature has been hours worked per worker, which has limited the ability of researchers to accurately measure labor productivity (Klenow and Rodríguez-Clare, 1997; Hall and Jones, 1999; Caselli, 2005; Hsieh and Klenow, 2010). In the absence of data on hours worked, virtually all previous studies in this literature have measured labor productivity as GDP per worker.

Our data suggest that GDP per worker underestimates the true labor productivity differences across countries. To investigate this quantitatively, Table 10 shows two different measures of labor productivity: GDP per worker and GDP per hour. The first three columns report the respective values for our three core country income groups, normalizing the value of the high-income countries for each measure to 100, and the last column presents the ratio of the respective variable for high- to low-income countries. Focusing on this ratio, GDP per worker is 14.3 times larger in high-income countries than in low-income countries. We can improve on this by adding our data on hours worked. Since workers in low-income countries work on average 3.4 hours per week more, the ratio of GDP per hour in high- over low-income countries is even higher than the ratio of GDP per worker, specifically amounting to 16.5 instead of 14.3. This corresponds to 15 percent larger labor productivity differences across countries than implied by GDP per worker. Middle-income countries are also less productive relative to high-income countries based on hours worked compared to based on employment alone.¹⁴

Our findings imply that development accounting rests even more on the residual TFP term once cross-country differences in hours are taken into consideration. This casts doubt on theories of development that operate through lower labor input in poorer countries. Landes (1999), for example, points to hot weather in the tropics as a cost of working there. In his theory, TFP differences across countries are in part explained by differences in labor effort, with high effort in economies with higher TFP.

7.2 Welfare Differences Across Countries

How do measured welfare differences between rich and poor countries change if differences in hours worked are taken into account in addition to consumption differences? To answer this question, we broadly follow the approach by Jones and Klenow (2016), who study welfare differences across countries taking into account a wide range of outcome variables. We focus on hours worked, since our data set has a broader coverage of hours in low- and middle-income countries than the data in Jones and Klenow (2016). We rely on the standard neoclassical growth model, as used in the literature analyzing U.S.-Europe hours differences, see e.g. Prescott (2004). The key modification is a non-homotheticity in preferences in the form of a subsistence consumption requirement \bar{c} , as in Ohanian et al. (2008), and Restuccia and Vandenbroucke (2014). This subsistence consumption requirement implies that the income effect dominates the substitution effect. As consumption rises, this dominance becomes weaker, and in the limit the two effects cancel out.

¹⁴Feenstra et al. (2015) recommend using *rgdpo* rather than *rgdpe* for productivity comparisons across countries. Online Appendix Table C.5 shows that using *rgdpo* does not alter our conclusions.

Each country has a representative, infinitely lived household which maximizes life-time utility:

$$\max_{\{c_t, h_t\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t \left(\log(c_t - \bar{c}) - \alpha \frac{1}{1 + \frac{1}{\phi}} h_t^{1 + \frac{1}{\phi}} \right), \quad (3)$$

where c_t is consumption, h_t are hours worked, utility is separable in consumption and leisure, α determines the relative disutility weight of work, and ϕ is the Frisch elasticity, which allows for a more flexible labor supply elasticity than does Prescott (2004). The household faces a standard inter-temporal budget constraint, and there is a representative firm with an aggregate Cobb-Douglas production function. Combining the intra-temporal first-order condition of the household and the first-order condition of the firm gives the following solution for optimal hours:

$$h_t = \left[\frac{1 - \theta}{\left(\frac{c_t}{y_t} - \frac{\bar{c}}{y_t} \right) \alpha} \right]^{\frac{\phi}{1 + \phi}}. \quad (4)$$

Taking aggregate consumption (*csh_c* in PWT 9.0) and aggregate output (*rgdpe* in PWT 9.0) from the data, we use this equation to calibrate the preference parameters. We set the Frisch elasticity equal to $\phi = 1$ (Keane and Rogerson, 2015), and the capital share equal to $\theta = 0.3224$ (Prescott, 2004). We then calibrate α to match average hours worked per adult in high income countries, and the subsistence consumption term \bar{c} to match average hours worked per adult in low income countries. The resulting annual subsistence consumption amounts to \$892 in 2011 US \$, which is close to the value of \$842, corresponding to 5 percent of the U.S. GDP per capita in 1956, chosen by Ohanian et al. (2008). Online Appendix Section B includes more details on the model setup, calibration, and fit.

We use the calibrated utility function to compute a simple welfare metric, building on the work of Jones and Klenow (2016). Conceptually, the welfare metric imagines giving the representative household of some country i a choice between two options: first, to work the average hours of individuals in high-income countries and to consume a fraction λ of the average consumption in these high-income countries. The second option is to stay in country i , and to work i 's average actual hours and enjoy i 's average consumption. Formally, the welfare metric in country i is the λ_i that solves

$$u(c_i, h_i) = u(\lambda_i \cdot c_{HI}, h_{HI}) \quad (5)$$

where c_{HI} and h_{HI} are the average consumption and hours of individuals in our sample of high-income countries.

