Innovation and Product-Level Trade*

Scott French†

November, 2009

Abstract

The exports of poor countries are highly concentrated in the common set of products for which they have a strong comparative advantage in the product-level trade data. Rich countries, on the other hand, spread their exports very evenly across the range of products, including those in which they have a comparative disadvantage vis-a-vis poorer countries. Theories in which comparative advantage is driven by a symmetric relationship (like relative factor endowments) or idiosyncratic productivity or product differences (as in theoretical gravity models) cannot account for both these facts.

To reconcile these facts, I develop a Ricardian model of trade in which average productivity is allowed to vary across products within a country as well as across countries. These product level productivity levels are determined by the equilibrium of an endogenous growth model in which (a) researchers have the choice between inventing a new good and developing a new way to produce an existing good and (b) developing a process for a good invented abroad requires researchers to first devote effort to learning about that good. I show that this model is able to quantitatively account for these facts and that, compared to a restricted version of the model that delivers an aggregate gravity equation, the full model better predicts aggregate trade flows.

The innovation and learning environment also provides a convenient way to nest models of trade based on increasing returns and those based on productivity differences, with elements of both present in equilibrium.

*I am thankful to Dean Corbae for his guidance and support. I am also thankful for advice and comments from Kripa Freitas, Kim Ruhl, Natalia Ramondo, Ken Hendricks, and Jason Abrevaya as well as discussions with Jonathan Eaton and seminar participants at the University of Texas and various conferences. This paper was previously circulated under the title “Innovation in Product Space and Trade”

†Department of Economics, University of Texas at Austin. stf223eco.utexas.edu

1 Introduction

It has been documented that rich countries tend to export a different set of products than poor countries.[1] In this paper, I extend this finding, showing that while poor countries' exports are highly concentrated in a relatively small, common set of products, rich countries' exports are nearly uniformly distributed across all products. Standard trade models cannot simultaneously account for both of these observations.

This paper provides a theoretical interpretation of these facts. Based on the structure of Eaton and Kortum (2001), patterns of bilateral trade are determined by innovative effort in each country over time. I show that if researchers have the choice between inventing a new product and developing a process for producing an existing one (as in Young (1998)) and if researchers in one country can learn to produce products invented abroad (as in Krugman (1979b)), then this feature of the data is present along a balanced growth path of the economy with parameter values chosen to match the distribution of manufacturing wages across countries. In addition, the model yields a tractable analytical representation of the distribution of technology across products within each country. And, to my knowledge, it is the first model of endogenous growth or product cycles in trade that can be seriously taken to the multiple country, bilateral, product level trade data that has become increasingly available in recent years.

Using data on bilateral trade in 4,157 Harmonized System manufacturing product codes, I show that, on average, 31% of the exports of a country with per capita GDP less than $10,000 fall in a set of products making up 5% of world trade. However, on average, no more than 6.5% of a rich country's exports fall into any 5% group of products. Standard models of international trade which deliver a gravity equation have been quite successful in accounting for aggregate bilateral trade flows, at least among rich countries.[2] However, they are silent on product level trade facts due to the common simplifying assumption that there is no correlation between characteristics of a country and the likelihood that it exports a particular good. Models of comparative advantage based on factor endowments can account for the fact that a definite pattern of product-level comparative advantage exists but predict that rich countries' exports would also be concentrated in the products which use their relatively abundant factors intensively.

[1] e.g. Schott (2002)

[2] e.g. Eaton and Kortum (2002), Anderson and van Wincoop(2003), and Helpman, Melitz & Rubenstein (2008)

In order to reconcile these facts, I propose a Ricardian trade model in which the average level of productivity in a country is the result of innovative effort, based on the framework of Eaton and Kortum (2001). My model differs from this one in three important ways. First, the average productivity level in a country is allowed to vary across goods.[3] Second, researchers have the option of performing two kinds of research: inventing a new product or developing a new way to produce an existing good. In addition, if a product was invented abroad, a researcher must spend time learning about the good before obtaining a production idea for it. This final assumption provides an intuitive and analytically tractable way to nest two environments that have been very important to both the endogenous growth and trade literatures – product variety and technological differences – and to allow interesting interaction between the two.

As is the case in Eaton and Kortum (2002) with aggregate bilateral trade volumes, expected product level trade flows in this model depend crucially on an index of technology levels in each country discounted by input costs and geographic barriers. The probability that an idea for producing a given product will be usable is decreasing proportionally to this index. However, expected expenditure on the product is increasing. The former effect is dominant, and the expected profit flow from an idea is higher if the idea applies to a good to which little research effort has been devoted or one for which the research devoted to it is concentrated in high wage, distant countries.

In equilibrium, countries that are very productive at research also have high wages. As a result, researchers in rich countries spend more time developing new products to avoid direct competition with low wage producers. However, since producers in poor countries do not have to be the most productive in order to produce a good at a lower cost (because of lower wages), they spend more time learning about goods invented abroad (for which most research has been done in countries with relatively higher wages) and developing processes to produce goods they have already learned about.

This pattern of research implies that, soon after a product's invention, ideas for producing it accumulate relatively slowly in poor countries, but the rate of accumulation increases over time as researchers there learn about the product. As a result, in the cross-section of

[3] Costinot & Komunjer (2008) performs a similar exercise to derive an industry level gravity equation and shows that a country exports relatively more in the industries in which its labor productivity is higher. Chor (2009) uses similar methodology, showing that trade patterns are related to countries' factor endowments and quality of institutions. Neither, however, considers how different average productivity levels across products and countries interact in equilibrium to determine trade patterns or their implications for innovation incentives, the mains focal points of this paper.

products at a point in time, rich countries have a large comparative (and absolute) advantage in producing newer goods. For older goods, on the other hand, since researchers in poorer countries have had time to learn about and begin developing ways to produce them, the technological advantage of rich countries is much smaller. However, because researchers in rich countries continue to accumulate ideas for producing existing goods, rich countries' technological advantage deteriorates slowly, ensuring that rich countries continue to export a significant proportion of older goods.

This paper contributes to several strands of the literature. First it documents a fact from bilateral, product level trade data that, to my knowledge, has never been explicitly studied. As such, it complements papers such as Hummels & Klenow (2004), Schott (2002 & 2004), and Baldwin & Harrigan (2007), which test the implications of trade theory using product level data. To document this fact, I use bilateral, product level trade data which covers trade among more than 200 countries in more than 5000 six digit Harmonized System product codes. Hummels & Klenow (2004) use a similarly comprehensive data set to evaluate the degree to which the size of the set of products and the quantity and price of products exported varies with exporter labor force size and GDP per worker. Several papers use less comprehensive product level trade databases to study the effect of importer or exporter characteristics on product unit prices.[4] Costinot & Komunjer (2008) and Chor (2009) use industry level trade data to evaluate the effect of productivity and factor endowment differences (respectively) on trade patterns. However, to my knowledge, this is is the first paper to study the link between income levels and patterns of product level trade using such a data set.

Second, it builds on the work of Eaton & Kortum (2001) which uses the endogenous growth model of Kortum (1997) to underpin a tractable and empirically relevant model of trade. It also draws on the approach of Young (1998), which was the first to combine the two major mechanisms of endogenous growth - the expanding product variety framework of Romer (1990) and quality ladders framework Grossman and Helpman (1991) - although for a very different purpose.

This paper also provides an analytically tractable and empirically relevant version of a model of product cycles in trade, an idea originating with Vernon (1966). In fact, the process of learning about goods invented abroad is an endogenous generalization of the main mechanism by which production of goods diffuses from rich to poor countries in Krugman

[4]e.g. Hummels & Skiba (2003); Choi, Hummels, & Xiang (2009); Schott (2004); and Fieler (2007)

(1979b). Like Eaton & Kortum (2006) and Grossman and Helpman (1991), it also allows for production to move back and forth between countries as rich countries continue to devote research to the product. However, unlike all these, this model extends the analysis to multiple countries and can be seriously taken to product level trade data. And, because of its relevance to the data, this paper shows how product level trade data can be used to inform models of endogenous innovation and diffusion of technology, such as Comin and Hobijn (forthcoming).

The model of this paper also nests models of trade based on increasing returns and endogenous product variety (e.g. Melitz (2003)) and those based on technology differences across a fixed set of products (e.g. Eaton & Kortum (2002)). The key element in moving between these extremes is the process of learning about products invented by others. If the learning process were shut down completely (even within a country), then the inventor of the good would be the only potential producer of the good. Similar to Grossman & Helpman (2001), the producer has no incentive to improve the process for producing his good, so process innovation ceases, and the model reduces to one of monopolistic competition, where innovation leads to expanded product variety. At the other extreme, if learning is free and instantaneous everywhere, then there is no incentive to invent new products, and the set of products will remain forever fixed at some initial level while innovation leads to increased productivity for these goods. In both of these cases, the probability that a given country exports a given product is independent of the identity of the good, which serves to highlight the necessity of a nondegenerate learning process for the theory to account for the product level trade facts discussed above.

That the prices of (and hence demand for) goods can systematically vary with their characteristics introduces an interesting interaction between the degree of dispersion in idiosyncratic productivity shocks and the CES elasticity of substitution in demand. Specifically, for newer goods, for which technology is concentrated in the inventing country, the elasticity of trade volume with respect to trade costs is closer to the elasticity of substitution of demand; for older goods, it is closer to the dispersion in idiosyncratic productivity. By contrast, gravity models typically predict that trade volumes depend of one of these two parameters, while the other is mostly irrelevant. The appendix also discusses how this interaction can reconcile the model of Fieler (2007) – which seeks to improve upon the ability of gravity models to predict the differing trade patterns across rich and poor countries observed in the data - with the evidence from Waugh (2009) contradicting the assumption that poor countries specialize in goods characterized by less dispersion in idiosyncratic productivity.

The next section presents the evidence for and discusses the implications of the facts mentioned above. Section 3 presents the static trade model with differences in average productivity across products within a country and then embeds it in the model of endogenous growth with research activity choice. Section 4 discusses the calibration of the model and explores the equilibrium implications of the calibrated model. Section 5 evaluates the model's ability to capture features of the trade data, both the product level facts presented here as well as aggregate trade flows, and it compares the latter predictions with a restricted version of the model in which all innovation is developing new processes, which reduces to a version of Eaton & Kortum (2002) and delivers a gravity equation in aggregate trade flows. Section 6 discusses the effects of counterfactual exercises and section 7 concludes.

2 Data

Before turning to theory, I present the product-level trade facts which a theory of product-level productivity determination should address.

The data used are from the BACI database of the CEPII described in Gaulier, et al, (2008). The major benefit of this dataset is the level of coverage and product-level detail. The data cover imports and exports of more than 200 countries in more than 5000 Harmonized System product codes. Data on PPP output per worker are from the World Development Indicators (WDI) database of the World Bank. Of the countries in the BACI database, 159 of them have GDP and per worker data in the WDI database for 1995, the year used for the calculations that follow. My sample, then is these 159 countries and trade in the 4157 product codes that correspond to ISIC manufacturing industries. Restricting the sample to manufacturing industries is done so that natural resource endowments (which may drive comparative advantage in the agriculture and mining/extraction industries) are not a large factor in the analysis that follows.

2.1 Exports and Income

The first fact suggests that the simplifying assumption of most empirical trade models – that all products are the same, save idiosyncratic productivity or preference shocks – is not borne out in the data and that development of a model of ex-ante (as opposed to idiosyncratic) comparative advantage is warranted.

1. The ranking of products by the intensity with which a country exports the product relative to other countries is positively correlated across countries within income groups but negatively correlated across income groups.

In other words, income per capita is an important determinant of which products a country exports. This finding is in line with Schott (2002), which finds that US intra-product trade – measured by the Grubel-Lloyd index – is much higher with high wage countries than with low wage countries.

To obtain this ranking of goods for each country, I employ the index of "revealed comparative advantage" from Balassa(1965):

$$RCA_i^j = \frac{X_i^j / X_i}{X^j / X}$$

where X_i^j is i's total volume exports of product j, X_i the total volume of exports of i, X^j total volume of world exports of j, and X the total volume of world trade. It should be noted that, despite its name, this ad-hoc index was selected not because it is thought to represent any theoretical formulation of comparative advantage, per se[5], but because it is a simple and intuitive way to compute the degree to which the exports of a country are concentrated in a particular product category, relative to other countries, while controlling for the considerable amount of dispersion in the sizes of product categories in the data.[6]

Having computed the Balassa index (RCA_i^j) for each exporting country and product code, I then used this index to rank the products for each country according to the relative intensity with which it exports each product. With this ranking in hand, one can immediately note that countries in the same income class tend to intensively export the same products. Table 1 shows the Spearman rank correlations between some countries of different income classes. The section in bold show that the correlations within the richest and poorest groups are quite high, while the correlations between members of these

[5]Costinot and Komunjer (2008) does provide a theoretical foundation for a bilateral version of this index using a generalization of Eaton and Kortum (2002). The use of the index in its present form is for expositional simplicity given that no particular trade model has yet been assumed.

[6]See Broda and Weinstein (2006) and Armenter and Koren (2008) for discussions about the perils of not properly accounting for this dispersion.

Table 1: Rank Correlation of RCA

	USA	Germany	Japan	S. Korea	Brazil	India	China	Nigeria
USA	1							
Germany	**0.1949**	1						
Japan	**0.2698**	**0.2825**	1					
S. Korea	-0.1443	-0.0482	0.2085	1				
Brazil	0.0681	0.0983	0.0583	0.0296	1			
India	**-0.2273**	**-0.1269**	**-0.1025**	0.2128	0.0903	1		
China	**-0.3366**	**-0.3112**	**-0.1615**	0.2801	-0.0884	**0.3785**	1	
Nigeria	**0.0173**	**-0.0787**	**-0.0543**	0.0887	0.0964	**0.1154**	**0.0899**	1

groups is quite low.[7] The table also shows that the rankings of products exported by the middle-income countries, South Korea and Brazil, also lie in the middle of those of the high- and low-income countries. Their correlation coefficients are slightly positive with respect to countries in both sets.

The evidence, then, suggests that there is meaningful *overall* ranking of products by the relative intensity with which they are exported by rich or poor countries. In order to obtain such a ranking, I then computed the average RCA for each j over two groups – the 14 countries in the sample for which 1995 PPP GDP per worker is greater than \$45,000 (group d)[8] and the 82 countries for which GDP per worker is less than \$10,000 (group l.[9] I then computed the relative RCA for each product code,

$$RRCA^j = \frac{RCA_d^j}{RCA_l^j}$$

where RCA_d^j is the average revealed comparative advantage for product j over the set of developed countries in the richest group, and RCA_l^j that over the set of countries in the least developed group.

[7] Also worth noting is that Nigeria does not follow this pattern as strongly. A casual look at the data seems to show that Nigeria is not an outlier, as the RCA ranking of African countries' exports is relatively uncorrelated with that of other countries, including other poor African countries. Perhaps this is a function of the fact that these countries export a relatively small amount of manufactures overall, and very few of these countries report trade data to the UN

[8] Excluding OPEC countries, whose exports tend to be focused in petroleum refining and byproducts industries.

[9] The use of a larger sample of poorer countries is to preclude both the economic size of the former group dwarfing the latter and the index being heavily influenced by a few small countries. Altering the size of each sample in various ways had little impact on the results that follow.

