

Inventory Management, Product Quality, and Cross-Country Income Differences[†]

By BERNARDO S. BLUM, SEBASTIAN CLARO, KUNAL DASGUPTA,
AND IGNATIUS J. HORSTMANN*

Previous research has documented that export shipments are “lumpy”—exporters make infrequent and relatively large shipments to any given export destination. This fact has been interpreted as implying that fixed, per shipment cost and inventory management decisions play a key role in international trade. We document here that exports from poor countries are considerably more lumpy—have higher fixed per shipment cost—than those from rich countries. Using a model of trade with inventory management, we estimate that the country at the ninetieth percentile of the distribution of per shipment costs has almost three times higher costs than the one at the tenth percentile. We show that these per shipment cost differences have a reduced-form representation given by an ad valorem trade cost that varies with export country income (as in Waugh 2010). A calibrated version of the model that incorporates these estimates and allows for endogenous product quality reveals that cross-country differences in per shipment costs explain almost 40 percent of the observed cross-country differences in income. It also shows that policies that lower per shipment costs can lead to significant welfare gains, mainly due to induced quality upgrading. (JEL F12, F14, F43, G31, O16, O19)

Recent analyses of micro-level international trade data have shown that export shipments are relatively infrequent and large. Researchers have interpreted this fact as indicative of the important, and previously unrecognized, role that inventory management decisions play in international trading behavior (see, for instance, (Alessandria, Kaboski, and Midrigan 2010—henceforth, AKM)). Specifically, researchers have interpreted the lumpiness of export shipments as indicative of a

*Blum: Rotman School of Management, University of Toronto, 105 St. George Street, Toronto, ON M5S3E6, Canada (email: bblum@rotman.utoronto.ca); Claro: Instituto de Economía, Pontificia Universidad Católica de Chile, Av. Vicuña Mackenna 4860, Macul, Santiago, Chile (email: sclaro@uc.cl); Dasgupta: Department of Economics, University of Toronto, 150 St. George Street, Toronto, ON M5S 3G7 (email: kunal.dasgupta@utoronto.ca); Horstmann: Rotman School of Management, University of Toronto, 105 St. George Street, Toronto, ON M5S3E6, Canada (email: ihorstmann@rotman.utoronto.ca). We thank the referees and Heski Bar-Isaac, Nate Baum-Snow, Marco Gonzalez-Navarro, Sebastian Heise, Matt Mitchell, Peter Morrow, Affonso Celso Pastore, Diego Restuccia, David Rivers, Olga Timoshenko, Trevor Tombe, Dan Trefler, and Xiaodong Zhu for comments. Bernardo S. Blum and Ignatius J. Horstmann thank the Social Sciences and Humanities Research Council of Canada for financial support (grant 410-2011-0475).

[†]Go to <https://doi.org/10.1257/mac.20170080> to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

world in which international transactions involve both a fixed, per shipment cost and a variable inventory cost, with the observed size and frequency of shipments the outcome of exporters economizing on these two cost margins. Seemingly unrelated to this line of research, and in a study of aggregate international trade flows and prices, Waugh (2010) argues that the observed patterns of trade and prices can only be understood if trade costs are asymmetric, even between two trading partners. Waugh (2010) further shows that these trade cost asymmetries disproportionately affect poor countries, so that one-third of the differences in real income per capita across countries can be attributed to the high trade costs faced by poor, but not rich, trading partners.

In this paper, we demonstrate that these two, apparently unrelated sets of observations can be reconciled within a single model of trade and inventory management. At the heart of this reconciliation is a new finding on differences between the degree of lumpiness in export shipments from poor countries relative to rich countries. Using shipment-level data for Chile, we show for the first time that, controlling for export volume, countries with higher per capita income (rich countries) have more frequent shipments of smaller size to Chile than do lower per capita income countries (poor countries). This is true even when the rich and the poor countries sell the same HS eight-digit product to the same Chilean importer. In addition, we show that the product exported by the rich country commands a higher price. In short, we establish that the level of a country's development, the frequency of trade transactions, and the quality/price of a country's exports are all correlated.

Based on these observations, we construct a multicountry, general equilibrium model of trade and inventory management in which fixed, per shipment costs vary across countries, and both the quality of an exporter's product and the size and frequency of shipments are determined endogenously. To maintain analytical tractability, and in contrast to AKM, our model assumes that all decisions are made in continuous time and that both demand and supply are deterministic (i.e., the model is essentially a standard Baumol-Tobin inventory model). These assumptions allow us to derive analytically an "all-in" cost of distribution—production, shipping, and inventory costs—and show that this cost has an "as if" ad valorem trade cost representation. This is important because it shows that, even in the face of inventory management concerns, a reduced-form, ad valorem trade cost assumption can consistently represent all trading costs. Because per shipment costs vary across export countries, it also shows that an asymmetric trade cost specification (as in Waugh) is necessitated by distribution and inventory management concerns. Finally, it also implies that quality choice must vary systematically with the cross-country variation in per shipment costs as an optimal response to inventory management problems: Countries with low per shipment costs make more frequent, smaller shipments, have lower inventory costs as a result, and so can more easily bear the increases in inventory costs associated with increased quality. In essence, low per shipment cost countries have a comparative advantage in the production and distribution of products that are most affected by inventory costs, namely, the high value/quality products.

Using our model, we generate estimates of the fixed per shipment export costs and their "as if" ad valorem representations across a sample of 74 countries.

These estimates employ only the shipment-level export data. Identification is obtained from variation, across export countries, in the frequency and size of shipments to Chile. We find that estimated per shipment cost vary widely across countries, with the ninetieth percentile value of the distribution of per shipment costs being almost three times larger than the tenth percentile value. The median value of the tariff-equivalent representation of the per shipment and inventory management costs equals 28 percent. We also find that our estimated per shipment costs are closely related to per capita incomes, even though our estimation procedure does not impose this relationship. The estimated fixed costs are unrelated to the physical distance between countries, however, suggesting that our estimates do indeed capture some fixed shipping cost component and are not spuriously capturing variable trade costs. Finally, we show that less than a quarter of the variation in the per shipment costs identified from the structure of our model and shipment data can be explained by the World Bank survey measures on the costs of exporting a container. This indicates the existence of important unobserved fixed per shipment costs of trading.