Table 11 presents our average λ_i s by income tercile, normalizing them to have an average of 100 for the group of high-income countries. In row one, we consider only cross-country differences in

consumption, and ignore differences in hours worked. Countries in the bottom third of the world income distribution have 8.4 percent of the welfare level of the richest third. The middle third features 35.7 percent of the welfare of the richest third. The final column shows that the ratio of the top to bottom third is 11.9, meaning, as expected, very sizable differences in welfare coming through consumption alone.

The second row adds differences in hours worked to the differences in consumption. Average welfare in the low-income countries falls to 5.3 percent of welfare in the high-income countries, leading to a welfare ratio of 18.8 between top and bottom. Thus, considering differences in hours worked between low- and high-income countries in addition to differences in consumption increases measured welfare differences by almost 60 percent. Measured welfare differences between middle- and high-income countries increase by a smaller degree, but still by a sizable 21 percent. Measuring welfare differences across countries is not an exact science, and our calculations leave out a lot of elements of reality that certainly matter for welfare, such as life expectancy and inequality, as emphasized by Jones and Klenow (2016). Our calculations make however clear that, all else equal, including cross-country differences in hours worked leads to substantially larger welfare differences across poor and rich countries than when ignoring differences in hours.¹⁵ This is a new point relative to Jones and Klenow (2016), whose hours data are largely from richer countries. Their main point regarding hours is that welfare differences between Western European countries and the United States are smaller than GDP differences, since Western Europeans enjoy more leisure on average.

8 Conclusion

In this paper, we document how hours vary with income across and within countries of all income levels. To do so, we compile and harmonize international survey data from countries of all income levels, focusing on a set of countries with the most scope for international comparisons. We document that, on average, adults in the developing world work about fifty percent more hours per week than adults in rich countries. Average hours worked are higher in developing countries both for men and for women, for all age and education groups, and along both the extensive and intensive margins. Within countries, hours are decreasing with income on average and particularly so in the poorest countries. In the richest countries, hours worked are flat or increasing in income. One implication of our findings is that aggregate labor productivity and TFP differences across countries are larger than previously thought. Moreover, ignoring hours worked also leads to misleading conclusions about the extent of welfare differences across countries. Put simply, residents of the

¹⁵Almas (2012) also argues that welfare of the poorest countries are overstated relative to the richest, though for a different reason, which is a bias in the PPP methodology used in the PWT.

poorest countries are not only consumption poor, but leisure poor as well.

Future work is needed to reconcile the patterns we document in a model that takes within-country heterogeneity seriously. The decrease in hours with income suggests preferences in which income effects dominate substitution effects. The flattening of the negative hours-income relationship within countries with GDP per capita points to subsistence preferences as a potential driving force. Replicating the fact that the hours-income relationship actually turns positive for the richest countries suggests additionally a role for a tax-transfer system that varies systematically with development.

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.
- Aguiar, M., E. Hurst, and L. Karabarbounis (2012). Recent Developments in the Economics of Time Use. *Annual Review of Economics* 4, 373–98.
- Aguiar, M., E. Hurst, and L. Karabarbounis (2013). Time Use During the Great Recession. *American Economic Review* 103(5), 1664–1696.
- Alesina, A., E. Glaeser, and B. Sacerdote (2005). Work and Leisure in the United States and Europe: Why So Different? *NBER Macroeconomics Annual* 20, 1–64.
- Almas, I. (2012). Measurement of Real Income, Inequality and Poverty. *American Economic Review* 102(2), 1093–1117.
- Alonso-Ortiz, J. (2014). Social Security and Retirement Across the OECD. *Journal of Economic Dynamics and Control* 47(300-316).
- Bick, A., B. Brüggemann, and N. Fuchs-Schündeln (2017). Hours Worked in Europe and the US: New Data, New Answers. Working Paper, Goethe University Frankfurt.
- Bick, A. and N. Fuchs-Schündeln (2017). Taxation and Labor Supply of Married Couples Across Countries: A Macroeconomic Analysis. *Review of Economic Studies*, forthcoming.
- Boppart, T. and P. Krusell (2016). Labor Supply in the Past, Present, and Future: A Balanced-Growth Perspective. Technical Report 22215, NBER Working Paper.
- Borjas, G. J. (1980). The Relationship between Wages and Weekly Hours of Work: The Role of Division Bias. *Journal of Human Resources* 15(3), 409–423.
- Bridgman, B., G. Duernecker, and B. Herrendorf (2017). Structural Transformation, Marketization, and Household Production Around the World. Unpublished Working Paper, Arizona State University.