Figure 1: Percentage of Exports by Product Group

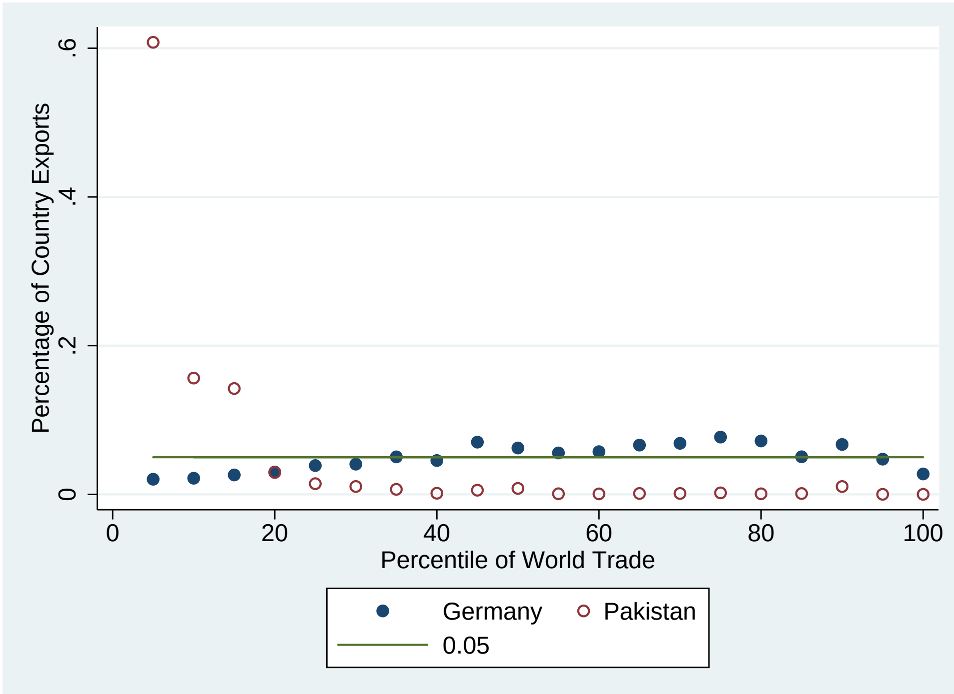

2.2 The Concentration of Exports

Now, having an overall ranking of the goods for which rich and poor countries, respectively, tend to have an ex-ante comparative advantage, I turn to the primary facts concerning the degree to which countries' exports are concentrated in their comparative advantage products.

2. The exports of poorer countries are concentrated in a few, common products.

3. Rich countries' exports, by contrast, are spread rather evenly across all products, including those intensively exported by poor countries.

That is, as a percentage of the country's total exports, a given low wage country tends to export much more in the common set of products in which the set of poor countries has a comparative advantage. On the other hand, rich countries export nearly as much in products in which they have a comparative disadvantage as they do in their comparative advantage products.

Figure 2: Percentage of Exports by Product Group

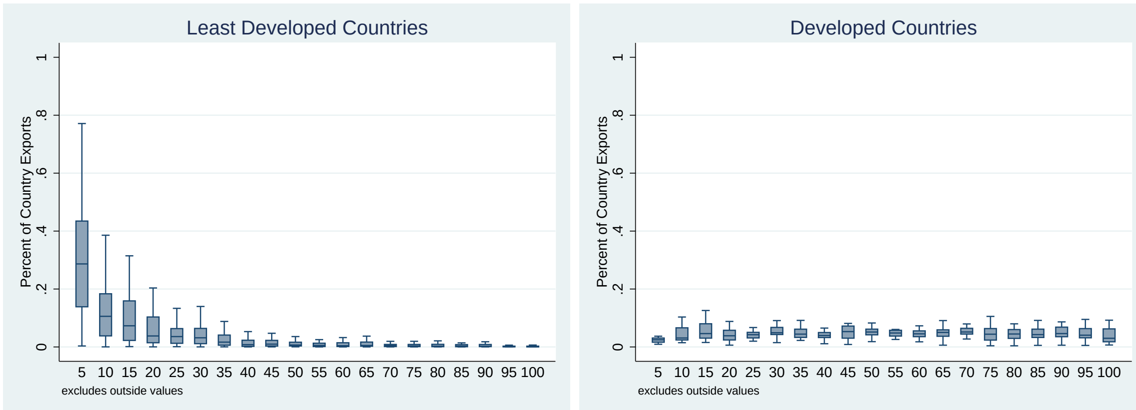

For expositional purposes, in what follows products are grouped into 5-percentile bins.[10] That is, each point on the x-axis represents a group of products that make up 5 % of world trade, and moving to the right, the group contains products that are exported relatively more intensively by richer countries. Given this set-up, if a country's exports are uniformly distributed across the set of products, its percentage of trade in each group would be exactly 5%. The degree to which more that 5% of a country's exports fall in one group of products is the degree to which it's exports are relatively more concentrated in those products than the average country. Figure 1 plots the percentage of exports that fall into each group for Germany and Pakistan. It becomes immediately apparent that, while Germany's exports are relatively evenly spread across the range of products, the vast majority of Pakistan's exports fall in just a couple groups, with the first group of 5% of products by volume making up nearly 60% of Pakistan's exports.

Figure 2 shows that this phenomenon is not restricted to these two countries. Each box represents the percentage of countries' exports that fall into the corresponding group, with the line being the median across countries, the box the representing the 25th to 75th percentile countries, and the "whiskers" representing the maximum and minimum values.[11] In general, as with the particular case of Germany, for the set of rich countries, exports in each group make up roughly the same proportion of exports. For the set of poor countries, however, one can see that their exports are very concentrated in the common, small set of

[10]This is also to deal with the afore mentioned issues of dealing with product categories in the data that may contain many or few varieties of a product. In this setup, a product category is implicitly weighted by its importance in world trade, similar to how categories are weighted in Broda and Weinstein's (2006) appropriate price index.

[11]A very small number of of outliers are omitted.

Table 2: Exports by Product and Income Group

GDP/Worker	Bin 1	Bin 20
High	4.32%	4.74%
Middle	12.6%	1.29%
Low	31.1%	0.70%

products in which the group has the highest overall comparative advantage.

To illustrate this point more clearly, table 2 presents the average percentage of exports in the first and last 5% bin for each group of the income distribution. It can be seen, then, that the average low-income country has a large portion of its exports in a single bin, while the exports of the average high-income country are not concentrated in any bin.

3 Model

Consider a world with N countries and a continuum of goods. At a given point in time, the set of goods and the state of technology for producing them in each country – along with trade costs – will determine wages and trade flows. The set of goods and the state of technology will evolve over time with research effort in the countries. The following section presents the static trade model, and the next section will embed it into the dynamic model of innovation.

3.1 Static Trade Equilibrium

The world at time t consists of N countries and a continuum of goods of measure J_t. Each country $i = 1, 2, ..., N$ is made up of a continuum of identical consumers of measure L_{it} who each inelastically supply one unit of labor. For the remainder of this section, time subscripts will be suppressed.

3.1.1 Demand

Each consumer has preferences over the set of goods given by the CES utility function

$$U = \left(\int_0^J y_j^{\frac{\sigma-1}{\sigma}} dj \right)^{\frac{\sigma}{\sigma-1}},$$

where y_j is the quantity of good j consumed, and $\sigma > 1$ is the elasticity of substitution between any two goods.

Maximization of this utility function implies that total expenditure by consumers in country i on good j is

$$x_i^j = X_i \left(\frac{p_i^j}{P_i} \right)^{1-\sigma}, \tag{1}$$

where X_i is total expenditure by country i, p_i^j is the price of good j in country i, and

$$P_i = \left(\int_0^J (p_i^j)^{1-\sigma} dj \right)^{\frac{1}{1-\sigma}} \tag{2}$$

is the CES price index in country i.

3.1.2 Production and Technology

The fundamental unit of technology is an idea. An idea is pair consisting of the quality of the idea, Q, and the good j to which it applies. An idea of quality Q allows a producer to produce a unit of good j using $\frac{1}{Q}$ units of labor[12]; and, the quality of each idea is assumed to be drawn from the Pareto distribution, that is

$$Pr(Q < q) = 1 - q^{-\theta}.$$

At a point in time, in each country, there are a number of producers, $K_i^j = 0, 1, 2, ...$ who have an idea for producing good j. The number of potential producers of good j in country i is assumed to be the realization of the Poisson distribution with parameter T_i^j. So,

$$Pr(K_i^j = k) = \frac{e^{-T_i^j} (T_i^j)^k}{k!}.$$

One can think of T_i^j as the amount of research effort that has been devoted to drawing ideas for producing good j by researchers in country i. At a point in time, the set $\{T_i^j\}_{i,j}$ is taken as given. The next section will explore how this set evolves over time.

Since consumers are indifferent as to who supplies a particular good, I assume that competition is Betrand, meaning that only the producer who can supply a particular good at the lowest cost will do so. Moreover, that producer will charge a price just low enough so that no other producer can profitably undercut this price. As a result, only the producer with the highest quality idea in country i, $Z_i^j = \max_k \{Q_i^{jk}\}$, will ever produce in equilibrium.

[12] It is worth noting that Q can be interpreted in two ways which are isomorphic in this specification. One can think of an idea with a higher Q as either the ability to produce a given amount of good j with fewer units of labor or as the ability to produce a higher quality version of good j with a given amount of labor. In the second case, y_j should be interpreted at quality units of good j.

I show in the appendix that the technological frontier in country i is distributed as follows:

$$F_i^j(z) = Pr(Z_i^j < z) = e^{-T_i^j z^{-\theta}} \quad (3)$$

where $z \in [0, \infty)$.[13] This is the Frechet distribution familiar from Eaton and Kortum 2002, etc.

It should be noted that the static trade model is very similar that of Eaton and Kortum (2001, 2002, etc.). The notable difference is that the parameters governing the state of the technological frontier in each country are allowed to vary across the range of goods, as opposed to most trade models of heterogeneous costs, in which all goods are ex ante identical.

3.1.3 Trade

The cost of delivering a unit of good j from the best producer in country i to a consumer in country n is

$$C_{ni}^j = \frac{w_i d_{ni}}{Z_i^j},$$

where w_i is the wage in country i, and $d_{ni} > 1$ is an iceberg cost of delivering a good from country i to country n, i.e. delivering one unit of good j from i to n requires shipping d_{ni} units.

This implies that C_{ni}^j is distributed as follows:

$$G_{ni}^j(c) = Pr(C_{ni}^j < c) = 1 - e^{-T_i(w_i d_{ni})^{-\theta} c^{\theta}}.$$

Then, the lowest cost way of delivering a unit of good j to country n from *anywhere* in the world, $C_n^j = min_i C_{ni}^j$, is distributed

$$G_n^j(c) = Pr(C_n^j < c) = 1 - \prod_{i=1}^{N} G_{ni}^j(c) = 1 - e^{-\Phi_n^j c^{\theta}} \quad (4)$$

where $\Phi_n^j = \sum_{i=1}^N T_i^j (w_i d_{ni})^{-\theta}$, which can be thought of as the total amount of research effort in the world that has been devoted to ideas for producing good j, down-weighted by each country's labor and trade costs. In addition, the probability that the producer from which consumers in country n purchase good j is from country i, is simply

$$\pi_{ni}^j = \frac{T_i^j (w_i d_{ni})^{-\theta}}{\Phi_{ni}^j}, \quad (5)$$

[13] The formulation of the Pareto distribution above implicitly assumes that Q is bounded below by 1, implying that $F_i^j(z)$ should similarly be truncated at 1. Eaton & Kortum (2006) shows that this approximation becomes arbitrarily close to the correct formulation as T becomes large. The appendix provides a more precise formulation under which this issue disappears.

or i's contribution to the parameter governing the overall state of technology for good j.

3.1.4 Profits and Prices

Given that a producer possesses the best technology for delivering good j to country n its profit from selling in country n will be

$$\Pi_n^j = X_n \left(\frac{p_n^j}{P_n} \right)^{1-\sigma} (p_n^j - C_n^j) \tag{6}$$

As noted above, under the Bertrand assumption, the producer with the best way of delivering good j to country n will set its price as high as possible while keeping the nearest competitor at bay. However, the producer will not set the price above the monopoly price, so the price of good j in country n will be

$$p_n^j = min\{C_n^{(2)j}, \bar{m}C_n^j\},$$

where $C_n^{(2)j}$ is the second lowest cost of delivering j to n and $\bar{m} = \frac{\sigma}{\sigma-1}$ is the monopoly markup.

Given the cost distribution and pricing rule, I show in the appendix that the price index in country n is

$$P_n = \gamma \Phi_n^{\frac{-1}{\theta}}, \tag{7}$$

where

$$\gamma = \left[\Gamma \left(\frac{2\theta + 1 - \sigma}{\theta} \right) \left(1 + \frac{\sigma - 1}{\theta - (\sigma - 1)} \bar{m}^{-\theta} \right) \right]^{\frac{1}{1-\sigma}},$$

and $\Gamma(\cdot)$ is the Gamma function. And,

$$\Phi_n = \left(\int_0^J (\Phi_n^j)^{\frac{\sigma-1}{\theta}} dj \right)^{\frac{\theta}{\sigma-1}} \tag{8}$$

is the CES aggregate of the technology parameters for all goods, a sort of world technology index from the perspective of country n. This price index is well defined only for $\theta > \sigma - 1$,[14] so this parameter restriction will be assumed henceforth.

An important result of this assumption can be immediately seen. Consider the expected expenditure by consumers in country n on a good j with technology parameter Φ_n^j.

$$E[X_n^j] = E\left[X_n \left(\frac{p_n^j}{P_n} \right)^{1-\sigma} \right] = X_n \left(\frac{\Phi_n^j}{\Phi_n} \right)^{\frac{\sigma-1}{\theta}}, \tag{9}$$

[14] Since a larger σ means that goods are more substitutable, and a smaller θ means that the dispersion of the quality of ideas is greater, if σ were too high relative to θ, then consumers would tend toward consuming only the tiny measure of varieties with nearly infinite quality, and the price index would explode toward infinity.

which is increasing in Φ_n^j. However, the expected value of expenditure by country n on good j supplied by a producer in country i – that is, the expected value of expenditure on j times the probability that a producer in i is the one supplying it – is

$$E[X_{ni}^j] = \pi_{ni}^j E[X_n^j] = X_n \frac{T_i^j (w_i d_{ni})^{-\theta}}{\Phi_n} \left(\frac{\Phi_n}{\Phi_n^j} \right)^{\frac{\theta-(\sigma-1)}{\theta}}, \tag{10}$$

which is decreasing in Φ_n^j.

The intuition for this result is as follows. If relatively more research has been devoted to good j in the world, then its price is expected to be low relative to other goods, so consumers will spend more on this good in a relative magnitude governed by the elasticity of substitution σ. However, as the technological frontier moves outward, the likelihood that an increase in the level of research that has been devoted to j will lead to a new best idea, and hence a lower price, decreases at a rate governed by θ. Since $\theta > \sigma - 1$, the latter effect dominates, and the overall relationship between Φ_n^j and $E[X_n^j]$ is an increasing but concave one. So, since the likelihood that a producer in i possesses the best idea for delivering good j to n is inversely proportional to Φ_n^j, this ensures that a given producer's expected revenue from selling good j is decreasing in Φ_n^j. This result will be important for the dynamic model.

