Human Capital and Development Accounting Revisited*
Job Market Paper

Hannes Malmberg†
Click Here for Latest Version

December 30, 2016

Abstract

I study the role of human capital in accounting for world income differences when skilled and unskilled labor services are imperfect substitutes. With imperfect substitutability, the relative price of skilled and unskilled labor services varies across countries. To measure the contribution of human capital to output, this relative price needs to be estimated. I develop a novel method for estimating this relative price using international trade data. My method relies on the negative relationship between relative prices of skilled labor services and relative export values in skill intensive industries. Using a gravity trade framework, I estimate that relative skilled labor service prices are approximately four log points lower in rich countries. Skilled wage premia are approximately one log point lower in rich countries, and I interpret my findings as reflecting a relatively high quality of skilled labor in rich countries. When I integrate my estimates into a development accounting exercise, the share of world income differences accounted for by human capital rises from 12% to 65%. The implied TFP differences between rich and poor countries shrink by 66%. I conclude that when accounting for world income differences, we cannot reject a dominant role for human capital.

*I am grateful to Timo Boppart, John Hassler, Karl Harmenberg, Per Krusell, Erik Öberg, and Torsten Persson for extensive comments on the project. Moreover, I am grateful for helpful discussions with Rodrigo Adao, Ingvild Almås, Francisco Buera, Timothy Besley, Konrad Burchardi, Saman Darougheh, Niels-Jakob Harbo Hansen, Benjamin Jones, Bo Malmberg, Kurt Mitman, Benjamin Moll, Jon de Quidt, Elias Papaioannou, Richard Rogerson, Klaus Schmidt, Jósef Sigurdsson, Vanessa Sternbeck Fryxell, Kjetil Storesletten, David Strömberg, Alwyn Young, and seminar participants at DEGIT XXI, SUDSWEC 2016, OXDEV 2016, and the IIES internal seminar. Thanks to Christina Lönnblad for proofreading. All errors are my own.

†Graduate student. Institute for International Economic Studies. Email: hannes.malmberg@iies.su.se.

1 Introduction

Are cross-country income differences due to differences in factor supplies, or due to differences in total factor productivity (TFP)? In particular, can human capital explain world income differences?

The question of the factor-TFP split is an important and perennial question in growth and development economics. This is the question behind Robert Lucas' theories on the role of human capital (Lucas, 1988) and the cross-country regression study by Mankiw et al. (1990) (MRW), which famously suggested that accumulation differences could explain most of world income differences.

Today, the state-of-the-art estimates of the factor-TFP split come from the literature on development accounting, which was initiated by Klenow and Rodriguez-Clare (1997) and Hall and C Jones (1999). The development accounting literature disciplined the estimated output effects of human capital using microeconomic returns to schooling, and found that MRW's large output effects from human capital were inconsistent with more modest microeconomic returns to schooling. Differences in factor supplies could only explain a small share of world income differences. Later contributions by Caselli (2005), Hsieh and Klenow (2010a), and C Jones (2015) find similar results.[1]

The development accounting literature has instead emphasized low TFP as the main culprit in causing low incomes. This view has had a considerable influence on the macroeconomic growth and development literature, leading the literature to put a strong focus on theories of TFP differences – such as theories of factor misallocation and technology diffusion barriers.[2]

In this paper, I revisit the issue of the factor-TFP split, and focus on the role of human capital. I propose a new way of measuring the supply of human capital when skilled and unskilled labor services are imperfect substitutes. Using my method, I find that human capital accounts for 65% of world income differences, compared to only 12% when I use traditional development accounting methods.

My analysis follows the traditional development accounting literature in positing an aggregate production function

$$Y = K^{\alpha}(LAh)^{1-\alpha},$$

where Y is output, K is the capital stock, L is the size of the labor force, A is total factor productivity, and h is the average human capital of the labor force (Hall and C Jones, 1999).

However, whereas the traditional development accounting literature assumes that labor services

[1] Even though the mainstream view is that low TFP plays a dominant role in explaining low incomes, there have been an ongoing debate about the robustness of development accounting. See, for example, Acemoglu and Zilibotti (2001), Erosa et al. (2010), Schoellman (2011), B Jones (2014a), B Jones (2014b), Manuelli and Seshadri (2014), and Schoellman and Hendricks (2016).

[2] Parente and Prescott (1999) and Acemoglu et al. (2007) discuss barriers to technology diffusion. Restuccia and Rogerson (2008), Hsieh and Klenow (2009), and Midrigan and Xu (2014) are a few contributions in the large literature on factor misallocation.

are perfectly substitutable and that h is a linear aggregator, I follow the labor economics literature in allowing for imperfect substitutability between unskilled and skilled labor services, and I assume that h is a constant elasticity of substitution aggregator of unskilled and skilled labor services. Countries differ in their shares of skilled and unskilled workers, and they also differ in the average amount of services provided by skilled and unskilled workers – which I will call the qualities of skilled and unskilled labor.[3]

Most of the traditional development accounting setup carries over to my setting. I use the same aggregate production function, the same capital share α, a similar measure of capital, a similar measure of the share of skilled workers, and a similar calibration of the quality of unskilled labor. However, I face two challenges when measuring human capital.

First, with imperfect substitutability, the relative wages of skilled and unskilled workers are not equal to the relative amounts of labor services provided by skilled and unskilled workers. Instead, the relative wage of skilled and unskilled workers is the product of the relative *amount* of services provided by skilled and unskilled workers, and the relative *price* of those services. For example, the relative wage of a programmer and a hairdresser is given by the relative number of lines of code and haircuts that they provide, times the relative price of a line of code and a haircut. This can be expressed as:

$$\frac{w_s}{w_u} = \frac{Q_s}{Q_u}\frac{r_s}{r_u}, \tag{1}$$

where $\frac{w_s}{w_u}$ is the relative wage of skilled and unskilled workers, $\frac{Q_s}{Q_u}$ is the relative quality of skilled and unskilled workers, and $\frac{r_s}{r_u}$ is the relative price of skilled and unskilled labor services. With perfect substitutability, the relative price of skilled and unskilled labor services is constant across countries, but with imperfect substitutability, it might vary across countries. Second, I need to estimate the elasticity of substitution between skilled and unskilled labor services.[4]

Both these issues are resolved by estimating the relative price of skilled and unskilled labor services, or the relative price of skilled services for short. With estimates of the relative price of skilled services, I can obtain relative qualities from wage data using equation (1). Furthermore, I can estimate the elasticity of substitution by regressing log relative prices of skilled services on log relative supplies of skilled labo rfaservices across countries.

A key innovation in my paper is a novel method for measuring the relative price of skilled services across countries. My method is based on two premises. First, low relative prices of skilled services imply low relative unit costs in skill-intensive industries. Second, low relative unit costs

[3]Even though my top-level aggregator only has two skill levels, my formulation is consistent with skilled labor services being an arbitrary constant returns aggregator of heterogeneous types of skilled labor services. See B Jones (2014a) for a discussion of the theory of development accounting under imperfect substitutability.

[4]There are US panel and time-series based estimates of the elasticity of substitution between skilled and unskilled labor services (Katz and Autor, 1999; Ciccone and Peri, 2005). I estimate this parameter between countries to ensure that it is the relevant long-run, cross-country estimate to be used in development accounting.

in skill-intensive industries imply high relative export values in skill-intensive industries. The first premise means that relative unit cost data are informative about the relative price of skilled services; the second premise means that export value data are informative about relative unit costs. The latter point is important given the lack of detailed industry unit cost data sets covering both rich and poor countries. To estimate the relative price of skilled services, I use a gravity trade model, and I derive a regression specification that combines a trade elasticity estimate with data on export values and industry factor shares.

My trade data analysis suggests that there is a strong negative relationship between country income level and the relative price of skilled services. My baseline specification estimates a rich-poor difference in log relative skilled service prices of about 4. Using these relative price estimates, I can complete my development accounting exercise.

In my development accounting exercise, I find that human capital can account for 65% of world income differences, whereas human capital only can account for 12% of income differences using traditional development accounting methods. The estimated rich-poor log TFP difference falls by 66%. To illustrate the difference: traditional development accounting implies that increasing Tanzanian human capital to US levels would make Tanzania as rich as Senegal, whereas my estimates imply that this would make Tanzania as rich as Russia. I conclude that human capital can play a dominant role in accounting for world income differences.

The key economic driver of my results is a high estimated quality of skilled labor in rich countries. This quality estimate comes from combining moderate rich-poor differences in the skilled wage premium with large rich-poor differences in the relative price of skilled services. The trade data estimates suggest that rich countries have 4-5 lower log relative prices of skilled services. Equation (1) shows that these differences can either be explained by rich countries having low skilled wage premia, or by rich countries having a high relative quality of skilled labor. Wage data suggest that skilled wage premia are indeed lower in rich countries, but not more than approximately one log point lower. Using equation (1), I impute a 3-4 log-point difference in the relative quality of skilled labor between rich and poor countries. Intuitively, to rationalize the high exports of engineering products from Germany, the productivity of German skilled labor needs to be high.

Why does traditional development accounting not detect these large quality differences? To understand this, we need to study in more detail how traditional development accounting measures human capital. Traditional development accounting constructs a measure of human capital by aggregating microeconomic returns to human capital. The method relies on all labor services being perfect substitutes, and wages reflecting human capital. These two assumptions imply that we can convert the workforce to unskilled equivalent labor units using relative wages. By assuming that unskilled labor has the same quality across countries (or by modelling unskilled quality separately, see Caselli, 2005), we can compare the supplies of human capital across countries.

Traditional development accounting underestimates differences in the quality of skilled labor, as it assumes that quality improvements of skilled labor are reflected one-to-one in the skilled wage premium. This is not true with imperfect substitutability. Instead, there are two counteracting effects. Quality improvements increase the productivity of skilled labor which pushes up the skilled wage premium. At the same time, quality improvements make skilled services more abundant which pushes down the relative price of skilled services, and, consequently, the skilled wage premium.

This mechanism can be illustrated in an economy where the human capital aggregator is Cobb-Douglas:

$$h = u^{1-\beta}(Q_s s)^{\beta}.$$

In this economy, the relative price of skilled and unskilled labor services is

$$\frac{r_s}{r_u} = \frac{\beta}{1-\beta}\frac{u}{Q_s s}.$$

The relative price of skilled services is inversely proportional to the quality of skilled labor. In this setting, the skilled wage premium is actually independent of the quality of skilled labor Q_s, as an increase in Q_s leads to a falling relative price of skilled services that precisely offsets the positive productivity effect. If a country increases its quality of skilled labor, traditional development accounting methods will not estimate any change in Q_s, and will attribute all output gains to TFP. The Cobb-Douglas functional form makes this effect stark, but the mechanism is general

Thus, I find that income differences in large part are explained by differences in skilled labor quality, Q_s. My main interpretation of this finding is that skilled workers in rich countries have higher human capital than in poor countries, which is missed by traditional development accounting due to the biases outlined above. Differences in Q_s could to some extent, of course, reflect differences in skill-augmenting technologies. Such technology differences would not be *total* factor productivity differences, but technology differences that selectively augment skilled labor. Theories that explain such technology differences would still have to account for the interaction between the economic environment and human capital.[5] In Section 5, I discuss the interpretation of Q_s further, and conclude that human capital differences are a natural interpretation of the data. An important implication of this paper is that future work should focus on the breakdown between human capital and skill augmenting technology in explaining Q_s.

The outline of the paper is as follows. Section 2 describes the model environment in the development accounting exercise. Section 3 describes how the variables and parameters in the development accounting exercise are measured and estimated. The measurement section is split into two parts: one part estimates r_s/r_u, and another part estimates all other variables assuming that r_s/r_u is

[5] For a discussion of skill-biased technology differences, see Caselli and Coleman (2006). In particular, approaches that have been used to understand TFP differences, such as misallocation and barriers to technology adoption, can still be used to understand how these effects disproportionately affect skilled labor.

known. This reflects that r_s/r_u is the most challenging variable to measure. Section 4 presents the development accounting results. Section 5 discusses the economic interpretation of my results, as well as potential alternative interpretations, including a discussion of skill-biased technology differences (Caselli and Coleman, 2006; Caselli, 2015) and migrant wage data (Hendricks, 2002). Section 6 discusses the relationship between my paper and B Jones (2014a) who also studies development accounting with imperfectly substitutable labor services. Section 7 performs a large number of robustness checks on the baseline results, and Section 8 concludes the paper.

Related literature. My paper is part of the development accounting literature, going back to Klenow and Rodriguez-Clare (1997) and Hall and C Jones (1999). This literature is surveyed in Caselli (2005), Hsieh and Klenow (2010a), and C Jones (2015). There has been a number of papers revisiting the contribution of human capital in development accounting, most often in a framework featuring perfect substitutability between different types of labor services. These papers include Hendricks (2002), Erosa et al. (2010), Schoellman (2011), Manuelli and Seshadri (2014), and Hendricks and Schoellman (2016).

A few papers have analyzed development accounting with imperfectly substitutable labor services. As mentioned, these papers include Caselli and Coleman (2006) and Caselli (2015) which I discuss in Section 5, as well as B Jones (2014a) and B Jones (2014b), which I discuss in Section 6. Another paper that studies development accounting with imperfectly substitutable labor services is Caselli and Ciccone (2013). They claim that traditional development accounting provides an upper bound for the contribution of human capital to income differences. The argument is that there are diminishing returns to moving people from low to high educational groups when different educational groups are imperfect substitutes. I find different results as I analyze quality changes within skilled groups, and not just movements between different skill groups.

Beyond development accounting, my paper builds on the gravity trade literature to estimate the relative prices of skilled services (Tinbergen, 1962; Anderson et al., 1979; Eaton and Kortum, 2002; Anderson and van Wincoop, 2003; Redding and Venables, 2004; Costinot et al., 2011; Head and Mayer, 2014). A number of papers have used trade data to obtain information about productivities, including Trefler (1993) and Levchenko and Zhang (2016). My paper also relates to the literature that uses industry data to obtain information about economic development, which include Rajan and Zingales (1998) and Ciccone and Papaioannou (2009). In the context of trade, papers that analyze the relationship between country variables and the industrial structure of trade include Romalis (2004), Nunn (2007), Chor (2010), Cuñat and Melitz (2012), and Manova (2013). This literature is reviewed in Nunn and Trefler (2015).

2 Environment

In this section, I present the modeling environment used in my development accounting exercise. I posit that output can be written as an aggregate production function

$$Y = \left(\frac{K}{Y}\right)^{\frac{\alpha}{1-\alpha}} ALh. \tag{2}$$

Here, Y is total output, K is the physical capital stock, A is a labor-augmenting technology term, L is the size of the labor force, and h captures the average human capital of the labor force. I express aggregate output as a function of the capital-output ratio. This approach follows Hall and Jones (1999) and Hsieh and Klenow (2010b), and takes into account the steady-state effects of human capital and technology differences on capital accumulation.

Human capital h is defined by

$$h = f(Q_u u, Q_s s) = \left((Q_u u)^{\frac{\eta-1}{\eta}} + a_s (Q_s s)^{\frac{\eta-1}{\eta}}\right)^{\frac{\eta}{\eta-1}}.$$

Here, u and s are the shares of unskilled and skilled workers, and Q_u and Q_s are the amount of unskilled/skilled services delivered by each unskilled/skilled worker. I will refer to Q_u and Q_s as the *qualities* of unskilled and skilled labor. The term a_s is the skilled service share and η is the elasticity of substitution between skilled and unskilled labor services.

The labor market is competitive, and the relative wage of skilled and unskilled workers is

$$\frac{w_s}{w_u} = \frac{Q_s}{Q_u}\frac{r_s}{r_u},$$

where

$$\frac{r_s}{r_u} = \frac{f_s}{f_u} = a_s\left(\frac{Q_s s}{Q_u u}\right)^{-1/\eta}$$

is the relative price of skilled and unskilled labor services.

The human capital term h has two potential interpretations. One interpretation of the aggregator is that there are two homogenous skill types, with respective qualities Q_u and Q_s. A second interpretation is that there are two aggregators $H_u = Q_u u$ and $H_s = Q_s s$, which combine heterogeneous types of services into an aggregate flow of unskilled and skilled services. With this interpretation, Q_u and Q_s represent the average flow of unskilled/skilled labor services per unit of unskilled/skilled labor, and w_s and w_u are the average wages of skilled and unskilled workers. In Appendix A.1, I show that the two interpretations give the same development accounting results.

The interpretation with two types of labor services is easier to discuss, whereas the aggregator interpretation is more realistic. I will derive my results in the language of the interpretation with two labor types. I will refer to u and s as the share of unskilled and skilled workers, and to Q_u and

Q_s as the (average) qualities of unskilled and skilled labor. When I analyze economic mechanisms, I will leverage the mathematical equivalence to interpret my results in light of the aggregator interpretation.

When estimating the aggregate production function, I assume that the economy consists of multiple industries and that it trades with the outside world. In light of this, the aggregate production function should be interpreted as reflecting substitution possibilities within and between industries, as well as substitution possibilities between domestic and foreign production. Appendix A.2 provides more details on how I aggregate the production side, allowing for multiple industries and trade opportunities. Appendix A.3 motivates my choice of functional form.

3 Measurement

In this section, I measure and estimate the variables and parameters I need for my development accounting exercise. The most challenging part is to estimate the relative price of skilled services r_s/r_u. Thus, the measurement exercise splits naturally into two parts: measuring r_s/r_u, which I do in Section 3.1, and measuring all other variables given that r_s/r_u is known, which I do in Section 3.2. I dicuss my measurement choices in more detail in Appendix B.

3.1 Estimating the relative price of skilled services

The aim of this section is to estimate how the relative price of skilled and unskilled services r_s/r_u varies across countries. For this purpose, I construct a method for estimating relative factor service prices in general. My method uses a trade elasticity estimate together with industry factor share data and bilateral trade data.

3.1.1 Estimation strategy

My estimation strategy is based on two premises. The first premise is that relative factor service prices influence relative unit production costs. To illustrate this, we can consider a case with two industries. Consider Table 1, which shows the factor shares for "Cut and Sew Apparel" (NAICS code 3152) and "Communications Equipment" (NAICS code 3342). Production of Communications Equipment is more skill intensive than production of Cut and Sew Apparel. If the relative price of skilled services rises, we expect a rise in the relative unit production cost of Communications Equipment compared to that of Cut and Sew Apparel.[6]

The second premise of my strategy is that relative unit production costs affect relative export flows, which is a version of the principle of comparative advantage. To illustrate this premise,

[6] The cost shares are defined as shares of gross output. In Appendix C.3, I desribe an alternative method where I decompose the non-tradable component of the intermediate input cost share into cost shares of other inputs using an input-output table. The final results are not affected by whether I use the basic cost shares or perform such a decomposition.

we can consider a set of export values. Consider Table 2, which presents a number of US and Indonesian export values to Japan. Relative Indonesian-US exports are much higher in Cut and Sew Apparel compared to Communications Equipment. I use the principle of comparative advantage to interpret this as evidence of Indonesia having a high relative unit production cost of Communications Equipment.

By combining the first and second premise, I can obtain information about relative factor service prices from trade data. For example, the trade data in Table 2 suggests that Indonesia has a high relative unit production cost of Communications Equipment. Furthermore, factor shares in Table 1 suggests that Communications Equipment production is more skill intensive than Cut and Sew Apparel production. These two facts together suggest that Indonesia has a high relative price of skilled services.

The key feature of my estimation strategy is this method of obtaining information about relative factor service prices using relative export values conditional on trade destination. To formalize and generalize this method, I rely on a gravity trade model. My main result is that using a gravity trade model, it is possible to identify relative factor service prices using:

1. Industry factor shares
2. Bilateral industry trade data
3. The price elasticity of export flows

One particular feature of my estimation strategy is that I use trade data to obtain indirect information about relative unit costs. This estimation choice reflects the lack of a data set that provides detailed cross-country comparable industry unit cost data, and which covers both rich and poor countries. The best available data set comes from the Groningen Growth and Development Center, which has done important work in constructing a data set of industry unit costs for cross-country comparisons (Inklaar and Timmer, 2008). However, their data set only covers 35 industries in 42 countries, with a limited coverage of poor countries. In contrast, trade data are recorded on a highly detailed industry level in both rich and poor countries. This makes trade data an attractive source of information for development accounting. In Section 7.3, I show that for countries where we have both unit cost data and trade data, analyses using unit cost data and trade data yield similar results.

3.1.2 Setup

This section describes the setup of my estimation exercise. The notation is summarized in Table 3.

There are $I = 103$ countries, and each country has $K = 84$ industries.[7] I observe the value of

[7] The countries correspond to the countries with available data on export values, output levels, capital stocks, schooling levels, and shares of workers in skilled occupations. The industries correspond to NAICS 4-digit manufacturing industries.

Table 1: Factor shares for Cut and Sew Apparel and Communication Equipment

	Cut and Sew Apparel	Communications Equipment
Factor services (f)	US factor shares	US factor shares
Unskilled labor	0.08	0.03
Skilled labor	0.05	0.11
Capital	0.32	0.39
Intermediate inputs	0.54	0.46
Energy	0.01	0.01
Sum	1.00	1.00

Table 2: Selected export values from Indonesia and USA to Japan (thousands of United States dollars)

Origin	Destination	Industry	Export value
Indonesia	Japan	Cut and Sew Apparel	565,993
USA	Japan	Cut and Sew Apparel	197,100
Indonesia	Japan	Communications Equipment	16,503
USA	Japan	Communications Equipment	236,103

trade flows $x_{i,j}^k$ from country i to country j in industry k. Each industry produces a good using $F = 5$ factor services. In my baseline analysis, these are services from unskilled labor, skilled labor, capital, intermediate inputs, and energy. $r_{i,f}$ denotes the price of factor service f in country i. The unit production cost c_i^k of industry k in country i is a function of factor service prices. The relationship is given by

$$c_i^k = \frac{C^k(r_{i,1}, \ldots, r_{i,F})}{Z_i}.$$

This assumption implies that there is an industry cost function C^k that is common across countries. In an individual country, the unit cost function c_i^k is derived by deflating the common industry cost function C^k with a country-specific productivity term Z_i, which is common across industries. This particular setup implies that cross-country differences in relative unit costs only stem from cross-country differences in relative factor service prices. However, my development accounting results are not affected if the setup is modified to allow for cross-country differences in technologies that augment non-labor inputs, or if the setup is modified to allow for cross-country differences in technologies that augment all types of labor services equally.[8]

[8] However, I interpret any specific shifter of the price of skilled labor services as reflecting quality differences. In Section 5, I discuss the quality vs technology interpretation of shifters of the cost of skilled services, and in Section 7.1, I discuss regression specifications that address other potential confounders in the specification of unit costs.

