Federal Reserve Bank of Minneapolis
Research Department Staff Report 435

September 2009

International Trade and Income Differences*

Michael E. Waugh

Federal Reserve Bank of Minneapolis
and New York University

ABSTRACT

I develop a novel view of the trade frictions between rich and poor countries by arguing that to reconcile bilateral trade volumes and price data within a standard gravity model, the trade frictions between rich and poor countries must be systematically asymmetric, with poor countries facing higher costs to export relative to rich countries. I provide a method to model these asymmetries and demonstrate the merits of my approach relative to alternatives in the trade literature. I then argue that these trade frictions are quantitatively important to understanding the large differences in standards of living and total factor productivity across countries.

*Email: mwaugh@stern.nyu.edu. Many people have provided helpful comments and suggestions at many different venues that improved this paper. Special thanks go to Timothy J. Kehoe, Narayana Kocherlakota, Samuel Kortum, and Kei-Mu Yi. Finally, I owe a debt of gratitude to my former advisors B. Ravikumar and Raymond G. Riezman for their invaluable encouragement and support.

I. Introduction

Standards of living between the richest and poorest countries differ by more than a factor of 30. A large literature has evolved that attempts to explain this fact within the context of a standard (closed-economy) neoclassical growth model. The consensus is that physical and human capital accounts for only 50 percent of the variation in income per worker; the rest is productivity differences. Given this finding, a growing literature has attempted to understand how various frictions result in large differences in measured productivity across countries.

In this paper, I develop a novel view of the frictions to trade between rich and poor countries by arguing that to reconcile bilateral trade volumes and price data within a standard gravity model, the trade frictions between rich and poor countries must be systematically asymmetric, with poor countries facing higher costs to export relative to rich countries. I then argue that these frictions to trade are quantitatively important to understanding why standards of living and measured total factor productivity between the richest and poorest countries differ by so much.

The starting point of my analysis is a multicountry model of trade combining a standard gravity model with elements from a neoclassical growth model. Each country has two sectors: a tradable goods sector and a final goods sector, both with constant returns technologies. Labor, capital, and tradable goods are used as factors of production. In the tradable goods sector there is a continuum of goods. As in Dornbusch, Fischer and Samuelson (1977), production technologies differ across goods on the continuum only in their productivity levels. As in Eaton and Kortum (2002), productivity levels are treated as random variables drawn from a parameterized distribution. Each country's distribution differs in its average productivity level. Trades occur only within intermediate goods, which are purchased from the country with the lowest price that includes "iceberg" costs to trade. The final goods sector produces a nontraded consumption good with a technology common to all countries.

Within this framework, I first ask: What trade costs between rich and poor countries are necessary to reconcile bilateral trade volumes and aggregate price data? The building block to answering this question is a relationship in the model between the degree of relative home bias, relative bilateral trade shares, relative prices, and relative trade costs. The first three are observable, whereas trade costs are not. Thus, I use data and the structure of the model to infer relative trade costs between country pairs. In the data, home bias and relative prices for tradable goods do not covary strongly with income per worker, yet relative bilateral trade shares do. In the model, these observable are monotonically related to trade costs, thus trade costs must covary with income per worker such that poor countries face higher costs to export relative to rich countries.

I propose a simple approach to model this asymmetry that retains the parsimony of estimable log-linear gravity equations. The idea is to allow trade costs to vary contingent upon the exporter. To illustrate this idea, consider the following facts and how the model might try to fit it. The U.S. import share from

Japan is larger than its import share from Senegal. Yet both Japan's and Senegal's import shares from the United States are similar. In the model, a country imports a larger share from countries able to supply goods at lower prices relative to domestic producers. Two factors influence the prices at which a country can supply goods: (i) the trade cost to export and (ii) the domestic unit cost of the goods. These are the free parameters picked to fit the trade data.

To fit these facts the result is that (i) Japan's cost to export to the United States is lower than Senegal's and (ii) the countries' domestic unit costs are similar. Result (i) reconciles differences in the United States' import share across the two countries. Result (ii) reconciles Japan's and Senegal's similar import shares from the United States because their cost to produce domestically and to import from the United States are the same. This result is important because the domestic unit cost function determines (approximately) the aggregate price of tradables in each country, since most goods are purchased from home in the data and in the model. Thus, the model-implied price of tradable goods is similar between Japan and Senegal—consistent with the data. Furthermore, I also show how the model correctly predicts the observed differences in income per worker across countries when using this approach.

These arguments contribute to the gravity literature because previous approaches have difficulties fitting bilateral trade volumes and/or price data when rich and poor countries are included in the analysis. Standard approaches to modeling trade costs usually assume that they are symmetric. For example, trade costs are assumed to be some function of distance, shared border, language, colonial relationship, and so on—all symmetric.[1] The merits of my approach can be seen by comparing my model's ability to fit bilateral trade volumes relative to standard approaches: my approach has an R^2 that is 36 percent higher than standard approaches.

An exception is Eaton and Kortum (2001, 2002), which allows trade costs to vary contingent upon the importer. This approach fits the trade data as well as my approach. But this approach also generates counterfactual implications relative to data on the aggregate price of tradables. To see this, consider how the approach of Eaton and Kortum (2001, 2002) reconciles the fact that the U.S. import share from Japan is larger than its import share from Senegal. The only way to reconcile this fact is if Japan's domestic unit cost is lower than Senegal's because the U.S. cost to import from Japan and Senegal must be the same given the restrictions imposed. Because of the relationship between the domestic unit cost function and the aggregate price of tradables, this implies that Japan should have a lower price of tradables than Senegal, which is inconsistent with the data. Finally, I should emphasize that my arguments are equally applicable to alternative structural gravity models such as Anderson and van Wincoop (2003) with no firm heterogeneity, or variants of Melitz (2003) with fixed costs and Pareto distributed productivity; thus, my arguments may be interpreted more generally.

With a model that can reconcile salient features of the data, I then ask: If trade costs changed, then how would cross-country income differences change?[2] In this paper, I focus on two counterfactual exercises

[1] For examples, see Anderson and van Wincoop's (2004) survey of this literature.

[2] Other papers have studied the welfare consequences of frictions to trade in similar models of trade, e.g., Eaton and Kortum

that emphasize the quantitative importance of the systematic asymmetry in trade frictions. In one experiment, trade costs are set so that two countries both face the minimum calibrated trade costs between them. Given the systematic asymmetry in trade costs, the premise is that costs above this minimum reflect some extra distortion one country faces while the other does not, and that these distortions, unlike distance, are not natural. In this experiment, providing countries with equal market access reduces cross-country income differences by up to 31 percent. In the second experiment, I endow all countries with the same effective trade costs that OECD countries face among themselves. The premise here is that the costs to trade between OECD countries are relatively closer to free trade than the entire sample. In this experiment, cross-country income differences are reduced by up to 23 percent. In all these cases, the reduction in income differences comes from the reallocation of goods production across countries according to comparative advantage.

To put these two exercises in context, consider their impact relative to either a complete elimination of trade costs or a move to autarky. Eliminating the asymmetry or moving toward OECD trade costs delivers 59 and 41 percent of the reduction in income differences relative to complete elimination of trade costs. Or at the other extreme, increasing trade costs so there is *no* trade changes income differences little. The systematic asymmetry in trade costs is so punitive that removing it takes the economy from basically autarky to over 50 percent of the way relative to frictionless trade.

These arguments contribute to the macro-development literature by demonstrating the quantitative importance of trade frictions for economic development. Klenow and Rodríguez-Cláre (1997), Hall and Jones (1999), Parente and Prescott (2002), and Caselli (2005) are examples that find cross-country income differences mostly result from differences in total factor productivity. With this well-established fact, the literature has moved toward a focus on how various frictions affect measured total factor productivity.[3] In this model, measured total factor productivity is endogenous and depends upon the pattern of trade and, hence, trade costs. I demonstrate that reductions in trade costs are quantitatively important to reducing differences in measured total factor productivity and, hence, cross-country income differences by allowing for specialization via comparative advantage. Admittedly, my counterfactual exercises are abstract because I follow the gravity literature and use the model and trade data to infer trade fictions. However, my arguments suggest some component of these frictions is not related to natural trade barriers, and this component is quantitatively important to understanding why standards of living and measured total factor productivity between the richest and poorest countries differ by so much.

(2002) for OECD countries and Alvarez and Lucas (2007) for a larger set of developed and undeveloped countries. Redding and Venables (2004) study similar questions in a structural model of economic geography. My paper's contribution is a new view of the trade frictions between rich and poor countries that can simultaneously reconcile both trade flows and prices that previous approaches cannot. Further, I argue that the asymmetric component is quantitatively important to understanding cross-country income differences.

[3]See Lagos (2006) on labor market frictions; Hsieh and Klenow (2007a) and Restuccia and Rogerson (2008), which focus on distortions between establishments within a country; Guner, Ventura and Yi (2008) study size-dependent policies; and Burstein and Monge-Naranjo (2009) focus on barriers to foreign direct investment.

II. The Model

Consider a world with N countries. Each country has two sectors: a tradable goods sector and a final goods sector. Only tradable goods are traded. Within each country i, there is a measure of consumers L_i. Each consumer has one unit of time supplied inelastically in the domestic labor market, and each is endowed with capital supplied to the domestic capital market. Furthermore, each consumer has preferences only over the final good, which is nontraded. In the following, all variables are normalized relative to the labor endowment in country i.

A. Tradable Goods Sector

As in Dornbusch et al. (1977), there is a continuum of tradable goods indexed by $x \in [0,1]$. To produce quantity $m_i(x)$ in country i, capital k_i, labor n_i, and the aggregate tradable good q_i are combined by the following nested Cobb-Douglas production function:

$$m_i(x) = z_i(x)^{-\theta}[k_i^{\alpha} n_i^{1-\alpha}]^{\beta} q_i^{1-\beta}.$$

Across goods x, production technologies differ only in their productivity $z_i(x)^{-\theta}$. Power terms α, β, and θ are common to all countries.[4] The representative firm's problem in country i is to minimize the cost of supplying $m_i(x)$ by choosing capital, labor, and the aggregate tradable good, given factor prices, r_i, w_i, and p_i. All firms in country i have access to the technology for any good x with the efficiency level $z_i(x)^{-\theta}$.

Individual tradable goods are aggregated according to a standard symmetric Dixit-Stiglitz technology producing the aggregate tradable good with elasticity of substitution $\eta > 0$:

$$q_i = \left[\int_0^1 m(x)^{\frac{\eta-1}{\eta}} dx\right]^{\frac{\eta}{\eta-1}}$$

Firms in this sector simply have the problem of supplying q_i at minimum cost by purchasing tradable goods $m(x)$ from the lowest cost producer across all countries.

1. How Does Trade Occur?

Trade occurs in the following manner. A firm producing the aggregate tradable good imports each good from the lowest cost producer across all countries. Three factors influence which country is the lowest

[4] It is worthwhile to contrast the use of tradable goods here with the model of Yi (2003) in which there are two stages of production. Individual goods x in the first stage of production are used directly in the second stage of production and then aggregated. It is this mechanism that is important for quantitatively explaining the growth in world trade.

cost producer: (i) factor prices (w, r, p), (ii) trade costs between countries (see below), and (iii) good-level productivity $z^{-\theta}$. Factor prices are determined in equilibrium. Trade costs are exogenous. Good-level productivity is modeled as an idiosyncratic random variable drawn from an exogenously given country-specific distribution as in Eaton and Kortum (2002).

Modeling productivity this way is particularly convenient. For example, consider two identical countries. Ex ante, there is no incentive to trade, but ex post—after z's are assigned (randomly) to each good—there is. Because of the luck of the draw, one country will be relatively more productive than the other country (and vice versa) at producing different goods, and hence there is room for trade.[5] Given the appropriate distributional assumptions on good-level productivity, calculating aggregate trade flows between countries boils down to calculating some probability statistics. Below I discuss the benchmark distributional assumptions.

2. The Distribution of Productivity Across Goods

I follow Alvarez and Lucas (2007) and assume that $z_i(x)$ is distributed independently and exponentially with parameter λ_i differing across countries. Because good-level productivity is $z^{-\theta}$, this formulation is equivalent to a Type II extreme value distribution or Fréchet distribution used by Eaton and Kortum (2002).[6] The λ_is and θ play the following roles regarding the distribution of productivity across goods:

λ_i governs each country's average productivity level. One can show that each country's mean productivity is proportional to λ_i^{θ}, with the constant of proportionality not depending upon the country. So a country with a relatively larger λ_i is, on average, more efficient at producing all tradable goods.

θ controls the dispersion of efficiency levels. Mechanically, a larger (smaller) θ yields more (less) variation in efficiency levels relative to the mean. Returning to the discussion above, as θ increases (decreases), it increases (decreases) the likelihood that the productivity of the two countries at producing the same good will be different, thus yielding more (less) incentives to trade. In this sense, θ controls the degree of comparative advantage.

B. Final Goods Sector

In each country, a representative firm produces a homogeneous good that is nontraded. Each firm has access to the following nested Cobb-Douglas production function combining capital, labor, and the ag-

[5] There is evidence that productivity within narrowly defined sectors can vary significantly even across developed countries for idiosyncratic reasons. See Baily and Gersbach (1995) and the associated studies of the McKinsey Global Institute.

[6] Kortum (1997) shows how a model of innovation and diffusion consistent with balanced growth can give rise to this distribution. Appendix F discusses results from alternative distributional assumptions.

gregate tradable good:

$$y_i = [k_i^{\alpha} n_i^{1-\alpha}]^{\gamma} q_i^{1-\gamma}.$$

Factor shares, α and γ, are the same across countries. The representative firm's problem is to minimize the cost of producing y_i, at price p_i^y, by selecting the amount of capital, labor, and aggregate tradable good, taking prices as given. Recall that each consumer has preferences only over this good.

C. Trade Costs

To model trade costs, the standard iceberg assumption is made, i.e., $\tau_{ij} > 1$ of good x must be shipped from country j for one unit to arrive in country i in which $(\tau_{ij} - 1)$ "melts away" in transit. Trade costs τ_{ij} are thought to be composed of both policy and nonpolicy related barriers. In addition, τ_{ii} is normalized to equal one for each country.

D. Equilibrium

A competitive equilibrium is characterized by a set of allocation rules, prices, and trade shares. The most important objects are the functions that determine the aggregate price of tradable goods, trade shares, and wages from which all other equilibrium objects are determined. These three functions also provide the basis for the mapping between the data and the model, as discussed throughout the rest of the paper.

Allocation Rules: With Cobb-Douglas production technologies, it is straightforward to show that a fraction γ of capital, labor, and β of the aggregate tradable good are allocated toward the final goods sector.

Price Index: Each country faces the following price of tradable goods for each country i:

$$p_i = \Upsilon \left\{ \sum_{j=1}^{N} \left[r_j^{(1-\alpha)\beta} w_j^{\alpha\beta} p_j^{(1-\beta)} \tau_{ij} \right]^{\frac{-1}{\theta}} \lambda_j \right\}^{-\theta}, \tag{1}$$

where Υ is a collection of constants.

Trade Shares: X_{ij} is country i's expenditure share on goods from country j. It is also the fraction of all goods that country i imports from country j. Since there is a continuum of goods, computing this fraction boils down to finding the probability that country j is the low-cost supplier to country i given the joint distribution of efficiency levels, prices, and trade costs for any good x. The expression for a trade share is,

$$X_{ij} \quad = \quad \frac{\left[r_j^{(1-\alpha)\beta} w_j^{\alpha\beta} p_j^{(1-\beta)} \tau_{ij}\right]^{\frac{-1}{\theta}} \lambda_j}{\sum_{\ell=1}^{N} \left[r_\ell^{(1-\alpha)\beta} w_\ell^{\alpha\beta} p_\ell^{(1-\beta)} \tau_{i\ell}\right]^{\frac{-1}{\theta}} \lambda_\ell}. \tag{2}$$

Note that the sum across j for a fixed i must add up to one.

