Why Doesn't Technology Flow to Poor Countries? An Analysis of Multinational Wage Patterns*†

Jonas Hjort
Columbia University
& CEPR & NBER

Hannes Malmberg
University of Minnesota

Todd Schoellman
Federal Reserve Bank
of Minneapolis

March 29, 2019

PRELIMINARY AND INCOMPLETE: DO NOT CITE

Abstract

In this paper, we ask: why don't more firms with frontier technologies use them to operate in poor countries, where input costs are presumably lower. We draw on a database documenting the salaries paid by multinational firms by occupation in cities from 149 different countries. We merge this database with nationally representative microdata on earnings by country. We document that the multinationals in our sample pay very high salaries, on the order of 30 times the average salaries in the poorest countries, implying a much smaller cost advantage than suggested by average wages. Accounting for the workforce in multinationals being highly skilled, there is still a residual gap of roughly a factor of 10. We interpret the multinational wage gap through the lens of a model where multinationals face high labor prices either because they require scarce and expensive (measured and unmeasured) skills or because of distortions to their input costs. Either force inhibits the reallocation of resources from the traditional sector. When we use standard elasticities from the misallocation literature, our results suggest that reducing the multinational wage costs, say through training or a reduction in distortions, can plausibly lead to substantial technology transfer, sectoral reallocation, and increases in output.

*hjort@columbia.edu, malmberg@stanford.edu, todd.schoellman@gmail.com

†The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

1 Introduction

Lucas (1990) famously asks why capital does not flow from rich to poor countries, introducing what has become known as the Lucas Paradox. In a simple neoclassical model, he shows that the low wages in poor countries should imply high returns to capital, making the lack of capital inflows paradoxical. A subsequent empirical literature confirms that net capital flows are small and points to cross-country differences in total factor productivity (TFP) as a leading candidate explanation (Alfaro *et al.* , 2008; Gourinchas & Jeanne, 2013). Likewise, a parallel literature constructs the return on capital around the world and finds small differences relative to Lucas' neoclassical benchmark (Caselli & Feyrer, 2007; Monge-Naranjo *et al.* , forthcoming). Again, cross-country differences in TFP are a leading explanation for why poor countries can simultaneously have low wages and modest returns to capital.[1]

This consensus has kicked off a search for the sources of the large cross-country differences in TFP (Prescott, 1998). A leading candidate is that poor countries do not utilize the best available technologies (Parente & Prescott, 1994; Cole *et al.* , 2016).[2] This idea has been confirmed by empirical studies showing that frontier technologies diffuse to poor countries with a substantial lag (Comin & Hobijn, 2004, 2010). However, these studies raise a *technology paradox*: if poor countries have cheap labor and capital, why don't the owners of the frontier technologies find it profitable to produce in poor countries?

In this paper we analyze this question by investigating the locational choices of multinational firms, which we assume throughout are engaged in production using frontier technologies. We analyze this question using a new dataset on the salaries paid by multinational firms. The firms in our dataset represents 621 of the largest and most technologically advanced firms in the global economy – the US public firms in the sample alone cover 32% of all the revenue and 44% of all the R&D investments in Compustat North America, which that covers the publicly traded firms on the continent. The special feature of our data is that it provides detailed occupational salary data for these firms in primarily capital cities of 149 countries, with GDP per worker levels stretching from Norway at $170,000 to Burundi at $500. These data allow us to speak directly to the implied cost savings from

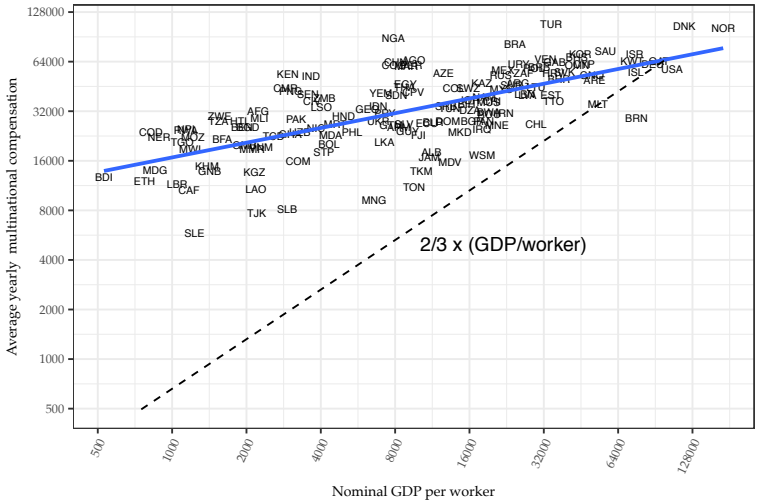
[1] This consensus also aligns with the main findings of the development accounting literature, which tends to find a small role for physical capital and a large role for human capital or TFP (Hall & Jones, 1999; Caselli, 2005; Hsieh & Klenow, 2010). Indeed, the finding of the two literatures can be linked through the dual approach to growth and development accounting (Hsieh, 2002).

[2] This is by no means the only explanation. Other leading candidates include institutions (Acemoglu *et al.* , 2005) and misallocation of resources (Hsieh & Klenow, 2009; Restuccia & Rogerson, 2008). Buera *et al.* (2011) asks a related question of why the best institutions do not spread across countries.

producing in poor countries.

Our main finding is that multinational firms in our sample pay very high salaries in poor countries. In countries with a nominal GDP per worker of 5,000 USD, the multinational firms in our sample have an average salary of 27,000 USD, that is, 5.4 times the local GDP per worker level. In the US, average multinational salaries are 55% of the GDP per worker, closer to the aggregate labor share. Thus, the multinational firms in poor countries pay about 8 times higher salaries than what would be expected given their income levels and a standard labor compensation share. The results are displayed in Figure 1. [3]

FIGURE 1: AVERAGE MULTINATIONAL WAGES AND NOMINAL GDP PER WORKER



The nominal GDP per worker is the 2010 local currency GDP per worker from the World Development Indicators, deflated by the market exchange rate. The market exchange rate is from the WDI, defined as the ratio of the local currency GDP and the GDP in current USD. The average wage is a weighted average of wage expressed in 2010 dollar equivalents, obtained by converting all local currency wages to 2010 levels using nominal income growth, and converting these wages to dollars using market exchange rates. The weighting is by the number of worker-year observations.

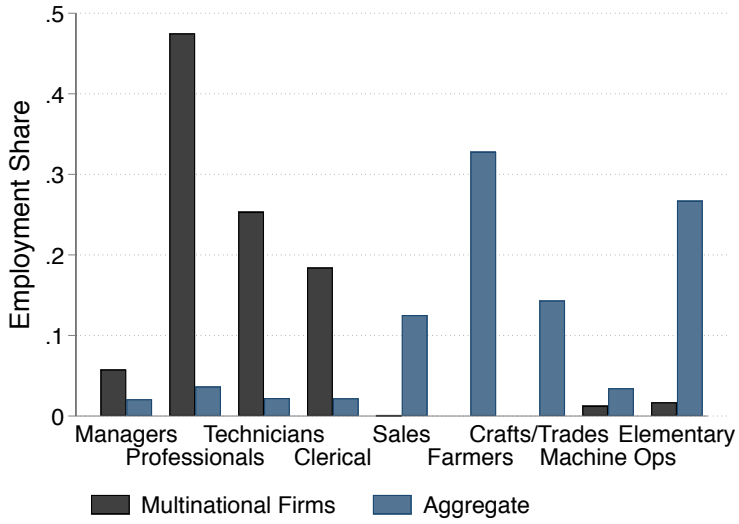
Our findings suggest a potentially simple explanation to the technology paradox. Multinational firms do not exploit low labor costs in poor countries, because, for them, there are no low labor costs to exploit. If true, this finding implies that understanding high wages in poor countries is central to resolving the technology paradox. Broadly, there are two candidate explanations for our finding. One is that multinational firms require an unusual and evidently expensive bundle of labor inputs to utilize their frontier technologies. This view is akin to the appropriate technology and technology

[3] This finding is related to existing work that shows that large firms pay higher wages in most countries, see for example Oi & Idson (1999). The fact that foreign firms pay a premium has been documented in several countries previously, although generally with a focus on production workers in manufacturing firms Aitken *et al.* (1996); Lipsey & Sjoholm (2004).

adoption work of Basu & Weil (1998). The other is that multinational firms pay a substantial premium for similarly skilled labor. In this case, this would point to the labor market frictions facing highly productive firms as an important source of misallocation in poor countries.