Table 3: Notation

Variable	Description
i	Origin country
j	Destination country
k	Industry
f	Factor service ($f = 1$ unskilled labor services)
$x_{i,j}^k$	USD export value of industry k from country i to country j
$r_{i,f}$	Factor service price of factor f in country i
$\alpha_{i,f}^k$	Cost share of factor f in industry k in country i
c_i^k	Unit cost of industry k in country i
σ	Price elasticity of trade

3.1.3 Key equations

My estimation builds on the following two equations:

$$\log(x_{i,j}^k) = \delta_{i,j} + \mu_j^k - (\sigma - 1)\log(c_i^k) \tag{3}$$

$$\log(c_i^k) = \log(c_{US}^k) - \log\left(\frac{Z_i}{Z_{US}}\right) + \log\left(\frac{r_{i,1}}{r_{US,1}}\right) + \sum_{f=2}^{F} \alpha_{US,f}^k \log\left(\frac{r_{i,f}/r_{i,1}}{r_{US,f}/r_{US,1}}\right). \tag{4}$$

The first equation (3) is a traditional gravity trade equation. The log export value from country i to country j in industry k depends on three terms. The first term is a bilateral fixed effect $\delta_{i,j}$. It captures determinants of overall bilateral trade flows such as the size of the two countries, their bilateral distance, common legal origins, shared language, etc. The second term is a destination-industry fixed effect μ_j^k that captures the demand for good k in destination j, as well as how good access country j has to industry k given its other trading partners. The third term captures that conditional on the two fixed effects, exports depend negatively on origin unit production costs, with a price elasticity $\sigma - 1$. In Appendix C.1, I show how equation (3) can be derived from both a trade model in the style of Eaton and Kortum (2002) where trade is driven by country-variety specific productivity shocks, and from an Armington model where each country produces a unique variety of each good k.

The second equation (4) is a log-linear approximation of industry unit costs around the US cost structure, where $f = 1$ indexes unskilled labor services. I obtain the the approximation in two steps. I first note that

$$\begin{aligned}\log(C^k(r_{i,1},\ldots,r_{i,F})) &\approx \log(C^k(r_{US,1},\ldots,r_{US,F})) + \sum_{f=1}^{F} \frac{\partial C^k}{\partial r_f}\frac{r_{US,f}}{C^k}\log\left(\frac{r_{i,f}}{r_{US,f}}\right)\\ &= \log(C^k(r_{US,1},\ldots,r_{US,F})) + \sum_{f=1}^{F}\alpha^k_{US,f}\log\left(\frac{r_{i,f}}{r_{US,f}}\right).\end{aligned}$$

where C^k is the common cost function of industry k, and $\alpha^k_{US,f}$ denotes the US factor share of factor f in industry k. The second line uses Shepherd's lemma applied to the cost function to conclude that $\alpha^k_{US,f} = \frac{\partial C^k}{\partial r_f}\frac{r_{US,f}}{C^k}$ when firms are price-takers.

Combining this expression with $c^k_i = \frac{C^k(r_{i,1},\ldots,r_{i,F})}{Z_i}$ gives me

$$\log(c^k_i) = \log(c^k_{US}) - \log\left(\frac{Z_i}{Z_{US}}\right) + \sum_{f=1}^{F}\alpha^k_{US,f}\log\left(\frac{r_{i,f}}{r_{US,f}}\right). \tag{5}$$

I re-arrange this equation to equation (4), as my aim is to find the *relative* price of factor services compared to unskilled labor services, $\log\left(\frac{r_{i,f}/r_{1,f}}{r_{US,f}/r_{US,1}}\right)$. This makes it useful to normalize equation (5) with the price of unskilled labor services. I use that factor shares sum to 1 to express the unskilled cost share $\alpha^k_{US,1}$ in terms of the other cost shares: $\alpha^k_{US,1} = 1 - \sum_{f=2}^{F}\alpha^k_{US,f}$. Substituting this expression into (5) gives me equation (4).

Equation (4) decomposes log unit cost differences from the US into one term capturing absolute productivity differences, one term capturing differences in the cost of unskilled labor, and a linear combination of relative factor service price differences times US factor shares. Equation (4) shows that countries with a relatively high factor service price in factor f (high $\log\left(\frac{r_{i,f}/r_{i,1}}{r_{US,f}/r_{US,1}}\right)$) will have relatively high unit costs in sectors intensive in factor f (relatively high $\alpha^k_{US,f}$)).

As explained, equation (4) comes from a log linear approximation around the US cost structure. If industry production functions are Cobb-Douglas, this approximation is exact. If industry production functions are not Cobb-Douglas, there is a second-order bias. In Section 7.1, I analyze the effect of relaxing the Cobb-Douglas assumption.

3.1.4 Regression specification

To derive my regression specification, I combine the gravity equation (3) and the unit cost equation (4). I obtain

$$\log(x^k_{i,j}) = \tilde{\delta}_{i,j} + \tilde{\mu}^k_j - (\sigma-1)\sum_{f=2}^{F}\alpha^k_{US,f}\log\left(\frac{r_{i,f}/r_{i,1}}{r_{US,f}/r_{US,1}}\right).$$

Here, $\tilde{\delta}_{i,j} = \delta_{i,j} - (\sigma - 1)\left(\log\left(\frac{r_{i,1}}{r_{US,1}}\right) - \log\left(\frac{Z_i}{Z_US}\right)\right)$ denotes a modified fixed effect that includes the trade bilateral fixed effect, the origin absolute advantage, and the origin unskilled factor service prices. The term $\tilde{\mu}_j^k = \mu_j^k - (\sigma - 1)\log(c_{US}^k)$ denotes a modified fixed effect that includes the trade destination-industry fixed effect μ_j^k and US industry unit costs.

I can use this equation to derive a regression specification. For this purpose, I note that I can measure $x_{i,j}^k$ from international trade data, that I can measure $\alpha_{US,f}^k$ from American industry data, and that I can use the trade literature to obtain estimates of σ.[9] Thus, $\log(x_{i,j}^k)$ is my left-hand variable, and $(\sigma-1)\alpha_{US,f}^k$ for $f = 2, \ldots, F$ are my explanatory variables. My aim is to estimate the relative factor service price differences $\log\left(\frac{r_{i,f}/r_{i,1}}{r_{US,f}/r_{US,1}}\right)$. This quantity varies on a country-factor basis. Therefore, I want to estimate one parameter for each factor-country combination, and I write $\beta_{i,f}$ for this set of parameters. Given the interpretation of $\beta_{i,f}$ as differences in relative factor service prices compared to those in the US, I normalize $\beta_{i,f}$ by setting $\beta_{US,f} = 0$ for all f.

I obtain the following specification:

$$\log(x_{i,j}^k) = \tilde{\delta}_{i,j} + \mu_j^k - \sum_{f=2}^{F}\left[(\sigma - 1)\alpha_{US,f}^k\right] \times \beta_{i,f} + \varepsilon_{i,j}^k, \quad \beta_{US,f} = 0 \;\; f = 2, \ldots, F. \tag{6}$$

I regress log bilateral trade flows on a bilateral fixed effect, a destination-industry fixed effect, and $-(\sigma - 1)\alpha_{US,f}^k$ for $f = 2, \ldots, F$, allowing for country-factor specific parameters $\beta_{i,f}$. In total, I estimate $(5 - 1) \times 103 = 412$ parameters: one for each country-factor combination, excluding unskilled labor services. With this regression specification, $\beta_{i,f} = \log\left(\frac{r_{i,f}/r_{i,1}}{r_{US,f}/r_{US,1}}\right)$ identifies the difference between country i and the US in the log relative price of factor service f compared to unskilled labor services. The difference to the US in the log relative price of skilled labor services is identified by $\beta_{i,skill}$.

3.1.5 Data in trade regression

The regression equation (6) requires data on bilateral trade flows $x_{i,j}^k$, US factor shares $\alpha_{US,f}^k$, and a parameter estimate for the trade elasticity σ.

For trade flows, I use the BACI data set which is compiled by CEPII and based on COMTRADE (Gaulier and Zignago, 2010). For each country-destination pair, it reports export values on the HS 2007 6-digit industry level. I use data for 2010.

I measure factor shares by combining data from the NBER-CES Manufacturing Industry Database (Bartelsman and Gray, 1996) with data from the Occupational Employment Statistics (OES) survey. I use the NBER-CES database to obtain the cost shares of capital, labor, materials, and energy. I define the shares of labor, materials, and energy as factor outlays divided by industry

[9] Some papers estimate σ directly from trade data (Broda et al., 2004; Soderbery, 2015), exploiting short run variations in trade prices and quantities. As I am interested in the long-run elasticity of trade, I choose a calibration approach to select σ.

gross output, and I define the capital share as 1 minus the other factor shares. To find the shares of skilled and unskilled services, I use the OES to calculate the share of payroll in each industry that goes to workers in occupations with skill levels 3 and 4 in the ISCO-08 classification. This corresponds to the major occupational groups "Managers", "Professionals", and "Technicians and Associate Professionals". I calculate the skill share as the labor share from the NBER CES times the share of payroll going to skilled workers, and the unskilled share as the labor share times the share of payroll going to unskilled workers. Note that in my regression, I include the materials and energy shares in the regression. Appendix C.3 provides a more detailed discussion of different choices of intermediate input measurement and their effects.

The regression is performed using NAICS 4-digit coding, which is the coding scheme of the OES industry data. The trade data is recorded using HS6 codes and the NBER-CES data is recorded using NAICS 6-digit codes. The OES occupational data is recorded according to SOC, and it is converted to ISCO-08 to calculate the share of payroll going to skilled workers. All factor share and trade data are converted between coding schemes using a concordance procedure described in Appendix F.

I take my value of the trade elasticity σ from the literature. I look for an estimate of the *long-run* elasticity *between different foreign varieties in the same industry*. This choice reflects the nature of my regression. The regression is run between countries on different parts of the world-income distribution, and aims to capture persistent cross-country differences. Furthermore, the regression explains a source country's exports conditioned on the total industry imports of a destination country. Thus, the relevant elasticity is the long-run elasticity between different foreign varieties.

I select $\sigma = 10$ as my baseline elasticity. This is a reasonably high estimate of trade elasticity and reflects a conservative choice for estimating the importance of human capital. A higher σ shrinks the importance of human capital as it reduces the estimated differences between countries: differences in relative trade flows translate into smaller unit cost differences. Even though Eaton and Kortum (2002) open up for estimates as high as $\sigma = 14$, my estimate is higher than $\sigma = 5$ found in Simonovska and Waugh (2014), $\sigma = 7.2$ found in Costinot et al. (2011), and the baseline $\sigma = 9.2$ found in Eaton and Kortum (2002).[10]

My baseline estimate $\sigma = 10$ corresponds to the higher range estimates found in Romalis (2007) when he estimates the trade effects of NAFTA. He calculates a pooled trade elasticity by investigating how differential reductions of tariffs due to NAFTA affected trade in the quadrangle USA, Canada, Mexico, and the EU. I select this high estimate to be conservative, and due to the fact that the long-run effects of NAFTA studied by Romalis (2007) reflect the type of long-run, foreign-to-foreign substitution that my regression specification seeks to capture. In Section 7.1, I

[10]Note that the trade elasticity θ in Eaton and Kortum-style models represents the elasticities of export value with respect to price changes, whereas σ represents the elasticity of quantity with respect to price changes. Hence, $\sigma = \theta + 1$ when we convert between the two type of parameters.

discuss the effects of making different assumptions about σ.

3.1.6 Results from trade regression

My main results are displayed in abridged form in Table 4. The table presents log relative factor service price estimates for different factors, and for six randomly selected countries in each World Bank Income group. Standard errors are calculated by clustering on industry-country level.

The table shows that poor countries in general have higher relative factor service prices for skilled services, capital services, and intermediate input services. The pattern is especially pronounced for skilled services. There is some tendency for relative energy service prices to be higher in poor countries, but this pattern is less clear. Relative energy service prices vary more between similar countries and are less precisely estimated.

My primary interest is in the relative prices of skilled services, as these are used in my development accounting exercise. In Figure 1, I provide a graphical illustration of the relationship between estimated relative skilled service prices and log GDP per worker. There is a strong negative relationship, and poor countries have approximately 4-5 log points higher relative prices of skilled services. If I take standard errors into account, the results are consistent with a stable, almost linear, relationship between log GDP per worker and the log relative price of skilled services. There is a less clear relationship between log GDP per worker and skilled service prices for the very poorest countries, which could reflect that manufacturing exports are relatively unimportant in these countries. In Section 7.3, I analyze the effect of excluding the poorest countries from the analysis, and I show that this increases the estimated importance of human capital.

Even though I do not use the other factor service price estimates in my development accounting exercise, they can be used to evaluate the estimation method. Appendix C.2 discusses the results for other factors in greater detail.

3.2 Measurement given known relative skill prices

Section 3.1 estimated the relative price of skilled services r_s/r_u. In this section, I describe how I measure and estimate the other variables used in my development accounting exercise, given that the relative price of skilled services is known.

Data on real output Y, labor force size L, and physical capital stock K are from the Penn World Table Version 8.1. I use data from 2010, and I set the capital share α to 1/3.

I use data from ILO to measure the share of skilled workers s. I define the share of skilled workers as the share of workers having an occupation requiring skill level 3 or 4. According to ILO, occupations require skill level 3 or 4 when they "typically involve the performance of [...] tasks that require an extensive body of [...] knowledge in a specialized field". In the International Standard Classification of Occupations 2008 (ISCO-08), these are "Managers", "Professionals", and "Technicians and Associate Professionals". Figure 2 shows the relationship between the share of

Table 4: Regression estimates of log relative factor service price parameters (US = 0)

	Factor services			
	Skilled labor	Capital	Intermediate inputs	Energy
Low income: Gambia	5.64 (1.01)	2.53 (0.68)	2.56 (0.57)	−1.17 (1.01)
Low income: Liberia	4.56 (1.09)	2.73 (0.77)	2.28 (0.64)	3.03 (1.43)
Low income: Nepal	5.89 (1.15)	2.83 (0.77)	2.85 (0.74)	5.35 (1.78)
Low income: Rwanda	4.54 (1.31)	1.40 (0.75)	1.55 (0.72)	3.04 (1.78)
Low income: Tanzania	3.72 (1.03)	1.45 (0.69)	1.54 (0.55)	0.21 (1.25)
Low income: Uganda	3.09 (0.96)	0.82 (0.64)	0.93 (0.52)	1.44 (1.32)
Lower middle income: Indonesia	3.78 (0.98)	1.41 (0.64)	1.65 (0.57)	0.47 (1.20)
Lower middle income: Pakistan	4.92 (1.08)	2.14 (0.76)	2.35 (0.68)	2.12 (1.61)
Lower middle income: Philippines	1.52 (1.06)	0.57 (0.73)	1.02 (0.58)	1.91 (1.20)
Lower middle income: Tunisia	2.62 (1.03)	1.61 (0.66)	1.54 (0.55)	0.81 (1.29)
Lower middle income: Ukraine	2.45 (0.92)	0.92 (0.63)	0.96 (0.52)	−2.30 (1.49)
Lower middle income: Vietnam	3.60 (1.21)	2.15 (0.78)	2.39 (0.67)	3.10 (1.35)
Upper middle income: Colombia	3.74 (0.96)	1.01 (0.59)	1.37 (0.50)	−0.87 (1.11)
Upper middle income: Dominican Republic	3.27 (1.17)	0.65 (0.73)	1.40 (0.64)	0.75 (1.38)
Upper middle income: Paraguay	5.67 (1.18)	1.21 (0.71)	1.27 (0.67)	1.59 (1.82)
Upper middle income: Russia	1.12 (0.95)	0.001 (0.65)	−0.10 (0.55)	−4.57 (1.19)
Upper middle income: South Africa	1.67 (0.90)	0.46 (0.58)	0.43 (0.47)	−1.61 (1.08)
Upper middle income: Turkey	3.59 (0.97)	1.95 (0.60)	2.09 (0.49)	0.40 (1.18)
High income: Chile	4.13 (1.09)	0.54 (0.65)	0.65 (0.56)	−1.20 (1.46)
High income: Ireland	−0.10 (0.99)	−1.05 (0.68)	−0.53 (0.61)	0.21 (1.66)
High income: Netherlands	0.59 (0.88)	−0.45 (0.56)	−0.17 (0.45)	−0.24 (1.00)
High income: New Zealand	1.54 (0.91)	0.62 (0.62)	0.32 (0.58)	1.15 (1.19)
High income: Taiwan	−0.10 (1.07)	1.37 (0.67)	1.27 (0.58)	0.58 (1.31)
High income: United States	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)
Observations	453,147			
R^2	0.69			

Note: Standard errors are clustered on origin-industry level

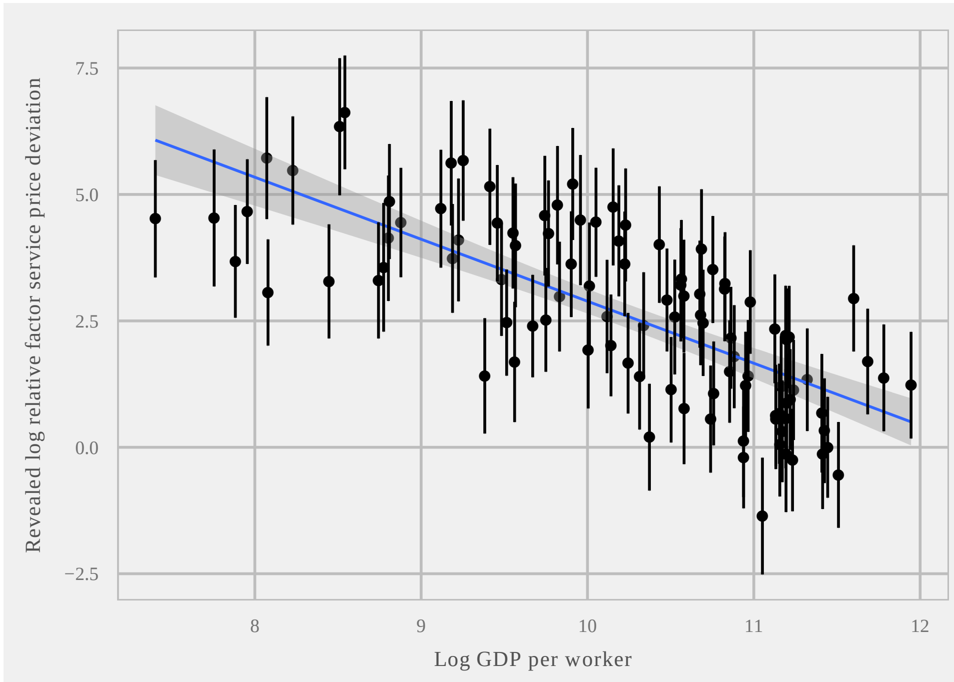

Figure 1: Log GDP per worker and log $\left(\frac{r_{i,skill}/r_{i,unskill}}{r_{US,skill}/r_{US,unskill}}\right)$

skilled workers and log GDP per worker. There is a strong positive relation, and a linear regression of the skill share on log GDP per worker has an R^2-value of 0.75. My skill definition differs from the literature in being occupation-based instead of schooling-based. I discuss this choice in Appendix B.1.

I also use ILO data to measure the skilled wage premium w_s/w_u. ILO summarizes wage data from multiple sources, and I restrict attention to countries where data are available from administrative records, a labor-focused establishment survey, and/or a labor force survey. I use the measure of mean nominal monthly earnings of employees. I combine data on wages and employment across occupations, and I calculate the relative average wage between workers with skill levels 3 or 4 and workers with skill levels 1 or 2. Figure 3 shows the relationship between log skilled wage premia and log GDP per worker. Apart from two outliers (Vietnam and Qatar), there is a strong negative relationship. The ILO data only covers a limited set of countries, and there are large variations between countries with similar levels of log GDP per worker. In my development accounting exercise, I want to use a large set of countries, and I am interested in systematic differences between rich and poor countries. Thus, to assign values of the skilled wage premium, I regress the log skilled premium on log GDP per worker (excluding outliers). I assign each country a skilled premium using the fitted value of this regression. This allows me to extend the country coverage beyond the limited set of countries covered in the ILO data, while capturing the systematic changes of

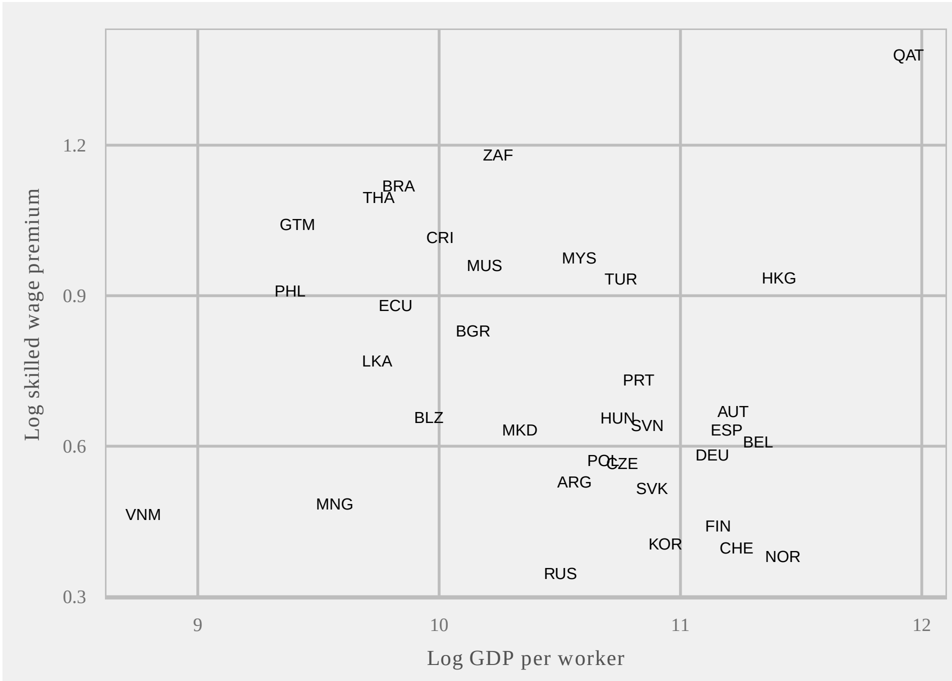

Figure 2: Log skilled wage premia versus log GDP per worker

the skilled wage premium across the GDP per worker distribution. In Section 7.3, I consider how changes in the measurement of the skilled wage premium change my results.

I calibrate the quality of unskilled labor Q_u using data on schooling levels and Mincerian returns. I define

$$Q_u = e^{\phi(S_u)}, \tag{7}$$

where S_u is the average schooling years of unskilled workers, and ϕ is a Mincerian return function capturing the relationship between schooling and wages. I measure S_u using the Barro-Lee schooling data for 2010. I assume that there is perfect postive sorting between years of schooling and working in a skilled profession, which means that unskilled workers correspond to the $1 - s$ share of the workforce with the least schooling. I assume that S_u is the average number of school years in this group.[11] I take the Mincerian return function $\phi(S)$ from Caselli (2005) and define it as a piecewise linear function with slope 0.13 for $S < 4$, slope 0.1 for $S \in [4, 8)$, and slope 0.08 for $S \geq 8$. This specification was introduced in the literature as a reduced form way of capturing that poor countries have higher Mincerian returns.

I measure the relative quality of skilled and unskilled labor using the equation

$$\frac{w_s}{w_u} = \frac{Q_s}{Q_u}\frac{r_s}{r_u} \iff \frac{Q_s}{Q_u} = \frac{w_s/w_u}{r_s/r_u}. \tag{8}$$

[11] See Appendix B.2 for details on how I calculate the average schooling of unskilled labor.

