Distribution Costs, Product Quality, and Cross-Country Income Differences

Bernardo S Blum   Sebastian Claro   Kunal Dasgupta
Ignatius Horstmann

Preliminary Draft – March 2016

Abstract

We show that the efficiency of countries' distribution systems helps determine the quality of goods produced and traded, i.e., is a source of comparative advantage in quality. Using the structure of our model and shipment-level imports data from Chile, we estimate the efficiency of trade distribution systems for a sample of 86 countries. We find that the implied efficiency of distribution systems vary widely across countries, with the 90th percentile value of per-shipment costs being almost 150 percent larger than the 10th percentile value. After calibrating the parameters of the model, we show that differences in the efficiency of distribution systems can generate more than half of the observed elasticity of export prices with respect to per capita income. Moreover, the welfare effects of reducing inefficiencies in distribution systems via quality upgrading are larger than the effects via trade volumes.

KEYWORDS: Distribution cost, per-shipment cost, quality, comparative advantage

JEL code: F10, F12

1 Introduction

Distribution costs, defined broadly as the costs of moving goods from the point of production to the point of consumption, help determine the organization of economic activity both within and across countries. Much is known about the ways that transportation networks – road, rail, ports – and their associated costs affect the extent of specialization in production across regions and so the extent and volume of trade (see Anderson and van Wincoop [2004], Hummels [2007], and Redding and Turner [2014] for surveys). Much less well understood is how, not only transport networks, but other aspects of the distribution system – inventory, logistics and distribution management costs – affect the characteristics of the goods produced in any location and the patterns of trade.[1] While it goes without saying that goods that are time sensitive – perishable goods, seasonal goods, fashion goods – will be difficult to produce in and export from countries with poor internal transport networks, long delays at ports and high shipping consolidation costs, little in the way of systematic analysis has been done on this problem.

In this paper, we tackle this problem head on, examining how the efficiency of countries' distribution systems helps determine, not just the volumes and directions of trade, but the quality of goods produced and traded as well. We show that the efficiency of a country's distribution system can serve as a source of comparative advantage in quality: countries with very efficient distribution systems have a comparative advantage in producing high quality goods relative to those with inefficient systems. Moreover, we show that, while reductions in distribution costs have a direct welfare benefit, due to the usual reasons, this benefit is small relative to the indirect benefit that arises due to the quality upgrading that the enhanced distribution system facilitates.

Our point of departure for this analysis is evidence from transaction-level data on all imports into Chile. Not surprisingly, we find in the data that the size and frequency of import shipments vary considerably across HS 8-digit products. What is more surprising is that, within a given HS 8-digit product category, the size and frequency of shipments vary in a systematic way with the per-capita income of the exporting country. In particular, controlling for total export volume, countries with higher per-capita income (rich countries) have smaller shipment sizes and more frequent shipments than do countries with lower per-capita income (poor countries). Indeed, we establish this relationship even when the same Chilean importer buys an HS 8-digit product from rich and poor countries. Moreover, this importer pays higher prices for the product imported from the rich country than the poor country: the level of a country's development, the quality / price of

[1] One exception is the "shipping the good apples out" mechanism proposed by Alchian and Allen [1969] and taken to trade data by Hummels and Skiba [2004]. In contrast to this mechanism, in our paper distribution systems will be a source of comparative advantage and will affect the the quality of goods countries produce.

exports and size and frequency of transactions appear to be linked.

To uncover how (or if) this link can arise, we build a model of international distribution that we imbed in a full, general equilibrium model of multi-country trade. Our model focuses on three features of international product distribution. First, we observe in the data that i) importers make infrequent shipments [Alessandria et al., 2010] and ii) the size and frequency of these infrequent shipments vary across exporting countries, even within importer and traded product (see this paper). The former suggests that a key feature of distribution systems is a *fixed per-shipment cost* associated with exporting, while the latter suggests that at least some part of the per-shipment cost is export-country specific. These (country-specific) per-shipment costs can represent a variety of (quantity independent) distribution costs, ranging from the costs of administrative activities associated with exporting, to delay costs at ports, to the costs of shipment consolidation and the like. These costs represent one kind of distribution cost. Next, we observe that trade is often not carried out in a single, large shipment occurring once a quarter, say [Alessandria et al., 2010]. This suggests that another key feature of distribution systems are *inventory costs* borne by the importer that are increasing in the size of the shipment. These costs can range from the costs of unsold inventory, to carrying costs, to costs of maintaining the product while in storage. These costs are likely product and importing country specific and represent a second kind of distribution cost. Finally, distribution also results in *transportation costs*, which we assume are of the usual iceberg variety.

Confronted with these three sorts of distribution costs, an exporter from any given country, wishing to sell a given amount of a product in some other country, chooses the size and frequency of export shipments in order to minimize total distribution costs. Distribution cost minimization results in exporters from countries with high per-shipment costs bearing the increased inventory costs that large shipments incur in order to reduce the number of shipments and so save on per-shipment costs. The opposite occurs for exporters from low per-shipment cost countries. As a result, for a given total volume of exports of a given product, exporters from countries with high per-shipment cost make large, infrequent shipments while exporters from low per-shipment cost countries make small, frequent shipments.

Under this optimal distribution policy – the cost minimizing shipment size and frequency – we show that the combined per-shipment, inventory management and transport costs yield an equivalent ad valorem distribution cost that is export-import country pair and product specific. An important feature of this cost is that countries with lower per-shipment cost, all else equal, have a lower ad valorem distribution cost. When embedded in a trade model in which exporters choose both product price and quality, exporters from countries with low per-shipment costs have a lower cost of increasing quality than do exporters from countries with high per-shipment costs. As a result, we expect to observe that countries with low per-shipment costs produce higher quality products and export these products in small but frequent shipments. Ultimately, more efficient distribution systems at exporting countries give countries a comparative advantage in producing high quality goods.

Although not directly observable, we are able, nevertheless, to estimate the efficiency of distribution systems, specifically per-shipment costs, at exporting countries for a sample of 86 countries using the model structure and our shipment level imports data. The identification of this efficiency parameter is obtained from variation, across export countries, in the frequency and size of shipments to Chile. We find that the implied efficiency of distribution systems at exporting countries varies widely across countries, with the 90th percentile value of per-shipment costs being almost 150 percent larger than the 10th percentile value. Moreover, we find that distribution efficiency is closely related to per capita incomes, even though our estimation procedure does not impose this relationship.

After calibrating the remaining parameters of the model, we show that differences in the efficiency of distribution sectors can generate more than half of the observed (in the data) elasticity of export prices with respect to per capita income, and a similar response of export quality with respect to per capita income. The model also generates between twenty and forty percent of the income inequality between countries observed in data, depending on how one measures inequality.

With the calibrated model in hand, we can also quantify the effects of hypothetical changes in the efficiency of countries' distribution systems. For instance, we study the implications of a policy that makes the efficiency of the distribution system in all countries equal to the efficiency of the US distribution system. We find that this reduces the elasticity of price with respect to per capita income by more than half, and reduces income inequality across countries by more than 20 percent. The median welfare gains for this exercise is 2.7 percent, with the poor countries gaining much more than the rich. More importantly, most of these welfare gains are a consequence of the quality upgrading that takes place following the lowering of distribution costs.

Our results are related to a few papers in different literatures. Hummels and Klenow [2005] estimate that quality differences could be the proximate cause of around 9 percent of country differences in real income per worker. We show that quality differences across countries are, in part, due to differences in the efficiency of distribution systems. Waugh [2010] shows that one requires asymmetric, ad-valorem trade costs to explain trade flow data and that these asymmetric costs account for a sizeable part of observed income differences across countries. Our model provides a micro-founded mechanism to Waugh's story and the evidence to support it. We confirm that the microeconomic inefficiency captured in Waugh's work and made explicit in our work leads to large macroeconomic effects, in particular on countries' incomes.

Our paper also contributes to the literature on comparative advantage in quality, identifying a new source of comparative advantage: the efficiency of a country's distribution system. Our explanation stands in contrast to the existing literature that argues that comparative advantage in quality is the result of: i) quality being intensive in physical and human capital, factors that are abundant in rich countries [Markusen, 1986, Bergstrand, 1990]; ii) economies of scale in the production of quality, giving rich countries an advantage [Linder, 1961, Fajgelbaum et al., 2011]; iii) high quality being complementary with high productivity [Flam and Helpman, 1987, Matsuyama, 2000, Baldwin and Harrigan, 2011, Johnson, 2012, Feenstra and Romalis, 2014].

Our paper has implications for the literature that attempts to disentangle unobserved product quality from prices [Khandelwal, 2010, Hallak and Schott, 2011, Feenstra and Romalis, 2014]. In particular, it highlights that price reflects the efficiency of distribution systems as well, at the same time that the efficiency of distribution systems help determine quality. We add, as well, to research that documents systematic patterns in the size and frequency of shipments from the perspective of exporters [Kropf and Sauré, 2014, Hornok and Koren, 2015]. Our work is fundamentally different than these papers in that, by looking at imports data from multiple export countries, we identify differences in the efficiency of distribution systems *located at exporting countries* and show how these differences could be a source of comparative advantage in quality. Alessandria et al. [2010] has a distribution system with features similar to ours, which the authors show to have a sizeable impact on trade volumes and use to study the dynamic effects of large devaluation episodes. In contrast, we show that the efficiency of distribution systems vary across exporting countries and how this has implications for production, income and inequality across countries.

Finally, our paper adds to the literature that documents the different components of distribution systems and measure their impacts on trade. Martincus et al. [2015], for instance, uses Uruguayan data to document the time it takes for export shipments to be processed. Blonigen and Wilson [2008] and Feenstra and Ma [2014] study the effects of port efficiency on trade flows and on the extensive margin of trade. The World Bank Doing Business Indicators provide survey metrics for the time, number of documents, and dollar cost of exporting a container for a large sample of countries. In contrast, we use the structure of our model to infer the total distribution costs incurred at exporting countries that would be consistent with the size and frequency of shipments observed in trade data.

The remainder of the paper is organized as follows. The next two sections present the data and the evidence on shipment sizes, frequency and prices. Sections 4 and 5 develop a full, general equilibrium model of trade and distribution. This section derives the mechanism linking the choices of product quality and shipments' frequency and size. In Section 6, we calibrate the model and quantify the extent to which the mechanism proposed by the model can generate the export price variation observed in data. We also quantify the implications of the mechanism in our model to country differences in per capita incomes. Lastly, we run a number of counterfactual exercises to quantify the impacts of changes in the efficiency of distribution systems in exporting countries.

2 Data Description

This paper uses data on the universe of Chilean import shipments in 2006. Collected by the Chilean Customs Office, the data contain detailed shipment information, such as the product traded, prices and quantities, country of origin, and an identifier for the importer (buyer) and the exporter (seller) of the shipment. In our analysis, we exclude import transactions in mineral fuels (HS27),[2] and shipments in which the importer is an individual, as opposed to a firm. Table 1 shows summary statistics on these data.

Chilean Imports (000's US$)	24,251,469
Number of Importers	19,238
Number of HS 8-digit codes imported	6,587
Number of Source countries	133
Number of Shipments	1,460,847
Average Shipment Value (US$)	16,660
Median Shipment Value (US$)	1,954

Table 1: Summary statistics - Chilean import data (2006)

Table 2 begins to describe some of the features of interest in the data. The top panel of the table shows the per importer distribution of total imports (column 1), number of HS 8-digit products imported (column 3), number of countries purchased from (column 4) and number of import shipments per importer (column 5). As is typical in trade data, the distribution of total

[2]Mineral fuels account for about 15 percent of all Chile's imports in 2006. Import shipments of mineral fuels are markedly different than the rest of Chilean import shipments. For instance, they are significantly larger than the typical shipment.

imports is highly skewed: the average Chilean importer purchases close to US$1.2 million from abroad while the median importer purchases less than US$ 45,000. The same kind of skewness is present in the other distributions as well. Note that, in terms of HS 8-digit products, the median firm imports 5 products while the top one percent importers purchase 156 products. Similarly, if we focus on source country, the median firm imports products from 2 countries while the top one percent importers purchase products from more than 20 countries.

	Imports (000's US$) (1)	# Importers (2)	# of HS8 Codes (3)	# Source Countries (4)	# of Shipments (5)
Distribution over Importers: 19,238 firms					
P25	6.9	–	2	1	2
P50	44.3	–	5	2	8
P75	265.4	–	14	3	38
P90	1,388.6	–	38	7	135
P99	22,468.5	–	156	21	1,071
Mean	1,260.0	–	15.1	3.2	75.9
Distribution over HS 8-digit products: 6,587 codes					
P25	73.5	4	–	3	11
P50	492.2	13	–	8	46
P75	2,261.2	41	–	14	177
P90	7,098.1	111	–	24	527
P99	53,665.7	447	–	42	2,756
Mean	3,681.7	44.1	–	10.4	221.7
Distribution over Importer-HS 8-digit product pairs: 290,348 pairs					
P25	0.5	–	–	1	1
P50	2.5	–	–	1	1
P75	14.7	–	–	1	3
P90	75.8	–	–	2	9
P99	1,203.1	–	–	6	59
Mean	83.5	–	–	1.4	5.0

Table 2: Distribution of key variables - Chilean imports (2006)

The second panel of Table 2 shows the per HS 8-digit product distributions of total imports, number of importers, number of countries purchased from and number of import shipments carried out. Note that the average HS 8-digit product accounts for more than US $ 3.5 million in imports and is imported by almost 45 Chilean firms from more than 10 source countries. Again these distributions are skewed, with the median number of importers and source countries being much smaller.

Finally, the last panel of Table 2 shows distributions over importer-HS 8-digit product pairs. These distributions are particularly skewed so that the vast majority of cases involve an importer buying a given product from a single country and in small dollar amounts. For more than 10 percent of the importer-product pairs (close to 40,000 pairs), however, the importer buys the same HS 8-digit product from multiple countries and in large dollar amounts. These are the cases that provide the ideal identification for the main stylized facts discussed in the next section and so we describe them in more detail below.