Wages: An equilibrium wage vector is computed given trade shares and imposing balanced trade. Imports are defined as

$$\text{Imports} = L_i p_i q_i \sum_{j \neq i}^{N} X_{ij},$$

which is the total value of all goods that country i consumes from abroad. Similarly, exports are defined as

$$\text{Exports} = \sum_{j \neq i}^{N} X_{ji} L_j p_j q_j,$$

which is the total value of all goods that countries abroad purchase from country i.

Imposing balanced trade and including each country i's consumption of goods produced at home implies the following relationship must hold,

$$L_i p_i q_i \sum_{j=1}^{N} X_{ij} = \sum_{j=1}^{N} L_j p_j q_j X_{ji},$$

which says the aggregate value of tradable goods purchased by country i is equal to the value of tradable goods all N countries purchase from country i.

Using the observation that each country allocates $(1-\gamma)$ of capital and labor to the production of the tradable goods sector and the relationship between factor payments and total revenue (see Alvarez and Lucas (2007)), the equilibrium wage rate for each country i is

$$w_i \quad = \quad \sum_{j=1}^{N} \frac{L_j}{L_i} w_j X_{ji}. \tag{3}$$

At this point, the three key pieces of the model have been derived. Equation (1) describes the equilibrium price of tradable goods, equation (2) describes the share of goods countries purchase from each other, and equation (3) describes the equilibrium wage rate for each country. From these functions, all other prices and quantities are determined and an equilibrium constructed.

III. Trade Data, Price Data, and Model $\Rightarrow$ Asymmetric Trade Costs

In the model, a country is a labor endowment, capital endowment, technology parameter λ_i, and collection trade costs τ_{ij}. The first two I can observe, the last two I cannot. Here, I use the model and some data to understand the structure of trade costs across countries. Standard approaches to modeling trade costs usually assume they are symmetric, i.e., some function of distance, shared border, language, colonial relationship, and so on—all symmetric. In this section, I show that when price data are brought into the analysis, the model-implied trade costs systematically deviate from symmetry by covarying with level of development.

A. Trade Data

As a benchmark, I consider the model year to be 1996.[7] Seventy-seven countries are in the sample and represent over 90 percent of world GDP in 1996. I assume that the tradable goods sector corresponds to manufactures.[8] I constructed trade shares X_{ij} following Eaton and Kortum (2002) and Bernard, Eaton, Jensen and Kortum (2003) in the following manner:

$$
\begin{aligned}
X_{ij} &= \frac{\text{Imports}_{ij}}{\text{Gross Mfg. Production}_i - \text{Total Exports}_i + \text{Imports}_i}, \\
X_{ii} &= 1 - \sum_{j \neq i}^{N} X_{ij}.
\end{aligned}
$$

In the numerator is the aggregate value of manufactured goods that country i imports from country j. These data are from Feenstra, Lipsey and Bowen (1997), and manufactures are defined to be the aggregate across all 34 BEA manufacturing industries. In the denominator is gross manufacturing production minus total manufactured exports (for the whole world) plus manufactured imports (for only the sample). Basically, this is simply computing an expenditure share by dividing the value of inputs consumed country i imported from country j divided by the total value of inputs in country i. Gross manufacturing data are from either UNIDO, OECD, or imputed from value added data from the World Bank.

Table 1 presents a matrix of trade shares for selected countries. A row denotes the exporting country, and a column denotes the importing country. Note two important features:

[7] I also considered data from 1985, and all the results are quantitatively consistent with the results discussed throughout the paper. Details are available upon request.

[8] This is a simplification, but it is reasonable as a first-order approximation to reality because, for all countries in the sample, this represents on average 80 percent of all merchandise imports; the median is 94 percent. The more relevant concern is whether there is a systematic bias in the amount of trade not included and level of development. I considered this by regressing the percent of trade not included on a bilateral basis on income level of importer and exporter. There is a slight relationship between the difference in income per worker of importing and exporting countries and the amount of trade not included; however, the magnitude is negligible.

Table 1: 1996 Trade Shares X_{ij} in Percent for Selected Countries

	U.S.	Can.	Japan	Mexico	China	Senegal	Malawi	Zaire
U.S.	83.25	39.73	2.27	31.62	3.63	2.16	1.57	2.93
Can.	3.78	49.21	0.21	0.72	0.32	0.56	0.67	0.51
Japan	3.04	2.01	92.56	1.59	6.99	1.34	2.65	0.82
Mexico	1.88	1.33	0.02	61.09	0.057	0.01	0	0.007
China	1.78	1.41	1.44	0.30	77.61	2.69	2.50	6.81
Senegal	0^*	0^*	0^*	0	0^*	52.68	0	0
Malawi	0^*	0^*	0^*	0	0	0	41.52	0
Zaire	0.003	0.005	0.003	0^*	0^*	0	0	51.53

Note: Entry in row i, column j, is the fraction of goods country j imports from country i. Zeros with stars indicate the value is less than 10^{-4}. Zeros without stars are zeros in the data.

O.1. Home bias for both rich and poor countries. First, by home bias I mean that countries purchase most of their goods from home, i.e., X_{ii} data. Home bias in the data is seen by noticing the large values lying along the diagonal of Table 1 relative to off-diagonal entries. The important observation is that there is little variation in the X_{ii}s relative to a country's income per worker. A regression of the logarithm of X_{ii} on the logarithm of purchasing power parity (PPP) adjusted GDP per worker in 1996 has a slope coefficient of 0.12 and is barely different from zero statistically. Rich countries purchase slightly more from home than poor countries, but the difference in magnitude is small.

O.2. Systematic correlation between bilateral trade shares and relative level of development. To see this correlation, notice the values in the upper right quadrant (encompassing poor countries' imports from rich countries) are large relative to those in the lower left quadrant of Table 1 (encompassing rich countries' imports from poor countries). To illustrate this point for all countries, I regressed the logarithm of $\frac{X_{ij}}{X_{ji}}$ on the logarithm of $\frac{y_j}{y_i}$. Here, y_i is PPP adjusted GDP per worker in 1996. The intercept is approximately zero and the slope coefficient is 2.40. Both are precisely estimated. The regression illustrates that the larger the difference in relative incomes, the larger the disparity in bilateral trade shares between the two countries.

B. Price Data

In the model, the prices p_i are the aggregate price indices of tradable goods. These are *tradable* goods, not *traded* goods, since in equilibrium some goods may not be traded. Furthermore, since the bundle of tradable goods is the same for all countries, a key concern is that the data are comparable across countries. To construct data on these prices, I employed price data from the United Nations International Comparison Program (ICP). This program collects prices on goods and services in various countries and benchmark years, with the most relevant feature being the explicit goal of comparability during their collection. That is, prices are supposed to be for the same or similar goods, and the baskets of goods are the same across countries.

To construct tradable price indices, benchmark data for the year 1996 were obtained from the Penn World Table (PWT) website (http://pwt.econ.upenn.edu). Data on prices are provided at disaggregate categories for each benchmark year in local currencies. From this data, only categories that best correspond with the bilateral trade data are included.[9] I then constructed the appropriate price indices of tradable goods and then converted to U.S. dollar prices.[10] This results in the third key feature of the data:

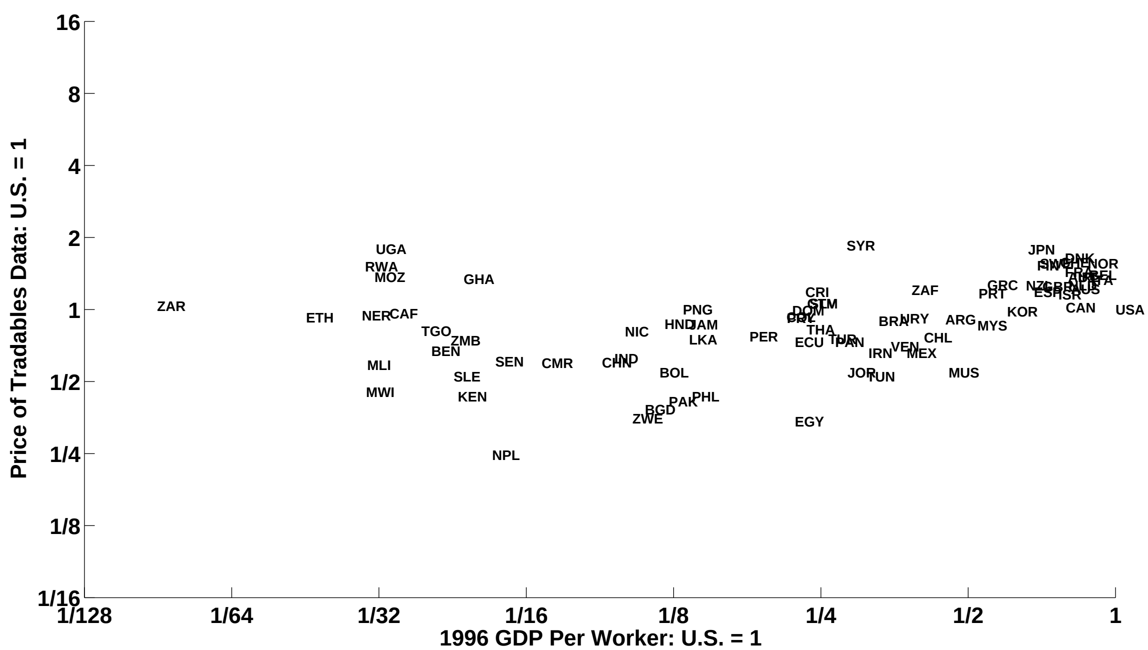

Figure 1: Price of Tradable Goods: Similar Between Rich and Poor Countries

O.3. Aggregate tradable goods prices are similar between rich and poor countries. Figure 1 plots the price of tradable goods for 1996 versus PPP adjusted GDP per worker data for that year. As the figure illustrates, poor countries have slightly lower prices of tradable goods with an elasticity of the price of tradable goods with respect to income level of 15 percent. My results are consistent with Kravis

[9]There is no one-to-one mapping, so discretion is involved. Some categories seem to include items that are inherently non-traded, e.g., “Footwear and Repairs”, which was included. At the current aggregation provided for 1996, this is unavoidable, and PWT administrators are unwilling to provide me with any finer level of aggregation.

[10]Not all 77 countries are benchmark countries. The trade-off I face is between a large sample of bilateral trade shares versus using only benchmark countries. I opted for more trade data. To construct price indices for nonbenchmark countries, I imputed their values from information in the PWT. To do so, I regressed the constructed prices for the entire benchmark table on the price of consumption (pc) and price of investment (pi), which is available for all countries directly from the PWT. Given the estimated coefficients, I imputed the price of tradable goods for nonbenchmark countries by using the observed price of consumption and price of investment.

and Lipsey (1988), who document similar relationships between the price of tradable goods and level of development, and Hsieh and Klenow (2007b), who study similar price indices for only investment goods.

C. The Implications of Trade Data, Price Data, for Trade Costs

In the model, **O.1-O.3** yield a straightforward implication for trade costs. Manipulation of equations (1) and (2) yields the following relationship between home trade shares, bilateral trade shares, aggregate prices, and trade costs:

$$\frac{X_{ij}}{X_{jj}} = \tau_{ij}^{-\frac{1}{\theta}} \times \left(\frac{p_j}{p_i}\right)^{-\frac{1}{\theta}}. \tag{4}$$

Equation (4) is basically an arbitrage condition.[11] It says that if $p_i > p_j$, then country i has incentives to purchase relatively more goods from country j because they are cheaper. Or if trade costs between country i and j are large, then country i has fewer incentives to purchase a good from country j. Dividing equation (4) by the opposing expression relating country j and i yields:

$$\left(\frac{X_{ij}}{X_{ji}}\frac{X_{ii}}{X_{jj}}\right) \times \left(\frac{p_j}{p_i}\right)^{\frac{2}{\theta}} = \left(\frac{\tau_{ij}}{\tau_{ji}}\right)^{-\frac{1}{\theta}}. \tag{5}$$

Consider the term $\left(\frac{X_{ij}}{X_{ji}}\frac{X_{ii}}{X_{jj}}\right)$, which incorporates data from both **O.1** and **O.2**. These observations imply that this term is positively correlated with exporters' income per worker relative to importers' income per worker. For example, the elasticity of this term with respect to relative income per worker $\frac{y_j}{y_i}$ is 2.30 and statistically different from zero.

In a symmetric world, the term $\left(\frac{X_{ij}}{X_{ji}}\frac{X_{ii}}{X_{jj}}\right)$ equals one always. Equation (5) states that deviations from the symmetric benchmark occur for only two reasons: (i) aggregate prices of tradable goods are different or (ii) relative trade costs between the two countries are different. **O.3** informs us on point (i). Aggregate tradable goods prices are similar between rich and poor countries, i.e., there is little correlation between $\left(\frac{p_j}{p_i}\right)^{\frac{2}{\theta}}$ and relative income level.

Thus **O.1-O.3**, equation (5), and a $\theta > 0$ imply that relative trade costs $\frac{\tau_{ij}}{\tau_{ji}}$ are negatively correlated with exporters' income per worker relative to importers' income per worker. This means for a poor country trading with a rich country, it is relatively more difficult for the poor country to export its goods to the rich country than import goods from the rich country. Furthermore, this result conflicts with standard approaches to modeling trade costs which assume that $\tau_{ij} = \tau_{ji}$.

[11] There is nothing unique about equation (4) to the Eaton and Kortum (2002) framework. As I show in Appendix A, the structural gravity models of Anderson and van Wincoop (2003) or variants of Melitz (2003) all generate this relationship. Thus the implications of **O.1-O.3** and the arguments throughout the paper apply under alternative models of international trade.

IV. Modeling Asymmetry: Some Examples

One would like an approach to model asymmetric trade costs, yet retain the parsimony that structural gravity equations provide when estimating trade costs from trade flows. The key to any approach is to do so without resulting in counterfactual implications for prices and other objects of interest. In this section, I provide two examples illustrating different approaches to fitting the bilateral trade share data, their different implications, and how the trade data map into the estimated parameters. Example 1 provides the motivation for the benchmark approach I use in Section V. Example 2 is an alternative approach, similar in spirit to Eaton and Kortum (2001, 2002), and with which I will contrast my results.