With our per shipment cost estimates in hand, we calibrate our model for a sample of 68 low-, medium-, and high-income countries to get some sense of how significant shipping and inventory costs are in explaining cross-country differences in product quality/price and incomes. The calibrated model explains about a quarter of the observed elasticity of export prices with respect to per capita income and close to 40 percent of the income inequality between countries observed in the data. These results are consistent with Waugh (2010), who shows that asymmetric trade costs, with the asymmetry being export-country specific, explain both the trade flow and price data, and a sizable part of the observed income differences across countries.¹

Finally, using the calibrated model, we quantify the effects of policies that aim to enhance the efficiency of countries' export systems. We study, in particular, the impact of a policy that makes the per shipment cost of exporting in all countries equal to the estimated US per shipment cost. We find that this change reduces income inequality across countries by 27 percent. This result is in keeping with Waugh who finds that the elimination of trade cost asymmetries reduces cross-country income inequality by 34 percent. The world welfare gains from this policy are about 1.6 percent but are highly unevenly distributed. Poor countries, having much higher per shipment costs, gain much more than do the rich countries. The median country in the sample experiences a welfare gain of 13.7 percent. Finally, we find that virtually all of these welfare gains are a consequence of the quality upgrading that takes place following the lowering of distribution costs.

This last result highlights the importance of understanding the structure of trade cost asymmetries. In our model, variable trade costs are endogenous and depend, among other things, on the size and frequency of transactions firms choose to make. For this reason, a reduction in per shipment cost will affect variable trade costs differently for different countries. For poor countries that export little—they have low productivity and lower quality products—inventory costs do not loom large

¹ In this way, our model and the empirical evidence supporting it provide a micro-founded mechanism for Waugh's asymmetric trade costs.

in their decision making. Consequently, a reduction in per shipment cost for poor countries, quality fixed, has little effect on the size and number of shipments, and therefore has little effect on variable trade costs, exports, and incomes. Put simply, a reduction in per shipment costs to the level of rich countries eliminates almost none of the variable trade cost asymmetry when quality is fixed. As quality upgrading occurs, however, poor countries export more and can now exploit the benefits of lower per shipment costs by increasing the number of shipments and so dramatically reducing effective variable trade cost. This leads to the predicted increase in real income.

The notion that trade frictions have an important per shipment component is in line with a literature that documents systematic patterns in the size and frequency of international shipments (Kropf and Sauré 2014; Hornok and Koren 2015; Békés et al. 2017). We contribute to this literature by linking these patterns to comparative advantage in the production of quality and by quantifying the effect of these frictions on cross-country income inequality. In this way, our work also contributes to the literature linking trade costs to the quality of exports. The Alchian-Allen hypothesis states that per unit trade costs reduce the relative price of high-price (high-quality) varieties, and so countries ship out the good apples to high-cost destinations (see, for example, Alchian and Allen 1964, Hummels and Skiba 2004, Martin 2012, and Lugovskyy and Skiba 2016). In our model, the combination of per shipment costs and inventory management costs create an “as if,” ad valorem trade cost and push in the direction of increasing the relative price of high-quality products. As a result, countries with lower per shipment costs export higher quality products.²

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 III and IV develop a full, general equilibrium model of trade and distribution. These sections provide the mechanism linking exporters’ choices of product quality, shipment size, and frequency. Section V provides a calibrated version of the model used to quantify the extent to which the mechanism proposed by the model can generate the export price variation observed in data as well as the implications of the mechanism for cross-country differences in per capita incomes. It also contains counterfactual exercises to quantify the impacts of changes in the efficiency of distribution systems on exporting countries. Finally, it provides robustness results with respect to alternative modeling assumptions and parameter values.

I. 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

²Our work is also related to Evans and Harrigan (2005) who argue that products for which timely delivery is important will be produced by countries that are close to the source of final demand.

TABLE 1—SUMMARY STATISTICS: CHILEAN IMPORT DATA (2006)

Chilean imports (000's US\$)	24,251,469
Number of importers	19,238
Number of HS eight-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

Note: Excludes mineral fuel imports and imports by individuals.

TABLE 2—DISTRIBUTION OF KEY VARIABLES: CHILEAN IMPORTS (2006)

	Imports (000's US\$) (1)	Number of importers (2)	Number of HS8 codes (3)	Number of source countries (4)	Number of shipments (5)
<i>Panel A. Distribution over importers: 19,238 firms</i>					
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
<i>Panel B. Distribution over HS eight-digit products: 6,587 codes</i>					
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
<i>Panel C. Distribution over importer-HS eight-digit product pairs: 290,348 pairs</i>					
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

analysis, we exclude import transactions in mineral fuels (HS27),³ and shipments in which the importer is an individual, as opposed to a firm. Table 1 shows summary statistics for these data.

Table 2 begins to describe some of the features of interest in the data. Panel A of the table shows the per importer distribution of total imports (column 1), number of

³Mineral 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.

HS eight-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 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 eight-digit products, the median firm imports 5 products, while the top 1 percent of importers purchase 156 products. Similarly, if we focus on source country, the median firm imports products from 2 countries while the top 1 percent of importers purchase products from more than 20 countries.

Panel B of Table 2 shows the per HS eight-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 eight-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, panel C of Table 2 shows distributions over importer-HS eight-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 eight-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.⁴

Table 3 shows characteristics of importer-HS eight-digit product pairs by the number of countries from which an importer buys the HS eight-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 2 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 \$2 in imports within this group of importer-product pairs, there is a 53 percent chance that these \$2 will come from the same country, and thus a 47 percent chance they will come from different countries.

⁴ Antràs, Fort, and Tintelnot (2017) show that the majority of US importers buy a given product from a single country only. This is certainly the case in our data as well. For instance, they report that the average, the median, and the ninety-fifth percentile US importer buys a given product, on average, from 1.11, 1.03, and 1.78 countries, respectively. The equivalent numbers for Chilean importers are 1.15, 1.00, and 1.75. Yet, around 20 percent of the importer-product pairs have the importer buying the same product from multiple countries, and these importer-product pairs account for close to 70 percent of Chilean imports.

TABLE 3—CHARACTERISTICS OF IMPORTER-HS EIGHT-DIGIT PRODUCT PAIRS

Number of source countries (1)	Number of importer-HS8 pairs (2)	Share of importer-HS8 pairs (3)	Share of imports (4)	HHI imports across countries (5)	Mean absolute deviation 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

Note: Column 6 reports the mean absolute deviation from the mean per capita income of the countries the importer buys the product from.

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 Mexico and Germany—roughly US\$20,000.

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 after controlling for import volume and the identities of the importer and exporter.