Now, assuming that the distribution of Φ_n^j across the space of goods is reasonably well behaved, integrating over all goods gives the total expenditure by consumers in n on goods from i,

$$X_{ni} = X_n \frac{\tilde{T}_{ni} (w_i d_{ni})^{-\theta}}{\Phi_n}, \tag{11}$$

where

$$\tilde{T}_{ni} = \Phi_n^{\frac{\theta-(\sigma-1)}{\theta}} \int_0^J \frac{T_i^j}{(\Phi_n^j)^{\frac{\theta-(\sigma-1)}{\theta}}} dj \tag{12}$$

is an index of the overall level of technology in country i in which the level of research devoted to each good is down-weighted by the overall level of research in the world devoted to that good. So, all else equal, a country will have a higher $\tilde{T}_{ni}$ if it tends to have had more research devoted to goods that have been the target of research less often in other countries.

Also, in the appendix, I show that the Bertrand competition assumption implies that a producer's expected profit is a constant share, $\frac{1}{1+\theta}$ of expected revenue, so the expected profit of producers in i from selling j in n is

$$E[\Pi_{ni}^j] = \frac{X_n}{1+\theta} \frac{T_i^j (w_i d_{ni})^{-\theta}}{\Phi_n} \left(\frac{\Phi_n}{\Phi_n^j}\right)^{\frac{\theta-(\sigma-1)}{\theta}}, \tag{13}$$

and so the overall profit earned by producers in i from selling in n is

$$\Pi_{ni} = \frac{X_n}{1+\theta} \frac{\tilde{T}_{ni}(w_i d_{ni})^{-\theta}}{\Phi_n}. \tag{14}$$

3.1.5 Labor Market

Suppose that in each country i, there are L_i^p workers available for production. Labor market clearing implies that total spending on labor in country i is

$$w_i L_i^p = \frac{\theta}{1+\theta} \sum_{n=1}^N X_{ni} = \frac{\theta}{1+\theta} \sum_{n=1}^N X_n \frac{\tilde{T}_{ni}(w_i d_{ni})^{-\theta}}{\Phi_n}. \tag{15}$$

Balanced trade requires that

$$X_i = \sum_{n=1}^N X_{ni}, \tag{16}$$

which, with the above equation, implies that

$$\frac{\theta}{1+\theta} X_i = w_i L_i^p, \quad \forall i.$$

Substituting this for all X_n into (15), then, gives

$$w_i L_i^p = \sum_{n=1}^N \frac{(w_i d_{ni})^{-\theta} \tilde{T}_{ni}}{\Phi_n} w_n L_n^p, \tag{17}$$

which implicitly determines wages in every country. A static trade equilibrium, then, is a set of wages $\{w_i\}_{i=1}^N$ that satisfy this condition, given that Φ_n and $\tilde{T}_{ni}$ are as defined above.

3.2 Research

The previous section described a world equilibrium of production and trade given the distribution of technology parameters T_i^j over the set of countries and the range of goods. This section will describe how the distribution of parameters is determined over time.

3.2.1 Types of Research

There are three activities in which a researcher in a given country can engage. First, a researcher can endeavor to create a new kind of product – that is, a new j. If he is successful, the return to this activity is three-fold. Upon "inventing" a product, he will possess, foremost, a concept for a new kind of product that consumers will demand. In addition, while in the process of developing this new concept, he will also have gained knowledge that will be useful in guiding the development of a particular design and production technique for the product.[15] And, in turn, this knowledge may immediately yield a few ideas for producing the product – that is, some Q's that apply to the j. A key assumption for what follows is that this knowledge obtained during the invention process immediately becomes publicly known in the country of invention but can only be obtained through effort in other countries. On the other hand, particular production ideas are entirely proprietary.

The second activity, then, is that of attempting to obtain the product specific knowledge known in other countries. Again, the act of obtaining this knowledge has the two-fold effect of immediately endowing the learner with a handful of ideas for producing the good as well focusing future research in the country of the learner. And, the third activity is that of simply attempting to develop an idea for a way to produce an existing type of good, taking advantage of the knowledge concerning the good that already is accessible in the country.

In terms of the total yield of a successful activity, the first is obviously more fruitful than the second, which, in turn, is more fruitful than the third. However, the level of difficulty (or likelihood of failure) is decreasing from the first to the third activity so that the particular option that is most attractive to a given researcher is ambiguous at this point and will depend on market forces. It is also useful to note that researching new products benefits other producers worldwide by increasing the space of goods which can be produced. Likewise, obtaining product specific knowledge that exits in a foreign country benefits other potential producers in the same country, while the benefits of researching new production ideas using domestic knowledge accrue entirely to the researcher.

Research effort in country i of type $r = n, f, d$ (corresponding to research in new goods and research in existing goods taking advantage of foreign and domestic knowledge, respectively) produces a flow of research of

$$R_{it}^{r} = \alpha_i^r (s_{it}^r)^\beta L_{it}, \tag{18}$$

where s_{it}^r denotes the share of the labor force in country i at time t devoted to research of

[15]This idea is similar in spirit to Nelson (1982), which discusses how knowledge of particular properties of a product can focus the search for production ideas. "...assume that the decision-maker knows more than merely the probability distribution of economic payoffs over the entire set of candidate techniques... [E]xpected economic payoff may vary with weight... Blue projects may be better than yellow projects. These correlates in general will not be foolproof guides, but they can enable the decision-maker to do better than he could merely by sampling randomly."

type r, $\alpha_i^r = \alpha_i \alpha^r$ is a measure of productivity of researches in country i devoted to research of type r, and $\beta \leq 1$.[16]

3.2.2 The Value of Research

Since all ideas are drawn from the same quality distribution regardless of when it was drawn, the expected profit at time t of selling in country n from a particular idea in country i that applies to good j, $h_{ni}^j(t)$, (not conditioning on quality) is simply the expected profit of all ideas – that is the best idea – in i pertaining to j, $E[\Pi_{ni}^j]$, relative to the expected number of ideas, T_i^j. Thus, we have that

$$h_{nit}^j = \frac{X_{nt}}{1+\theta} \frac{(w_i d_{ni})^{-\theta}}{\Phi_{nt}} \left(\frac{\Phi_{nt}}{\Phi_{nt}^j} \right)^{\frac{\theta-(\sigma-1)}{\theta}} = \frac{X_{nt}}{1+\theta} \frac{(w_i d_{ni})^{-\theta}}{\Phi_{nt}^{\frac{\sigma-1}{\theta}}} (\Phi_{nt}^j)^{\frac{(\sigma-1)-\theta}{\theta}}. \tag{19}$$

And, the total expected profit of a given idea in i pertaining to good j from selling anywhere in the world is then

$$h_{it}^j = \sum_{n=1}^{N} h_{nit}^j. \tag{20}$$

The expected discounted value of an idea in country i for producing j at time t is then

$$V_{it}^j = \int_t^\infty e^{-\rho(s-t)} \frac{P_{it}}{P_{is}} h_{is}^j ds, \tag{21}$$

where ρ is a common discount rate, and P_{is} is the price level in i at time s.

Since labor can move freely among production and the different types of research activity, it must be the case that, in equilibrium, if all activities are being undertaken, the marginal researcher must be indifferent between his research activity and production. So, we have the following conditions

$$\begin{aligned} \beta\alpha_i^n (s_{it}^n)^{\beta-1} V_{it}^n &= w \\ \beta\alpha_i^f (s_{it}^f)^{\beta-1} V_{it}^f &= w \\ \beta\alpha_i^d (s_{it}^d)^{\beta-1} V_{it}^d &= w \end{aligned} \tag{22}$$

where V_{it}^n is the expected value of an idea for a good that is new at time t, and V_{it}^f and V_{it}^d are the expected values of an idea calculated over the space of existing goods, weighted by the probability that an idea of type f or d research, respectively, applies to each good j.

[16] More formally, α_i^r can be thought of as the probability of success of a researcher in i engaged in activity r, while $(s_{it}^r)^\beta L_{it}$ is aggregate research intensity. So, R_{it}^r is the aggregate flow of successful research of type r in i.

3.2.3 Evolution of the Product Space

Suppose that there is a unit measure of good specific "knowledge" associated with each product j that is useful in developing ideas for producing j. As such, J_t is not only the measure of the set of goods that can be produced in the world but also the measure of the world stock of product-specific knowledge associated with this set of goods. Now, let J_{it} denote the measure of product-specific knowledge that is available to researchers in country i at time t. Similarly, let $J_{-it} = J_t - J_{it}$ denote the stock of product-specific ideas available only outside country i.

Ideas for new types of goods arrive in country i at rate R_{it}^n and are accompanied by the full measure of product-specific knowledge associated with the good. Similarly, a unit of type f research in country i brings a unit of knowledged randomly sampled from set J_{-it}, so a measure of existing knowledge equal to R_{it}^f arrives in i at t. So, at t the total measure of product-specific knowledge available in i increases by

$$\dot{J}_{it} = R_{it}^n + R_{it}^f.$$

However, since only type n research produces new product knowledge, the total measure of this knowledge in the world increases by

$$\dot{J}_t = \sum_{i=1}^{N} R_{it}^n.$$

Turning to a particular product j, denote the measure of product knowledge about j available in country i at t by $J_{it}^j \in [0, 1]$. Since a unit of learned product knowledge is equally likely to be any particular unit of knowledge not previously available in i, the probability that a new unit of learned knowledge applies to good j is $\frac{1-J_{it}^j}{J_{-it}}$. Therefore, the total measure of learned knowledge applying to good j at time t increases by

$$\dot{J}_{it}^j = R_{it}^f \frac{1 - J_{it}^j}{J_{-it}}. \tag{23}$$

3.2.4 Production Technology

As mentioned above, the number of production ideas in i that apply to good j is governed by the parameter T_{it}^j, where T_{it}^j is a measure of the stock of research effort that has been devoted to ideas for producing j. To formalize the evolution of T_{it}^j, suppose that an endeavor to obtain an idea is undertaken when a researcher combines a kernel of product-specific knowledge from the set J_{it} with research effort of type d. Then, T_{it}^j will increase by the intensity of these research effort if the kernel of knowledge was associated with good j, which is the case with probability $\frac{J_{it}^j}{J_{it}}$, the fraction of all product knowledge that applies to good j.

In addition, as mentioned above, the act of successfully developing a new product or of successfully acquiring product knowledge known elsewhere will endow the inventor or learner with the equivalent of some research effort applied to ideas for producing j. More formally, I assume that successfully inventing a new type of good or acquiring knowledge about it, endows the researcher with one unit of this effort (the equivalent of an expected value of one production idea).

Now, the level of T_{it}^j will be affected by whether the good was first developed in i, but its evolution will depend only on the aggregate intensities of the other two types of research as well as the measure of knowledge associated with it that is available in j at t. That is, recalling that a newly learned piece of product-specific knowledge applies to good j with probability $\frac{1-J_{it}^j}{J_{-it}}$, T_{it}^j increases by

$$\dot{T}_{it}^j = R_{it}^f \frac{1 - J_{it}^j}{J_{-it}} + R_{it}^d \frac{J_{it}^j}{J_{it}}. \tag{24}$$

Of course, if good j was first developed in i, then $J_{it}^j = 1$, so the above equation would reduce to $\dot{T}_{it}^j = \frac{R_{it}^d}{J_{it}}$.

3.2.5 Steady State

For simplicity, I consider a steady state in which each country devotes a constant share of its labor force, s_i^r to each type of research. In order for this steady state to emerge, I assume that the labor force everywhere grows at constant rate $\frac{\dot{L}_{it}}{L_{it}} = n$. This implies that growth in this model is of the semi-endogenous form first described in Jones (1995).

Define $\tau_{it} \equiv \frac{J_{it}}{L_{it}}$. In steady state, τ_i must be constant. That is,

$$\dot{\tau}_i = \frac{R_{it}^n + R_{it}^f}{L_i t} - \tau_i n = 0,$$

so it must be the case that

$$\tau_i = \frac{\alpha_i^n (s_i^n)^\beta + \alpha_i^f (s_i^f)^\beta}{n},$$

which implies that the stock of product-specific knowledge in each country grows at rate

$$\frac{\dot{J}_{it}}{J_{it}} = \frac{(\alpha_i^n (s_i^n)^\beta + \alpha_i^f (s_i^f)^\beta) L_{it}}{J_{it}} = n. \tag{25}$$

In addition, the overall stock of product-specific knowledge must grow at the same rate, $\frac{\dot{J}_t}{J_t} = n$. This combined with (25) implies that fraction of the total stock of product ideas that is available in i at any time is

$$\frac{J_i}{J} = \frac{\dot{J}_{it}}{\dot{J}_t} = \frac{R^n_{it} + R^f_{it}}{\sum_{i=1}^N R^n_{it}} = \frac{\alpha^n_i (s^n_i)^\beta + \alpha^f_i (s^f_i)^\beta}{\sum_{m=1}^N \alpha^n_m (s^n_m)^\beta \frac{L_m t}{L_i t}}, \tag{26}$$

since the world stock of product knowledge expands as new products are developed in each country (but not as previously existing knowledge in learned).[17] Similarly, the proportion of the world stock of product knowledge that was first developed country i (that is, the fraction of goods that were developed in i), denoted η_i, is

$$\eta_i = \frac{R^n_{it}}{\sum_{i=1}^N R^n_{it}} = \frac{\alpha^n_i (s^n_i)^\beta}{\sum_{m=1}^N \alpha^n_m (s^n_m)^\beta \frac{L_m t}{L_i t}} \tag{27}$$

Since the labor force – and, hence, research output – in each country is growing at the same rate as the stock of product knowledge, it follows that the evolution of the stock of product knowledge in i for a particular good j (given that j was not invented in i) becomes

$$\dot{J}^j_{it} = \frac{R^f_{it}}{J_{-it}}(1 - J^j_{it}) = \delta_i (1 - J^j_{it}),$$

where

$$\delta_i = \frac{R^f_{it}}{J_{-it}} = n \frac{\alpha^f_i (s^f_i)^\beta}{\left(\sum_{m \neq i} \alpha^n_m (s^n_m)^\beta \frac{L_{it}}{L_{mt}}\right) - \alpha^f_i (s^f_i)^\beta}. \tag{28}$$

This implies that the measure of knowledge pertaining to good j that is *not* available in i, $1 - J^j_{it}$ experiences exponential decay. So, if j was invented at time $t - a^j$ in a country other than i, at time t,

$$1 - J^j_{it} = e^{-\delta_i a^j}.$$

Since this value depends only on the *age* of the good, a^j, – as well as where it was invented – and not on the identity of the good, per se, it is useful to refer to the goods by age rather than by their index label, j. As a result, the stock of product knowledge available to a researcher in country i as a function of the age of the good is given by

$$J_i(a, m) = \begin{cases} 1 & \text{if } m = i \\ 1 - e^{-\delta_i a} & \text{otherwise} \end{cases} \tag{29}$$

where m is the country in which the product originated.

[17]Note that the time subscripts have been dropped to emphasize the stationarity of the ratio.