# Source Countries (1)	# Importer-HS 8 Pairs (2)	Share of Importer-HS 8 Pairs (3)	Share of Imports (4)	HHI Imports across countries (5)	Abs. Dev. Country Per-Capita Income (6)
1	232,467	0.801	0.327	1	0
2	35,921	0.124	0.207	0.71	8,699
3	11,411	0.039	0.138	0.60	9,622
4	4,907	0.017	0.117	0.53	9,767
5	2,421	0.008	0.059	0.49	10,136
6-10	2,874	0.010	0.115	0.43	10,304
11-15	287	0.001	0.030	0.38	10,738
16+	60	0.0002	0.008	0.33	11,561

Table 3: Characteristics of Importer-HS 8-digit pairs

Table 3 shows characteristics of importer-HS 8-digit product pairs by the number of countries from which an importer buys the HS 8-digit product. For 80 percent of all importer-product pairs, the importer buys the product from one source country only. These cases, however, account for less than 33 percent of Chile's imports (see column 4), meaning that almost 70 percent of Chile's imports are carried out by importers buying the same product from at least two different countries. In fact, more than 20 percent of imports are accounted for by importers who buy the same product from at least 5 countries.

The purchase volumes are also relatively dispersed across countries. This fact is shown in column 5, which gives the Herfindahl-Hirschman Index (HHI) of importers' purchases of the same product across source countries. We see, for instance, that in cases in which importers buy a given product from 4 countries, the average HHI is 0.53. This implies that, if one randomly selects two dollars in imports within this group of importer-product pairs, there is a 53 percent chance that these two dollars will come from the same country, and thus a 47 percent chance they will come from different countries.

We also find that importers that buy the same product from different countries tend to buy the product from countries at different levels of development. This is shown in the last column in Table 3 which gives the mean absolute deviation in the per capita incomes of the countries from which importers buy the same product. In essence, what this column shows is that the average difference in per capita income between the countries a Chilean importer buys a given product from is of the same magnitude as the difference in per capita income between Brazil and the Czech Republic - roughly US$ 10,000.

In the next section, we exploit this cross-country variation in import purchase amounts and per-capita incomes to derive a relationship between source country per-capita income and characteristics of import shipments. In particular, we show that, for a given product, shipments are smaller and more frequent and the average price is higher when the product is imported from a high per-capita income country than from a low per-capita income country. This is true even when controlling for import volume, the identity of the importer, as well as the exporter.

3 The frequency, size and price of import shipments

To derive the relationship between the characteristics of import shipments and source country per-capita income, we start by decomposing the value purchased of product h from country c by Chilean importer i (M_{hci}) as:

$$M_{hci} = N_{hci}^{m} \times \frac{\sum_{s=1}^{N_{hci}^{m}} m_{hci}(s)}{N_{hci}^{m}} = N_{hci}^{m} \times \bar{m}_{hci},$$

where $m_{hci}(s)$ is the value of the s-th shipment carried out by the importer, N_{hci}^{m} is the number of shipments importer i carries out when purchasing product h from country c, and thus $\bar{m}_{hci}$ is the average value of these shipments. This average shipment value can be further decomposed as the product of the average physical quantity shipped ($\bar{q}_{hci}$) and the weighted average per-unit price of the goods in the shipment ($\hat{p}_{hci}$):

$$\bar{m}_{hci} = \frac{1}{N^m_{hci}} \sum_{s=1}^{N^m_{hci}} \Big(q_{hci}(s) \times p_{hci}(s) \Big) = \hat{p}_{hci} \times \bar{q}_{hci},$$

where

$$\hat{p}_{hci} = \frac{\sum_{s=1}^{N^m_{hci}} \Big(q_{hci}(s) \times p_{hci}(s) \Big)}{\sum_{s=1}^{N^m_{hci}} q_{hci}(s)},$$

and

$$\bar{q}_{hci} = \frac{\sum_{s=1}^{N^m_{hci}} q_{hci}(s)}{N^m_{hci}}.$$

Therefore we can writte:

$$M_{hci} = N^m_{hci} \times \bar{m}_{hci} = N^m_{hci} \times \hat{p}_{hci} \times \bar{q}_{hci}.$$

Each part of the decomposition above, as well as the total amount imported, is the outcome of optimal decisions made by the Chilean importer and its foreign exporters. To study how these decisions are affected by characteristics of the exporting countries, especially per-capita income, we regress each component of the decomposition on a number of country characteristics, most notably the "gravity" variables. In particular, for each variable $X_{hci} = \big[M_{hci}, N^m_{hci}, \bar{m}_{hci}, \hat{p}_{hci}, \bar{q}_{hci}\big]$, we estimate the equation below:

$$\ln(X_{hci}) = \delta_{hi} + \beta_1 \ln(gdp_c) + \beta_2 \ln(gdp\_pc_c) + \beta_3 \ln(dist_c) + \epsilon_{hci},$$

where gdp_c, gdp_pc_c, and $dist_c$ measure the exporting country's GDP, per capita GDP, and physical distance to Chile, respectively. The terms δ_{hi} capture importer-product fixed-effects, so that the effects of source country characteristics on each of the import variables are identified off importers that buy the same product from multiple countries.[3] Finally, when estimating the equations above, the error term (ϵ_{hci}) is clustered at the country level.[4]

The first column in Table 4 shows the gravity effects, i.e., the effect of GDP, per capita GDP, and distance on dollar values traded at the importer-product level. Interestingly, within importer-

[3]We obtain similar results in a specification that includes product and importer effects, as opposed to product-importer effects. We do not report these estimates to conserve space, given that our preferred specification is the one that accounts for importer-product pair heterogeneity.

[4]In a few cases, the quantity purchased by the same importer of a given HS 8-digit product is recorded in different units in different shipments. This makes it impossible for us to compare physical amounts and unit prices across these shipments. Given that this is an issue for 0.15% of the shipments only, we drop these shipments from the sample.

product pairs, Chilean firms buy more from larger countries but not from richer or closer-by countries. This suggests that the distance effect commonly observed in aggregate- and product-level trade data, and present in the Chilean data, is driven by importer and product composition effects. The second and third columns in the same table show that exporting country characteristics affect the decision on how best to structure import shipments. Chilean importers buy more from larger countries (high GDP) both because they receive more shipments from these countries and because they have larger shipments, both in values as well as in physical quantities. Exporters in larger economies do not tend to sell higher-priced varieties, however.

A very different picture emerges when we look at the effect of the exporting country's level of development (per capita GDP). Although Chilean importers do not tend to buy more from richer countries, they make more frequent and physically smaller transactions when buying from these countries. As previously established in the literature, they also buy higher-priced varieties from rich countries. Lastly, physical distance from Chile only affects the frequency of shipments importers make.

Dep. variable:	$\ln(M_{hci})$ (1)	$\ln(N^m_{hci})$ (2)	$\ln(\bar{m}_{hci})$ (3)	$\ln(\bar{q}_{hci})$ (4)	$\ln(\hat{p}_{hci})$ (5)
$\ln(gdp_c)$	0.212***	0.162***	0.049***	0.070**	-0.022
	(0.035)	(0.028)	(0.014)	(0.029)	(0.025)
$\ln(gdp\_pc_c)$	0.031	0.062*	-0.031	-0.287***	0.258***
	(0.059)	(0.037)	(0.029)	(0.065)	(0.045)
$\ln(dist_c)$	-0.064	-0.109***	0.045	0.039	0.006
	(0.068)	(0.037)	(0.038)	(0.074)	(0.052)
Importer-HS 8 FE	Yes	Yes	Yes	Yes	Yes
Std. Error Cluster	Country	Country	Country	Country	Country
N	395,004	395,004	395,004	394,522	394,522
R^2	0.84	0.69	0.88	0.92	0.95

Note: Standard errors in parentheses. ***,** and *denote significance at the 1%, 5% and 10% levels respectively.

Table 4: Importer-HS 8-digit product fixed effect regressions

While Chilean importers receive frequent and small shipments when importing a product from a high per-capita income country, it is unclear whether this is because the importer deals more frequently with each exporter (and so imports less each time) or because the importer deals with more exporters (and so imports less from each one). We can disentangle these two effects using information from Chile's customs forms on the identity of the foreign exporters from whom each Chilean importer buys. We do this in a sub-sample of the data including the 30 HS 8-digit codes with the largest number of importers sourcing that product from multiple countries.[5] This is obviously a selected sample built with the intention of maximizing the power of our identification strategy, which is based on importers that buy the same HS 8-digit product from multiple countries. In the Appendix, we describe this sub-sample and show that it shares many of the main characteristics of the full sample. Importantly, the main findings described in Table 4 hold in this sub-sample as well, as Table 5 shows. The noticeable difference in this sub-sample is that importers not only structure their shipments differently when buying from richer countries, but also tend to buy more from richer countries.

Dep. variable:	$\ln(M_{hci})$ (1)	$\ln(N^m_{hci})$ (2)	$\ln(\bar{m}_{hci})$ (3)	$\ln q_{hci})$ (4)	$\ln(\hat{p}_{hci})$ (5)
$\ln(gdp_c)$	0.228***	0.180***	0.048***	0.085***	-0.040
	(0.048)	(0.038)	(0.016)	(0.032)	(0.025)
$\ln(gdp\_pc_c)$	0.245***	0.211***	0.033	-0.245***	0.283***
	(0.094)	(0.056)	(0.044)	(0.081)	(0.042)
$\ln(dist_c)$	-0.161*	-0.139**	-0.021	-0.026	0.004
	(0.091)	(0.062)	(0.042)	(0.0081)	(0.049)
Importer-HS 8 FE	Yes	Yes	Yes	Yes	Yes
Std. Error Cluster	Country	Country	Country	Country	Country
N	41,102	41,102	41,102	41,011	41,011
R^2	0.75	0.59	0.80	0.86	0.86

Note: Standard errors in parentheses. ***,** and *denote significance at the 1%, 5% and 10% levels respectively.

Table 5: Importer-HS 8-digit product fixed effect regressions. Top 30 Sample

To identify the importance of the two effects on shipment frequency, we decompose the total number of shipments carried out by a given importer when purchasing product h from country c

[5]The reason we do this only on a subsample of HS 8-digit products is that the process for cleaning the names of foreign exporters in each shipment is quite cumbersome and resource intensive. It is worth noting, however, that we have access to the actual name of the exporting company as reported in the Chilean customs form. This information allows us to clean the data for name mis-spellings, a problem that plagues much of the empirical analysis in the literature using matched importer-exporter data.

(N^m_{hci}) into the number of exporter partners the importer transacts with and the average number of shipments it receives from each of its export partners:

$$N^m_{hci} = n^x_{hci} \times \frac{N^m_{hci}}{n^x_{hci}} = n^x_{hci} \times \bar{n}^m_{hci},$$

where n^x_{hci} is the number of export partners that importer i buys product h from in country c and $\bar{n}^m_{hci}$ is the average number of shipments that importer i has in product h per export partner in country c. Table 6 shows how country characteristics, especially per-capita income, impact these two shipment size effects.

Dep. variable:	$\ln(N^m_{hci})$	$\ln(n^x_{hci})$	$\ln(\bar{n}^m_{hci})$
	(1)	(2)	(3)
$\ln(gdp_c)$	0.180***	0.108***	0.072***
	(0.038)	(0.021)	(0.019)
$\ln(gdp\_pc_c)$	0.211***	0.082***	0.128***
	(0.056)	(0.027)	(0.038)
$\ln(dist_c)$	-0.139**	-0.047*	-0.093*
	(0.062)	(0.027)	(0.052)
Importer-HS 8 FE	Yes	Yes	Yes
Std. Error Cluster	Country	Country	Country
N	41,102	41,102	41,102
R^2	0.59	0.55	0.60

Note: Standard errors in parentheses. ***, ** and *denote significance at the 1%, 5% and 10% levels respectively.

Table 6: Seller's Margin. Top 30 Sample.

We see that, when buying from larger countries, Chilean importers buy from more exporters, as well as has more frequent shipments with each exporter. Close to two-thirds of the variation, however, is due to importers buying from more exporters - the sellers margin of trade. Chilean importers also buy from more exporters and receive more frequent shipments when buying from richer countries. In this case, however, the shipments-per-exporter margin accounts for more than half of the variation. Finally, distance from Chile affects the importers' decision both on how many sellers to trade with and on the frequency of shipments.

The overall message from the data is this: Much of Chile's trade volume involves a given importer buying the same HS 8-digit product from both rich (high per-capita income) and poor (low per-capita income) countries. This trade is structured such that the importer receives smaller and more frequent shipments when buying from the rich country, and pays higher prices. More than half of this increased shipment frequency is due to the importer purchasing smaller amounts on a more frequent basis from a given set of exporters. And given that importers buy, in total, the same or more from richer countries, this suggests that trade from rich countries differ from trade from poor countries in that they carry higher quality products in smaller and more frequent shipments.

The existing literature has mechanisms to explain why rich countries export (and import) high quality goods. It also has mechanisms for why the size and frequency of trade shipments may vary with country characteristics. In the next section, we provide a model of trade and distribution in which these features of the data are jointly determined. The model endogenizes the choice of export shipment size and frequency, as well as the quality of the product exported, and shows that trade costs of a certain nature will be source of comparative advantage in quality for countries, and will drive the structure of the country's trade shipments.

4 A model of trade and distribution with endogenous product quality and shipment frequency

Consider a world in which there are C countries indexed by $j = 1, 2, ..., C$. There is a single factor of production in each country, labor, with the (fixed) endowment of labor in country j denoted by L_j. In any given country, there is a single, monopolistically competitive manufacturing sector producing a differentiated intermediate good. Varieties of this good, which we denote by ω, are horizontally differentiated both within and across countries, as well as vertically differentiated across countries. These intermediate varieties, in turn, are used to produce a non-traded final consumption good. We describe both the final goods market and the intermediates market below. All variables are expressed in per capita terms.