- Caselli, F. (2005). Accounting for Cross-Country Income Differences. In P. Aghion and S. Durlauf. (Eds.), *Handbook of Economic Growth*.
- 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.
- De Magalhães, L., D. Koh, and R. Santaaulàlia-Llopis (2017). The Costs of Consumption Smoothing: Less Schooling and Less Nutrition. Unpublished Working Paper, University of Bristol.
- Duernecker, G. and B. Herrendorf (2016). On the Allocation of Time: A Quantitative Analysis of the US and France. Unpublished Working Paper, University of Mannheim and Arizona State University.
- Erosa, A., L. Fuster, and G. Kambourov (2012). Labor Supply and Government Programs: A Cross-Country Analysis. *Journal of Monetary Economics* 59, 84–107.
- Feenstra, R. C., R. Inklaar, and M. P. Timmer (2015). The Next Generation of the Penn World Table. *American Economic Review* 105(10), 31503182.
- Francis, N. and V. A. Ramey (2009). Measures of Per Capita Hours and Their Implications for the Technology-Hours Debate. *Journal of Money, Credit and Banking* 41, 1071–1098.
- Gollin, D. (2008). Nobody’s Business But My Own: Self-Employment and Small Enterprise in Economic Development. *Journal of Monetary Economics* 55(2), 219–233.
- Gollin, D., D. Lagakos, and M. E. Waugh (2014). The Agricultural Productivity Gap. *Quarterly Journal of Economics* 129(2), 939–993.
- Hall, R. E. and C. I. Jones (1999). Why Do Some Countries Produce So Much More Output per Worker than Others? *Quarterly Journal of Economics* 114(1), 83–116.
- Hofman, A. A. (1998). *Latin American Economic Development : A Causal Analysis in Historical Perspective*. Groningen Growth and Development Centre: Groningen Growth and Development Centre Monograph series ; 3. Groningen: Groningen Growth and Development Centre.
- Hsieh, C.-T. and P. J. Klenow (2010). Development Accounting. *American Economic Journal: Macroeconomics* 2(1), 207–223.
- Huberman, M. and C. Minns (2007). The Times They Are Not Changin’: Days and Hours of Work in Old and New Worlds, 1870-2000. *Explorations in Economic History* 44(4), 538–567.
- Jones, C. I. and P. J. Klenow (2016). Beyond GDP? Welfare Across Countries and Time. *American Economic Review* 106(9), 2426–2457.
- Keane, M. and R. Rogerson (2015). Reconciling Micro and Macro Labor Supply Elasticities: A Structural Perspective. *Annual Review of Economic* 7, 89–117.
- Klenow, P. J. and A. Rodríguez-Clare (1997). The Neoclassical Revival in Growth Economics: Has it Gone Too Far? In B. S. Bernanke and J. Rotemberg (Eds.), *NBER Macroeconomics Annual 1997*. Cambridge, MA: MIT Press.

- Landes, D. S. (1999). *The Wealth and Poverty of Nations*. New York: W.W. Norton and Company.
- Lee, S., D. McCann, and J. C. Messenger (2007). *Working Time Around the World: Trends in Working Hours, Laws and Policies in a Global Comparative Perspective*. Routledge.
- Lehmann, E. and J. P. Romano (2005). *Testing Statistical Hypotheses*. Springer Texts in Statistics.
- Maddison, A. (2001). *The World Economy: A Millennial Perspective*. Development Centre Studies. Paris: OECD.
- McDaniel, C. (2011). Forces Shaping Hours Worked in the OECD, 1960-2004. *American Economic Journal: Macroeconomics* 3, 27–52.
- McGrattan, E. R. and R. Rogerson (2004). Changes in Hours Worked. *Federal Reserve Bank of Minneapolis Quarterly Review* 28(1), 14–33.
- Minnesota Population Center (2015). Integrated Public Use Microdata Series, International: Version 6.4. Minneapolis: University of Minnesota.
- Ngai, R. L. and C. A. Pissarides (2008). Trends in Hours and Economic Growth. *Review of Economic Dynamics* 11(2), 239–56.
- Ngai, R. L. and C. A. Pissarides (2011). Taxes, Social Subsidies and the Allocation of Work Time. *American Economic Journal: Macroeconomics* 3(4), 1–26.
- 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.
- Parente, S. L., R. Rogerson, and R. Wright (2000). Homework in Development Economics: Household Production and the Wealth of Nations. *Journal of Political Economy* 108(4), 680–687.
- 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. (2009). Time Spent in Home Production in the 20th Century United States: New Estimates from Old Data. *Journal of Economic History* 69(1), 1–47.
- Ramey, V. A. and N. Francis (2009). A Century of Work and Leisure. *American Economic Journal: Macroeconomics* 1(2), 189–224.
- Rendall, M. (2015). Female Market Work, Tax Regimes, and the Rise of the Service Sector. Working paper, University of Zurich.
- Restuccia, D. and G. Vandenbroucke (2014). Explaining Educational Attainment across Countries and over Time. *Review of Economic Dynamics* 17(4), 824–841.
- Rogerson, R. (2006). Understanding Differences in Hours Worked. *Review of Economic Dynamics* 9(3), 365–409.
- Wallenius, J. (2013). Social Security and Cross-Country Differences in Hours: A General Equilibrium Analysis. *Journal of Economic Dynamics and Control* 37, 2466–2482.

