Red Nacional de Investigadores en Economía

Trade, Income and Heterogeneous Labor Supply

Nicolas Depetris-Chauvin (HES-SO)

Agustin Velasquez (IMF)

DOCUMENTO DE TRABAJO N° 261

Agosto de 2023

Citar como:
Depetris-Chauvin, Nicolás y Agustin Velasquez (2023). Trade, Income and Heterogeneous Labor Supply. *Documento de trabajo RedNIE N°261*.

Trade, Income and Heterogeneous Labor Supply

Nicolas Depetris-Chauvin
Haute École de Gestion de Genève (HES-SO)

Agustin Velasquez*
International Monetary Fund

Abstract

Workers in developing countries tend to spend more time at work than those in developed countries. This can be explained by preferences with prevalent income effects: as income rises, workers reduce their supply of labor hours to consume more leisure. However, not all workers benefit alike. In this study, we estimate the heterogeneous effects of trade, as a shifter of aggregate income, on workers' labor supply by age, education, and gender. We find that all workers benefit from more leisure caused by the income boost triggered by trade. However, young and elder workers benefit significantly more than prime-age workers. In addition, following increased trade openness women and less educated workers tend to reduce their labor supply relatively more.

Keywords: Hours worked, leisure, labor supply, international trade

JEL classification: F16, J22

*This research was supported by the Research for Development (R4D) initiative in the module "Employment in the context of sustainable development", funded by the Swiss National Science Foundation and the Swiss Development Cooperation. Contact: avelasquez@imf.org.

1. Introduction

The number of hours worked varies widely across countries. For example, the average person in Ghana works 478 more hours per year than in Germany. Averages, however, hide even larger differences among population sub-groups. For example, an average elder person in Ghana works 905 more hours per year than the average elder German (Bick et al., 2018).[1]

A recent study suggests that income differences alone can explain up to 77 percent of cross-country differences in hours worked (Bick et al., 2019).[2] As countries grow and develop, households adjust their supply of labor to rising wages depending on their preferences. On one hand, they could reduce the number of hours worked as they can sustain previous income levels working less hours (income effects). On the other hand, they might prefer to work more hours to take advantage the higher compensation for labor (substitution effects). If preferences are such that income effects outweigh substitution effects, then a rise in labor income reduces labor supply (Boppart and Krusell, 2020). In other words, households prefer to sacrifice higher consumption to enjoy more leisure.

In a recent study, Velasquez (2021) provides causal evidence that income effects dominate. He exploits exogenous variation in trade openness as a shifter of aggregate income to assess the response of aggregate labor supply. Using a cross-country panel, Velasquez shows that trade openness leads to a lower supply of labor hours. This indicates that workers in high-income countries enjoy, on average, more leisure than those in low-income countries. However, this aggregate result does not mean that leisure gains are equally distributed across workers.

In this study, we investigate the distribution of the leisure gains among the working population. We assess how the labor supply response to an increase in aggregate income varies across workers of different age, education and gender.

First, we rely on a heterogeneous worker model (as in Boppart et al. (2017)) to conceptualize differences in hours worked across individuals. The model describes

[1] Considering as elder anyone aged 65 and above.

[2] With government policies, institutions and cultural norms explaining the remaining variation.

an economy with a representative household composed by many individuals. Each of them can either stay out of the labor market, or work any number of hours subject to an upper limit, i.e. full-time employment. While all members share the same preferences over consumption and leisure, they are heterogeneous in terms of their idiosyncratic labor efficiency. This efficiency commands the relative potential earnings between household members. Finally, real wages also depend on the state of technology in the aggregate economy. Improvements in technology boost the level of *real* wages, benefiting all workers alike.

There is a household head who decides which household members should work and how many hours. We assume all members have the same utility weight and consume equal amounts, regardless of how much they work. The equilibrium indicates that individual labor supply is increasing in labor efficiency. It is optimal for more efficient workers to be employed and work more hours. Individual labor supply is a function of the state of technology and preferences. Technological improvements allow all workers to reduce their labor supply to consume leisure, but not all respond by the same amount.

We then link the unobservable labor efficiency from the model to the observable characteristics of age, education and gender. For differences across age cohorts, we assume a time-varying evolution of labor efficiency for each individual over its life-cycle. For education, we proxy efficiency with schooling levels. Regarding gender, we discuss possible heterogeneous labor supply responses subject to biased preferences.

Next, we empirically assess the labor supply response of workers of different ages, levels of education and gender to changes in aggregate income. Following the approach of Velasquez (2021), we exploit exogenous variation in trade openness as a shifter of aggregate income. When countries reduce their trade barriers and integrate into the global economy, they benefit from technological transfers and access to cheaper inputs and consumption goods (Waugh, 2010; Donaldson, 2015). This boosts wages' purchasing power and raises households' real income. Opening to trade is equivalent to lifting the state of technology in the model.[3]

[3]Trade also has heterogeneous effects across workers within a country, creating both winners and

We use a novel dataset of hours worked from Bick et al. (2018), which are collected from household surveys across countries. To the best of our knowledge, this is the only dataset which includes hourly data for low- and high-income countries where hours are decomposed by population characteristics. Unfortunately, it does not include a time dimension, which prevents controlling for country fixed effects. To alleviate endogeneity concerns, we employ an instrumental variable (IV) strategy. Following Frankel and Romer (1999), we build an instrument for trade using the geographic distance between trading partners as a proxy for trade costs. This IV strategy allows us to eliminate reverse causality between labor supply, income, and trade, and limit the scope of omitted variable biases.

Which workers gain the most leisure from the income boost triggered by trade? We document three main findings. First, trade reduces hours worked across all age cohorts. In other words, all workers benefit from more leisure. Across age cohorts, the effect of trade on hours worked has an inverted U-shape. Prime-age workers adjust their labor supply only marginally, while the young —those aged 15-25 —and the elder —65 and older—benefit most from leisure gains. For those 55+, the leisure gains are exponentially growing in age.

Second, we find that less educated workers benefit relatively more than highly educated workers, but only when they become senior. We find that workers aged 25-55 of all education levels have a similar elasticity of hours to trade. However, for those aged 55+ workers, labor supply is more responsive than for those who have less education. This gap widens as workers become older, implying that income effects have a stronger impact on less educated senior workers.

Finally, the increase in income triggered by trade benefits women more than men. We document that women's elasticity of hours to trade to be larger than men's. This evidence is apparent across all age cohorts and along the intensive and extensive margin of hours.

---

losers. This is typically driven by a country's comparative advantage in some sector and skill. However, the heterogeneous effect of trade on labor market outcomes should cancel out when accounting for countries of all income levels. We are implicitly assuming that the impact of trade on labor supply is driven by higher income and not by trade itself. In other words, we assume that the heterogeneity in the response of labor supply of workers of different population characteristics is orthogonal to the degree of trade openness.