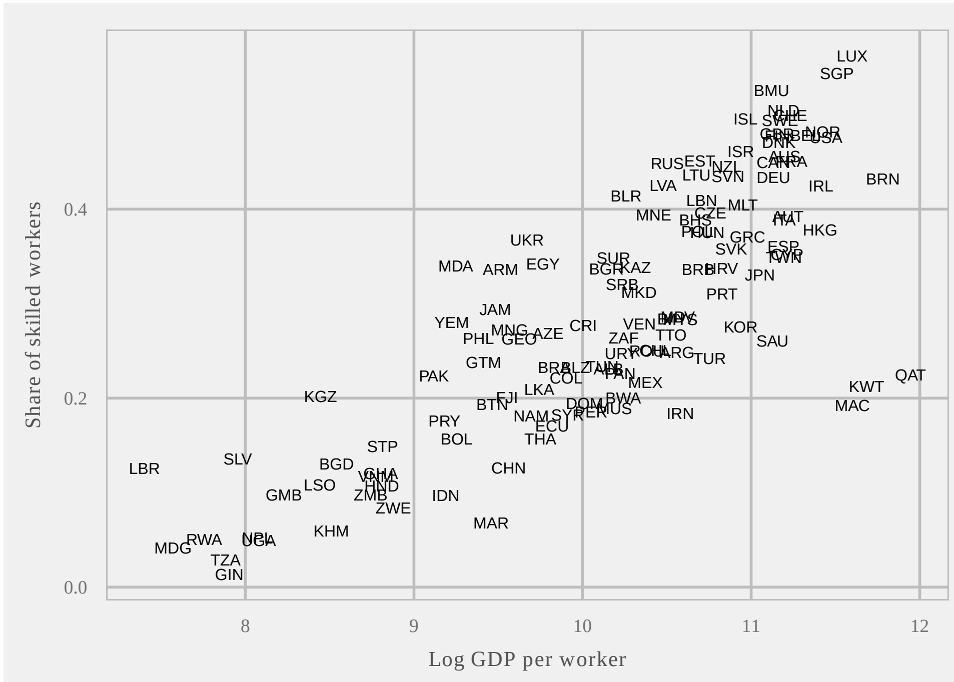

Figure 3: Share of workers in skill level 3+4 vs log GDP per worker

The skilled wage premium w_s/w_u is observable, and the relative price of skilled services $\frac{r_s}{r_u}$ was estimated in Section 3.1.[12] This equation states that the skill premium equals the relative *amount* of services provided by skilled and unskilled workers, times the relative *price* of those services. Re-arranging the equation shows that there will be a high estimate of relative skilled labor quality if either a) the skill premium is high given the price of skilled services, since this reflects a large amount of services being delivered, or b) if observed skilled service prices are low given the skill premium, since this reflects a high quality of skilled labor, thus bringing down the quality adjusted price.

The estimates of relative skilled labor quality Q_s/Q_u can be used to estimate the human capital aggregator parameters a_s and η. In a competitive labor market, the wage premium can be expressed as

$$\frac{w_s}{w_u} = \frac{a_S Q_s^{1-1/\eta}}{Q_u^{1-1/\eta}} \left(\frac{s}{u}\right)^{-1/\eta} \Leftrightarrow \log\left(\frac{r_s}{r_u}\right) = \log(a_s) - \frac{1}{\eta}\log\left(\frac{Q_s s}{Q_u u}\right). \tag{9}$$

I recover $\log(a_s)$ and $-1/\eta$ as the intercept and slope from a cross-country regression of log relative

[12]In Section 3.1, I estimated $\frac{r_s/r_u}{r_{US,s}/r_{US,u}}$. To find r_s/r_u, I need $r_{US,s}/r_{US,u}$. I find this by normalizing US skilled labor quality to $Q_{US,s} = 1$. This implies:

$$\frac{r_{US,s}}{r_{US,u}} = \left(\frac{1}{Q_{US,u}}\right)^{-1} \frac{w_{US,s}}{w_{US,u}}.$$

service prices $\log\left(\frac{r_s}{r_u}\right)$, on log relative service supplies, $\log\left(\frac{Q_s s}{Q_u u}\right)$. This specification is a close cross-country analogue of the regression specification introduced in Katz and Murphy (1992).[13] I estimate a skill share $a_s = 2.06$, and an elasticity of substitution $\eta = 1.27$.

4 Results

4.1 Baseline results

In this section, I perform development accounting using the measurements from Section 3. My main outcome variables are the shares of world income differences accounted for by physical capital, human capital, and TFP differences. To evaluate how my human capital measurement method affects development accounting, I compare my results to those obtained when human capital is measured using traditional development accounting methods.

To decompose income differences into contributions from factors and technology, I use the aggregate production function (2). The aggregate production function admits a linear decomposition of log output per worker:

$$\log(y) \equiv \log\left(\frac{Y}{L}\right) = \frac{\alpha}{1-\alpha}\log\left(\frac{K}{Y}\right) + \log(h) + \log(A).$$

Using this decomposition, I define the shares of income differences attributable to different factors:

$$\begin{aligned}
\rho^K &= \frac{Cov\left(\frac{\alpha}{1-\alpha}\log\left(\frac{K}{Y}\right), \log(y)\right)}{Var(\log(y))} \\
\rho^h &= \frac{Cov\left(\log(h), \log(y)\right)}{Var(\log(y_i))} \\
\rho^A &= 1 - \rho^K - \rho^h.
\end{aligned}$$

In addition to share parameters, I define a summary measure of TFP-differences between rich and poor countries. To define this measure, I regress log TFP on log GDP per worker which gives me predicted log TFP as a function of log GDP per worker. My definition of the rich-poor log TFP difference is the change of this predicted value between the 10^{th} and the 90^{th} percentile of the GDP per worker distribution. I write $\Delta\log(A)$ for this difference.

I also calculate the share parameters and the TFP differences using an alterative measure of human capital h_{trad}, which is constructed in line with traditional development accounting methods. It is measured by converting skilled workers to unskilled equivalents using the skilled wage

[13] The difference is that I use the trade data to obtain an independent estimate of the labor-augmenting terms Q_s/Q_u, whereas Katz and Murphy (1992) obtain identify η by assuming that there is a log-linear time trend in $\frac{Q_s}{Q_u}$, and they estimate the elasticity by deviations around this trend.

premium.[14] I define h_{trad} as

$$h_{trad} = Q_u \left(u + s \frac{w_s}{w_u} \right), \tag{10}$$

where unskilled labor quality Q_u is defined in equation (7).

To compare my measure h_{new} with the traditional development accounting measure h_{trad}, I compare how the share of world income differences explained by human capital – ρ^h – changes when I change the human capital measure from h_{trad} to h_{new}. Furthermore, I estimate the reduction in log TFP differences between rich and poor countries when I change the human capital measure from h_{trad} to h_{new}. To measure this reduction, I define the share of TFP differences explained as

$$TFP_{share} = 1 - \frac{\Delta \log(A_{new})}{\Delta \log(A_{trad})}.$$

To interpret this measure, recall that $\Delta \log(A)$ refers to the difference in log TFP between rich and poor countries. If there are no TFP differences left between rich and poor countries with my method of measuring human capital, $TFP_{share} = 1$. If TFP differences between rich and poor countries are the same with my method of measuring human capital as with the traditional development accounting method, $TFP_{share} = 0$.

Table 5 presents the baseline results of my development accounting exercise. Capital-output variations explain 8% of income differences. This share does not depend on the method of measuring human capital. The traditional development accounting method attributes 12% of world income differences to human capital, and 79% to TFP.[15] My method attributes 65% of world income differences to human capital, and only 26% to TFP. Estimated log TFP differences between rich and poor countries shrink 67% when I change the human capital measurement method.

4.2 Intuition from country example: Tanzania

To make the development accounting results more concrete, I focus on what they mean for one poor country: Tanzania. In 2010, Tanzania had a GDP per worker of $2650, which made it the 17^{th} poorest country among the 165 countries in the Penn World Table. I ask the following question: how do different human capital measurement methods predict that Tanzanian GDP per worker

[14] My calculation method is analogous to traditional development accounting as it calculates human capital using unskilled equivalents estimated using relative wages. The standard references in development accounting Hall and C Jones (1999) and Caselli (2005) use a slightly different implementation as they use years of schooling as their skill measure instead of occupation, and they use Mincerian returns instead of occupation-based skilled wage premia to calculate wage differences. They define human capital as $h_i = \exp(\phi(S_i))$ where ϕ is a Mincerian return function and S_i is the average years of schooling in country i. In my setting, their method yields very similar results to using equation (10).

[15] This estimate is slightly above the $50\% - 70\%$ interval discussed in the review article by Hsieh and Klenow (2010a) and the 70% in the latest handbook chapter written by C Jones (2015). Four percentage points of the difference can be explained by the Mincerian method attributing 14% to human capital. I also use a later version of the Penn World Table and updated data.

	Baseline
Capital	0.08
Human capital – traditional method	0.12
Human capital – my method	0.65
TFP – traditional method	0.79
TFP – my method	0.26
Log TFP diff – traditional method	2.54
Log TFP diff – my method	0.85
Percent of TFP differences explained	67%
Elasticity of substitution η	1.27

Table 5: Contribution of factors and TFP to income differences: baseline parametrization

would change if the skill levels of the Tanzanian workforce were increased to the levels of the US workforce, keeping the Tanzanian capital-output ratio and TFP constant?

I answer this question using both the traditional development accounting method of aggregating microeconomic returns to schooling as in Hall and C Jones (1999), and by using my way of measuring human capital.[16] Granted, it is a complex counterfactual to ceteris paribus increase the skill levels of Tanzanian workers to those of US workers – including specialized computer engineers, world-class researchers, the whole range of the US medical profession, financial experts, corporate lawyers, and so forth. However, the exercise illustrates the effect of varying the method of measuring human capital.

I start with the traditional development accounting approach. For 2010, the Barro-Lee data estimates Tanzanian average schooling levels to be 5.81 years, and US average schooling levels to be 13.18 years, a difference of approximately 7.5 years. Using the Mincerian return function from Hall and C Jones (1999) and Caselli (2005), these schooling differences translate into an approximately 0.6 log point difference in human capital. Using the aggregate production function (2), log Tanzanian GDP per worker increases by the same amount.

This example illustrates that traditional development accounting does not attribute a dominant role to human capital in explaining world income differences. Even if Tanzania increases the skill levels of its workforce all the way to US skill levels, GDP per worker only increases 0.6 log points, or to \$4675. After this change in skill levels, Tanzanian income levels would not move higher than somewhere between Senegal and Bangladesh.

In contrast, my method estimates that there is an approximate 2.6 log point difference in human capital between the US and Tanzania. After increasing the skill levels of the Tanzanian workforce, Tanzania would have a GDP per worker of approximately $\$36,000$, making it approximately as rich

[16] I use the method of Hall and C Jones (1999) instead of equation (10). In this setting, they yield very similar results, but it is easier to explain the method of Hall and C Jones (1999) in this context.

as Russia. The lower TFP of Tanzania would still make it substantially poorer than the US (with a GDP per worker of $\$93,000$), but the change in its skill levels would make it an upper middle income country.

5 Interpretation of mechanism: High quality of skilled labor

5.1 Mechanism

Section 4 showed that my method of measuring human capital attributed a much larger share of income differences to human capital differences than traditional development accounting did. The key mechanism driving this result is that my method estimates a high quality of skilled labor in rich countries. Figure 4 shows the relationship between log GDP per worker and the quality of skilled labor according to the traditional development accounting method which equates relative skilled labor quality with the skilled wage premium, and according to my method, which also allows for differences in the relative price of skilled services (in both cases, I normalize log US skilled labor quality to 0). The figure shows that traditional development accounting actually estimates that poor countries have a somewhat higher quality of skilled labor than rich countries. This reflects higher skilled wage premia in poor countries. My method paints a different picture. With my method, the quality of skilled labor is about four and a half log points lower in poor countries compared to rich countries. My large estimated quality differences reflect large estimated differences in relative skilled service prices. The relative price of skilled services and the relative quality of skilled labor are related through

$$\frac{w_s}{w_u} = \frac{Q_s}{Q_u}\frac{r_s}{r_u},$$

where $\frac{w_s}{w_u}$ is the skilled wage premium. My trade data estimates suggest that the relative price of skilled services $\frac{r_s}{r_u}$ is 4-5 log points lower in rich countries. Skilled wage premia are also lower in rich countries, but only approximately one log point lower. This means that the relative quality of skilled labor is 3-4 log points higher in rich countries. My results follow from combining this finding with the 0.5 rich-poor log difference in the quality of unskilled labor. Intuitively, large quality differences are needed to reconcile moderate differences in skilled wage premia with large differences in trade patterns.

Large skilled labor quality differences lead me to attribute more importance to human capital than traditional development accounting does. Indeed, traditional development accounting will in general underestimate the importance of human capital differences when these differences take the form of rich countries having a higher quality of skilled labor. The reason is that traditional development accounting relies on the skilled wage premium to capture the output effect of improved quality of skilled labor. However, when skilled and unskilled labor services are imperfect substitutes, improved quality of skilled labor will not increase the skilled wage premium one-for-one. Instead,

Figure 4: Log GDP per worker and log quality of skilled workers

skilled labor quality improvements lead to two counteracting effects. One effect is a standard productivity effect which increases the skilled wage premium. A second effect is a relative price effect, whereby improvements in the quality of skilled labor increase the supply of skilled services and push down the relative price of skilled services. This second effect ensures that the skilled wage premium increases less than one-for-one when skilled labor quality improves.[17]

5.2 Interpretation of skilled-labor quality differences

The previous section showed that the log quality of skilled labor is approximately 4 higher in rich countries than in poor countries. This corresponds to rich countries having approximately 50 times higher quality of skilled labor. In this section, I discuss my interpretation of these quality differences.

I first note that there is circumstantial evidence for at least some quality differences in skilled labor between rich and poor countries. Hanushek and Woessmann (2015) document that rich countries have a much larger share of top performers in standardized secondary school tests. OECD countries have 17.5% of total world population, but 82.5% of the 1000 top-ranked universities in the world (Center of World University Rankings).[18] Lagakos et al. (2016) document that on-the-job

[17]For further discussions of the role of quality differences and human capital accounting, see B Jones (2014a) and B Jones (2014b).

[18]In the university data, China is an outlier in the non-high income world. China has almost 50% of top-ranked

learning is more rapid in rich countries, and large differences in managerial quality between rich and poor countries found in Bloom et al. (2016) is suggestive of quality differences in skilled labor (even though one should be careful with interpreting management behavior as directly reflecting skills as it is a composite of skills and adoption incentives).

A more interesting question concerns the size of quality differences, and whether there are economic mechanisms that could make quality differences large. Even though I do not perform a calibration in this paper, we can note that there are a number of such mechanisms when the quality of skilled labor is interpreted as resulting from aggregation of heterogeneous skilled services (as noted in Section 2, this interpretation of Q_s is consistent with my development accounting procedure). For example, if there are large complementarities between skilled tasks, then low productivity in a small range of tasks can lead to large productivity losses (C Jones, 2011). In particular, we know from Kremer's O-ring theory (Kremer, 1993) that small differences in error probabilities in complex production processes can lead to large productivity losses. Another mechanism that could potentially make skill quality differences large is different levels of specialization across countries. B Jones (2014a) provides a simple calibration where modest difference in the degree of specialization can lead to large differences in the average quality of skilled labor. Intuitively, task productivities can differ enormously between experts and non-experts, and increased specialization means that a large share of tasks are produced by experts.

5.3 Alternative interpretation: skill-biased technology differences

The trade data evidence suggested that the relative price of skilled services is low in rich countries. I interpreted this as evidence for rich countries having a high quality of skilled labor. An alternative interpretation is that there are skill-biased technological differences (SBTD) between rich and poor countries. Caselli and Coleman (2006) and Caselli (2015) argue that there is evidence for SBTD between rich and poor countries. And if technology in rich countries disproportionately augments the services of skilled workers, this lowers the revealed relative price of skilled services.

With a flexible specification of variation in technology and quality across countries, it is actually not possible to distinguish SBTD from quality differences using only price and quantity data. Indeed, human capital quality and factor augmenting technology terms appear in the same way in production functions. Thus, they have the same implications for quantity and price data. Intuitively, price and quantity data alone cannot tell whether a worker is good at hammering, or has a good hammer.

On one hand, as my estimates are based on price and quantity data, this equivalence means that my estimates can be explained by flexible SBTD. But conversely, the evidence for SBTD in Caselli and Coleman (2006) and Caselli (2015) is also based on price and quantity, which means that existing evidence for SBTD equally can be interpreted as evidence for skill-biased quality

universities outside of the OECD.

differences.

To discriminate between the hypotheses, more theoretical structure or other sources of evidence are needed. In the Appendix, I conduct two tests to distinguish between a human capital and a technology interpretation of Q_s. In Appendix D.1, I show that contrary to earlier results in the literature (Hendricks, 2002), migration data does not reject the existence of large differences in human capital. The reason is that with imperfect substitutability, there is no longer a simple mapping between human capital and pre- and post-migration wages. In Appendix D.2, I also show that when technology bias is constrained to be endogenous to factor prices (as in, for example, Caselli and Coleman (2006) or Acemoglu (2007)), it is not possible to explain the large differences in Q_s without positing large differences in skilled human capital quality.

The potential technology differences underpinning Q_s suggest that Q_s could be treated either as technology or as human capital when doing development accounting. In my development accounting exercise, I treat Q_s as human capital. I make this choice for two reasons. First, this interpretation is closer to the assumption made in traditional development accounting, which is that skill bias is the same across countries (Caselli and Ciccone, 2013). Thus, my results estimate how the role of human capital changes when I modify traditional development accounting by changing the human capital aggregator keeping technology bias constant.

Second, and more importantly, I think that human capital is a more appropriate designation than technology of the role that Q_s plays in accounting for world income differences. The human capital interpretations discussed in Section 5.2 are all reflected in Q_s. Treating Q_s as technology in development accounting conceals that these standard human capital explanations of income differences are consistent with my development accounting results. Furthermore, given that Q_s plays a large role in accounting for world income differences, explanations of income differences should have a relatively large focus on human capital even if technology differences does play a part in explaining Q_s. Indeed, the technology differences discussed in this section are not TFP differences but skill-specific technology differences, and explanations of them need to explain the interaction between the economic environment and human capital.

6 Relationship to B Jones (2014)

The paper most closely related to mine is B Jones (2014a), who constructs a theory of development accounting under imperfect substitutability. His key claim is that with a general human capital aggregator, you have to scale traditional development accounting results with the marginal product of unskilled labor to obtain the full value of the human capital aggregator. A general aggregator satisfies

$$\begin{aligned} G(H_1, \dots, H_N) &= G_1 \times \left(H_1 + \sum_{i=2}^{N} \frac{G_i}{G_1} H_i \right) \\ &= G_1 \times \left(H_1 + \sum_{i=2}^{N} \frac{w_i}{w_1} H_i \right) \end{aligned}$$

where $G_i = \frac{\partial G}{\partial H_i}$, and where the second line uses a competitive market assumption. The terms in brackets on the second line represent the traditional development accounting aggregator, which has to be scaled up by the marginal product of unskilled labor G_1.

Although using a different formulation than in my paper, Jones also highlights that traditional development accounting misses quality improvements in skilled labor. In my formulation, traditional development accounting underestimates improvements in the quality of skilled labor as an increased abundance of skilled services depresses the relative price of skilled services. In Jones' formulation, an improvement in the quality of skilled labor increases the marginal product of unskilled labor, which increases the appropriate scaling on the results of traditional development accounting.

Furthermore, Jones recognizes that the quality of skilled labor can be interpreted as resulting from an aggregation of heterogeneous skilled services, which opens up for large quality differences. He emphasizes specialization which is more fully developed in an unpublished paper (Jones, 2014b). As discussed in Section 5, other potential mechanisms that can lead to large quality differences include strong complementarities between different skill types (C Jones, 2011), in particular O-ring effects (Kremer, 1993).

However, Jones' positive argument for large quality differences is less strong than his conceptual points. His quantitative argument relies on applying rich country time-series and panel estimates of the elasticity of substitution η to cross-country data. If this elasticity is globally valid, the low supply of skilled labor in poor countries must imply a very high price of skilled services. As these high skilled service prices are not observed in skilled wage premia, the quality of skilled labor has to be very low in poor countries.

The challenge to this argument is that most estimates of the elasticity of substitution are medium-run estimates done on time series data from rich countries.[19] We do not know a priori whether existing estimates are the relevant long-run cross-country elasticity estimates to be used in development accounting, or whether it is appropriate to assume a constant elasticity of substitution when analyzing cross-country data. Furthermore, the estimated importance of human capital is

[19] One of few papers that take a long-run perspective is Ciccone and Peri (2005), which estimates long-run elasticities using compulsory schooling reforms and US cross-state data on a decadal level. The estimation method is closer to my desired parameter as it is a long-run estimate, and it uses an instrument to deal with the endogeneity of state-level supply of skilled labor. Their preferred estimate is 1.5 with a standard error of 0.44. The study is unfortunately somewhat limited by weak instruments (the first-stage using the most credibly exogenous instrument has an F-value of 2.56 with 6 instruments). Furthermore, there are five observations for every state and the standard errors are not clustered on the state level, which opens up for larger standard errors.

sensitive to this elasticity parameter. Using my definition of the share of skilled workers, an elasticity of substitution $\eta = 2$ using Jones' method would bring down then share of world income differences accounted for by human capital to approximately 25% of world income differences, whereas an elasticity of substitution of approximately $\eta = 1.2$ would mean that all world income differences would be explained by human capital.

Thus, Jones' quantitative argument is difficult to evaluate if we do not have independent estimates of relative skilled service prices in poor countries. My trade data method provides such estimates, and I find that relative prices of skilled services are indeed very high in poor countries. My estimated elasticity of substitution η is 1.27. In Appendix A.3, I also provide suggestive evidence that a constant elasticity of substitution is appropriate to model cross-country data. My paper thus provides quantitative backing to Jones' conceptual points.

7 Robustness and consistency checks

Here, I present various robustness and consistency checks of my results. In Section 7.1, I analyze how sensitive my estimates of relative skilled service prices are to varying underlying assumptions and parameters. In Section 7.2, I test whether my estimates of relative skilled service prices are consistent with estimates based on unit production cost data when such data is available. In Section 7.3, I analyze how my development accounting exercise is affected when I change the measurement of skilled wage premia, and how it is affected when I exclude very poor countries and oil producing countries from the analysis. The discussion of each robustness check is brief, and Appendix E provides more detailed descriptions and discussions of the robustness checks.

Across a wide range of specifications and parameter values, the conclusion holds that the role of human capital is considerably expanded compared to findings based on traditional development accounting methods. Furthermore, for countries where both trade data and unit cost data are available, the two types of analyses give similar results. Excluding the poorest countries and oil producing countries increases the estimated importance of human capital.

7.1 Sensitivity of relative skilled service price estimates

I estimate the relative price of skilled services using the regression specification (6). In this section, I test the sensitivity of my relative price estimates to variations in the price elasticity of trade, the set of control variables, the functional form of the underlying industry production functions, and the presence of zero trade flows.

Table 6 shows how my development accounting results change when I change the elasticity of trade σ. Variations in σ are quantitatively important, and a larger σ means a lower importance of human capital. The intuition is that a larger σ means that less relative unit cost differences are needed to explain the trade data. This reduces the estimated differences in relative skilled service prices, which in turn imply a reduction in the estimated quality differences of skilled labor. Even though a larger σ implies a smaller role for human capital, the estimated importance of human capital for $\sigma = 15$ is still 4.5 times as large as that found using traditional development accounting methods. When the trade elasticity is $\sigma = 5$, human capital explains more than 100% of world income differences. In Appendix E.1, I discuss the effect of allowing trade elasticities to be different across industries.