Tables and Figures

Table 1: Average Hours Worked Per Adult

Panel A: Means

	Country Income Group			# Countries
	Low	Middle	High	
Core Countries	28.5 (11)	21.7 (15)	19.0 (23)	(49)
Core + Partial-Year	26.7 (16)	22.5 (31)	19.4 (26)	(73)
All Countries	26.9 (18)	22.7 (34)	19.7 (28)	(80)

Panel B: Tests of Differences in Means

	Country Income Group		
	Low-High	Low-Middle	Middle-High
Core Countries	9.5 (0.00)	6.8 (0.00)	2.7 (0.01)
Core + Partial-Year	7.3 (0.00)	4.2 (0.01)	3.1 (0.00)
All Countries	7.2 (0.00)	4.2 (0.00)	3.0 (0.00)

Note: Panel A reports average weekly hours worked per adult by country income group in the core countries, the core countries plus those with partial-year surveys, and in our full set of countries. The number of countries in each group is in parentheses. Panel B reports differences in mean hours among pairs of country income groups. p-values of permutation tests are reported in parentheses.

Table 2: Employment Rates and Hours Per Worker

	Country Income Group		
	Low	Middle	High
Hours Per Adult	28.5	21.7	19.0
Employment Rate	75.3	52.7	54.5
Hours Per Worker	38.4	41.1	35.0

Note: This table reports average weekly hours worked per adult, average employment rates, and average weekly hours worked per worker in the core countries by country income group.

Table 3: Hours Spent in Production of Home Services

	Country Income Group		
	Low	Middle	High
Cooking	8.8 (4)	8.3 (6)	5.8 (9)
Childcare	5.5 (6)	6.2 (6)	2.8 (9)
Cleaning	6.0 (5)	7.2 (6)	5.8 (9)
Collecting Water	3.0 (7)	2.1 (4)	– (0)
Shopping	2.0 (5)	2.0 (6)	3.8 (9)
Total	25.3	25.8	18.2

Note: Average weekly hours for each activity are computed only over countries in which data have been collected. The number of countries is in parentheses.

Table 4: Average Hours Worked Per Adult by Sex

Sex	Country Income Group		
	Low	Middle	High
All	28.5	21.7	19.0
Women	24.4	16.6	14.4
Men	32.7	27.3	23.8

Note: This table reports average weekly hours worked per adult among the core countries by sex and country income group.

Table 5: Average Hours Worked Per Adult by Education Level

Education	Country Income Group		
	Low	Middle	High
All Ages	28.5	21.7	19.0
Ages 25+ (<i>Non-missing Educ.</i>)	33.0	25.2	20.7
Ages 25+			
Less than Secondary	31.8	18.3	12.5
Secondary Completed	37.3	28.3	23.6
More than Secondary	39.5	31.2	27.0

Note: This table reports average weekly hours worked per adult among the core countries by education and country income group. The sample is restricted to individuals aged 25 or more for whom the education status is known. For comparison, the first row shows the data for all ages including also observations with a missing education status.

Table 6: Average Hours per Worker by Sector

Sector	Country Income Group		
	Low	Middle	High
All	38.4	41.1	35.0
All (<i>Non-missing Sec.</i>)	40.0	41.1	35.3
Agriculture	36.0	38.3	39.7
Manufacturing	44.9	42.5	37.0
Services	47.7	41.8	34.7

Note: This table reports average weekly hours worked per worker among the core countries by sector of main job and country income group. The sample is restricted to individuals for whom the sector of employment is known, and excludes Switzerland, for which sectoral data are unavailable. For comparison, the first row shows the data including observations with a missing sector of employment.

Table 7: Hypothetical Hours with U.S. Composition

	Country Income Group		
	Low	Middle	High
Actual Hours per Adult	28.5	21.7	19.0
Hypothetical Hours: U.S. Age Composition	29.5	22.0	19.5
Actual Hours per Adult (<i>Ages 25+, Non-miss. Educ.</i>)	33.0	25.2	20.7
Hypothetical Hours (<i>Ages 25+</i>): U.S. Educ. Comp.	38.3	27.6	24.7
Hypothetical Hours (<i>Ages 25+</i>): U.S. Age & Educ. Comp.	34.9	24.8	22.8
Actual Hours per Worker (<i>Non-miss. Sec.</i>)	40.0	41.1	35.3
Hypothetical Hours: U.S. Sectoral Comp.	46.9	41.9	35.2

Note: This table reports hypothetical mean hours using the U.S. composition rather than the actual country-specific composition. The first row shows actual hours per adult, and the second hypothetical hours imposing the U.S. age structure (5 year age groups from 15-19 to 95+). The third row shows actual hours per adult aged 25 or older with education information, the fourth one hypothetical hours imposing the U.S. education structure, and the fifth hypothetical hours imposing the U.S. age-education structure (10 year age groups combined with three education groups). The sixth row shows actual hours per worker with sectoral information, and the seventh hypothetical hours imposing the U.S. sectoral structure.