Literature review. We contribute to the growing literature that studies the response of labor supply to changes in aggregate income. This link has been underexplored due to empirical challenges in disentangling income effects from other determinants, such as government policies and institutions. In addition, data on hours worked in developing countries is scarce (Bick et al., 2018). Recent research, however, increasingly aims at filling these gaps. Bick et al. (2019) build a model to explain differences in hours across developing and developed countries, which isolates income effects from taxation effects. Velasquez (2021) uses exogenous income variation, based on trade openness, to estimate the elasticity of hours to income and finds a positive income effect. Finally, Boppart and Krusell (2020) show that preferences with income effects outweighing substitution effects are also consistent with balanced growth.

Our research is particularly close to Boppart and Ngai (2017) and Boppart et al. (2017). Boppart and Ngai (2017) develop a growth model to track the rise in aggregate leisure in the US, and its distribution across workers of different education. Boppart et al. (2017) build a model of heterogeneous agents with adjustments of labor along the intensive and extensive margin. They focus on the link to balanced growth to answer the question: who will work in the future? In our study, we provide an empirical answer to this question focusing on the population characteristics of age, education, and gender. While Boppart and Ngai (2017) and Boppart et al. (2017) center on intertemporal labor supply adjustments, we provide a static approach. Using trade as a shifter of aggregate technology, we estimate the response of hours to a one-off permanent income boost that equally impacts all workers.

This paper is also complementary to the literature linking trade and child labor. Edmonds and Pavcnik (2005) show that trade liberalization in Vietnam caused a drop in child labor rates, despite the increase in wages for both adults and children. Similar evidence has been documented in India (Edmonds et al., 2010) and across countries (Edmonds and Pavcnik, 2006a). Our research supports these empirical findings by providing a theoretical link between trade, income and heterogeneous labor supply adjustments by age. We find that young workers (aged 15-19) significantly reduce hours worked as a response to a trade boost. We could expect income effects, and the consequent reduction in labor supply, to be even stronger among children.

Roadmap. Section 2 presents stylized facts of hours worked for different subgroups of the population. Section 3 develops the model, provides the mapping into workers' characteristics and presents the theoretical link between trade and income. Section 4 describes the empirical approach used to estimate the effect of trade on hours worked. Section 5 shows the estimation results and section 6 discusses them. Finally, section 7 concludes.

2. Differences in Hours Worked across Age, Education and Gender

We show how hours worked by age cohorts, education levels and gender vary with income. For this we compare hours devoted to market work across countries with different income levels.

Figure 1: Hours worked and income

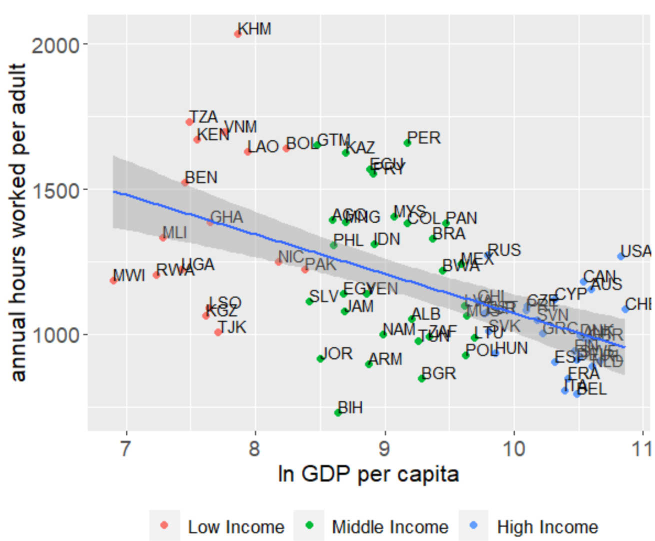

Source: Bick et al. (2018). Adults are is anyone aged 15 and older.

We begin by presenting the cross-country correlation between annual hours worked per adult and income. Figure 1 shows a clear negative correlation. Workers in low-income countries spend, on average, more hours working than those in high-income countries. This is consistent with preferences where income effects dominate substitution effects. As countries develop and workers earn higher wages, labor supply falls to boost leisure consumption.

Figure 2 decomposes the correlation of hours worked and income by age cohorts. Prime-age workers work, on average, 1500 hours per year. We observe that this number is roughly similar in both high- and low-income countries. This suggests no (or weak) correlation with income. On the other hand, a negative correlation between hours and income is noticeable for younger and older workers. Labor supply differences among the young and old explain a large share of the negative correlation between hours per capita and income described in Figure 1.

Figure 2: Hours worked and income by age

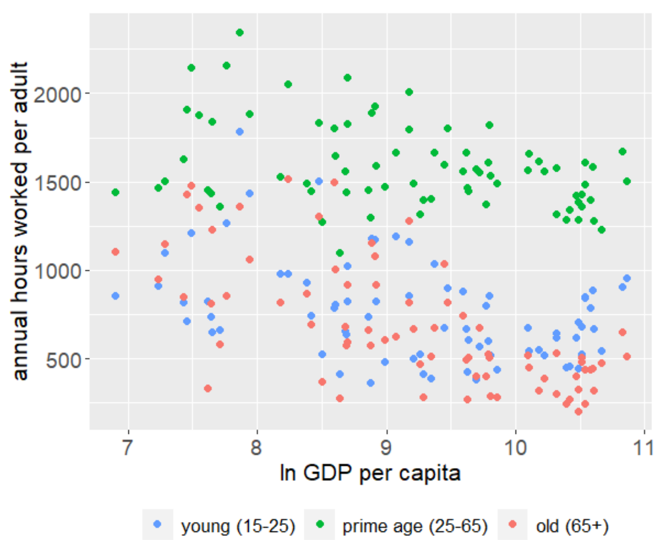

Source: Bick et al. (2018).

Figure 3, depicts the correlation of hours and income by education levels for all adults aged 25 and older. Education is split into: less than secondary, secondary completed, and more than secondary completed. The negative correlation of hours and income is observed across all education levels. Within countries, however, we observe that more educated workers tend to spend more time working than less educated workers.

Figure 3: Hours worked and income by education level

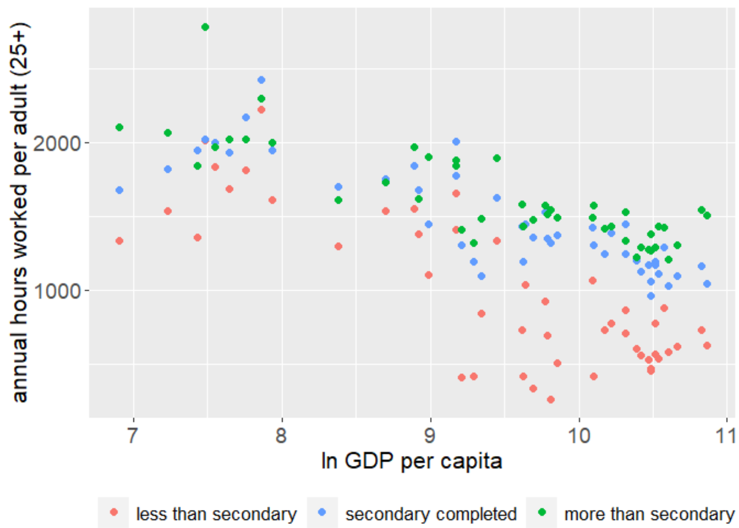

Source: Bick et al. (2018).