II. 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 i by Chilean importer l (V_{ihl}) as

$$V_{ihl} = N_{ihl} \times \frac{\sum_{k=1}^{N_{ihl}} s_{ihl}(k)}{N_{ihl}} = N_{ihl} \times \bar{s}_{ihl},$$

where $s_{ihl}(k)$ is the value of the k th shipment carried out by the importer, N_{ihl} is the number of shipments the importer carries out when purchasing product h from country i , and thus $\bar{s}_{ihl}$ is the average value of these shipments. This average shipment value can be further decomposed as the product of the average physical

quantity shipped ($\bar{q}_{ihl}$) and the weighted average per unit price of the goods in the shipment ($\hat{p}_{ihl}$):

$$\bar{s}_{ihl} = \frac{1}{N_{ihl}} \sum_{k=1}^{N_{ihl}} (q_{ihl}(k) \times p_{ihl}(k)) = \hat{p}_{ihl} \times \bar{q}_{ihl},$$

where

$$\hat{p}_{ihl} = \frac{\sum_{k=1}^{N_{ihl}} (q_{ihl}(k) \times p_{ihl}(k))}{\sum_{k=1}^{N_{ihl}} q_{ihl}(k)},$$

and

$$\bar{q}_{ihl} = \frac{\sum_{k=1}^{N_{ihl}} q_{ihl}(k)}{N_{ihl}}.$$

Therefore, we can write

$$V_{ihl} = N_{ihl} \times \bar{s}_{ihl} = N_{ihl} \times \hat{p}_{ihl} \times \bar{q}_{ihl}.$$

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 $z_{ihl} = [V_{ihl}, N_{ihl}, \bar{s}_{ihl}, \hat{p}_{ihl}, \bar{q}_{ihl}]$, we estimate the equation below:

$$\ln(z_{ihl}) = \delta_{hl} + \beta_1 \ln(gdp_i) + \beta_2 \ln(pcgdp_i) + \beta_3 \ln(dist_i) + \epsilon_{ihl},$$

where gdp_i , $pcgdp_i$, and $dist_i$ measure the exporting country’s GDP, per capita GDP, and physical distance to Chile, respectively. The terms δ_{hl} 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. Finally, when estimating the equations above, the error term (ϵ_{ihl}) is clustered at the country level.⁵

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-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 next

⁵ In a few cases, the quantity purchased by the same importer of a given HS eight-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 percent of the shipments only, we drop these shipments from the sample.

TABLE 4—IMPORTER-HS EIGHT-DIGIT PRODUCT FIXED EFFECT REGRESSIONS

Dependent variable:	$\ln(V_{iht})$ (1)	$\ln(N_{iht})$ (2)	$\ln(\bar{s}_{iht})$ (3)	$\ln(\bar{q}_{iht})$ (4)	$\ln(\bar{p}_{iht})$ (5)
$\ln(gdp_i)$	0.211 (0.023)	0.162 (0.018)	0.049 (0.009)	0.071 (0.019)	−0.022 (0.016)
$\ln(pcgdp_i)$	0.034 (0.039)	0.063 (0.024)	−0.029 (0.019)	−0.287 (0.042)	0.258 (0.029)
$\ln(dist_i)$	−0.066 (0.044)	−0.110 (0.024)	0.044 (0.025)	0.039 (0.048)	0.006 (0.034)
Importer-HS eight FE	Yes	Yes	Yes	Yes	Yes
Standard error cluster	Country	Country	Country	Country	Country
Observations	162,388	162,388	162,388	162,388	162,388
R^2	0.63	0.45	0.70	0.81	0.85

Note: Standard errors are in parentheses.

three columns in the same table show that exporting country characteristics affects the decision on how best to structure import shipments. Chilean importers buy more from larger countries (high GDP) because they receive more shipments from these countries and because they have larger shipments, both in values as well as in physical quantities. Finally, the last column shows that 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 affects the frequency and size of shipments importers make.⁶

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 subsample of the data including the 30 HS eight-digit codes with the largest number of importers sourcing that product from multiple countries.⁷ This is obviously a selected sample built with the intention of

⁶ We obtain similar results when we include observations with zero trade using the pseudo-maximum-likelihood estimation technique proposed in Santos-Silva and Tenreyro (2006). The results are not reported to conserve space but are available upon request. Results are also similar when we limit the analysis to those observations in which the importer receives at least two shipments when buying a product from a given country.

⁷ The reason we do this only on a subsample of HS eight-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 misspellings, a problem that plagues much of the empirical analysis in the literature using matched importer-exporter data.

TABLE 5—IMPORTER-HS EIGHT-DIGIT PRODUCT FIXED EFFECT REGRESSIONS. TOP 30 SAMPLE

Dependent variable:	$\ln(V_{ihl})$ (1)	$\ln(N_{ihl})$ (2)	$\ln(\bar{s}_{ihl})$ (3)	$\ln(\bar{q}_{ihl})$ (4)	$\ln(\bar{p}_{ihl})$ (5)
$\ln(gdp_i)$	0.225 (0.036)	0.180 (0.029)	0.045 (0.012)	0.085 (0.024)	−0.040 (0.019)
$\ln(pcgdp_i)$	0.249 (0.072)	0.212 (0.042)	0.038 (0.034)	−0.245 (0.060)	0.283 (0.032)
$\ln(dist_i)$	−0.164 (0.069)	−0.141 (0.047)	−0.024 (0.033)	−0.026 (0.061)	0.004 (0.037)
Importer-HS eight FE	Yes	Yes	Yes	Yes	Yes
Standard error cluster	Country	Country	Country	Country	Country
Observations	22,586	22,586	22,586	22,586	22,586
R^2	0.57	0.42	0.64	0.68	0.73

Note: Standard errors are in parentheses.

TABLE 6—SELLER’S MARGIN. TOP 30 SAMPLE

Dependent variable:	$\ln(N_{ihl})$ (1)	$\ln(N^x_{ihl})$ (2)	$\ln(\bar{n}_{ihl})$ (3)
$\ln(gdp_i)$	0.180 (0.029)	0.109 (0.016)	0.071 (0.015)
$\ln(pcgdp_i)$	0.212 (0.042)	0.083 (0.021)	0.129 (0.029)
$\ln(dist_i)$	−0.141 (0.047)	−0.047 (0.021)	−0.094 (0.039)
Importer-HS eight FE	Yes	Yes	Yes
Standard error cluster	Country	Country	Country
Observations	22,586	22,586	22,586
R^2	0.72	0.42	0.42

Note: Standard errors are in parentheses.

maximizing the power of our identification strategy, which is based on importers that buy the same HS eight-digit product from multiple countries. This subsample shares many of the main characteristics of the full sample. Importantly, the main findings described in Table 4 hold in this subsample as well, as Table 5 shows. The noticeable difference in this subsample is that importers not only structure their shipments differently when buying from richer countries, but also tend to buy more from richer countries.