Table 8: Elasticities of Hours to Aggregate and Individual Income

Panel A: Both Sexes				
	ln Hours	ln Hours	ln Hours	ln Hours
ln (GDP per Hour)	-0.120 (0.037)	-	-0.040 (0.060)	-
ln (Hourly Wage)	-	-0.096 (0.035)	-0.073 (0.055)	-0.087 (0.041)
Country Fixed Effects	No	No	No	Yes
R^2	0.086	0.100	0.103	0.225
Obs.	690,479	690,479	690,479	690,479
Panel B: Men				
	ln Hours	ln Hours	ln Hours	ln Hours
ln (GDP per Hour)	-0.096 (0.035)	-	0.018 (0.060)	-
ln (Hourly Wage)	-	-0.094 (0.032)	-0.105 (0.052)	-0.116 (0.039)
Country Fixed Effects	No	No	No	Yes
R^2	0.075	0.118	0.119	0.245
Obs.	405,431	405,431	405,431	405,431
Panel C: Women				
	ln Hours	ln Hours	ln Hours	ln Hours
ln (GDP per Hour)	-0.138 (0.044)	-	-0.093 (0.043)	-
ln (Hourly Wage)	-	-0.093 (0.041)	-0.040 (0.050)	-0.076 (0.037)
Country Fixed Effects	No	No	No	Yes
R^2	0.092	0.084	0.096	0.222
Obs.	285,048	285,048	285,048	285,048

Note: This table reports the coefficients from an estimation of a variant of equation 2 on a data set containing individual observations with only wages from paid employment from 46 countries. The dependent variable is the logarithm of individual hours worked per worker. The explanatory variables are the ones listed in each row, plus age and age squared. Standard errors are clustered at the country level and given in parentheses.

Table 9: Average Employment Rates by Education Level and Age

Panel A: Prime Aged (25-54)

	Country Income Group					
	Low		Middle		High	
	Men	Women	Men	Women	Men	Women
Less than Secondary	94.9	80.5	73.6	49.4	74.1	53.5
Secondary Completed	92.9	74.3	82.7	62.6	86.9	71.5
More than Secondary	95.1	82.1	88.7	79.3	92.0	83.6

Panel B: Old (55+)

	Country Income Group					
	Low		Middle		High	
	Men	Women	Men	Women	Men	Women
Less than Secondary	77.3	63.1	38.0	20.0	23.2	11.7
Secondary Completed	77.4	60.7	46.8	24.1	35.5	23.2
More than Secondary	74.3	62.8	52.2	39.6	47.1	36.9

Note: Panel A reports average employment rates of prime aged men and women (ages 25-54) among the core countries by education group and country income group. Panel B reports average employment rates of old men and women (ages 55+) among the core countries by education group and country income group. The sample is restricted to individuals for whom the education status is known, and excludes Turkey, for which education data are unavailable.

Table 10: Labor Productivity Differences Across Countries

	Country Income Group			
	Low	Middle	High	High/Low
GDP per Worker	7.0	40.2	100.0	14.3
GDP per Hour Worked	6.0	33.6	100.0	16.5

Note: Labor productivity is computed as the average labor productivity within each country income group relative to the average labor productivity of the high-income group, which is normalized to 100. Only core countries are included in the analysis.

Table 11: Consumption-Equivalent Welfare Differences Across Countries

	Country Income Group			
	Low	Middle	High	High/Low
Consumption	8.4	35.7	100	11.9
+ Hours	5.3	29.5	100	18.8

Note: Average consumption-equivalent welfare is depicted for each country income group relative to the welfare of the average high-income country, which is normalized to 100. The first row includes only consumption and ignores differences in hours. The second row includes in addition differences in hours worked. Only core countries are included in the analysis.