The rest of the paper digs further into the main finding. We first evaluate the role of the workforce composition in explaining the high multinational salaries. Our dataset contains detailed occupational data for multinationals, and we use a large set of nationally representative labor force surveys with occupation and income data to compare our firms to the broader economy. Doing this, we find that multinationals employ a much higher share of white-collar, high-skilled workers than the overall economy. The results are displayed in Figure 2. Almost all multinational workers are in skilled and semi-skilled occupations. This is not too far from the pattern for the total US labor force. In contrast, almost all workers in the broader economy in poor countries are in less skilled occupations.

FIGURE 2: OCCUPATIONAL DISTRIBUTION IN MNFS VS THE FULL ECONOMY



The set of countries are those with a GDP/worker of less than $5,000 for which we have both data on multinationals and a nationally representative labor force survey. The categories correspond to the 1-digit ISCO-08 categories.

By using the nationally representative surveys, we can also decompose the sources of the high multinational salaries. For countries with a nominal GDP of $5,000, we find that the average multinational salary is 11 times the average national salary. given the occupational distribution of the multinational firms, we should expect their workers to have 1.8 times higher salaries than the average salary, and since our firms are also located in the capital city, this factor grows to 2.4. Hence, if we focus on workers in similar occupations in the capital city, we can shrink the multinational salary

gap to $11/2.4 \approx 5$.

In sum, we find that multinational firms recruit a more skill-intensive bundle than the average firm in the economy. Furthermore, we find that this explains some, but not all, of the higher salaries paid by multinational firms.

Next, we provide a framework to understand the macroeconomic consequences of the multinational earnings gap, and that allows us to conduct counterfactuals. We assume that there are two types of firms: "modern firms" and "traditional firms", with our multinational firms being a subset of the modern firms. Modern firms have a more productive technology and use a different occupational mix than traditional firms. They might also have to pay a markup on prevailing occupational earnings, and we assume that this markup reflects either

- Unobserved heterogeneity of tasks within occupations between multinational firms and the overall economy, meaning that multinational firms have to pay more because they are using more expensive tasks. For example, multinationals might need to hire a foreigner with specialized expertise.
- Rents to workers from working in multinational firms due to fairness norms, holdup problems, and other profit sharing arrangements.

This markup acts as a wedge that misallocates resources between the modern and traditional sectors. This characterization relates to a long literature that views resources in poor countries as being misallocated between urban/non-agricultural/modern sectors and rural/agricultural/traditional sectors (Harris & Todaro, 1970; Restuccia *et al.* , 2008; Gollin *et al.* , 2014). By focusing on the wage distortion facing the modern sector where multinational firms operate, our analysis allows us to focus on the role of labor market factors in explaining the limited size of multinationals in poor countries.

We calibrate our model to replicate key aspects of the data. Our focus is on replicating the workforce composition and annual earnings paid for the the modern sector (taken from our data on multinational firms) and the traditional sector (taken from nationally representative surveys) by occupation. Our model fits these moments well by design.

Finally, we evaluate counterfactuals. As in most work on misallocation, we want to explore the effects of reducing the wedge that misallocates resources between sectors. Qualitatively, it is clear that doing so will move us closer to the efficient allocation. As the wedge decreases, modern sector wages fall, traditional sector wages increase, and there is sectoral reallocation out of the traditional

sector. Since the marginal revenue product is higher in the modern sector initially, reducing the wedge increases overall output. Quantitatively, the size of the effect depends on the elasticity of substitution between the modern and the traditional sector. A high elasticity of substitution implies extensive sectoral reallocation, a large increase in output, a large increase in the wages in the traditional sector, and only a modest decrease of the wages in the modern sector.

While this elasticity is hard to pin down, using standard values from the misallocation literature suggests that the effects can be large. For example, with an elasticity of 4, more than 75% of the workforce would work in the modern sector if we removed the residual wage gap, even if only 0.5% of the workforce started out in the modern sector. Furthermore, output would increase by a factor of almost 4. Hence, for conventional elasticity estimates, the wage gaps are large enough to be macroeconomically relevant, even if very few workers are currently in these types of firms.

The rest of the paper proceeds as follows. Section 2 describes the key features of the dataset on salaries paid by multinational firms. Section 3 uses both average and within-firm variation to show that multinational firms pay high salaries in poor countries. Section 4 decomposes the high salaries into a portion that can be accounted for by the characteristics of the multinational labor force and a residual that cannot. Section 5 provides a model of residual earnings gaps and sectoral misallocation that allows us to conduct counterfactual experiments. Finally, Section 6 offers a brief conclusion.

2 Data

The dataset we use comes from a company that gathers information on compensation at establishments around the world (the "Company"). The Company specialized in compensation in developing and emerging markets. It collects data annually from large firms with an international presence that extends into their countries of interest. Most firms in the dataset are based in North America (predominantly the United States), followed by Africa and Europe. Employers themselves choose to report data to the Company. Most are well-known to the public, and most of the private firms are publicly listed.

Each firm selects which of its establishments that will participate in the survey each year. Exactly which establishments participate varies over time so that the panel is unbalanced from the establishment$\times$year perspective. The establishments span 170 cities around the world, the vast majority of which are capital cities. On average, we observe each multinational operating establishments

in 3.7 different countries.

The parent directs the human resources personnel at each chosen establishment to report the full list of positions present in the establishment that year, how many workers are employed in each position, and their pay. Pay is reported as gross and net income in three categories: base wage, bonus, and other income. Firms report 298 different position titles, which we refer to as occupations. Company analysts conduct interviews to group these occupations into 25 occupational categories and 16 skill levels that are defined globally. Examples of low-skill occupations include cleaner, guard, and data entry clerk. Middle-skill occupations include administrative assistant, systems analyst, and finance officer, while high-skill occupations include senior legal counsel, regional office manager, and human resources director.

We have access to the Company's dataset spanning the years 2005—2015. Specifically, an observation for us is a reported wage for a firm-affiliate-occupation-year. For each observation, we know these four variables, as well as the Company's recoding of occupational information; the location (city, country) of the affiliate; the year of the survey; and the pay. We also have some information on sector and can link the data by parent name to other data on firm characteristics.

Table 1 displays summary statistics of the overall sample. In Panel A we see that the mean wage they pay firm-wide is around USD 20,000, with a standard deviation of around USD 25,000. The average employer employs workers in 28 different occupations (s.d. = $\sim$ 18) that belong to 11 different skill levels (s.d. = $\sim$ 1.5), and has 3.7 foreign establishments (s.d. = $\sim$ 11). Panel B shows that 1,100 of the 1,800 employers in the full sample are private firms; these pay somewhat higher mean wages than public employers and NGOs.

In Panel C of Table 1 we display employers' mean wage at each of the four quartiles of the skill-level distribution. The mean wage is around USD 10,000 in the lowest quartile and around USD 40,000 in the highest quartile.

For our analysis, we restrict attention to for-profit foreign affiliates of multinational firms. Hence, we remove all local firms, NGOs, educational institutions, embassies, international organizations, and public firms. We also eliminate affiliates operating in the firm's home country because these affiliates do not transfer technologies across national borders.