Figure 4 analyzes hours worked and income by gender. It presents two observations. First, men work more hours than women in all countries. Second, the negative correlation between hours worked and income is present in both genders. Both men and women work fewer hours in higher income countries. There is not a clear trend difference between the two.

3. Theoretical Background

In this section we employ Boppart et al. (2017)'s model to study the link between the labor supply of heterogeneous individuals and aggregate income. The model derives labor supply choices of individuals of different labor efficiency within a representative household. We then link the predictions of the model to the differences in labor supply across age, education and gender presented in the previous section. Finally, we display the link between trade and aggregate income.

Figure 4: Hours worked and income by gender

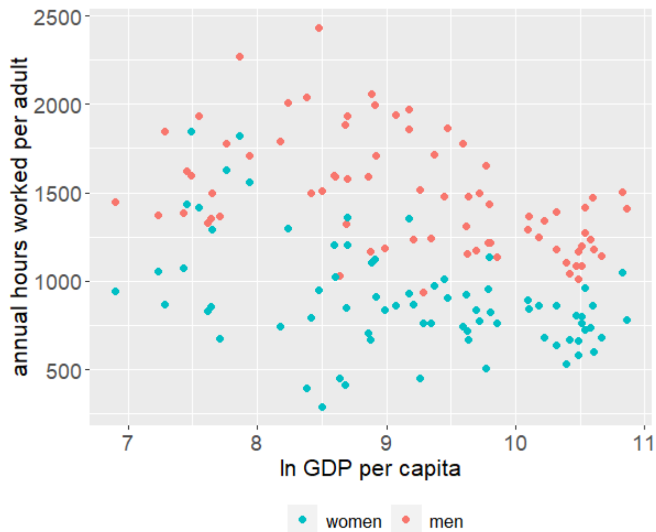

Source: Bick et al. (2018). Here adults refers to anyone aged 25 and older.

3.1 The Environment

The model describes the behavior of a representative household with multiple individuals. To avoid modelling within household bargaining, we assume there is a household head that dictates which individuals work, and how much. When taking this decision, the household head weighs the utility of each individual. For simplicity, we assume the utility of each individual has the same weight.

Each household member is endowed with a measure of labor efficiency. We can think of this efficiency measure as the hourly labor productivity irrespective of job, task or occupation. We assume that each household member is born with this idiosyncratic efficiency $\omega \geq 0$ independently withdrawn from a specific probability distribution. We define its probability density function (PDF) $f(\omega)$, and normalize the total population in the economy to unity. This implies

$$\int_{\omega} f(\omega) d\omega = 1$$

We also assume that all household members share the same preferences over consumption and leisure. These preferences take the form

$$u(c(\omega), h(\omega)) = \frac{c(\omega)^{1-\eta} - 1}{1-\eta} - \frac{h(\omega)^{1+\frac{1}{\varepsilon}}}{1+\frac{1}{\varepsilon}}$$

where $c(\omega)$ denotes the amount of consumption, and $h(w)$ the number of hours worked. The parameter $\eta > 0$ is the inverse of the inter-temporal elasticity of substitution, and $\varepsilon > 0$ the Frisch elasticity. This class of preferences, first proposed by MaCurdy (1981), allow income and substitution effects to vary depending on the magnitude of η. If $\eta > 1$ income effects outweigh substitution effects. If $\eta = 1$ the utility function is logarithmic in consumption, with income and substitution effects cancelling out. In this case, hours worked do not respond to changes in real income (as shown below).

All labor income is shared within the household and allocated in equal amounts for consumption across all individuals, regardless of the number of hours worked. That is, $c(\omega)$ is ex-post the same for all household members. Taking this into consideration, the household head decides who works and how many hours. The possible number of hours worked belongs to the discontinuous set

$$h \in \{0, [\underline{h}, 1]\} \in \mathcal{H},$$

where 0 reflects not working. If employed, there is a continuum of time, from working minimum hours $\underline{h}$ up to an upper bound of time represented by 1. A full-time worker with allocated time 1 cannot work more time, even if she wished to. In the margin, some households may be indifferent between working $\underline{h}$ and not working.

To properly represent the two margins of hours, we define $\pi(h, \omega)$ as the employment share of household members of efficiency ω that are working $h \in \mathcal{H}$ units of time.[4] By definition

$$\pi(h, \omega) \in [0, 1] \tag{1}$$

[4]We assume that at least one individual in the representative household is always employed. This allow us to avoid a corner solution. We also assume that efficiency ω is independent from the number of hours worked.

with

$$\int_{h\in\mathcal{H}} \pi(h,\omega)dh = 1 \tag{2}$$

The production side of the economy involves a continuous of firms employing labor as the sole input.[5] All final goods markets are competitive, as is the labor market. There are no labor market frictions and both firms and households have perfect information. Under these assumptions, labor is paid its marginal productivity. Each employed worker's efficiency ω is also her nominal hourly wage.

Finally, the model includes an aggregate income shifter. The parameter $T > 0$ represents the state of technology, or total factor productivity, in the aggregate economy. An increase in T boosts the productivity of all workers in all firms, raising labor income of all employees in the same magnitude. Hence, it does not modify the relative labor efficiency among household members. We can think the real hourly wage w/p of any worker with efficiency ω as $w/p = \omega T$. A country with a larger T has, *ceteris paribus*, higher real wages and higher aggregate income.

3.2 The Optimal Allocation of Individual Labor Supply

The representative household head faces the following problem

$$\max_{(c(\omega),\pi(h,\omega)\forall h\in\mathcal{H})_\omega} \left[\int_\omega \frac{c(\omega)^{1-\eta}-1}{1-\eta} - \int_{h\in\mathcal{H}} \frac{\pi(h,\omega)h(\omega)^{1+\frac{1}{\varepsilon}}}{1+\frac{1}{\varepsilon}} dh\right] f(\omega)d\omega \tag{3}$$

subject to the feasibility constraints (1) and (2), and the budget constraint

$$nT = \int_\omega c(\omega)f(\omega)d\omega, \tag{4}$$

with

$$n = \int_\omega \omega \int_{h\in\mathcal{H}} \pi(h,\omega)h(w)dhf(\omega)d\omega$$

[5]The predictions can be generalized to include capital. However, this does not modify the main takeaways of the model.

The household head evaluates ω for each member and decide whether she should work or not, and if so, the optimal number of hours. The decision is made by comparing the marginal increase in labor income to the marginal disutility of working.