A. Common Elements to Both Examples

Consider a world with three countries. Think of country 1 as the United States and countries 2 and 3 as middle income and poor countries. Assume throughout the example that $X_{12} > X_{13}$ but $X_{21} = X_{31}$ and countries 2 and 3 do not trade with each other. These assumptions resemble Table 1 and the data generally because the United States exports similar shares to most countries, yet the United States imports a monotonically increasing share of goods as a function of the exporting country's level of development, i.e., country 2 is richer than country 3, so the Unites States imports a larger share from country 2 relative to country 3. With columns denoting importers and rows as exporters, the matrix in equation (6) depicts the set of bilateral trade shares normalized by the importing countries' home trade share:

$$\begin{pmatrix} 1 & \frac{X_{21}}{X_{22}} & \frac{X_{31}}{X_{33}} \\ \frac{X_{12}}{X_{11}} & 1 & 0 \\ \frac{X_{13}}{X_{11}} & 0 & 1 \end{pmatrix} \quad \begin{array}{l} \text{where } \frac{X_{21}}{X_{22}} = \frac{X_{31}}{X_{33}} \\ \text{and } \frac{X_{12}}{X_{11}} > \frac{X_{13}}{X_{11}} \end{array} \tag{6}$$

To simplify the following examples, assume that (i) all countries have the same labor endowment, (ii) labor is the only factor of production, and (iii) country 1's w_1 and λ_1 are normalized to one. Equation (2) implies the following relationship between (6) and six unknown structural parameters $\{\lambda_2, \lambda_3, \tau_{12}, \tau_{21}, \tau_{13}, \tau_{31}\}$:

$$\begin{pmatrix} 1 & \frac{1}{(w_2\tau_{21})^{\frac{-1}{\theta}}\lambda_2} & \frac{1}{(w_3\tau_{31})^{\frac{-1}{\theta}}\lambda_3} \\ \frac{(w_2\tau_{12})^{\frac{-1}{\theta}}\lambda_2}{1} & 1 & 0 \\ \frac{(w_3\tau_{13})^{\frac{-1}{\theta}}\lambda_3}{1} & 0 & 1 \end{pmatrix} \quad \begin{array}{l} \text{where (6)} \Rightarrow \frac{1}{(w_2\tau_{21})^{\frac{-1}{\theta}}\lambda_2} = \frac{1}{(w_3\tau_{31})^{\frac{-1}{\theta}}\lambda_3} \\ \text{and } \frac{(w_2\tau_{12})^{\frac{-1}{\theta}}\lambda_2}{1} > \frac{(w_3\tau_{13})^{\frac{-1}{\theta}}\lambda_3}{1} \end{array} \tag{7}$$

Notice there are six parameters yet only four informative moments, $\left\{\frac{X_{12}}{X_{11}}, \frac{X_{13}}{X_{11}}, \frac{X_{21}}{X_{22}}, \frac{X_{31}}{X_{33}}\right\}$. Thus, one needs additional restrictions on the parameter space before taking the model to the data. In the N country case, this identification problem remains with only $N^2 - N$ moments yet N^2 parameters of interest. Below I consider two alternative restrictions. Both fit the trade data equally well, but have qualitatively different implications for the aggregate price of tradable goods and productivity/income differences.

B. Example 1: Export Effects

This setup restricts the parameter space so $\{\tau_{21}, \tau_{31}\} = \bar{\tau}$, i.e., the cost for countries 2 and 3 to import is the same. The free parameters are the export costs that countries 2 and 3 face $\{\tau_{12}, \tau_{13}\}$ and their technology parameters. With these restrictions, I can make the following inferences:

A.1 Because both countries import similar shares from the United States ($X_{21} = X_{31}$), the assumption on trade costs implies that $w_2^{\frac{-1}{\theta}} \lambda_2 = w_3^{\frac{-1}{\theta}} \lambda_3$. The interpretation is that both countries 2 and 3 can produce a good at the same unit cost *on average*.[12] Thus, because both countries import similar amounts from the United States, their cost to produce a good domestically must be the same, on average.

B.1 Because the United States imports more from country 2 than country 3 ($X_{12} > X_{13}$), but both countries can produce a good at similar cost, then this pattern implies that $\tau_{12} < \tau_{13}$. That is, *exporting* to the United States is more difficult for country 3 than it is for country 2.

These restrictions have the following implications for prices and income differences:

Example 1: Implications for Prices. Using equation (1), the aggregate price index of tradable goods in countries 2 and 3 is

$$p_2 = \Upsilon \left\{ w_2^{\frac{-1}{\theta}} \lambda_2 + \bar{\tau}^{\frac{-1}{\theta}} \right\}^{-\theta} \quad = \quad p_3 = \Upsilon \left\{ w_3^{\frac{-1}{\theta}} \lambda_3 + \bar{\tau}^{\frac{-1}{\theta}} \right\}^{-\theta},$$

with both countries facing the same price indices. Two forces are present: First, the prices paid for domestically produced goods in country 3 are (on average) the same as in country 2. Second, the cost to import a good from the United States is the same across both countries 2 and 3, so the prices paid for imported goods are (on average) the same. This example is qualitatively consistent with the data: comparable price indices for tradable goods are similar across countries.

Example 1: Implications for Income Differences. This example implies that unit costs are the same across countries 2 and 3, on average. Unit costs are a function of both wages and productivity. So to infer differences in productivity, I use balanced trade, which pins down wage rates and then allows for the recovery of λ_i. Balanced trade from equation (3) implies that

$$\frac{X_{12}}{X_{21}} = w_2 \quad > \quad \frac{X_{13}}{X_{31}} = w_3.$$

[12] To see this interpretation, take the expectation of $z_i w_i^{\frac{-1}{\theta}}$ over all goods and then take this term to the power $-\theta$. The value is $w_i \lambda_i^{-\theta}$, which is the unit cost of a producer with the average productivity level.

Notice that these wages are independent of the assumed structure of trade costs. This observation implies that country 2 is more productive than country 3 ($\lambda_2 > \lambda_3$). Conditional on fitting the pattern of trade, λ's map directly into cross-country income differences (see equation (14)), so country 2 must be richer than country 3. Nothing deep has been revealed here (yet), but below I contrast the magnitude of the productivity differences across the two scenarios demonstrating how these implications are informative.

C. Example 2: Import Effects

Now consider an alternative restriction on the parameter space so $\{\tau_{12}, \tau_{13}\} = \bar{\tau}$, i.e., countries 2 and 3 face the same cost to export to the United States. The free parameters here are the import costs for countries 2 and 3 $\{\tau_{21}, \tau_{31}\}$ and their technology parameters. With these restrictions and the data, I can make the following inferences:

A.2 Because the United States imports more from country 2 than from 3 ($X_{12} > X_{13}$), but both countries 2 and 3 face the same cost to export to the United States ($\{\tau_{12}, \tau_{13}\} = \bar{\tau}$), then $w_2^{\frac{-1}{\theta}} \lambda_2 > w_3^{\frac{-1}{\theta}} \lambda_3$. In contrast to the previous example, the interpretation is that country 2 is a lower cost producer than country 3, on average.

B.2 Abstracting from trade costs, because country 2 is a lower cost producer than country 3, country 3 should import a larger share from the United States than country 2. In the example, the two countries import the same share, $X_{21} = X_{31}$ from the United States. Thus, *importing* a good from the United States is more difficult for country 3 than it is for country 2, i.e. $\tau_{21} < \tau_{31}$.

These restrictions have the following implications for prices and income differences:

Example 2: Implications for Prices. Using equation (1), the price index for prices paid for goods in countries 2 and 3 are

$$p_2 = \Upsilon \left\{ w_2^{\frac{-1}{\theta}} \lambda_2 + \tau_{21}^{\frac{-1}{\theta}} \right\}^{-\theta} \quad < \quad p_3 = \Upsilon \left\{ w_3^{\frac{-1}{\theta}} \lambda_3 + \tau_{31}^{\frac{-1}{\theta}} \right\}^{-\theta}.$$

The middle income country has a lower aggregate price of tradables relative to the poor country. Two forces are present: First, the prices paid for domestically produced goods in country 3 are (on average) higher than in country 2 because country 2 is a lower cost producer. Second, because it is more costly for country 3 to import goods from the United States, it is also paying higher prices on goods it does import, thus further increasing the aggregate price index. This example is *not* qualitatively consistent with the data.

Example 2: Implications for Income Differences. This example implies that country 2 is a lower cost producer than country 3, on average. Because wages are higher in country 2 than country 3 (see above), the implication is that country 2 is more productive than country 3, which is similar to Example 1. However, between the two examples there is a distinct difference in magnitude; Example 2 is larger than Example 1:

$$\underbrace{\frac{\lambda_2}{\lambda_3}}_{\text{Import Effects}} > \underbrace{\frac{\lambda_2}{\lambda_3}}_{\text{Export Effects}} > 1.$$

Conditional on fitting the pattern of trade (which both scenarios do), productivity differences map into cross-country income differences. Thus, the variance in income per worker in the model helps cross-validate the restrictions imposed on the model.

V. Estimating Technology and Trade Costs from Trade Data

This section outlines my approach to jointly recover estimates of the trade costs and technology parameters from trade data consistent with the arguments of Sections III and IV.

A. Benchmark Approach

As discussed in Eaton and Kortum (2002), the framework here nests a structural log-linear "gravity" equation. To derive this relationship, simply divide each country i's trade share from country j by country i's home trade share. Taking logs yields $N-1$ equations for each country i:

$$\log\left(\frac{X_{ij}}{X_{ii}}\right) = S_j - S_i - \frac{1}{\theta}\log\tau_{ij}, \tag{8}$$

in which S_i is defined as $\log\left[r_i^{\frac{\alpha\beta}{\theta}} w_i^{\frac{(1-\alpha)\beta}{\theta}} p_i^{\left(\frac{1-\beta}{\theta}\right)} \lambda_i\right]$. Recall from Example 1 an interpretation of the value of S_i is that it is a monotonic and decreasing function of the unit cost of a producer with the average productivity level in country i. The S_is are recovered as the coefficients on country-specific dummy variables given the imposed restrictions on how trade costs can covary across countries.

Similar to Section IV, there are only $N^2 - N$ informative moments and N^2 parameters of interest. Thus, restrictions on the parameter space are necessary. To do so, I build on the arguments in Section IV and assume that trade costs take the following functional form:

$$\log(\tau_{ij}) = d_k + b_{ij} + ex_j + \varepsilon_{ij}. \tag{9}$$

Here, trade costs are a logarithmic function of distance, where d_k with $k = 1, 2, ..., 6$ is the effect of distance between country i and j lying in the kth distance intervals. Intervals are in miles: $[0, 375)$;

$[375,750)$; $[750,1500)$; $[1500,3000)$; $[3000,6000)$; and $[6000,\text{maximum}]$. b_{ij} is the effect of a shared border in which $b_{ij}=1$, if country i and j share a border and zero otherwise. I assume ε_{ij} reflects barriers to trade arising from all other factors and is orthogonal to the regressors. These features of the trade cost function are the same as in Eaton and Kortum (2002), and, more generally, consistent with the entire literature on estimating trade costs from bilateral trade flows (see Anderson and van Wincoop (2004)).

An important difference lies in the term ex_j which I will call an exporter fixed effect. Example 1 in Section IV motivates this term. To see the connection with Example 1, consider the matrix of trade costs (abstracting from distance and shared borders) in a three-country world:

$$\begin{pmatrix} 1 & \exp(ex_1) & \exp(ex_1) \\ \exp(ex_2) & 1 & \exp(ex_2) \\ \exp(ex_3) & \exp(ex_3) & 1 \end{pmatrix}.$$

The functional form assumption in equation (9) generates the same pattern of trade costs discussed in Example 1, i.e., $\{\tau_{21},\tau_{31}\}=\bar{\tau}=\exp(ex_1)$.

To summarize, equations (8) and (9) provide the basis for the benchmark estimation of trade costs τ_{ij}s and S_is for which I use ordinary least squares.[13]

B. Alternative Approach

I will contrast my approach with an alternative trade costs specification:

$$\log(\tau_{ij}) = d_k + b_{ij} + m_i + \varepsilon_{ij}. \tag{10}$$

In contrast to an exporter fixed effect, this specification has an importer fixed effect m_i in its place. Here, the estimate m_i reflects the extra cost country i faces to import a good from any country j. Of note, this is the specification of Eaton and Kortum (2001, 2002).

To connect this specification with Example 2, consider the matrix of trade costs (abstracting from distance and shared borders) in a three-country world:

$$\begin{pmatrix} 1 & \exp(m_2) & \exp(m_3) \\ \exp(m_1) & 1 & \exp(m_3) \\ \exp(m_1) & \exp(m_2) & 1 \end{pmatrix}.$$

The functional form assumption in equation (10) generates the same pattern of trade costs discussed in Example 2, i.e., $\{\tau_{12},\tau_{13}\}=\bar{\tau}=\exp(m_1)$.

[13] I also experimented with the Poisson pseudo-maximum-likelihood estimator advocated by Santos-Silva and Tenreyro (2006). I employed their technique of estimating the gravity equation in levels and including zero observed trade flows, and found that the substantive contributions of this paper do not differ relative to using ordinary least squares. Appendix D contains a more detailed discussion.

I must emphasize that both of these specifications will fit the pattern of bilateral trade equally well. However, as argued in Section IV, price and income data help in distinguishing between the two.

C. Recovering Technology

I recover the λ's in the following way. Given the estimated $\hat{S}_i$s and $\hat{\tau}_{ij}$s, the estimated aggregate price of tradable goods is then computed as

$$\hat{p}_i = \Upsilon \left\{ \sum_{j=1}^{N} e^{\hat{S}_j} \hat{\tau}_{ij}^{\frac{-1}{\theta}} \right\}^{-\theta}. \tag{11}$$

Then, given the $\hat{p}_i$s computed from equation (11), one can recover the convolution of $r_i^{\frac{\alpha\beta}{\theta}} w_i^{\frac{(1-\alpha)\beta}{\theta}} \lambda_i$ from the estimates of $\hat{S}_i$. I then use wages from observed bilateral trade shares X_{ij}, each country's observed labor endowment, and the balanced trade condition in equation (3):

$$w_i = \left(\sum_{j=1}^{N} \frac{L_j}{L_i} w_j X_{ji} \right).$$

Wages in combination with aggregate capital-labor ratios determine rental rates. Then with all prices computed, each country's technology parameter λ_i is recovered.[14]

VI. Estimating θ

The parameter θ is important for two reasons. First, the size of θ controls the gains from trade. Second, one may speculate whether results in Section III are driven by the assumption that θ is common to all countries. In this section, I estimate θ consistent with the model and from data which includes both rich and poor countries.[15] I also apply my estimation strategy to subsamples of rich and poor countries to evaluate the possibility that different θ's between rich and poor countries could explain the results in Section III.

A. Estimation Approach

To estimate θ, I employ the following arguments. First, if one had data on trade costs τ_{ij}, aggregate price data p_i, and trade data X_{ij}, then equation (4) can be used to estimate θ. The difficulty is that the trade

[14] My approach differs substantially relative to Alvarez and Lucas (2007). Their baseline calibration assumed that each country's λ_i is proportional to an unobservable endowment L_i. This assumption, in combination with balanced trade, output data, and some proxies for trade costs such as average tariff rates, allowed them to calibrate each country's λ_i and L_i jointly.

[15] Eaton and Kortum (2002) provided point estimates for θ for only OECD countries. Several papers have estimated parameters resembling θ from limited samples of developed countries. See Anderson and van Wincoop (2004) for a survey.

costs τ_{ij} are unobserved. However, an estimate of τ_{ij} is possible with a simple arbitrage argument and *disaggregated* price data. The idea is that it must be the case that for any given good ℓ at a disaggregate level, $\frac{p_i(\ell)}{p_j(\ell)} \leq \tau_{ij}$, otherwise there would be an arbitrage opportunity. This implies an estimate of τ_{ij} is the maximum of relative prices over goods ℓ.

In logs, my estimate of τ_{ij} will use the second-order statistic rather than the maximum:

$$\log \hat{\tau}_{ij} = \max_{\ell} 2 \left\{ \log\left(p_i(\ell)\right) - \log\left(p_j(\ell)\right) \right\}, \tag{12}$$

where max 2 denotes the second highest value. Eaton and Kortum (2002) use this approach generating their preferred estimate of θ. They argue that this approach helps alleviate any measurement error and find that their estimates of τ_{ij} when computed with the second-order statistic are more correlated with the normalized bilateral trade shares $(\frac{X_{ij}}{X_{jj}})$ than when computed using the first-order statistic. Consistent with their results, I also find the same outcome.