The remaining firms come from a wide variety of sectors, including banking, consulting, health care, mining and other natural resources, technology, telecommunications, and transport (using Standard Industrial Classification codes). In the end, this leaves us with 621 firms in 149 countries.

We construct total gross pay as the sum of gross base salary, gross bonus, and other gross income. All amounts are reported to us in contemporaneous US dollars; original data were either reported in US dollars or were converted as such using market exchange rates. We make several adjustments to make sure that these amounts can be averaged and compared across countries and years, which is complicated by the fact that some emerging markets grow rapidly and hence have rapid wage increases.

Our approach is to first convert all earnings back into local currency units using contemporaneous market exchange rates. We then adjust all amounts to year 2010 local currency units by adjusting for the average rate of nominal wage growth between year t and year 2010, inferred from the growth rate of nominal GDP per worker.[4] This adjustment makes salaries comparable over time by assuming that each occupation would have experienced the aggregate average wage growth; it misses any occupation-specific wage growth over this decade. Finally, we convert year 2010 wages in local currency to year 2010 dollars again using the market exchange rate.

This procedure utilizes market exchange rates and so may be biased by high inflation or devaluations. We exclude all country-year observations that precede a devaluation and explore excluding all country-year observations that precede an episode of high inflation in our analysis.

3 Main results

We first document the basic relationship between country income levels and the compensation in multinational firms. The main result is shown in Figure 3 which plots the average nominal GDP per capita on the horizontal axis and the average multinational wages on the y-axis.[5] The dotted line indicates 2/3 times the GDP per worker as a reference wage, since this would be the average wage if the labor share was 2/3, and if multinational firms were representative of the overall economy.

We note that, in the richest countries, average wages are quite close to the reference wage implied by a labor share of two-thirds. However, in the poor countries, wages are dramatically higher. For countries with a nominal GDP per worker of $1,000, dollar wages in multinationals are approximately $16,000, i.e., about 24 times higher than what would be expected given a labor share of two-thirds.

In the next step, we test whether this could be due to different countries having different compo-

[4] Exchange rates and GDP per worker both from World Development Indicators.

[5] For each country, the wage is the average dollar wage of all workers in the sample, and the average nominal GDP per worker is a weighted average across all years in the sample, using the total number of workers in each year as weight

TABLE 1: SUMMARY STATISTICS

Panel A: Full Sample	Mean (1)	SD (2)	Min (3)	Max (4)
Net Wage (USD)	20,242.8	24,897.2	571.6	173,029.2
Number of Occupations	28.2	18.3	2	126
Number of Skill Levels	10.7	1.4	1.0	16.0
Number of Establishments	3.6	11.2	1.0	152.0
Observations	8,334	8,334	8,334	8,334

Panel B: By Sector	Public Sector Orgs. (1)	Private Sec (2)	NGOs (3)
Net Wage (USD)	19,845.07	26,242.22	20,857.35
# Employers	148	1098	560

Panel C: Distribution of Wages	HQ-Quart1 (1)	HQ-Quart2 (2)	HQ-Quart3 (3)	HQ-Quart4 (4)
Headquarter				
Wage	10,100.53	15,647.67	23,367.12	39,720.93
Establishments				
Wage as % of HQ Wage	66.5	74.1	80.3	83.1

Note: All variables are measured at the employer-year level. There are 5,894 employer-year observations and 1,806 total employers. "Net Wage" is the average wage of all employees in a firm in a given year, and is measured in 2005 USD. "Number of Occupations" is the number occupations a firm has (across all branches) in a given year. "Skill levels" is average number of skill levels that exist in a firm. "Number of Establishments" is the number of establishments that a firm has operating in countries outside of the headquarter country. In Panel B, we separate employers into public sector organizations, private sector employers, and NGOs. All numbers in Panel B are means. In Panel C, we show the average wages at a firm's headquarter within a quartile. We then show the average wage in the firm's establishments as a percentage of headquarter wages for each quartile.

FIGURE 3: MEAN MULTINATIONAL WAGES AND NOMINAL GDP PER WORKER: FIXED EFFECTS

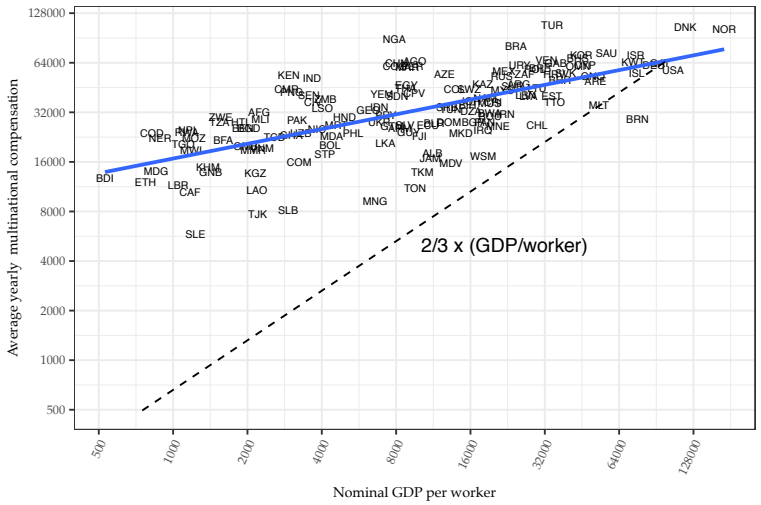
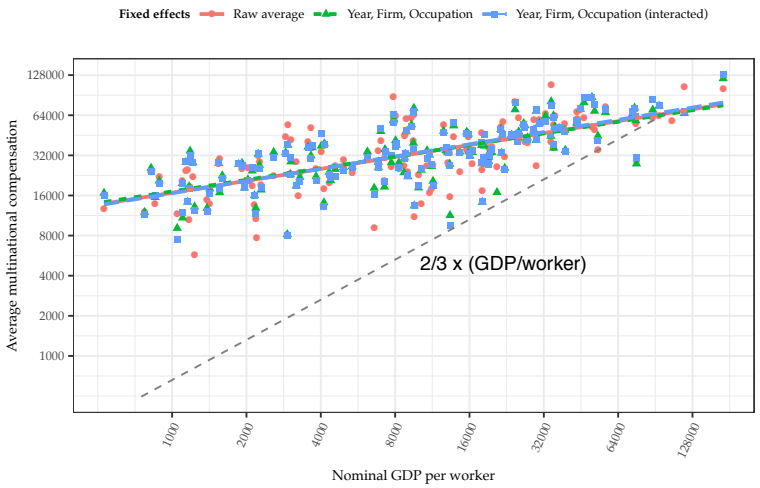


FIGURE 4: MEAN MULTINATIONAL WAGES AND NOMINAL GDP PER WORKER: FIXED EFFECTS



sition of firms, occupation, or years of observation. To this end, we run the regression

$$\begin{aligned}\log\left(\frac{w_{c,o,f,t,i}}{y_c}\right) &= \mu_c^{add} + \eta_o + \phi_f + \gamma_t + \varepsilon_{c,o,f,t,i} \\ \log\left(\frac{w_{c,o,f,t,i}}{y_c}\right) &= \mu_c^{int} + \chi_{o,f} + \tilde{\gamma}_t + \tilde{\varepsilon}_{c,o,f,t,i},\end{aligned}$$

where $w_{c,o,f,t,i}$ denotes the wage of individual i in country i, occupation o, firm f at year t, and y_c is the nominal GDP per worker in country c in 2010. all wages are converted to 2010 equivalents using the growth in output per worker. In the first regression, μ_c^{add} captures the overall wage level relative to GDP per workere when we remove occupation, firm, and time fixed effects additively. In the second regression, μ_c^{int} captures the overall wage level relative to GDP per worker when we have interacted firm-occupation fixed effects.