The model's solution is presented in full detail in A. The equilibrium yields that individual labor supply is increasing in ω. Within the household, the most efficient individuals are the ones employed and the ones who supply the most hours. The equilibrium displays different allocation of hours worked across individuals depending of their ω and three cutoffs $\omega^1 < \omega^2 < \omega^3$. The least efficient individuals, those with $\omega < \omega^1$, stay out of the labor force and do not work. Those individuals with $\omega^1 < \omega < \omega^2$ are employed, but only work minimum hours $\underline{h}$. Then, those with $\omega^2 < \omega < \omega^3$ work some hours, but not full time. The number of hours is increasing in the magnitude of ω. Finally, those individuals with $\omega^3 < \omega$ work full time. These cutoffs, and each labor supply decision are depicted in Figure 5, assuming ω follows a Fréchet distribution.

The explicit cutoffs of the model (without assuming a specific distribution for ω) are

$$\omega^1 = \frac{\underline{h}^{1/\varepsilon}}{1+1/\varepsilon} T^{\eta-1}\varphi^{\eta} \qquad \omega^2 = \underline{h}^{1/\varepsilon} T^{\eta-1}\varphi^{\eta} \qquad \omega^3 = T^{\eta-1}\varphi^{\eta}$$

with $\varphi = \left[\int_{\omega^1}^{\omega^2} \omega \underline{h} f(\omega) d\omega + \int_{\omega^2}^{\omega^3} \omega(\mu\omega T)^{\varepsilon} f(\omega) d\omega + \int_{\omega > \omega^3} \omega f(\omega) d\omega\right]$, and μ being the Lagrange multiplier. Note that the actual aggregate labor income earned by the household is $T\varphi$.

The key feature of the model is the link between individual labor supply and T. If income effects outweigh substitution effects $\eta > 1$, the three ω cutoffs are strictly increasing in T. This means that a more advance state of technology reduces individual and aggregate labor supply while raising aggregate income. In Figure 5, an increase in T is depicted as right-ward shift of ω^1, ω^2 and ω^3. Note that if $\eta = 1$, then income and substitution effects cancel out and individual and aggregate labor supply are independent from changes in the state of technology.

Figure 5: Distribution of efficiency and labor supply decisions

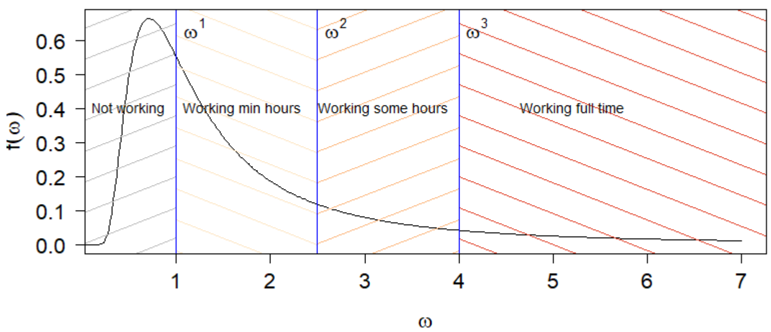

Note: Left axis depicts the PDF of a Fréchet distribution for ω.

3.3 Labor Efficiency by Age, Education and Gender

The model exploits the heterogeneity in labor efficiency to predict which individuals are employed, and which supply the most hours. While labor efficiency is not directly observable, it provides an useful guide to understand different patterns of labor supply across workers of different age, education levels and gender.

We begin by analyzing the link between labor efficiency and age. While the household lives forever, its members do not. Imagine the representative household receives a newborn. When would it be optimal for the newborn to start working? and when should she stop and retire? We can think of labor efficiency as a quadratic function of age. For example: $\omega = a + b * age - c * age^2$, where a, b, c are positive constants and $b > c$. In the first years, age increases efficiency thanks to cognitive and physical development. At some point in life a maturity is reached and ω is maximized. After this maximum is reached, more years cause cognitive skills and physical stamina to deteriorate, reducing labor's efficiency. This decline is accelerates with seniority.

This assumption of labor efficiency over the life-cycle predicts an inverse U-shaped relationship between labor supply and age. We can expect children and very old individuals to stay out of the labor force, with young and senior working part time, and prime age adults working full time. If income rises, then we expect to observe the young and old workers to reduce their labor supply, with some of them stopping work altogether. This pattern would explain the differences in hours worked in age sub-groups between high-income and low-income countries in Figure 2.

When thinking about education, we can think of it as a proxy for labor efficiency if there is a positive sorting between the two. That is, as long as education increases labor efficiency, or individuals born with high ω self-select to obtain more schooling, we can think of education as a good proxy for efficiency.[6]. Then, workers with more education are expected to work more hours relative to those with less education. This evidence is present in all countries in Figure 3. The supply of labor of highly educated workers should be relatively larger throughout the life-cycle. High ω individuals are expected to start working before low ω individuals, and are also expected to retire later. Given that acquiring education requires time, workers who obtained more education will substitute fewer hours when young to work more hours when old, e.g. PhDs. Consequently, they will tend to retire later.

The model in its simple form does not provide a clear guidance on the differences of labor supply and income by gender. The fact that men perform more market work than women in Figure 4 could be the outcome of two factors. The first one is that men have higher ω relative to women. The second factor is that the household attaches different weights to the utility of men and women. In other words, cultural norms which are ignored in the model, could lead to a slight household preference towards men supplying more labor. The fact that the ratio of women's supply of labor hours to men's is relatively constant across countries of different income suggests that the second factor may be more prevalent.

3.4 The link between Trade and Hours

Now we discuss the link between trade and labor supply. The model shows that aggregate labor supply is decreasing in the state of technology. Here we show that a

[6]The model has perfect information so everybody (firms and individuals) knows everyone's ω.

rise in trade affects aggregate labor supply by improving the state of technology T.

International trade boosts aggregate income. When countries open to trade, they are able to import cheaper consumption goods, inputs for production, access to more modern and sophisticated technology. These factors expand the production possibility frontier and raise wages' purchasing power. In this sense, opening to trade can be understood as an upward shift in labor's efficiency (a rise in T) in the model described above.

In an influential study, Arkolakis et al. (2012) show that the impact of trade on wages from a broad class of trade models can be summarized by the expression

$$\frac{w}{p} = (1 - \frac{M}{Y})^{-\frac{1}{\theta}} \tag{5}$$

where w denotes the nominal wage and p is the price index of domestically consumed goods. $\frac{M}{Y}$ is the import share of gross output, and $\theta \geq 1$ the elasticity of imports with respect to variable trade costs, more widely known as the trade elasticity. Equation (5) shows how labor income increases as a function of trade openness $\partial \frac{w}{p} / \partial M > 0$. An economy in autarky, $M = 0$, is strictly smaller than an economy with trade. As shown in Arkolakis et al. (2012), θ is a sufficient statistic that embodies all the factors shifting the production possibility frontier leading to the income gains from trade. Equation (5) shows that the effect of trade on real wages is akin to a rise in the state of technology.

Combining (5) with the cutoff values $\omega^1, \omega^2, \omega^3$ we have a mapping between trade and individual and aggregate labor supply. Trade leads to a rise in the state of technology T, which increases real wages. In turn, households react to this rise in income by reducing their labor supply to enjoy more leisure. The two key parameters influencing the elasticity of hours worked to trade are the η, which drives the income effects, and θ.