With an estimate of τ_{ij} from equation (12), I can use equation (4), bilateral trade data, and aggregate price data to estimate θ. This is my benchmark estimation approach.

B. Disaggregated Price Data

To construct estimates of equation (12), benchmark price data for the year 1985 were used from the Penn World Table website (http://pwt.econ.upenn.edu). I use the year 1985 because it provides disaggregated price data at a level higher than publicly available data for the other benchmark year 1996. The assumption is that θ is time invariant. For tradable goods, this data set has 76 different good categories for 43 of the countries in my data set. The bilateral trade data are also from 1985 and constructed in the same manner as described in Section III.

C. A Benchmark Estimate of θ

To summarize, I estimated equation (4) with proxies for trade costs from equation (12) with ordinary least squares with no intercept term as the theory predicts.[16] Least squares yields an estimate of 5.5. This implies a θ of 0.18 and it is my benchmark estimate of θ throughout the paper.[17]

[16]The estimated trade costs from equation (12) further confirm my arguments. I ran the regression $\log \hat{\tau}_{ij} = d_k + b_{ij} + \rho_{im,y} \log(y_i) + \rho_{ex,y} \log(y_j)$, where d_k and b_{ij} control for distance and shared border as described in Section V. The last two variables are importer income and exporter income per worker. The estimates are $\rho_{im,y} = -0.06$ and $\rho_{ex,y} = -0.14$ and are precisely estimated. Because the estimates are such that $\rho_{ex} < \rho_{im}$, they imply that poor countries face systematically higher costs to export their goods relative to rich countries.

[17]Method of moments generates similar results as well. Eaton and Kortum (2002) use this approach to generate their preferred estimate of $\theta = 0.12$ — lower than the estimates here. However, for only OECD countries, the method of moments estimate of θ is 0.118 and is very similar to their preferred estimate.

1. Rich Country θ Higher Than Poor Country θ?

One may speculate whether the results in Section III are driven by the assumption that θ is common to all countries. A scenario that might be able to rationalize the results in Section III would have rich countries with larger θ's (i.e., more variation in productivity and more incentives to trade) and poor countries with smaller θ's (i.e., less variation in productivity and fewer incentives to trade).[18] To evaluate this claim, I divided the sample into rich (OECD) and poor (non-OECD) countries and estimated a separate θ for each one. Because trade between rich countries will be driven by higher θ's and trade between poor countries will be driven by lower θ's, this strategy should help address this possibility.

When non-OECD countries and OECD countries are considered separately, the estimates of $1/\theta$ are 5.5 and 7.9 respectively. Notice that among non-OECD countries the estimate of $1/\theta$ is nearly the same as the estimate from the sample with all countries. Furthermore, when only OECD countries are considered, the point estimate for θ is *smaller* than the estimate for only non-OECD countries. This evidence is contrary to the argument that rich countries have higher θ's relative to poor countries as an explanation for the results in Section III.[19] In fact, these results would seem to further deepen the puzzle.

VII. Measurement and Common Parameter Values

A. Measuring Income Per Worker

Throughout the rest of this paper, real GDP per worker in my model—as measured in the PWT—is a central object of interest. Appendix B provides a detailed discussion about the mapping between real GDP as measured in the PWT and the model, but the end result is simply the following:

$$y_i = \frac{w_i}{p_i^y} + \frac{r_i k_i}{p_i^y}, \tag{13}$$

in which income from wages and capital is deflated by each country's final goods price.

B. Factor Shares

Given the model's structure resulting in equation (14), I want α to be consistent with the exercises in the income accounting literature. To do so, I set α equal to 1/3. Gollin (2002) provides an argument for

[18]Fieler (2007) is a formal articulation of this idea in combination with nonhomothetic preferences aiming to explain the observed positive correlation between aggregate trade and income per capita.

[19]The micro-evidence on productivity variation suggests this view as well. Banerjee and Duflo (2005) discuss industry-level evidence from India and they argue that the best firms are basically using frontier technologies. Lewis (2004) argues that this same pattern prevails in other developing countries such as Brazil, Russia, and Korea. Lagakos (2008) provides evidence for the retail sector in developing countries that have similar compositional patterns. Thus, the coexistence of high productive firms and extremely low productive firms in developing countries suggests that the variance in productivity is higher in poor countries than it is in rich countries.

setting α equal to 1/3 by calculating labor's share for a wide crosssection of countries and finding it to be around $2/3$.

Value added in tradable goods production is controlled by β. Since tradable goods are assumed to correspond with manufactures, one measure of β is manufacturing value added relative to gross manufacturing production. Using manufacturing value added and gross production data from UNIDO (1996), I calculated that 0.33 is the average across 61 of the countries with data available. Across all countries, there is no correlation between the calculated β and a country's level of development. Based on this evidence, 0.33 seems to be a reasonable value for β.

Value added in nontradable goods production is controlled by γ. It also determines the allocation of labor and capital between tradables and nontradables. Because of the variations in interpretation, there is wider variation on the appropriate value for γ. Alvarez and Lucas (2007) discuss plausible values for γ ranging from 0.80 to 0.70 depending on the source. As a baseline value, I use their baseline value with γ equal to 0.75.

I followed Alvarez and Lucas (2007) in selecting the value for η. Other than satisfying the necessary assumptions detailed in Alvarez and Lucas (2007), this value plays no quantitative role. i

C. Capital, Labor, and Distance Data

To compute equilibrium prices, one needs measures of endowments. For this purpose, I used aggregate capital-labor ratios from Caselli (2005). They were constructed using the perpetual inventory method with PPP adjusted investment rates in Heston, Summers and Aten (2002). I used labor endowments from Caselli (2005), which are based on information in Heston et al. (2002) as well.

The distance measures used to estimate trade costs are in miles from capital city in country i to capital city in country j, calculated by the great circle method. These measures and border data are from Centre d'Etudes Prospectives et d'Informations Internationales (http://www.cepii.fr).

VIII. Estimation Results

In this section, I discuss the estimated parameters from the benchmark model with exporter fixed effects and contrast the results with a model estimated with importer fixed effects.

A. Benchmark Results

In terms of fitting bilateral trade flows, my approach performs substantially better than standard approaches in the gravity literature with symmetric trade costs. For example, my model's R^2 is 0.83.

Table 2: Estimation Results

Summary Statistics			
No. Obs	TSS	SSR	σ_ε^2
4242	4924	851	2.08
Geographic Barriers			
Barrier	Parameter Estimate	S.E.	%effect on cost
$[0,375)$	-4.66	0.21	133.3
$[375,750)$	-5.60	0.14	177.1
$[750,1500)$	-6.16	0.09	206.3
$[1500,3000)$	-7.22	0.06	271.3
$[3000,6000)$	-8.44	0.04	363.9
$[6000,\text{maximum}]$	-9.37	0.05	449.7
Shared border	0.69	0.16	-13.0

Note: All parameters were estimated by OLS. For an estimated parameter $\hat{b}$, the implied percentage effect on cost is $100 \times (e^{-\theta\hat{b}} - 1)$ with $\theta = 0.18$.

Without the exporter fixed effect, the R^2 declines to 0.61. The simple suggestions in Section V improve the performance of standard approaches to estimating gravity models by 36 percent.

My model does not suffer from the criticisms of Fieler (2007) and captures the positive relationship between aggregate trade and income per worker. Fieler (2007) argues that structural gravity models predict that the share of trade in GDP is strongly decreasing in income per worker, yet in the data the share of trade in GDP is weakly increasing. A regression of the logarithm of imports relative to GDP predicted by my model on log income per worker yields a slope coefficient of 0.045. In the data, this same regression yields a slope coefficient of 0.057.

Note two features about the estimated trade costs. First, consistent with the gravity literature, distance is an impediment. The estimates reported are consistent with those in Eaton and Kortum (2001), which considers a similar sample of countries but only trade data on machinery and equipment. The overall size of the trade costs for a developed country are only slightly larger than those reported in Anderson and van Wincoop (2004).[20]

Second, the exporter fixed effect is negatively correlated with income per worker. Figure 2 plots each country's exporter fixed effect, expressed in terms of the percentage effect on cost, versus income per worker data. As Figure 2 depicts, poor countries appear to have a serious disadvantage at exporting goods relative to rich countries. For example, a good *exported* from the United States costs 62 percent less than a good exported from the average country. In contrast, a good *exported* from Zimbabwe will

[20] Anderson and van Wincoop (2004) report that for a representative developed country, trade barriers fall in a range between 40 and 90 percent depending on the study and elasticities of substitution. I find that the median trade cost between OECD countries is equivalent to a 118 percent tariff. This is above the upper range, but I am using a slightly lower value of θ than looking at only OECD countries would imply.

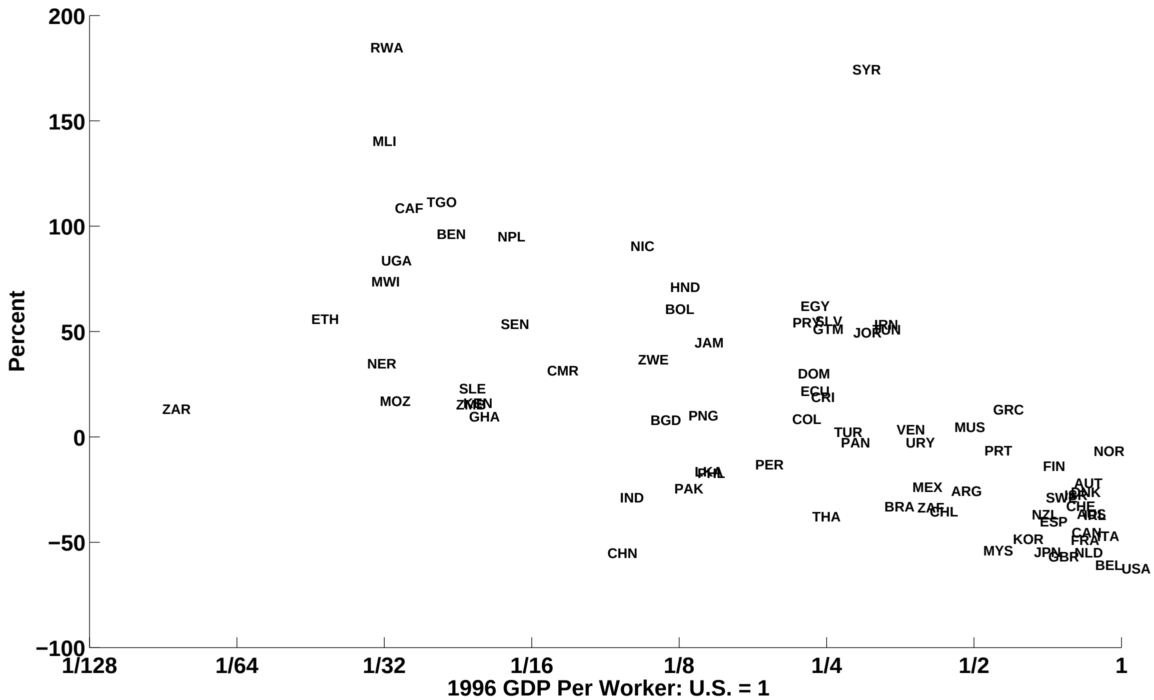

Figure 2: **Exporter Fixed Effect: Easy for Rich Countries to Export, Difficult for Poor Countries**

cost 35 percent more than a good exported from the average country. These results are consistent with the arguments of Section III, which imposed no restrictions on the structure of trade costs.

Table 3 presents the estimated S_i terms and the recovered technology parameters. These values have more meaning relative to the values when the model is estimated with importer effects. Hence, I will delay my discussion.

B. A Comparison to the Model with Importer Fixed Effects

The alternative approach to generating asymmetries in trade costs is to use the functional form assumptions with an importer fixed effect. This is the same approach used in Eaton and Kortum (2001, 2002). Regarding the overall fit of the trade data, the model with importer fixed effects performs as well as the model with exporter fixed effects. In fact, it has the same fitted values.

Using importer fixed effects yields the same effects of distance and shared borders. Furthermore, the estimated value of a country's importer fixed effects is the same as that country's estimated exporter fixed effect. The interpretation is different. Now a good *imported* by the United States costs 62 percent less than a good imported by the average country. In contrast, a good *imported* from Zimbabwe will cost 35 percent more than a good imported by the average country. These results are inconsistent with the arguments of Section III, which impose no restrictions on the form that trade costs can take.

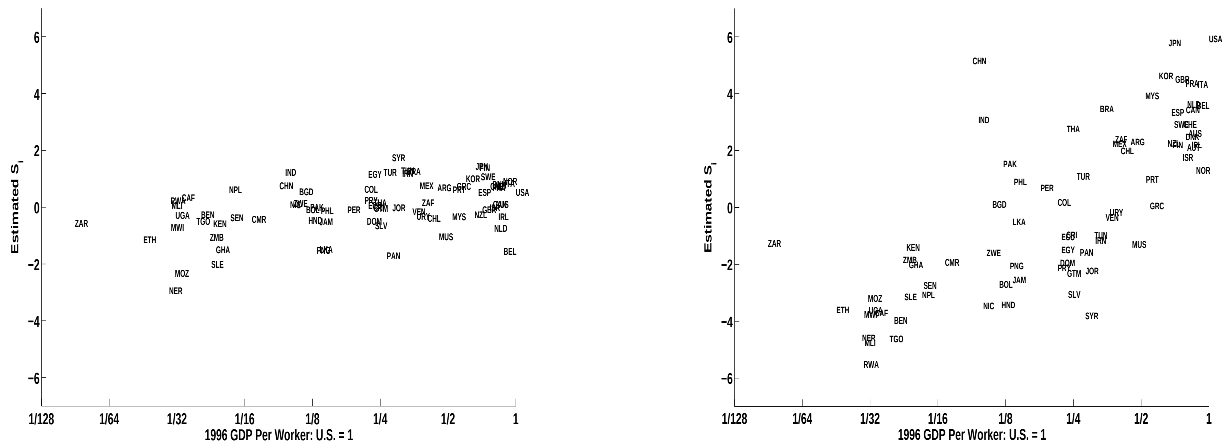

(a) S_i from Model with Exporter Fixed Effect

(b) S_i from Model with Importer Fixed Effect

Figure 3: Estimated S_i versus GDP Per Worker.

The key difference between the two approaches is the estimated S_is. Figure 3(a) plots the S_is when estimated from a model with exporter fixed effects versus income per worker data. Figure 3(b) plots the S_is when estimated from a model with importer fixed effects versus income per worker data. When the model is estimated with importer fixed effects, the S_is covary strongly with income per worker data.

Recall the interpretation of S_i: they are a decreasing and monotonic function of the unit cost of a producer with the average productivity level in each country. For example, Figure 3(a) implies that the unit costs of a producer with the average productivity level are similar across countries. In contrast, Figure 3(b) implies that the unit costs of a producer with the average productivity level is strongly decreasing with the level of development. As discussed in Section IV, the differences in how unit costs covary across countries maps into differences in the aggregate price of tradables and income differences, which I will discuss next.

IX. The Quantitative Implications of the Estimated Model

As an assessment of the model, I compare the implications of the model for moments not explicitly used in the estimation of the model. In all the results below, I used the estimated $\hat{\lambda}_i$s and $\hat{\tau}_{ij}$s, computed an equilibrium, and compared the model-implied aggregate price of tradables and variation in income per worker relative to the data.