4.1 Final goods

Consumers have identical taste both within and across countries and consume a single non-traded final good. The utility of a representative consumer in country i is simply the quantity of the final good consumed and is given by

$$U_j = (c_j^f)^{\beta}(l_j^f)^{1-\beta},$$

where l_j^f is the amount of labour used to produce the final good, and c_j^f is the amount of a CES aggregate of the manufactured intermediates available in country j:

$$c_j = \left[\sum_{i=1}^{C} \mathbf{1}_{ij} \int_{\omega} q_{ij}(\omega) x_{ij}(\omega)^{\frac{\sigma-1}{\sigma}} d\omega \right]^{\frac{\sigma}{\sigma-1}},$$

where $\mathbf{1}_{ij}$ is an indicator function that takes on a value of 1 if a variety produced in country i is sold in country j. $q_{ij}(\omega)$ and $x_{ij}(\omega)$ gives the quality and quantity of variety ω produced in country i that is used to produce the manufactured aggregate in country j. Hence, we allow a producer to sell goods of different quality to different countries. Denoting the price of a representative variety ω sold in country j by $p_{ij}(\omega)$, the quantity demanded by an individual consumer in country j of that variety is given by [6]

$$x_{ij}(\omega) = \frac{q_{ij}(\omega)^{\sigma} \left(p_{ij}(\omega)\right)^{-\sigma}}{P_j^{1-\sigma}} \beta y_j, \tag{1}$$

where y_j is the income (and expenditure) of an individual in country j and P_j is the price index of the manufactured aggregate in country j:

$$P_j^{1-\sigma} = \sum_{i=1}^{C} \mathbf{1}_{ij} \int_{\omega} q_{ij}(\omega)^{\sigma} p_{ij}(\omega)^{1-\sigma}. \tag{2}$$

4.2 Intermediate goods

A manufacturer of an intermediate variety in country i produces using labor and the same composite of all the available manufactured intermediates, c_i. The production technology for intermediates is described by a constant returns to scale, Cobb-Douglas production function:

$$x_{ij}(\omega) = e^{-aq_{ij}} \left(h_i l_{ij}^m(\omega)\right)^{\alpha} \left(c_{ij}^m(\omega)\right)^{1-\alpha},$$

where l_{ij}^m is the quantity of labor used while c_{ij}^m is the amount of the manufactured aggregate used. h_i is a labour-augmenting productivity parameter. The share of value-added in intermediate good production, α, is assumed to be the same across countries. Under this assumption, the unit cost of producing a variety of quality $q = 0$, the lowest possible quality, in country i is given by (up to a constant)

$$r_i = \left(\frac{w_i}{h_i}\right)^{\alpha} P_i^{1-\alpha}, \tag{3}$$

[6] We can instead assume that there is a non-traded sector that combines the manufactured aggregate with labour to produce the final good. In this case, the demand for intermediates will come from the non-traded sector.

where w_i is the wage rate in Country i. The unit cost of producing a variety of quality q_{ij} is then $r_i e^{aq_{ij}}$ – unit cost is increasing in quality.

4.3 Trade

A manufacturer in country i exporting to country j incurs the usual iceberg trade cost, so that one unit landing in country j requires $\tau_{ij} > 1$ units of production in country i. In addition, the manufacturer also incurs a fixed export cost of K_i each time that the manufacturer ships any product from country i to any other country, including itself. This per-shipment cost reflects such things as the administrative costs of sending an international shipment from country i, the efficiency of the underlying consolidation process and other exit costs that do not depend on the size of the shipment.

Products from country i are imported into country j via a perfectly competitive import / distribution sector in country j. For simplicity, we assume that the only cost that this sector incurs, beyond the cost of the imported product, is an inventory management cost. This cost takes the form of a proportional (to import shipment value) depreciation cost, δ, which may represent anything from pure wastage, to interest costs, to potential obsolescence or to direct storage costs. Together, these latter two costs create a decision problem for the manufacturer in country i as regards how many shipments to make in any given period and the size of each shipment. This choice affects the overall cost of selling its product in country j and hence, the quality and price of its product in j. In the next section, we describe this problem and determine the optimal shipping strategy, the choice of quality and the associated price.

5 Equilibrium

To determine the amount of trade from any country i to country j, the quality of products manufactured in country i and the price of products from i sold in j, we first must solve the shipment problem for a manufacturer in country i. The solution of this problem will determine an implicit trading cost between country i and country j and so determine the choice of quality and the equilibrium price. We turn, then, to the manufacturer's shipment decision.

5.1 Shipping Problem

To create both a tractable and non-trivial shipping choice problem, we think of the above model as defining demand for a single time period of length $T = 1$. Within this time period, we assume that consumption is a flow that occurs uniformly over the period. Because all varieties are symmetric, manufacturers in country i will charge the same price and sell the same amount to country j in equilibrium. Hence, we can drop ω from any future analysis. This implies that, for any total (for the entire period T) quantity demanded in country j for a variety produced in country i, $X_{ij} = x_{ij}L_j$, instantaneous demand in Country j (quantity demanded for an interval of time dt) is $X_{ij}dt$.[7]

The manufacturer's shipping problem for any total quantity demanded, X_{ij}, is to allocate export shipments over the entire time interval $[0, 1]$ in a way that meets flow demand in a cost minimizing way; that is, for any X_{ij}, the exporter must choose a frequency of export shipments, N_{ij}, and size of shipment arriving in Country j, S_{ij}, to minimize shipping and inventory costs.[8] This shipping decision will depend on the level of the distributor's inventory at any date, $INV_{ij}(t)$ (recall from above that shipments to country j are inventoried and distributed by a local import distributor). The law of motion of inventory, given total quantity demanded X_{ij}, is

$$\frac{dINV_{ij}(t)}{dt} = -X_{ij} - \delta INV_{ij}(t).$$

Solving for the change in inventory between some initial time t_0 and $t > t_0$ gives us

$$INV_{ij}(t) = e^{\delta(t_0-t)}INV_{ij}(t_0) - \frac{X_{ij}}{\delta}(1 - e^{\delta(t_0-t)}).$$

Under any optimal inventory policy, an export shipment should arrive just as inventories in j go to zero. If we define this point as $t = t_0$, then inventory at $t = t_0$ is just the quantity of goods in the shipment: $INV_{ij}(t_0) = S_{ij}$. Because the time interval between two successive shipments is $\frac{1}{N_{ij}}$, the inventory as $t \to t_0 + \frac{1}{N_{ij}}$ must be zero. This is shown in Figure 1. Utilizing these facts, we can solve the above equation for the optimal shipment size, conditional on any given shipment frequency, N_{ij}, and X_{ij}:

$$S_{ij} = \frac{X_{ij}}{\delta}(e^{\delta/N_{ij}} - 1). \tag{4}$$

The above equation suggests that for fixed total quantity demanded, the size of shipments is

[7] In this paper, we assume deterministic demand. With stochastic demand, it is usually not possible to get analytical solutions for the inventory management problem, and one has to resort to numerical methods [See Alessandria et al., 2010]. In the appendix, we present a simple two-period model with demand uncertainty that permits closed-form solutions. The main predictions of the theory go through.

[8] Our assumption that shipment sizes will be the same is without loss of generality. As shown by Arrow et al. [1951], given the nature of demand, orders of equal size at equal intervals are optimal.

decreasing in shipment frequency (N_{ij}).

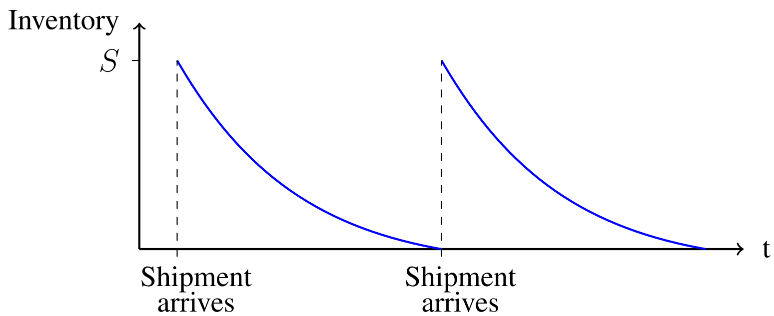

Figure 1: Evolution of inventory over time

Because a cost of K_i is incurred each time a shipment is sent, the exporter faces a trade-off between number of shipments, N_{ij}, and shipment size, S_{ij}. On the one hand, less frequent shipments mean that the importer has to incur K_i fewer times. This reduces the total ordering cost $N_{ij}K_i$. On the other hand, the importer then has to hold a bigger inventory on average, which leads to more wastage due to depreciation. This added wastage raises the inventory cost. Because the higher cost is passed on to the consumers, this negatively affects the exporter's sales. Accordingly, the exporter internalizes the inventory cost of choosing more/less frequent shipments.

The exporter's total *distribution cost*, DC_{ij}, can then be expressed as

$$DC_{ij} = \underbrace{N_{ij}K_i}_{\text{ordering cost}} + \underbrace{\tau_{ij} r_i e^{aq_{ij}} \left(N_{ij} S_{ij} - X_{ij} \right)}_{\text{inventory cost}}.$$

The first term on the right-hand side gives the total fixed shipment costs. To understand the second term, note that, for the exporter to sell X_{ij} units of a variety in country j, it must have $N_{ij}S_{ij}$ units of the product land in country j. Hence, $N_{ij}S_{ij} - X_{ij}$ is the total physical wastage. Because the exporter incurs a cost of $\tau_{ij} r_i e^{aq_{ij}}$ for each unit landing in country j, the second term measures the added cost that the exporter incurs due to inventory management. Using (4), the delivery cost can be expressed only as a function of N_{ij} and X_{ij} :

$$DC_{ij} = N_{ij}K_i + N_{ij}(\tau_{ij} r_i e^{aq_{ij}} X_{ij}) \Big(\frac{e^{\delta/N_{ij}} - 1}{\delta} - \frac{1}{N_{ij}} \Big).$$

For any value of X_{ij}, the exporter chooses N_{ij} to minimize DC_{ij}. The following lemma provides the solution for N_{ij}.

Lemma 1. *The equilibrium value of N_{ij} solves the following implicit equation:*

$$\frac{1}{N_{ij}} e^{\delta/N_{ij}} + \frac{1}{\delta}(1 - e^{\delta/N_{ij}}) = \frac{K_i}{\tau_{ij} r_i e^{aq_{ij}} X_{ij}}. \tag{5}$$

It is straightforward to show that the left-hand side of the above equation is monotone decreasing in N_{ij} while the right-hand side is independent of N_{ij}. Thus, the optimization problem of the exporter yields a unique N_{ij} and DC_{ij} corresponding to every X_{ij}. The next proposition discusses two important comparative statics results.

Proposition 1. *An exporter*
(i) sends more shipments when selling more, for a given per shipment cost.
(ii) sends more shipments when facing a lower per shipment cost, for a given demand.

Results (i) and (ii) of Proposition 1 are a direct consequence of the inventory management problem. At the margin, the cost of an additional shipment must equal the savings on inventory management. Hence, when per shipment cost declines, efficiency requires that the exporter increase the number of shipments, while reducing their size. On the other hand, when total demand rises, there is a scale effect. Higher demand allows the exporter to spread the per shipment costs over more units, thereby allowing him to choose more shipments.[9]

The value of DC_{ij} at the optimal N_{ij} is given by

$$DC^*_{ij} = (e^{\delta/N_{ij}} - 1)\tau_{ij} r_i e^{aq_{ij}} X_{ij}. \tag{6}$$

Notice that as long as $\delta > 0$, $e^{\delta/N_{ij}} > 1$. Hence, (6) suggests that the inventory management problem results in a distribution cost that is proportional to X_{ij}. The factor of proportionality is not a constant, however. With endogenous shipment frequency, as demand starts to increase, the exporter wants to send more shipments to save on inventory costs. As a result, DC^*_{ij} increases less than proportionately relative to X_{ij}.

The interaction between DC^*_{ij}, K_i and X_{ij} is exhibited in Figure 2. When X_{ij} increases, the way DC^*_{ij} changes depends on the value of K_i. This is shown in Figure 2a. For $K_i = 0$, a distributor does not have to carry any inventory; the exporter simply sends a shipment every

[9] The effect of increased demand on shipment size is less straightforward. Demand has both a direct positive effect on S and an indirect negative effect through N. Intuitively, a higher demand means that the importer should have larger shipments, other things remaining constant. But a higher demand also raises the optimal frequency of shipments, which tends to reduce the size of shipments. Which effect dominates depends on the magnitude of δ. When δ is small, there is less wastage and the former effect dominates, resulting in bigger shipments.

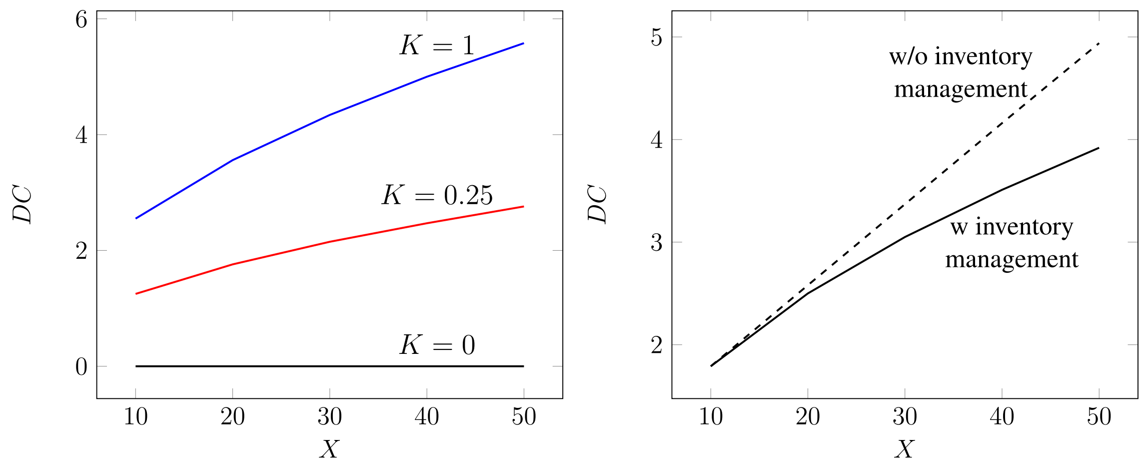

(a) DC for different K

(b) Model with and without inventory management

Figure 2: Relation between DC and output

time there is some demand for the product. In this case, DC^*_{ij} is equal to zero. For $K_i > 0$, following an increase in demand, the exporter re-optimizes by increasing both the size, as well as the number of shipments – DC^*_{ij} becomes an increasing function of X_{ij}. For a given value of X_{ij}, DC^*_{ij} also increases with K_i as the exporter reduces the number of shipments, causing an increase in inventory cost.