4. Empirical Framework

In this section, we present an empirical framework to assess the causal impact of aggregate income on the supply of hours of different population subgroups. With this aim, we use trade as a shifter of aggregate income.

We set the following generic specification

$$\log H_i^j = \alpha^j + \beta^j \log Imports_i + \varepsilon_i^j, \tag{6}$$

where H_i^j is the sum of the annual number of hours worked by all employed workers of characteristic j divided by total adult population j in country i. α^j is a constant and $Imports_i$ denotes the total CIF value of imports. β^j is the elasticity of hours worked to trade for population in group j and our main parameter of interest. It measures the sensitivity of the hours worked with respect to variations in total imports. ε_i^j is the error term.[7]

The parameter β^j represents the percentage point adjustment in hours worked in population j following a one percentage point increase in the volume of trade. The superscript j refers to either age cohort, level of education, gender or a combination of these characteristics. Notice that the income shock generated by trade has the same magnitude on all j groups. Therefore, our main interest is placed on how β^j varies across these groups. With income effects dominating substitution effects, we expect to see a decrease in labor supply following a income boost from trade, i.e. $\beta^j \leq 0 \quad \forall j$. In (6) we use adults as a denominator to allow adjustments in both the intensive and extensive margin of hours.

Estimating (6) with ordinary least squares (OLS) presents possible endogeneity concerns due to reverse causality and omitted variable bias. To limit both, we use an instrumental variable (IV) strategy. Following Frankel and Romer (1999), we build an instrument for trade using bilateral geographic distance between trading partners as a proxy for trade costs.

The construction of the instrument is theoretically supported by the gravity equa-

[7] As a measure of trade, we use imports because it presents less errors and omissions than exports data. Nonetheless, the same results hold when using exports on the right-hand side.

tion. It consists in two steps. First, we regress bilateral imports on the distance between trading partners, a dummy for shared border, and importer and exporter fixed effects. Second, we predict the imports for each bilateral relationship and aggregate the value of all imports for each country. We denominate $PredictedTrade$ the amount of trade that countries are expected to engage in, given their geographic fundamentals. Countries that are more isolated are predicted to trade less than countries that are more proximate to high-income countries. The details on the construction of $PredictedTrade$ are presented in B.

The first stage of the IV strategy is

$$\log Imports_i = \lambda + \gamma \log PredictedTrade_i + \epsilon_i \tag{7}$$

where λ is a constant, $\gamma > 0$ the effect of $PredictedTrade$ on imports, and ϵ_i the error term. This IV strategy eliminates reverse causality, because the measure of trade between country pairs is based on geographic distance between them. Instrumented imports are expected to have a positive impact on income, and an indirectly negative effect on hours worked.

This instrument, however, has been criticized by Rodriguez and Rodrik (2000) arguing that the effect of trade on income might be subject to omitted variable bias. If geographic distance is correlated with geographic characteristics such as weather conditions and natural endowments (oil reserves, soil fertility), then the impact of trade on income might not be driven only from trade, but through these channels as well.

As long as the determinants of labor supply between j groups other than income are not correlated with geography, i.e., $Cov(PredictedTrade_i, \varepsilon_i^j) = 0$, then our estimated parameter should not suffer from omitted variable bias. Put differently, we assume that the household's distribution of labor supply between its members is orthogonal to the geographic location and geographic characteristics of each country. This is also why we emphasize that our main interest is on how the β^js vary over the j groups rather than estimating their exact magnitude.

4.1 Data Description

We use data on hours worked from Bick et al. (2018). The authors compile data on average hours from household and labor force surveys, which includes hours reported by individuals in paid employment, self-employment and unpaid family work. Actual hours towards market activities are recorded, i.e., those towards non-market activities, such as cooking and cleaning, are not recorded. The number of hours refers to the actual hours worked in the past week in all jobs performed by adults (define by any person 15 years old or older).

Our dependent variable is annual hours worked per adult in the group j in the year 2005.[8] This variable is the product of the employment rate time hours per employed worker for adults with characteristic j. The employment rate refers to the share of all adults that report positive hours worked or being employed. The number of hours per worker is defined as the average number of hours worked in the past week in all jobs reported by those who are employed. Cross-country differences in holidays, paid leave and seasonal adjustments are accounted for in the hourly data.

Trade data are from the International Monetary Fund's Direction of Trade Statistics (DOTS).[9] Imports are measured as the CIF value of all goods and services purchased by each country to the rest of the world. For the instrument for trade, we employ geographic data from CEPII. We use kilometers weighted by population of main cities as a proxy for bilateral distance between countries.

5. Results

In this section we present the estimation results on the impact of trade and hours worked per adult and disaggregated by population characteristics.

We begin by estimating the effects of trade on aggregate labor supply. Using

[8]Bick et al. (2018)'s data report hours worked in 2005 or closest available year. 75 percent of the sample's reporting year lies between 2005 and 2009. Given the slow adjustments in labor supply over time, we use 2005 as the reference year for all countries.

[9]We use trade data from year 2005 for all countries in the sample.

specification (6), we regress hours worked per adult on total imports. Here j refers to the total adult population (aged 15+) in country i. Table 1 shows the results. Column 1 presents the elasticity of hours per adult to imports at the country level employing OLS, with Huber-White standard errors in parentheses. The estimated coefficient is -0.044 and statistically significant at the one percent level. Column 2 presents the IV results. The estimated coefficient employing the IV is -0.05, and is not significantly different from the OLS one.[10]

This result provides evidence that households have preferences with income effects dominating substitution effects. To show that trade only has an effect on hours through income, we include GDP per capita as a control. Column 3 shows the results. Once income levels are accounted for, the coefficient of instrumented imports looses its significance. The coefficient for GDP per capita has a negative and statistically significant effect on hours. Although it has the expected sign and magnitude, we do not draw causal claims. These results in line with the findings in Velasquez (2021). He employs a similar specification, but in a country-year panel and an IV that allows to control for country fixed effects. He estimates that trade leads to a -0.10 percent decline in hours per worker.

Next, we study how the decline in hours worked varies across age cohorts. In specification (6) we set j to be 5-year age cohorts. We present the IV estimation results in Figure 6. This figure shows the point estimate of the elasticity of hours per adult to trade β^j for each age group j and its respective 95 percent confidence interval. The first noticeable finding is that the response of hours to an increase in trade is negative throughout the age distribution. That is, income effects dominate substitution effects across all age cohorts, even for prime age workers. The second finding is that the elasticity of hours to trade takes an inverted U-shape across age cohorts. Hours decline the most for those aged 15-19 as well as those who are 60+. Interestingly labor supply declines by about 0.03 percent for workers aged between 20 and 54 years old. Besides presenting a small elasticity, the point estimate is marginally insignificant for those workers aged 20-29.

[10]This can be explained by the great explanatory power of our instrument. The first-stage R^2 shows that predicted trade explains up to 90 percent of the cross-country variation of imports.