A. The Benchmark Model

Figure 4(a) plots the prices from the benchmark model and the data. Broadly speaking, the prices from the model are in line with the data: both rich and poor countries have similar prices of tradable goods.

There is room for improvement. The elasticity with respect to income level is approximately -0.04, whereas in the data this elasticity is 0.15. There is also more variance in the data than in the model. I should emphasize the parsimony of my approach; there are many more degrees of freedom available to improve the model on a variety of dimensions that I did not use. Furthermore, the benefits of my approach should be compared to the alternative discussed below.

As another assessment, I considered the model's ability to quantitatively replicate the cross-country income differences seen in the data. Figure 5 depicts the model's income levels versus the data relative to the United States along with the 45° line. In Figure 5, the ordered pairs lie around the 45° line, suggesting that the model accurately captures the variation in income per worker across countries. For example, the model predicts that Uganda has an income level 1/27 of the U.S. level. In the data, Uganda has an income level 1/32 of the U.S. level. As a measure of dispersion for all countries, the variance of log income per worker in the model is 1.30 relative to 1.38 in the data, and the 90/10 percentile ratio is 25.6 in the model relative to 25.7 in the data.

B. The Implications of the Estimated Model with Importer Fixed Effects

Figure 4(b) plots the prices from a model estimated with importer fixed effects. These prices systematically deviate from the data, with poor countries facing systematically higher prices relative to rich countries. For example, the elasticity with respect to income level is -0.29.[21] This is seven times larger than my model with an exporter fixed effect.

The model estimated with importer fixed effects does poorly at replicating the variation in income per worker: the variation of log income per worker is 2.46, and the 90/10 percentile ratio is 89.5. Recall that the variance of log income per worker and the 90/10 percentile ratio in the data are 1.38 and 25.6. The reason for the differing implications for cross-country income differences relates to the results regarding prices. The model with an importer fixed effect results in systematically higher prices of traded goods for poor countries mapping into higher prices of nontraded goods p_i^y. Because wages and rental income are deflated by p_i^y to compute real income per worker, this systematically lowers real income per worker for poor countries, resulting in an overprediction of income differences.

C. A Brief Discussion

To clarify the forces driving these outcomes, recall that the model with importer fixed effects reconciles the fact that the United States imports more from Japan than Senegal by making unit costs of production

[21] This outcome is similar to the results in Eaton and Kortum's (2001) study of investment goods. They find that the estimated prices systematically deviate from the data, with poor countries facing higher prices relative to rich countries—similar to the outcome here. It is also similar to the criticism in Balistreri and Hillberry (2006) that Anderson and van Wincoop's (2003) estimation predicts that the cost of living in Canada is 24 percent higher than in the United States and shows evidence to the contrary. In my estimation, the price of tradables in Canada is 2 percent lower than the Unites States.

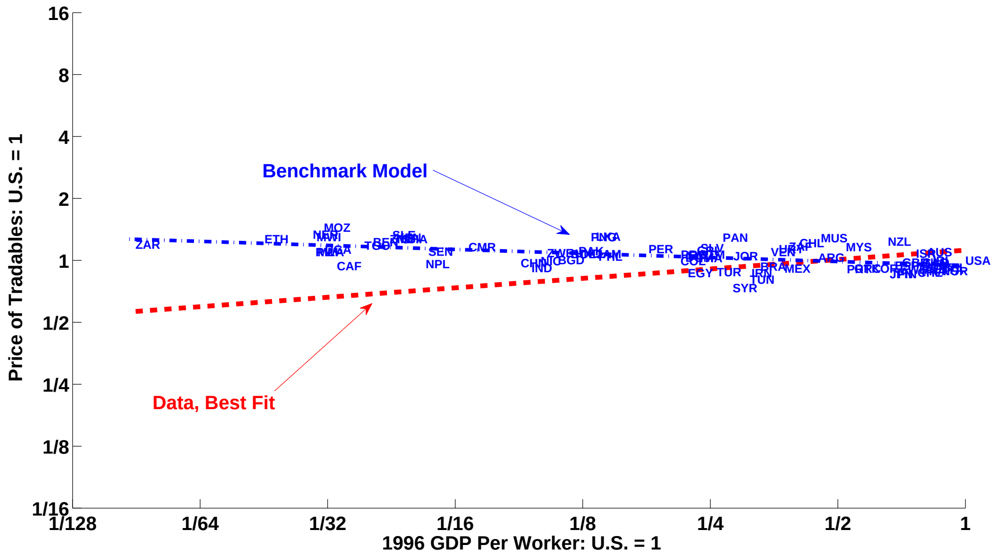

(a) Price Data and Benchmark Model

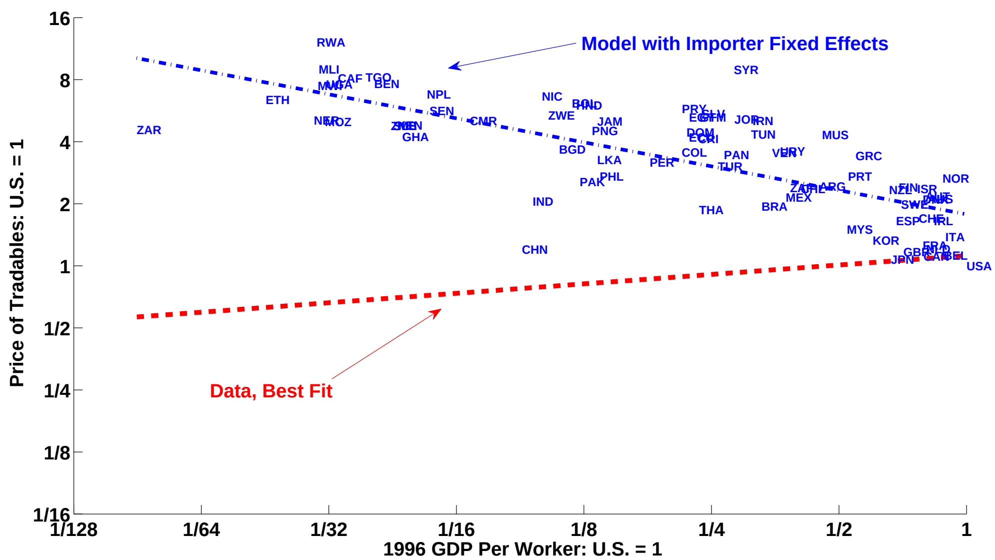

(b) Price Data and Model with Importer Fixed Effects as in Eaton and Kortum (2001, 2002)

Figure 4: Price Data and Model versus GDP Per Worker.

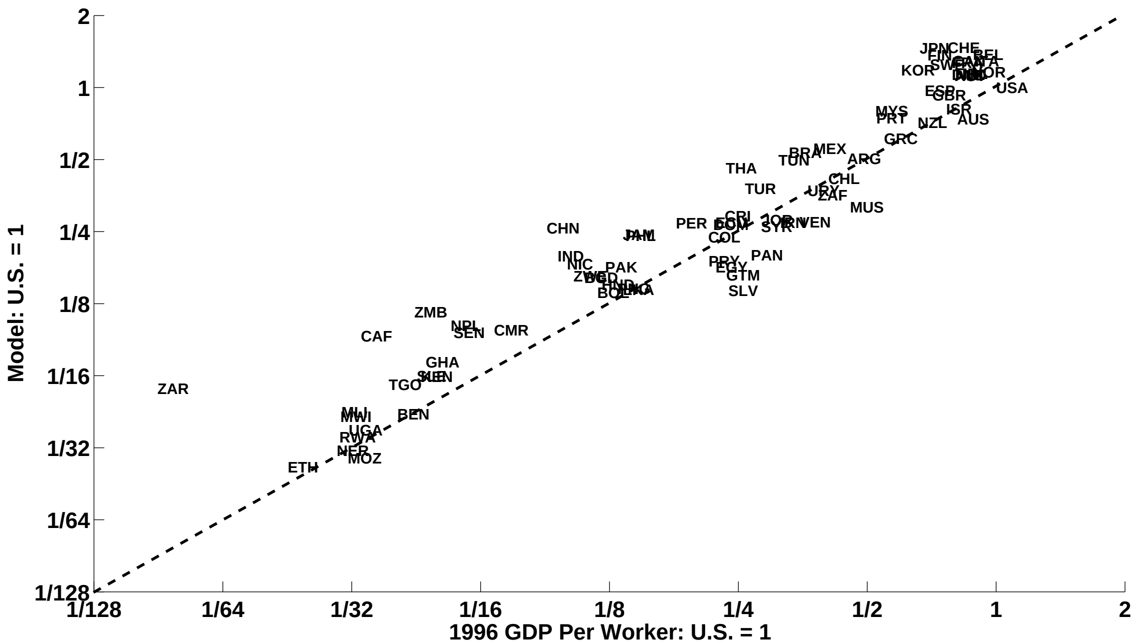

Figure 5: Income Per Worker: Data and Benchmark Model

(on average) lower in Japan than in Senegal. Figure 3(b) illustrates this. So prices paid for domestically produced goods are lower in Japan than in Senegal. But the data also say that Senegal imports a similar share of goods from the United States as does Japan. The model reconciles this fact by increasing the cost for Senegal to import relative to Japan's cost to import. Thus, the prices paid for imported goods are higher in Senegal as well. Together, these implications mean the price index in Senegal is higher than in Japan—inconsistent with the data.

In contrast, my model with exporter fixed effects reconciles the differences in the United States' import share from Japan relative to Senegal by manipulating each country's export cost. The model reconciles the similarities in Japan's and Senegal's import share from the United States with similar unit costs of production. Figure 3(a) illustrates this. Thus, prices paid for domestic goods and imported goods in Japan and Senegal are similar, implying that the price index of tradables is similar in the two countries.

In general, the model's implications for the price of tradable goods is highly robust to various parameterizations. I should note that the model's implications for the variation in income per worker are sensitive to the calibration of γ. A larger γ in both models results in less variation in income per worker. I do not view this as a problem because the model should probably underpredict the variation in income per worker. Hence, with larger values of γ, my model is more in line with the data than alternative approaches. I say that because I assumed no productivity differences in the nontradable sector. Allowing for this source of variation would increase the variation in income per worker.

X. How Do Trade Costs Affect Income Differences?

In this framework, income differences are driven by differences in technologies, endowments, and trade costs. Trade costs prevent countries from specializing in their comparative advantage and hence amplify productivity differences relative to those in a world with no trade costs. Given that my model can replicate important features of the data, I now explore the quantitative importance of trade costs to understanding the large measured productivity differences across countries.

A. Eliminating Asymmetries in Trade Costs Reduces Income Differences

In this section, I will study two exercises that focus on eliminating the systematic asymmetry in trade costs. A way to view this systematic asymmetry is like a "wedge" or deviation from a benchmark model with only symmetric trade costs, in the spirit of Chari, Kehoe and McGrattan (2007). And these exercises quantify the importance of these wedges for understanding income differences across countries.

In the first experiment, I adjusted trade costs so the new costs to trade between two countries are $\hat{\tau}_{ij} = \min(\tau_{ij}, \tau_{ji})$. Given how the recovered trade costs impact poor countries relative to rich countries, the premise is that costs above this minimum reflect some extra distortion one country faces while the other does not and that these distortions, unlike distance, are not natural. With the new trade costs, the variation in log income per worker declines to 1.05 and the 90/10 ratio is only 17.3. Cross-country income differences decline by up to 31 percent relative to the baseline model. In this exercise, the new trade costs are still large. For example, the median trade cost is 192 percent for all countries, which is twice the value reported in Anderson and van Wincoop (2004) for a developed country. Table 4 also reports the average gain in income per worker. All countries gain—but poor countries gain relatively more than rich countries.

As another experiment, I endow all countries with the same effective distance costs OECD countries face among themselves. The premise is that the costs to trade between OECD countries are relatively closer to free trade than the entire sample. Table 4 presents the results. In this experiment, the variance in log income per worker is reduced by 13 percent and the 90/10 ratio is reduced by 23 percent.

Table 4: Income Differences with Counterfactual Trade Costs

	Baseline	Autarky	$\min(\tau_{ij}, \tau_{ji})$	OECD τ	$\tau_{ij} = 1$
$\mathrm{var}[\log(\mathrm{y})]$	1.30	1.35	1.05	1.13	0.76
$\mathrm{y}_{90}/\mathrm{y}_{10}$	25.7	23.5	17.3	19.8	11.4
Mean change in y, percent	—-	–10.5	24.2	10.0	128.0

To put these two exercises in perspective, compare these gains relative to a world with no trade costs. In this world, the variation in log income per worker is only 0.76 and the 90/10 ratio is 11.4. Eliminating the asymmetry or moving toward OECD trade costs delivers 59 and 41 percent of the reduction in income differences relative to complete elimination of trade costs. Or at the other extreme, if trade costs were set to infinity, the variance of log income per worker increases to only 1.35. The systematic asymmetry in trade costs is so punitive that removing it takes the economy from basically autarky to over 50 percent of the way relative to frictionless trade.

B. On The Mechanics Behind Reductions in Income Differences

Differentials in the pattern of specialization between rich and poor countries drive reductions in income differences. To see this, I show in Appendix C that starting from (13) the definition of PWT real GDP per worker in the model can be expressed as

$$y_i = A_i k_i^{\alpha}, \quad \text{with } A_i = X_{ii}^{\frac{-\theta(1-\gamma)}{\beta}} \lambda_i^{\frac{\theta(1-\gamma)}{\beta}}. \tag{14}$$

Real GDP per worker is similar to a standard one-sector growth model with a TFP term and capital-labor ratio taken to a power term. Here, measured TFP is decomposed into an endogenous trade factor, $X_{ii}^{\frac{-\theta(1-\gamma)}{\beta}}$, and an exogenous domestic factor, $\lambda_i^{\frac{\theta(1-\gamma)}{\beta}}$.[22]

When trade costs are infinite, countries are unable to specialize and must produce all goods domestically. Hence, X_{ii} = 1 and TFP becomes $\lambda_i^{\frac{\theta(1-\gamma)}{\beta}}$ Given how efficiency levels are distributed in the production of tradable goods, each country's average efficiency level is λ_i^{θ} taken to the power $\frac{(1-\gamma)}{\beta}$.[23] When trade costs are finite, countries are able to specialize and import some goods from relatively more efficient producers. Hence, $X_{ii} < 1$ and each country's gain from trade in the form of increased TFP is $X_{ii}^{\frac{-\theta(1-\gamma)}{\beta}}$.

Equation (14) implies that the only way for income differences to decline is with X_{ii} changing more for poor countries relative to rich countries. Observation **O.1** argued that X_{ii} is similar across rich and poor countries. Asymmetries in trade frictions are the reason for this. When these asymmetries are removed, poor countries' X_{ii} decline more than rich countries' X_{ii}. This means that poor countries drop more unproductive activities and scale up the productive activities relative to rich countries. This differential change in activities results in a relative increase in measured TFP and a reduction in income differences.

[22]Independently, Finicelli, Pagano and Sbracia (2007) developed a similar approach to measure how much international competition raised manufacturing TFP for OECD countries.

[23]These power terms reflect the fact that tradables operate like intermediate goods. For example, if $\gamma = 0$, then this expression is similar to an open-economy version of Jones (2008), with technology differences amplified to the power $1/\beta$ reflecting a multiplier because intermediates are used to produce intermediates. In my model, this multiplier is dampened because tradables are only $1-\gamma$ of total output.