In Figure 4, we plot the fixed effects μ_c^{add} and μ_c^{int} against nominal GDP per worker, where the fixed effects are normalized to fit the average wage levels in the raw data. Accounting for fixed effects does very little to reduce the high wages in poor countries. If anything, the relationship steeper in the raw data than in the fixed effects.

Table 2 shows the result of the regression

$$\log w_{c,o,f,t,i} = \alpha + \beta \log y_{c,t} + \varepsilon_{c,o,f,t,i}$$

with different fixed effects added. Regardless of the specification, we obtain a parameter estimate of β between 0.25 and 0.3. This means that as nominal GDP per worker falls 10%, multinational only fall by between 2.5% and 3%. Cumulatively, this means that poor countries have much higher relative wages in multinationals than rich countries do.

-

TABLE 2: RELATIONSHIP BETWEEN GDP PER WORKER AND MULTINATIONAL WAGES

	Log full compensation			
	(1)	(2)	(3)	(4)
Log nominal GDP per worker	0.316***	0.325***	0.286***	0.286***
	(0.061)	(0.042)	(0.023)	(0.023)
Year FE?	No	Yes	Yes	Yes
Firm FE?	No	No	Yes	No
Occupation FE?	No	No	Yes	No
Firm-Occuation FE?	No	No	No	Yes
N	128,856	128,856	128,856	128,856
R^2	0.109	0.253	0.757	0.789

Notes: Full compensation is the sum of gross base pay, bonues, and "other payments", converted at the market exchange rate. The nominal GDP is the output-based real GDP times the output-based price index, both from the Penn World Table 9.0. The number of workers is the size of the engaged population in the Penn World Table 9.0. Standard errors are clustered on a country level.

4 Accounting Results

Our main finding so far is that multinational firms pay high salaries in poor countries. We dig into this finding further, with two goals. First, we assess whether such high pay in developing countries is plausible by comparing our database to other data on the pay of skilled workers in poor countries. Second, we assess how much of the high pay in developing countries can be accounted for by worker characteristics by comparing our database to nationally representative data on earnings by worker skill level in twenty countries worldwide.

Throughout, we focus on a common metric: the residual earnings gap, which we define as the ratio of average earnings paid by multinationals in country i to a benchmark of earnings in country i. Implicitly, we have already constructed one such benchmark. Using the familiar rule that average

earnings should be proportional to GDP per worker, our first measure of residual earnings is:

$$r_i^1 = \frac{w_i^{MNF}}{(1-\alpha)y_i}$$

We assume the labor share is roughly two-thirds and does not vary with development, consistent with Gollin (2002).

FIGURE 5: RESIDUAL EARNINGS GAPS AND DEVELOPMENT

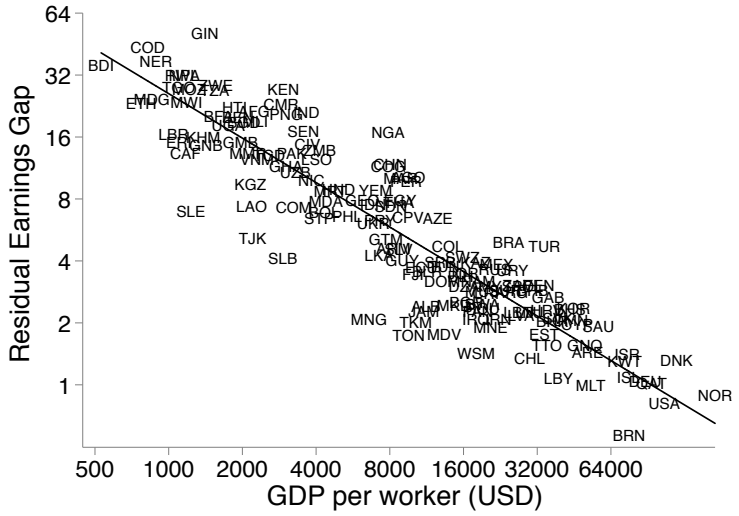


Figure 5 plots this measure of residual earnings against GDP per worker, with both axes in log scale. Standard 3-digit ISO codes are used to denote each country. Residual earnings vary substantially with development status: they are large in poor countries but approximately 1 (earnings equal to two-third of GDP per worker) in rich countries such as the United States. We introduce two statistics to help summarize this pattern. The first is the residual earnings elasticity, which is simply the slope of the line plotted in Figure 5 (since both axes are in log-scale). If multinational firms pay local average wages in all countries (e.g., two-thirds of GDP per worker), the elasticity would be 0. If multinational firms pay a constant USD wage regardless of countries, the elasticity would be -1. For r^1 the elasticity is -0.72, indicating that the raw data are much closer to the equal pay benchmark. A second statistic is simply the predicted residual earnings gap for a hypothetical country at \$5,000 GDP per worker in 2010 dollars, which is a factor of 8.2 in this case.

Each metric suggests large gaps in poor countries. The question for the remainder of this sec-

tion is whether these earnings gaps are plausible and whether they are accounted for by observable characteristics of workers. Re:

4.1 Benchmark Earnings from Recruiting Firms

Our first goal is to compare the Company's reported earnings to those from alternative sources. For this, we turn to data provided by Robert Walters, a self-described "global, specialist professional recruitment consultancy."[6] Robert Walters focuses on skilled disciplines and works with both sides of the labor market: it helps companies recruit workers and workers match with appropriate firms. It is particularly useful for us because it helps match high skilled workers to firms in developing countries.

Robert Walters uses its business expertise to produce an annual Salary Survey, which provides a typical earnings range for the current and previous year by country or region and role (occupation). Because their expertise is in recruiting and placing highly skilled workers, the report differs from our benchmark data in four ways. First, it includes only data for skilled workers. Second, the amount of detail varies substantially by country or region, reflecting the number and type of positions that they have helped to fill in each area. Third, the data reflect placements of new workers including expatriates, rather than payments to all local workers. Finally, they report salaries exclusive of bonuses and other benefits.

Robert Walters operates in countries and regions throughout the world. We focus on their data for Africa, which contains most of the poorest countries in the Company's sample. The geographic detail in the Salary Survey increases over time; we collect data from the 2017 survey, which was the first to decompose Africa into South Africa plus four geographic regions: North Africa, East Africa, West Africa, and Central-South Africa (Robert Walters, 2017). We focus for now on the four regions. The Salary Survey includes earnings for 65 roles spread across these four regions along with a salary range for each.

We match this data with our database in four steps. First, we associate regions with countries. We suspect that not all countries are actually represented in the Salary Survey, so we read the last four years of reports and marked every country or city mentioned specifically by name and associate the regions to these countries.[7] Second, we replace each salary range with the midpoint. Third, we

[6] https://www.robertwalters.com/news/about-us.html.

[7] This gives us the following mapping: North Africa includes Morocco, Egypt, Algeria, Libya, and Tunisia. East Africa includes Kenya, Uganda, Mauritius, Mozambique, Madagascar, Ethiopia, Rwanda, and Tanzania. West Africa includes Nigeria, Cote d'Ivoire, Ghana, and Senegal. Central-South Africa includes the Democratic Republic of the Congo, Angola, and

convert salaries to year 2010 US dollars using the same algorithm as the rest of the paper. Finally, we match the listed roles from the Salary Survey with occupations from our database using several instances of actual crosswalks provided by the Company.