Table 1: Baseline results

	Dependent variable:		
	log hours per adult		
	OLS	IV	IV
	(1)	(2)	(3)
log imports	-0.044^{***}	-0.050^{***}	-0.005
	(0.011)	(0.012)	(0.021)
log GDP per capita			-0.102^{***}
			(0.038)
Observations	75	75	75
R^2	0.159		
First Stage	*Dependent variable:*		
	log Imports		
log $PredictedTrade$	0.642^{***}		
	(0.024)		
First-Stage R^2	0.91		
F-stat	705.2		

Note: This table reports the OLS and IV estimates of regressing hours per worker on the CIF value of imports. The third column includes GDP per capita as control. The construction of the instrument log Predicted Trade is presented in B. The F-stat is the Kleiberg-Paap Wald F-statistics for weak identification. Huber-White standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Figure 6: Elasticity of hours to trade by age cohort

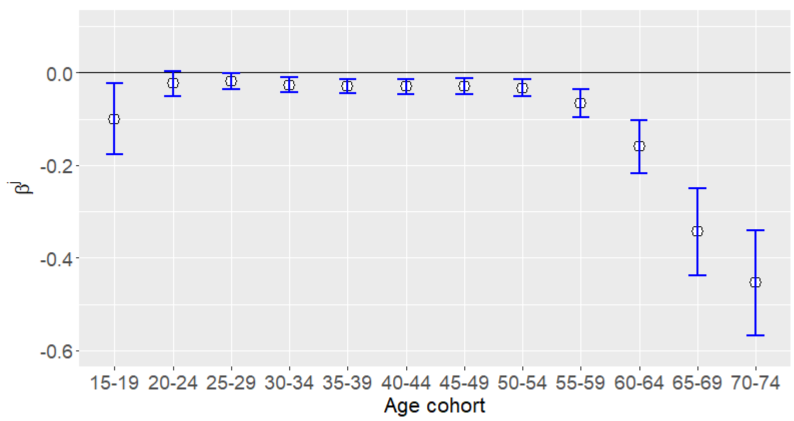

Note: Each dot represents the point estimate of the elasticity of hours to trade in each age cohort.
It is surrounded by its 95 percent confidence interval.

We expand our results by age to study if there are heterogeneous effects by level of education. We define j jointly by age and three education levels: less than secondary, secondary completed, more than secondary completed. Therefore, for each combination of the latter and age cohort, we estimate β^j using (6). Figure 7 presents the IV estimation results of hours worked per adult jointly by age and education level. For clarity, we omit showing the confidence intervals but all estimates are significantly different from zero at the 5 percent level. The elasticity of hours to trade is stable between ages 25 to 54 at values close to -0.06 across all levels of education. From age 55 and above, however, we do estimate differences in β^j across education levels. Following a one percent increase in trade, hours for elderly workers decline, but they do so more for workers with less education. In other words, among elder workers, those with more education are less sensitive to the income gains from trade.

We also analyze if there are gender differences in the response of labor supply to changes in trade. For this we employ the annual number of hours worked for men and women as dependent variable. Once again, we consider all adults regardless of age and education. Table 2 presents the IV results. The estimated β^j shows that women's supply of hours to be more sensitive to changes in trade that men's.

Figure 7: Point estimates by education level

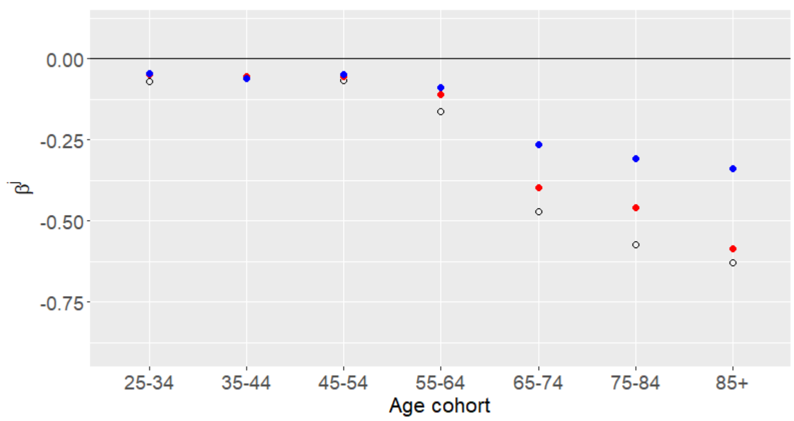

Note: Each dot represents the point estimate of the elasticity of hours to trade for the type of worker in the age cohort.

White: less than secondary, red: secondary completed, blue: more than secondary.

The point estimate for women is -0.069 compared to that of -0.037 of men's. This difference is statistically significant at the 5 percent level.

Table 2: IV results by gender

	Dependent variable:	
	log hours per adult	
	men	women
	(1)	(2)
log imports	-0.037^{***}	-0.069^{***}
	(0.013)	(0.020)
Observations	75	75
R^2	0.09	0.12

Note: This table reports the IV estimates of regressing hours per worker of men and women on the CIF value of imports. The construction of the instrument log Predicted Trade is presented in B. Huber-White standard errors are reported in parentheses. $^{*}\ p < 0.10$, $^{**}\ p < 0.05$, $^{***}\ p < 0.01$

Finally, we analyze the adjustment margins of labor supply. For this, we decompose the effects of trade on hours by age and gender by margin of hours. That is, for each population group of certain age and gender, we decompose the impact of trade on hours per worker and on the employment rate. These become our dependent variables in (6).

Table 3 presents the results. The first row shows the effect of trade on hours worked per adult. Once again, it shows that the elasticity of hours to trade is largest for women and older workers. The second row presents the results of trade on hours per worker. We find that a one percentage point increase in trade is associated with a 0.04 percent decline in the intensive margin. This is common for workers of all ages and gender, with the exception of older men. The third and last row presents the effect of trade along the extensive margin of hours. We observe a marked heterogeneity across ages. We do not find significant reductions in employment rates for young or prime age workers. However, we do find significant effects for the older working population. When income rises, these workers tend to retire and leave the labor force. In terms of differences by gender, we observe that there women are more sensitive to changes in income. This is the case along the intensive margin as well as the extensive margin.

Table 3: IV results by margin of hours, age and gender

	young (aged 15-24)			prime age (aged 25-64)			old (aged 65+)		
	all	male	female	all	male	female	all	male	female
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Dependent variable: log Hours per adult								
log imports	-0.035	-0.038*	-0.032	-0.026***	-0.019**	-0.043**	-0.168***	-0.146***	-0.204***
	(0.022)	(0.022)	(0.027)	(0.009)	(0.009)	(0.019)	(0.027)	(0.024)	(0.040)
	Dependent variable: log Hours per employed person								
log imports	-0.040***	-0.036***	-0.046***	-0.031***	-0.023***	-0.044***	-0.020*	-0.014	-0.037***
	(0.013)	(0.011)	(0.014)	(0.008)	(0.007)	(0.010)	(0.010)	(0.009)	(0.012)
	Dependent variable: log employment rate								
log imports	0.005	-0.003	0.014	0.005	0.005	0.001	-0.148***	-0.133***	-0.167***
	(0.024)	(0.022)	(0.032)	(0.009)	(0.005)	(0.020)	(0.028)	(0.023)	(0.043)

Note: The first row reports the IV estimates of regressing hours per worker by age and gender on the CIF value of imports. The second and third row repeats the analysis using hours per worker and the employment rate as dependent variable, respectively. All regressions have 75 observations. The construction of the instrument log Predicted Trade is presented in B. The F-stat is the Kleiberg-Paap Wald F-statistics for weak identification. Huber-White standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

6. Discussion of Findings

The empirical results reveal who captures the leisure gains triggered by higher aggregate income. Here we discuss how do the predictions of the model relate to the empirical findings.