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 i (N_{ihl}) into the number of export partners the importer transacts with and the average number of shipments it receives from each of its export partners:

$$N_{ihl} = N^x_{ihl} \times \bar{n}_{ihl},$$

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

We see that, when buying from larger countries, Chilean importers buy from more exporters, as well as have 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 eight-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 imports from rich and poor countries are structured differently, with imports from rich countries featuring higher quality products in smaller and more frequent shipments.

The existing literature has offered mechanisms to explain why rich countries export (and import) high-quality goods. It has also offered mechanisms for why the size and frequency of trade shipments may vary with product and importing country characteristics. In the next section, we provide a model of trade and inventory management in which exporting country characteristics jointly determine the quality of exports, the size, and the frequency of export shipments. By endogenizing the choice of export shipment size and frequency, as well as the quality of the product exported, the model shows that trade costs of a certain nature incurred at exporting countries can be a source of comparative advantage in quality for these countries, and drive the structure of the country's trade shipments.

III. A Model of Trade and Distribution

We consider a continuous-time, finite-horizon world in which, without loss of generality, time, t , is an element of the interval $[0, 1]$. In this world there are J countries labeled $1, 2, \dots, J$. In each country j , there is a single factor of production, labor. The total endowment of labor in country j is exogenous and denoted by L_j . The flow of labor supplied at date t , $l_j(t)$ is determined to meet demand for labor services at t . Each country has a single, monopolistically competitive manufacturing sector producing a differentiated intermediate good. Each firm in each country produces a different variety of the intermediate good using a production technology that is identical across firms as in Krugman (1980). Firms may also differentiate their product on a quality dimension. This quality choice is made once and for all at $t = 0$, but can be tailored to the particular market in which the intermediate is

sold (firms can customize their product quality to each local market). Finally, in each country there is a non-traded, final consumption good made from intermediate varieties and labor. Complete details are provided below.

A. Final Goods Sector

There is a continuum of consumers in each country who consume a single, non-traded final good. For simplicity, we assume that any consumer in a given country consumes only once in the period $[0, 1]$. The probability of a given consumer consuming at $t \in [0, 1]$ follows some known distribution with probability density function $\phi(t)$. This function is assumed to be the same across countries. By the Law of Large Numbers, the measure of consumers with a positive demand for the final good in country j at time t is simply $\phi(t) L_j$.

Consumer preferences are assumed identical both within and across countries. Utility of a representative consumer in country j is the quantity of the final good consumed:

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

where l_j^f and c_j^f are the amount of labor and the CES aggregate of the manufactured intermediates (both in per capita terms), respectively, used to produce the final good. The total amount of the CES aggregate used for final consumption at time t is given by

$$\phi(t) L_j c_j^f = \left[\sum_{i=1}^C \mathbf{1}_{ij} M_i q_{ij} [x_{ij}^f(t)]^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}},$$

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 ; M_i is the measure of country i firms; and $x_{ij}^f(t)$ is the quantity of a representative variety produced in country i and used at time t to produce the aggregate intermediate for final good consumption in country j . The variable q_{ij} is the (endogenous) quality of each intermediate variety.⁸ (Recall that quality, which is chosen once-for-all at date zero, can be different for different countries).

Given this specification, and denoting the purchase price of a representative country i variety sold in country j at time t by $p_{ij}(t)$, the quantity demanded of that variety by country j consumers at time t is given by⁹

$$(1) \quad x_{ij}^f(t) = \frac{q_{ij}^\sigma p_{ij}(t)^{-\sigma}}{P_j(t)^{1-\sigma}} \beta \phi(t) Y_j,$$

⁸ We conceive of quality as a measure of the efficiency units of production/consumption services embodied in a particular unit of an intermediate good. As will become apparent, cross-country quality variation will be valuable for explaining the price and income data.

⁹ We can instead assume that there is a non-traded sector that combines the manufactured aggregate with labor to produce the final good. In this case, the demand for intermediates will come from the non-traded sector.

where Y_j is aggregate income (and with balanced budget, expenditure) in country j earned over the entire interval $[0, 1]$ and $P_j(t)$ is the price index of the manufactured aggregate in country j at time t . This price index is given by the expression:

$$(2) \quad P_j(t)^{1-\sigma} = \sum_{i=1}^C \mathbf{1}_{ij} M_i q_{ij}^\sigma p_{ij}(t)^{1-\sigma}.$$

B. Intermediate Goods Sector

A manufacturer of an intermediate variety in country i produces using labor and a composite of all the available manufactured intermediates. We assume that this composite is the same as the one used for final good production. The production technology for intermediates is described by a constant returns to scale, Cobb-Douglas production function. For a given choice of quality, q_{ij} , by a representative firm in country i , output of the manufactured intermediate for sale in country j at time t is

$$z_{ij}(t) = \frac{A_i}{q_{ij}^{\theta_1}} \left[\frac{l_{ij}^m(t)}{\alpha} \right]^\alpha \left[\frac{c_{ij}^m(t)}{1-\alpha} \right]^{1-\alpha},$$

where $\theta_1 > 0$, and $l_{ij}^m(t)$, $c_{ij}^m(t)$ are the amounts of labor and the manufactured aggregate, respectively, used by a firm in country i in the production of output for country j at date t . The variable A_i captures total factor productivity (TFP). The share of value-added in intermediate good production, α , is assumed to be the same across countries.

Under these assumptions, the unit cost of producing a variety of quality q_{ij} in country i at time t is given by $r_i(t) q_{ij}^{\theta_1}$, where

$$(3) \quad r_i(t) = \frac{1}{A_i} w_i^\alpha P_i(t)^{1-\alpha},$$

and w_i is the wage rate in country i , which will be constant over the interval $[0, 1]$. In addition to the unit cost of production, the manufacturer also incurs a once-for-all, quality-dependent, sunk cost of establishing a production facility. This cost is given by $f q_{ij}^{\theta_2}$, where $f, \theta_2 > 0$. Under this specification, the variables θ_1 and θ_2 govern, respectively, the responsiveness of marginal and fixed cost to a change in product quality.

Finally, we assume that the cost for any manufacturer to buy directly from other manufacturers or sell directly to consumers and manufacturers is prohibitive. Instead, all shipments pass through a distribution sector. Because of inventory management issues in this sector, manufacturers will not, in general, sell continuously—shipments may be lumpy. It is to this sector, and these issues, that we turn to next.