A second potential problem in regression (6) is omitted variables in the specification of unit costs. The regression specification assumes that variations in relative unit costs are only driven by variations in relative factor service prices. If there are other determinants of unit costs correlated with relative factor service prices, there will be an omitted variable bias. I test for the importance of an omitted variable bias by controlling for potential determinants of unit costs apart from relative factor service prices. In particular, I allow there to be a country-specific penalty on external financing and/or contracting. These penalties increase the log unit cost of an industry in proportion to the financial dependence and/or contracting dependence of the industry. To measure financial dependence and contracting dependence at an industry level, I use measures similar to those developed by Rajan and Zingales (1998) and Nunn (2007), respectively. The results are presented in Table 7. Including a term for contracting sensitivity does not affect the importance of human capital, and including a term for financial sensitivity decreases the importance of human capital from 65% to 51%. In Appendix E.2, I describe the definition of industry financial and contracting sensitivities, and how they are included in my regression.

A third potential problem in regression (6) is a second-order bias in the log-linearization of unit costs. The regression specification is based on log-linearizing unit costs around the US cost structure. This log-linearization is exact if industry production functions are Cobb-Douglas. If industry production functions are not Cobb-Douglas, there will be a second-order bias as industry factor shares vary with relative factor service prices. I analyze how my results change if industry production functions are CES with a common elasticity of substitution $\xi \neq 1$. I test for this bias by creating model generated unit costs from a model where industry production functions are CES. I run my regression specification (6) on the model generated data and look for the price differences in the model such that my regressions yield similar results on actual and model generated data. This procedure allows me to gauge the bias in my baseline estimates. Table 8 shows the development accounting results for different assumed values of ξ. Appendix D.3 explains the environment, the estimation method, the results, and the economic intuition in greater detail.

A fourth potential problem in regression (6) is zero trade flows. Approximately 62% of bilateral trade flows on the NAICS 4-digit level are zero. Given that regression (6) is defined for log trade flows, export flows of value zero are dropped, which risks biasing my estimates. One way of gauging the effects of excluding zeros is to run the regression on a higher level of aggregation, which reduces the numbers of zeros. Figure 5 shows estimated relative skilled service prices when

I run the regression on 4-digit and 3-digit manufacturing industries. The 3-digit estimates are less precisely estimated as there are only 21 industries instead of 84. However, there is a very similar relationship between log income per worker and log estimated relative skilled service prices.

	Baseline ($\sigma = 10$)	$\sigma = 5$	$\sigma = 15$
Capital	0.08	0.08	0.08
Human capital – traditional method	0.12	0.12	0.12
Human capital – my method	0.65	1.33	0.45
TFP – traditional method	0.79	0.79	0.79
TFP – my method	0.26	-0.4	0.46
Log TFP diff – traditional method	2.54	2.54	2.54
Log TFP diff – my method	0.85	-1.3	1.48
Percent of TFP differences explained	67%	153%	42%
Elasticity of substitution	1.27	1.10	1.46

Table 6: Contribution of factors and TFP to income differences: different σ

	Baseline	Contracting	Financing	Both
Capital	0.08	0.08	0.08	0.08
Human capital – traditional method	0.12	0.12	0.12	0.12
Human capital – my method	0.65	0.63	0.51	0.51
TFP – traditional method	0.79	0.79	0.79	0.79
TFP – my method	0.26	0.27	0.40	0.40
Log TFP diff – traditional method	2.54	2.54	2.54	2.54
Log TFP diff – my method	0.85	0.89	1.28	1.28
Percent of TFP differences explained	67%	66%	50%	50%
Elasticity of substitution	1.27	1.28	1.35	1.35

Table 7: Contribution of factors and TFP to income differences: different control variables

7.2 Consistency between trade data and unit cost data

In Section 3.1, I used trade data to substitute for missing unit cost data. However, the Groningen Growth and Development Center has constructed a unit cost measure for 34 industries across 42 countries. A natural consistency check is whether my trade data method yields similar conclusions as a unit cost based method on this set of countries.

	$\xi = 0.6$	$\xi = 0.8$	$\xi = 1$	$\xi = 1.2$	$\xi = 1.4$
Capital	0.07	0.07	0.07	0.07	0.07
Human capital – traditional method	0.12	0.12	0.12	0.12	0.12
Human capital – my method	0.65	0.62	0.63	0.68	0.84
TFP – traditional method	0.80	0.80	0.80	0.80	0.80
TFP – my method	0.26	0.29	0.28	0.23	0.07
Log TFP diff – traditional method	2.56	2.56	2.56	2.56	2.56
Log TFP diff – my method	0.85	0.95	0.92	0.75	0.24
Percent of TFP differences explained	67%	63%	65%	71%	91%
Elasticity of substitution η	1.36	1.39	1.38	1.34	1.26

Table 8: Contribution of factors and TFP to income differences: different ξ.

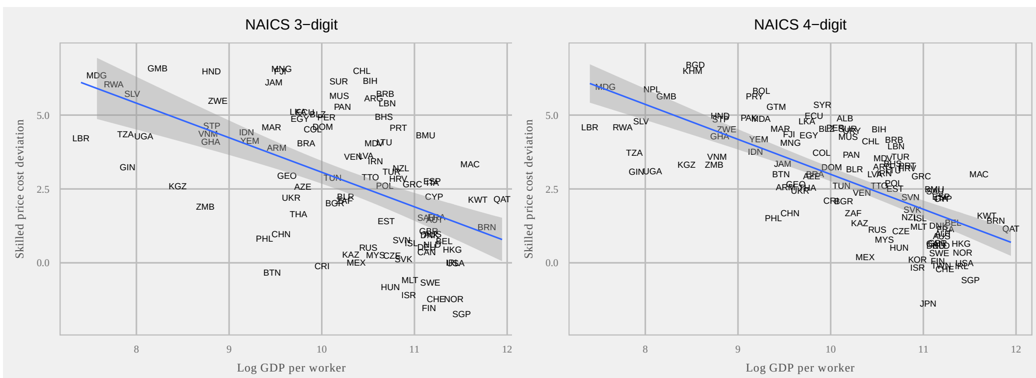

Figure 5: Estimated skill price differences with $\sigma = 10$. NAICS 3-digit and 4-digit.

The GGDC index covers both tradable and non-tradable industries, and manufacturing as well as services. Using the GGDC data set, I can run a unit cost regression to estimate relative factor service prices.[20]

$$\log(c_i^k) = \delta_i + \mu_k + \sum_{f=2}^{F} \alpha_{US,f}^k \tilde{\beta}_{i,f}.$$

Here, δ_i is a country fixed effect, μ_k is an industry fixed effect, and $\tilde{\beta}_{i,f}$ identifies the country-factor relative factor service price differences. In Figures 6 and 7, I plot the relationship between estimated log relative skilled service prices and log GDP per worker, both with country names and with error bars. The results have larger standard errors than the trade based estimates. This reflects the lower number of industries. However, just like the trade based estimates, they exhibit a strong negative correlation with log GDP per worker. The slope parameter of log relative skilled service prices on log GDP per worker is -1.19 using the unit cost data, and -1.53 using the trade data

[20] In Appendix E.3 I derive this regression specification, and provide more details on all measurements.

Figure 6: Skilled price deviation estimates vs log GDP per worker using unit cost data

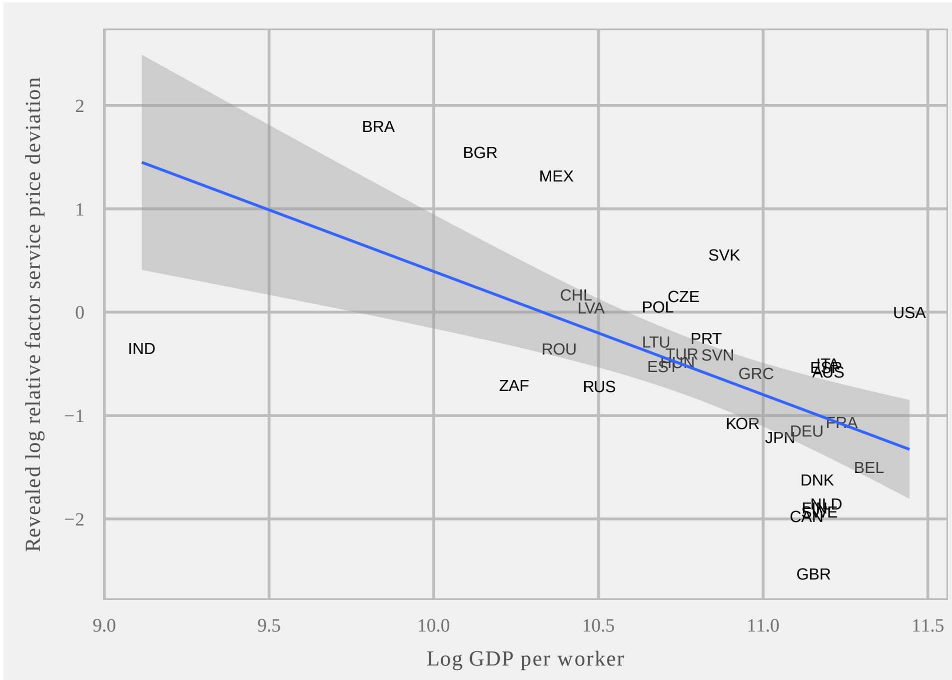

method for the same set of countries. These estimates are similar, and I cannot reject that the two slopes are equal, even when I do not take into account the large standard errors on the unit cost based estimates. Thus, when both types of data exist, the trade data method and the unit cost method paint a similar picture of the relationship between relative skilled service prices and income per worker.

7.3 Further robustness tests of development accounting

In this section, I consider further robustness tests of my development accounting exercise. I analyze how my results change when I exclude the poorest countries and when I exclude oil countries, and I analyze how my results change when I change the measurement of skilled wage premia.

My baseline analysis includes all countries with available trade data, ILO data, and PWT data. Hence, my analysis includes very poor countries and countries with significant oil revenues. Including these countries can be problematic as I use manufacturing trade data to estimate the relative price of skilled services prices. Very poor countries have limited manufacturing trade, and the trade patterns of oil countries is primarily determined by their oil endowment. In Table 9, I show the results when I exclude oil countries and countries with log GDP per worker of less than 9 in 2010 (corresponding approximately to Ghanaian income levels). Excluding these countries considerably expands the role of human capital, and when both sets of countries are excluded, no

Figure 7: Skilled price deviation estimates vs log GDP per worker using unit cost data

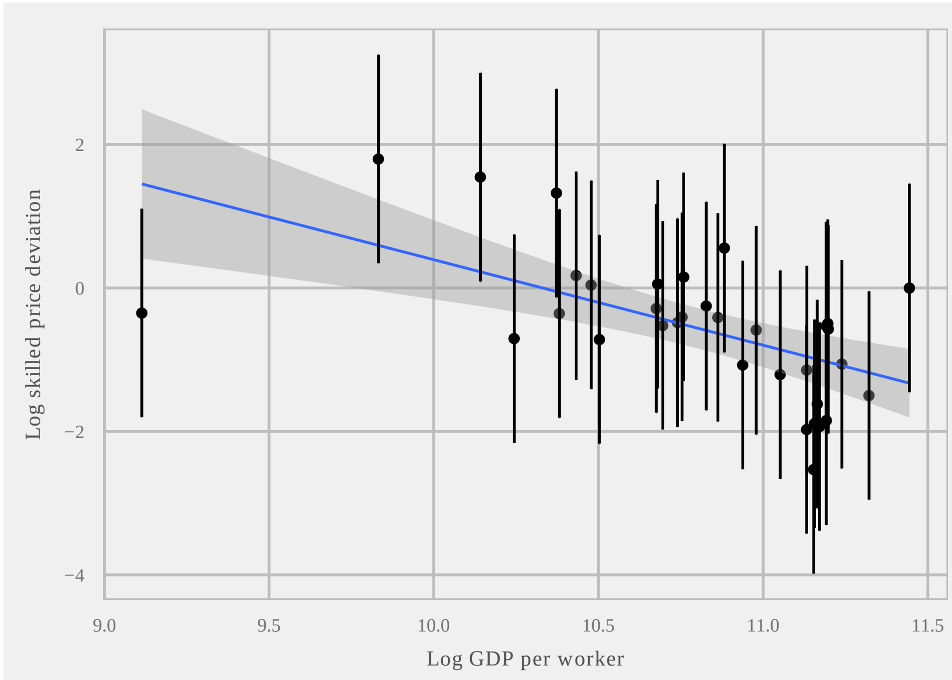

TFP differences are needed to explain income differences among the remaining countries.

I also analyze the robustness of my results to different measurements of skilled wage premia. My skilled wage premia measures are based on limited ILO data, and I want my estimates to be robust to systematic errors in the data on skilled wage premia in poor countries. I am particularly concerned that my measures understate skilled wage premia in poor countries due to the difficulty of measuring wages of self-employed workers and subsistence farmers. My skilled wage premia measures are based on using a linear relation between log GDP per worker and log skilled wage premia. To test how my results depend on skill premia, I consider how my results change if I allow for a steeper relation between country income and log skilled wage premia keeping rich country skilled wage premia constant. I redo my analysis for different values of the income-skill premia slope $\gamma \leq 0$. The results are presented in Table10.

Variations in the posited slope between skilled wage premia and country income have little effect on the estimated importance of human capital. The reason is that two effects counteract each other. Higher skilled wage premia in poor countries reduce the estimated skill-biased quality differences, but they simultaneously reduce the estimated elasticity of substitution between skilled and unskilled workers. These two effects have opposite consequences for the importance of human capital, and they approximately offset each other. Intuitively, there are two cases. If skilled wage premia are very high in poor countries, it suggests that skilled services are difficult to replace, and poor countries are poor because they have few skilled services. If skilled premia are very low in poor countries, large quality differences in human capital are needed to fit the trade data. Again, the conclusion is that human capital is important to account for world income differences.

	Baseline	No v. poor	No oil	Neither
Capital	0.08	0.05	0.08	0.06
Human capital – traditional method	0.10	0.07	0.11	0.10
Human capital – my method	0.63	0.83	0.67	0.92
TFP – traditional method	0.81	0.86	0.79	0.82
TFP – my method	0.27	0.11	0.23	0.01
Log TFP diff – traditional method	2.61	2.77	2.54	2.64
Log TFP diff – my method	0.89	0.35	0.76	0.03
Percent of TFP differences explained	66%	88%	71%	99%
Elasticity of substitution	1.28	1.21	1.28	1.21

Table 9: Contribution of factors and TFP to income differences: different excluded countries

	Baseline (slope =-0.12)	Slope = 0	Slope = -0.2	Slope = -0.4
Capital	0.08	0.08	0.08	0.08
Human capital – traditional method	0.10	0.14	0.10	0.03
Human capital – my method	0.63	0.60	0.63	0.64
TFP – traditional method	0.81	0.77	0.81	0.87
TFP – my method	0.27	0.30	0.28	0.26
Log TFP diff – traditional method	2.61	2.48	2.60	2.81
Log TFP diff – my method	0.89	0.99	0.89	0.86
Percent of TFP differences explained	66%	61%	66%	70%
Elasticity of substitution	1.28	1.38	1.28	1.19

Table 10: Contribution of factors and TFP to income differences: different wage premia

8 Concluding remarks

What share of world income differences can be explained by differences in human capital? The development accounting literature has studied this question by aggregating microeconomic returns to schooling. The overall assessment of the importance of human capital has been negative. Even though there are large human capital differences between countries, they cannot explain more than a small fraction of world income differences.

I have revisited the role of human capital in development accounting, using a framework that allows for imperfect substitutability between skilled and unskilled labor services. I have shown that development accounting is possible in this framework if one can estimate the relative price of skilled and unskilled services, and I have developed a method for estimating this relative price using international trade data. My question has been: does development accounting give us sufficient ground to reject a dominant role for human capital in explaining world income differences?

My results suggest that the answer is no. Using trade data, I find that rich countries have substantially lower relative prices of skilled services. Combining these estimates with data on skilled wage premia suggests that the quality of skilled labor is substantially higher in rich countries. When I include these quality differences in my development accounting exercise, my estimates imply that human capital differences explain 65% of world income differences.

Moving beyond development accounting, there is also a broader takeaway from my results: trade data suggests that there are large productivity differences in skilled labor across rich and poor countries, and that these productivity differences are large enough to explain a dominant share of world income differences. This conclusion holds regardless of whether these productivity differences are due to skill-biased technology differences or skill-biased quality differences.

Thus, my paper supports the conclusions of Caselli and Coleman (2006) and B Jones (2014a), who have argued for large cross-country differences in skilled labor productivity. Their results build on a different method than my results. They note that even though skill premia are somewhat higher in poor countries, skill premia are not as high as they should be given the low relative supply of skilled labor in poor countries, at least not if the elasticity of substitution between skilled and unskilled labor is in line with rich country estimates. Both their papers explain this observation by positing that skilled labor productivity is relatively low in poor countries. Even though they differ in their interpretation of these productivity differences – Caselli and Coleman argue that they reflect skill-biased technology differences and B Jones argues that they reflect skill-biased human capital differences – they agree on the importance of productivity differences in skilled labor.

My findings suggest that their results are not just an artifact of assuming that rich country estimates of substitution elasticities are globally valid. When I analyze trade data, a similar pattern emerges. Relative skilled service prices diverge more sharply between countries than skilled wage premia, suggesting large differences in the productivity of skilled labor. Furthermore, the estimated productivities of skilled workers are strongly and positively correlated with GDP per worker. Quantitatively, productivity differences among skilled workers account for a substantial share of the variation in per capita output. By combining my results with the observations made by Caselli and Coleman (2006) and B Jones (2014), we see how skilled labor productivity differences can provide a unified perspective of the relationship between country income levels, trade patterns, skilled labor supply, and skilled wage premia.

If output differences are primarily driven by productivity differences in skilled labor, this can influence the research agenda of growth and development economics. First, it means that skilled labor human capital differences can drive a large share of output differences, which in turn warrants a greater focus on theories of skill acquisition. Potentially interesting areas include the quality of higher education and the incentives and efficiency of on-the-job learning. Second, if the productivity of skilled labor is driven by skill-specific technology shifters, our technology explanations should put a larger emphasis on why technology differences selectively make skilled labor more productive. This suggests a shift away from general TFP explanations toward more specific theories of technology differences. For example, when we study misallocation, it might be warranted to focus more on how the productivity of skilled labor is harmed by misallocation – potentially by looking at intersectoral patterns of misallocation. Similarly, when studying technology diffusion, it is warranted to study whether barriers to technology diffusion specifically prevents the diffusion of technologies that are complementary to skilled workers.

References

Acemoglu, D. (2007). Equilibrium Bias of Technology. *Econometrica*, 75(5):1371–1409.

Acemoglu, D., Antràs, P., and Helpman, E. (2007). Contracts and Technology Adoption. *American Economic Review*, 97(3):916–943.

Acemoglu, D. and Autor, D. (2011). Chapter 12 Skills, Tasks and Technologies: Implications for Employment and Earnings. In *Handbook of Labor Economics*, volume 4, pages 1043–1171.

Acemoglu, D. and Ventura, J. (2002). The World Income Distribution. *The Quarterly Journal of Economics*, 117(2):659–694.

Acemoglu, D. and Zilibotti, F. (2001). Productivity Differences. *The Quarterly Journal of Economics*, 116(2):563–606.

Anderson, J. E., Anderson, and James (1979). A Theoretical Foundation for the Gravity Equation. *American Economic Review*, 69(1):106–16.

Anderson, J. E. and van Wincoop, E. (2003). Gravity with Gravitas: A Solution to the Border Puzzle. *American Economic Review*, 93(1):170–192.

Barro, R. J. and Lee, J. W. (2013). A new data set of educational attainment in the world, 19502010. *Journal of Development Economics*, 104:184–198.

Bartelsman, E. J. and Gray, W. (1996). The NBER Manufacturing Productivity Database.

Bloom, N., Lemos, R., Sadun, R., Scur, D., and Van Reenen, J. (2016). International Data on Measuring Management Practices. *American Economic Review*, 106(5):152–156.

Broda, C., Weinstein, D., Broda, C., and Weinstein, D. (2004). Globalization and the Gains from Variety. *Quarterly Journal of Economics*, 121:541–585.

Caselli, F. (2005). Chapter 9 Accounting for Cross-Country Income Differences. *Handbook of Economic Growth*, 1:679–741.

Caselli, F. (2015). *Technology Differences Over Space and Time.*

Caselli, F. and Ciccone, A. (2013). The contribution of schooling in development accounting: Results from a nonparametric upper bound. *Journal of Development Economics*, 104:199–211.

Caselli, F. and Coleman, W. J. (2006). The World Technology Frontier. *American Economic Review*, 96(3):499–522.

Chor, D. (2010). Unpacking sources of comparative advantage: A quantitative approach. *Journal of International Economics*, 82(2):152–167.

Ciccone, A. and Papaioannou, E. (2009). Human Capital, the Structure of Production, and Growth. *Review of Economics and Statistics*, 91(1):66–82.

Ciccone, A. and Peri, G. (2005). Long-Run Substitutability Between More and Less Educated Workers: Evidence from U.S. States, 19501990. *Review of Economics and Statistics*, 87(4):652–663.

Costinot, A., Donaldson, D., and Komunjer, I. (2011). What Goods Do Countries Trade? A Quantitative Exploration of Ricardo's Ideas. *The Review of Economic Studies*, 79(2):581–608.

Cuñat, A. and Melitz, M. J. (2012). Volatility, Labor Market Flexibility, and the Pattern of Comparative Advantage. *Journal of the European Economic Association*, 10(2):225–254.

Eaton, J. and Kortum, S. (2002). Technology, Geography, and Trade. *Econometrica*, 70(5):1741–1779.

Erosa, A., Koreshkova, T., and Restuccia, D. (2010). How Important Is Human Capital? A Quantitative Theory Assessment of World Income Inequality. *Review of Economic Studies*, 77(4):1421–1449.

Gaulier, G. and Zignago, S. (2010). BACI: International Trade Database at the Product-Level (the 1994-2007 Version). *SSRN Electronic Journal.*

Gollin, D. (2002). Getting Income Shares Right. *Journal of Political Economy*, 110(2):458–474.

Hall, R. E. and Jones, C. I. (1999). Why do Some Countries Produce So Much More Output Per Worker than Others? *The Quarterly Journal of Economics*, 114(1):83–116.

Hanushek, E. A. and Woessmann, L. (2015). *The Knowledge Capital of Nations: Education and the Economics of Growth*. MIT Press, Cambridge, MA.

Head, K. and Mayer, T. (2014). Gravity Equations: Workhorse,Toolkit, and Cookbook. In *Handbook of International Economics*, volume 4, pages 131–195. Elsevier B.V.

Hendricks, L. (2002). How Important Is Human Capital for Development? Evidence from Immigrant Earnings. *American Economic Review*, 92(1):198–219.

Hendricks, L. and Schoellman, T. (2016). Human Capital and Development Accounting: New Evidence from Wage Gains at Migration.

Hsieh, C.-T. and Klenow, P. J. (2009). Misallocation and Manufacturing TFP in China and India *. *Quarterly Journal of Economics*, 124(4):1403–1448.

Hsieh, C.-T. and Klenow, P. J. (2010a). Development Accounting. *American Economic Journal: Macroeconomics*, 2(1):207–23.

Hsieh, C.-T. and Klenow, P. J. (2010b). Development Accounting. *American Economic Journal: Macroeconomics*, 2(1):207–223.

Inklaar, R. and Timmer, M. P. (2008). GGDC Productivity Level Database: International Comparisons of Output, Inputs and Productivity at the Industry Level. *GGDC Research Memorandum*.

Jones, B. F. (2014a). The Human Capital Stock: A Generalized Approach. *American Economic Review*, 104(11):3752–3777.

Jones, B. F. (2014b). The Knowledge Trap: Human Capital and Development Reconsidered.