Figure 2b shows how DC^*_{ij} differs between a model with a fixed number of shipments (the standard model) and one with the optimal shipping decision. The figure suggests that when X_{ij} is small, DC^*_{ij} in the two cases are almost identical - for small values of X_{ij}, the optimal number of shipments is close to one. But as X_{ij} starts to increase, by forcing the exporter to have a fixed number of shipments, the standard model increases the delivery cost relative to our model.

5.2 Export Equilibrium

The export equilibrium is defined by a two-stage game. In the first stage, each exporter in country i simultaneously chooses a quality level for its variety. In the second stage, these exporters simultaneously choose a price at which to sell its variety in country j, p_{ij}, conditional on the quality choices made in stage one. The export equilibrium is given by the subgame-perfect Nash equilibrium of this two stage game. We solve for this equilibrium in the usual way, first determining the pricing equilibrium for arbitrary quality choices and then solving for the quality equilibrium given the pricing equilibrium.

Choosing price: Denoting the profit to an exporter from country i selling a variety in country j

as π_{ij}, we can define the exporter's pricing problem as

$$\max_{p_{ij}} \pi_{ij} = \big(p_{ij} - \tau_{ij} r_i e^{aq_{ij}}\big) x_{ij} L_j - DC^*_{ij}.$$

where $x_{ij} = X_{ij}/L_j$ is given by (1). Note that since DC^*_{ij} is a function of X_{ij}, in maximizing profit an exporter simultaneously solves for the optimum frequency and size of shipments. To see this, replace the value of DC^*_{ij} from (6) in the above equation to get

$$\max_{p_{ij}} \pi_{ij} = \big(p_{ij} - e^{\delta/N_{ij}} \tau_{ij} r_i e^{aq_{ij}}\big) x_{ij} L_j.$$

The resulting profit-maximizing price is

$$p_{ij} = \Big(\frac{\sigma}{\sigma - 1} \frac{e^{\delta/N_{ij}} - 1}{\delta/N_{ij}}\Big) \tau_{ij} r_i e^{aq_{ij}}. \tag{7}$$

In the presence of the inventory management problem, the exporter loses a part of the output due to depreciation. This output loss effectively raises the cost of the units that the exporter eventually sells - inventory management ends up imposing an ad-valorem tax on exports. We can re-write the above equation as

$$p_{ij} = \frac{\sigma}{\sigma - 1}\big[d_{ij} + \tau_{ij} r_i e^{aq_{ij}}\big],$$

where $d_{ij} = (\frac{e^{\delta/N_{ij}-1}}{\delta/N_{ij}} - 1)\tau_{ij} r_i c e^{aq_{ij}}$ is the *marginal distribution cost*. This cost, which is an equilibrium object, is a function of per-shipment cost and final demand, among other things. The following proposition characterizes d_{ij}:

Proposition 2. *The marginal distribution cost has the following properties:*
(i) $d_{ij} \geq 0$,
(ii) $\frac{\partial d_{ij}}{\partial X_{ij}} < 0$,
(iii) $\frac{\partial d_{ij}}{\partial K_i} > 0$.

When $\delta = 0$, there is no inventory management problem and $d_{ij} = 0$. In this case, the price coincides with the standard mark-up over marginal cost of production. But as δ becomes positive, $d_{ij} > 0$. In this case, part (ii) of Proposition 2 reveals that, even if the marginal cost of production is a constant, distribution costs decline with total sales in country j. This is a consequence of exporter optimization: as total demand rises, the exporter sends shipments more frequently, thereby economizing on the inventory cost. The implication is that overall marginal costs decline with output: there are effectively scale economies.

This last fact has some interesting implications. First, even though σ is a constant, the mark-up over the marginal cost of production plus ad-valorem trade cost is variable. In particular, as part (ii) of Proposition 2 suggests, the mark-up is decreasing in output. As demand rises, the distributor becomes more efficient in managing inventories, which allows the exporter to lower the price, even without any change in the marginal cost of production. Second, the inventory management problem affects the choice of quality in a systematic way, as we explore below. Also important for the quality choice problem is the fact that the marginal distribution cost is increasing in per shipment cost, as shown in part (iii) of Proposition 2. A higher per shipment cost forces the exporter to send larger shipments and consequently waste more output, thereby pushing up the distribution cost at the margin.

Finally, because the import / distribution sector in country j is perfectly competitive, the distributors are compensated only for the cost they incur – the cost of inventory management. A distributor pays an exporter a *FOB* price for the product, p_{ij}^{FOB}, that is equal to the final price, p_{ij}, net of *only* the per unit inventory cost and τ_{ij}. This gives us,

$$p_{ij}^{FOB} = r_i e^{aq_{ij}} \left[1 + \frac{1}{\sigma - 1} \frac{e^{\delta/N_{ij}} - 1}{\delta/N_{ij}} \right].$$

The following proposition discusses an important property of the FOB price:

Proposition 3. *For a given marginal cost of production, an exporter receives a lower FOB price when there are more shipments.*

Let us hypothesize that manufaturers in rich countries face lower per-shipment costs of export. Proposition 3 then suggests that other things being the same, the FOB price received by exporters in rich countries should actually be lower. This is because the FOB price is inversely related to shipment frequency, and as already shown in Section 3, the latter happens to be higher for rich country exporters. But we have also shown the FOB price is higher when a Chilean importer purchases a good from a richer country, suggesting that other things are not the same. This is an important observation that we exploit below.

Choosing quality: In stage one of the game, each firm determines its profit maximizing quality. Below, we consider the case in which a firm in country i selling in country j can customize quality for country j so that q_{ij} is not necessarily equal to q_{ik}, $j \neq k$. In an appendix, we derive results for the case in which a firm in country i must deliver the same quality product to all countries: $q_{ij} = q_{ik} = q_i$. We consider only symmetric equilibria in which all firms in country i make the same quality choice. For the current case, the value of quality for a firm in country i selling to country j, q_{ij}, is defined by:

$$\max_{q_{ij}} \pi_{ij} = (p_{ij} - e^{\delta/N_{ij}} \tau_{ij} r_i e^{a q_{ij}}) X_{ij},$$

where p_{ij} is given by (7) and the number of shipments is defined implicitly by (5) above. Because the values of both p_{ij} and N_{ij} are given by their profit maximizing levels, the solution to the above problem is defined by the condition

$$\begin{aligned} 0 = \frac{\partial \pi_{ij}}{\partial q_{ij}} &= \left(p_{ij} - e^{\delta/N_{ij}} \tau_{ij} r_i e^{a q_{ij}}\right) \frac{\partial X_{ij}}{\partial q_{ij}} \\ &- a e^{\delta/N_{ij}} \tau_{ij} r_i e^{a q_{ij}} X_{ij} \\ &+ \frac{\delta}{N_{ij}^2} e^{\delta/N_{ij}} \tau_{ij} r_i e^{a q_{ij}} X_{ij} \frac{\partial N_{ij}}{\partial q_{ij}}. \end{aligned} \tag{8}$$

The first two terms on the right-hand side of the above equation give the usual quality effects. The first term represents the variable profit gain due to the fact that an increase in quality results in an increase in demand for the product ($\frac{\partial X_{ij}}{\partial q_{ij}} = \frac{\sigma X_{ij}}{q_{ij}} > 0$) and the product sells at a positive price-cost margin ($p_{ij} - e^{\delta/N_{ij}} \tau_{ij} r_i e^{a q_{ij}} > 0$). The second term gives the variable cost increase that results from an increase in quality. In the usual setting, the optimal quality represents a trade-off between these two effects. In our setting, however, there is an additional effect represented by the third term in the above equation. This effect is due to the fact that when quality increases, the value of the shipments go up allowing the firm to reduce marginal costs arising from the inventory management problem ($\frac{\partial N_{ij}}{\partial q_{ij}} > 0$, implying that the third term is positive). This additional, profit-enhancing effect creates an additional impetus for the firm to enhance quality.

Substituting in (8) for the values of p_{ij}, X_{ij}, $\frac{\partial X_{ij}}{\partial q_{ij}}$ and $\frac{\partial N_{ij}}{\partial q_{ij}}$, we obtain an implicit solution for q_{ij} given by[10]

$$q_{ij} = \frac{\sigma}{\sigma - 1} \frac{N_{ij}}{a\delta} (1 - e^{-\delta/N_{ij}}). \tag{9}$$

Because N_{ij} is a function of the fixed, per-shipment cost, K_i, the value of q_{ij} depends on the size of the per-shipment cost. To generate some intuition on the impact of K_i on quality, we plot, in Figure 3, the profit function for different levels of per shipment cost. There are two features of these plots that deserve mention. First, an exporter with a lower per shipment cost has higher profit at all levels of quality. The reason is not just that the value of K_i is lower but also that a decrease in per shipment cost increases the number of shipments, N_{ij}, and so lowers the marginal

[10]The equation (9) defines an implicit solution for q_{ij} because N_{ij} is a function of q_{ij}.

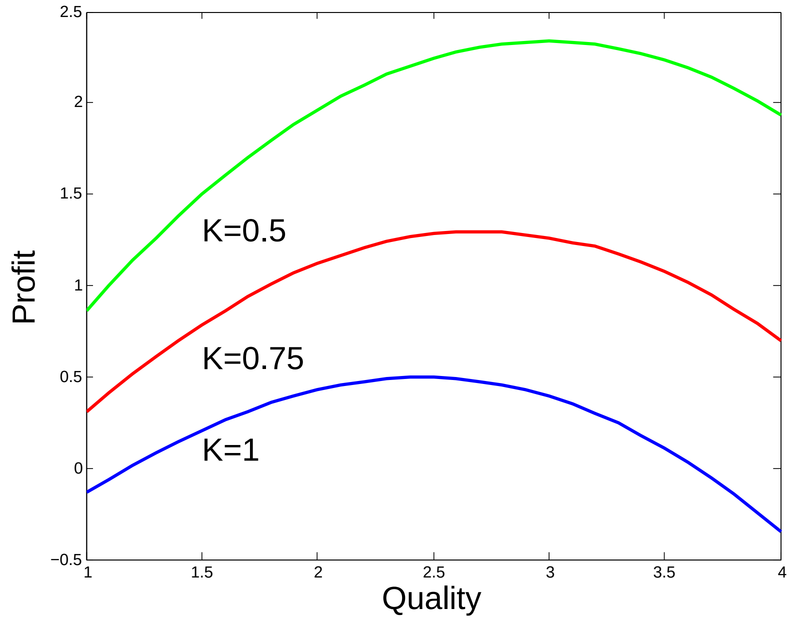

Figure 3: Profit as a function of quality

distribution cost. A lower per shipment cost thus acts as a source of *absolute advantage*.

The second feature of these plots is that the profit-maximizing level of quality is actually decreasing in per shipment cost. As a result, all else equal, exporters facing lower per-shipment costs tend to choose a higher quality product. This is a consequence of the third effect mentioned above. More specifically, when a manufacturer tries to raise the quality of his exports, the value of the goods he ships rises. *Ceteris paribus*, this causes a bigger wastage due to depreciation. The wastage, however, can be lowered by increasing the frequency of shipments and thereby reducing inventory. Now, the countries with lower per-shipment costs are the ones that can raise shipment frequency at a lower cost. A lower per-shipment cost therefore also acts as a source of *comparative advantage*. The following proposition formalizes this statement:

Proposition 4. *The value of* q_{ij} *is decreasing in* K_i*.*

Recall that one of the features of the data is that a product exported from a rich country sells for a higher price than an identical product exported from a poor country. One interpretation of this fact is that, even within narrow product categories, rich countries sell higher quality versions of the product than do poor countries. This finding is not novel. Our model, and specifically the above result, suggests something more, however: the fact that rich countries have lower per-shipment costs and that they also sell higher quality products are not unrelated. Rather, lower per-shipment costs, and so the ability to make more frequent and smaller shipments, give the rich countries a comparative advantage in the production of high quality. In equilibrium, the rich countries end up exporting high quality products *precisely because of their low per-shipment costs*.

Of course, the result in the above proposition is not a full general equilibrium result and so is, at this point, suggestive only. In addition, we would not want to claim that low per-shipment costs are the only reason that rich countries produce higher quality products. In the next section, we calibrate our model to understand the extent to which it can explain the variations we observe in shipment size and frequency and in product price across rich and poor countries. We also undertake several counterfactual exercises to quantify the impact that high per-shipment costs have on the distribution of prices, qualities and incomes across rich and poor countries.

6 Quantitative exercise

In this section, we use the multi-country general equilibrium model developed in Sections 4 and 5 to quantitatively assess the extent to which per-shipment costs incurred at exporting countries can generate the observed price difference in varieties exported by rich and poor countries. Also, we quantify the effect these per-shipment costs and the trade distortions they cause have on the distribution of income across countries. To do so, first we need to close the model developed in the previous section.