C. Distribution Sector

Products from country i are sold in country j (including $j = i$) via a perfectly competitive distribution sector in j . There are three costs involved in the distribution of a product. First, there are the usual iceberg trade costs, so that one unit landing

in country j requires $\tau_{ij} > 1$ units of production in country i . Second, each time a shipment is sent from a producer in country i to a distributor in country j (where $j = i$ represents an internal shipment), a fixed shipping cost of K_i is incurred. This per shipment cost reflects such things as the administrative costs of sending a shipment either domestically or internationally, the efficiency of any underlying consolidation process, and other exit or supply chain distribution costs that do not depend on the size of the shipment. Third, if distributors choose to carry an inventory, there is an inventory depreciation cost, δ , that is proportional to shipment value. This cost may represent pure wastage or potential obsolescence. Together with the iceberg trade cost, this inventory cost determines the fraction of goods leaving the factory gate that “melts away” before reaching the customer (either an intermediate good producer or consumer). In all that follows, we assume that the value of δ is independent of quality.¹⁰ Of course, being proportional to value, higher priced/higher quality products will have a larger inventory cost in absolute terms.

Because the distribution sector is perfectly competitive, the purchase price (to both consumers and manufacturing customers) for a variety produced in country i and distributed in country j , $p_{ij}(t)$, is simply the factory gate (or what we will refer to as the wholesale) price in country i plus per unit shipping and inventory costs in j .¹¹ As a result, in making its production and shipping decisions, a manufacturer in country i will internalize all production and distribution (trade plus inventory) costs for j . This means that, given any price path $p_{ij}(t)$, a country i manufacturer shipping to country j will select a production and shipping policy to minimize total costs. In the next section, we define this cost minimization problem and solve for a manufacturer's equilibrium production and shipping policies as well as its quality and price choices.

IV. 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 must first solve the shipment problem. The solution to 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 path. It is to this shipping problem that we now turn.

¹⁰ We have been unable to find any literature suggesting that higher quality products have a higher depreciation rate (higher value of δ) than lower quality ones and can imagine arguments suggesting the opposite. High-quality items may be less subject to fashion whims than low-quality ones. Were one to assume that higher quality products come with higher values of δ , however, this assumption would directly imply that countries with low per shipment cost should export high-quality products.

¹¹ As suggested by a referee, we could alternatively assume that each manufacturer in country i either owns or contracts efficiently with a distributor in country j . In all cases, the purchase price for a good manufactured in country i and sold in country j at time t will be the same: the wholesale price at t plus unit shipping and inventory costs. For any of these arrangements, the essential feature is that the manufacturer actually sets the purchase price path, either directly in the case of ownership or efficient contracting, or indirectly via setting the wholesale price path in the case of a competitive distribution sector. Of course, in all cases, the wholesale price path must still be determined via a Nash equilibrium pricing outcome among the sellers in country i selling to country j .

A. Shipping Problem

To solve for a manufacturer's cost minimizing shipping policy, let $x_{ij}(t)$ be any arbitrary demand (from both consumers and other manufacturers) at date t in country j for a representative intermediate good manufacturer in country i . Given there is no uncertainty in our model, total demand is deterministic and is given by $x_{ij}(t) = x_{ij}^f(t) + M_j x_{ij}^m(t)$, where $x_{ij}^m(t)$ is the demand coming from a representative intermediate good producer.¹² The inventory of a representative country i intermediate good held by a country j distributor at date t , $m_{ij}(t)$, will decrease due to sales, $x_{ij}(t)$, and depreciation. Assuming that the distributor receives no shipments from country i at t , the law of motion of inventory is given by

$$(4) \quad \frac{dm_{ij}(t)}{dt} = -x_{ij}(t) - \delta m_{ij}(t).$$

When a shipment arrives in country j from country i in period t , inventory jumps up by $m_{ij}(t^+) - m_{ij}(t^-) > 0$. Letting $\mathcal{T}_h$ ($h = 1, 2, \dots$) denote the time periods when inventory is replenished (i.e., when new shipments arrive in country j from country i) and $c_{ij} = \tau_{ij} r_i q_{ij}^{\theta_1}$, the cost of each unit of inventory, then, the manufacturer's cost minimization decision is given by

$$(5) \quad D(m_{ij}) = \min_{m_{ij}(t), \mathcal{T}_h} \sum_h [K_i + c_{ij}(m_{ij}(\mathcal{T}_h+) - m_{ij}(\mathcal{T}_h-))],$$

subject to (4) and $m_{ij}(t) > 0$. Because the decision problem is a finite horizon problem, we assume above that the manufacturer's rate of time preference is one.

With a cost of K_i incurred each time a shipment is sent, the manufacturer faces a trade-off between how often to replenish inventory and how much to ship each time inventory is replenished. On the one hand, less frequent shipments mean that the manufacturer has to incur K_i fewer times, thus reducing costs. On the other hand, the distributor then has to hold a bigger inventory, on average, which leads to more wastage due to depreciation. This added wastage raises the inventory cost.¹³ For arbitrary paths $x_{ij}(t)$, the cost-minimizing shipping policy for a manufacturer does not generally permit a simple, closed-form solution. To obtain a closed-form solutions for this policy, we make two simplifying assumptions. First, we assume that $\phi(t)$ follows a uniform distribution on $[0, 1]$. This assumption implies that, at constant prices p_{ij} , final consumption demand for intermediate varieties is spread uniformly over the time interval. Second, we assume, as seems natural,

¹² With stochastic demand, the inventory management problem is typically not solvable analytically and numerical methods must be employed (see Alessandria, Kaboski, and Midrigan 2010). In Appendix C, we present a very simple two-period model with demand uncertainty that permits closed-form solutions. The main predictions of our model go through in this case.

¹³ Note the similarity of the distributor's problem with the standard problem of determining money demand. Essentially, an individual, in deciding how much money (cash) to hold, faces a trade-off. On the one hand, holding more cash implies fewer trips to the bank, where there is a cost per trip. On the other hand, holding more cash means forgoing interest paid by the bank. See Baumol (1952) and Tobin (1956) for early exposition of this problem, and Alvarez and Lippi (2009) for a recent extension.

that if manufacturers inventory intermediates, they face the same inventory cost as distributors. This latter assumption implies that, at constant prices, only the total amount purchased of intermediates over some time interval affects the manufacturer's cost but not the path of purchases: the manufacturer's cost minimizing path is indeterminate.¹⁴ In this case, we can assume, without loss of generality, that, when faced with constant prices, the manufacturer also distributes demand for intermediates uniformly over the entire time interval.¹⁵ Together, these two assumptions imply that total demand, $x_{ij}(t)$, is uniform on $[0, 1]$ should prices be constant. This fact simplifies the shipping problem considerably and, for the rest of this section, we maintain the constant price assumption. We show subsequently that this is the equilibrium pricing outcome.