Jones, C. I. (2011). Intermediate Goods and Weak Links in the Theory of Economic Development. *American Economic Journal: Macroeconomics*, 3(2):1–28.

Jones, C. I. (2015). The Facts of Economic Growth.

Katz, L. F. and Autor, D. H. (1999). Chapter 26 Changes in the Wage Structure and Earnings Inequality. In *Handbook of Labor Economics*, volume 3, pages 1463–1555.

Katz, L. F. and Murphy, K. M. (1992). Changes in Relative Wages, 1963-1987: Supply and Demand Factors. *The Quarterly Journal of Economics*, 107(1):35–78.

Klenow, P. and Rodriguez-Clare, A. (1997). The Neoclassical Revival in Growth Economics: Has It Gone Too Far? Technical report.

Kremer, M. (1993). The O-Ring Theory of Economic Development. *The Quarterly Journal of Economics*, 108(3):551–75.

Krugman, P. (1988). Differences In Income Elasticities and Trends in Real Exchange Rates. Technical report, National Bureau of Economic Research, Cambridge, MA.

Lagakos, D., Moll, B., Porzio, T., Qian, N., and Schoellman, T. (2016). Life-Cycle Human Capital Accumulation Across Countries: Lessons From U.S. Immigrants.

Levchenko, A. A. and Zhang, J. (2016). The evolution of comparative advantage: Measurement and welfare implications. *Journal of Monetary Economics*, 78:96–111.

Lucas, R. E. (1988). On the mechanics of economic development. *Journal of Monetary Economics*, 22(1):3–42.

Mankiw, N. G., Romer, D., and Weil, D. N. (1990). A Contribution to the Empirics of Economic Growth. Technical report.

Manova, K. (2013). Credit Constraints, Heterogeneous Firms, and International Trade. *The Review of Economic Studies*, 80(2):711–744.

Manuelli, R. E. and Seshadri, A. (2014). Human Capital and the Wealth of Nations . *American Economic Review*, 104(9):2736–2762.

Mattsson, L.-G., Weibull, J. W., and Lindberg, P. O. (2014). Extreme values, invariance and choice probabilities. *Transportation Research Part B: Methodological*, 59:81–95.

Midrigan, V. and Xu, D. Y. (2014). Finance and Misallocation: Evidence from Plant-Level Data . *American Economic Review*, 104(2):422–458.

Nunn, N. (2007). Relationship-Specificity, Incomplete Contracts, and the Pattern of Trade. *The Quarterly Journal of Economics*, 122(2):569–600.

Nunn, N. and Trefler, D. (2015). Chapter 5: Domestic Institutions as a Source of Comparative Advantage. In *Handbook of International Economics*, volume 4, pages 263–315.

Oberfield, E. and Raval, D. (2014). Micro Data and Macro Technology. Technical report, National Bureau of Economic Research, Cambridge, MA.

Parente, S. L. and Prescott, E. C. (1999). Monopoly Rights: A Barrier to Riches. *American Economic Review*, 89(5):1216–1233.

Rajan, R. G. and Zingales, L. (1998). Financial Dependence and Growth. *American Economic Review*, 88(3):559–86.

Rauch, J. E. (1999). Networks versus markets in international trade. *Journal of International Economics*, 48(1):7–35.

Redding, S. and Venables, A. J. (2004). Economic geography and international inequality. *Journal of International Economics*, 62(1):53–82.

Restuccia, D. and Rogerson, R. (2008). Policy distortions and aggregate productivity with heterogeneous establishments. *Review of Economic Dynamics*, 11(4):707–720.

Romalis, J. (2004). Factor Proportions and the Structure of Commodity Trade. *American Economic Review*, 94(1):67–97.

Romalis, J. (2007). NAFTA's and CUSFTA's Impact on International Trade. *http://dx.doi.org/10.1162/rest.89.3.416*.

Schoellman, T. (2011). Education Quality and Development Accounting. *The Review of Economic Studies*, 79(1):388–417.

Schoellman, T. and Hendricks, L. (2016). Human Capital and Development Accounting: New Evidence from Wage Gains at Migration. *2016 Meeting Papers*.

Simonovska, I. and Waugh, M. E. (2014). The elasticity of trade: Estimates and evidence. *Journal of International Economics*, 92(1):34–50.

Soderbery, A. (2015). Estimating import supply and demand elasticities: Analysis and implications. *Journal of International Economics*, 96(1):1–17.

Tinbergen, J. (1962). Shaping the World Economy; Suggestions for an International Economic Policy. *Books (Jan Tinbergen)*.

Trefler, D. (1993). International Factor Price Differences: Leontief was Right! *Journal of Political Economy*, 101(6):961–987.

A Appendix: Environment

A.1 Heterogeneous skill type aggregator interpretation of Q_u and Q_s

Here, I show that my estimation of the relative quality Q_s/Q_u is consistent with a nested structure where the quality terms Q_u and Q_s arise from aggregation of heterogeneous unskilled and skilled services.

My human capital aggregator is

$$h = \left((Q_u u)^{\frac{\eta-1}{\eta}} + a_s (Q_s s)^{\frac{\eta-1}{\eta}}\right)^{\frac{\eta}{\eta-1}}.$$

Before proving the result, I will provide a formal statement of what equivalence means in this context. Assume that the true human capital aggregator is

$$h = \left((H^u)^{\frac{\eta-1}{\eta}} + a_s (H^s)^{\frac{\eta-1}{\eta}}\right)^{\frac{\eta}{\eta-1}}.$$

where H^u and H^s are arbitrary constant returns to scale aggregators of heterogeneous unskilled and skilled services. I say that my relative quality estimation is consistent with an aggregator interpretation if the following holds. Given the definition of quality

$$\begin{aligned} Q_u &\equiv \frac{H^s}{s} \\ Q_s &\equiv \frac{H^u}{u}, \end{aligned}$$

the relative quality of skilled and unskilled labor Q_s/Q_u satisfies the equation

$$\frac{w_s}{w_u} = \frac{Q_s}{Q_u}\frac{r_s}{r_u}, \tag{11}$$

where $\frac{w_s}{w_u}$ is the relative average wage of skilled and unskilled workers, and $\frac{r_s}{r_u}$ satisfies

$$\frac{r_s}{r_u} = a_s \left(\frac{H^s}{H^u}\right)^{-1/\eta}.$$

This quality definition defines the quality of unskilled and skilled labor as the average amount of services provided by each worker in each skill category.

I will now prove the equivalence result. I assume that there are $N_u \geq 1$ types of unskilled labor services and $N_s \geq 1$ types of skilled labor services. A share u_{t_u} of the workforce performs unskilled services of type t_u where $t_u = 1, \ldots, N_u$, and a share s_{t_s} of the workforce performs skilled services of type t_s where $t_s = 1, \ldots, N_s$. The average quality of an unskilled worker of type t_u is Q_{u,t_u} and the average quality of a skilled worker of type t_s is Q_{s,t_s}. The workforce shares sum to the aggregate share of skilled and unskilled workers

$$\begin{aligned}\sum_{t_u=1}^{N_u} u_{t_u} &= u\\ \sum_{t_s=1}^{N_s} s_{t_s} &= s.\end{aligned}$$

With this formulation, the quality of unskilled and skilled labor are defined as

$$\begin{aligned}Q_u &\equiv \frac{H^u(Q_{u,1}u_1,\ldots,Q_{u,N_u}u_{N_u})}{u} = H^u\left(Q_{u,1}\tilde{u}_1,\ldots,Q_{u,N_u}\tilde{u}_{N_u}\right)\\ Q_s &\equiv \frac{H^s(Q_{s,1}s_1,\ldots,Q_{s,N_u}s_{N_u})}{s} = H^s\left(Q_{s,1}\tilde{s}_1,\ldots,Q_{s,N_s}\tilde{s}_{N_s}\right)\end{aligned}$$

where a tilde ($\sim$) denotes that we normalize the unskilled and skilled worker shares u_{t_u} and s_{t_s} with the total supply of unskilled and skilled workers s and u.

Now consider an arbitrary unskilled service type t_u and an aribtrary skilled service type t_s. Assuming that the labor market is competitive, these two types of workers have a relative wage

$$\frac{w_{s,t_s}}{w_{u,t_u}} = \left(\frac{H^s}{H^u}\right)^{-1/\eta}\frac{H^s_{t_s}Q_{s,t_s}}{H^u_{t_u}Q_{u,t_u}} = \frac{r_s}{r_u}\frac{H^s_{t_s}Q_{s,t_s}}{H^u_{t_u}Q_{u,t_u}},$$

where H^s_t and H^u_t denote the partial derivatives of the human capital aggregator functions with respect to their t^{th} elements. The relative wage is a product of i) the relative marginal product of the two aggregators, and ii) the relative marginal contributions of the two skill types to their respective aggregators.

I can use this equation to prove that (11) holds. First, I multiply both sides with $\tilde{s}_{t_s}$ and sum over $t_s = 1,\ldots,N_s$ to obtain

$$\frac{w_s}{w_{u,t_u}} = \frac{r_s}{r_u}\frac{Q^s}{H^u_{t_u}Q_{u,t_u}} \tag{12}$$

where I use Euler's theorem to obtain

$$Q^s = \sum_{t_s=1}^{N_s} Q_{s,t_s}\tilde{s}_{t_s}H^s_{t_s},$$

and use that average skilled wages are defined by

$$w_s = \sum_{t_s=1}^{N_s}\tilde{s}_{t_s}w_s.$$

I obtain equation (11) by applying the same procedure to unskilled labor. I start with equation (12), invert the equation, multiply both sides with $\tilde{u}_{t_u}$, sum over $t_u = 1, \ldots, N_u$, and lastly I re-invert the equation.

This proves that an aggregator interpretation of the quality terms is equivalent to a two labor type interpretation when estimating the relative quality of skilled labor Q_s/Q_u. When doing development accounting, I make one further restriction in assuming that the unskilled aggregator is a linear aggregator. This allows me to estimate Q_u from Mincerian return data, and together with my estimation of Q_s/Q_u, I can complete the development accounting exercise.

A.2 Supply-side aggregation with multiple industries and trade

I express output with an aggregate production function

$$Y = K^{\alpha}(ALh)^{1-\alpha}.$$

When estimating the aggregate production function, I assume that the economy consists of multiple industries and that it trades with the outside world. In light of this, the aggregate production function should be interpreted as reflecting substitution possibilities within and between industries, as well as substitution possibilities between domestic and foreign production. Here, I discuss the assumptions needed to have a constant returns to scale aggregate production function with multiple industries and trade. In Appendix A.3, I motivate my particular choice of functional form.

I show that a CRS aggregate production function exists under fairly general conditions when countries are price takers in the world market. However, there are more stringent conditions for the existence of a CRS aggregator in variety trade models such as Eaton and Kortum and Armington models. In these models, being small compared to the rest of the world is not sufficient to make a country a price-taker, as every country is a large producer of its own varieties. This means that the terms of trade move against countries as they expand factor supplies. Given that my estimation exercise relies on variety models, this is a potential problem.

However, I show that a CRS aggregate production function is possible under a reasonable modification of variety models. The modification is to assume that quality in an Armington model (and absolute productivity advantage in an Eaton and Kortum style model) is homogenous of degree one in aggregate or industry factor supplies. I demonstrate how this modification yields a CRS representation in an Armington model with many small countries, and a similar mechanism applies to the Eaton and Kortum framework.

To motivate my modification, I first argue that the terms of trade effect is unlikely to be a long+run phenomenon. In particular, if such a long-run effect existed, terms of trade would be sensitive to subdivision of countries. For example, if Scotland and UK were formally separated, a long-run terms of trade effect from size would imply that both English and Scottish terms of trade should improve with respect to the rest of the world if they split. This feature is unrealistic, and it suggests that whatever scarce resource makes the global demand curve for a country's goods slope downward – restricted number of varieties in an Armington framework, or restricted idea generation in an Eaton and Kortum framework – this scarce resource should scale with size.[21]

Once I modify the Armington model such that qualities scale with factor supplies, a CRS aggregate production function representation is possible. Furthermore, allowing quality to scale with inputs does not affect the key feature of the model: that relative exports across countries and goods are determined by relative trade costs and relative production costs.

A.2.1 Setup

To study the conditions needed for the existence of a CRS representation, I study a general multi-industry model of a country with K industries and F factor services in an open economy $i \in I$. I use a dual formulation. The production technology in country i for each industry is CRS and represented by the unit cost function $c_i^k(r_{i,1}, \ldots, r_{i,F})$. Factor service supplies are $v_{i,f}$. I write y_i^k for production in industry k and x_i^k for consumption in industry k (these two quantities might differ due to trade). I write p_i^k for the domestic price of good k. There exists a representative consumer whose preferences are defined by an expenditure function $e(\mathbf{p}_i, u_i)$. I assume that these preferences are homothetic, which means that there exists a utility representation of preferences such that the expenditure function can be written

$$e(\mathbf{p}_i, u_i) = \tilde{e}(\mathbf{p}_i)u_i$$

for some function $\tilde{e}$. Throughout this section, I assume that preferences are homothetic and I will write $\tilde{e}$ without a tilde going forward.

A CRS aggregator representation exists if prices are unchanged and output and consumption scale linearly when we scale factor inputs. Formally, I say that a CRS aggregator representation exists if the following condition holds. Let $x_i^k, y_i^k, u_i, r_{i,f}, p_i^k, c_i^k$ be an arbitrary equilibrium given factor supplies $v_{i,f}$. A CRS representation exists if for each such equilibrium, a factor supply $\lambda v_{i,f}$ implies that $\lambda x_i^k, \lambda y_i^k, \lambda u_i, r_{i,f}, p_i^k, c_i^k$ is an equilibrium.

[21] This modification is related to Krugman (1988) who show that growing countries do not face deterioriating terms of trade, and he explains this with a variety model of growth. For a contrasting perspective, see Acemoglu and Ventura (2002) who argue that a country's terms of trade deteriorates when it grows through capital accumulation.

A.2.2 CRS representation when country is price-taker

I first consider a model where each country is a price-taker in the world market. In this case, the equilibrium conditions can be written as:

$$
\begin{aligned}
\sum_{k=1}^{K} \frac{\partial c_i^k}{\partial r_{i,f}} y_i^k &= v_{i,f} \quad f = 1, \ldots, F \\
\frac{\partial e}{\partial p_i^k} u_i &= x_i^k \quad k = 1, \ldots, K \\
c_i^k &\geq p_i^k \quad = 0 \text{ if } y_i^k > 0 \\
e(\mathbf{p}_i) u_i &= \sum_{f=1}^{F} r_{i,f} v_{i,f}
\end{aligned}
$$

The first equation gives clearing conditions for the factor markets, where the left-hand side uses Shepherd's lemma applied to the unit cost function to derive factor demands for each factor f and for industry k. The second equation expresses consumer demand, applying Shepherd's lemma to the expenditure function. The third equation is a zero-profit condition, where the inequality constraint reflects that I allow for zero production. The fourth equation is the budget constraint for the representative consumer.

By inspection, this system of equations allows for a CRS aggregator representation. If there exists a set of prices such that $y_i^k, x_i^k, u_i, v_{i,f}$ solve the system, then any scaling $\lambda y_i^k, \lambda x_i^k, \lambda u_i, \lambda v_{i,f}$ for $\lambda > 0$ solves the system for the same set of prices.

A.2.3 CRS representation with an Armington model

To study the Armington case, I retain the assumption that the country is small in the aggregate world economy. However, the country is large in its own varieties. I represent this with an Armington model with a continuum of countries and K goods. I write $i \in [0, 1]$ for the country on which I focus.

There are K final goods. Each final good is assembled domestically using a composite of country-industry specific intermediate varieties that are traded between countries. To produce good k, one needs an input variety from each country in the world. I assume that there are no trade costs so that the unit cost C_i^k of assembling final good k in country i is the same in every country and equal to

$$C_i^k \equiv C^k = \left(\int_0^1 a_j^k (c_j^k)^{1-\sigma} dj \right)^{\frac{1}{1-\sigma}} \qquad \sigma > 1.$$

I normalize a_j^k so that the unit production costs are $c_j^k = 1$ for all countries $j \neq i$ (our unit of analysis). This means that

$$C^k = 1 \quad k = 1, \ldots, K.$$

Write $q_{i,j}^k$ for the amount of input to industry k that is produced in country i for use in country j. As there are no trading costs and countries are symmetric, $q_{i,j}^k$ does not depend on destination j. Furthermore, using Shepherd's lemma,

$$q_{i,j}^k = \frac{\partial C^k}{\partial c_i^k} x_j^k,$$

where x_j^k is the country j consumption of final goods in industry k.

I can now write down the equilibrium definition.

$$\begin{aligned}
q_{i,j}^k &= a_i^k (c_i^k)^{-\sigma} x_j^k \\
p_{i,k} &= c_{i,k} \\
x_i^k &= \frac{\partial e(1,\ldots,1)}{\partial P^k} u_i \\
\sum_{f=1}^F r_{i,f} v_{i,f} &= e(1,\ldots,1) u_i \\
\sum_{k=1}^K \int_0^1 q_{i,j}^k \frac{\partial c_i^k}{\partial r_{i,f}} &= v_{i,f}
\end{aligned}$$

The first equation gives country j's demand for industry k goods produced in country i. The formulation uses that the price index $P_j^k = C_j^k = 1$ for all j. The second equation is a non-profit condition for production in country i. There is no inequality constraint, reflecting that with a CES specification of production technology from intermediates, production of each variety is always positive. The third equation applies Shepherd's lemma to the consumer's expenditure function. It is evaluated at $(1,\ldots,1)$ as all prices $P^k = 1$. The fourth and fifth equations give the consumer budget constraints and the factor market clearing condition.

By inspection, there does not exist a CRS aggregator representation of this system. In the first equation, we see that scaling output will change prices, violating the assumption that there exist scaled equilibria with the same prices. This reflects a terms of trade effect whereby scaling output depresses the terms of trade.

However, there exists a simple modification of the system to obtain a CRS aggregator. If I define $a_i^k = \Phi_i^k(v_{i,1}^k,\ldots,v_{i,F}^k)$ for some CRS aggregator Φ_i^k, there exists a CRS representation of the equilibrium. Allowing the quality term a_i^k to scale linearly with factor supply captures the intuition that subdivision of observation units should not affect trade patterns with third parties. Even with this modification, *relative* trade patterns across industries are still shaped by *relative* costs, and if we were to add trade costs, then trade costs would affect the distribution between domestic uses and exports, and trade costs would also affect relative exports to different countries.

A.3 Functional form of aggregate production function

My aggregate production function has the form

$$Y = K^{\alpha}(ALh)^{1-\alpha}.$$

As discussed in Appendix A.2, this represents an aggregation taking into account the existence of multiple industries and opportunities for international trade. In this section, I discuss my choice of functional form.

I choose a Cobb-Douglas aggregator between capital and labor services. This is standard in the development accounting literature, and can be motivated by there being constant labor shares across countries (Gollin, 2002).[22]

For the human capital aggregator, I use a CES aggregator of skilled and unskilled labor services, which is standard in the labor economics literature (Acemoglu and Autor, 2011). Ideally, I should have a skill aggregator that was formally aggregated from production functions on the industry level together with a trade model. Unfortunately, there is no straightforward aggregation to a CES representation from industries with heterogeneous factor shares. Thus, the constant elasticity assumption should be interpreted as an approximation to a more freely specified underlying aggregator.

One way of testing my assumption of constant elasticity of substitution is by plotting the cross-country relationship between log relative factor service prices and log relative factor supplies. Theoretically, these should be related by

$$\log\left(\frac{r_s}{r_u}\right) = \log(a_s) - \frac{1}{\eta}\log\left(\frac{Q_s s}{Q_u u}\right).$$

If the CES assumption is true, the relationship should be linear. The test is not ideal, as my estimated relative quality $\log\left(\frac{Q_s}{Q_u}\right)$ is implicitly present in the relative price of factor services, and thus it appears on both sides of the equation, which biases the relationship towards being linear. However, if the log relative supply of skilled and unskilled workers $\log\left(\frac{s}{u}\right)$ was not linearly related to the relative quality $\log\left(\frac{Q_s}{Q_u}\right)$, the relationship would not be linear. Thus, testing the linearity of this relationship offers an opportunity to falsify the CES assumption. The results are plotted in Figure 8, which suggests that the linearity assumption is appropriate.

Looking ahead, potential extensions include modifying the functional form to allow for capital-skill complementarities and non-unitary elasticity of substitution between labor and capital.

[22] Recent studies cast doubt on the Cobb-Douglas assumption (Oberfield and Raval, 2014), and Caselli (2005) suggests that the elasticity of substitution between capital and labor can be a crucial parameter in development accounting. I do not pursue this line of inquiry further here, but it is an interesting avenue of future research. The Cobb-Douglas specification of labor and capital also precludes capital-labor complementarities.

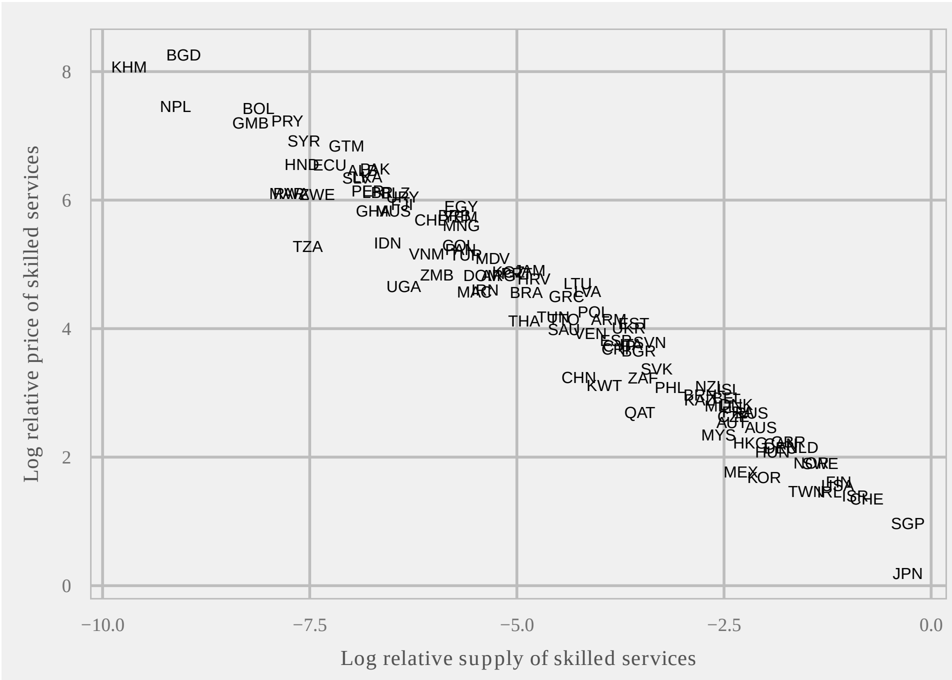

Figure 8: Testing constant elasticity of substitution

B Appendix: Measurement given known relative skill prices

B.1 Occupational vs schooling based skill cutoff

I define the share of unskilled and skilled workers u and s as the shares of people working in an unskilled and skilled occupation, respectively. This contrasts to the approach taken in Caselli and Coleman (2006), B Jones (2014a), and Caselli (2015) who define the share of skilled workers as the share of individuals having an educational attainment above a pre-specified threshold (for example, primary education and above, high school and above, or college and above).