6.1 Closing the model

The total quantity sold of any country i variety to country j is $L_j x_{ij}$ and total expenditure on that variety by country j consumers is $p_{ij} x_{ij} L_j$. Because the importers in country j incur the inventory cost, the exporter receives only $\tau_{ij} p_{ij}^{FOB} x_{ij} L_j$ for every variety. Recall that the number of varieties in country i is denoted by M_i. From symmetry, total expenditure in country j on all country i varieties, T_{ij}, is then given by $M_i \tau_{ij} p_{ij}^{FOB} x_{ij} L_j$. Using (1), (3) and (7), we can write T_{ij} as

$$T_{ij} = \Delta M_i q_{ij}^{\sigma} \phi(N_{ij}) \left[\tau_{ij} \left(\frac{w_i}{h_i} \right)^{\alpha} P_i^{1-\alpha} e^{a q_{ij}} \right]^{1-\sigma} \frac{\beta Y_j}{P_j^{1-\sigma}}, \tag{10}$$

where

$$\phi(N_{ij}) = \left[1 + \frac{1}{\sigma - 1} \frac{e^{\delta/N_{ij}} - 1}{\delta/N_{ij}} \right] \left(\frac{e^{\delta/N_{ij}} - 1}{\delta/N_{ij}} \right)^{1-\sigma},$$

and $\Delta = (\frac{\sigma}{\sigma-1})^{1-\sigma}$, N_{ij} solves (5), q_{ij} solves (9) and P_j is given by (2). $Y_j = y_j L_j$ is the aggregate income (and expenditure) in country j. Replacing the value of p_{ij} in (2), we have

$$P_j^{1-\sigma} = \sum_{k=1}^{C} \mathbf{1}_{ij} M_i q_{ij}^{\sigma} \Delta \Big[\Big(\frac{e^{\delta/N_{ij}} - 1}{\delta/N_{ij}} \Big) \tau_{ij} \Big(\frac{w_i}{h_i} \Big)^{\alpha} P_i^{1-\alpha} e^{a q_{ij}} \Big]^{1-\sigma}. \tag{11}$$

We assume that the number of varieties produced in each country is fixed, and set $M_i = L_i$. Accordingly, there are profits in the manufacturing sector. Following Chaney [2008], we assume that the total profits from each country are pooled together in a global mutual fund that pays out dividends to each country in proportion to its nominal income. Furthermore, each individual in country j supplies one unit of labour inelastically. It can then be shown that $Y_j = (1+s) w_i L_i$, where s is a function of only the parameters of the model.[11]

Trade in each country j is balanced, i.e. $\sum_{i=1}^{C} T_{ij} = \sum_{i=1}^{C} T_{ji}$. This equation can be manipulated to obtain a system of equations in wages:

$$w_j = \sum_{k}^{C} \frac{L_k}{L_j} \lambda_{jk} w_k, \tag{12}$$

where $\lambda_{kj} = T_{jk}/(\beta w_k L_k)$ is the expenditure share of country k on country j's products. Finally, the market for labour and the manufacturing composite must clear. This completes our description of the equilibrium.

The standard solution strategy of such a model is as follows: for a given vector of w, use equations similar to (11) to solve for the P-s. Equipped with w and P, one can then solve for T_{ij}. Finally, using the trade shares, check whether the trade balance equation (12) is satisfied or not and keep iterating on w. In this model, however, (10) and (2) are not block-recursive. Rather, P_j depends on T_{ij} through N_{ij}. Hence, for a given vector of wages, the price indices and trade flows have to be solved simultaneously.

To solve the model, we need values for the model parameters σ, α, β, δ, L_i, h_i, τ_{ij}, a and K_i. In what follows, we discuss how we obtain the values for other parameters.

6.2 Parameter calibration

For the model parameters σ, α, β, H_i, and d_{ij} we use values from the literature. The value for the elasticity of substitution (σ) in the literature ranges from 2 to 10 [Anderson and van Wincoop,

[11] $s = \frac{\beta/\sigma}{1-\beta/\sigma}$.

2004], and we pick $\sigma = 5$.[12] Following Waugh [2010], we set value-added in production (α) at 0.3. The share of expenditure on manufactures (β) is set at 0.25, which is the median value for our sample of countries (CIA World Factbook).

The labour-augmenting productivity parameter in country i, h_i, is equal to the average human capital in the economy and L_i is the size of the workforce. Following Hall and Jones [1999], we assume that h_i takes the following functional form: $h_i = exp(\gamma(e)e)$, where e is average years of schooling and $\gamma(e)$ is piece-wise linear with slope 0.13 for $e \leq 4$, 0.10 for $4 < e \leq 8$, and 0.07 for $8 < e$. The data for average years of schooling is for 2005 and come from Barro and Lee [2013], while the data for workforce is for 2006 and come from the World Development Indicators (WDI).

Parameter	Symbol	Value	Source
Elasticity of substitution	σ	5	Anderson and van Wincoop [2004]
Value-added in production	α	0.3	Waugh [2010]
Expenditure share of manufacturing	β	0.25	CIA World Factbook
Average depreciation rate	δ	0.3	Alessandria et al. [2010]
Returns to education	$\gamma(e)$	multiple	Hall and Jones [1999]
Distance elasticity of trade	ρ	0.3	Hummels [2001]
Cost elasticity of quality	a	10	match share non-zero bilateral trade

Table 7: Calibration of various parameters

We assume a parsimonious specification for bilateral trade costs: $\tau_{ij} = dist_{ij}^{\rho}$, where $dist_{ij}$ is the distance between countries i and j. Thus, we ignore all other conventional barriers to trade typically used in the gravity literature such as tariffs, contiguity, common language, common colonial origin, etc. For $dist_{ij}$, we use the great circle distance between capital cities provided by CEPII. Hummels [2001] uses actual freight rates to compute the distance elasticity of trade costs. We use his estimate of 0.3 for ρ.

The value of the parameter a is obtained by calibrating it to match the fraction of non-zero bilateral trade flows in the data. This moment is a measure of the extent of globalization. Recall that a is the elasticity of the marginal cost function with respect to quality. As a declines, producers in every country choose higher quality, which tends to increase foreign demand for their products (as well as domestic demand), thereby allowing them to overcome the fixed cost of the

[12]Results are not very sensitive to variations in σ.

very first shipment (recall that the only "fixed" cost of exporting in the model is the per-shipment cost). For our sample of countries, the fraction of non-zero trade flows in 2000 was 0.6 (NBER-UN trade database). Thus we choose a value for a to match this number. We summarize the set of parameter values to be used in the quantitative exercise in Table 7.

It is worth noting that, by assuming that the value of a is the same across countries, we have neutralized any direct comparative advantage in quality effects in our calibrated model. We assume that the value of δ is independent of quality as well. Were the value of δ higher for higher quality products, say due to higher costs of maintaining quality during storage or to greater sensitivity to fashion trends, this would also give countries with lower per-shipment costs a comparative advantage in quality. Instead, in the calibrated model all the variation in shipment frequency and size and in price / quality is driven solely off the differences between rich and poor countries in efficiency of distribution systems at the exporting countries: the per-shipment cost.[13]

The only remaining parameters are the values for K_i and the depreciation rate while in inventory δ. We exploit the structure of our model to estimate these values from the shipment data. Specifically, recall that the values of K_i and δ are related to the number of shipments and domestic sales by equation (5):

$$\frac{e^{\delta/N_{ij}}}{N_{ij}} + \frac{1}{\delta}(1 - e^{\delta/N_{ij}}) = \frac{K_i}{\tau_{ij} r_i e^{a q_{ij}} x_{ij}}.$$

We can re-write this relationship in terms of the *import* value of variety, T_{ij}, and number of shipments by noting that $T_{ij} = \tau_{ij} p_{ij}^{FOB} \times N_{ij} \times S_{ij}$. Doing so yields the following relationship:

$$\frac{\delta/N_{ij}\Delta}{N_{ij}(e^{\delta/N_{ij}} - 1)(\delta/N_{ij} + \Delta)} = \frac{K_i}{T_{ij}},$$

where $\Delta = e^{\frac{\delta}{N_{ij}}}(\frac{\delta}{N_{ij}} - 1) + 1$.

When taking this equation to the data, we must adjust it to reflect that our data *a)* have export flows into Chile only, thus $j = Chile$; and *b)* have transactions for all HS 8-digit products and all Chilean importers. Using the subscripts h and l to index products and importers, respectively, we can re-write the equation above as:

$$\frac{\delta_{ihl}/N_{ihl}\Delta}{N_{ihl}(e^{\delta_{ihl}/N_{ihl}} - 1)(\delta_{ihl}/N_{ihl} + \Delta)} = \frac{K_{ihl}}{T_{ihl}},$$

where the subscript $j = Chile$ is ommitted to simplify notation. When $(\frac{\delta_{ihl}}{N_{ihl}})$ is small, the

[13] Two additional simplifying assumptions we make are: i) the share of expenditure on manufactures is the same across countries; and ii) we ignore physical capital in production.

equation above can be approximated by:

$$ln(T_{ihl}) - 2ln(N_{ihl}) = ln(K_{ihl}) + ln(\delta_{ihl}),$$

Assuming further that $ln(K_{ihl}) = \gamma_i + \epsilon_{ihl}$ and $ln(\delta_{ihl}) = \gamma_h + \upsilon_{ihl}$, we arrive at the estimable equation:

$$ln(T_{ihl}) - 2ln(N_{ihl}) = \gamma_i + \gamma_h + \eta_{ihl},$$

With both T_{ij} and N_{ij} coming from the Chilean imports data and $\eta_{ihl} = \epsilon_{ihl} + \upsilon_{ihl}$ arguably orthogonal to the country and product effects. The equation above highlights the fact that γ_i and δ_h are only identified up to a normalization. In line with Alessandria et al. [2010], we assume that the average value for γ_h is 0.3,[14] and this pins down all the values of γ_i and γ_h.

An important modeling assumption we are making is that per-shipment costs are only exporting country-specific. The reality may be that these costs are exporting-importing country-pair specific. That is to say, the per-shipment cost incurred in Canada when selling to Chile may be different than the one incurred when selling to the U.S. Should this be the case, then our estimated per-shipment costs, based entirely on shipments to Chile, will not be representative of the per-shipment costs countries face when trading with countries other than Chile.

As a check for this possibility, we investigate whether, after controlling for exporting country characteristcis, our estimated per-shipment costs vary systematically with exporting-importing country pair characteristics, and specifically with exporting country distance to Chile. Column 1 in Table 8 shows that, after controlling for GDP and per capita GDP of the exporting country, our estimate of per-shipment costs is correlated with the exporting country's distance to Chile. Further analysis reveals that this correlation is driven by Chile's trading partners in Latin America. These countries' estimated per-shipment costs to Chile are substantially lower than would be predicted by their sizes and levels of development. For countries outside of Latin America, distance to Chile does not effect estimated per-shipment costs. To illustrate this point, column 2 in Table 8 shows the result of a regression of our estimated per-shipment cost values on the gravity variables and on a Latin-America dummy-variable. In this specification, the effect of distance is reduced by around 80% and is no longer statistically significant.

We can use the coefficient on the Latin-America dummy variable to quantify the effect that being in the same region has on per-shipment costs. On average, Latin American countries in our sample have Ln(per-shipment cost) equal to 8.39. This fact and the results in column 2 mean

[14] Alessandria et al. [2010] assume a monthly value of 0.025 for the depreciation rate. The unit of observation in our model is a year. With this depreciation rate, a fraction $(1 - 0.025)^{12} = 0.73$ of a good is left at the end of an year. This translates into an annual depreciation rate of about 0.27

Dep. variable: $\ln(K_i)$				
	(1)	(2)	(3)	(4)
$\ln(gdp_i)$	-0.082**	-0.065*		-0.069*
	(0.037)	(0.037)		(0.037)
$\ln(gdp\_pc_i)$	-0.295***	-0.334***		-0.180**
	(0.067)	(0.067)		(0.083)
$\ln(dist_i)$	0.295***	-0.084		-0.091
	(0.108)	(0.201)		(0.204)
Latin America Dummy		-0.654**		-0.678**
		(0.295)		(0.287)
ln(# Docs to Export)			0.044	0.032
			(0.042)	(0.046)
ln(# Days to Export)			0.025***	0.009
			(0.008)	(0.001)
ln($ Cost to Export)			-0.000	0.000
			(0.000)	(0.000)
N	85	85	83	83
R^2	0.33	0.37	0.32	0.47

Note: Standard errors in parentheses. ***,** and *denote significance at the 1%, 5% and 10% levels respectively.Reported standard errors are bootstrapped.

Table 8: Per-shipment costs and country characteristics

that countries in the region, Latin America in this case, have almost 10% lower Log per-shipment costs than comparable countries–in terms of GDP, per capita GDP, and distance to Chile–from outside the region. It is worth noting that the estimated in-region reduction in per-shipment cost is quite large, with the estimated 0.65 reduction in Log per-shipment cost translating into a 48% reduction in absolute per-shipment costs for Latin American countries. The estimated in-region corrected measure of the per-shipment costs is reported in appendix Table 15. We use this in-region corrected measure of per-shipment cost in the quantitative exercises below, allowing for a 10% cost reduction for trade between countries in the same region.

Lastly, for the quantitative exercises we must take a stand on the per-shipment costs firms incur when selling in the domestic market. We assume firms face the cost to sell domestically that they face when selling to in-region foreign countries. We run sensitivity checks on this assumption and the main results we find do not depend on it.