Given constant prices and so a uniform distribution of total demand, the optimal shipping policy must satisfy two properties (Arrow, Harris, and Marschak 1951). First, a shipment from a manufacturer in i should arrive at a distributor in j just as inventories go to zero. Second, shipments of equal size at equal intervals are optimal. The policy is shown in Figure 1. The first property is a consequence only of the deterministic demand assumption: In the absence of uncertainty, there is no reason for the distributor to hold inventory beyond what is required to meet demand and provide for depreciation. The second property follows from the uniformity of demand faced by the distributor.

This symmetry in the shipping decision implies that, at constant prices, the optimal shipping problem essentially reduces to a static problem. To reduce notational complexity, we drop, for now, the time subscript. Then, given the above two properties, the manufacturer's cost-minimization problem reduces to

$$\tilde{D}(m_{ij}) = \min_{\{n_{ij}, s_{ij}\}} n_{ij} (K_i + c_{ij} s_{ij}),$$

where n_{ij} and s_{ij} denote the number and size of shipments from country i to country j . Using (4) and properties of the optimal inventory policy, we can solve for the optimal shipment size, conditional on any given shipment frequency, n_{ij} , and total demand, x_{ij} :¹⁶

$$(6) \quad s_{ij} = \frac{x_{ij}}{\delta} (e^{\delta/n_{ij}} - 1).$$

¹⁴ To see this, suppose the manufacturer produces at time t_1 and t_2 . One possibility is to purchase the intermediate exactly at those time periods. Another possibility is to slowly build up inventory of the intermediate between t_1 and t_2 . In both cases, the inventory cost is the same.

¹⁵ Strictly speaking, this is true only if the manufacturer ships at $t = 0$ and $t = 1$. We will have more to say about this issue below.

¹⁶ Given constant prices, only total demand matters. Solving for the change in inventory between some initial time t_0 and $t > t_0$ gives us

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

As discussed above, a shipment should arrive just as inventories in j go to zero. Suppose a shipment arrives at t_0 . Then inventory at t_0 is just the quantity of goods in the shipment. Because the time interval between two successive shipments is $1/n_{ij}$, the inventory as $t \rightarrow t_0 + 1/n_{ij}$ must be zero. Using $m_{ij}(t_0) = s_{ij}$ and $m(t) = 0$ in the above equation, we get (6).




FIGURE 1. EVOLUTION OF INVENTORY OVER TIME

Conditional on demand, the size of shipments is decreasing in shipment frequency. Using the above equation, the manufacturer's objective function can be expressed only as a function of n_{ij} (for a given x_{ij}):

$$\tilde{D}(m_{ij}) = \min_{n_{ij}} n_{ij} K_i + c_{ij} x_{ij} \left(\frac{e^{\delta/n_{ij}} - 1}{\delta/n_{ij}} \right).$$

Note that, when $K_i = 0$, the objective function is minimized by choosing n_{ij} as large as possible. Intuitively, in the absence of a per shipment cost, the manufacturer would want to send as many shipments as possible in order to avoid the losses due to depreciation of inventory. By contrast, when $\delta = 0$, the second term on the right-hand side of the above expression becomes independent of n_{ij} and the objective function is minimized by choosing $n_{ij} = 1$. Intuitively, in the absence of depreciation, the manufacturer would want to send one shipment at the beginning of the period and have the distributor hold a very large inventory. Therefore, we need both K_i , and δ to be nonzero in order to obtain a nontrivial solution for n_{ij} . The following lemma provides the solution for n_{ij} .

LEMMA 1. *The cost-minimizing value of n_{ij} solves the following implicit equation:*

$$(7) \quad \frac{1}{n_{ij}} e^{\delta/n_{ij}} + \frac{1}{\delta} (1 - e^{\delta/n_{ij}}) = \frac{K_i}{c_{ij} x_{ij}}.$$

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 manufacturer yields a unique n_{ij} corresponding to every x_{ij} . The next proposition discusses two important comparative statics results.

PROPOSITION 1. *A manufacturer:*

- (i) *makes more shipments when selling more, for a given per shipment cost.*
- (ii) *makes 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 the per shipment cost declines, efficiency requires that the manufacturer 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 manufacturer to spread the per shipment costs over more units, thereby making it optimal to choose more shipments.¹⁷

We call the value of $\tilde{D}(m_{ij})$ at the optimal n_{ij} , *total distribution cost*, and denote it by D_{ij} :

$$(8) \quad D_{ij} = e^{\delta/n_{ij}} c_{ij} x_{ij}.$$

It is clear from the above equation that, in the presence of an inventory management problem, total distribution cost has the following features: First, it is bounded below by the total cost of replenishing inventory, $c_{ij}x_{ij}$ ($e^{\delta/n_{ij}} > 1$ as long as $\delta > 0$ and n_{ij} is finite). Second, it is proportional to x_{ij} . The factor of proportionality is not a constant, however. With endogenous shipment frequency, as demand starts to increase, the manufacturer wants to send more shipments to save on inventory costs. As a result, δ/n_{ij} declines with x_{ij} .

The interaction between D_{ij} , K_i , and x_{ij} is exhibited in Figure 2. When x_{ij} increases, the way D_{ij} changes depends on the value of K_i . Figure 2, panel A shows the total distribution cost net of the cost of replenishing inventory, $D_{ij} - c_{ij}x_{ij}$, as a function of x_{ij} . For $K_i = 0$, a distributor does not have to carry any inventory; the distributor simply receives a shipment every time there is some demand for the product. In this case, δ/n_{ij} is arbitrarily small and $D_{ij} - c_{ij}x_{ij} = 0$. For $K_i > 0$, following an increase in demand, the manufacturer re-optimizes by increasing both the size as well as the number of shipments. Hence, $D_{ij} - c_{ij}x_{ij}$ increases, but at a decreasing rate. For a given value of x_{ij} , D_{ij} also increases with K_i as the manufacturer reduces the number of shipments, causing an increase in inventory cost.

Figure 2, panel B shows how D_{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, D_{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 manufacturer to have a fixed number of shipments, the standard model increases the distribution cost relative to our model.

B. Price and Quality

To complete the model, we need to solve for the equilibrium quality choices and price paths. Recall from above that manufacturers can customize their quality for each country to which they sell, but that this quality choice is made once-for-all

¹⁷ 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 manufacturer 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.




FIGURE 2. RELATION BETWEEN DISTRIBUTION COST AND OUTPUT

at $t = 0$. By contrast, price can be adjusted over the interval $[0, 1]$. As a result, we first solve for the equilibrium price path for an arbitrary collection of quality choices.¹⁸ Because all manufacturers in country i shipping to country j are identical, we solve for a common value of $p_{ij}(t)$ along the path.¹⁹ Subsequently, we solve for the equilibrium quality choices.