XI. Robustness Checks and Alternative Evidence

A. Price Data: A Robustness Check

My arguments depend critically on how aggregate prices covary with a country's level of development. There are concerns regarding the accuracy of prices collected from the International Comparison Programme (ICP); see the discussion in Hsieh and Klenow (2007b). This issue is that if these prices are mismeasured, then for it to matter there must be a systematic measurement error with respect to a country's level of development. Given this observation, at minimum I can ask: What degree of systematic measurement error is necessary to reverse my results?

A simple way to answer this question is in terms of the elasticity of prices with respect to the level of development. Recall equation (5): $\frac{\tau_{ij}}{\tau ji} = \left(\frac{X_{ij}X_{ii}}{X_{ji}X_{jj}}\right)^{-\theta}\left(\frac{p_i}{p_j}\right)^2$. Define ρ_X as the elasticity of $\frac{X_{ij}X_{ii}}{X_{ji}X_{jj}}$ with respect to relative income level $\frac{y_j}{y_i}$ and ρ_p as the elasticity of $\frac{p_i}{p_j}$ with respect to $\frac{y_j}{y_i}$. Equation (5) implies that ρ_τ, the elasticity of $\frac{\tau_{ij}}{\tau_{ji}}$ with $\frac{y_j}{y_i}$, is

$$\rho_\tau = -\theta\rho_X + 2\rho_p. \tag{15}$$

Given equation (15), I can ask the following questions: What price elasticity, ρ_p, is necessary to change my argument, and what is the implied magnitude of systematic measurement error?

Table 5 presents the results. The first column ρ_p is the data and the implied ρ_τ. The second column considers the case if prices were equalized. Here, trade costs would still be systematically asymmetric with respect to a country's level of development. The third column considers the case necessary to yield symmetric trade costs. In this case, the elasticity of prices with respect to income level would have to be 0.21—almost the complete opposite value seen in the data. To illustrate the magnitude, consider a country with 1/30 the income level of the United States. In the data, this country would have a tradable goods price approximately 1/2 of the U.S. level. This thought experiment suggests that a country with 1/30 the U.S. income level should face a tradable goods price two times larger than the U.S. value. That is, these prices must be systematically mismeasured by more than 337 percent for this result to disappear. This suggests the magnitude of systematic measurement error by the ICP would have to be dramatic.[24]

[24] A related issue is that the price data are not producer prices but retail prices. This fact would affect my arguments only if distribution costs were positively correlated with level of development, i.e., poor countries have lower distribution costs than rich countries. In this sense, unobserved distribution costs are isomorphic to systematic measurement error discussed here. The evidence regarding distribution costs is limited, but it suggests the opposite pattern: Burstein, Neves and Rebelo (2003) find that the distribution costs in Argentina are 50 percent higher than in the United States.

Table 5: Robustness of Asymmetric Trade Costs

	Data	Counterfactual Price Elasticities	
		$p_i = p_j$	Symmetry
ρ_p	-0.23	0	0.21
ρ_τ	-0.87	-0.41	0

Variable ρ_p is the elasticity of the relative price of tradables with respect to relative income level. Variable ρ_τ is simply computed from equation (15) and ρ_x is 2.30 and the same under all experiments.

B. Alternative Evidence on Asymmetric Trade Costs

Matched bilateral trade flows from the International Monetary Fund's (IMF) Direction of Trade Statistics provides alternative and model free evidence consistent with my arguments. This database provides bilateral trade flows for the f.o.b. (free on board) value of exports between countries and the the c.i.f. (cost, insurance, freight) value of imports between countries.[25] This is useful because the ratio of the c.i.f. value relative to the f.o.b. value provides a direct (though imperfect) measure of trade costs between countries, which I will denote as t_{ij}.[26]

To explore these direct measures of trade costs, I ran the regression: $\log(t_{ij}) = d_k + b_{ij} + \rho_{im,y}\log(y_i) + \rho_{ex,y}\log(y_j)$. Variables d_k and b_{ij} control for distance and shared border as described in Section V. The key variables are the last two: importer income and exporter income per worker.

The estimates of $\rho_{im,y}$ and $\rho_{ex,y}$ are 0.05 and -0.08, and both are statistically different than zero. If these estimates were the same, then there would be no systematic asymmetry in trade costs. Since $\rho_{ex} < \rho_{im}$, the implication is that poor countries face systematically higher costs to export their goods relative to rich countries. For concreteness, Zimbabwe has 1/10th the income level of the United States. These estimates imply that Zimbabwe will face a trade cost that is 37 percent higher to export to the United States than for the United States to export to Zimbabwe. To summarize, these data provide direct evidence consistent with my arguments.

[25] The f.o.b. value is basically the quantity times the price of the goods if sold in the exporting country. The c.i.f. value is basically the quantity times the price of the goods sold in the importing country. In theory, these values should be based on border prices, i.e., the prices of the goods at each country's border. In reality, these values are based on final destination prices as Limao and Venables (2001) argue.

[26] Hummels and Lugovskyy (2006) raise concerns about the possibility of measurement error in these data. To help alleviate these concerns, I computed t_{ij} as the average over the years 1994 to 1996 for the set of countries in my sample. Furthermore, I eliminated any values less than 1 implying negative trade costs, and I dropped values equal to 1.10 because this is the value the IMF uses to impute missing data on either the f.o.b. or c.i.f. side.

1. Other Evidence

Data on policy barriers to trade also support my view of trade costs. Kee, Nicita and Olarreaga (2006) estimate trade restrictiveness indices from data on both tariff and nontariff barriers for a large set of developed and undeveloped countries. These indices summarize the multitude of barriers that exist in a compact format. They find that poor countries systematically face the highest trade barriers on their export bundle—similar to the asymmetry in trade costs here.

Moreira, Volpe and Blyde (2008) provide evidence on transportation costs that supports my view of trade costs. They study transportation costs in Latin America and find that the shipping cost to export from a Latin American country to the United States is 70 percent higher than the shipping cost to export from the Netherlands to the United States. This is in spite of the fact that the Latin American countries considered are closer to the United States than the Netherlands. More important, they argue that more than 40 percent of this difference can be attributed to variables such as infrastructure, which policy makers have some control over.

Finally, policies in poor countries could create this effect as well. Export marketing boards are one possible source of this distortion. These boards place a wedge between the price at which producers sell goods and the price at which the good is exported. Krueger, Schiff and Valdes (1988) argue that these costs can constitute up to 50 percent of more of the border price of exportables and that these marketing boards are prevalent in developing countries. Furthermore, any differential treatment of exporters versus domestic producers could create this effect as well; see Krueger et al. (1988) for a variety of examples.

C. Evidence from Trade Liberalizations

To provide more support for my model, I considered the model's implications relative to an observed trade liberalization. I considered Chile because it had a large unilateral trade liberalization; beginning in March 1974, they reduced the average tariff rate from 105 percent to 20 percent in 1977. For the time period 1983-1985, Chile briefly deviated and increased its uniform tariff rate to 35 percent, but by 1996 Chile had a uniform tariff rate of 10 percent. Chile continued to reduce its tariff rate further with a simple average tariff rate of just under 5 percent in 2007.[27]

Figure 6 plots total imports relative to GDP data for Chile over this same time period from 1960 to 2007. Note that imports relative to GDP increased from about 8 percent in 1973 (prior to the liberalization) to nearly 25 percent in 1996 and 28 percent in 2007.

As a test of my model, I performed the following experiment. First, I computed the equilibrium for the model calibrated to 1996 data as described throughout the paper. I then changed the calibrated trade costs

[27] See Corbo (1997) and Edwards and Lederman (1998) for pre-1996 tariff data. See the World Bank's World Trade Indicators (www.worldbank.org/wti2008) for more recent data.

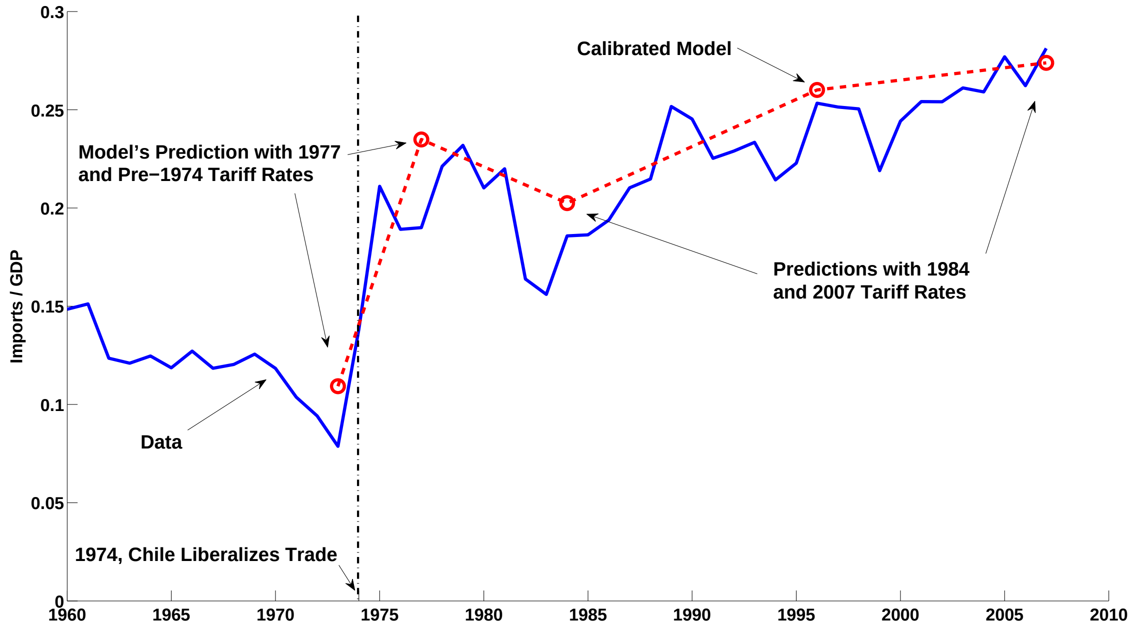

Figure 6: Chilean Trade Liberalization and Predictions from the Model

for Chile to import goods by the same amount that Chile's average tariff rate in 2007, 1984, 1977, and 1973 changed relative to the 1996 level. All other calibrated and estimated parameters are kept constant. With the changes in trade costs, I then recomputed the equilibrium and studied how much trade changed. The circles in Figure 6 summarize these results by plotting Chile's import to GDP ratio in the calibrated model for the year 1996 and the model's predicted import to GDP ratio when tariffs change.

Overall, the model does a reasonable job predicting the changes in trade when actually changes in policies are studied. The importance of this exercise is that it speaks to my model's elasticity of trade with respect to changes in trade barriers, and this same elasticity maps into the response of income with respect to trade barriers. Because of the model's performance, it provides further support for my results.

XII. Conclusion

I have argued that systematically asymmetric trade frictions are necessary to reconcile both price and quantity data in a standard model of international trade. Furthermore, these asymmetries are quantitatively important to understanding cross-country income differences. In a sense, my arguments outline a puzzle that has nontrivial implications: What are these asymmetries in the pattern of trade between rich and poor countries? Though direct evidence supports this pattern of trade costs, I will suggest two routes to further understand this puzzle that are complements rather than substitutes. One is better theory, i.e., some of these frictions may be reduced-form representations of equilibrium responses to the fundamentals faced by agents; e.g., the model of Fieler (2007) with nonhomothetic preferences is an example. An alternative route is better measurement. Measuring bilateral trade flows in value added, exploiting disaggregate trade flows, and disaggregate measures of tariff and non-tariff barriers are all possible avenues for future research as well.

References

Alvarez, Fernando and Robert Jr. Lucas, "General Equilibrium Analysis of the Eaton-Kortum Model of International Trade," *Journal of Monetary Economics*, September 2007, *54* (6), 1726–1768.

Anderson, James and Eric van Wincoop, "Gravity with Gravitas: A Solution to the Border Puzzle," *American Economic Review*, March 2003, *93*, 170–192.

____ **and** ____ , "Trade Costs," *Journal of Economic Literature*, September 2004, *42* (3), 691–751.

Baily, Martin Neil and Hans Gersbach, "Efficiency in Manufacturing and the Need for Global Competition," *Brookings Papers on Economic Activity: Microeconomics*, 1995, *2*, 307–358.

Balistreri, Edward and Russell Hillberry, "Trade Frictions and Welfare in the Gravity Model: How Much of the Iceberg Melts ?," *Canadian Journal of Economics*, February 2006, *39* (1), 247–265.

Banerjee, Abhijit and Esther Duflo, "Growth Theory through the Lens of Development Economics," in Philippe Aghion and Steven Durlauf, eds., *Handbook of Economic Growth*, Vol. 1, Amsterdam: Elsevier, 2005.

Bernard, Andrew, Jonathan Eaton, J. Bradford Jensen, and Samuel Kortum, "Plants and Productivity in International Trade," *American Economic Review*, September 2003, *93* (4), 1268–1290.

Burstein, Ariel and Alexander Monge-Naranjo, "Foreign Know-How, Firm Control, and the Income of Developing Countries," *Quartarly Journal of Economics*, February 2009, *124* (1), 149–195.

Burstein, Ariel T., Joao C. Neves, and Sergio Rebelo, "Distribution Costs and Real Exchange Rate Dynamics During Exchange-Rate-Based Stabilizations," *Journal of Monetary Economics*, September 2003, *50* (6), 1189–1214.

Caselli, Francesco, "Accounting for Cross-Country Income Differences," in Philippe Aghion and Steven Durlauf, eds., *Handbook of Economic Growth*, Vol. 1, Amsterdam: Elsevier, 2005.

Chaney, Thomas, "Distorted Gravity: The Intensive and Extensive Margins of International Trade," *American Economic Review*, September 2008, *98* (4), 1707–1721.

Chari, V. V., Patrick J. Kehoe, and Ellen R. McGrattan, "Business Cycle Accounting," *Econometrica*, May 2007, *75* (3), 781–836.

Corbo, Vittorio, "Trade Reform and Uniform Import Tariffs: The Chilean Experience," *American Economic Review*, May 1997, *87* (2), 73–77.

Dornbusch, Rudiger, Stanley Fischer, and Paul Samuelson, "Comparative Advantage, Trade, and Payments in a Ricardian Model with a Continuum of Goods," *American Economic Review*, December 1977, *67* (5), 823–839.

Eaton, Jonathan and Samuel Kortum, "Trade in Capital Goods," *European Economic Review*, 2001, *45* (7), 1195–1235.

____ **and** ____ , "Technology, Geography, and Trade," *Econometrica*, September 2002, *70* (5), 1741–1779.

Edwards, Sebastian and Daniel Lederman, "The Political Economy of Unilateral Trade Liberalization: The Case of Chile," NBER Working Paper 6510, April 1998.

Feenstra, Robert C., Alan Heston, Marcel P. Timmer, and Haiyan Deng, "Estimating Real Production and Expenditures Across Nations: A Proposal for Improving the Penn World Tables," NBER Working Paper 10866, November 2004.

Feenstra, Robert, Robert Lipsey, and Henry Bowen, "World Trade Flows, 1970-1992, with Production and Tariff Data," NBER Working Paper 5910, March 1997.

Fieler, Ana Cecilia, "Non-Homotheticity and Bilateral Trade: Evidence and a Quantitative Explanation," Working Paper, New York University, November 2007.

Finicelli, Andrea, Patrizio Pagano, and Massimo Sbracia, "Trade-Revealed TFP," Technical Report, Bank of Italy, June 2007.