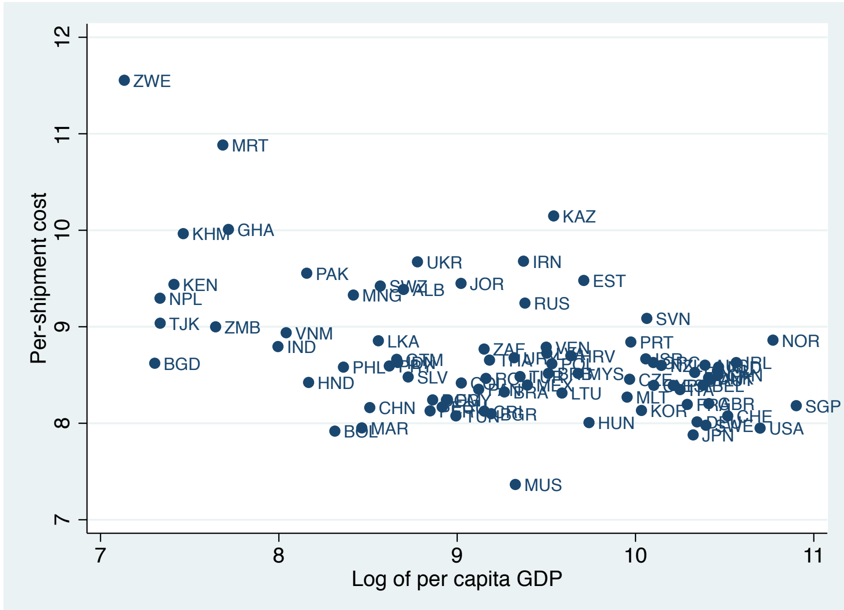

Note: The per capita GDP (PPP adjusted) data is for the year 2006 and obtained from World Development Indicators (WDI).

Figure 4: Relation between per-shipment cost and per capita GDP of the exporting country

As an additional check on the validity of our estimated per-shipment cost values, we investigate the correlation between these estimates and survey measures produced by the World Bank on the costs of doing business across borders.[15] In particular, we correlate our estimates to measures of a) the number of documents that need to be processed for exporting, b) the number of days for the release of an export shipment and c) the monetary cost incurred for the release of an export shipment in each country of our sample. These three survey metrics are highly correlated among themselves and each metric, taken on its own, is strongly correlated with our per-shipment cost estimates. Column 3 in Table 8 shows the results of a regression of our per-shipment cost estimates on the three survey measures together. Combined, these variables explain 32% of the variation in the calibrated per-shipment costs. This result indicates that, while our estimates are correlated with standard survey measures of export costs incurred at the exporting country, the estimates contain additional information on trading costs not captured by the standard measures. The in-region effects identified above are one example. As a rough comparison, while the average monetary cost incurred for the release of an export shipment for countries in our sample is equal

[15]For more information on these data, go to http://www.doingbusiness.org/methodology/trading-across-borders.

to US $1,203, in the World Bank data, the average per-shipment cost based on our methodology is almost seven times larger, at US $8,316.

We note, finally, that our model proposes that high income countries have lower per-shipment costs of trading than do low-income countries. The regression results in Table 8 show this fact, with per capita GDP negatively correlated to our estimate of per-shipment cost under the multiple specifications. Figure 4 provides a visual display of this fact, plotting (log of) per-shipment costs against (log of) per capita GDP of the countries in our sample. It shows a clear negative relation between these two variables, with a correlation of -0.49. We should emphasize here that this correlation is a feature of the data and not necessitated by our procedure for estimating per-shipment costs. In particular, no characteristics of the exporting countries entered in the calculation of the per-shipment costs other than the size and frequency of the country's export shipments to Chile.

6.3 Model fit

With the parameters of the model now fixed, we can evaluate how well the quantitative model fits the data. Observe that wage, and hence aggregate income is endogenous in our model. In the benchmark scenario (i.e., using the calibrated per-shipment costs), the correlation between (real) GDP and (real) per capita GDP in the model and the data is around 0.9.

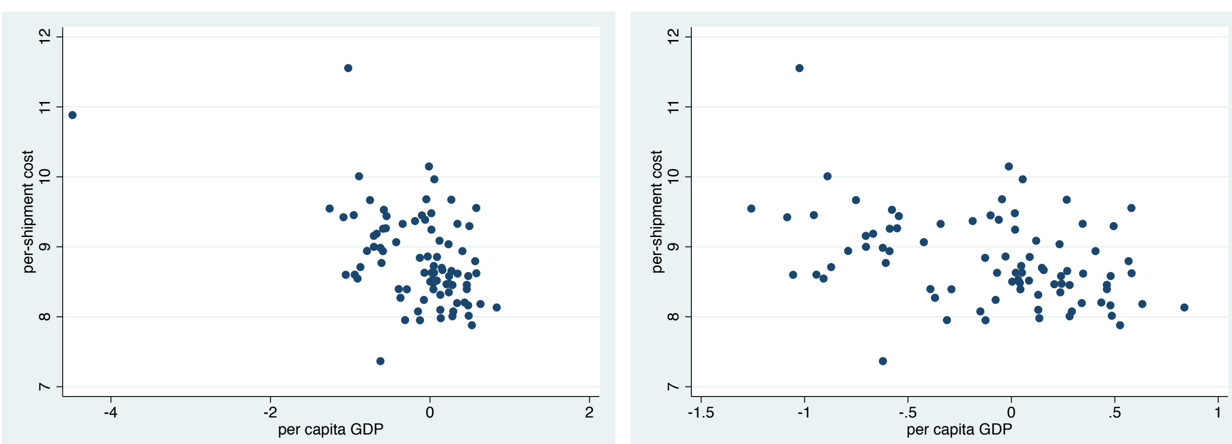

(a) Per shipment costs (b) Per shipment costs-no outliers

Note: Values for per capita GDP are in logs and expressed as deviations from the mean.

Figure 5: Distribution cost and exporter income

A key insight of our analysis is that rich countries have lower per-shipment costs than poor countries and these lower costs generate a comparative advantage in high-quality products for rich countries. Since income levels are an equilibrium outcome in our model while per-shipment costs are parameters, we first check whether or not per-capita income and per-shipment costs are negatively correlated in the calibrated model. Figure 5a reveals that this is indeed the case. In Figure 5a, one of the countries – Mauritania – shows up as an outlier with very small income per capita, so Figure 5b shows the same information excluding Mauritania.[16]

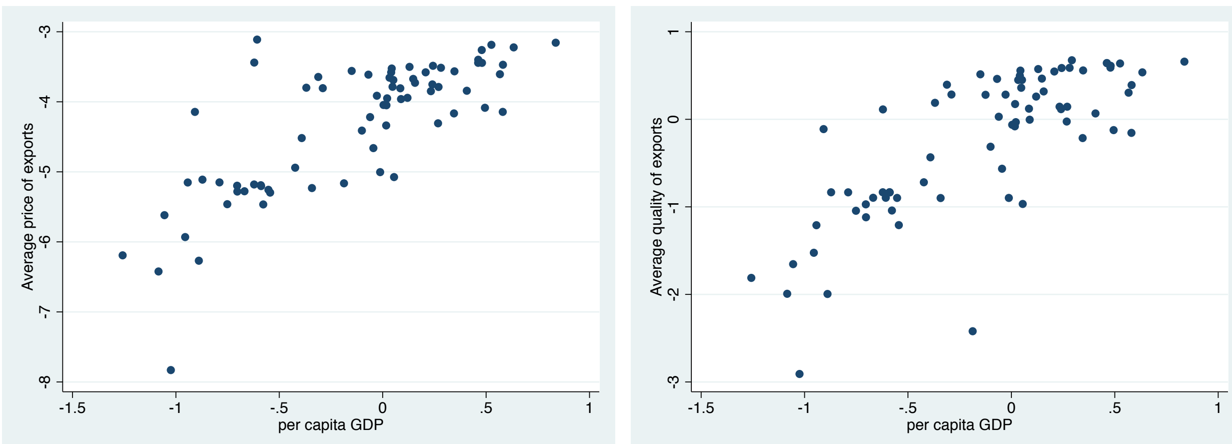

(a) Export prices and per capita income

(b) Quality and per capita income

Note: Export prices and quality are trade-weighted averages across all destinations. Values for per capita GDP are in logs and expressed as deviations from the mean.

Figure 6: Price, quality and exporter income

Next, consistent with other studies, we use price as an observed proxy for quality and explore how well the model predicts the correlation between price and per-capita income. Figure 6a plots prices p_{ij} averaged across importer j against the model-generated per-capita GDP of exporting country i. Consistent with data, the model produces a clear positive relationship between export price and per capita income. The model also produces a similar relationship between quality and exporter income, and this is reported in Figure 6b. Finally, Figure 7 shows how price varies with quality in the cross-section.

But how much of the relationship between price and per capita GDP can be generated by our mechanism? To answer this question, we use the model data to estimate the income elasticity of prices and compare it to a similar metric in data. The result is reported in Table 9. Recall from Section 3 that, within Chilean importer-product pairs, a 10 percent increase in the per capita

[16]Note that this country remains in the analysis that follows, but is excluded from figures for ease of presentation.

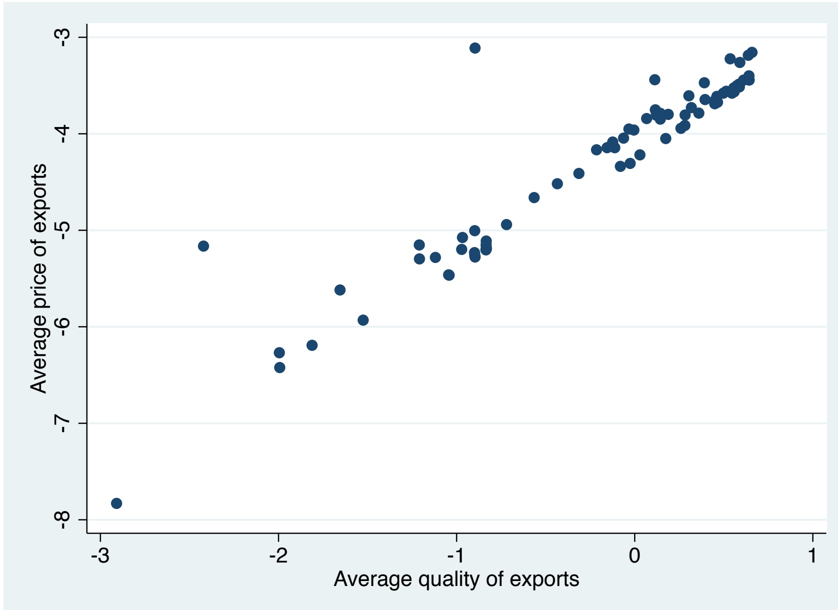

Note: Export prices and quality are averages across all destinations.

Figure 7: Export price and quality

A. Data

	All HS8 codes	Top 30 HS8 codes
	0.26***	0.28***
	(0.05)	(0.04)

B. Model

Benchmark	$K_i = K_{US}$	$K_i = K_{region}$	$K_i = 0$
0.23***	0.10***	0.10***	0.07***
(0.081)	(0.001)	(0.002)	(0.001)

Note: The numbers are coefficients on per capita GDP when (log) price is regressed on (log) GDP and (log) per capita GDP of exporter and (log) bilateral distance. Standard errors in parentheses. *** denotes significance at the 1% level.

Table 9: Price elasticities

income of the source country is related to a 2.6 percent average increase in the prices of the varieties purchased from this country when we include all HS8 codes, and a 2.8 percent average increase in the prices when looking at only the Top 30 HS8 codes. Table 9 shows the analogous elasticity for our model, derived from a regression of log p_{ij} on log GDP, log per capita GDP and log bilateral distance. Because all importing firms in the model are identical and all products are symmetric, the regressions already "control" for importing firms and products. As the model has multiple importing countries, the regression includes importer country fixed effects.

Panel A of Table 9 reproduces the price elasticity found in the Chilean import data. In the benchmark scenario reported in Panel B, Column 1, i.e., the calibrated model, the variety price elasticity with respect to exporters' per capita income is equal to 0.23 (statistically significant at the one percent level). The model, therefore, generates more than 80 percent of the sensitivity of variety price to exporter income observed in the data. This is in spite of the fact that all avenues of comparative advantage in quality, other than those arising from a lower per-shipment cost, are neutralized in the calibrated model.[17]

As our analysis emphasized the link between comparative advantage in quality, due to low per-shipment costs, and income, the calibrated model also makes predictions about cross-country income inequality. In our calibrated model, variance of log-income across countries is 0.46, while the ratio of the 90th to the 10th percentile of the country income distribution is 4.1. This is shown in Table 11. As a reference, for our sample of 85 countries, the value of the same metrics are 0.96 and 16 respectively.

Finally, Waugh [2010] argues that, to explain trade flow and price data, export costs must take an ad-valorem form and must be decreasing in the exporting country's level of development. As discussed in Section 4, in the presence of an inventory management problem, the per-shipment costs show up as ad-valorem export costs in equilibrium. From our model, we can calculate the implied ad valorem export costs and correlate them with the per-capita GDP of the exporting countries. The model produces a strong negative correlation, as required by Waugh (2010).[18] This is reported in Figure 8.

6.4 Counterfactuals

Using the calibrated version of the model, we can quantify the effects that changes in the per-shipment costs have on distribution (shipping plus inventory) costs, on the quality / price of coun-

[17]Recall that we calibrated the only free parameter, a, by matching the share of non-zero bilateral trade flows.

[18]While Waugh [2010] does not specify what exactly these ad-valorem export costs are, our model provides one possible micro-foundation for them.