Being a finite horizon problem, the subgame perfect Nash equilibrium price path for the pricing game in our model will be a sequence of one-shot Nash equilibrium price outcomes, as long as there is a unique Nash equilibrium in the one-shot pricing game.²⁰ As the basic structure of the model—endowments, preferences and technologies—is time invariant, the one-shot Nash equilibrium price is also time invariant as long the total demand function is time invariant. As discussed above, this will be the case if manufacturer demand is uniformly distributed on $[0, 1]$. On the interior of the interval, this is without loss of generality. The endpoints pose a potential problem since, in a neighborhood of $t = 1$, the manufacturers will make no further shipments. To address this problem, we assume that, at both $t = 0$ and $t = 1$, the manufacturer must hold a sufficient inventory of intermediate goods to produce the optimal shipment amount.²¹ Under this endpoint condition, the total

¹⁸ The price referred to here, and subsequently, is the purchase or final price. Because the distribution sector is perfectly competitive, an exporter internalizes all distribution costs and, by setting the wholesale price path, sets the purchase price path. As a result, we can solve for the purchase price path by solving the manufacturer's profit maximizing problem when costs are total distribution costs.

¹⁹ Note that we allow for international price discrimination, a common assumption in models of this sort.

²⁰ Because strategies for continuous time games are not clearly defined, the game we have in mind is one in which the interval $[0, 1]$ is partitioned into a finite set of subintervals $[0, \omega]$, $[\omega, 2\omega]$, $\dots$, $[N\omega, 1]$. Manufacturers choose a price for each subinterval and we solve for the subgame perfect Nash equilibrium price path. We define the pricing equilibrium for our continuous time game as the limit of these price paths as the value of ω goes to zero. If the one-shot game has a unique Nash equilibrium, then the price path will simply be the sequence of one-shot Nash equilibria for all values of ω . With constant elasticity demands, the one-shot pricing game involves a price that is a constant markup on marginal cost. Appropriate restrictions on the values of σ and δ ensure that the equilibrium is unique.

²¹ The idea here is that, given the time-invariant nature of the model, we want to reproduce a steady-state outcome while avoiding excess complexity. If the world were one in which manufacturers operated in intervals $\dots, [-T, -T+1], \dots, [-1, 0], [0, 1], [1, 2], \dots$, then, under the optimal policy, manufacturers would accumulate

demand function will be time-invariant and so the subgame perfect Nash equilibrium price path will involve a constant price. Below, we solve for this steady-state pricing equilibrium.

Choosing Price.—Given the steady-state assumption, the profit of a manufacturer in country i selling to country j is simply total revenue minus total distribution cost. Using (8), we can define the manufacturer's pricing problem as

$$\max_{p_{ij}} \pi_{ij} = (p_{ij} - e^{\delta/n_{ij}} c_{ij}) x_{ij}.$$

where π_{ij} is the profit to a manufacturer from country i selling a variety in country j . The resulting profit-maximizing price (see Appendix B) is

$$(9) \quad p_{ij} = \frac{\sigma}{\sigma-1} (1 + d_{ij}) c_{ij},$$

where

$$1 + d_{ij} = \frac{e^{\delta/n_{ij}} - 1}{\delta/n_{ij}}$$

denotes the *marginal distribution cost*. In the presence of the inventory management problem, the manufacturer loses a part of output due to depreciation. This output loss effectively raises the cost of the units that the manufacturer eventually sells: inventory management ends up imposing an ad valorem cost on exports. This cost, which depends not just on the depreciation rate but on the size (and hence, frequency) of the shipments, is endogenous. The following proposition characterizes d_{ij} , the ad valorem export cost.

PROPOSITION 2. d_{ij} 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; as δ becomes positive, $d_{ij} > 0$, so price is higher. In this case, part (ii) of Proposition 2 reveals that, even if the marginal cost of production is a constant, marginal cost

sufficient inventory in the intervals $[-1, 0]$ and $[0, 1]$ to be able to make optimal shipments at $t = 0$ (leading into $[0, 1]$) and $t = 1$ (leading into $[1, 2]$).

of distribution is not. This is a consequence of manufacturer optimization: as total demand rises, the manufacturer makes more frequent shipments, 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 manufacturer 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 manufacturer to send larger shipments and consequently waste more output, thereby pushing up the distribution cost at the margin.

Finally, because the distribution sector in country j is perfectly competitive, a manufacturer in country i receives a FOB price for the product, p_{ij}^{FOB} , that is equal to the final price, p_{ij} , net of the per unit inventory cost and τ_{ij} . This gives us,

$$p_{ij}^{FOB} = \frac{c_{ij}}{\tau_{ij}} \left[1 + \frac{1}{\sigma - 1} \frac{e^{\delta/n_{ij}} - 1}{\delta/n_{ij}} \right].$$

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

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

If manufacturers in rich countries face lower per shipment cost, then Proposition 3 suggests that, other things being the same, the FOB price received by manufacturers 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 II, 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 narrowly-defined good from a richer country, implying that *other things are not the same*. This observation leads us quite naturally to the issue of quality choice.

Choosing Quality.—At $t = 0$, each firm determines its profit maximizing quality, understanding how Nash equilibrium pricing will play out subsequently. Because we restrict all firms in country i selling to country j to choose the same quality level, we can write the representative firm's quality choice problem in the steady state equilibrium as:

$$\max_{q_{ij}} \pi_{ij} = (p_{ij} - e^{\delta/n_{ij}} \tau_{ij} r_i q_{ij}^{\theta_1}) x_{ij} - f q_{ij}^{\theta_2},$$

where p_{ij} is given by (9), $c_{ij} = \tau_{ij} r_i q_{ij}^{\theta_1}$, and the number of shipments is defined implicitly by (7). The solution to the above problem is defined by the condition

$$(10) \quad 0 = \frac{\partial \pi_{ij}}{\partial q_{ij}} = \left(p_{ij} - e^{\delta/n_{ij}} \tau_{ij} r_i q_{ij}^{\theta_1} \right) \frac{\partial x_{ij}}{\partial q_{ij}} \\ - \theta_1 e^{\delta/n_{ij}} \tau_{ij} r_i q_{ij}^{\theta_1-1} x_{ij} \\ - f \theta_2 q_{ij}^{\theta_2-1} \\ + \frac{\delta}{n_{ij}^2} e^{\delta/n_{ij}} \tau_{ij} r_i q_{ij}^{\theta_1} x_{ij} \frac{\partial n_{ij}}{\partial q_{ij}},$$

where the values of both p_{ij} and n_{ij} are given by their profit-maximizing levels. The first three 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 ($\partial x_{ij}/\partial q_{ij} = \sigma x_{ij}/q_{ij} > 0$) and the product sells at a positive price-cost margin ($p_{ij} - e^{\delta/n_{ij}} \tau_{ij} r_i q_{ij}^{\theta_1} > 0$). The second and third terms capture the variable and fixed cost increases that result from an increase in quality. In the usual setting, the optimal quality represents a trade-off between these effects. In our setting, however, there is an additional effect represented by the fourth term in the above equation. This effect captures the fact that, when quality increases, the value of the shipments go up and the manufacturer can reduce marginal costs arising from the inventory management problem ($\partial n_{ij}/\partial q_{ij} > 0$, implying that the fourth term is positive). This profit-enhancing effect creates an additional impetus for the manufacturer to enhance quality.