The distinction between the share of workers with a skilled occupation and the share of workers with a certain educational level does not matter if all countries have the same mapping between educational attainment and occupational skill level. However, there is no a priori reason to believe that this mapping should be the same across countries. Acemoglu and Autor (2011) have highlighted the importance of distinguishing between educational attainment and tasks when analyzing US time series data as the allocation of skills to tasks is an equilibrium outcome. Their point is more relevant when analyzing differences between countries with very large differences in educational systems. When educational attainment does not map to occupational skill content in the same way across countries, this modeling choice matters.

I choose an occupational definition for two reasons. First, there are multiple ways of acquiring skills, and education is only one of them. Many people learn skilled occupations outside the educational system, and poor quality of schooling increases the risk that schooling does not fully reflect skill acquisition. When skills are not equal to educational attainment, the complexity of the occupation is a proxy for skill. Indeed, as long as there is a positive skilled wage premium, barring compensating differential concerns, people will work in the most complex occupations that they can perform. Second, occupation is closer to the definitions used for skill shares in my trade data exercise, where I define the skill share as the share of gross output that goes to the payroll of workers in certain occupations.

Thus, I measure the share of skilled workers in line with the ILO's ISCO-08 definitions of skill requirements and major occupational groups. The ILO defines 10 major occupational groups and four skill levels. The occupational groups and their respective skill levels are presented in Figure 9. I use the ILOSTAT database to obtain s as the share of the labor force working as managers, professionals, or technicians and associated technicians, i.e. skill categories 3 and 4 (I define the armed forces as primarily unskilled). I define the unskilled share as $u \equiv 1 - s$.

Figure 10 compares the results from an education based and occupation based definition of the skill share. Figure 10 shows that for poor countries, the share of high school educated workers and the share of skilled workers approximately coincide. For rich countries, there are much more high school educated workers than skilled workers. This is evidence that the mapping between educational attainment and skill level is different in rich and poor countries, and that the educational cutoff for being in a skilled occupation is lower in poor countries.

These results suggest that education based ratios of skilled and unskilled workers will exaggerate rich-poor differences in the relative supply of skilled and unskilled workers. Overall, my method is therefore more conservative when it comes to finding an important role for human capital. I find that this difference matters when I apply the method in B Jones (2014) using my data definitions. He defines a skilled worker as someone having any education above primary education, and finds that even with an elasticity of substitution of 2, human capital is very important in explaining world income differences. With my definition of skilled labor, an elasticity of substitution of 2 means that human capital is only modestly more important than what is found when using traditional development accounting methods.

B.2 Measurement of unskilled labor quality Q_u

I define the quality of unskilled labor Q_u using a Mincerian definition. The quality of unskilled labor is defined as

$$Q_u = \exp(\phi(S_u)).$$

Here, S_u is defined as the average years of schooling of unskilled workers. ϕ is a function capturing the Mincerian returns to education. I use a functional form from Hall and C Jones (1999) and

Figure 9: Mapping of ISCO-08 major groups to skill levels

ISCO-08 major groups	Skill level
1 Managers	3 + 4
2 Professionals	4
3 Technicians and Associate Professionals	3
4 Clerical Support Workers 5 Services and Sales Workers 6 Skilled Agricultural, Forestry and Fishery Workers 7 Craft and Related Trades Workers 8 Plant and Machine Operators, and Assemblers	2
9 Elementary Occupations	1
0 Armed Forces Occupations	1 + 2 + 4

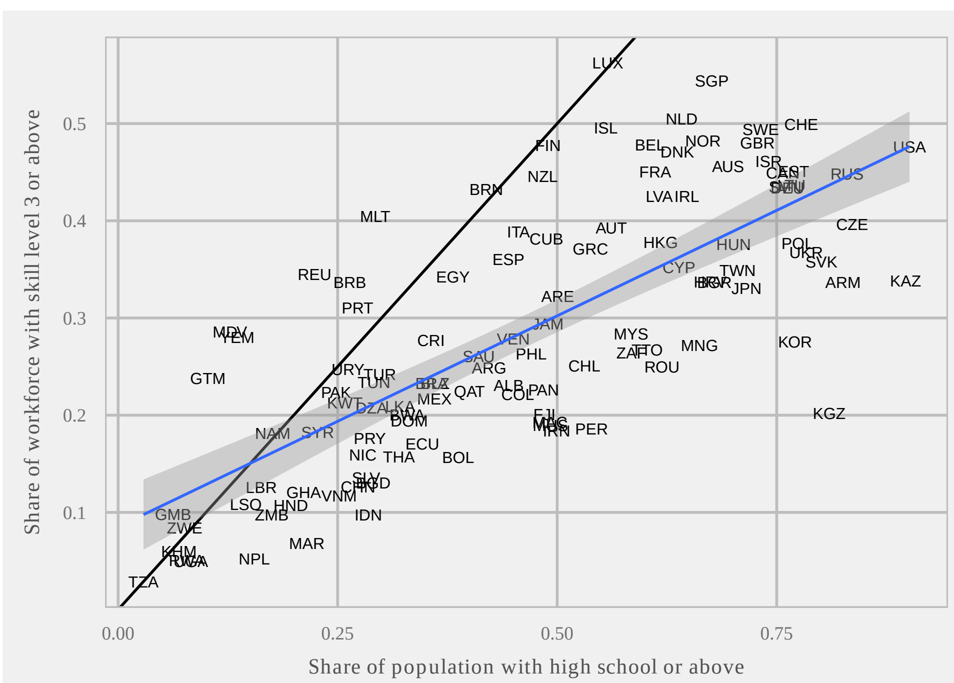

Figure 10: High school and above and share of skilled occupations

Caselli (2005) where $\phi(S)$ is a piecewise linear function with slope 0.13 for $S < 4$, a slope 0.1 for $S \in [4, 8)$, and a slope 0.08 for $S \geq 8$.

I measure S_u by using the data from Barro and Lee (2013). I assume that there is positive sorting between education and skill levels in occupation, and that S_u represents the average years of schooling of the share u of the population working in unskilled occupations. The Barro-Lee data does unfortunately not record the cumulative distribution of years of schooling, but only total schooling attainment within different levels of schooling. It records the number of schooling years at the primary level, the secondary level, and the higher level.

To calculate S_u, I first note that in the vast majority of countries, the cutoff between skilled and unskilled workers goes below the college level, and I attribute none of the schooling years in higher education to unskilled workers. To calculate the years of schooling in primary and secondary school that should be attributed to low skilled workers, I subtract 7 times the share of skilled workers from both primary and secondary school years, using the approximation that all skilled workers have finished high school and that primary and secondary school both are both 7 years. The results are not sensitive to details in this specification. After this subtraction, I divide the remaining primary and secondary school years with the share of unskilled workers to obtain S_u.

C Appendix: Estimating the relative price of skill

C.1 Theoretical derivation of gravity equation

In this section, I show how my gravity specification can be derived from theoretical trade models. Appendix C.1.1 derives the specification from an Armington style trade model, and Appendix C.1.2 derives the specification from an Eaton and Kortum style trade model.

C.1.1 Armington model

There are K industries and I countries, indexed i for source countries and j for destination countries. Each country admits a representative household with preferences

$$U_j = \left(\sum_{i=1}^{I} \sum_{k=1}^{K} (a_j^k)^{1/\sigma} (q_{i,j}^k)^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}} \qquad j = 1, \ldots, I; \ \sigma > 1 \tag{13}$$

where $q_{i,j}^k$ are goods from industry k produced in country i and consumed in country j, σ captures the elasticity of substitution between different varieties, and a_j^k is a country-specific taste term. The taste term is a reduced form way of capturing differences in tastes across countries, including potential non-homotheticities in preferences. The representative consumer maximizes (13) subject to a constraint

$$\sum_{i=1}^{I}\sum_{k=1}^{K} P_{i,j}^{k} q_{i,j}^{k} \leq Y_j$$

where $P_{i,j}^{k}$ is the price of good k produced in country i and bought in country j. Y_j is income in country j.

Each variety is produced using a constant returns to scale production function with the unit cost function

$$c_i^k = C^k\left(r_{i,1}, \ldots, r_{i,F}\right) \tag{14}$$

where $r_{i,f}$ is the price of factor service f in country i.

Trade costs take an iceberg form and to consume one unit of a good from country i, a country j consumer has to buy $d_{i,j} \geq 1$ goods from country i. The cost term $d_{i,j}$ satisfies

$$\begin{aligned} d_{i,j} &\geq 1 \\ d_{i,i} &= 1 \quad \forall i = 1, \ldots, I \\ d_{i,j} d_{j,l} &\geq d_{i,l}. \end{aligned}$$

Output markets are competitive, which implies that prices satisfy

$$P_{i,j}^{k} = c_i^k d_{i,j}. \tag{15}$$

Each country has a supply of factor service flows

$$e_{j,f} \geq 0 \quad i = 1, \ldots, I; \; f = 1, \ldots, F,$$

and country income is given by

$$Y_j = \sum_{f=1}^{F} r_{j,f} e_{j,f} \tag{16}$$

An equilibrium is a set of consumption quantities $q_{i,j}^{k}$, production quantities Q_i^k, factor service prices $r_{i,f}$, unit costs c_i^k, output prices $P_{i,j}^{k}$, and incomes Y_j such that:

1. $\{q_{i,j}^{k}\}$ solves the consumer problem given output prices and incomes.

2. Output market clears
$$Q_i^k = \sum_{j=1}^{I} q_{i,j}^{k} d_{i,j} \forall i, k$$

3. c_i^k and $P_{i,j}^{k}$ satisfy (14) and (15) respectively

4. Income is given by (16)

5. Factor markets clear

$$e_{i,f} \quad = \quad \sum_k Q_i^k \frac{\partial c_i^k}{\partial r_{i,f}}$$

I will not solve the complete equilibrium, but will only solve for the regression specification relating industry export values to unit costs. In the data, export values between i and j in industry k are presented excluding trade costs (FOB). This corresponds to $P_{i,i}^k q_{i,j}^k$, i.e. the domestic price in i of good k produced in i. Using the competitive output market assumption, this quantity is $c_i^k q_{i,j}^k$.

Consumer optimization implies that for any country-industry pairs (i,k), (i',k')

$$\frac{(a_j^k)^{1/\sigma}(q_{i,j}^k)^{-1/\sigma}}{(a_j^{k'})^{1/\sigma}(q_{i',j}^{k'})^{-1/\sigma}} \quad = \quad \frac{P_{i,j}^k}{P_{i',j}^{k'}}$$

$$\sum_{i=1}^{I}\sum_{k=1}^{K} q_{i,j}^k P_{i,j}^k = Y_j$$

Re-arranging the terms gives us

$$P_{i,i}^k q_{i,j}^k = Y_j \frac{a_j^k (P_{i,j}^k)^{1-\sigma}}{\sum_{j',k'} a_j^{k'} (P_{i,j'}^{k'})^{1-\sigma}} \frac{P_{i,i}^k}{P_{i,j}^k}.$$

Taking logarithms, writing total exports $x_{i,j}^k = P_{i,i}^k q_{i,j}^k$, and substituting in (14) for prices gives me

$$\log(x_{i,j}^k) = \delta_{i,j} + \mu_j^k - (\sigma - 1)\log(c_i^k) \tag{17}$$

where

$$\delta_{i,j} \quad = \quad \log(Y_j) - \log\left(\sum_{i',k'} a_j^{k'} (c_{i'}^{k'} d_{i',j})^{1-\sigma}\right) - \log(d_{i,j})$$

$$\mu_j^k \quad = \quad \log(a_j^k).$$

Here, $\delta_{i,j}$ captures all terms that only depend on the bilateral relationship: the income of the buying country, the market access term of the buying country, and all bilateral trading costs between the two countries. μ_j^k captures industry-specific demand effects in the buying country.

C.1.2 Eaton and Kortum model

To derive an industry based gravity equation using an Eaton and Kortum framework, I construct a model close to Chor (2010), who analyzed industry-level trade in an Eaton and Kortum setup. There are I countries where i is an index for a source country and j is an index for a destination country. The model has K goods which are produced domestically, and the production of each good k uses a range of internationally traded intermediate good varieties.

Each country has a representative consumer with preferences

$$U_j = \left(\sum_{k=1}^{K} a_j^k (Q_j^k)^{\frac{\xi-1}{\xi}} \right)^{\frac{\xi}{\xi-1}} \quad \xi > 1.$$

Each final good k is a composite of internationally traded varieties $q_i^k(z)$ with $m \in [0,1]$. The price of final good k in country i is

$$P_j^k = \left(\int_0^1 p_j^k(m)^{1-\eta} dm \right)^{\frac{1}{1-\eta}}, \quad \eta > \xi > 1,$$

where $p_j^k(m)$ is the country j price of variety m in industry k. The assumption on the elasticity of substitution means that different varieties are more substitutable than goods from different industries.

As varieties are internationally traded, the price $p_j^k(m)$ paid for a variety will reflect the cheapest available variety for country j. When I specify the cost function for varieties, I am therefore interested in the unit cost of *offered* varieties from country i to country j, which I write $p_{i,j}^k(m)$. The price $p_j^k(m)$ is obtained by minimizing over potential source countries i.

The offered price $p_{i,j}^k(m)$ will depend on a deterministic component of costs in country i and industry k, on trade costs between country i and j, and on a stochastic productivity shock to this particular variety. The deterministic component of costs is

$$c_i^k = C^k(r_{i,1}, \ldots, r_{i,F}) \tag{18}$$

where $r_{i,f}$ denotes the factor service price of factor f in country i. Trade costs take an iceberg form and to obtain one unit of an intermediate good from country i, a country j producer has to buy $d_{i,j} \geq 1$ intermediate goods from country i. The cost term $d_{i,j}$ satisfies

$$\begin{aligned} d_{i,j} &\geq 1 \\ d_{i,i} &= 1 \quad \forall i = 1, \ldots, I \\ d_{i,j} d_{j,l} &\geq d_{i,l}. \end{aligned}$$

The offered price is

$$p_{i,j}^k(m) = \frac{c_i^k d_{i,j}}{z_i^k(m)} \tag{19}$$

where $z_i^k(m) \sim Frechet(\theta)$ is a country-industry-variety specific productivity shock which is Frechét distributed with a parameter θ. A random variable Z is Frechét-distributed with parameter θ if

$$P(Z \leq z) = e^{-z^{-\theta}}.$$

I will not solve a full equilibrium for this model, but only derive the gravity trade equation that results from the model. For each variety m in industry k, country j obtains an offer $p^k_{i,j}(m)$ from each country i given by equation (19). The probability distribution of this offer is

$$\begin{aligned}
P(p^k_{i,j}(m) \leq p) &= P\left(\frac{c^k_i d_{i,j}}{p} \leq z^k_i(m)\right) \\
&= 1 - e^{-\left(\frac{c^k_i d_{i,j}}{p}\right)^{-\theta}} \\
&= 1 - e^{-\left(c^k_i d_{i,j}\right)^{-\theta} p^\theta}
\end{aligned}$$

The best price $p^k_i(m)$ for country i is the minimum of all offers $\min_i p^k_{i,j}(m)$ and has distribution

$$\begin{aligned}
G(p) &= P\left(\min_i p^k_{i,j}(m) \leq p\right) \\
&= 1 - P(\max_i p^k_{i,j}(m) > p) \\
&= 1 - \prod_i P(p^k_{i,j}(m) > p) \\
&= 1 - \prod_i (1 - P(p^k_{i,j}(m) \leq p) \\
&= 1 - e^{-\sum_i \left(c^k_i d_{i,j}\right)^{-\theta} p^\theta}
\end{aligned}$$

I write

$$\Phi^k_j = \sum_i \left(c^k_i d_{i,j}\right)^{-\theta}. \tag{20}$$

This expression summarizes country j's access to industry k. It is decreasing in production costs in industry k and in the bilateral trading costs $d_{i,j}$.

Country j chooses to buy a variety from the country with the lowest price. The probability that country i offers the lowest price is

$$\begin{aligned}
\pi^k_{i,j} &\equiv P(p^k_{i,j}(z) \leq \min_i p^k_{i,j}(z)) \\
&= \frac{(c^k_i d_{i,j})^{-\theta}}{\Phi^k_j}.
\end{aligned}$$

If x^k_j is the total amount of intermediate inputs bought by country j in industry k, the trade flow matrix is

$$x_{i,j}^k = \pi_{i,j}^k x_j^k = \frac{(c_i^k d_{i,j})^{-\theta}}{\Phi_j^k} x_j^k \tag{21}$$

Equation (21) requires that the share of import value coming from country i only depends on the share of inputs for which i is the supplier. This property holds as the Frechet distribution has a desirable property called max-stability, which ensures that the best offered price $p_{i,k}(z)$ to country i is independent of the source of the best offer (see Eaton and Kortum (2002) for a derivation in this particular case, and Mattsson et al. (2014) for a more general discussion of this property of random variables) . This means that the total expenditure on imports from one country will be fully determined by the share of varieties $\pi_{n,i}^k$ bought from that country. The reason is that all countries offer identical distributions of variety prices conditioned on them offering the best prices.

Taking the logarithm of both sides of equation (21) gives me

$$\log(x_{i,j}^k) = \delta_{i,j} + \mu_j^k - \theta \log(c_i^k)$$

where $\delta_{i,j} = -\theta \log(d_{i,j})$ and $\mu_j^k = \log(X_j^k) - \log(\Phi_j^k)$. Thus, the model implies a gravity equation of the right form. Note that when using Eaton and Kortum elasticity estimates θ, there needs to be added a 1 to convert them to the corresponding Armington elasticity estimates σ.

C.2 Results for other factors than skilled labor

In Section 3.1, I estimated regression (6) to obtain estimates of relative factor service prices across countries. My main interest was in the relative price of skilled services, as this relative price is used directly in development accounting. However, my estimation procedure also yields relative factor service price estimates for capital, intermediate inputs, and energy. Even though I do not use these directly in my development accounting exercising, they are useful to check the plausibility of my factor service price estimation method.

In particular, as capital, intermediate inputs, and energy are partly tradable, we should expect the relative price of these factors compared to unskilled labor to fall with GDP per worker. The reason is that tradable services should have similar prices across countries, whereas we expect the price of unskilled labor services to rise with GDP per worker.

It is possible to quantify how much unskilled service prices should fall with GDP. If we assume that the labor share of output is constant at $1-\alpha$, the unskilled wage satisfies equation

$$\begin{aligned} w_u &= \frac{w_u}{w_u u + w_s s} \times (w_u u + w_s s) \\ &= \frac{1}{u + \frac{w_s}{w_u} s} (1-\alpha) y \end{aligned}$$

where y in the second line denotes output per worker. Using that the price of unskilled labor services is $r_u = w_u/Q_u$ where Q_u is the quality of unskilled workers, I obtain

$$\log(r_u) = \log(1-\alpha) + \log(y) - \log(h_{trad})$$

where $\log(h_{trad}) = \log(Q_u) + \log(u + \frac{w_s}{w_u}s)$ is human capital according to traditional development accounting methods, as defined in equation (10). Letting r_t be the price of any tradable input service, its relative price compared to unskilled labor services will be

$$\log\left(\frac{r_t}{r_u}\right) = \log(r_t) - \log(1-\alpha) - \log(y) + \log(h_{trad}).$$

If $\log(r_t)$ is constant across countries, we can make the following observation: constant $\log(h_{trad})$ across countries implies that relative tradable factor prices decrease one-to-one with GDP per capita. If $\log(h_{trad})$ is positively correlated with GDP, relative tradable factor service prices will fall slower than one-for-one. And even though it is not explicitly modeled in the equation, we can also note that a non-tradable component of t will also make the relative price/GDP-slope less negative.

In my data, $\log(h_{trad})$ increases at approximately $0.15 - 0.2$ with GDP per capita. Thus, if capital, intermediate inputs, and energy services are fully tradable, they should have a negative slope of between 0.8 and 0.85 with respect to GDP per worker. If they are not fully tradable, the negative relationship should be weaker. The results are presented in Figures 11-13. The negative relationship between capital and intermediate input service prices and log GDP per worker is similar at -0.6, which is close to what is predicted by my previous reasoning. The conclusions are less stable for energy prices. Here, there is also a negative relationship, but the data is less precise. This is due to energy having a very small factor share in most industries, and the results for energy are more driven by outliers. Reassuringly, large energy producers such as Saudi Arabia, Kuwait, Russia, and Iran have low revealed energy service prices.

C.3 Treatment of intermediate inputs

C.3.1 Baseline specification

In my main specification, I include the cost share of intermediate inputs $\alpha^k_{US,int}$. The corresponding estimate $\beta_{i,int}$ identifies $\log\left(\frac{r_{i,int}/r_{i,1}}{r_{US,int}/r_{US,1}}\right)$. This estimate gives the difference between the US and country i in the relative cost of intermediate input and unskilled labor services.

In my interpretation of this parameter, I assume that intermediate inputs are traded. I interpret $r_{i,int}$ as a product of an international price of intermediate inputs r_{int}, which is constant across countries, and a country-specific barrier to international intermediate input markets τ_i, which varies across countries.