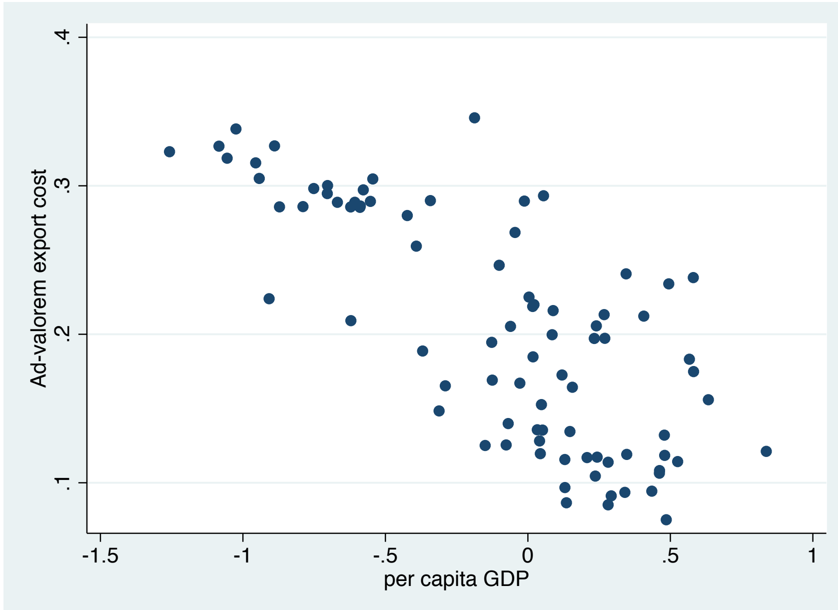

*Note:*Ad-valorem export costs are averages across all destinations. Values for per capita GDP are in logs and expressed as deviations from the mean.

Figure 8: Ad-valorem export cost and exporter income

tries' exports, on income and income inequality, and on welfare. It is worth noting that much emphasis has been placed in the implementation of trade facilitation policies, and most of these initiatives should impact per-shipment costs of exporting. To illustrate, in December of 2013, the WTO concluded the Trade Facilitation Agreement. The agreement "...contains provisions for expediting the movement, release and clearance of goods,...".[19] Among other measures, these negotiations have spurred a wave of interest in "Single Window" initiatives aimed at reducing the paper work and bureaucracy involved in international trade transactions. The quantification of these effects is, to our knowledge, the first general equilibrium treatment of the welfare implications of this type of policy.

Our first experiment looks at the impacts of a policy that changes (in almost all cases, lowers) the per-shipment costs of all countries in our sample to U.S. levels.[20] The effect on shipment size and frequency is shown in Figure 9. As expected, most countries increase shipment frequency and reduce shipment size in response to the now lower per-shipment costs they face, with the median exporter increasing shipment frequency by 14 percent and reducing shipment size by 48

[19] https://www.wto.org/english/tratop_e/tradfa_e/tradfa_e.htm

[20] Note that only two of the 85 countries in our sample have per-shipment costs lower than the U.S - Japan and Mauritius.

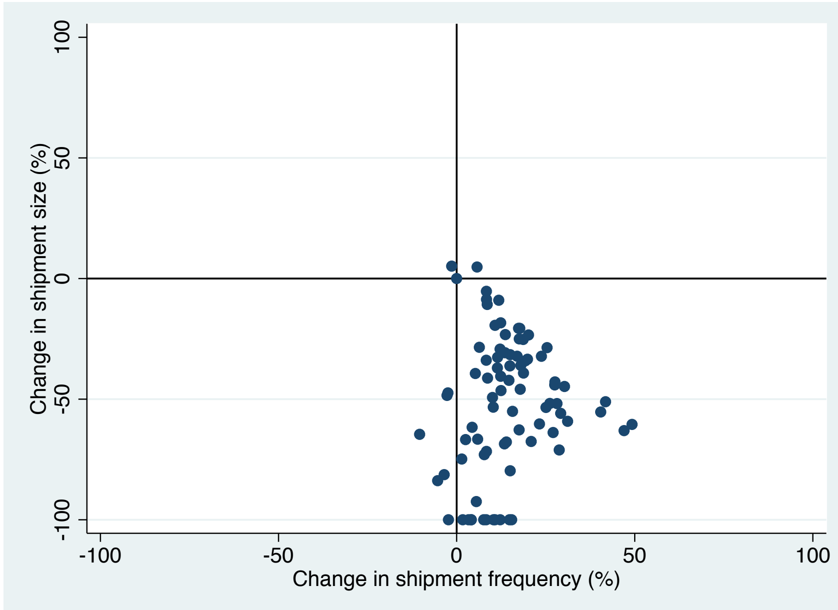

Note: In this exercise, we reduce per-shipment costs from the benchmark values to K_{US}.

Figure 9: Change in shipment frequency and size

percent. The same policy also causes the export price elasticity with respect to exporting country per capita GDP to drop by more than 50 percent (from 0.23 to 0.1). This is reported in Table 9. In essence, the elimination of the differences in per-shipment cost between rich and poor countries diminishes rich countries' comparative advantage in quality and leads to convergence in quality across countries.

Table 10 reports the effects of changing per-shipment costs to U.S. levels on countries' real income (i.e., welfare). Such a policy results in a 3.4% increase in the real income of the median country in the sample. There is considerable heterogeneity in welfare change across countries, however. For example, looking at the distribution of welfare change, the country at the 25th percentile experiences only a 1 percent change while the country at the 75th percentile experiences more than a 10 percent.

Furthermore, this effect is larger for poor countries, with the median change among the poorest 10 countries in the sample being a 27.8% increase in real income. The median change among the 10 richest countries is actually a 0.3% decline in real income. Figure 10 plots the resulting welfare effect of changing per-shipment costs to US levels against the countries' per capita GDP and confirms the heterogenous effects across the country income distribution. Finally, this policy also reduces inequality across countries; Table 11 reports that income inequality across countries is reduced by between 20 and 60 percent depending on how it is measured.

	$K_i = K_{US}$	$K_i = K_{region}$	$K_i = 0$
Real Income:			
Median country	3.4%	3.7%	9.5%
Median rich country	-0.3%	-0.3%	0.1%
Median poor country	27.8%	18%	67.1%

Table 10: Income Effects of Counterfactual Exercises

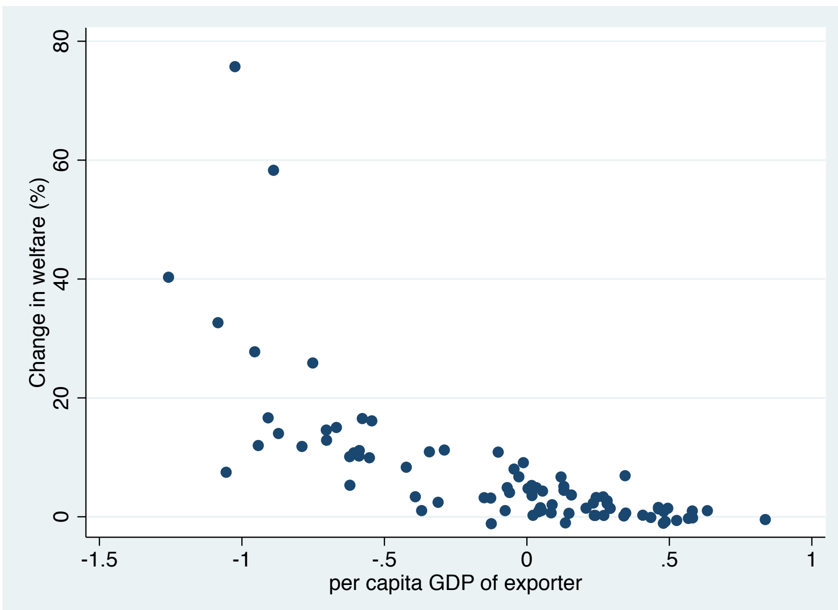

Note: In this exercise, we reduce per-shipment costs from the benchmark values to K_{US}.

Figure 10: Welfare Effect of $K_i = K_{US}$

Lastly, Table 12 reports the tariff-equivalent of the export specific per-shipment costs. For the benchmark scenario, the tariff-equivalent of per-shipment costs for the median country is equal to 20%. Once again, there is considerable heterogeneity in this number. The median tariff equivalent for the 10 richest countries is 14% while for the 10 poorest countries, this value is 32%. The policy that changes per-shipment costs to US levels reduces the ad-valorem cost for the median country by more than half. The reductions are similar for the poorest and the richest countries as well.

For comparison purposes, Anderson and van Wincoop [2004] estimates a 170% total trade cost for the representative rich country, almost half of it being due to local distribution costs.

	Benchmark	$K_i = K_{US}$	$K_i = K_{region}$	$K_i = 0$
$Var[Log(Income)]$	0.46	0.19	0.23	0.15
y_{90}/y_{10}	4.1	3.4	3.6	2.8

Note: y_{90}/y_{10} denotes the ratio of the 90th to the 10th percentile real wage.

Table 11: Inequality Effects of Counterfactual Exercises

	Benchmark	$K_i = K_{US}$	$K_i = K_{region}$	$K_i = 0$
Tariff equivalent:				
Median country	20%	12%	14%	0%
Median rich country	14%	10%	9%	0%
Median poor country	32%	25%	29%	0%

Note: The tariff equivalent of distribution cost is defined in Section 4 as $(e^{\delta/N_{ij}} - 1)$.

Table 12: The tariff equivalent of distribution cost

Finally, we can decompose the welfare effects of a change in per-shipment costs into that due to changes in shipment costs, in inventory costs, and in direct utility (consumption changes due to changes in quality adjusted prices). The effect on shipment costs can be further decomposed into the part directly due to the change in per shipment cost, i.e., holding fixed the number of shipments, and the part due to the implied changes in the number of shipments. In a similar way, the effect on inventory cost (IC) can be decomposed into the part directly due to the change in the number of shipments and the part due to changes in the amount traded. The latter accounts for the fact that both physical quantities and quality (and thus prices) endogenously adjust to the new per-shipment costs.

Table 13 reports the decomposition for the policy that changes per shipment costs in all countries to US levels. The table reports results for the country experiencing the median welfare

	Total	Delivery Costs				Direct Utility
		Shipping Cost		Inventory Cost		
		$\frac{\partial(K\times N)}{\partial K}\|_N$	$\frac{\partial(K\times N)}{\partial K}\|_{N(K_{US})}$	$\frac{\partial(IC)}{\partial N}\|_{X_{ij}}$	$\frac{\partial(IC)}{\partial X_{ij}}\|_{N(K_{US})}$	
Real Income:						
Median country	3.4%	0.10%	-0.09%	0.06%	-0.05%	3.38%
Median rich country	-0.3%	0.10%	-0.05%	0.05%	-0.02%	-0.36%
Median poor country	27.8%	0.22%	-0.21%	0.14%	-0.12%	27.5%

Note: IC is inventory cost and is defined in Section 4 as $IC = \tau_{ij} r_i e^{aq_{ij}} \left(N_{ij} S_{ij} - T_{ij}\right)$

Table 13: Decomposition of Welfare Effects (variable quality)

impact, where the median is either among all countries in the model, the 10 richest countries in the model and the 10 poorest countries in the model. For the country experiencing the median change among all countries, Mexico, the reduction in its per-shipment cost to US levels amounts to a reduction of almost 40%. This reduction has various impacts on overall distribution costs. First, there is a direct per-shipment cost saving amounting to 0.1% of real income due to the fact that Russian exporters now incur lower per shipment costs on every export shipment. Second, there is an indirect cost effect amounting to minus 0.09% of real income, resulting from the fact that the lower per-shipment cost induces Danish exporters to send more shipments. Counteracting this effect is a third effect due to the fact that the additional shipments result in smaller shipments and so reduced inventory costs. This reduced inventory cost effect amounts to an offsetting 0.06% increase in real income. The fourth and final effect occurs because the lower distribution costs result in a larger volume of exports and exports of a higher quality. This value effect leads to added inventory costs amounting to minus 0.05% of real income. Hence, for Mexico at least, the reduced per-shipment costs result in a distribution cost savings that is negligible.

The major impact of the reduction in per-shipment costs is clearly not the savings in distribution costs. Rather, it is the quality upgrading that reduced distribution costs permit. For the median country, this impact amounts to 3.38% of real income. For the median among the poor countries, it is substantial – 27.5% of real income. Observe that for the rich countries, the savings in shipping and inventory cost actually reduce the negative effect on welfare.

We confirm this fact in Table 14 that shows the effects of the same policy - changes in per-shipment costs to US levels - when quality is not allowed to vary. As the table shows, the total effect of the policy change is much smaller. For the median country, the welfare gains are about a

	Total	Delivery Costs				Direct Utility
		Shipping Cost		Inventory Cost		
		$\frac{\partial(K\times N)}{\partial K}\|_N$	$\frac{\partial(K\times N)}{\partial K}\|_{N(K_{US})}$	$\frac{\partial(IC)}{\partial N}\|_{X_{ij}}$	$\frac{\partial(IC)}{\partial X_{ij}}\|_{N(K_{US})}$	
Real Income:						
Median Country	0.4%	0.04%	-0.02%	0.02%	-0.00%	0.36%
Median rich country	0.1%	0.10%	-0.04%	0.06%	-0.00%	-0.02%
Median poor country	1.3%	0.22%	-0.08%	0.15%	-0.00%	1.01%

Note: IC is inventory cost and is defined in Section 4 as $IC = \tau_{ij} r_i e^{a q_{ij}} \left(N_{ij} S_{ij} - T_{ij}\right)$

Table 14: Decomposition of Welfare Effects (fixed quality)

tenth of what they would be if quality were flexible. As before, there is a lot of heterogeneity in the welfare gains. Interestingly, when quality is fixed, the median among the rich countries actually gains from the reduction in per-shipment cost, with most of this gain coming from savings in distribution costs. This implies that the welfare reduction that the rich countries experience when quality is endogenous is being driven entirely by quality change.

We run two more counterfactual exercises. Recall that in the model, the per-shipment cost faced by German firms exporting to France are lower than what they face when exporting to Brazil. In the first counterfactual, we assume that the per-shipment cost faced by firms in a country is the *within-region* per-shipment cost. The rationale behind such an exercise is that reducing per-shipment costs to the level of the U.S. might be difficult for many poor countries with weak institutions, while setting the per-shipment costs at the within-region level, although modest, is feasible. As shown in Tables 11, 10, 11 and 12, the effect of this reduction on prices, welfare, inequality and distribution costs are actually quite similar to the previous counterfactual exercise. This suggests that even a small change in per-shipment costs can have a large impact.