We merge this survey with the Company's data on pay at multinational firms. The matched dataset covers more than 12,000 incumbent workers in 19 countries. Each matched observation includes both the reported salary from the Company and the matched salary from newly hired workers reported in the Robert Walters Salary Survey. We use these observations to construct a second measure of residual earnings, which is earnings in our data compared to the Robert Walters (RW) survey,

$$r_i^2 = \frac{w_i^{MNF}}{w_i^{RW}}.$$

We find that, by this benchmark, there is no residual earnings gap. Actually, workers in our survey are paid only 78 percent of what we would be suggested by the Salary Survey. This lower pay may be attributed in part to the fact that the Salary Survey includes expatriates, who are presumably paid more conditional on occupation, or by the fact that it excludes bonuses. Figure 6 plots the residual earnings gap by country and shows that there is no strong trend. These findings are at a strong contrast with the high level and strong trend in residual earnings gaps displayed in Figure 5. They validate the high reported pay in the poorest countries in our sample. Both recruiting consultancies and firms themselves report that it is extremely expensive to hire skilled labor in poor countries. Next, we ask how much of this fact can be accounted for by the observed characteristics of workers.

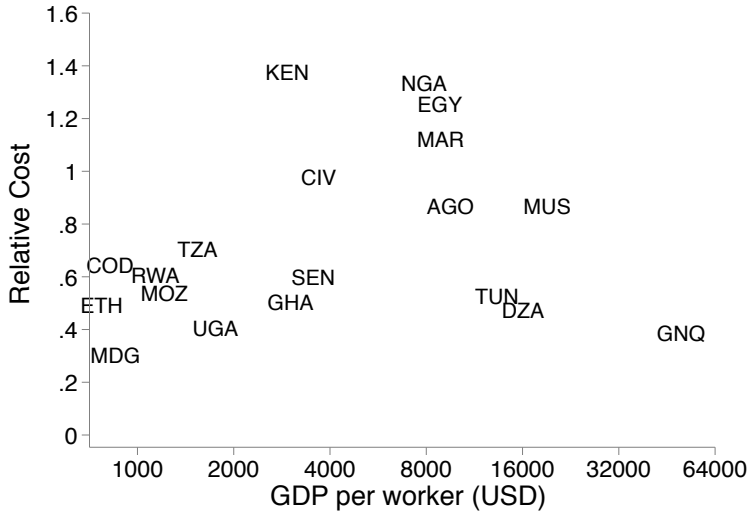
4.2 Benchmark Earnings from Microdata

We construct a dataset combining censuses and labor force surveys that contain both earnings and occupations for as many countries as possible. Census data come mostly from Minnesota Population Center (2014). Labor force surveys are collected from governments around the world and include those used in Donovan *et al.* (2019), plus several others harmonized in a similar fashion. The advantage of these datasets is that they give nationally representative data on earnings paid by occupation.

We have so far identified 20 countries for which we can access representative microdata with workers' occupations and wages. In all countries we focus on employed wage workers aged 16–70. We construct annual earnings in local currency units. As a first step, we explore using the average

Equatorial Guinea. Other African countries in the Company's data are assumed to be uncovered by the Robert Walters data.

FIGURE 6: RESIDUAL EARNINGS GAPS: COMPARISON WITH RECRUITMENT DATA



earnings of wage workers as an alternative benchmark to two-thirds of GDP per worker. To do so, we construct the average year-2010 equivalent earnings following the procedure outlined in Section 2. We then merge these year-2010 equivalent earnings onto our database of the earnings at multinational firms.

We construct our third measure of residual earnings,

$$r_i^3 = \frac{w_i^{MNF}}{\bar{w}_i},$$

the ratio of earnings paid by multinationals to average earnings in the same country, adjusted for currency and year. Figure 7 plots r_i^2 against GDP per worker in black. The resulting residual earnings gaps are noisier than those constructed using GDP per worker, but the overall pattern is similar. We report the two accounting metrics discussed above in column (3) of Table 3. The same figures for the case where we use two-thirds of GDP per worker as our benchmark is repeated in column (1). Using average earnings reduces the residual earnings elasticity but increases the residual earnings gap, showing no clear pattern on the overall magnitude of the results. Overall, we conclude that using microdata on earnings does not substantially change our results.

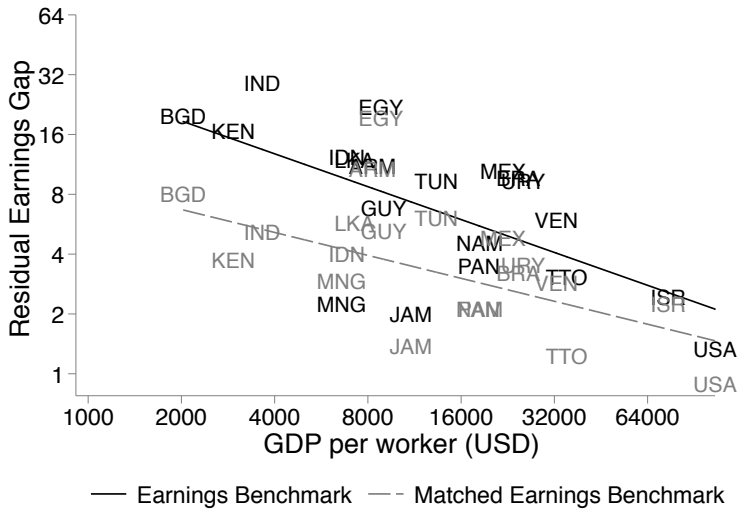
We then study earnings by occupation. We convert all occupations in both datasets to 2-digit ISCO-08 occupational codes. Most censuses and labor force surveys use modified versions of ISCO-

TABLE 3: ACCOUNTING FOR COMPOSITION EFFECTS

	Benchmark			
	GDP p.w. (1)	Earnings (3)	+Occupation (4)	+Capital (5)
Metric				
Residual Earnings Elasticity	-0.72	-0.55	-0.43	-0.38
Residual Earnings Gap at $5,000 p.w.	8.2	11.3	6.2	4.7

88 or ISCO-08 codes.[8] Our multinational wage database uses a non-traditional system that codes each worker into an occupation type and level, so that we know for example that a worker has a high level in finance. We use the combination of the two to convert these codes to ISCO-08.[9]

FIGURE 7: ACCOUNTING FOR RESIDUAL EARNINGS GAPS



Our third benchmark is the average earnings of workers in the same occupation in a given country-year. We match earnings by occupation in the multinational survey to the average earnings of workers with the same occupation in the representative survey. We collapse each figure to the country level and take the ratio , which we denote r_i^4. The results are shown in column (4) of Table 3. Matching on occupation reduces the residual earnings elasticity and the residual earnings gap by a significant amount, but a sizable gap remains.

[8] We build on standard crosswalks to translate ISCO-68 or ISCO-88 to ISCO-08. In some cases we need to assign country-specific codes by hand. We also borrow some crosswalks from Donovan *et al.* (2019).

[9] Specifically, we interpret the levels as conveying information akin to the ISCO-08 distinction between senior officials (for finance, finance managers, 1211), professionals (for finance, financial professionals, group 241), and associate professionals (for finance, finance and mathematical associate professionals, group 331).

Our microdata allow us to refine the sample of workers in the benchmark sample further. All of our multinational earnings are reported from capital cities, which generally have higher earnings overall. For our fourth benchmark we construct mean earnings by occupation for workers who live in the capital city or surrounding area, constructed using the most detailed geography available in the data set. Column (5) of Table 3 shows that this reduces both the residual earning elasticity and the residual earnings gap still further. Figure 7 plots r_i^5 against GDP per worker in blue.

Both Table 3 and Figure 7 show that accounting for worker characteristics such as occupation and work location can reduce the residual earnings gap in poor countries. However, the gap is reduced by less than one-half, with a sizable residual earnings gap of a factor of 4.7 remaining after adjusting.

5 Model

Our data set shows that multinational firms pay high salaries in poor countries -- sometimes up to 30 times the typical earnings, based on either a simple GDP per worker benchmark or average income measured in select countries. We find that a a substantial residual earnings gap remains even while accounting for the characteristics of the workforce (occupation, location). Furthermore, even the gap that can be accounted for by occupation and location might be possible to affect via policy, for example by lowering education costs and migration barriers to the capital.