Young and elder workers benefit most. Across age cohorts, we observe that following a boost in income, the supply of hours from young and elder workers declines significantly more than that of prime-age workers. Regarding the elder population, we find that the decline in hours worked is increasing with age. Overall, this response follows an inverted U shape, in line with our assumptions on the evolution of labor efficiency throughout the life-cycle. After some age threshold, labor efficiency begins to decay, and is accelerating in seniority. Therefore, when household income rises elder workers reduce their labor supply the most. When decomposing the response of hours to trade by margin of hours, it is noticeable that a large adjustment in their supply of labor is through the employment rate. Elder workers simply retire.

For young workers, a higher permanent income boost reduces the need to supply labor in the present period. From a life-cycle perspective, the household head finds it optimal for individuals to postpone working when young and less efficient until they reach their prime age. In addition, this effect is strengthened by the prospects of acquiring human capital. A permanent income boost reduces the net present value of education from a household's intertemporal perspective.[11] Unlike elder workers, young workers adjust most of their hours worked through the intensive margin and not the extensive margin. Prime-age workers only have modest leisure gains compared to young and elder workers.

Less educated workers work fewer hours and retire earlier. Our results show that, for workers in their prime age, the elasticity of hours to trade does not vary with education. Once they reach the 60 years-old threshold, the elasticity of hours to trade becomes relatively larger for those with less education. In line with the

[11] Assuming a fixed cost of schooling and wage education premium. Our results are in line with Edmonds and Pavcnik (2005, 2006b) who find that a rise in parents and children wages through trade reduces child labor rates.

theory, workers who are more efficient should work more hours. Therefore, it is optimal, from the household perspective, that more educated individuals work more hours in their senior years and retire later. In addition, from a life-cycle perspective, individuals with more education sacrificed more wage income when young while acquiring human capital. Given the upper bound in working time, they work more when old to compensate for these forgone earnings.

Women capture more leisure gains than men. In terms of differences by gender, we observe that the response of labor supply to trade-led income changes is larger for women than for men. This is the case across all age cohorts and along the intensive and extensive margin. One explanation of this finding is that norms and traditions cause households to attach more weight on women's utility than men's. Then higher income would bias the distribution of leisure towards women.

7. Conclusions

Given the prevalent cross-country preferences with dominating income effects, households react to higher income by reducing their labor supply and increasing leisure consumption. However, the leisure gains are not evenly distributed across household members. In this study, we use exogenous variation in income triggered by trade to study the heterogeneous response of hours worked across workers of different age, level of education and gender. We find that most of the leisure gains are allocated to the young and elder workers. In addition, women and workers with less education also benefit relatively more.

Our findings are especially relevant for workers in developing countries. In these countries, young and elder workers spend numerous hours engaged in low-wage and low-productivity jobs. Our study shows that opening to trade is equivalent to raising the productivity of these jobs. Reducing the household's need to work also provides leisure time that can be used in education, child care and other activities. In addition, international trade is often limited in developing countries, suggesting that the leisure gains from trade can be substantial.

References

Arkolakis, Costas, Arnaud Costinot, and Andrés Rodríguez-Clare, "New Trade Models, Same Old Gains?," *American Economic Review*, February 2012, *102* (1), 94–130.

Bick, Alexander, Nicola Fuchs-Schündeln, and David Lagakos, "How do hours worked vary with income? cross-country evidence and implications," *American Economic Review*, 2018, *108* (1), 170–99.

— , — , — , and Hitoshi Tsujiyama, "Why are Average Hours Worked Lower in Richer Countries?," Technical Report w26554, National Bureau of Economic Research, Cambridge, MA December 2019.

Boppart, Timo and L Rachel Ngai, "Rising Inequality and Trends in Leisure," *CEPR DP 12325*, 2017.

— and Per Krusell, "Labor Supply in the Past, Present, and Future: a Balanced-Growth Perspective," *Journal of Political Economy*, 2020, *120* (1), 118–157.

— , — , and Jonna Olsson, "Labor Supply in the Future: Who Will Work?," 2017 Meeting Papers 157, Society for Economic Dynamics 2017.

Donaldson, Dave, "The Gains from Market Integration," *Annual Review of Economics*, 2015, *7* (1), 619–647.

Edmonds, Eric V and Nina Pavcnik, "Child Labor in the Global Economy," *Journal of Economic Perspectives*, February 2005, *19* (1), 199–220.

Edmonds, Eric V. and Nina Pavcnik, "International trade and child labor: Cross-country evidence," *Journal of International Economics*, 2006, *68* (1), 115 – 140.

— and — , "Trade liberalization and the allocation of labor between households and markets in a poor country," *Journal of International Economics*, July 2006, *69* (2), 272–295.

Edmonds, Eric V, Nina Pavcnik, and Petia Topalova, "Trade Adjustment and Human Capital Investments: Evidence from Indian Tariff Reform," *American Economic Journal: Applied Economics*, October 2010, *2* (4), 42–75.

Frankel, Jeffrey A and David Romer, "Does Trade Cause Growth?," *The American Economic Review*, 1999, *89* (3), 21.

MaCurdy, Thomas, "An Empirical Model of Labor Supply in a Life-Cycle Setting," *Journal of Political Economy*, 1981, *89* (6).

Rodriguez, Francisco and Dani Rodrik, "Trade Policy and Economic Growth: A Skeptic's Guide to the Cross-National Evidence," *NBER Macroeconomics Annual*, 2000, *15*, 261 – 338. OCLC: 990306665.

Velasquez, Agustın, "The Leisure Gains from International Trade," *Working Paper*, 2021.

Waugh, Michael E, "International Trade and Income Differences," *American Economic Review*, December 2010, *100* (5), 2093–2124.

WTO, "A Practical Guide to Trade Policy Analysis," Technical Report 2013.

A. The Model's Equilibrium

Here we briefly present the solution to the model. For a more in-depth discussion see Boppart et al. (2017).