In the second counterfactual, we consider a scenario that eliminates per-shipment costs at exporting countries altogether. As expected, this exercise produces more pronounced effects, leading to less price sensitivity to exporter's level of development, higher real incomes, lower inequality across countries and lower tariff equivalent of distribution costs.

7 Conclusion

For a long time economists abstracted from the nuances of distribution systems in order to focus on the key determinants of trade and incomes, such as productivity, capital accumulation and the sources of comparative advantage. Distribution systems were part of the background, often represented by iceberg-type transportation costs and fixed costs of selling abroad. Recent evidence, however, has uncovered that one cannot understand incomes and trade flows without understanding how goods move from producers to consumers. This paper confirms these findings and shows that the ability to move goods in the domestic market is a key piece for understanding specialization in quality and its impacts on income differences across countries.

References

Armen Albert Alchian and William Richard Allen. *Exchange and production: theory in use*. Wadsworth Belmont, CA, 1969.

George Alessandria, Joseph Kaboski, and Virgiliu Midrigan. Inventories, lumpy trade, and large devaluations. *American Economic Review*, 100(5):2304–39, 2010.

James E. Anderson and Eric van Wincoop. Trade costs. *Journal of Economic Literature*, 42: 691–751, 2004.

Kenneth J Arrow, Theodore Harris, and Jacob Marschak. Optimal inventory policy. *Econometrica*, pages 250–272, 1951.

Richard Baldwin and James Harrigan. Zeros, quality and space: Trade theory and trade evidence. *American Economic Journal: Microeconomics*, 3(2):60–88, 2011.

Robert J Barro and Jong Wha Lee. A new data set of educational attainment in the world, 1950–2010. *Journal of development economics*, 104:184–198, 2013.

Jeffrey H Bergstrand. The heckscher-ohlin-samuelson model, the linder hypothesis and the determinants of bilateral intra-industry trade. *The Economic Journal*, 100(403):1216–1229, 1990.

Bruce A Blonigen and Wesley W Wilson. Port efficiency and trade flows*. *Review of International Economics*, 16(1):21–36, 2008.

Thomas Chaney. Distorted gravity: the intensive and extensive margins of international trade. *The American Economic Review*, 98(4):1707–1721, 2008.

Pablo Fajgelbaum, Gene M Grossman, and Elhanan Helpman. Income distribution, product quality, and international trade. *Journal of Political Economy*, 119(4):721–765, 2011.

Robert C Feenstra and Hong Ma. Trade facilitation and the extensive margin of exports. *Japanese Economic Review*, 65(2):158–177, 2014.

Robert C Feenstra and John Romalis. International prices and endogenous quality. *The Quarterly Journal of Economics*, 129(2):477–527, 2014.

Harry Flam and Elhanan Helpman. Vertical product differentiation and north-south trade. *The American Economic Review*, 77(5):810–822, 1987.

Robert E Hall and Charles I Jones. Why do some countries produce so much more output per worker than others?. *The Quarterly Journal of Economics*, 114(1):83–116, 1999.

Juan Carlos Hallak and Peter K. Schott. Estimating cross-country differences in product quality. *The Quarterly Journal of Economics*, 126(1):417–474, 2011.

Cecília Hornok and Miklós Koren. Per-shipment costs and the lumpiness of international trade. *Review of Economics and Statistics*, 97(2):525–530, 2015.

David Hummels. Toward a geography of trade costs. GTAP Working Papers 1162, Center for Global Trade Analysis, Department of Agricultural Economics, Purdue University, 2001.

David Hummels. Transportation costs and international trade in the second era of globalization. *The Journal of Economic Perspectives*, 21(3):131–154, 2007.

David Hummels and Peter J. Klenow. The variety and quality of a nation's exports. *American Economic Review*, 95(3):704–723, June 2005.

David Hummels and Alexandre Skiba. Shipping the good apples out? an empirical confirmation of the alchian-allen conjecture. *Journal of Political Economy*, 112(6):1384–1402, 2004.

Robert C Johnson. Trade and prices with heterogeneous firms. *Journal of International Economics*, 86(1):43–56, 2012.

Amit Khandelwal. The long and short (of) quality ladders. *The Review of Economic Studies*, 77: 1450–1476(27), 2010.

Andreas Kropf and Philip Sauré. Fixed costs per shipment. *Journal of International Economics*, 92(1):166–184, 2014.

Staffan B Linder. An essay on trade and transformation. 1961.

James R Markusen. Explaining the volume of trade: an eclectic approach. *The American Economic Review*, pages 1002–1011, 1986.

Christian Volpe Martincus, Jerónimo Carballo, and Alejandro Graziano. Customs. *Journal of International Economics*, 96(1):119–137, 2015.

Kiminori Matsuyama. A ricardian model with a continuum of goods under nonhomothetic preferences: Demand complementarities, income distribution, and north-south trade. *Journal of Political Economy*, 108(6):1093–1120, 2000.

Stephen J Redding and Matthew A Turner. Transportation costs and the spatial organization of economic activity. Technical report, National Bureau of Economic Research, 2014.

Michael Waugh. International trade and income differences. *The American Economic Review*, 100(5):2093–2124, 2010.

Appendix A

Table 15: Per shipment costs

Country	Cost U.S. Dollars	Country	Cost U.S. Dollars	Country	Cost U.S. Dollars
Japan	2,641	Barbados	4,996	Zambia	8,095
United States	2,834	Finland	5,045	Panama	8,149
Morocco	2,840	Iceland	5,143	Tajikistan	8,404
Sweden	2,918	Bolivia	5,282	Mexico	8,515
Hungary	3,000	Philippines	5,330	Colombia	8,697
Germany	3,022	Netherlands	5,333	Honduras	8,749
Switzerland	3,215	New Zealand	5,428	Slovenia	8,832
Tunisia	3,216	Australia	5,440	El Salvador	9,259
Bulgaria	3,290	Poland	5,521	Russia	10,351
South Korea	3,403	Bangladesh	5,549	Paraguay	10,375
China	3,506	Ireland	5,589	Nepal	10,890
Singapore	3,575	Greece	5,593	Guatemala	11,115
France	3,623	Indonesia	5,594	Mongolia	11,253
United Kingdom	3,654	Thailand	5,733	Uruguay	11,299
Egypt	3,792	Israel	5,810	Albania	11,927
Malta	3,907	Croatia	5,997	Swaziland	12,361
Lithuania	4,076	Latvia	6,162	Kenya	12,566
Italy	4,223	South Africa	6,429	Venezuela	12,612
Spain	4,410	Costa Rica	6,496	Jordan	12,699
Cyprus	4,414	Peru	6,516	Estonia	13,092
Belgium	4,421	India	6,598	Pakistan	14,112
Austria	4,690	Ecuador	6,783	Ukraine	15,883
Czech Republic	4,702	Portugal	6,916	Iran	15,985
Romania	4,743	Sri Lanka	7,005	Cambodia	21,265
Denmark	4,774	Norway	7,050	Ghana	22,211
Turkey	4,835	Dominican Republic	7,301	Kazakhstan	25,546
Canada	4,921	Vietnam	7,615	Mauritania	53,281
Malaysia	4,995	Brazil	7,927	Zimbabwe	104,225

Appendix B

Proof of Proposition 1. Note that we can write the implicit solution to the inventory management problem as:

$$e^{\delta/N}\Big(\frac{1}{\delta}-\frac{1}{N}\Big)=\frac{1}{\delta}-\chi.$$

The variable χ is increasing in K and declining in x. Differentiating the above equation with respect to χ, we have

$$\frac{dN}{d\chi}=-\frac{N^3}{\delta e^{\delta/N}}.$$

It follows that N is increasing in x and decreasing in N. □

Proof of Proposition 3. The FOB price depends on

$$\phi(s)=e^s+1-\frac{e^s-1}{s},$$

where $s=\delta/N$. Differentiating with respect to s yields

$$\frac{\partial\phi(s)}{\partial s}=\frac{1}{s^2}[se^s(s-1)+(e^s-1)].$$

Now, the value of the above derivative, as $s\rightarrow 0$ is not defined. Hence, we have to apply L'Hospital's rule, which yields

$$\begin{aligned}\frac{\partial\phi(s)}{\partial s}\Big|_{s\rightarrow 0}&=\frac{e^s(s+1)}{2}\Big|_{s\rightarrow 0}\\&=\frac{1}{2}.\end{aligned}$$

In the neighbourhood of $s=0$, the function $\phi(s)$ is increasing in s. Now, because poorer countries have a lower shipment frequency N, s is higher for these countries. Therefore, $\phi(s)$, and hence p^{FOB} is higher for poorer countries. □

Proof of Proposition 2. Note that d_{ij} can be written as $d_{ij}=f(s)\times A$, where A is independent of both x_{ij} and K_i, and

$$f(s)=\frac{e^s-1}{s},$$

where $s=\delta/N$. As long as $s>0$, $f(s)>0$ and (i) follows. Differentiating $f(s)$ with respect to s yields

$$\frac{\partial f(s)}{\partial s}=\frac{1}{s^2}[e^s(s-1)+1].$$

Using a Taylor Series expansion, $e^s > s + 1$. Replacing this in the above equation, we have

$$\begin{aligned} \frac{\partial f(s)}{\partial s} &> \frac{1}{s^2}[(s+1)(s-1)+1] \\ &= 1. \end{aligned}$$

Hence, when s increases, so does d_{ij}. But s is decreasing in x_{ij} and increasing in K_i. Hence, (ii) and (iii) follow. □

Proof of Proposition 4. Note that q_{ij} can be written as $q_{ij} = g(s) \times B$, where B is exogenously given and

$$g(s) = \frac{1 - e^{-s}}{s},$$

where $s = \delta/N$. Differentiating $g(s)$ with respect to s yields

$$\frac{\partial g(s)}{\partial s} = \frac{1}{s^2}[e^{-s}(s+1) - 1].$$

From $e^s > s + 1$, we have

$$e^{-s}(s+1) - 1 < 0.$$

Therefore, $\frac{\partial g(s)}{\partial s} < 0$. Now, conditional on q, an increase in K reduces N, and hence, reduces $g(s)$. It can be shown that $g(s)$ is a concave function of q. It follows then that an increase in K will reduce q. □

Appendix C

In this section, we examine the relation between the number (and size) of shipments and (i) per shipment costs, and (ii) depreciation rate of goods when a firm solves a two period dynamic inventory management problem under uncertainty. Demand in the two periods, q_1 and q_2, is uncertain and is drawn independently from a distribution $\Phi(q)$ with pdf $\phi(q)$. We solve the firm's problem recursively.

Suppose the firm starts period 2 with an inventory of x. It has to decide how much inventory to hold over this period, y_2, which will also determine how much, if at all, it has to order at the begin of period 2. The good does not depreciate between periods 1 and 2. At the end of period 2, the firm can sell any excess inventory at a price of h per unit. h captures how fast goods depreciate, with lower h corresponding to higher depreciation rate. Each unit normally sells for p, costs c and the per shipment cost is K. Net expected profit in period 2 is given by

$$E[\pi_2] = \begin{cases} E[r_2] - c(y_2 - x) - K & \text{if } y_2 > x, \\ E[r_2] & \text{otherwise.} \end{cases}$$

where $E[r_2]$ is expected revenue in period 2. The firm sells an amount equal to period 2 demand, q_2, if it has enough inventory. Otherwise, it just sells y_2. It also earns h on every unit of excess inventory $y_2 - q_2$. Therefore,

$$E[r_2] = p \int_0^1 min[y_2, q]\phi(q)dq + h \int_0^{y_2} (y_2 - q)\phi(q)dq.$$

Assuming that $\Phi(q)$ follows a uniform distribution with support $[0, 1]$, we have

$$E[\pi_2] = py_2 - \frac{1}{2}py_2^2 - c(y_2 - x) - K.$$

Maximizing π_2 with respect to y_2, we have $y_2 = \frac{p-c}{p-h}$, i.e., period two optimal inventory is independent of x. Note that we are assuming here that $y_2 > x$. If $y_2 < x$, the firm chooses y_2 to maximize $E[r_2]$ only. Then, the optimal inventory in period 2 is $y_2 = 1$.

In period 1, the firm starts with nothing. It has to choose how much inventory to hold over this period, y_1. Whatever is left unsold, $y_1 - q_1$, carries over to period 2. As before, the firms's period 1 expected profit is given by

$$E[\pi_1] = py_1 - \frac{1}{2}py_1^2 - cy_1 - K.$$

where we have assumed that the firm starts with zero inventory. The expected period 2 cost is given by $Prob(x < y_2).E[c(y_2 - x) + K|x < y_2]$. Replacing x with $y_1 - q_1$ and simplifying, the expected cost is $(1 - y_1 + y_2)(\frac{c(y_2-y_1+1)}{2} + K)$. Assuming that the firm does not discount its period 2 profits, the expected total profits of the firm are given by

$$E[\pi_1 + \pi_2] = py_1 - \frac{1}{2}py_1^2 - cy_1 - K + (1 - y_1 + y_2)\big(\frac{c(y_2 - y_1 + 1)}{2} + K\big) + \xi.$$

where ξ is independent of y_1. Maximizing the above expression with respect to y_1 yields

$$y_1 = 1 + \frac{K + c\frac{h-c}{p-h}}{p + c}.$$

Because y_1 is the size of the shipment in period one, a higher K increases the size of the shipment. Furthermore, shipments are ordered in period 2 with probability $Prob(q_1 > y_1 - y_2) = 1 - y_1 + y_2$. Because y_1 is increasing in K while y_2 is independent of K, this probability is decreasing in K. Applying the Law of Large Numbers, the number of period two shipments, and hence the total number of shipments, ordered when K is high is low.

At the same time, goods with low h (higher depreciation rate) have a low y_1 and accordingly, higher shipments on average.