The goal of this section is to provide a logical framework that explains the macroeconomic consequences of the multinational earnings gap, and that allows us to conduct counterfactuals. We focus on the wage paying sector, and assume that it consists of two types of firms: "modern firms" and "traditional firms". Modern firms have a more productive technology, use a different occupational mix, and might have to pay a markup on the prevailing market wage rate for their workers. We assume that the modern firms contain our multinational firms as a subset. Our main focus is how the modern-traditional wage gap affects the size of the modern sector and aggregate output.

In our model, we think of developing countries as facing a wedge that misallocates resources between these two sectors, consistent with a long tradition in development economics that analyze misallocation between different sectors, firms, and locations in developing markets (Harris & Todaro, 1970; Hsieh & Klenow, 2009; Gollin *et al.* , 2014) Here, we quantify a particular element, which is the cost distortion faced by firms with frontier technologies.

Our counterfactual exercise will reduce the size of this wedge. In general equilibrium, this will

reduce the wages in the modern sector, increase the wages in the traditional sector, and induce sectoral reallocation out of the traditional sector. Since the marginal revenue product is higher in the modern sector initially, reducing the wedge will increase overall output.

The extent of structural transformation and the increase in output increase will depend on the elasticity of substitution between the modern and the traditional sector. A high elasticity of substitution means extensive sectoral reallocation, a large increase in output, large increases in the wages in the traditional sector, and only modest decreases of the wages in the modern sector.

5.1 Environment

Total output is given by

$$Y = \left(\int_0^1 y(i)^{\frac{\eta-1}{\eta}} di \right)^{\frac{\eta}{\eta-1}}.$$

We assume that a share $m \in (0,1)$ of firms are frontier firms, and that the set of frontier firms and the set of traditional firms are internally homogeneous. If ourpur is optimally allocated across firms within each category, this implies an aggregate production function

$$Y = \left[(1-m) y_T^{\frac{\eta-1}{\eta}} + m y_M^{\frac{\eta-1}{\eta}} \right]^{\frac{\eta}{\eta-1}}.$$

The two types of firms have the production technologies

$$y_x = A_x \min \left\{ \frac{\ell_{x,1}}{a_{x,1}}, \ldots, \frac{\ell_{x,F}}{a_{x,F}} \right\} \quad x = t, m. \tag{1}$$

A_x is a labor-augmenting technology factor, and $\ell_{x,f}$ denotes the amount of labor of type f used in firms of type x, and the Leontieff aggregator implies a linear cost aggregator for labor

$$\hat{w}_t = \sum_f a_{t,f} w_f \tag{2}$$

$$\hat{w}_m = \sum_f a_{m,f} w_f \times e^{\tau_f}, \tag{3}$$

where e^{τ_f} is an occupation-specific wage premium paid by modern firms, and w_f is the market wage for labor of type f. The wage premium could be due to labor market frictions, or unobserved heterogeneity in the tasks collected under an occupation.[10] We normalize the weights ($\sum_f a_{x,f} = 1$), in

[10]The key assumption is that reducing τreduces the gap between $\hat{w}_m$ and $\hat{w}_t$. This is not a vacuous assumption. A coun-

which case the Leontieff formuation means that average wages map to the average unit cost of of a labor bundle.

Firms are monopolistically competitive, and hence charge a constant markup $\frac{1}{\eta-1}$ on unit costs. The price of the final output is normalized to 1

$$1 = \left((1-m)p_t^{1-\eta} + mp_m^{1-\eta}\right)^{\frac{1}{1-\eta}}, \tag{4}$$

where

$$p_t = \frac{\eta}{\eta-1}\frac{\bar{w}_t}{A_t} \tag{5}$$

$$p_m = \frac{\eta}{\eta-1}\frac{\bar{w}_m}{A_m}. \tag{6}$$

There is a representative household with preferences

$$U = c - v\left(\frac{\ell_1}{L}, \ldots, \frac{\ell_f}{L}\right),$$

subject to the budget constraint

$$c \leq \sum_{i=1}^{f} \ell_f w_f + T,$$

where $T = \Pi + \sum_f (e^{\tau_f} - 1)\pi_f w_f \ell_f$ is the sum of profits and rents from working in the modern sector, where π_f is the share of labor of type f that is employed in the modern sector.

The disutility of labor v is given by

$$v\left(\frac{\ell_1}{L}, \ldots, \frac{\ell_F}{L}\right) = \left(\sum_f \bar{w}_f \bar{p}_f^{-1/\phi}\left(\frac{\ell_f}{L}\right)^{\frac{\phi+1}{\phi}}\right)^{\frac{\phi}{\phi+1}} \qquad \phi > 0,$$

and the households choose ℓ_f subject to

$$\sum_f \ell_f = L. \tag{7}$$

By constraining aggregate labor supply to L we abstract from the choice of aggregate hours, and we

terexample is if the wage gap only was driven by differences in the number of efficiency units supplied by the workers in different sectors. In that case, reducing the wedge would not change the price of their labor bundles, since their wages per efficiency unit would already have been equalized.

do the normalization $L = 1$. Household optimization implies

$$\frac{\ell_{f'}}{\ell_f} = \frac{\bar{\ell}_{f'}}{\bar{\ell}_f} \left(\frac{w_{f'}/w_f}{\bar{w}_{f'}/\bar{w}_f} \right)^{\phi}. \tag{8}$$

This formulation is a reduced form way of creating an upward-sloping supply curve of workers in different occupations. The preference formulation is a stand-in for the cost of education/training, as well as for heterogeneous abilities within the population. Note that $\bar{\ell}_f$ is a reference labor supply, and that $\bar{w}_f$ is a reference wage. One of the parameters is redundant, but for interpretation, it is useful to have both. Shifts in the reference labor supply $\bar{\ell}_f$ is a natural way to think about expansions in education, whereas a reference wage is a useful parametrization in cases when there is a perfectly elastic supply of different types ($\phi = \infty$). Labor supply choices are based on the wages in the traditional sector, since the worker takes the aggregate rents T as given in their choice of labor supply.

5.2 Equilibrium

An equilibrium is a set of wages $\{w_f\}$, a labor allocation $\{\ell_{x,f}\}$, sector-specific prices p_t, p_m, and sector outputs y_t and y_m such that

1. The prices p_t and p_m satisfy (2)-(6)
2. Sector outputs satisfy (1) and
$$\frac{y_t}{y_m} = \left(\frac{1-m}{m} \right)^{\eta} \left(\frac{p_t}{p_m} \right)^{-\eta}$$
3. The labor market clears
$$\ell_{m,f} + \ell_{t,f} = \frac{y_t a_{t,f}}{A_t} + \frac{y_m a_{m,f}}{A_m}$$
4. Household labor choices satisfy (7) and (8)

When labor supply is infinitely elastic, the traditional sector wages are fixed at $\bar{w}_t$, condition 4 drops out, and condition 3 is replaced by

$$L = \frac{y_t}{A_t} + \frac{y_m}{A_m}$$

if we normalize to $\sum_f a_{t,f} = \sum_f a_{m,f} = 1$. Furthermore, for an equilibrium to exist, we need that the references wages $\bar{w}_f$ are consistent with the final goods having the price 1. Since only relative wages determine the allocation across sectors, this can be achieved by solving the full model for a given set of wages $\bar{w}_f$, and then rescale the wages to ensure that the final good has the price 1.

5.3 Calibration

We calibrate our parameters to match a typical country with a nominal GDP per worker of \$5,000 in 2010. Since the modern sector is small for this set of countries, we assume that the traditional sector can be approximated by the wage data on the full economy.