The household's maximization problem is

$$\max_{(c(\omega),\pi(h,\omega)\forall h\in\mathcal{H})_\omega}\left[\int_\omega \frac{c(\omega)^{1-\eta}-1}{1-\eta}-\int_{h\in\mathcal{H}}\frac{\pi(h,\omega)h(\omega)^{1+\frac{1}{\varepsilon}}}{1+\frac{1}{\varepsilon}}dh\right]f(\omega)d\omega \tag{8}$$

subject to

$$\pi(h,\omega)\in[0,1],\quad \int_{h\in\mathcal{H}}\pi(h,\omega)dh=1$$

and the budget constraint

$$nT=\int_\omega c(\omega)f(\omega)d\omega \tag{9}$$

with

$$n=\int_\omega \omega\int_{h\in\mathcal{H}}\pi(h,\omega)h(w)dhf(\omega)d\omega$$

The first order conditions (FOCs) are

$$\bar{c}^{-\eta}=\mu,$$

where μ denotes the Lagrange multiplier. Ex-post consumption is equalized within the household, regardless of ω. Then, $c(\omega)=\bar{c}$.

For the household members that work the FOC is

$$h(\omega)^{1/\varepsilon}=\mu\omega T$$

Next, we derive the cutoff points using this las FOC. For those individuals that work and select the maximum number of hours possible, $h(\omega)=1$, this cutoff point is $\omega^3=\frac{1}{\mu T}$.

Then, workers that are a bit less productive, will work positive hours supplying

$h(\omega) = (\mu\omega T)^{\varepsilon}$. Note how labor supply is linearly increasing in ω. Therefore, lower efficiency worker will supply less hours. In the limit, there will be a mass of individuals working minimum hours. These workers supply $h(\omega) = \underline{h}$. Which yields a cutoff of $\omega^2 = \underline{h}^{1/\varepsilon}\frac{1}{\mu T}$. Finally, some individuals who have low ω will be marginally indifferent between between working minimum hours and staying out of the labor force. They will opt to work if the utility of working $\underline{h}$ is larger than not working, $\mu\omega T\underline{h} - \frac{\underline{h}^{1+1/\varepsilon}}{1+1/\varepsilon} > 0$. This decision yields the cutoff $\omega^1 = \frac{\underline{h}^{1/\varepsilon}}{1+1/\varepsilon}\frac{1}{\mu T}$

In sum, the three cutoffs are

$$\omega^1 = \frac{\underline{h}^{1/\varepsilon}}{1+1/\varepsilon}\frac{1}{\mu T}$$

$$\omega^2 = \underline{h}^{1/\varepsilon}\frac{1}{\mu T}$$

$$\omega^3 = \frac{1}{\mu T}$$

Individuals with $\omega > \omega^3$ will work full time. Those with $\omega^3 > \omega > \omega^2$ will work $h = (\mu\omega T)^{\varepsilon}$. Those with $\omega^2 > \omega > \omega^1$ will work minimum hours $\underline{h}$ and those with $\omega < \omega^1$ will not work.

We can rewrite the budget constraint as:

$$T\left[\int_{\omega^1}^{\omega^2} \omega\underline{h} f(\omega)d\omega + \int_{\omega^2}^{\omega^3} \omega(\mu\omega T)^{\varepsilon} f(\omega)d\omega + \int_{\omega>\omega^3} \omega f(\omega)d\omega\right] = \mu^{-1/\eta}$$

using this expression, we can show that each cutoff value of ω is increasing in T

$$\omega^1 = \frac{\underline{h}^{1/\varepsilon}}{1+1/\varepsilon}T^{\eta-1}\varphi^{\eta} \qquad \omega^2 = \underline{h}^{1/\varepsilon}T^{\eta-1}\varphi^{\eta} \qquad \omega^3 = T^{\eta-1}\varphi^{\eta}$$

with $\varphi = \left[\int_{\omega^1}^{\omega^2} \omega\underline{h} f(\omega)d\omega + \int_{\omega^2}^{\omega^3} \omega(\mu\omega T)^{\varepsilon} f(\omega)d\omega + \int_{\omega>\omega^3} \omega f(\omega)d\omega\right]$. Alternatively, we could express the ω cutoffs as a function of total income, which is $T\varphi$.

The model shows that the most productive individuals will work, and will supply the largest number of hours. As T rises, all individuals will supply less labor.

The cutoff points depend of the distribution of ω.

B. Building an Geographic Instrument for Trade

Here we present the steps to build an instrument for trade using the gravity equation, as in Frankel and Romer (1999) and WTO (2013).

The gravity equation that explains bilateral trade flows takes the form

$$Imports_{ij} = GS_iM_j\psi_{ij} \tag{10}$$

where $Imports_{ij}$ refer to total value of goods and services country i purchased from country j measured in current US dollars. S_i represents importer-specific factors that denote importer's total demand (e.g. importing country's GDP), and M_j comprises exporter-specific factors (e.g. exporter's GDP). G represent common factors that affect both exporters and importers equally. Finally, ψ_{ij} represents the inverse of the bilateral trading costs. It measures the ease of importer i to purchase goods from j including transport costs.

Given its multiplicative nature, we can express (10) as a log-linear specification

$$\log Imports_{ij} = \lambda_i + \lambda_j + \log \tau_{ij} + error_{ij}, \tag{11}$$

where the λ_i and λ_j are importer and exporter fixed effects that control for S_i and M_j. τ_{ij} represents the bilateral trade costs between both countries (the inverse of ψ_{ij}). We assume that τ_{ij} takes the form

$$\tau_{ij} = d_{ij}^{\vartheta_1} exp\{\vartheta_2 border_{ij}\}$$

with d_{ij} representing the geographic distance between i and j, and a dummy that equals one if the trading partners share a border.

We present the OLS estimation results of (11) in Table A1. All the estimated parameters are in line with those found in the literature. The amount of imports declines with distance, a one percentage point increase in bilateral distance is associated with a 1.9 percent decline in imports. Countries that share a border also tend to trade more. As is usual in the literature, these few variables explain a high percentage of the bilateral variations in trade (R^2 is 0.73).

To build a predicted value of trade, we generate fitted values for imports for each country pair. Under the assumption of balanced trade, the prediction of total imports should be equal to the prediction of total exports for any country. We simply refer to this amount as *trade*. The predicted trade value of country i is the result of aggregating its predicted fundamentals over all bilateral trading partners

$$PredictedTrade_i = \sum_{i \neq j} e^{\hat{\lambda}_i + \hat{\lambda}_j + \hat{\vartheta}_1 \log dist_{ij} + \hat{\vartheta}_2 border_{ij}}$$

$$PredictedTrade_i = e^{\hat{\lambda}_i} \sum_{i \neq j} e^{\hat{\lambda}_j + \hat{\vartheta}_1 \log dist_{ij} + \hat{\vartheta}_2 border_{ij}} \tag{12}$$

Equation (12) defines the predicted trade variable that is used as an instrument for actual imports in our empirical exercise.

Table A1: Gravity estimation results

	(1)
	log imports
log distance	-1.881***
	(0.0231)
shared border	0.849***
	(0.110)
constant	6.615***
	(0.951)
Importer FE	yes
Exporter FE	yes
Observations	22967
R^2	0.725

Note: This table reports the OLS estimates of regressing imports on bilateral distance and a dummy of shared border. It includes importer and exporter fixed effects. Huber-White standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$