Similarly, in steady state, the amount of research devoted to obtaining ideas for producing a good of age a reduces to

$$\frac{R_{it}^d}{J_{it}} J_{it}^j = \kappa_i(1 - e^{-\delta_i a}),$$

where

$$\kappa_i = \frac{R_{it}^d}{J_{it}} = n \frac{\alpha_i^d (s_i^d)^\beta}{\alpha_i^n (s_i^n)^\beta + \alpha_i^f (s_i^f)^\beta}. \tag{30}$$

This, together with (24) implies that

$$\dot{T}_i(a,m) = \begin{cases} \kappa_i & \text{if } m = i \\ \delta_i e^{-\delta_i a} + \kappa_i(1 - e^{-\delta_i a}) & \text{otherwise} \end{cases} \tag{31}$$

which, in turn, implies that, at a given point in time, the parameter governing the distribution of technology for producing a good of age a in country i is

$$T_i(a,m) = \int_0^a \dot{T}_i(a,m)da = \begin{cases} 1 + \kappa_i a & \text{if } m = i \\ (1 - e^{-\delta_i a})(1 - \frac{\kappa_i}{\delta_i}) + \kappa_i a & \text{otherwise} \end{cases} \tag{32}$$

This expression is key for what follows, so it is useful to briefly explore its properties. Note first that the inventing country of a good gets an immediate technological advantage; it possesses an idea (in expectation) for producing the good before researchers in other countries know anything about the good. As time progresses, $T_i(a,i)$ grows at rate κ_i – the rate at which ideas arrive for goods for which researchers in i possess knowledge of the product. This is the result of the assumption that researchers in the inventing country possess the full measure of product knowledge for a particular good. If a good was invented in another country, $T_i(a,m)$ is initially equal to 0 and grows at a rate which is a convex combination of κ_i and δ_i – the rate at which researchers acquire knowledge of products invented abroad.

Figure 3 illustrates T_i for different values of δ_i for goods invented in i and elsewhere ($\kappa_i = 0.3$ in all cases).[18] This illustration demonstrates how δ_i controls the rate at which the slope of T_i converges to κ_i.

Figures 4 and 5 illustrate how the ratio $\frac{T_i}{T_j}$ evolves over time for countries with differing values of δ and κ. In both figures, $\delta_i = 0.01$, $\delta_j = 0.05$, $\kappa_i = 0.3$, and $\kappa_j = 0.6$. Figure 4 shows the case in which the good was invented in neither country. The country with a higher value of δ has a technological advantage early in the life of the good (converging to $\frac{\delta_i}{\delta_j}$ as $a \to 0$) with this advantage eroding over time and converging to $\frac{\kappa_i}{\kappa_j}$. Figure 5 illustrates

[18] In the calibration below $\delta_i < \kappa_i$ for all countries.

Figure 3: Examples of $T_i(a, m)$

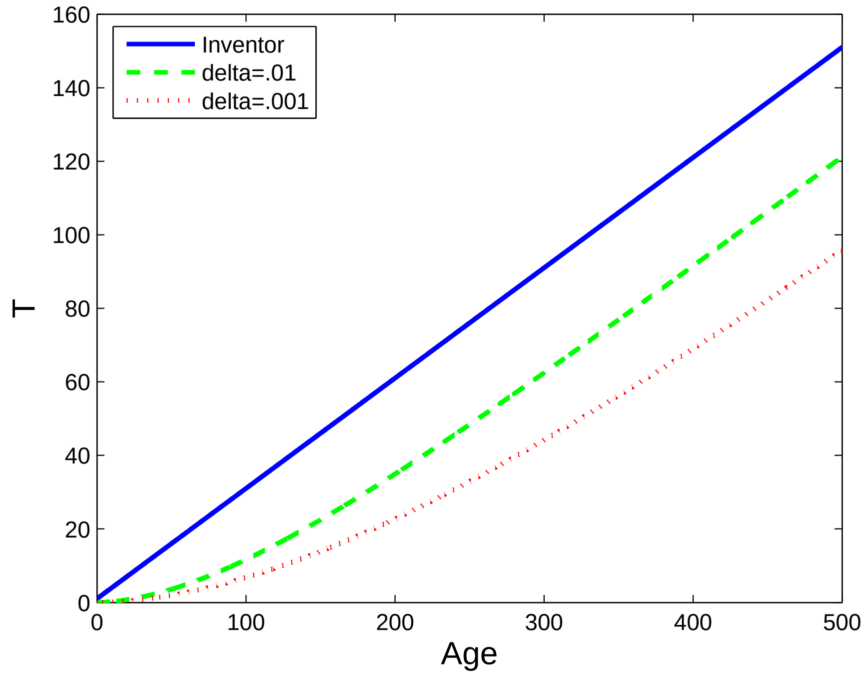

Figure 4: Ratio of T_i's for Different δ_i and κ_i

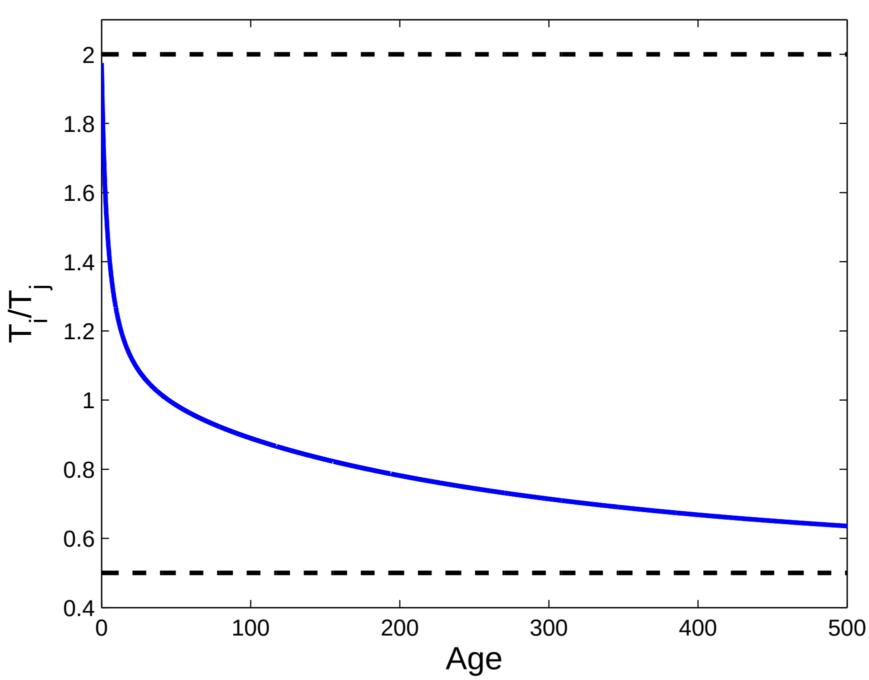

Figure 5: Ratio of T_i's; Invented vs. Learned

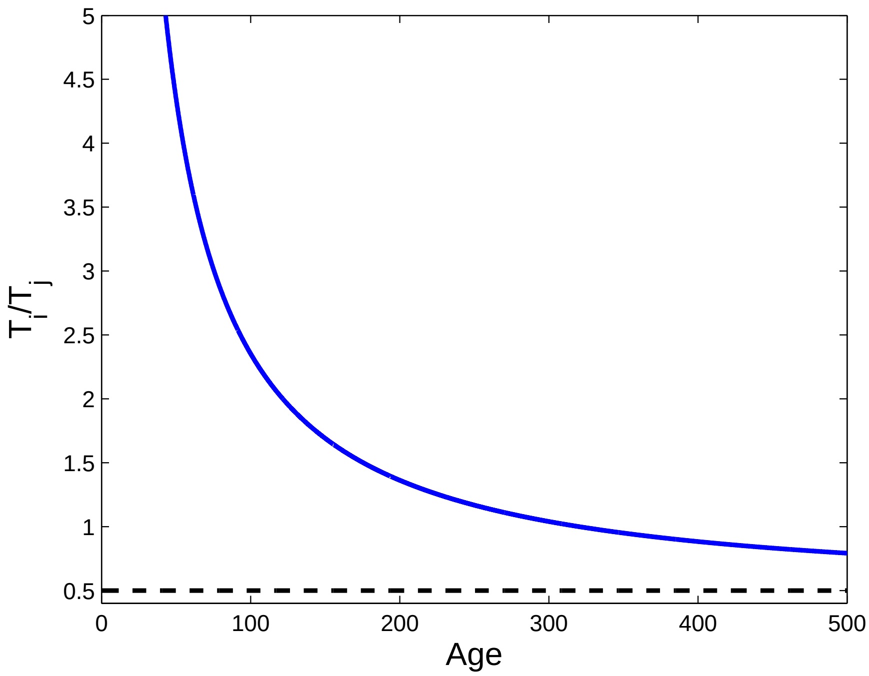

the case in which the good was invented in country i. In this case the early technological advantage is much higher (converging to inf as $a \to 0$), but as the good ages the plot looks very similar that of figure 4, with the ratio similarly converging to $\frac{\kappa_i}{\kappa_j}$.

3.2.6 Steady State Research

These results greatly simplify the form of the value of each type of research. First of all, note that, in a steady state, while research effort is increasing the quality of production ideas for a given good over time, research effort is also increasing the space of goods by introducing new goods with lower quality production ideas. As a result, the distribution of technology parameters T_i^j is invariant over time. An immediate result of this is that Φ_{nt}, and hence the price index in each country, changes only with the world measure of product knowledge, J_t, over time.

More formally, note that $T_i(a, m)$ depends only on the age of a given product and where it was invented, with the distribution of the latter depending on the value η_i, which does not change in steady state. Next, note that

$$\Phi_{nt}^j(a, m) = \Phi_n(a, m) = \sum_{i=1}^{N} (w_i d_{ni})^{-\theta} T_i(a, m) \tag{33}$$

depends only on the values of the T_i's across the range of product ages. And, finally note that since J_t grows at a constant rate, the the mass of goods that are of a particular age a at any point in time is $ne^{-na}J_t$, so that a constant proportion of goods

$$f(a) = ne^{-na} \tag{34}$$

is a given age at every point in time.

As a result, we can rewrite the equation for Φ_{nt} as follows:

$$\Phi_{nt} = \left(\int_0^{J_t} (\Phi_{nt}^j)^{\frac{\sigma-1}{\theta}} dj \right)^{\frac{\theta}{\sigma-1}} = \left(J_t \int_0^{\infty} \sum_{i=1}^{N} (\Phi_{nt}(a,i))^{\frac{\sigma-1}{\theta}} \eta_i f(a) da \right)^{\frac{\theta}{\sigma-1}} = J_t^{\frac{\theta}{\sigma-1}} \Phi_n,$$

where Φ_n is the time invariant component if Φ_{nt}. This implies that the price index in n at t can be expressed as

$$P_{nt} = \gamma \Phi_{nt}^{\frac{-1}{\theta}} = \gamma J_t^{\frac{1}{1-\sigma}} \Phi_n^{\frac{-1}{\theta}} = J_t^{\frac{1}{1-\sigma}} P_n. \tag{35}$$

Now, the value of an idea that applies to a particular good j in i becomes

$$V_{it}^j = \int_t^{\infty} e^{-\rho(s-t)} \frac{P_{it}}{P_{is}} h_{is}^j ds = \int_t^{\infty} e^{\left(\frac{n}{\sigma-1} - \rho\right)(s-t)} h_{is}^j ds = \int_0^{\infty} e^{\left(\frac{n}{\sigma-1} - \rho\right)s} h_{is}^j ds = V_i^j$$

Furthermore, since labor market clearing and the balanced trade requirement imply that $X_n = \frac{1+\theta}{\theta} w_n L_n^p$, h_{nit}^j can be written as a function of a and m as

$$h_{ni}(a,m) = \frac{X_{nt}}{1+\theta} \frac{(w_i d_{ni})^{-\theta}}{\Phi_{nt}^{\frac{\sigma-1}{\theta}}} \left(\Phi_n^j(a,m)\right)^{\frac{(\sigma-1)-\theta}{\theta}} = \frac{w_n s_n^p}{\tau_i \theta} \frac{(w_i d_{ni})^{-\theta}}{\Phi_n^{\frac{\sigma-1}{\theta}}} \left(\Phi_n^j(a,m)\right)^{\frac{(\sigma-1)-\theta}{\theta}} \tag{36}$$

So, then the value of selling in n of an idea in i that applies to a particular good becomes

$$V_{ni}(a,m) = \frac{w_n s_n^p}{\tau_i \theta} \frac{(w_i d_{ni})^{-\theta}}{\Phi_n^{\frac{\sigma-1}{\theta}}} \int_0^{\infty} e^{\left(\frac{n}{\sigma-1} - \rho\right)s} \left(\Phi_n(a+s,m)\right)^{\frac{(\sigma-1)-\theta}{\theta}} ds. \tag{37}$$

We now have all the elements needed to express the expected value of an idea that results from a particular type of research by a researcher in country i, V_i^r. The simplest of these is V_i^n. Given that an idea for a producing a good was obtained from research into new products, the product the idea pertains to must be new (that is, age $a = 0$) when it arrived, and it must have originated in country i, where the researcher is. Therefore, the value of this idea is

$$V_i^n = \frac{1}{\theta} \sum_{n=1}^{N} \frac{(w_i d_{ni})^{-\theta}}{\Phi_n^{\frac{\sigma-1}{\theta}}} \frac{w_n s_n^p}{\tau_n} \frac{J_n}{J} \int_0^{\infty} e^{\left(\frac{n}{\sigma-1} - \rho\right)s} \left(\Phi_n^j(s,i)\right)^{\frac{(\sigma-1)-\theta}{\theta}} ds \tag{38}$$

The expressions for the value of ideas resulting from learning foreign product-specific knowledge and from utilizing existing domestic knowledge are a bit more complex because they can potentially apply to goods of any age and origin. First, consider the case of learning foreign knowledge. Recall that the probability that a researcher acquires a kernel of knowledge pertaining to a particular good is $\frac{1-J_i(a,m)}{J_{-it}}$, given that the good is of age a and origin m, and the measure of goods that are of age a is $f(a)J_t$. Then, recalling that the probability that a given product originated in country m is η_m,

$$V_i^f = \frac{1}{\theta} \sum_{n=1}^{N} \frac{(w_i d_{ni})^{-\theta}}{\Phi_n^{\frac{\sigma-1}{\theta}}} \frac{w_n s_n^p}{\tau_n} \frac{J_n}{J} \sum_{m=1}^{N} \eta_m \times \int_0^\infty \int_0^\infty e^{\left(\frac{n}{\sigma-1}-\rho\right)s} \left(\Phi_n^j(a+s,m)\right)^{\frac{(\sigma-1)-\theta}{\theta}} ds \left(\frac{1-J_i(a,m)}{1-(J_i/J)}\right) f(a)da. \quad (39)$$

Similarly, recalling that the probability that probability that an idea for a producing a good obtained using existing domestic knowledge applies to a particular good of age a and origin m is $\frac{J_i(a,m)}{J_{it}}$, the expected value of such an idea is

$$V_i^d = \frac{1}{\theta} \sum_{n=1}^{N} \frac{(w_i d_{ni})^{-\theta}}{\Phi_n^{\frac{\sigma-1}{\theta}}} \frac{w_n s_n^p}{\tau_n} \frac{J_n}{J} \sum_{m=1}^{N} \eta_m \times \int_0^\infty \int_0^\infty e^{\left(\frac{n}{\sigma-1}-\rho\right)s} \left(\Phi_n^j(a+s,m)\right)^{\frac{(\sigma-1)-\theta}{\theta}} ds \left(\frac{J_i(a,m)}{J_i/J}\right) f(a)da. \quad (40)$$

While these values do not have tractable analytic expressions, they can be computed numerically using standard techniques.