To calibrate the relative wage of the modern and the traditional sector $\frac{\hat{w}_m}{\hat{w}_t}$, we use the relative salaries of multinational firms relative to all wage earners at a hypothetical country with \$5,000 in GDP per worker. We obtain this value from column (2) in Table 3. For the wedge τ_f, we assume that there is a common wedge τ across all occupations. We identify this wedge by finding how much more multinationals pays conditional on their occupational distribution. This corresponds to the value r^2 in Section 4, and we find it for a hypothetical \$5,000 GDP per worker country in column (3) in Table 3. We obtain $\tau = \log(6.2)$.

In our main experiment, we assume that labor supply is perfectly elastic at the the traditional sector wages $\mathbf{w}_t$. Hence, we select $\phi = \infty$ and $\bar{w}_f = (\mathbf{w}_t)_f$. We later plan to explore the consequences of making labor supply imperfectly elastic. For the elasticity of demand η, we use different values between 2 and 6 in line with standard values used in the misallocation literature (Hsieh & Klenow, 2009).

The relative sectoral productivities and the share of modern firms m are not separately identified, but jointly identified by

$$\frac{A_m \pi}{A_t(1-\pi)} = \left(\frac{m}{1-m}\right)^{\eta} \left(\frac{\mathbf{w}_m \cdot \mathbf{a}_m / A_m}{\mathbf{w}_t \cdot \mathbf{a}_t / A_t}\right)^{-\eta} \iff \frac{A_m^{\frac{\eta-1}{\eta}}}{A_t^{\frac{\eta-1}{\eta}}} \frac{m}{(1-m)} = \left(\frac{\pi}{1-\pi}\right)^{1/\eta} \frac{\mathbf{w}_m \cdot \mathbf{a}_m}{\mathbf{w_t} \cdot \mathbf{a}_t}, \quad (9)$$

where π is the share of workers in the modern sector, and we use that $y_m = A_m \pi$ and $y_t = A_t(1-\pi)$. The ratio $\frac{\mathbf{w}_m \cdot \mathbf{a}_m}{\mathbf{w_t} \cdot \mathbf{a}_t}$ is simply the relative average wages in the two sectors, which was calibrated above.

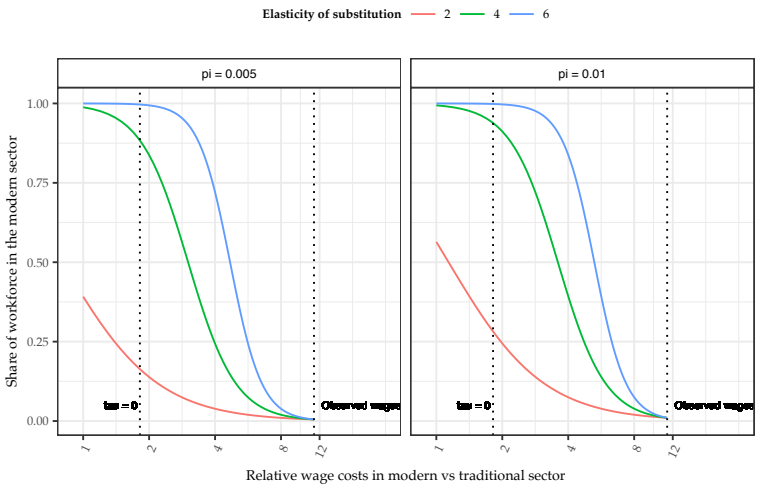
Lastly, we calibrate π, the share of workers in the modern sector. This is challenging to calibrate since our definition of a "modern" firm does not map directly to definitions used by statistical agencies. To be conservative in our counterfactuals, we assume that only 1% of workers in the capital city are in the modern sector, defined as being similar to the multinational firms in our sample. However, we also experiment with other values for π.

All calibration moments are listed in Table 4.

TABLE 4: PARAMETER VALUES

Parameter	Values
$\hat{w}_m/\hat{w}_t$	11.32
$\exp(\tau)$	6.24
$\hat{w}_m^{US}/\hat{w}_t^{US}$	1.76
ϕ	Inf
η	4.00
π	0.01

FIGURE 8: COUNTERFACTUAL MODERN SECTOR EMPLOYMENT SHARES



5.4 Experiments

We use the model to find the effect of changing the wage gap $\hat{w}_m/\hat{w}_t$ between the modern and the traditional sector. In our baseline experiments, we keep the sectoral productivities A_t and A_m constant, which means that $\hat{w}_m/\hat{w}_t$ is manipulated by changing the modern sector wedge τ and/or the reference wages $\bar{w}_f$ in the labor supply equation (8).

Even though there are multiple dimension of manipulations, they only affect the aggregate output and the size of the modern sector through the relative cost of labor in the two sectors. In particular, for any shocks to $\bar{w}_f$ and τ, the relative labor shares are given by

$$\frac{L_m}{L_t} \Big/ \frac{L_m^0}{L_t^0} = \left(\frac{\hat{w}_m/\hat{w}_t}{\hat{w}_m^0/\hat{w}_t^0} \right)^{-\eta}, \tag{10}$$

where

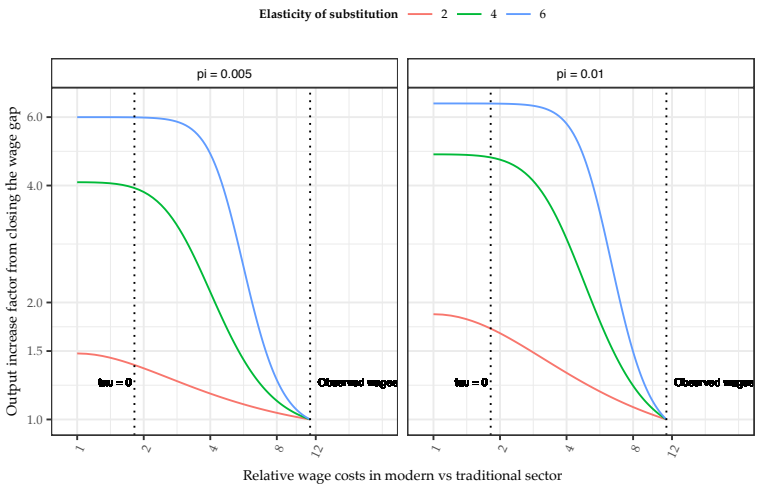
$$\hat{w}_x = \sum_f a_{x,f} w_f \exp(\tau \mathbb{I}(x = m)) \quad x = m, t.$$

If we write π_0 for the share of workers in the modern sector at the relative wage level $\hat{w}_t^0/\hat{w}_m^0$, we can use (9), and that total labor supply is normalized to 1, to express aggregate output as

$$Y = A_t m L_t \left(1 + \frac{m}{1-m} \left(\frac{A_m}{A_t} \frac{L_m}{L_t} \right)^{\frac{\eta-1}{\eta}} \right)^{\frac{\eta}{\eta-1}} = A_t m \frac{1}{1 + L_m/L_t} \left(1 + \left(\frac{\pi_0}{1-\pi_0} \right)^{1/\eta} \frac{\hat{w}_m^0}{\hat{w}_t^0} \left(\frac{L_m}{L_t} \right)^{\frac{\eta-1}{\eta}} \right)^{\frac{\eta}{\eta-1}}.$$

Using (9), output is only a function of $\hat{w}_m/\hat{w}_t$.

FIGURE 9: COUNTERFACTUAL OUTPUT LEVELS



Figures 8 and 9 show the main results. On both figures, the horizontal axes is the wage gap between the modern and traditional sector, plotted between 1 meaning no gap, and 11.3, the gap observed in the data. The vertical dashed line dashed line shows the wage gap that would prevail with $\tau = 0$. Since $\exp(\tau) = 6.2$, this value is $11.3/6.2 = 1.8$. The different colors denote different elasticities of substitution, and the different panels represent different assumptions about π, the initial share of modern employment.