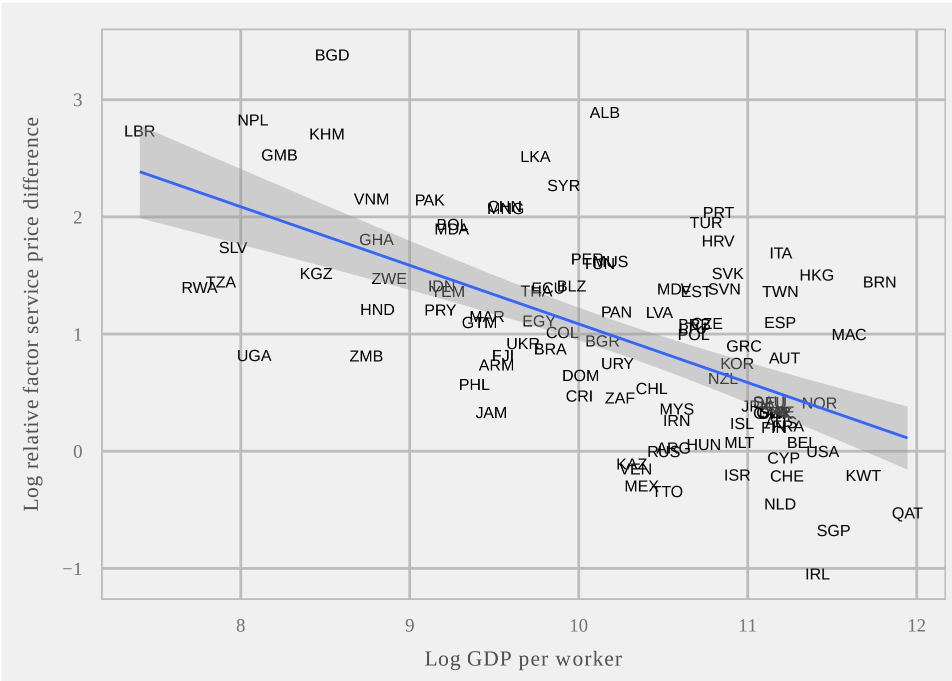

Figure 11: Log relative capital services prices and log GDP per worker

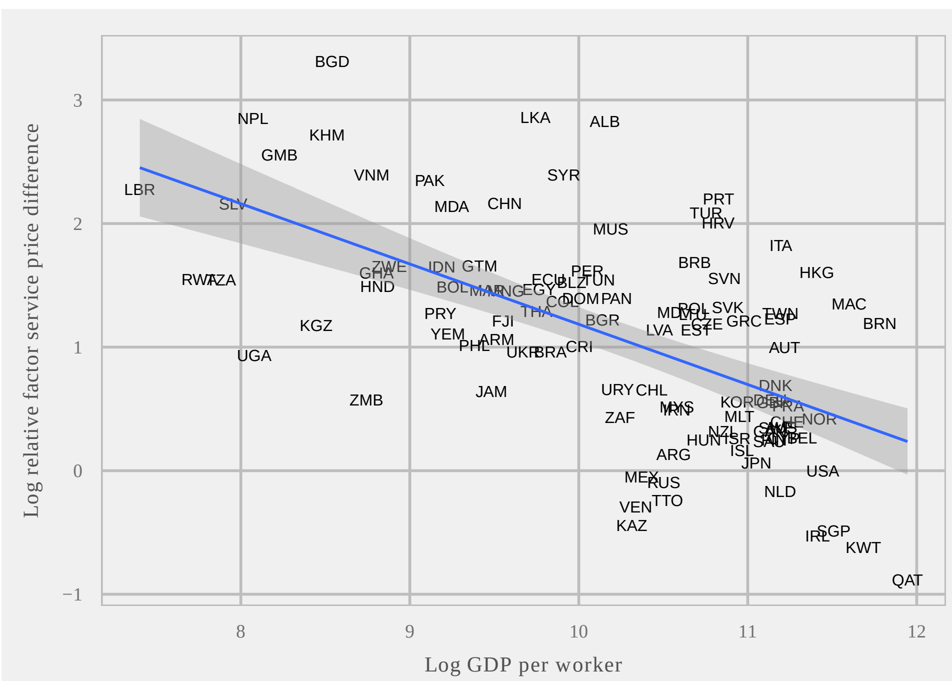

Figure 12: Log relative intermediate input services prices and log GDP per worker

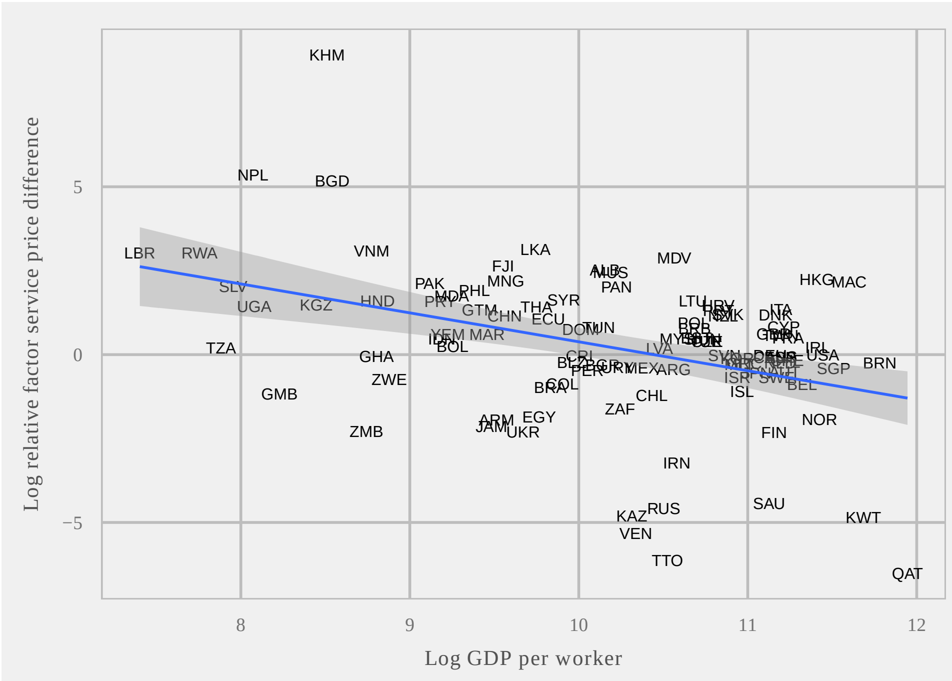

Figure 13: Log relative energy services prices and log GDP per worker

With this interpretation,

$$\beta_{i,int} = \log(\tau_i/\tau_{US}) - \log\left(\frac{r_{i,1}}{r_{US,1}}\right).$$

$\beta_{i,int}$ varies across countries for two reasons. First, countries differ in their access to international intermediate goods markets τ_i. Bad access to international markets (high τ_i) gives a high revealed price of intermediate input services (high $\beta_{i,int}$). Second, countries differ in their prices of unskilled labor services $\log\left(\frac{r_{i,1}}{r_{US,1}}\right)$. Countries with a low price of unskilled services have a high revealed price of intermediate input services. This has an intuitive interpretation: relatively inexpensive unskilled labor services make internationally traded intermediate inputs relatively expensive.

C.3.2 Robustness to non-traded intermediate inputs

If intermediate inputs are not traded and the aim is to identify factor service price differences, a different approach is called for. In this case, there is an indirect effect of factor service price differences via input prices. To reflect this, the intermediate input share in an industry k should be resolved into contributions from different factor services, using the input-output structure to determine the factor shares of industry k's intermediate inputs.

To check the robustness of my baseline specification, I develop an approach that allows for both traded and non-traded intermediate inputs. To implement my approach, I use the US input-output table and assume that services are non-traded and that other goods are traded.[23] I use the EU-KLEMS data together with Occupational Employment Survey data to obtain factor shares in service sectors.

I write N_T for the number of traded goods and N_{NT} for the number of non-traded goods. The input-output table $L$is an $(N_T + N_{NT}) \times (N_T + N_{NT})$ matrix. For each good $k = 1, \ldots, N_T + N_{NT}$, I measure its factor shares including its intermediate input share, and I use these measured factor shares to define the *first-stage* factor shares $\tilde{\alpha}_f^k$:

$$\tilde{\alpha}_f^k = \begin{cases} \text{measured factor share} & \text{if } f \neq \text{intermediate inputs} \\ \text{measured intermediate share} \times \text{share of tradeable intermediates} & \text{if } f = \text{intermediate inputs} \end{cases} .$$

This expression defines the first-stage factor shares $\tilde{\alpha}_f^k$. In the first stage, I am interested in the cost shares of different factors and of tradable inputs. For each industry, $1 - \sum_{f=1}^{F} \tilde{\alpha}_f^k$ gives the share of costs in industry k going to nontraded factor inputs. These first-stage factor shares are the building blocks of the factor shares α_f^k that will be obtained by resolving the cost share of nontraded intermediate inputs into conventional factors and tradable inputs.

I find the factor shares α_f^k of tradable goods recursively by first finding the factor shares of nontradable goods. I define two matrices L_T and L_{NT} where L_T is an $N_T \times N_{NT}$ matrix giving the input uses of nontraded intermediate inputs in the traded sector, and L_{NT} is an $N_{NT} \times N_{NT}$ matrix giving the cost shares from nontraded inputs in the nontraded sector.

I solve the system recursively. The factor shares of nontraded goods are

$$\alpha_{NT} = \tilde{\alpha}_{NT} + (L_{NT})\alpha_{NT} \Longleftrightarrow \alpha_{NT} = (I - L_{NT})^{-1}\tilde{\alpha}_{NT}$$

where α_{NT} is an $N_{NT} \times F$ matrix, $\tilde{\alpha}_{NT}$ is an $N_{NT} \times F$ matrix, and L_{NT} is an $N_{NT} \times N_{NT}$ matrix. The final matrix α_{NT} gives the factor shares of nontraded services in terms of standard factor shares and traded input shares. All nontraded input shares have been resolved into these constituent parts. Having solved for the factor shares of nontraded goods, the factor shares of traded goods are

$$\alpha_T = \tilde{\alpha}_T + (L_T)\alpha_{NT}.$$

Using this modified definition of factor shares, I can re-estimate my baseline specification. In Figure 14, I compare the estimates for the estimated skilled service coefficient to my baseline estimation. The new results are very similar to my baseline estimates. The reason is that even though resolving the nontraded factors increases the skilled share in all industries (as I move the

[23]There is moderate trade in some services such as entertainment, financial services, and transportation, but the distinction captures the large differences in traded shares between services and other goods in the US input-output table.

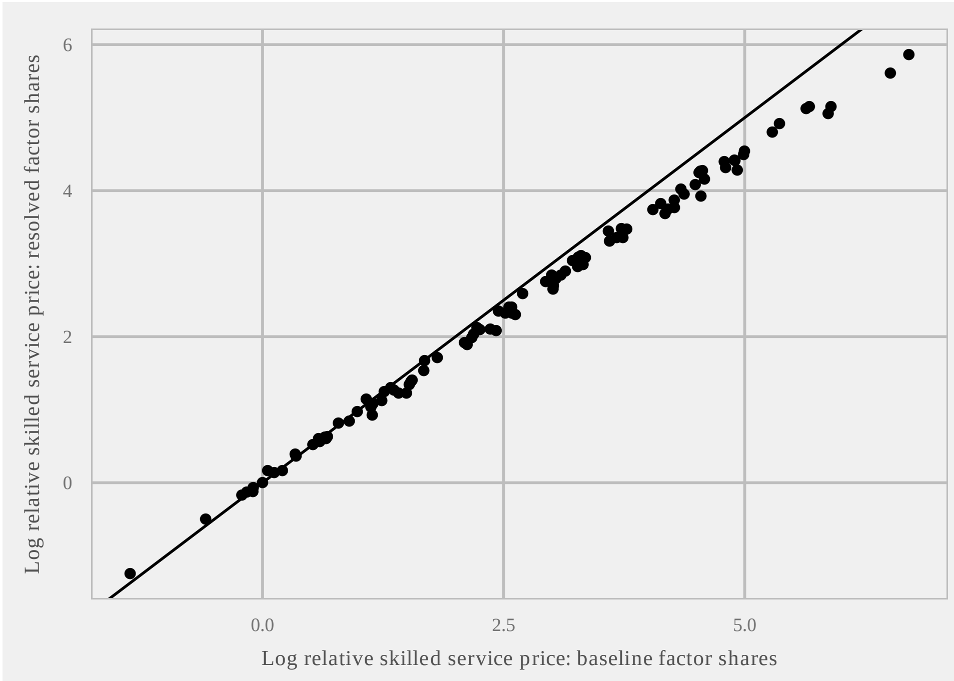

Figure 14: Comparison of estimated relative skilled service prices with different input measurements

skilled component of inputs from the intermediate input share to the skill share), the resolving of nontraded factors does little to alter the *relative* skill shares across industries, which are the bases of my estimation.

D Appendix: Mechanisms

D.1 Interpretation of migration data

In this section, I analyze the relationship between my results and data on migrant wages. Migrant wage data has been an important source of information in development accounting since Hendricks (2002).[24] Ideally, migration provides a natural experiment to distinguish between human capital based and technology based explanations of world income differences. If selection issues are appropriately addressed, migration data allows us to compare similar workers in two different environments. Human capital is kept constant, and wage differences have to be attributed to some other factor in the environment. Under some conditions, this other factor can be interpreted as technology.

In particular, migrant wage data has been used to argue against a dominant role for human capital in accounting for world income differences. This was the main argument in Hendricks

[24]In addition to Hendricks (2002), papers that use migrant data include Schoellman (2011), Lagakos et al. (2016), and Hendricks and Schoellman (2016).

(2002). He showed that migrants from poor countries in the US had dramatically higher wages than workers in their native countries. He argued that this was inconsistent with human capital differences being large enough to explain world income differences. Even though later contributions have tempered this conclusion by using individual data to account for selection (Hendricks and Schoellman, 2016), it remains important for human capital based explanations of world income differences to be consistent with migrant wage data.

Given that I argue for a dominant role for human capital in explaining world income differences, it is natural to ask how my results relate to migrant wage data. In this section, I show that my results are consistent with existing evidence from migrant wage data. The key difference between my analysis and that of Hendricks (2002) is that I relax the assumption of perfect substitutability and allow for imperfect substitutability between labor services. This leads to a different interpretation of migrant wage data. In particular, imperfectly substitutable labor services imply that human capital is multidimensional and that there is no longer a simple mapping from human capital to pre- and post-migration wages. A worker's wage is the product of the amount of labor services that the worker provides, and the price of those labor services. Even though wage changes at migration can be due to technology differences, they can also be due to labor service price differences.

I analyze the implications of my results for migrant wage data, and I discuss the implications both for unskilled and skilled migrants. When migrants are unskilled workers, my development accounting results imply that there are limited human capital quality differences between rich and poor countries. I am interested in whether these limited quality differences are consistent with large wage gains for unskilled migrants going from poor to rich countries. When migrants are skilled workers, my development accounting results suggest that there are substantial quality differences between rich and poor countries. In this situation, I am interested in whether these large estimated quality differences necessarily mean that skilled migrants going from rich to poor countries should have much higher wages than local workers, and conversely if skilled migrants going from poor to rich countries necessarily should have much lower wages than local workers.

Starting with unskilled migrants, I note that it is consistent with my results that unskilled workers going from poor to rich countries experience large wage gains. The mechanism is that in rich countries, the high relative supply and quality of skilled workers increase the relative price of unskilled labor services. This relative scarcity of unskilled labor services in rich countries makes unskilled wages higher. Wage gains for unskilled migrants are thus consistent with similar quality of unskilled labor across rich and poor countries, and these wage gains do not rely on large cross-country differences in technology.

For skilled migrants, I begin by analyzing skilled migrants going from rich to poor countries. According to my estimates, these skilled migrants have a much higher quality of human capital than their local counterparts. Does this imply that they will necessarily get much higher wages than local skilled workers? The answer is no. The reason is that high US quality of skilled labor is the result of an aggregation of heterogeneous skilled services. To make sharp predictions of how wages change at migration for a worker with particular skills, we need to know the complementarity and substitutability patterns implicit in the aggregator of skilled services. More concretely, the question is not whether a standard US engineer migrating to Tanzania gets a much higher salary than the average Tanzanian engineer, but whether a US hydraulic engineer specializing in sediment transportation migrating to Tanzania gets a much higher salary than an average Tanzanian engineer. The latter question is not possible to answer without knowing the details of the skilled service aggregator.[25]

Conversely, we can analyze what my results predict about skilled migrants going from poor to rich countries. According to my results, poor countries have a substantially lower quality of skilled labor than rich countries. Do my results predict that a skilled migrant going from a poor country to a rich country necessarily should have a much lower salary than local skilled workers? The answer again is no. To begin with, the argument about complementarity and substitutability patterns that I made concerning skilled workers migrating from rich to poor countries still applies to this situation. Furthermore, even if we neglect potential heterogeneity among skilled and unskilled services, my explanation is still consistent with skilled migrants to rich countries not having dramatically lower wages than local counterparts. The reason is that there is potential for occupational switching at migration. Indeed, note that my results suggest that the relative price of skilled services is lower in rich countries than in poor countries. Thus, a worker that has a comparative advantage in a skilled occupation in a poor country might have a comparative advantage in a low skill occupation in a rich country. For example, a moderately competent computer programmer from a poor country might find it profitable to work in an unskilled profession in the US. If the scarcity of unskilled services in the US has driven up unskilled wages, this is consistent with skilled migrants to rich countries only having moderately lower wages than their local skilled counterparts, compared to the large estimated quality differences in skilled labor. Even though B Jones (2014b) discusses the potential importance of occupational switching for migrant wage data, there has not been any full empirical examination of this mechanism. However, B Jones (2014a) and Hendricks and Schoellman (2016) provide suggestive evidence that occupational downgrading is more common for migrants from poor countries.

In conclusion, my development accounting results are consistent with migrant wage data. This consistency is not due to migrant wage data confirming sharp predictions derived from my results.

[25]Here, I use a low-dimensional representation of labor force heterogeneity to analyze cross-country differences, and a high-dimensional representation of labor force heterogeneity to analyze migration data. This procedure is analogous to the treatment of capital in aggregative growth models. The Solow model and the neoclassical growth models use a one-dimensional representation of capital, and these models are appropriate for capturing broad patterns of growth, output and marginal returns to capital. However, capital aggregation hides an underlying heterogeneity. This means that model predictions from these models are commonly not tested by comparing cross-country differences in rental prices of specialized pieces of equipment. Such comparisons are outside the domain of validity of the aggregate model setup. The same applies to my setup.

Apart from predicting wage gains for unskilled migrants going to rich countries, my results put weak restrictions on migrant wage data. Given the natural experiment aspect of migration that makes migrant wage data an attractive source of information about human capital differences, an important avenue of future work is to place restrictions on my setup to derive sharper predictions for migrant wage data.[26]

D.2 Endogenous skill-biased technology differences and human capital quality

Even though SBTD and quality differences are observationally equivalent with respect to price and quantity data, they are not equivalent in general. If quality differences in skilled labor explain why countries are rich, development theory needs to explain why countries differ in their quality of skilled labor. If SBTD explain why countries are rich, development theory needs to explain how similar qualities of skilled labor can result in very different levels of skilled labor productivity.

Thus, I try to move beyond price and quantity data to gauge the relative merits of SBTD and quality difference interpretations of my estimates.[27] To this end, I examine whether SBTD reduce estimated quality differences when I put theoretical structure on how technology varies across countries. In particular, I analyze a standard mechanism from the literature where SBTD arise endogenously in response to relative factor service price differences (Caselli and Coleman, 2006; Acemoglu, 2007; Caselli, 2015). In Appendix D.3, I test in general for estimation errors arising from second order erros and endogenous technology differences. There, I provide a detailed description of the environment, my measurement procedure, and my results. Below is a summary.

I set up a simple model of endogenous technological choice in line with Caselli and Coleman (2006) and Acemoglu et al. (2007). Technological bias varies on an industry-country level as a function of factor service prices. For each set of relative factor service prices, I generate unit cost data from the model and run my baseline regression specification on the model generated data. I find the relative factor service prices such that my regression specification gives the same results when applied to model data as when applied to actual data. This gives me the relative service factor prices that are consistent with my regression estimates given that there is endogenous SBTD. By comparing these relative factor service prices with those found under my baseline assumptions, I can test whether my results in Section 3.1 overstates rich-poor quality differences in skilled labor.

The results are mixed but there is no overall tendency for the endogenous SBTD based model to imply lower quality difference than my baseline setup. In many cases, estimated quality differences are actually higher when I allow for SBTD. The exact results depend on parameters and the effects are non-monotone in the size of relative skilled service price differences.

[26] In an unpublished paper, B Jones (2014b) discusses migrant wage data with imperfect substitutability. He also argues for the importance of relative scarcity in accounting for the wages of unskilled migrants, and for occupational switching in accounting for the wages of skilled migrants. The points about complementarity and substitutability patterns are to my knowledge original to this paper.

[27] In Section 5.2 I did this by discussing circumstantial evidence for quality differences and economic mechanisms that could make them large.

These results can be viewed as somewhat surprising: if there are SBTD, a reasonable expectation is that they would lower the need for price differences in skilled labor services to explain the trade data. One mechanism that helps explain my findings is that there are to two opposing tendencies. SBTD reduce the need for quality differences as they increase the relative productivity of skilled labor. However, when SBTD are endogenous, they also increase the effective elasticity of substitution between skilled and unskilled labor services. In the context of my estimation procedure, this can sometimes mean that there are larger quality differences to explain away. The net effect is ambiguous.

Thus, accounting for SBTD becomes complex when SBTD are endogenous. I have not resolved all issues, and a more thorough investigation of endogenous SBTD in my context is an important avenue for future research. However, in the case when SBTD arise endogenously from relative factor service price differences, they do not unambiguously obviate the need for large quality differences to explain my estimates.

D.3 Robustness to industry function specification and endogenous technology bias

My baseline estimates relied on the assumption that it was possible to approximate unit cost differences from the US by log-linearizing around the US cost structure. In terms of assumptions on industry production functions, this assumption amounts to assuming that industry production functions are Cobb-Douglas. Furthermore, to interpret estimates r_s/r_u in terms of human capital, I needed to assume that there were no skill-biased technology differences between countries. This section tests the robustness of my results to deviations from these two assumptions.

The section has three subsections. In Appendix D.3.1, I describe an environment featuring CES industry production functions, and endogenous technology bias in response to relative factor service price variations along the lines of Caselli and Coleman (2006), Acemoglu et al. (2007), and Caselli (2015). In Appendix D.3.2, I show how it is possible to quantify the extent of bias introduced by varying production function assumptions. Appendix D.3.3 describes the results of the quantification exercise.

D.3.1 Environment

I assume that industry cost functions satisfy

$$c_i^k\left(\frac{r_{i,1}}{Z_{i,1}^k}, \ldots, \frac{r_{i,F}}{Z_{i,F}^K}\right) = \left(\sum_{f=1}^{F} a_f^k \left(\frac{r_{i,f}}{Z_{i,f}^k}\right)^{1-\xi}\right)^{\frac{1}{1-\xi}} \qquad \xi > 0,$$

where $r_{i,f}$ is the factor service price of factor f in country i, a_f^k is the factor share of factor f in industry k, $\xi > 0$ is the elasticity of substitution, and $Z_{i,f}^k$ is a factor-augmenting technology term.

The technology terms vary endogenously across countries in response to changes in relative factor prices. In modeling this choice, I follow Acemoglu (2007) and assume that there exists a cost function $G^k(Z^k_{i,1}, \ldots, Z^k_{i,F})$ capturing the cost of acquiring a technology bundle. I assume that G^k is convex and homogenous of degree $\gamma > 1$. A country's technology bundle in an industry is the solution to

$$\tilde{c}^k_i = \min_{\{Z^k_{i,1},\ldots,Z^k_{i,F}\}} \left\{ c\left(\frac{r_{i,1}}{Z^k_{i,1}}, \ldots, \frac{r_{i,F}}{Z^k_{i,F}}\right) + \frac{P^k_i G^k(Z^k_{i,1}, \ldots, Z^k_{i,F})}{\bar{Z}_i} \right\} \tag{22}$$

where $\tilde{c}^k_i$ is the unit cost of good k in country i taking into account technology acquisition costs, $\frac{1}{\bar{Z}_i}$ is a country specific technology diffusion barrier, and P^k_i is the price of good k in country i. In equilibrium, $P^k_i = \tilde{c}^k_i$.

This specification aims at capturing a mechanism higlighted in the literature: the possiblity of endogenous technology bias in response to variations in relative factor service prices (Caselli and Coleman, 2006; Acemoglu, 2007; Caselli, 2015). Even though other mechanisms might be active, I have chosen a model specification that allows me to focus on this particular mechanism, and exclude other potential mechanisms. By defining technology choice as minimizing a unit cost, I preclude scale effects as my unit cost specification implies that the cost of acquiring technology scales with total industry production. By assuming that technology acquisition costs in an industry are denominated in industry output (which is implicit by including the price P^k_i), I preclude that technology choices are affected by the relative price of output and technology acquisition. Lastly, I assume that technology barriers $\frac{1}{\bar{Z}_i}$ are common across factors and industries. This precludes that technology choices are affected by industry specific technology diffusion barriers, and it precludes that factor biases in technology arise due to factor specific technology diffusion barriers.