Figure 1: Average Hours Worked per Adult in Core Countries

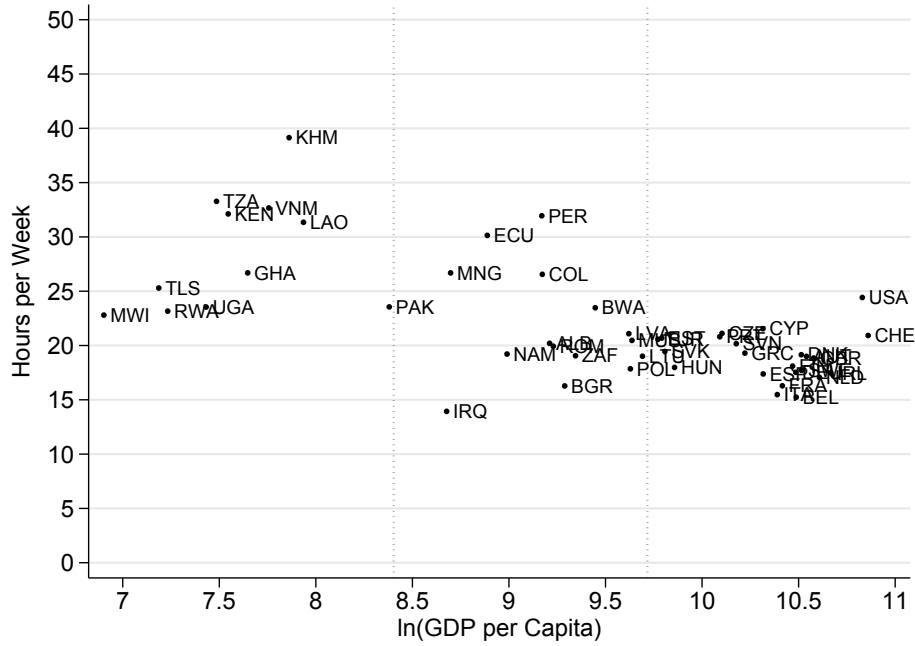
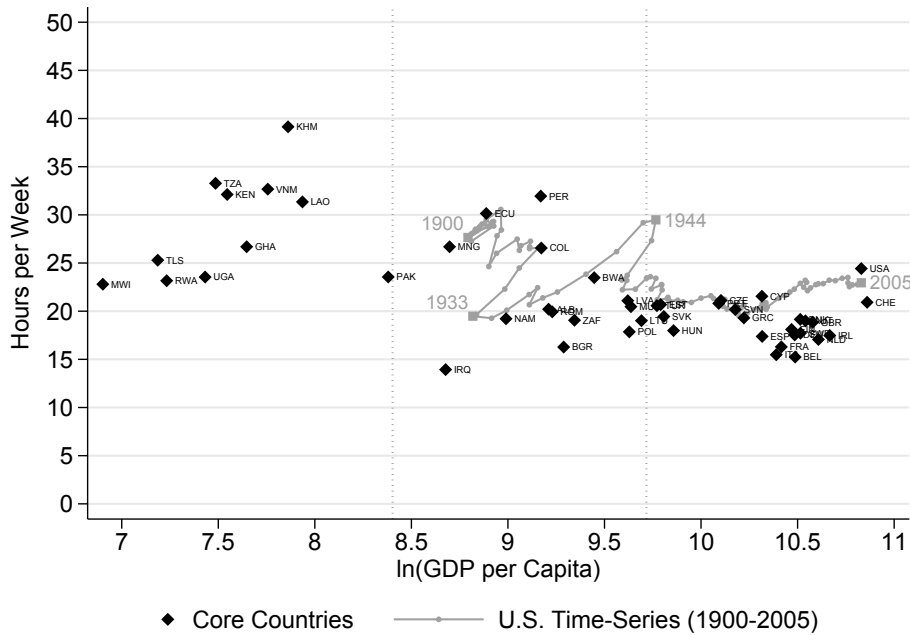


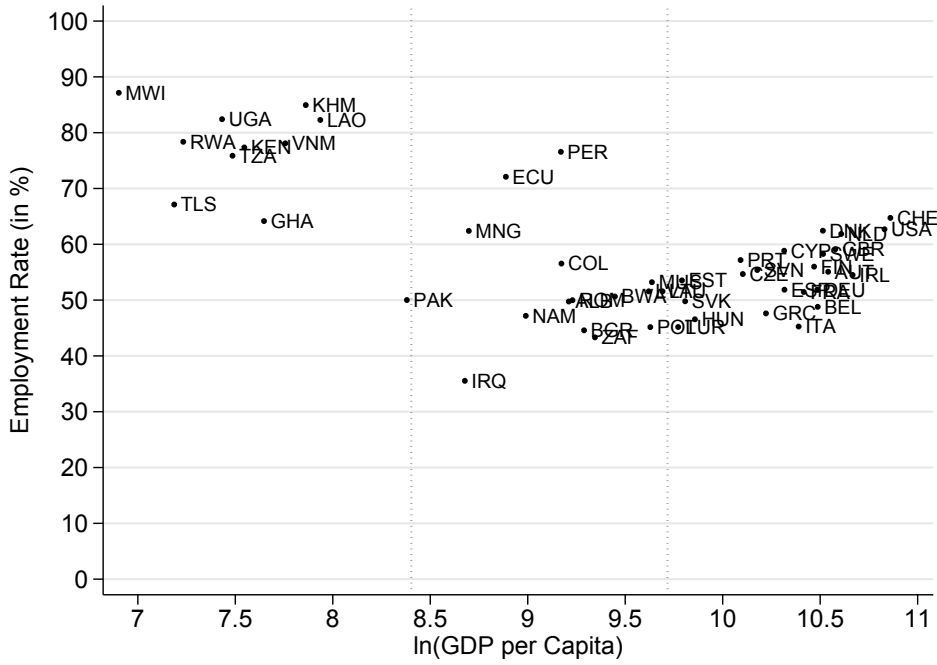
Figure 2: Average Hours per Adult – Core Countries vs. U.S. Time-Series



Note: The hours data for the U.S. Time-Series come from Ramey and Francis (2009).