4 Equilibrium

A steady state research equilibrium is a set of time-independent research intensities $\{\{s_i^r\}_{i=1}^n\}_{r\in\{n,f,d\}}$ that satisfy the research labor market incentive compatibility conditions given that wages $\{w_i\}_{i=1}^N$ comprise a static trade equilibrium, and the value of each form of research is as defined above.

4.1 Calibration

In order to bring the model to the data, I must assign values to the set of 8 universal parameters $\{r, n, \beta, \theta, \sigma, \{\alpha^r\}_{r\in\{d,f,n\}}\}$ as well as to the set of country-specific parameters, α_i and L_i, and the bilateral trade cost parameters, d_{ni}. Table 3 lists the values given to the 9 universal parameters.

I chose r to match the average rate of return on US treasury bills and n to match the current world population growth rate. θ is set to 6.66, the median value used in Alvarez &

Table 3: Parameter Values

Parameter	Value	Description	Source
r	0.02	Discount rate	data
n	0.0117	Population growth rate	data
β	0.25	Diminishing returns to research	OECD R&D Exp.
θ	6.66	Technology distribution parameter	Alvarez & Lucas (2006)
σ	4.0	Elasticity of substitution	Broda & Weinstein (2006)
α^d	1.0		Normalization
α^f	0.02	Research productivity parameters	Mansfield & Romeo (1980)
α^n	0.04		Mansfield, Schwarts, & Wagner (1981)

Lucas (2006) based on their survey of estimates in the literature. σ is set to 4.0, a value between the median and mean values estimated by Broda & Weinstein (2006) for imports to the U.S. at the 5 digit SITC product level. β is chosen to match expenditure on R&D as a percentage of GDP in the OECD of 2.4%. α^d is normalized to unity. α^n is chosen so that κ_{USA} matches the percentage of successful innovations that were imitated within one year in the US (23%) from Mansfield, Schwarts, & Wagner (1981). α^d is chosen so that $\sum_i \neq USA\delta_i$ is equal to the .25 to match the average time of 4 years that it took a non-US competitor to obtain a US technology in Mansfield & Romeo (1980).

The country specific populations, Li, are the taken from the size of the labor force from the World Bank's World Development Indicators. The country specific research productivity parameters, α_i, are chosen to match manufacturing wages from the UNIDO INDSTAT database. Data on manufacturing output, used to compute trade shares used in the estimation of iceberg trade costs, below, is also takend from the INDSTAT database. Availability of this data reduces the sample to 60 countries.[19]

In models that deliver an aggregate gravity equation, a common method of estimating the value of iceberg trade costs is to estimate the log-linear form of this equation using proxies for trade costs such as bilateral distance.[20] The model of this paper does not, in general, deliver a log-linear gravity equation for aggregate trade flows, but it does for product level trade trade flows.

This equation is based on the expression of the expected value of expenditure on good j from country i by country n.

$$E[X_{ni}^j] = \frac{T_i^j (w_i d_{ni})^{-\theta}}{\Phi_n^j} \left(\frac{\Phi_n^j}{\Phi_n} \right)^{\frac{\sigma-1}{\theta}} X_n$$

[19]Table 8 provides a list of countries.

[20]See, e.g., Eaton & Kortum (2002) and Waugh(2009).

Taking logs, we have the following expression

$$\ln(E[X_{ni}^j]) = S_i^j + S_n^j - \theta \ln(d_{ni})$$

where $S_i^j = \ln(T_i^j w_i^{-\theta})$ and $S_n^j = \ln\left(\frac{X_n}{\Phi_n^{\frac{\sigma-1}{\theta}}}\left(\Phi_n^j\right)^{\frac{(\sigma-1)-\theta}{\theta}}\right)$.

In principle, one can estimate this equation using product level trade flows, importer-product and exporter-product fixed effects and common proxies for trade costs from the gravity literature. However, using data from 60 countries in over 4,000 product codes requires the use of nearly 500,000 dummy variables, making the estimation extraordinarily computationally intensive. So, in order to simplify the estimation, I take an alternative, though admittedly more approximate, approach. Dividing the value of aggregate trade between a pair of countries by the importing country's absorbtion of domestic output gives the following expression.

$$\frac{X_{ni}}{X_{nn}} = \frac{\tilde{T}_{ni}}{\tilde{T}_{nn}}\left(\frac{w_i}{w_n}\right)^{-\theta} d_{ni}^{-\theta} = \frac{\tilde{T}_{ni}}{\tilde{T}_{nn}}\frac{\tilde{T}_{ii}}{\tilde{T}_{nn}}\left(\frac{w_i}{w_n}\right)^{-\theta} d_{ni}^{-\theta} \tag{41}$$

Taking logs, this can be rewritten as

$$\ln\left(\frac{X_{ni}}{X_{nn}}\right) = S_i + S_n - \theta \ln(d_{ni}) + S_{ni} \tag{42}$$

where $S_i = \ln(\tilde{T}_i i) - \theta \ln w_i$ and $S_{ni} = \ln \frac{\tilde{T}_{ni}}{\tilde{T}_{ii}}$. Neglecting the last term, this equation can be estimated via OLS using source and destination effects and proxies for trade barriers, as in Eaton & Kortum (2002) and Waugh (2009).[21] Of course, neglecting this term introduces an omitted variable bias. However, since S_{ni}, which is a function of how country i's accumulated research is distributed across the product space, is not likely to be highly correlated with the gravity variables used to proxy for $d_n i$, the extent of this bias is expected to be small, so I proceed with this estimation strategy.[22] An added benefit of this method is that it makes the results presented in the next section comparable with Eaton & Kortum (2002).

I parameterize d_{ni} by

$$\ln(d_{ni}) = d_k + b + l + c$$

where d_k $(k = 1, ..., 6)$ is the effect of the distance between n and i lying in the the kth interval, b is the effect of n and i sharing a border, and l is the effect of n and i have a

[21] French (2009) uses another strategy to estimate this equation, using highly disaggregated product level U.S. trade data to estimate the value of S_{ni}.

[22] S_{ni} is expected to be highly correlated with S_i, making the estimated value of this term much more unreliable. However, while Eaton & Kortum (2002) and Waugh (2009) use this value to calibrate countries' average productivity parameters, T_i^j, here, is endogenously determined after α_i is calibrated to match data on wages. So, this bias is not a problem in that respect.

Table 4: Trade Cost Parameter Estimates

Parameter	Value	s.e.
$-\theta d_1$	2.86	0.29
$-\theta d_2$	3.54	0.25
$-\theta d_3$	4.31	0.24
$-\theta d_4$	5.24	0.23
$-\theta d_5$	6.07	0.23
$-\theta d_6$	6.71	0.22
θb	0.44	0.12
θl	0.70	0.08
θc	0.45	0.13

common language, and c the effect of n and i having a colonial relationship. Since my model allows for endogenous asymmetries in trade flows, I consider only a symmetric specification of trade costs, similar to Fieler (2007) and unlike Waugh (2009) and Eaton & Kortum (2002), who include an exporter and importer fixed effect, respectively, in the specification of trade costs in order to account for asymmetries in aggregate trade flows. The parameter estimates are given in table 4.[23]

4.2 Endogenous Variables

Figures 6 - 8 illustrate the relationship between the research decisions made in equilibrium and countries' equilibrium wage rate. Figure 6 shows that more researchers in poor countries devote effort to developing processes for producing existing goods for which they already have sufficient product knowledge. Figure 7 shows that middle income countries devote relatively more research effort to learning about products invented abroad, while figure 8 shows that rich countries are overwhelmingly the inventors of new goods. The intuition for this result is the following. Any researcher would prefer to have an idea for a product with the lowest possible Φ_n^j. Since the technology levels, T_i^j, that make up Φ_n^j are down-weighted by production wages, researchers prefer to target goods for which technology is concentrated in rich countries. For researchers in rich countries, inventing a new good ensures this, since it will take time for researchers in poor countries to learn enough about the good to get a good production idea for it. For researchers in poor countries, this can be done by learning

[23] Whether because of bias from omitting S_{ni} (as assumed here) or from an importer or exporter specific trade cost shifter (as in Eaton & Kortum (2002) and Waugh (2009)), the estimates of S_i from the source and destination effects differ. The choice of which effect to take as the "true" value shifts the estimates of d_k. Following Eaton & Kortum (2002), I normalize the average difference between the two effects among OECD countries to 0.

Figure 6: Domestic Process Research

about goods invented abroad.

Table 8 in the appendix lists η_i, the percentage of products invented in country i. As rich countries are the ones doing most new goods research, it is not surprising that the large industrial economies account for most new goods. For example, the United States and Germany account for nearly 60% of inventions in the world, where, by contrast, China and India together account for less than 1%. This outcome of the model is also consistent with the observation that these countries account for the vast majority of international patents.

Also of note from this table is that κ_i does not vary greatly (as opposed to η_i and δ_i) across countries. This is because, while η_i and δ_i are measures of research output relative to the size of the world stock of product knowledge, κ_i is a measure of research output relative to the domestic measure of product knowledge. This is responsible for poor countries catching up in relative productivity for a given good over time.

5 Results

5.1 Product Level Trade

In order to assess how well the model fits the data above, I compute model analogues to figure 2 and table 2, above, using moments from the computed model equilibrium. To this

Figure 7: Foreign Learning

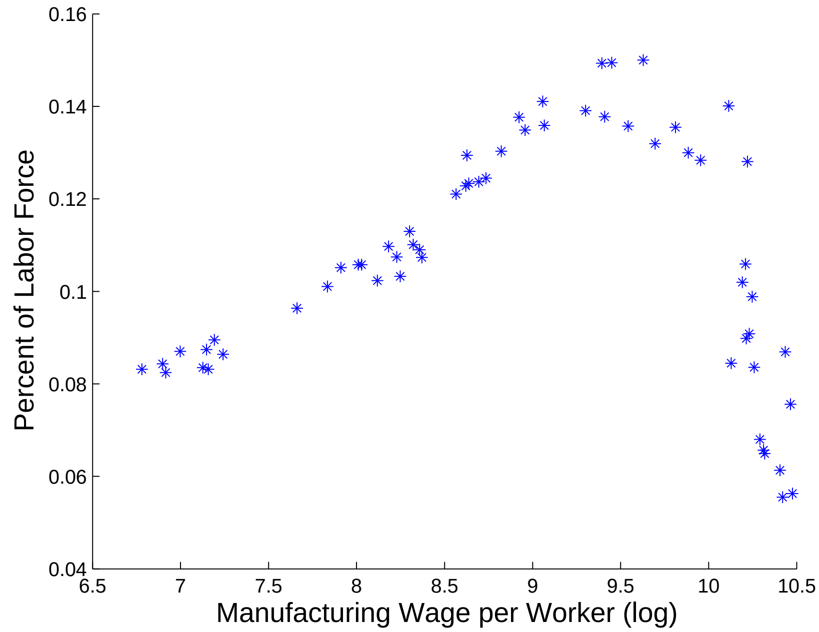

Figure 8: New Product Research

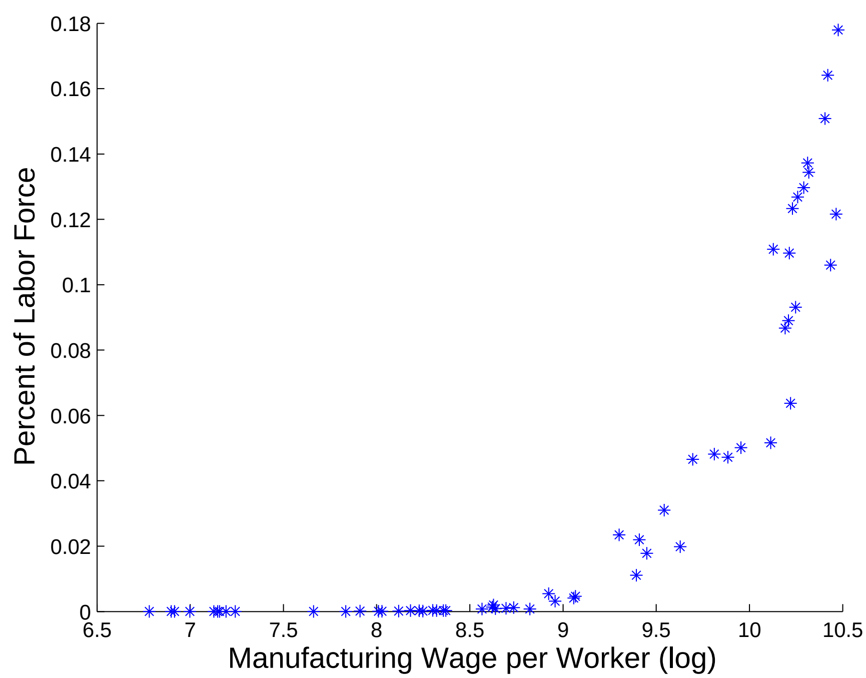

Figure 9: Percentage of Exports by Product Group

Developed Countries (Median)

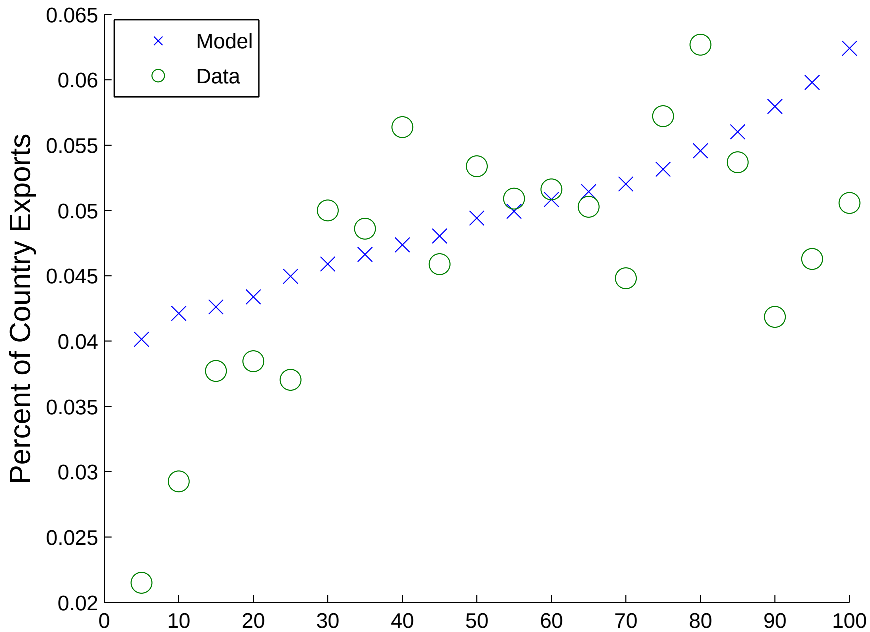

end, I performed a Monte Carlo simulation using 5000 products whose age is drawn from the model age distribution $f(a)$. Once the age of a product is known, the probability that country i will export it to country n, $\pi_{ni}(a)$, can be computed, as well as the expected expenditure on the product by country n, $X_n(a)$. Then, the expected value of exports from country i to country n of a product of age a is $X_{ni}(a) = \pi_{ni}(a)X_n(a)$. Summing $X_{ni}(a)$ over $\{n, i|n \neq i\}$ gives the total world trade volume of the product, so using this information to construct bins representing 5% of world trade analogously to the ones used above, where now the first bin contains the 5% of world trade made up by the oldest products and the last bin that of the the newest.