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

Guner, Nezih, Gustavo Ventura, and Xu Yi, "Macroeconomic Implications of Size-Dependent Policies," *Review of Economic Dynamics*, October 2008, *11* (4), 721–744.

Hall, Robert and Charles Jones, "Why Do Some Countries Produce So Much More Output Per Worker Than Others?," *Quarterly Journal of Economics*, February 1999, *114* (1), 83–116.

Helpman, Elhanan, Marc Melitz, and Yona Rubinstein, "Estimating Trade Flows: Trading Partners and Trading Volumes," *Quarterly Journal of Economics*, May 2008, *123* (2), 441–487.

Heston, Alan, Rober Summers, and Bettina Aten, "Penn World Table Version 6.1," Technical Report, Center for International Comparisons at the University of Pennsylvania (CICUP), October 2002.

Hsieh, Chang-Tai and Peter J. Klenow, "Misallocation and Manufacturing TFP in China and India," NBER Working Paper 13290, August 2007.

——— **and** ——— , “Relative Prices and Relative Prosperity,” *American Economic Review*, June 2007, *97* (3), 562–585.

Hummels, David and Volodymyr Lugovskyy, “Are Matched Partner Trade Statistics a Usable Measure of Transportation Costs?,” *Review of International Economics*, February 2006, *14* (1), 69–86.

Jones, Charles I., “Intermediate Goods, Weak Links, and Superstars: A Theory of Economic Development,” NBER Working Paper 13834, March 2008.

Kee, Hiau Looi, Alessandro Nicita, and Marcelo Olarreaga, “Estimating Trade Restrictiveness Indices,” Policy Research Working Paper Series 3840, World Bank, February 2006.

Klenow, Peter and Andres Rodríguez-Cláre, “The Neoclassical Revival in Growth Economics: Has It Gone Too Far?,” in Ben S. Bernanke and Julio J. Rotemberg, eds., *NBER Macroeconomics Annual 1997*, Vol. 12, Cambridge, MA: MIT Press, 1997.

Kortum, Samuel, “Research, Patenting and Technological Change,” *Econometrica*, November 1997, *65* (6), 1389–1419.

Kravis, Irving B. and Robert E. Lipsey, “National Price Levels and the Prices of Tradables and Nontradables,” *American Economic Review*, May 1988, *78* (2), 474–478.

Krueger, Anne O, Maurice Schiff, and Alberto Valdes, “Agricultural Incentives in Developing Countries: Measuring the Effect of Sectoral and Economywide Policies,” *World Bank Economic Review*, September 1988, *2* (3), 255–71.

Lagakos, David, “Superstores or Mom and Pops? Market Size, Technology Adoption, and TFP Differences,” Working Paper, UCLA, January 2008.

Lagos, Ricardo, “A Model of TFP,” *Review of Economic Studies*, October 2006, *73* (4), 983–1007.

Lewis, William W., *The Power of Productivity*, Chicago, IL: University of Chicago Press, 2004.

Limao, Nuno M. and Anthony J. Venables, “Infrastructure, Geographical Disadvantage, Transport Costs, and Trade,” *World Bank Economic Review*, 2001, *15* (3), 451–479.

McKinsey Global Institute, “India: The Growth Imperative,” Technical Report, McKinsey Global Institute 2001.

Melitz, Marc J., “The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity,” *Econometrica*, November 2003, *71* (6), 1695–1725.

Moreira, Mauricio Mesquita, Christian Volpe, and Juan S. Blyde, *Unclogging the Arteries: A Report on the Impact of Transport Costs on Latin American and Caribbean Trade*, Washington D.C.: Inter-American Development Bank, 2008.

Parente, Stephen L. and Edward C. Prescott, *Barriers to Riches*, Cambridge, MA: MIT Press, 2002.

Redding, Stephen and Anthony J. Venables, "Economic Geography and International Inequality," *Journal of International Economics*, January 2004, *62* (1), 53–82.

Restuccia, Diego and Richard Rogerson, "Policy Distortions and Aggregate Productivity with Heterogeneous Plants," *Review of Economic Dynamics*, October 2008, *11* (4), 707–720.

Santos-Silva, J.M.C. and Silvana Tenreyro, "The Log of Gravity," *Review of Economics and Statistics*, November 2006, *88* (4), 641–658.

UNIDO, *International Yearbook of Industrial Statistics*, United Nations Industrial Development Organization, 1996.

Yi, Kei-Mu, "Can Vertical Specialization Explain the Growth of World Trade?," *Journal of Political Economy*, February 2003, *111* (1), 52–102.

Table 3: Country-Specific Estimates (Benchmark)

Country	ex_i	S.E.	Percent Cost	$\hat{S}_i$	S.E.	$\left(\frac{\lambda_{us}}{\lambda_i}\right)^{\theta}$	Country	ex_i	S.E.	Percent Cost	$\hat{S}_i$	S.E.	$\left(\frac{\lambda_{us}}{\lambda_i}\right)^{\theta}$
United States	5.40	0.24	−62.5	0.54	0.17	1.00	Republic of Korea	3.64	0.24	−48.3	1.00	0.17	0.76
Argentina	1.62	0.26	−25.5	0.69	0.19	1.60	Sri Lanka	0.98	0.30	−16.3	-1.48	0.22	4.44
Australia	2.50	0.25	−36.4	0.11	0.18	1.42	Mexico	1.49	0.25	−23.6	0.76	0.18	1.34
Austria	1.35	0.24	−21.8	0.77	0.17	0.93	Mali	−4.83	0.37	140	0.08	0.24	9.16
Belgium	5.13	0.24	−60.7	-1.55	0.17	1.21	Mozambique	−0.87	0.36	17.0	-2.32	0.23	13.71
Benin	−3.71	0.41	96.3	-0.25	0.23	10.40	Mauritius	−0.26	0.29	4.68	-1.04	0.20	2.47
Bangladesh	−0.43	0.27	8.03	0.54	0.21	2.92	Malawi	−3.04	0.36	73.7	-0.71	0.24	10.42
Bolivia	−2.61	0.31	60.7	-0.09	0.21	3.83	Malaysia-Singapore	4.25	0.24	−53.8	-0.33	0.17	1.10
Brazil	2.21	0.25	−33.0	1.27	0.18	1.30	Niger	−1.64	0.39	34.7	-2.94	0.25	21.29
Central African Republic	−4.04	0.52	109	0.33	0.24	3.46	Nicaragua	−3.55	0.34	90.5	0.09	0.21	3.03
Canada	3.32	0.24	−45.2	0.11	0.17	0.99	Netherlands	4.38	0.24	−54.8	-0.75	0.17	1.17
Switzerland	2.19	0.24	−32.8	0.75	0.17	0.75	Norway	0.38	0.25	−6.6	0.92	0.18	0.94
Chile	2.40	0.26	−35.2	-0.39	0.18	1.89	Nepal	−3.68	0.34	95.1	0.62	0.24	4.68
China-Hong Kong	4.40	0.24	−55.0	0.76	0.17	1.85	New Zealand	2.52	0.27	−36.7	-0.27	0.19	1.39
Cameroon	−1.50	0.30	31.4	-0.43	0.20	4.75	Pakistan	1.55	0.25	−24.5	-0.01	0.19	3.01
Colombia	−0.45	0.26	8.51	0.63	0.19	2.62	Panama	0.14	0.30	−2.5	-1.71	0.20	6.17
Costa Rica	−0.96	0.28	19.0	0.01	0.20	2.51	Peru	0.77	0.27	−13.1	-0.08	0.20	2.62
Denmark	1.67	0.24	−26.1	0.81	0.17	0.88	Philippines	1.03	0.26	−17.0	-0.12	0.18	2.44
Dominican Republic	−1.45	0.29	30.1	-0.49	0.21	2.30	Papua New Guinea	−0.53	0.38	10.1	-1.51	0.25	4.20
Ecuador	−1.09	0.29	21.9	0.06	0.20	2.76	Portugal	0.37	0.25	−6.4	0.61	0.18	1.17
Egypt	−2.66	0.27	62.0	1.17	0.19	2.84	Paraguay	−2.38	0.31	54.2	0.26	0.21	3.53
Spain	2.82	0.24	−40.1	0.53	0.17	1.03	Rwanda	−5.76	0.42	184	0.24	0.25	10.17
Ethiopia	−2.45	0.33	55.9	-1.15	0.23	12.76	Senegal	−2.37	0.33	53.7	-0.36	0.22	4.05
Finland	0.82	0.25	−13.8	1.39	0.18	0.65	Sierra Leone	−1.14	0.38	23.0	-2.01	0.26	6.37
France	3.69	0.24	−48.8	0.68	0.17	0.83	El Salvador	−2.41	0.31	55.0	-0.64	0.21	4.86
United Kingdom	4.60	0.24	−56.6	-0.08	0.17	1.12	Sweden	1.86	0.25	−28.6	1.07	0.18	0.70
Ghana	−0.51	0.32	9.78	-1.50	0.22	6.01	Syrian Arab Republic	−5.55	0.31	174	1.75	0.22	2.32
Greece	−0.68	0.25	13.1	0.75	0.18	1.66	Togo	−4.12	0.37	111	-0.49	0.23	7.49
Guatemala	−2.28	0.29	51.2	-0.03	0.20	3.83	Thailand	2.61	0.26	−37.7	0.15	0.18	1.70
Honduras	−2.96	0.32	71.2	-0.46	0.20	4.07	Tunisia	−2.26	0.29	50.9	1.29	0.21	1.23
India	1.86	0.25	−28.6	1.24	0.18	2.11	Turkey	−0.13	0.25	2.41	1.23	0.18	1.75
Ireland	2.54	0.24	−36.9	-0.33	0.18	0.90	Uganda	−3.35	0.36	83.6	-0.27	0.24	6.29
Iran	−2.35	0.31	53.4	1.20	0.23	2.98	Uruguay	0.14	0.28	−2.51	-0.31	0.20	2.04
Israel	1.78	0.27	−27.6	-0.01	0.20	1.30	Venezuela	−0.19	0.29	3.43	-0.16	0.20	3.38
Italy	3.48	0.24	−46.9	0.85	0.17	0.76	South Africa	2.24	0.24	−33.4	0.16	0.17	1.97
Jamaica	−2.04	0.30	44.8	-0.50	0.19	2.77	Zaire (DRC)	−0.68	0.32	13.2	-0.57	0.22	4.59
Jordan	−2.22	0.32	49.5	-0.01	0.22	2.81	Zambia	−0.79	0.34	15.3	-1.05	0.24	4.23
Japan	4.35	0.24	−54.6	1.44	0.17	0.59	Zimbabwe	−1.73	0.29	36.8	0.14	0.20	3.64
Kenya	−0.82	0.27	16.0	-0.58	0.19	7.34							

The parameters were estimated by OLS. For an estimated parameter $\hat{b}$, the implied percentage effect on cost is $100 \times (e^{-\theta\hat{b}} - 1)$. Technology parameters, λ_i, are recovered as detailed in Section V and $\theta = 0.18$.

Appendix

A. Alternative Trade Models and Asymmetric Trade Costs

1. Anderson and van Wincoop (2003) Model

The model of Anderson and van Wincoop (2003) generates equation (4). To do so, assume that each country has constant returns technologies with competitive firms producing a good which is defined by its country of origin, i.e., the Armington assumption. These assumptions imply the unit cost (and price) to deliver a country j good to country i is $p_{ij} = \tau_{ij} T_j^{\frac{1}{1-\sigma}} c_j$. Here c_j is the cost of inputs to produce one unit of the country j good and $T_j^{\frac{1}{\sigma-1}}$ is total factor productivity in country j.

Preferences are equally simple. Each country has symmetric constant elasticity preferences over all the (country-specific) goods with common elasticity of substitution σ. The key result from this simple model are the expenditure shares

$$X_{ij} = \frac{T_j(\tau_{ij}c_j)^{1-\sigma}}{\sum_{\ell=1}^{N} T_j(\tau_{ij}c_j)^{1-\sigma}}. \tag{16}$$

The right-hand side is country i's imports from country j divided by country i's expenditure on all traded goods. The left-hand side relates the trade cost country i faces to import a good from country j and country j's unit cost of production relative to the sum of the prices paid for imported goods.[28]

Given preferences, each country faces the following price of tradable goods for each country i:

$$P_i \quad = \quad \Upsilon \left[\sum_{\ell=1}^{N} T_\ell \left(c_\ell \tau_{i\ell}\right)^{1-\sigma} \right]^{\frac{1}{\sigma-1}}. \tag{17}$$

Dividing equation (16) with the analogous equation for country j's expenditure on country j goods and noting the relationship between the denominator of equation (16) and the price index in equation (16) results in the following relationship:

$$\frac{X_{ij}}{X_{jj}} = \tau_{ij}^{1-\sigma} \times \left(\frac{P_j}{P_i}\right)^{1-\sigma}. \tag{18}$$

This is the same expression as in (4) relating the bilateral trade shares to trade costs and the relative aggregate price of tradables.

[28] Anderson and van Wincoop (2003) call the term P_i inward multilateral resistance because it is a summary measure of the difficulty for country i to import.

2. Melitz (2003)/ Chaney (2008) Model

As shown in Chaney (2008) and Helpman, Melitz and Rubinstein (2008), the Melitz (2003) framework can easily generate a "gravity-like" expression. As I show here, these frameworks generate a relationship similar to equation (4), which would yield the same conclusions as Section III, even though these frameworks have fixed costs and firms have market power. Below I generate the results from the model of Chaney (2008).[29]

The key components of this model are as follows: Each country has CES preferences over a measure of differentiated goods (determined in equilibrium) with a common elasticity of substitution σ. To produce a good for country i in country j, firms face a variable cost per quantity shipped $\tau_{ij}\frac{c_j}{z}$ and a fixed cost f_{ij} in units of the numeraire. The productivity levels z are specific to the firm and modeled as an idiosyncratic random variable drawn from a Pareto distribution with shape parameter γ,

$$G(z) = 1 - T_i z^{-\gamma}.$$

As in Helpman et al. (2008) and Chaney (2008), I assume that $\gamma > \sigma - 1$. Finally, firms are monopolistic competitors, and there is an unbounded measure of potential entrants.

This formulation generates the following expression for the share of country i expenditures on imports from country j:

$$X_{ij} = \frac{T_j \left(c_j \tau_{ij}\right)^{-\gamma} \left(\frac{f_{i,j}}{Y_i}\right)^{-\gamma\left(\frac{1}{\sigma-1}-\frac{1}{\gamma}\right)}}{\sum_{\ell=1}^{N} T_\ell \left(c_\ell \tau_{i\ell}\right)^{-\gamma} \left(\frac{f_{i,\ell}}{Y_i}\right)^{-\gamma\left(\frac{1}{\sigma-1}-\frac{1}{\gamma}\right)}},$$

where Y_i is total expenditures in country i. Similar to the approach above, note that the aggregate price of tradables in country i is

$$P_i = \kappa \left[\sum_{\ell=1}^{N} T_\ell \left(c_\ell \tau_{i\ell}\right)^{-\gamma} \left(\frac{f_{i,\ell}}{Y_i}\right)^{-\gamma\left(\frac{1}{\sigma-1}-\frac{1}{\gamma}\right)}\right]^{-\frac{1}{\gamma}}. \tag{19}$$

Using this fact and dividing the analogous equation for country j's share of expenditures on country j

[29] Helpman et al. (2008) generate nearly the same expression except for the complications associated with the truncated Pareto distribution for productivity.

goods yields the following expression:

$$\frac{X_{ij}}{X_{jj}} = \hat{\tau}_{ij}^{-\gamma} \times \left(\frac{P_j}{P_i}\right)^{-\gamma}, \tag{20}$$

$$\text{where } \hat{\tau}_{ij} = \tau_{ij} \times \left(\frac{f_{ij}}{f_{jj}}\right)^{\left(\frac{1}{\sigma-1}-\frac{1}{\gamma}\right)} \times \left(\frac{Y_j}{Y_i}\right)^{\left(\frac{1}{\sigma-1}-\frac{1}{\gamma}\right)},$$

which is the similar to equation (4), but now what was one represented by variable trade costs is now a mixture of the variable trade costs, the relative fixed costs between the two markets, and market size.