Figure 8 shows how the modern sector employment share varies with the relative wages. We see that modern sector employment expands dramatically if we lower the wage gap to that between the capital sectors for high-skilled occupations versus the overall economy. If the initial employment share is $\pi = 0.01$, it expands to 25%-100% depending on the elasticity of substitution. If the initial employment share is $\pi = 0.005$, the corresponding figures are 20% and 100%.

Figure 9 shows how the aggregate output changes as a function of relative wages. When the initial share of the modern sector is $\pi = 1\%$, output increases by a factor of about 1.75 to more than 6 when we remove the residual wage gap. For $\pi = 0.005$, the corresponding figures are approximately 1.5 to 6 times.

6 Conclusion

In this paper we investigate why more firms with leading technologies do not move production to poor countries, given the apparently low input costs there. An important step was to provide results

from a new dataset on salaries paid by foreign affiliates of multinationals around the world. These salaries are high in poor countries. We show that a portion of this can be accounted for by the fact that multinationals hire observably different workers, with a strong emphasis on white-collar occupations such as managers and professionals. However, a residual earnings gap remains unexplained.

As with other literatures documenting wage or earnings gaps, the interpretation of this gap is unclear. It is possible that multinationals pay high salaries to attract workers who are different on unobservable (to us) dimensions. On the other hand, it could be that multinationals have to pay high salaries for the exact same worker to compensate them for working long hours, working in a large firm, or simply because labor markets are distorted. In all cases, the implication for multinationals remains the same: they have reduced incentives to transfer technologies to poor countries.

We investigate the quantitative importance of this finding in a dual sector model where multinational firms are part of a modern sector where leading technologies are in use, but poor countries also have a sizable traditional sector with outdated technologies. High input costs restrict the size of the modern sector. Reducing these costs induces a substantial reallocation of resources across sectors and a sizable increase in output when using conventional elasticities.

References

ACEMOGLU, DARON, JOHNSON, SIMON, & ROBINSON, JAMES A. 2005. Institutions as a fundamental cause of long-run growth. *Chap. 6, pages 385–472 of:* AGHION, PHILIPPE, & DURLAUF, STEVEN N. (eds), *Handbook of Economic Growth*, vol. 1A. Elsevier.

AITKEN, BRIAN, HARRISON, ANN, & LIPSEY, ROBERT. 1996. Wages and Foreign Ownership: A Comparative Study of Mexico, Venezuela, and the United States. *Journal of International Economics*, **40**(3-4), 345–371.

ALFARO, LAURA, KALEMLI-OZCAN, SEBNEM, & VOLOSOVYCH, VLADIMIR. 2008. Why Doesn't Capital Flow from Rich to Poor Countries? An Empirical Investigation. *Review of Economics and Statistics*, **90**(2), 347–368.

BASU, SUSANTO, & WEIL, DAVID N. 1998. Appropriate Technology and Growth. *Quarterly Journal of Economics*, **113**(4), 1025–1054.

BUERA, FRANCISCO J, KABOSKI, JOSEPH, & SHIN, YONGSEOK. 2011. Finance and development: A tale of two sectors. *American Economic Review*, **101**(5), 1964–2002.

CASELLI, FRANCESCO. 2005. Accounting for Cross-Country Income Differences. *Chap. 9, pages 679–741 of:* AGHION, PHILIPPE, & DURLAUF, STEVEN N. (eds), *Handbook of Economic Growth*, vol. 1A. Elsevier.

CASELLI, FRANCESCO, & FEYRER, JAMES. 2007. The Marginal Product of Capital. *The Quarterly Journal of Economics*, **122**(2), 535–568.

COLE, HAROLD L., GREENWOOD, JEREMY, & SANCHEZ, JUAN. 2016. Why Doesn't Technology Flow from Rich to Poor Countries? *Econometrica*, **84**(4), 1477–1521.

COMIN, DIEGO, & HOBIJN, BART. 2004. Cross-Country Technology Adoption: Making the Theories Fit the Facts. *Journal of Monetary Economics*, **51**(1), 39–83.

COMIN, DIEGO, & HOBIJN, BART. 2010. An Exploration of Technology Diffusion. *American Economic Review*, **100**(5), 2031–2059.

DONOVAN, KEVIN, LU, JIANYU, & SCHOELLMAN, TODD. 2019. *Labor Market Flows and Development*. mimeo, Yale University.

GOLLIN, DOUGLAS. 2002. Getting Income Shares Right. *Journal of Political Economy*, **110**(2), 458–474.

GOLLIN, DOUGLAS, LAGAKOS, DAVID, & WAUGH, MICHAEL E. 2014. The Agricultural Productivity Gap. *Quarterly Journal of Economics*, **129**(2), 939–993.

GOURINCHAS, PIERRE-OLIVIER, & JEANNE, OLIVIER. 2013. Capital Flows to Developing Countries: The Allocation Puzzle. *Review of Economic Studies*, **80**(4), 1484–1515.

HALL, ROBERT E., & JONES, CHARLES I. 1999. Why Do Some Countries Produce So Much More Output Per Worker Than Others? *Quarterly Journal of Economics*, **114**(1), 83–116.

HARRIS, JOHN R., & TODARO, MICHAEL P. 1970. Migratoin, Unemployment and Development: A Two-Sector Analysis. *American Economic Review*, **60**(1), 126–142.

HSIEH, CHANG-TAI. 2002. What Explains the Industrial Revolution in East Asia? Evidence from the Factor Markets. *American Economic Review*, **92**(3), 502–526.

HSIEH, CHANG-TAI, & KLENOW, PETER J. 2009. Misallocation and Manufacturing TFP in China and India. *The Quarterly Journal of Economics*, **124**(4), 1403–1448.

HSIEH, CHANG-TAI, & KLENOW, PETER J. 2010. Development Accounting. *American Economic Journal: Macroeconomics*, **2**(1), 207–223.

LIPSEY, ROBERT E., & SJOHOLM, FREDRIK. 2004. Foreign direct investment, education and wages in Indonesian manufacturing. *Journal of Development Economics*, **73**(1), 415–422.

LUCAS, JR., ROBERT E. 1990. Why Doesn't Capital Flow From Rich to Poor Countries? *American Economic Review*, **80**(2), 92–96.

MINNESOTA POPULATION CENTER. 2014. *Integrated Public Use Microdata Series — International: Version 6.3 [Machine-readable database].* Minneapolis: University of Minnesota. http://www.ipums.org.

MONGE-NARANJO, ALEXANDER, SANCHEZ, JUAN, & SANTAEULALIA-LLOPIS, RAUL. forthcoming. Natural Resources and Global Misallocation. *American Economic Journal: Macroeconomics*.

OI, WALTER Y, & IDSON, TODD L. 1999. Firm size and wages. *Pages 2165–2214 of: Handbook of Labor Economics*, vol. 3. Elsevier.

PARENTE, STEPHEN L., & PRESCOTT, EDWARD C. 1994. Barriers to Technology Adoption and Development. *Journal of Political Economy*, **102**(2), 298–321.

PRESCOTT, EDWARD C. 1998. Needed: A Theory of Total Factor Productivity. *International Economic Review*, **39**(3), 525–551.

RESTUCCIA, DIEGO, & ROGERSON, RICHARD. 2008. Policy distortions and aggregate productivity with heterogeneous establishments. *Review of Economic dynamics*, **11**(4), 707–720.

RESTUCCIA, DIEGO, YANG, DENNIS T., & ZHU, XIAODONG. 2008. Agriculture and Aggregate Productivity: A Quantitative Cross-Country Analysis. *Journal of Monetary Economics*, **55**(2), 234–250.

ROBERT WALTERS. 2017. *Salary Survey*.