Figures 9 and 10 compare the median percentage of rich and poor countries' exports in each bin to that from the data as presented in figure 2.[24] Figure 9 shows that the model does quite well in accounting for the distribution of the exports of rich countries across the space of products. Figure 10 shows that the model correctly accounts for the fact that poor countries' exports are concentrated in a common set of products; however,

[24]Values from the data are recalculated using the smaller sample of countries for which data was available for the calibrated model.

Figure 10: Percentage of Exports by Product Group

Less Developed Countries (Median)

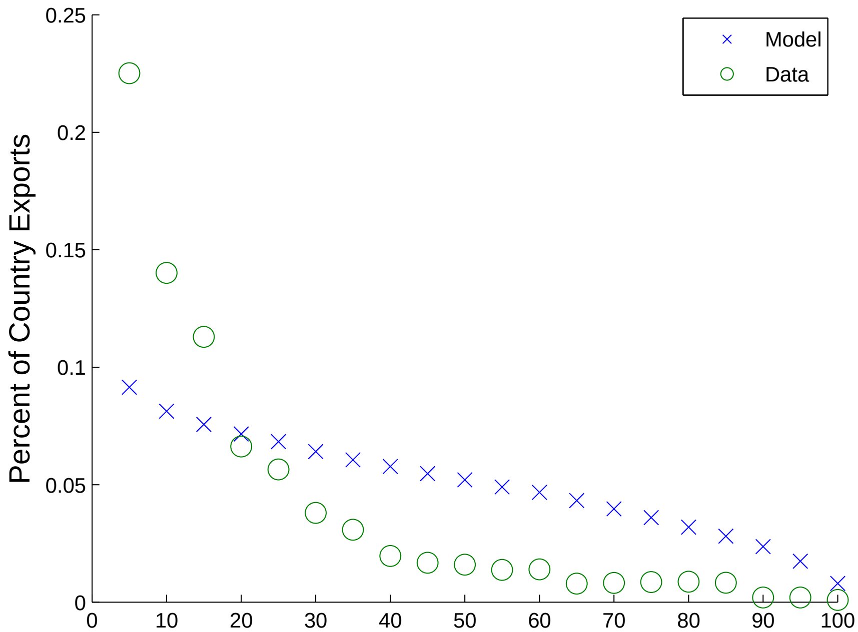

Table 5: Model Exports by Product and Income Group

	Model		Data	
GDP/Worker	Bin 1	Bin 20	Bin 1	Bin 20
High	4.01%	6.24%	2.15%	5.06%
Middle	8.85%	0.96%	8.45%	0.67%
Low	9.15%	0.79%	22.51%	0.10%

the predicted level of concentration is not nearly as extreme as that observed in the data. Table 5 demonstrates this point in the table 2. Again, the median percentage of exports that fall into the first and last bins is quite close to the data for rich as well as middle income countries, and the exports of poor countries are more concentrated in a single bin than those of other countries, but the degree on concentration is this bin is below that in the data.

There are two main reasons for this outcome. First, though poor countries have an extremely small chance of being the inventor of a new good and learn slowly about new goods invented elsewhere, their low wages imply that, were a poor country to develop in idea for producing a very new product, they would export so much of it that the expected value of exports of new goods is not negligible, so poor countries export more in the bin of newest goods in the model than in the data. Second, the assumption on the learning process that a unit of product specific knowledge acquired by a researcher is randomly sampled from the set of knowledge outside the researcher's country (similar to the setup of Krugman (1979b)) implies that fraction of the knowledge applying to a particular product that is available in a country, $J_i(a,m)$, takes the form of an exponential CDF as a function of the age of the good. Since the level of technology in a country for a particular good, $T_i(a,m)$, grows at the rate of learning, δ_i, early in the life of the good, the level of technology in a country relative to the rest of the world, $\frac{T_i(a,m)}{\Phi_n(a,m)}$, which governs comparative advantage, inherits the shape of the exponential CDF for newer products. This comes through in the shape of the plot in figure 10, as the percentage of poor countries' exports falling into a 5% bin is concave at first, moving from right to left, implying that the model predicts more exports for these countries of the products for which they have a comparative disadvantage than is seen in the data.

5.2 Aggregate Trade

The model does quite well in accounting for another shortcoming of theoretical gravity models of trade, that poor countries do not trade nearly as much, relative to their economic size, as these models models predict. The two most prominent attempts to reconcile the data with these models are Feiler (2007) and Waugh (2009). Feiler argues that different levels of dispersion in the idiosyncratic productivity draws across products combined with non-homothetic preferences account for the differences between the model and data. Waugh argues that for these models to be consistent with the asymmetric levels of trade across income groups and the small differences in prices of tradable goods across income groups requires systematically asymmetric trades costs such that it is much more costly for poor countries to export their products than for rich countries.

In this section, I will show that my model also goes very far in reconciling these models with the data. In addition, this model does so while providing a deep theory for the modification to the established trade models, rather than imposing trade cost or technology parameters based solely on their ability to match aggregate trade flow data. The appendix also contains a discussion of how the main mechanism of the models is similar to that of Feiler's but without being subject to the critique of Waugh (2009) that its predictions about the trade cost elasticity are not borne out in the data.

Figure 11: Total Trade and Income

Data: regression slope = 0.0081

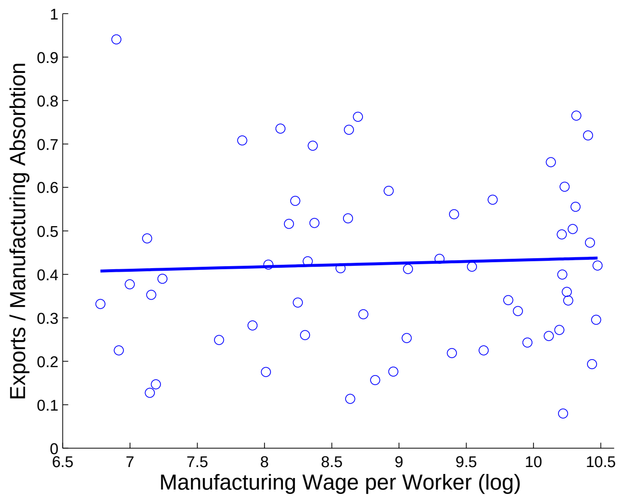

New Model: regression slope = -0.0011

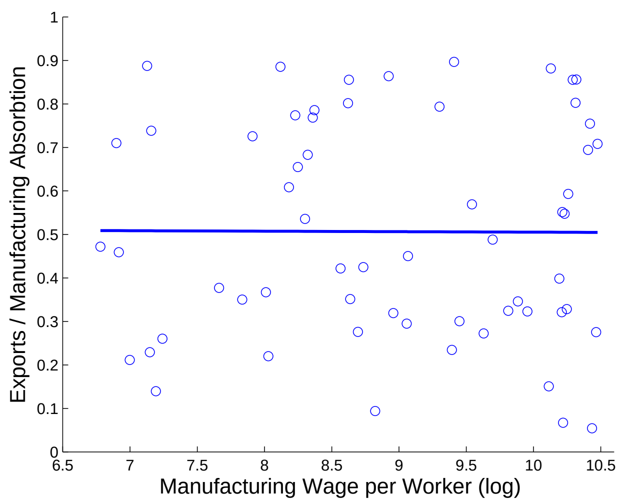

Eaton & Kortum: regression slope = -0.0325

Figure 6 shows that, using trade costs that are symmetric between importer-exporter pairs, as in Feiler (2007), that the model of Eaton & Kortum (2002) overpredicts the percentage of output which is traded by poor countries.[25] The predictions of my model come much closer to the data. The model predicts a very slight negative relationship between trade and income per worker, while the relationship is slightly positive in the data, which is much closer than the clear negative relationship predicted by Eaton & Kortum (2002).

The intuition for this result is that, since rich countries are the major inventors of new goods, they will have a strong comparative advantage in producing the goods invented in their country. As a result, the level of trade costs will have little effect on whether they are the low cost producers of these goods. Since the CES preferences admit a love of variety, all countries will demand these goods from the inventor country. So, rich countries will trade a great deal among each other even though they have a comparative advantage in the same types of goods because their comparative advantage will be quite idiosyncratic across all the varieties of these goods, as they were invented in different countries. Poor countries, on the other hand, are very good at producing a small set of goods that is common across the set of poor countries. This set is common because it consists of the goods that were invented in rich countries sufficiently long ago. The set is small relative to the total mass of products because the product space is expanding exponentially, so a given cohort makes up a much smaller proportion of the product space as it ages. As a result, since everyone is very good at producing these goods, even a small trade costs swamps the gains due to idiosyncratic productivity differences. So, poor countries consume much more of their own output.

6 Counterfactuals

Having in hand a model that determines technology levels and trade flows and is calibrated to salient moments of the data, I can now perform a series of counterfactual experiments. I, first, consider the effects of moving to a world of "free trade" in which there are no trade costs ($d_n i = 1$). Second, I consider the opposite extreme of prohibitively high trade costs that lead to a world of autarky in trade (but in which countries can still learn about one

[25]More precisely, Feiler (2007) points out that EK cannot simultaneously account for the level of trade of both rich and poor countries. In her paper, when trade costs are calibrated to match the data as closely as possible, EK comes closer to matching trade for poor countries, but then predicts that rich countries trade much too little.

another's products).

6.1 Free Trade

I consider the effect of a permanent change in trade costs at two points in time. First, I evaluate the effect in the static trade model, keeping levels of technology equal to the baseline steady state values. Second, I consider effect at the new steady state levels of research and technology levels implied by the model. This is indicative of effects of the change in the short and very long run, respectively.

In the case of the move to free trade, in the short run, not surprisingly, the real wage $\frac{w_i}{P_i}$ rises dramatically in all countries. Furthermore, the real wage increases to a greater extent in smaller and poorer countries relative to larger, richer ones. For example, the real wage of the average developed country increases by 87%, while that of the average developing country increases by 171%. Further, the real wage for the Germany grows by 42% while that of Belgium grows by 90%. Similarly, the real wage in China grows by 87% while that of Thailand grows by 184%. By contrast the model of Eaton & Kortum (2002) calibrated above predicts a smaller rise in the real wage and a smaller degree to which poor countries gain relative to rich countries. That model predicts, for example, that the real wage of the average developed developing country increases by 69% and 114%, respectively.

The intuition for this difference in predictions lies in the same mechanism by which this model better predicts the relationship between income level and total trade. Poor countries comparative advantage lies in products for which there is less variation in average productivity levels across countries, so even small trade costs can impede the gains from trade associated with idiosyncratic productivity differences across countries. Rich countries, on the other hand, export newer goods which many countries are nearly incapable of producing. So, lowering trade costs disproportionately benefits poorer countries.

The long run effects on levels of the three types of research of a move to free trade can be summarized as follows. In all countries, more researchers devote themselves to inventing new goods and fewer learn about goods invented abroad. Rich countries continue to perform about the same level of research on production techniques for domestic goods, while poor countries perform less. Table 6 summarizes these results.

The major reason for this result is that, in the presence of trade costs, a large portion of the expected profit from learning about a product is in serving the domestic market, while the majority of profit from a new idea is from exporting. Consider the expression for the

Table 6: Change in # of Researchers; Free Trade

GDP/Worker	s_i^d	s_i^f	s_i^n
High	0.0%	-58.6%	45.8%
Low	-1.2%	-14.4%	72.1%

expected flow of profit for a particular idea in i from selling in n

$$h_{ni}(a,m) = \frac{(w_i d_{ni})^{-\theta}}{\Phi_n(a,m)} \left(\frac{\Phi_n(a,m)}{\Phi_n}\right)^{\frac{(\sigma-1)}{\theta}} \frac{X_n}{1+\theta}$$

where the first term represents that probability that country n buys the good produced with this idea, and the second is the expected by n on the good. When a good is invented, it tends to be produced relatively inefficiently, so the expected price of a new good is high, so the expected expenditure on the good by a given country will be relatively low. However, since, initially, the inventor of the product is the only one who can produce it, the probability of being the low cost producer is equal to one (i.e. $\Phi_n(0,i) = (w_i d_{ni})^{-\theta}$). However, if the good was invented abroad, and producers in other countries have accumulated ideas for producing it, the probability of being the low cost provider of the good to distant markets is relatively low, but high trade costs increase the probability of serving the home and nearby markets. So lowering trade costs increases demand for newly invented goods abroad, while making it less likely that a new idea for producing a good invented abroad can compete with that of the technological leader, so research in all countries shifts from learning to invention. However, poor countries still have the advantage of relatively low wages, so they are not as reliant on trade barriers in order to be competitive with producers in the inventing country. As a result, the effect is dampened for these countries. Instead, the additional researchers are drawn from process development.

Interestingly, welfare is actually slightly lower in all countries in the new steady state than levels in the static model after the change in trade costs, with the average country's real wage about 1.5% less. This result is due to the properties of the CES price index. The elasticity of P_n with respect to a uniform change in the level of research that has accumulated to each good is $\frac{-1}{\theta}$, while its elasticity with respect to a change in the size of the product space is $\frac{-1}{\sigma-1}$. Since $\theta > \sigma - 1$, this implies that, all else equal, consumers would prefer that a given unit of research output be devoted to increasing the technology associated with existing goods than inventing new ones. Since lowering trade costs induces

Table 7: Change in # of Researchers; Autarky

GDP/Worker	s_i^d	s_i^f	s_i^n
High	0.7%	86.8%	-63.4%
Low	0.5%	-0.8%	-94.0%

a shift in research toward invention, welfare is lower after this adjustment.

6.2 Autarky

The move to zero trade, predictably, yields essentially the opposite result. In the short run, the real wage of the average country falls by about 15%, with smaller countries being affected more than larger ones. In the long run, researchers in rich countries flood out of invention into learning, using the only avenue now available of obtaining products invented abroad. Poor countries essentially abandon what tiny amount of invention they were doing in favor of a slight increase in domestic process research. The intuition for this result is very similar to that above. Now that goods cannot be exported, and producers in rich countries do not have to compete with low wage producers from abroad, there is much less incentive to invent new goods. On the other hand, it is now relatively much easier to become the low cost producer domestically of a product invented abroad, since it cannot be purchased from the country with the technological advantage.

Also conversely to the previous case, total welfare increases slightly from the short run to steady state equilibrium, mitigating some of the static welfare loss. In fact, in four countries – the United States, Japan, Australia, and Brazil; large countries that are fairly geographically isolated already – the real wage is actually higher than in the baseline calibration but by only 1% on average. This is a result of the fact that these countries were already so close to autarky in the calibration that the dynamic gain of shifting research to improving the techniques for producing existing goods outweighs the loss due to restricting trade. In addition, except for the case of Brazil, these countries possess some of the world's most productive researchers who, when their efforts are focused on learning about all other countries products, effectively mitigate the lost gains from trade by matching or surpassing the levels of technology with which these goods are produced elsewhere. It should also be noted that, were prohibitive barriers to trade correlated with impediments on the ability to learn about goods invented in other countries, this result would vanish.

7 Conclusion

Poor countries' exports are concentrated in a common small set of products, while rich countries' exports are mostly evenly divided among all products, even those for which they have strong comparative disadvantage. A model of endogenous growth in which researchers can choose between inventing a new product and developing a more efficient process for producing an existing product, when calibrated to match salient features of the data, predicts that researchers in rich countries spend more effort inventing new goods in order to avoid direct competition with low-wage producers in poor countries. It also predicts that researchers in poor countries devote relatively more effort learning about products invented in rich countries and developing ideas for producing products about which they have already learned because their low wage implies that they have a high probability of becoming the low cost producer of these goods.