Note that in this model, countries consume different varieties of goods. Hence, the price index in equation (19) is not comparable across countries. The data that I use in this paper are a common basket price index, which is the same across countries. Thus, an adjustment needs to be made because the objects in the model and the data do not correspond.

I will make an adjustment in the model by asking, if country n purchases the same basket of goods as the United States, than what would its price index be? I will use the basket in the United States as the reference basket, but using any country will generate the same results. The key is that we want to compute a *common* basket price index for all countries.

To implement this concept, I will assume that the fixed costs are multiplicatively decomposable, i.e., $f_{ij} = f_i \times f_j$ and $f_{ii} = f_i \times f_i$. With this assumption, one can show how country-specific price indices relate to country-common basket price indices:

$$P_n = (P_n^{us})^{\frac{\sigma-1}{\gamma}} \times \left(\frac{f_n Y_{us}}{f_{us} Y_n}\right)^{\left(\frac{1}{\sigma-1}-\frac{1}{\gamma}\right)}. \tag{21}$$

Then using equation (21) and substituting it into equation (21) yields the following expression:

$$\frac{X_{ij}}{X_{jj}} = \tau_{ij}^{-\gamma} \times \left(\frac{P_j^{us}}{P_i^{us}}\right)^{1-\sigma}, \tag{22}$$

which relates bilateral trade shares to relative aggregate prices of a common basket, variable trade costs, and bilateral trade shares. The only difference is that the power term on the variable trade costs is different from the power term on relative prices. Despite this difference, as long as $\sigma > 1$ and $\gamma > \sigma - 1$ (which the model requires), the same conclusions from Section III apply: to account for both bilateral trade shares and aggregate prices, trade costs must be systematically asymmetric with respect to development. Thus, my arguments about the structure of trade costs throughout this paper may be interpreted more generally.

B. Measuring Income per Worker in the Model as in the Data

Real GDP or income per worker is an important object of interest. This section describes my concept of real GDP in the model as measured in the data— specifically, real GDP in benchmark years of the Penn World Table (PWT). Feenstra, Heston, Timmer and Deng (2004) provide a very useful description of the mechanics behind the construction of real GDP in the PWT, and my presentation borrows from their analysis. First, the PWT collects GDP in current prices from each country's statistics agency. In the model, nominal GDP (in per worker terms) is

$$p_i^y c_i + ex_i - im_i, \tag{23}$$

where p_i^y is the price of the final good produced in country i, c_i is the quantity of the final good, and ex_i and im_i are nominal exports and imports, i.e., nominal final expenditures plus net trade balance at current prices. The PWT then uses the Geary-Khamis system to compute "reference prices" π_ℓ for each good ℓ and the purchasing power parities PPP_i for each country i used to deflate (23). Since there is only one final good, the Geary-Khamis system is the following set of equations:

$$\pi = \left(\sum_{j=1}^{N} \frac{p_j^y c_j}{PPP_j^e} \right) \left(\sum_{j}^{N} c_j \right)^{-1} \text{ and } PPP_j = \left(\frac{p_j^y c_j}{\pi c_j} \right).$$

For a unique solution, some normalization is necessary. Setting $\pi = 1$, it is straightforward to show that PPP_i for country i equals p_i^y, or the price of the final good. This results in the flowing definition of real GDP per worker:

$$c_i + \frac{ex_i - im_i}{p_i^y}.$$

With balanced trade, the last term equals zero. Notice that the net trade balance is deflated by the price index for domestic absorbtion. This is actual contrary to the advice of the United Nations 1993 System of National Accounts. They advise that the net trade balance be deflated with both export and import price indices. Hence, the PWT measure of real GDP is more closely related to gross domestic income or command-basis GDP as defined in the United Nations 1993 System of National Accounts. Imposing market clearing and balanced trade, real GDP per worker in my model—as measured in the PWT—is,

$$y_i = \frac{w_i}{p_i^y} + \frac{r_i k_i}{p_i^y},$$

in which income from wages and capital are deflated by each country's final goods price.

C. Derivation of a Useful Representation of Income per Worker

Suppressing some notation and rearranging (2), using (1), and the fact that the rental rate on capital is pinned down by the expression $r_i = \frac{\alpha}{1-\alpha} w_i k_i^{-1}$ provides an expression for each country's home trade share:

$$X_{ii} = \frac{\left[k_i^{-\alpha\beta} w_i^{\beta} p_i^{q(1-\beta)}\right]^{\frac{-1}{\theta}} \lambda_i}{p_i^{q\left(\frac{-1}{\theta}\right)} \Psi}. \tag{24}$$

Further rearrangement of (24) provides the expression

$$\left(\frac{w_i}{p_i^q}\right) = \Psi \left(\frac{\lambda_i}{X_{ii}}\right)^{\frac{\theta}{\beta}} k_i^{\alpha}, \tag{25}$$

in which wages, deflated by the intermediate goods price, are a function of each country's home trade share and its capital-labor ratio.

Given the definition of real income per worker defined above and using a representative firm's first-order conditions determining the rental rate as a function of the wage, I express income per worker as a function of only the wage and the final goods price:

$$y_i = \frac{1}{1-\alpha} \frac{w_i}{p_i^y}. \tag{26}$$

Since my interest is only in relative income differences, constant terms are abstracted from. Combining the expression for the price of final goods and (26), real income per worker is expressed as

$$y_i = \left(\frac{w_i}{p_i^q}\right)^{1-\gamma} k_i^{\alpha\gamma}. \tag{27}$$

Combining equations (25) and (27), real income per worker is now

$$y_i = X_{ii}^{\frac{-\theta(1-\gamma)}{\beta}} \lambda_i^{\frac{\theta(1-\gamma)}{\beta}} k_i^{\alpha}. \tag{28}$$

Here real income per worker is only a function of each country's home trade share X_{ii}, its technology parameter λ_i, and its capital-labor ratio.

D. Zeros

An implication of the Eaton and Kortum (2002) framework is that, in aggregate, every country should purchase some nonzero amount of goods from all other countries. In fact, the bilateral trade matrix has many recorded zeros. For the sample considered there are 5,929 possible trading combinations; 1,610 (27 percent) show no trade at all. This presents both an estimation issue and a computational issue.

Regarding the estimation, I will omit any zero observed trade flows from the estimation of equation (8). This has been a standard approach in the empirical trade literature. Santos-Silva and Tenreyro (2006) propose a Poisson pseudo-maximum-likelihood (PPML) estimator to alleviate any bias from the log-linear specification in equation (8) due to the presence of heteroskedasticity and the omission of zero observed trade flows. I employed their technique of estimating the gravity equation in levels and including zero observed trade flows, and found that the quantitative results and counterfactual exercises do not differ dramatically relative to using ordinary least squares. This does not contradict their findings. Consistent with their results, I find that OLS exaggerates the distance elasticity, suggesting that the bias they emphasize is present. For example using PPML, the percentage effect on cost is 129, 140, 141, 177, 223, and 263 percent for each distance category. Compared to Table 2, shorter distances are more costly and longer distances are less costly relative to OLS. This is consistent with the lower distance elasticity Santos-Silva and Tenreyro (2006) find when using PPML relative to OLS.

Helpman et al. (2008) particularly focus on zero trade flows, building on the model of Melitz (2003) with fixed costs and firm heterogeneity. When firm-level productivity is drawn from a truncated Pareto distribution, they can deliver zero trade flows between country pairs. Their results suggest that any bias arising from the omission of zero trade flows is quantitatively small.

Regarding the computation, when computing equilibrium prices and counterfactuals, I will set trade costs for the instances in which X_{ij} is zero to an arbitrarily large value to approximate what appears to be a trade cost of infinity.

E. Technology Heterogeneity: Evidence

A concern is that the distributional assumptions generate implausible differences in productivity at the micro or good level. This is important because the degree of dispersion in productivity affects the response of aggregate TFP and income differences as trade costs change. In this section, I argue that the dispersion in productivity implied by my model is reasonable and conservative based on the available empirical evidence.

To make this argument, I first performed the following exercise. I assumed that there were 100,000 goods and generated a productivity term $z^{-\theta}$ for each good from the calibrated distribution for the United States, United Kingdom, and Uganda. I then asked two questions: (i) how much does $z_i^{-\theta}$ vary within a country? and (ii) For a given good, how much does $z_{us}^{-\theta}$ differ relative to $z_{uga}^{-\theta}$ or $z_{uk}^{-\theta}$? After recording various measures of dispersion, I repeated this process 500 times. Table 4 reports the means of these measures across the simulations.

Table 4: Dispersion in Productivity

Model: Variation Within Countries, Across Goods			
	U.S.	U.K.	Uganda
Estimated $\lambda_{us}^{\theta}/\lambda_i^{\theta}$	—	1.10	6.30
Simulated mean relative to U.S.	—	1.12	6.30
$99-\;1$ ratio of $z_i^{-\theta}$	3.04	3.04	3.04

Model: Variation Within Goods, Across Countries		
	U.S. / U.K.	U.S. / Uganda
99th percentile of $(z_{us}/z_j)^{-\theta}$	2.58	14.5
1st percentile of $(z_{us}/z_j)^{-\theta}$	0.48	2.73

1. Variation Within Countries

The top panel of Table 4 reports the degree of variation in productivity within a country. For all countries, the ratio of productivity in the top 99th percentile over the bottom 1th percentile is about a factor of 3. This value is only a function of θ, with increases in θ increasing the difference between percentiles.

Relative to the available evidence, Table 4 shows my model implies a degree of dispersion in productivity that is conservative. For example, Hsieh and Klenow (2007a) employ plant-level data from China, India, and the United States and construct estimates of TFP at the plant level. My point of comparison is what they call "TFPQ", which is most closely related to how I would measure TFP given my model. They report a dramatic amount of dispersion in TFP. For example, the ratio of the 90th percentile to the 10th percentile in China in 1998 is 15.18—10 times the amount of dispersion in my model. In India in 1998, this same ratio is 31.18—20 times the amount of dispersion in my model. Measuring TFP at this level is difficult and comes with many caveats; however, this evidence suggests that amount the of dispersion in productivity implied by my model is very conservative.

2. Variation Across Countries

The bottom panel of Table 4 reports the variation across countries in productivity to produce a particular good. For example, compare the productivity to produce blue tennis shoes in the United States, United Kingdom, and Uganda in the model. The top row presents the 99th percentile of the distribution of relative productivities between countries to produce the same good. For the 99th percentile, the United States is 2.58 times more productive than the United Kingdom and 14.5 times more productive than Uganda to produce the same good. The bottom row present the 1st percentile. Here, the United Kingdom is 2 times more productive than the United States and Uganda is 2.7 times *less* productive than the U.S. to produce the same good.

Again, this degree of dispersion in TFP across countries is conservative relative to available empirical evidence on the variation in productivity within industries across developed countries. Baily and Gersbach (1995) show that value added per worker within the same manufacturing industries varies by as much as a factor of 3 between the United States, Japan, and Germany in the early 1990s.[30] And they argue that differences in the use of physical capital played little role in explaining these differences. Relative to the bottom panel of Table 4, this suggests that between rich countries I am slightly understating the dispersion in productivity differences.

Because even the best producer in the poor country is still nearly 3 times *less* productive than the United States, I am again understating the dispersion in productivity differences relative to the evidence. There is less micro-evidence regarding poor countries, but a theme that emerges is that some firms in poor countries are as productive as firms in rich countries. Banerjee and Duflo (2005) discus evidence from India on TFP at the industry level from the McKinsey Global Institute (2001). Banerjee and Duflo (2005) highlight the fact that the best firms in several of the manufacturing industries studied are basically using the global best practice technologies. Lewis (2004) argues that this same pattern prevails in other developing countries such as Brazil, Russia, and Korea. Though the focus of his paper is not in manufacturing, Lagakos (2008) provides evidence for the retail sector that the most productive firms in developing countries are nearly as productive as the firms in developed countries.

F. Alternative Distributional Assumptions

In this section, I consider alternative distributional assumptions and their quantitative implications.

Without making the distributional assumptions, the model loses analytical expressions mapping the data to parameters of the model, i.e., trade costs and technology parameters. To solve the model under alternative distributional assumptions, I employed the following approach. First, I assumed there was a large number (100,000) of potentially tradable goods. For each country, good-level efficiencies were drawn from the country-specific distribution and assigned to the production technology for each good. Then, for each importing country, the low-cost supplier across countries is found for each good and the aggregate bilateral pattern of trade is computed. In the examples below, I assume there are only 10 countries and adjusted the data under the assumption that these are the only countries in the economy. I did not consider the full 77-country example because of the computational burden associated with estimating/calibrating the model.

I considered two alternative distributional assumptions: (i) a log-normal distribution with country-specific center parameter μ_i and common parameter σ and (ii) Pareto distribution with country-specific center

[30] See Parente and Prescott (2002) for a nice discussion of this evidence.

term κ_i and common shape parameter ν. Because the common parameters σ and ν play roles similar to that of θ, I calibrated them for these examples such that all models have the same coefficient of variation. I calibrated the country-specific parameters (μ_i or κ_i) and ex_i for each country and common parameters relating to the effects of distance and shared borders to best fit the data. For comparison purposes, I used the same approach in the benchmark case with a Fréchet distribution rather than exploiting its analytical convenience.

Table 1: Alternative Distributional Assumptions

Model Fit			
Distribution	M.S.E.	$\text{var}(\log(y))$	y_{90}/y_{10}
Data	—	1.98	25.11
Fréchet (benchmark)	1.99	1.43	17.04
Log-normal	2.12	1.47	17.96
Pareto	2.25	1.53	19.25
Gain from Trade, $\tau = 1$			
	Mean $\Delta\, y$ (%)	$\text{var}(\log(y))$	y_{90}/y_{10}
Fréchet (benchmark)	68	0.82	8.10
Log-normal	43	1.12	11.58
Pareto	69	0.86	8.82

The top panel of Table 1 presents some summary measures of the fit. The first column is the mean squared error between the data and model in logs. The second and third columns report the implied difference in income per worker across the two scenarios. In terms of the fit of the data and measures of dispersion in income per worker, all three models perform similarly.[31]

The bottom panel of Table 1 presents the average increase in income per worker after removing all trade costs, the variance in log income per worker, and the 90/10 ratio in income per worker. Between the Fréchet and Pareto case there is little difference in the change across all three measures. Both generate reductions in cross-country income differences of approximately 44 percent. The log-normal distribution results in less reduction in cross-country income differences—approximately a 24 percent decline in the variance of log income per worker. As the log-normal case illustrates, the distributional assumptions obviously play a role. More research is needed on the implications of alternative distributions and evidence supporting these assumptions. But in all these cases, poor countries gained the most from reductions in trade costs, thus reducing cross-country income differences.

[31]For reference, had I used the analytical convenience when calibrating the economy with the Fréchet distribution, the mean square error would have been 1.31, the variance in log income per worker 1.62 and the 90/10 ratio 25.76.