To solve for the technology choice, I take the first-order conditions associated with problem (22).

$$(c^k_i)_f \frac{r_{i,f}}{(Z^k_{i,f})^2} = \frac{P^k_i G^k_f}{\bar{Z}_i} \quad f = 1, \ldots, F \tag{23}$$

where the subscripts f on c^i_k and G^k denote partial differentiation with respect to argument number f. Multiplying both sides by $Z^k_{i,f}$, summing over f, and using that c^k_i and G^k are homogenous of degree 1 and $\gamma > 1$ respectively, I obtain

$$c^k_i = \frac{\gamma P^k_i G^k}{\bar{Z}_i}.$$

This means that

$$\begin{aligned} P_i^k &= \tilde{c}_i^k \\ &= c_i^k + \frac{P_i^k}{\bar{Z}_i} G^k \\ &= c_i^k \left(1 + \frac{1}{\gamma}\right). \end{aligned}$$

Substituting this back into the first-order condition (23), I obtain

$$\frac{(c_i^k)_f r_{i,f}}{c_i^k Z_{i,f}^k} = \left(1 + \frac{1}{\gamma}\right) \frac{G_f^k Z_{i,f}^k}{\bar{Z}_i}. \tag{24}$$

Noting that the left-hand side is

$$\alpha_{i,f}^k = \frac{(c_i^k)_f r_{i,f}}{c_i^k Z_{i,f}^k},$$

where $\alpha_{i,f}^k$ is the factor share of factor f in industry k for country i, equation (24) captures the intuition that a country expands further in a factor-augmenting technology if it has a high share of its costs devoted to that factor.

I can provide a stronger characterization if I put more structure on G^k and assume that it is given by

$$G^k = \sum_{f=1}^{F} \tilde{a}_f^k (Z_{i,f}^k)^\gamma \quad \gamma > 0.$$

In this case, the factor bias can be expressed as

$$\frac{\alpha_{i,f}^k}{\alpha_{i,1}^k} = \frac{\tilde{a}_f^k}{\tilde{a}_1^k} \left(\frac{Z_{i,f}^k}{Z_{i,1}^k}\right)^\gamma$$

I normalize $\tilde{a}_f^k = \alpha_{US,f}^k$ to ensure that the US has no technological bias. In this case, the relative factor bias is

$$\left(\frac{Z_{i,f}^k}{Z_{i,1}^k}\right) = \left(\frac{\alpha_{i,f}^k / \alpha_{US,f}^k}{\alpha_{i,1}^k / \alpha_{US,1}^k}\right)^{\frac{1}{\gamma}} \tag{25}$$

The relative factor bias is uniquely determined by the relative factor shares compared to the US. The smaller is γ, the more strongly relative factor technologies react to relative factor service prices.

D.3.2 Quantification

In this section, I quantify how my baseline estimation is affected by the modified assumptions on the industry production functions. In particular, I test how well my baseline method estimates relative factor service prices $\log\left(\frac{r_{i,f}/r_{i,1}}{r_{US,f}/r_{US,1}}\right)$ in this new environment.

For this purpose, I solve for the technology choice $Z_{i,f}^k$ and for unit costs c_i^k given factor prices $r_{i,f}$. I then run a regression

$$\log(c_i^k) = \delta_i + \mu^k + \sum_{f=2}^{F} \tilde{\beta}_{i,f}\alpha_{US,f}^k \quad \tilde{\beta}_{US,f} = 0.$$

I am interested in which factor service price combinations $\tilde{r}_{i,f}$ that generate $\tilde{\beta}_{US,f}$ which are similar to the $\beta_{i,f}$ that I find in my baseline estimation (6). By comparing the baseline $\beta_{i,f}$ with $\log\left(\frac{\tilde{r}_{i,f}/\tilde{r}_{i,1}}{\tilde{r}_{US,f}/\tilde{r}_{US,1}}\right)$, I can test how well my baseline estimate $\beta_{i,f}$ captures the relative price of factor services $\log\left(\frac{r_{i,f}/r_{i,1}}{r_{US,f}/r_{US,1}}\right)$ in this new environment. I do not run the full trade regressions, as I only modify how factor prices map to unit costs, and I do not modify how unit costs map to trade flows. I perform the regression on c_i^k excluding technology acquisition costs. As equilibrium technology acquisition costs uniformly scale unit costs, they do not affect the regression.

The detailed implementation of my method is as follows. I assume that there are 84 industries corresponding to the NAICS 4-digit manufacturing industries used in the baseline specification, and that there are two countries: "Poor" and the US. I assume that there are only two countries to reduce the number of parameters to estimate, while still capturing broad differences between rich and poor countries. I normalize US factor prices $r_{US,f} \equiv 1$, and US unskilled technology $Z_{US,1}^k = 1$ for all $k = 1, \ldots, 84$. I set the technology choice parameters to $\tilde{a}_f^k = \alpha_{US,f}^k$ which normalizes US technologies to $Z_{US,f}^k = 1$ for all factors and industries. The normalization of the technology choice function implies a normalization of the unit cost function, which becomes

$$c_i^k = \left(\sum_{f=1}^{F} \alpha_{US,f}^k \left(\frac{r_{i,f}}{Z_{i,f}^k}\right)^{1-\xi}\right)^{\frac{1}{1-\xi}}. \tag{26}$$

Furthermore, as only relative factor service prices are relevant for my estimation exercise, I can without loss of generality normalize $r_{Poor,1} = Z_{Poor,1}^k = \bar{Z}_i = 1$.

My task is to find $\tilde{r}_{Poor,f}$ for $f = 2, \ldots, F$ that replicate my baseline results. First, I use the CES industry production form to derive that

$$\frac{\alpha_{Poor,f}^k}{\alpha_{Poor,1}^k} = \frac{(\tilde{r}_{Poor,f}/Z_{Poor,f}^k)^{1-\xi}}{(\tilde{r}_{Poor,1}/Z_{Poor,1}^k)^{1-\xi}} \frac{\alpha_{US,f}^k}{\alpha_{US,1}^k}.$$

By combining this expression with equation (25), I obtain

$$\left(\frac{Z_{Poor,f}^k}{Z_{Poor,1}^k}\right) = \left(\frac{\tilde{r}_{Poor,f}/Z_{Poor,f}^k}{\tilde{r}_{Poor,1}/Z_{Poor,1}^k}\right)^{\frac{1-\xi}{\gamma}} \iff \frac{Z_{Poor,f}^k}{Z_{Poor,1}^k} = \left(\frac{\tilde{r}_{Poor,f}}{\tilde{r}_{Poor,1}}\right)^{\frac{1-\xi}{\gamma+1-\xi}}.$$

Thus, for each set of $\tilde{r}_{Poor,f}$, I can solve for technologies $Z^k_{Poor,f}$ and for unit costs c^k_{Poor}. I run the regression

$$\log(c_i^k) = \delta_i + \mu_k + \sum_{f=2}^{F} \tilde{\beta}_{i,f} \alpha^k_{US,f} \qquad \begin{array}{l} \tilde{\beta}_{US,f} = 0 \\ i = Poor, US \\ k = 1, \dots, 84 \end{array} .$$

I solve for $\tilde{r}_{Poor,f}$ for $f = 2, \dots, F$ such that $\tilde{\beta}_{Poor,f}$ matches $\beta_{Poor,f}$ from the baseline specification (I define $\beta_{Poor,f}$ by regressing my estimated $\beta_{i,f}$ on log GDP per worker $\log(y_i)$ for each f, and I define $\beta_{Poor,f}$ as the fitted value for $\log(y) = 9$). By comparing $\log\left(\frac{\tilde{r}_{Poor,s}/\tilde{r}_{Poor,1}}{\tilde{r}_{US,f}/\tilde{r}_{US,1}}\right)$ with $\beta_{Poor,s}$, I can gauge how biased my baseline estimation is in estimating the log relative price of skilled services. I test the effect of this bias on my development accounting exercise by redoing the development accounting exercise using an estimate of relative skilled service prices

$$\log\left(\frac{r_{i,s}/r_{i,1}}{r_{US,f}/r_{US,1}}\right) = \left(\frac{\log(y_{US}) - \log(y_i)}{\log(y_{US}) - 9}\right) \log\left(\tilde{r}_{Poor,s}\right).$$

D.3.3 Results

Table 11 shows the share of income differences explained by human capital for different values of the elasticity of substitution ξ and the endogenous technology parameter γ. A large γ means that technology choices are insensitive to variations in relative factor service prices. Unsurprisingly, we see that the endogenous technology choice parameter γ does not matter when $\xi = 1$. In this case, the production function is Cobb-Douglas and all technology differences are neutral. Furthermore, when $\gamma = 5$, the results for different ξ are similar to those found in Table 8 when there were no endogenous technology differences. This reflects that with a large γ, technology choices respond weakly to changes in relative factor service prices.

Overall, there is no monotone effect of endogenous technological change on the importance of human capital. For $\xi = 0.6$ and $\xi = 0.8$, making endogenous technology choices more flexible (smaller γ) makes human capital less important. For $\xi = 1.4$, making technology choices more flexible makes human capital dramatically more important. Overall, no specification reduces the importance of human capital below 50%.

Table 11: Share of income differences explained by human capital for different ξ and γ

	$\gamma = 1.01$	$\gamma = 1.1$	$\gamma = 2$	$\gamma = 5$
$\xi = 0.6$	0.55	0.56	0.62	0.65
$\xi = 0.8$	0.61	0.61	0.62	0.63
$\xi = 1$	0.64	0.64	0.64	0.64
$\xi = 1.2$	0.69	0.69	0.69	0.69
$\xi = 1.4$	6.83	6.48	1.01	0.90

E Appendix: Robustness

E.1 Discussion: Industry-dependent trade elasticities

In my estimates, I assume that the elasticity of trade σ is common across industries. A number of papers in the trade literature has argued for σ varying at an industry level (Broda et al., 2004; Soderbery, 2015). I write σ_k to denote such an industry-varying trade elasticity. Looking ahead, an important extension of my paper is to redo the estimates with a serious treatment of industry-varying σ. However, I have performed a simple robustness check, and tested a number of ways of solving the problem. Here, I also outline which approaches to this that look relatively more promising.

E.1.1 Fitted-residual plots and σ_k

First, I note that it is possible to use residual plots to detect evidence for industry-varying σ_k. If σ_k is higher than average in an industry, a plot of fitted values and residuals will have a positive slope. Indeed, if a country has high fitted trade values in an industry, it suggests that it has low relative costs. If I use an elasticity for that industry which is too low, the fitted value will be low compared to the actual value. The opposite is true when an industry has a low fitted value of trade. If I have underestimated the trade elasticity, actual values will be even lower than fitted values. These effects mean that an underestimated σ_k leads to a positive relationship between fitted values and residuals on an industry level. Conversely, if I have overestimated σ_k, there will be a negative relationship between fitted values and residual values.

By considering industry-by-industry plots of residuals on fitted values, I can obtain information about industry-specific elasticities. I use this method to perform a simple robustness check by excluding all industries with an absolute value of the residual-fitted plot of more than 1 and I find similar results for this restricted set of industries.

E.1.2 Including σ_k in regression specification

I also run the regression specification

$$\log(x_{i,j}^k) = \delta_{i,j} + \mu_j^k - \sum_{f=2}^{F}[(\sigma_k - 1)\alpha_{US,f}^k]\beta_{i,f} + \varepsilon_{i,j}^k \quad \beta_{US,f} \equiv 0$$

and use different estimates of σ_k across industries. I first use the estimates of industry-specific trade elasticities in Broda et al. (2004). To test whether these help resolve the problem with varying trade elasticities, I analyze whether there is less evidence for industry-varying trade elasticities in the fitted-residual plots when I use the industry-specific estimates σ_k from Broda et al. (2004) compared to when I run the regression with a common elasticity of trade corresponding to their median estimate.

I find that using the industry-specific estimates of trade elasticity do not resolve the problem of correlation between fitted values and residuals on the industry level. If anything, using industry-specific elasticity estimates makes the problem worse.

In addition to using the estimates from Broda et al. (2004), I also try an iterative procedure to more directly bring the fitted-residual plots in line. I run the regression with a common $\sigma_k \equiv \sigma$. I iterate and increase the σ_k whenever the fitted-residual slope in industry k is positive, and decrease σ_k whenever the fitted-residual slope in industry k is negative. Unfortunately, the procedure does not converge.

E.1.3 Road ahead on trade elasticity

Using estimates from Broda et al. (2004) and the iterative procedure did not solve the problem with varying trade elasticities. One potential reason for this failure is that it is not theoretically correct to modify regression specification (6) by just changing σ_k. If trade elasticities vary across industries, they also interact with trade cost terms that are now included in the bilateral fixed effect $\delta_{i,j}$. Thus, this will partly depend on industry k, which means that a standard gravity specification with bilateral fixed effects will not work in this context.

Thus, looking ahead, a proper treatment of varying σ_k will require a way of jointly estimating σ_k across industries and modify the structural trade model to generate a regression specification that fully incorporates varying trade elasticities.

E.2 Specification with confounding variables

In Section 7.1, I discuss the effects of an omitted variable bias in my specification of unit costs. Here, I explain how I measure and include potential confounders in my regression specification.

I analyze two confounding variables: external financing sensitivity and contracting sensitivity. I assume that there are country-specific contracting and external financing penalties τ_{cont} and τ_{fin} which capture the general quality of a country's judicial and financial systems. Industries are characterized by a contracting intensity $\alpha^k_{US,cont}$ and a external financing intensity $\alpha^k_{US,fin}$. Country-level contracting and financial penalties change log unit costs of industries with $\tau_{cont} \times \alpha^k_{US,cont}$ and $\tau_{fin} \times \alpha^k_{US,fin}$, respectively. That is, contracting and financing penalties increase the unit costs of industries in proportion to their respective contracting and external financing intensities.

I define an industry's external financing intensity $\alpha^k_{US,fin}$ as the share of investment expenditure not covered by external financing (external financing share of investments) times the share of gross output devoted to investments. I take the external financing share from Rajan and Zingales (1998), and I measure the investment share of total output using NBER CES data. My definition differs from that in Rajan and Zingales (1998) as I multiply the external financing share of investments with the investment share. The reason is that I interpret the country financial penalty τ_{fin} as a markup on external financing needs. A financing penalty increases the unit costs of an industry in proportion to its external financing needs as a share of gross output. To obtain external financing needs as a share of gross output, I multiply the external financing share of investments with the investment share.

I define an industry's contracting intensity by multiplying two terms. The first term is the share of intermediate inputs expenditure that is sensitive to contracting. To measure this term, I follow Nunn (2007) and use the IO table to calculate the share of intermediate good expenditures that are spent on customized inputs, where I define an input as customized if it is not traded on an exchange nor referenced priced in a trade journal according to the classification of goods in Rauch (1999). The second term in my calculation is the intermediate input cost share, defined as total intermediate good expenditures divided by gross output. I measure this term using NBER CES data. I calculate the contracting intensity $\alpha^k_{US,cont}$ as the product of these two terms, i.e. as the product of the share of customized inputs and the intermediate input cost share. My contracting sensitivity method is a slight modification of the measure in Nunn (2007), which only uses the first of my two terms. My modification reflects that I interpret the country contracting penalty τ_{cont} as increasing the cost of contracting sensitive inputs due to the lack of relation-specific investments. The unit cost effect of this is proportional to the total cost of contracting sensitive inputs as a share of gross output. As Nunn's definition only gives the cost of contracting sensitive inputs as a share of intermediate input costs, I multiply his measure with the intermediate input share to obtain my final measure $\alpha^k_{US,cont}$.

I include $\alpha^k_{US,cont}$ and $\alpha^k_{US,fin}$ in the analysis by adding extra terms to the regression specification (6), and run the regression

$$\log(x^k_{i,j}) = \tilde{\delta}_{i,j} + \tilde{\mu}^k_j - \sum_{f=2}^{F}[(\sigma-1)\alpha^k_{f,US}] \times \beta_{f,i} - [(\sigma-1)\alpha^k_{fin.,US}] \times \beta_{fin,i} - [(\sigma-1)\alpha^k_{cont.,US}] \times \beta_{cont,i} + \varepsilon^k_{i,j}.$$

where $\alpha_{fin.,US}, \alpha_{contr.US}$ give the financial and contracting intensity measured on US data.

E.3 Comparison with unit costs

My unit cost analysis uses the Groningen Growth and Development Center's (GGDC) 2005 benchmark producer price index. This data set aims at providing a cross-country comparable producer price index for 34 industries across 42 countries. The index covers both tradable and non-tradable industries, and manufacturing as well as services (Inklaar and Timmer, 2008).

Following recommendations from a creator of the data set, I exclude financial services, business services, real estate, government, health services and education. For these industries, it is difficult to obtain data on output quantities which makes it difficult to make cross-country comparisons in unit costs. I also exclude "private households with employed persons" as this variable is missing for a large number of countries. After my exclusions, I am left with a total of 27 industries and 35 countries with a complete set of observations.

To obtain factor shares, I use the EU KLEMS data set for the US (as my analysis includes non-manufacturing industries, I cannot use the NBER CES database to obtain factor shares). For the US, EU KLEMS provides data on industry level gross output, labor compensation, and intermediate good compensation. I define the labor share as the labor compensation over gross output, and the intermediate share as the intermediate good compensation over gross output. I calculate the skill share by multiplying the labor share with the share of payroll going to skilled workers with an occupational skill level of 3 or 4. I define the capital share as one minus the other factor shares.

I run the regression

$$\log(c_i^k) = \delta_i + \mu_k + \sum_{f=2}^{F} \alpha_{US,f}^k \tilde{\beta}_{i,f} + \varepsilon_i^k$$

where $\tilde{\beta}_{i,f} = \log\left(\frac{r_{i,f}/r_{i,1}}{r_{US,f}/r_{US,1}}\right)$ captures the deviation of relative prices compared to the US.

I compare the results from the unit cost analysis with the trade data analysis by comparing the relationship between GDP per worker and $\tilde{\beta}_{i,f}$ with the relationship between GDP per worker and $\beta_{i,f}$, where $\beta_{i,f}$ comes from the trade data analysis.[28]

In Figures 6 and 7, I plot the results from the unit cost data analysis. The slope parameter of log relative skilled service prices on log GDP per worker is -1.19 using the unit cost data, and -1.53 using the trade data method for the same set of countries. I cannot reject that the two coefficients are equal, even without taking into account the large standard errors on the unit costs based parameters $\tilde{\beta}_{i,f}$. Thus, when both types of data exist, the trade data method and the unit cost method paint a similar picture of the relationship between relative skilled service prices and GDP per worker.

E.4 Differences in unskilled human capital quality Q_u

In the current setup, I estimate the quality of unskilled labor Q_u by assuming that unschooled labor is of equal quality and that improvements are reflected in Mincerian returns:

$$Q_{U,i} = \exp(\phi(S_{U,i}))$$

where ϕ is a Mincerian return function and $S_{U,i}$ is the average schooling time of unskilled labor.

A number of papers on human capital and development accounting have stressed that there

[28] An alternative way to compare the outcomes would be to regress $\beta_{i,f}$ on $\tilde{\beta}_{i,f}$ and test how close the results are to a 45 degree line. I have chosen my method as I am interested in broad correlations between skilled service prices and GDP per capita, and given the estimation errors in the skill price estimates, regressing them on each other biases the results down due to measurement error. Regressing $\beta_{i,f}$ on $\tilde{\beta}_{i,f}$ and regressing $\tilde{\beta}_{i,f}$ on $\beta_{i,f}$ both yield a regression coefficient of less than one.

might be uniform quality differences in human capital (Caselli, 2005; Manuelli and Seshadri, 2014). These quality differences might reflect differences in nutrition, health, or the quality of early schooling.

As my paper estimates Q_u and Q_s/Q_u any uniform increase in Q_u will also increase Q_s proportionally.

F Appendix: Concordance construction

To generate concordances and map data across coding systems, I create a general mathematical framework to treat the problem. Here, describe how the general system works, and then I show how I use it to convert our particular data.

The basic building block of our concordance system is a many-to-many concordance between coding systems A and B where I have weights on both A and B. I call such concordances two-weighted concordances. An example of such a concordance is provided in Table 12.

In Table 12, note that each code in system A can be converted to multiple B codes (in this example, code 2 in System A maps to both code b and c in System B). The converse is also true: both code 4 and 5 map to code e. The weights code how important the respective industries are. This could, for example, be the total value of shipments, total trade value, etc. Notice that the weights are both on A and B, and that they are constant whenever they stand for the same industry.

I can define this mathematically as there being two sets A,B with measures w_A, w_B giving the mass on each code, and a concordance being a correspondence

$$\phi : A \rightrightarrows B.$$

I will write results in terms of this mathematical definition, but also in terms of examples to show the working of the system.

I will go through three operations relating to two-weighted concordances:

1. How to transform quantity variables such as total industry sales using a two-weighted concordance

2. How to transform property variables such as capital share using a two-weighted concordance

3. How to create a two-weighted concordance using an unweighted concordance and a weighting scheme for one of the variables (e.g. when I want to create a two-weighted concordance between HS and SITC and only have total trade in HS codes).

Table 12: Example concordance table

A	B	A_w	B_w
1	a	10	70
2	b	20	50
2	c	20	100
3	c	15	40
4	d	5	70
5	d	25	70
6	e	30	90

F.1 Transform quantity variables using two-weighted concordances

Starting with quantity variables, suppose that I have export values denoted in industry code A. I then want to allocate it across different codes in industry code B given a weighting scheme on B. In this case, for each element $a \in A$, I allocate the export values in industry a across industries $b \in B$ in proportion to their weights w_b. The quantity attributed to element $b \in B$ is then the sum of the contributions from all elements in A to b.

I can write this in terms of the mathematical representation Φ as well, together with the weights μ_A and μ_B. If

$$f_A : A \to \mathbb{R}$$

is an arbitrary quantity measure on A I convert it to B by

$$f_B(b) = \sum_{a \in \Phi^{-1}(b)} f_A(a) \times \frac{\mu_B(b)}{\sum_{b' \in \Phi(a)} \mu_B(b')}.$$

F.2 Transform property variables using two-weighted concordances

The situation is different when I have so-called property variables, for example capital share, skill share or other industry-level properties. The difference can be illustrated with an example.

In the previous part, I considered the problem of mapping trade data from A to B. Then, the reasonable thing is to split it up the value a across $b \in \Phi(a)$ according to the weights w_b. However, suppose that I want to map the capital share from a to b. Then, we should not split up the capital share across $b \in \Phi(a)$. If b and b' have the same pre-image a, they should have the same capital share as a.

Thus, property variables translate across coding systems in a fundamentally different way from quantity variables. I define the transformation scheme for property variables by saying that for each code $b \in B$ in the target system, I define its property as a weighted average of the properties that its pre-images $a \in A$, where I use the weights on A as a weighting scheme. For example, in our example concordance, I would attribute c a property which is the weighted average of 2,3 in

System A, using the measures $\mu_A(\{2\}) = 20$ and $\mu_A(\{3\}) = 15$ as weights.

More formally, if I have a property measure

$$g_A : A \to \mathbb{R}$$

defined on A, then I translate it to B using ϕ by the equation

$$g_B(b) = \frac{\sum_{a\in\phi^{-1}(y)} g_A(a)\mu_A(a)}{\sum_{a\in\phi^{-1}(b)} \mu_A(a)}.$$

F.3 Construct a two-side weighted concordance from a one-sided weighted concordance

Above I defined how you translate between different coordinate systems if you have a two-sided weighted concordance. However, sometimes I only have a one-sided concordance. For example, if I have total trade data in HS 2007 six-digit and want to create a concordance between HS 2007 6-digit and NAICS 2007 it might be that I do not have data to create a natural weighting scheme for the NAICS 2007 coding scheme.

For this case, I have a procedure to create a two-sided weighted concordance from a one-sided weighted concordance. It is quite similar to the quantity transformation above. Suppose that I have a concordance ϕ and a measure μ_A on A and want to create a measure μ_B on B. Then I define the measure on B as.

$$\mu_B(b) = \sum_{x\in\phi^{-1}(b)} \frac{\mu_A(a)}{|\phi^{-1}(a)|}.$$

That is, I split the weights on $a \in A$ equally on all $b \in B$ to which a maps.