As a result, rich countries have a large comparative advantage in producing newer goods. As a good ages, researchers in poor countries learn enough to develop their own processes for producing it and catch up to rich countries in the ability to produce these products. But, the large technological advantage of rich countries erodes slowly, so they still export some products for which poor countries have an overall comparative advantage. This implies that poor countries' exports are highly concentrated in these older products, while rich countries export newer products almost exclusively while also exporting many older products as well.

Thus, the model explains the phenomenon that we observe in the product level data and quantitatively matches the distribution of rich and middle income countries' exports across the product space quite well. However, it fails to account for the full degree of the concentration of poor countries' exports. This is primarily due to the form of the learning process. The assumption that a researcher is equally likely to acquire any piece of product knowledge that he does not already possess (in the spirit of Krugman's (1979b) constant rate of diffusion of new goods from the North to the South) implies that the amount of product knowledge available in a country for a product invented elsewhere as a function of the product's age takes the form of an exponential CDF. For newer goods, the pattern of comparative advantage (as seen in figure 10) inherits the shape of this distribution. A learning process that implies a hazard rate that is increasing in the age of the good or the amount of knowledge already available in a country pertaining to the good would better capture the observed pattern of comparative advantage. However, introducing a more complex learning process would make the model much less analytically tractable. Introducing a correlation the age of a good and factor intensity could also amplify these results, as could the introduction of non-homothetic preferences. These are potential avenues of future research.

The calibrated model predicts that a decrease in trade barriers results in larger welfare gains than is predicted by a special case of the model that reduces to that of Eaton & Kortum (2002), with the difference greater for poor countries. This is because, in this model, the exports of poor countries are concentrated in products for which all countries are very productive, implying that poor country's exports are more responsive to changes in trade costs. The model further predicts that lower trade barriers lead to a shift in research activity toward invention, as a major incentive to invent is the ability to export it to a large world market, while trade barriers make it easier to become the low cost provider in the researcher's domestic market for a good that was invented in a distant country.

The model also does quite well in explaining a puzzle in the international trade literature – that theoretical gravity models cannot simultaneously explain the level of trade as a percentage of output for both rich and poor countries. It does this in an intuitive way. Since rich countries account for most of the world's invention, a large portion of their trade is variety driven, similar to trade models based on increasing returns (e.g. Krugman (1979a) and Melitz (2003)). Poor countries, on the other hand, are very good at producing the same, older products, so there is little scope for gains from comparative advantage among them. This implies that poor countries trade much less than would be predicted by a model where countries are equally productive, on average, for all goods, as most gravity models predict. In this sense, the mechanism of the model is similar to Feiler (2007); however, it avoids the criticism of Waugh (2009) that the elasticity of trade with respect to trade costs is not higher for poor countries in the data, a phenomenon discussed in the appendix.

References

[1] Alvarez, F. and R. Lucas. (2006). "General Equilibrium Analysis of the Eaton-Kortum Model of International Trade," Working Paper.

[2] Anderson, J. and E. van Wincoop. (2003). "Gravity with Gravitas: A Solution to the Border Puzzle," *The American Economics Review*, Vol. 93, No. 1: 170-192

[3] Armenter, R. and M. Koren. (2008). "A Balls-and-Bins Model of Trade," Working Paper.

[4] Balassa, B. (1965). "Trade Liberalization and Revealed Comparative Advantage," *The Manchester School of Economic and Social Studies*, Vol. 33: 99-123

[5] Bernard, A., J. Eaton, J.B. Jenson, and S. Kortum. (2003). "Plants and Productivity in International Trade," *The American Economic Review*, Vol. 93, No. 4.

[6] Broda, C. and D. Weinstein. (2006). "Globalization and the Gains from Variety," *Quarterly Journal of Economics*, Vol. 121, No. 2: 541-85.

[7] Comin, D, and B. Hobijn. (forthcoming) "An Exploration of Technology Diffusion," *The American Economic Review*.

[8] Costinot, A. and I. Komunjer. (2008). "What Goods Do Countries Trade?: A Structural Ricardian Model," Working Paper.

[9] Eaton, J. and S. Kortum (2001). "Technology, Trade, and Growth: A Unified Framework," *European Economic Review*, Vol. 45, No. 4-6: 742-755.

[10] Eaton, J. and S. Kortum (2002). "Technology, Geography, and Trade," *Econometrica*, Vol. 70, No. 5: 1741-1799.

[11] Eaton, J. and S. Kortum (2006). "Innovation, Diffusion, and Trade," NBER Working Paper 12385.

[12] Feenstra, R., R. Lipsey, H. Deng, A. Ma, and H. Mo. (2004). "World Trade Flows: 1962-2000," NBER Working Paper no. 11040.

[13] Feiler, A. (2007) "Non-Homotheticity and Bilateral Trade: Evidence and a Quantitative Explanation," Working Paper.

[14] French, S. (2009) "Export Composition and the Failure of Gravity for Deloping Countries," Working Paper.

[15] Gaulier, et al, (2008). "BACI: A World Database of International Trade Analysis at the Product-level," CEPII Working Paper.

[16] Grossman, G. and E. Helpman. (1991). *Innovation and Growth in the Global Economy.* The MIT Press.

[17] Helpman, E., M. Melitz, and Y. Rubenstein. (2008) "Estimating Trade Flows: Trading Partners and Trading Volumes," *Quarterly Journal of Economics* Vol. 123: 441-487

[18] Hummels, D. and P. Klenow. (2005). "The Variety and Quality of a Nation's Exports," *American Economic Review* Vol. 95: 704-723.

[19] Jones, C. (1995). "R&D-Based Models of Economic Growth", *The Journal of Political Economy* Vol. 103, No. 4: 759-784.

[20] Kortum, S. (1997). "Research, Patenting, and Technological Change," *Econometrica*, Vol. 65, No. 6: 1389-1419.

[21] Krugman, P. (1979a). "Increasing Returns, Monopolistic Competition, and International Trade," *Journal of International Economics* Vol. 9, No. 4: 469-479.

[22] Krugman, P. (1979b). "A Model of Innovation, Technology Transfer, and the World Distribution of Income", *The Journal of Political Economy* Vol. 87, No. 2: 253-266.

[23] Mansfield, E., M. Schwarts, and S. Wagner. (1981). "Imitation Costs and Patents: An Empirical Study," *The Economic Journal*, Vol. 91, No. 364: 907-918.

[24] Mansfield, E. and A. Romeo. (1980). "Technology Transfers to Overseas Subsidiaries by U.S.-Based Firms," *The Quarterly Journal of Economics*, Vol. 95, No. 4: 737-750.

[25] Melitz, M. (2003). "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica*, Vol. 71, No. 6: 1695-1725.

[26] Nelson, R. (1982). "The Role of Knowledge in R&D Efficiency", *The Quarterly Journal of Economics*, Vol. 97, No. 3: 453-470.

[27] Romer, P. (1990). "Endogenous Technical Change," *The Journal of Political Economy*, Vol. 98, No. 5: S71-102.

[28] Schott, P. (2002). "Moving Up and Moving Out: US Product-Level Exports and Competition from Low Wage Countries", Working Paper.

[29] Waugh, M. (2009) "International Trade and Income Differences," *Federal Reserve Bank of Minneapolis*, Research Department Staff Report 435.

8 Appendix

8.1 Products over Time

The data indicate two more phenomena which will be useful in explaining why poor countries's exports are concentrated in a small range of products. First, rich countries appear to have an advantage in exporting newer products. And, second, poor countries catch up to rich countries in the degree to which they export these products as they age.

There is no data that directly measures the age of a product, but is there is indirect evidence. United States import data is collected according to the 10-digit Harmonized Tariff Schedule (HTS), which is based on the 6-digit Harmonized System. While the 6-digit (HS) codes are updated every five years or so, the further disaggregated HTS codes are updated nearly continuously by the U.S. International Trade Commission. One reason for making changes to the codes is to accommodate the changing composition of products entering the United States. Using the concordance of Pierce and Schott (2009), I am able to follow changes in the HTS codes over, and I take codes that have been split in two or more new codes over time as evidence of their containing newer goods. The rationale for this is as follows. When a brand new type of good arrives on the US shore, it must be assigned to a HTS code. At first, a statistician assigns the good to a code that contains other types of goods with some similar characteristics. Over time, as imports of that good into the US grows, it may eventually be given its own category, either because of the goal of the USITC to keep codes relevant to changing trade patterns or because the producer is not happy with its classification for tariff or marketing purposes and petitions for the change. On the other hand, a good that has been imported into the U.S. for many years and undergoes no changes is not likely to be a candidate to be split out into a new code. As a result, codes that have been split over time are more likely to contain newer products.

Figure 3 shows the percentage of a country's exports to the U.S. in split categories relative to those categories' percentage of total U.S. imports. A value greater than 1 indicates that the split categories make up a larger proportion of that country's exports than they do overall U.S. importers. Developed countries export relatively more in these categories, indicating that they have a comparative advantage in producing newer goods.

Figures 4 and 5 illustrate the second point. Figure 4 shows that the groups of products exported more intensively by rich countries are growing faster over time. And, figure 5 shows that this growth is mostly due to the poor countries exporting relatively more of

Figure 12: Relative Exports in "Split" Categories

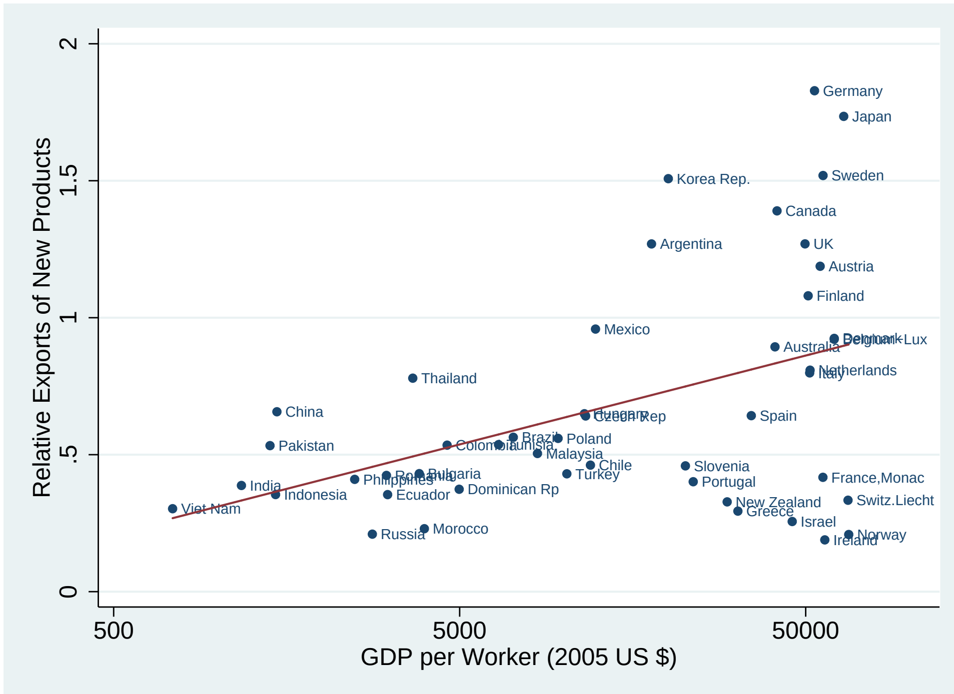

Figure 13: Growth of Trade by Product Group

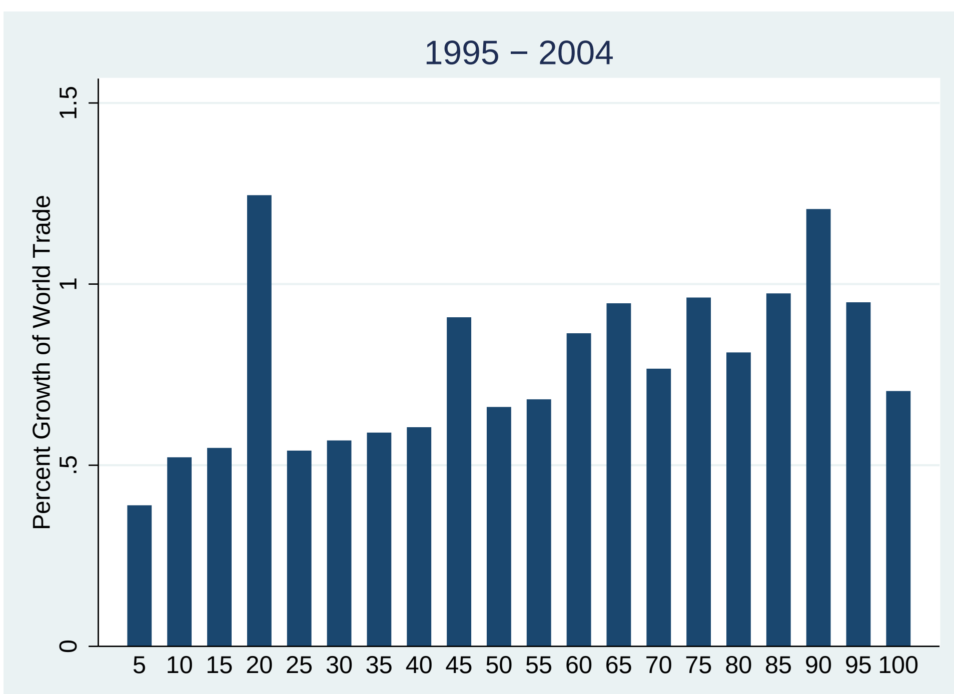

Figure 14: Growth of Exports by Product Group

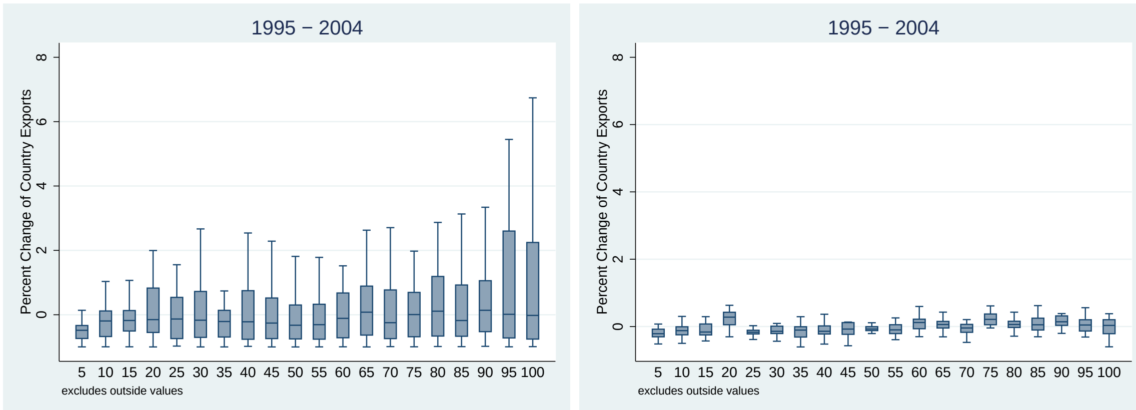

these products. This is consistent with the story of these product categories containing newer goods which the rich countries have a comparative advantage at producing, but over time, as these goods age, the poorer countries catch up. That the growth of trade of richer countries is quite evenly spread across goods may indicate that new products are being allocated fairly evenly across categories, to be split out later when the classification is revised.