Figure 3: Extensive and Intensive Margins in Core Countries

(a) Employment Rate



(b) Hours per Worker

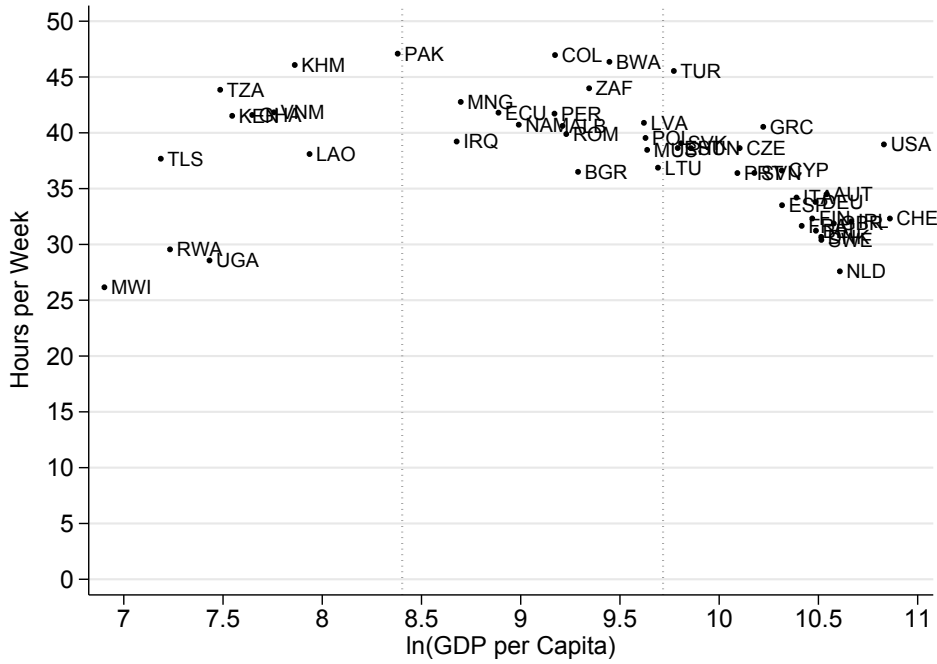


Figure 4: Average Hours per Adult over the Life Cycle

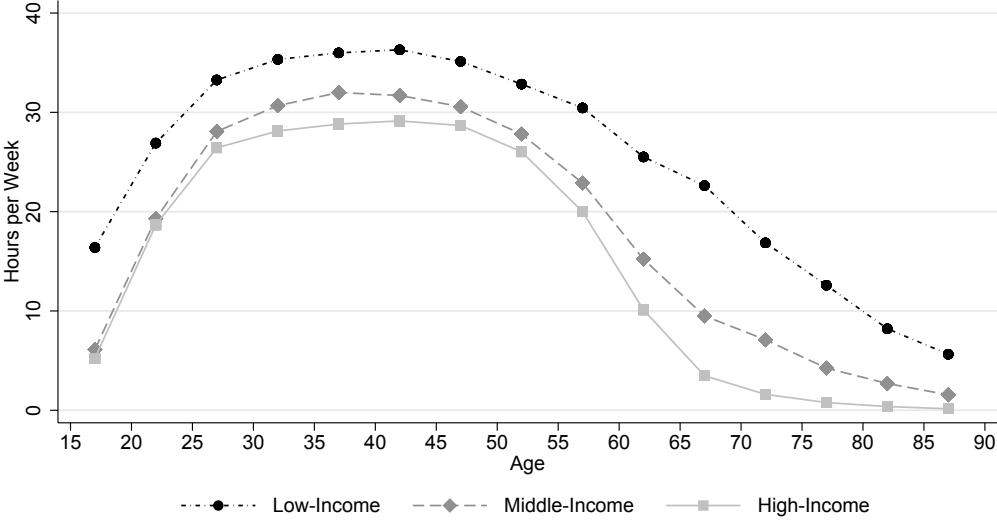


Figure 5: Hours By Wage Deciles For Paid Employees

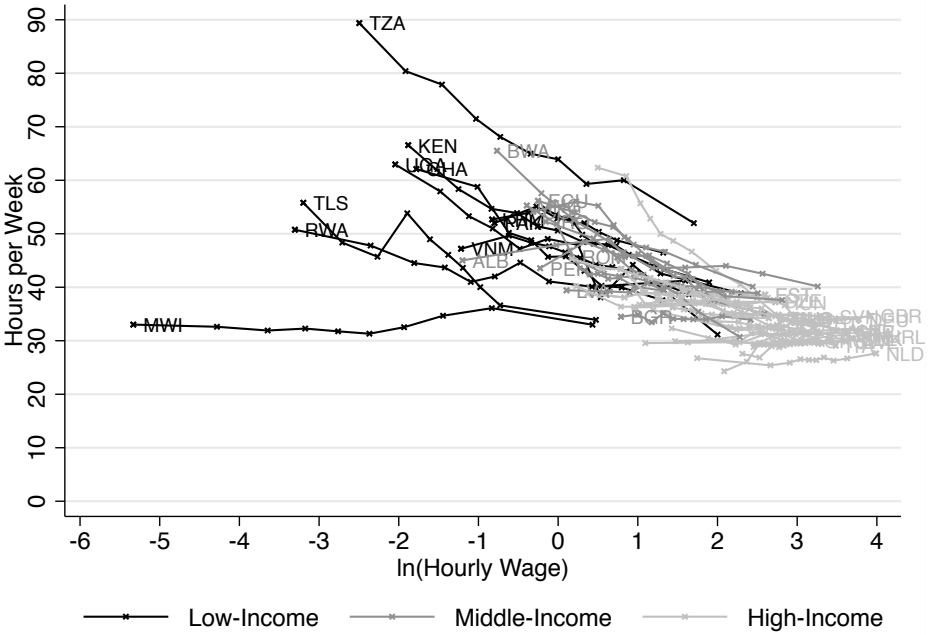
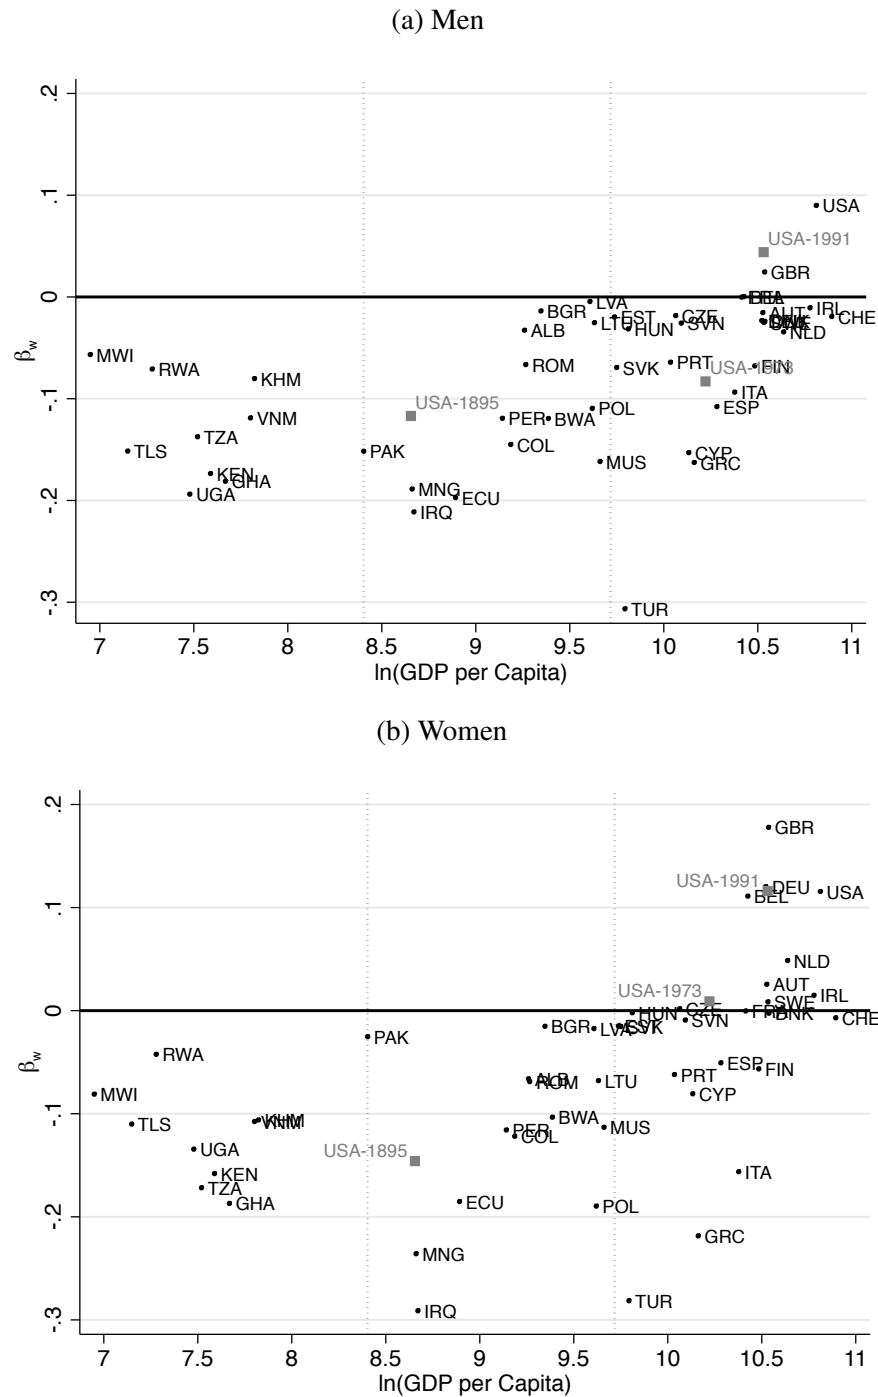


Figure 6: Country-Specific Elasticities of Hours to Wages



Note: The figures show the country-specific elasticities of hours to wages, represented by the coefficient β_w , from a regression of individual hours worked on individual wages, controlling for age and age squared. The upper panel shows results for a sample of men only, and the lower panel of women only. The gray data points are corresponding coefficients from US samples of different years (1890s, 1973, and 1991) reported in Costa (2000).