8.2 Math

8.2.1 Technological Frontier

An idea is pair of a quality level, Q, and a good to which it applies, j, where Q is drawn from the Pareto distribution:

$$Pr(Q < q) = 1 - \left(\underline{q}/q\right)^{\theta}$$

where $q \in (\underline{q}, \infty)$. Thus, $\underline{q}$ is the lower bound on Q, and θ is the dispersion parameter of the Pareto distribution. Now suppose that the number of ideas applying to a particular good, j, in a given country i, K_i^j is drawn from a Poisson distribution with parameter $\bar{a}T_i^j$:

$$Pr(K_i^j = k) = \frac{e^{-\bar{a}T_i^j}(\bar{a}T_i^j)^k}{k!}$$

Intuitively, one can think of T_i^j as the level of research that has been applied to drawing ideas for producing good j in country i and $\bar{a}$ as a measure of effectiveness of research in producing ideas. Denote the best idea for producing good j in country i by $Z_i^j = \max_k\{Q_i^{jk}\}$. Given that K_i^j ideas exist for producing j, Z_i^j is distributed as follows.

$$Pr(Z_i^j < z | K_i^j = k) = \bigcap_{m=1}^{k} Pr(Q_i < z) = \prod_{m=1}^{k} 1 - (\underline{q}/z)^\theta = \left(1 - (\underline{q}/z)^\theta\right)^k$$

Summing over all possible values of K_i^j

$$\begin{aligned} F_i^j(z) = Pr(Z_i^j < z) &= \sum_{k=0}^{\infty} \frac{e^{-\bar{a}T_i^j}(\bar{a}T_i^j)^k}{k!} \left(1 - (\underline{q}/z)^\theta\right)^k \\ &= e^{-\bar{a}T_i^j} \sum_{k=0}^{\infty} \frac{\bar{a}T_i^j \left(1 - (\underline{q}/z)^\theta\right)}{k!} \\ &= e^{-\bar{a}T_i^j} e^{\bar{a}T_i^j\left(1-(\underline{q}/z)^\theta\right)} \\ &= e^{-(\bar{a}\underline{q}^\theta)T_i^j z^{-\theta}} \end{aligned}$$

Assuming that if no idea has arrived pertaining to good j – i.e. $K_i^j = 0$ – it can be produced with efficiency $\bar{q}$, then there is a mass point at $Z = \bar{q}$ where $Pr(Z = \bar{q}) = e^{-\bar{a}T_i^j}$. Normalizing $\bar{a}\underline{q}^\theta = 1$ and letting $\underline{q} \to 0$ (therefore, $\bar{a} \to \infty$) allows $F_i^j(z)$ to be defined on $z \in [0, \infty]$ with $F_i^j(0) = 0$.[26] This allows the annoyance of a discontinuity in the distribution owing to the possibility of no ideas arriving – for very small values of T_i^j – to disappear while preserving the properties of the rest of the distribution.

8.2.2 Price Index

Available from the author upon request.

8.3 Is θ the Trade Cost Elasticity?

This model is similar to Feiler (2007) in that it implies that the elasticity of trade with respect to trade costs is higher for poor countries than for rich countries. However, this implication is the result of very different mechanisms. Because of this, the *estimated* elasticity of trade with respect to trade costs is the same across countries in this model as opposed to that of Feiler. This is important because Waugh (2009) finds that the implied relationship between this elasticity and income per worker is not present in the data. This appendix explains this result.

Recall from the text that the expenditure by country n on product j from country i is

$$E[X_{ni}^j] = X_n \frac{T_i^j (w_i d_{ni})^{-\theta}}{\Phi_n} \left(\frac{\Phi_n}{\Phi_n^j}\right)^{\frac{\theta-(\sigma-1)}{\theta}}$$

[26] Intuitively, this exercise assumes that any amount of research is infinitely productive at producing ideas but that ideas are almost surely of zero quality. However, we are only concerned with good ideas, so the addition of a many very bad ideas is of virtually no consequence to the distribution of the best idea.

It is easy to show that the elasticity of this expenditure with respect to the trade cost is the following.

$$\frac{\partial \ln(X_{ni}^{j})}{\partial \ln(d_{ni})} = (\sigma - 1)\pi_{ni}^{j} + \theta(1 - \pi_{ni}^{j})$$

The intuition for this result is as follows. If a given exporter has a very large comparative advantage over the rest of the world for a product, then a marginal change in the trade cost does not have a large effect on the probability that it will be the low cost provider of the good, which is governed by the dispersion idiosyncratic productivity, θ. What the change in the trade cost affects, then, is the expected price at which the importer will be able to purchase the good, which affects demand according to the elasticity of substitution, $(\sigma - 1)$. Conversely, if an exporter has a small probability of being the low cost provider of the good, a marginal change in the bilateral trade cost will have little impact on the expected price of the good but a significant relative affect on the probability of the exporter being the low cost provider.

In the model of this paper, in equilibrium, rich countries account for the vast majority of the invention of new goods, meaning that a large portion of their exports are in goods for which π_{ni}^{j} approaches 1. Poor countries' exports, however, are concentrated in the set of older goods, for which all countries have devoted a substantial amount of production research, implying that π_{ni}^{j} will be substantially less than 1 for most of poor countries' exports. In this way, this model makes a similar prediction as Feiler (2007), which predicts that rich countries specialize in products with a lower trade cost elasticity, while poor countries specialize in those with higher trade cost elasticity, a key component of the model that allows her to reconcile the trade model of Eaton and Kortum (2002) with the observation that poor countries do not trade as much as rich countries.

The major difference between Feiler (2007) and this model is that Feiler assumes differences in the dispersion of idiosyncratic productivity differences across products, while this model allows endogenous differences in the levels of average productivity across products. This difference is quite important when it comes to the implications of these models for gravity-type estimations.

Consider, for example, the following measure of trade for a given product normalized by the home trade share of the importing country.

$$\frac{X_{ni}^{j}}{X_{nn}^{j}} = \frac{T_{i}^{j}}{T_{n}^{j}} \left(\frac{w_i}{w_n}\right)^{-\theta} (d_{ni})^{-\theta}$$

Taking logs and grouping country specific terms together gives

$$\ln\left(\frac{X_{ni}^j}{X_{nn}^j}\right) = S_i^j - S_n^j - \theta\ln(d_{ni})$$

A common method for estimating such an equation (e.g. Eaton and Kortum (2002)) is to regress the left-hand side value calculated from the data on a set of importer and exporter dummy variables and typical proxies for trade costs from the gravity literature. It is easy to see, then, that if an appropriate trade cost estimate were available, the estimated elasticity would be θ. The reason for this is that the data used is typically a cross-section of countries, so identification is over a change in trade costs from country to country, not an actual change in trade costs for a given exporter-importer pair. This is important because, from the perspective of the importing country, there is no aggregate change to induce a change in the price or quantity demanded of good j, so the effect on the trade cost elasticity due to the elasticity of substitution across products in demand is missing.

In the case of Feiler (2007), however, the expression for the same value would be

$$\frac{X_{ni}^j}{X_{nn}^j} = \frac{T_i}{T_n}\left(\frac{w_i}{w_n}\right)^{-\theta}(d_{ni})^{-\theta^j}$$

where now T_i and T_n do not depend on j, but θ^j does. As a result, taking logs and grouping country terms as above gives

$$\ln\left(\frac{X_{ni}^j}{X_{nn}^j}\right) = S_i - S_n - \theta^j\ln(d_{ni})$$

Since the model of Feiler predicts that the goods most traded by rich countries are those for which θ is lowest, the model would also predict that the estimated value of θ for rich countries would be lower than the estimated value for poor countries, an implication contradicted by the empirical work of Waugh (2009).

8.4 Examples of Products

Table A.1

Major Exports of Poor Countries	
6871	Unwrought Tin
6592	Knotted Carpets and Rugs
7511	Typewriters & Check-writing Machines
8423	Trousers of Textile Fabric
7621	Radio Broadcast Receivers for Automobiles
8811	Photographic Cameras, Parts & Accessories
Major Exports of Rich Countries	
8996	Orthopedic Appliances
7764	Electronic Microcircuits
7810	Passenger Motor Cars for Transport
7523	Complete Digital Central Processing Units
7643	Radiotelegraphic & Radiotelephonic Transmitters
7415	Self-contained Air Conditioning Machines

8.5 Tables

Table 8: Exogenous and Endogenous Variables

Country	Wage	Labor Force	$(\frac{\alpha_i}{\alpha_{US}})^{\frac{1}{\theta}}$	s_d	s_f	s_n	κ_i	$(\frac{\delta_i}{\delta_{US}})^{\frac{1}{\theta}}$	η_i
Norway	35,458	290,492	1:38	2:13	0:05	0.17	0.22	0.61	3.83
Germany	35,079	7,472,387	1.08	2.13	0.07	0.12	0.23	0.90	24.17
USA	34,040	16,711,080	1.00	2.12	0.08	0.10	0.23	1.00	32.80
Denmark	33,526	466,469	1.24	2.13	0.05	0.16	0.22	0.59	3.21
Netherlands	33,020	829,334	1.17	2.13	0.06	0.15	0.23	0.62	4.27
Belgium-Lux	30,283	654,015	1.03	2.13	0.06	0.13	0.23	0.53	1.64
Austria	30,098	606,930	1.05	2.13	0.06	0.13	0.23	0.54	1.81
Switz.Liecht	29,478	673,840	0.99	2.13	0.06	0.12	0.23	0.52	1.43
Finland	28,519	419,515	1.08	2.13	0.08	0.12	0.23	0.53	1.55
UK	28,209	4,088,713	0.90	2.14	0.09	0.09	0.23	0.67	5.36
Israel	27,748	333,600	1.07	2.13	0.09	0.12	0.23	0.51	1.24
Japan	27,447	9,282,688	0.83	2.14	0.12	0.06	0.24	0.72	7.16
Sweden	27,276	767,092	0.98	2.13	0.08	0.10	0.23	0.54	1.67
Canada	27,163	1,888,719	0.92	2.14	0.10	0.08	0.23	0.60	2.95
France,Monac	26,686	3,887,712	0.84	2.14	0.10	0.08	0.23	0.63	3.65
Ireland	25,041	248,740	0.89	2.13	0.08	0.11	0.23	0.40	0.30
Australia	24,680	1,038,754	0.88	2.15	0.14	0.05	0.24	0.53	1.21
Italy	21,041	3,975,028	0.65	2.15	0.12	0.05	0.25	0.51	0.93
Spain	19,618	2,321,883	0.63	2.15	0.12	0.04	0.25	0.45	0.45
New Zealand	18,248	212,705	0.72	2.15	0.13	0.04	0.25	0.34	0.08
Untd Arab Em	16,249	209,335	0.62	2.15	0.13	0.04	0.25	0.30	0.00*
Korea Rep.	15,199	2,432,258	0.48	2.17	0.15	0.01	0.27	0.35	0.08
Greece	13,941	423,238	0.49	2.16	0.13	0.03	0.26	0.27	0.00*
Malaysia	12,703	338,885	0.46	2.17	0.14	0.01	0.27	0.25	0.00*
Slovenia	12,197	141,631	0.41	2.16	0.13	0.02	0.27	0.18	0.00*
Argentina	12,007	820,835	0.40	2.18	0.14	0.01	0.29	0.25	0.00*
Kuwait	10,945	69,024	0.41	2.16	0.13	0.02	0.27	0.16	0.00*
Portugal	8,664	948,313	0.27	2.19	0.13	0.00	0.31	0.17	0.00*
Chile	8,578	557,692	0.29	2.19	0.14	0.00	0.32	0.17	0.00*
Venezuela	7,767	835,114	0.25	2.20	0.13	0.00	0.32	0.15	0.00*
Qatar	7,501	39,760	0.27	2.19	0.13	0.00	0.31	0.10	0.00*
Brazil	6,785	4,812,166	0.19	2.22	0.13	0.00	0.37	0.16	0.00*
Turkey	6,216	1,108,831	0.19	2.21	0.12	0.00	0.36	0.12	0.00*
Mexico	5,970	1,454,707	0.18	2.21	0.12	0.00	0.36	0.12	0.00*
Iran	5,636	874,631	0.17	2.21	0.12	0.00	0.37	0.10	0.00*
Oman	5,581	32,958	0.20	2.20	0.12	0.00	0.34	0.07	0.00*
Tunisia	5,546	194,914	0.18	2.20	0.12	0.00	0.35	0.08	0.00*
Saudi Arabia	5,245	589,169	0.16	2.21	0.12	0.00	0.37	0.09	0.00*
Czech Rep	4,320	1,254,730	0.11	2.21	0.10	0.00	0.41	0.07	0.00*
Hungary	4,265	740,886	0.12	2.21	0.10	0.00	0.40	0.07	0.00*
Morocco	4,110	600,644	0.12	2.22	0.11	0.00	0.40	0.06	0.00*
Colombia	4,028	458,735	0.12	2.22	0.11	0.00	0.40	0.06	0.00*

* Value is less than 0.005.

Table 9: Exogenous and Endogenous Variables (cont.)

Country	Wage	Labor Force	$(\frac{\alpha_i}{\alpha_{US}})^{\frac{1}{\theta}}$	s_d	s_f	s_n	κ_i	$(\frac{\delta_i}{\delta_{US}})^{\frac{1}{\theta}}$	η_i
Poland	3,819	2,429,383	0:10	2:22	0:10	0.00	0.42	0.07	0.00*
Algeria	3,746	494,055	0.11	2.22	0.10	0.00	0.41	0.05	0.00*
Dominican Rp	3,573	251,197	0.11	2.22	0.10	0.00	0.40	0.05	0.00*
Slovakia	3,354	392,301	0.09	2.22	0.10	0.00	0.43	0.04	0.00*
Thailand	3,067	2,437,170	0.08	2.24	0.10	0.00	0.45	0.06	0.00*
Peru	3,011	650,828	0.09	2.23	0.10	0.00	0.44	0.05	0.00*
Ecuador	2,726	118,321	0.08	2.23	0.10	0.00	0.43	0.03	0.00*
Philippines	2,526	1,084,200	0.07	2.24	0.10	0.00	0.46	0.04	0.00*
Pakistan	2,126	1,630,607	0.05	2.24	0.09	0.00	0.48	0.03	0.00*
Russia	1,395	10,460,350	0.03	2.25	0.08	0.00	0.54	0.02	0.00*
China	1,328	30,461,310	0.02	2.26	0.08	0.00	0.55	0.02	0.00*
Romania	1,283	1,659,900	0.03	2.24	0.08	0.00	0.54	0.02	0.00*
India	1,270	7,903,693	0.02	2.25	0.08	0.00	0.55	0.02	0.00*
Bulgaria	1,245	615,829	0.03	2.24	0.08	0.00	0.54	0.01	0.00*
Indonesia	1,094	4,216,865	0.02	2.26	0.08	0.00	0.56	0.02	0.00*
Nigeria	1,008	1,498,020	0.02	2.25	0.08	0.00	0.55	0.01	0.00*
Angola	990	119,448	0.02	2.24	0.08	0.00	0.53	0.01	0.00*
Viet Nam	880	1,292,240	0.02	2.25	0.08	0.00	0.56	0.01	0.00*