Substituting in (10) for the values of x_{ij} , $\partial x_{ij}/\partial q_{ij}$, and $\partial n_{ij}/\partial q_{ij}$, and assuming that $f = r_i^{1-\sigma}$, we obtain an implicit solution for q_{ij} (see Appendix B):²²

$$(11) \quad q_{ij} = \left[\frac{\xi_1}{\theta_2} \left(\tau_{ij} \frac{\sigma}{\sigma-1} \frac{e^{\delta/n_{ij}} - 1}{\delta/n_{ij}} \right)^{1-\sigma} P_j^{\sigma-1} \beta Y_j \right]^{\frac{1}{\xi_2}},$$

where

$$\xi_1 = 1 - \theta_1(\sigma - 1)/\sigma,$$

and

$$\xi_2 = \theta_2 - \sigma \xi_1,$$

²² Equation (11) defines an implicit solution for q_{ij} because n_{ij} is a function of q_{ij} .




FIGURE 3. PROFIT AS A FUNCTION OF QUALITY

are constants. We assume that $\xi_1, \xi_2 > 0$. For a given σ , this would be true if the marginal cost of quality does not increase too fast with quality, i.e., θ_1 is small, and if the fixed cost of quality rises fast enough, i.e., θ_2 is large.²³

To see the role played by inventory management in determining quality, consider manufacturers from two different countries, i_1 and i_2 , that both ship to country j and for which $\tau_{i_1j} = \tau_{i_2j}$. It follows from (11) that

$$\frac{q_{i_1j}}{q_{i_2j}} = \left(\frac{e^{\delta/n_{i_1j}} - 1}{\delta/n_{i_1j}} \frac{\delta/n_{i_2j}}{e^{\delta/n_{i_2j}} - 1} \right)^{\frac{1-\sigma}{\xi_2}}.$$

The relative quality of products sold by countries i_1 and i_2 in country j depends only on the relative shipment frequency of their manufacturers. In other words, in the absence of an inventory management problem, the two countries would sell goods of the exact same quality in country j .²⁴

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, a manufacturer from a country 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 and so lowers

²³ In Appendix B, we show that the equilibrium quality, as defined by (11), is unique.

²⁴ The assumption that $f = r_i^{1-\sigma}$ is required for this result. This assumption allows us to focus on inventory management as the sole driver of comparative advantage in quality.

the marginal 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, manufacturers in countries with low per shipment costs tend to choose higher quality products. This is a consequence of the third effect. More specifically, when a manufacturer tries to raise product quality, the value of the goods shipped 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. Manufacturers from 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* in the sense that the relative cost of producing a high- versus low-quality product is lower in countries with low per shipment costs.²⁵ The following proposition formalizes this statement.

PROPOSITION 4. *For any collection of aggregate price indices, P_j fixed for all j , 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 may not be 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 Proposition 4 is not a full general equilibrium result (we hold aggregate prices and incomes fixed to prove this result) and so is, at this point, suggestive only.²⁶ In addition, we would not want to claim that low per shipment costs are the only reason that rich countries produce higher quality products. To examine the role of our proposed mechanism in generating differences in quality across countries, we undertake a quantitative analysis in the next section. In order to do so, however, first we need to close the model.

²⁵ That is, although we have one product in the model, there are varieties of the product that are differentiated vertically. When we talk about comparative advantage in the paper, we are treating varieties with different quality as distinct "products."

²⁶ One might wonder if this result is being driven by the straightforward observation that by raising trade volume, lower per shipment cost lead to higher quality due to the presence of fixed costs of upgrading quality. The answer is no. Proposition 4 is true even in a model with non-CES preference and no fixed cost of quality.

C. Closing the Model

The total quantity sold to country j of any representative country i product is x_{ij} , and total expenditure on that product by country j consumers is $p_{ij}x_{ij}$. Because part of the product melts due to inventory management, the manufacturer receives only $\tau_{ij}p_{ij}^{FOB}x_{ij}$ for its product. Recall that the number of manufacturers/varieties in country i is denoted by M_i . Therefore, total expenditure in country j on all country i products, V_{ij} , is given by $M_i\tau_{ij}p_{ij}^{FOB}x_{ij}$. Using (1), (3), and (9), we can then write V_{ij} as

$$(12) \quad V_{ij} = \Lambda M_i q_{ij}^\sigma \psi(n_{ij}) \left[\tau_{ij} w_i^\alpha P_i^{1-\alpha} q_{ij}^{\theta_1} / A_i \right]^{1-\sigma} \frac{\beta Y_j}{P_j^{1-\sigma}},$$

where

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

and Λ is a constant, n_{ij} solves (7), q_{ij} solves (11), and P_j is given by (2). Replacing the value of p_{ij} in (2), we have

$$(13) \quad P_j^{1-\sigma} = \sum_{k=1}^C \mathbf{1}_{ij} M_i q_{ij}^\sigma \Lambda \left[\left(\frac{e^{\delta/n_{ij}} - 1}{\delta/n_{ij}} \right) \tau_{ij} w_i^\alpha P_i^{1-\alpha} q_{ij}^{\theta_1} / A_i \right]^{1-\sigma}.$$

We assume that the number of products 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 labor inelastically. It can then be shown that $Y_j \approx w_i L_i$.

Trade in each country j is balanced, i.e., $\sum_{i=1}^C V_{ij} = \sum_{i=1}^C V_{ji}$. This equation can be manipulated to obtain a system of equations in wages (Alvarez and Lucas 2007):

$$(14) \quad w_j = \sum_k^C \frac{L_k}{L_j} \lambda_{jk} w_k,$$

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

The standard solution strategy for such a model is as follows. For a given vector of w , use equations similar to (13) to solve for the aggregate price indices—the P s. Equipped with w and P , one can then solve for V_{ij} . Finally, using the trade shares, check whether the trade balance equation (14) is satisfied or not and keep

iterating on w . In this model, however, (12) and (13) are not block-recursive. Rather, P_j depends on V_{ij} through n_{ij} . Hence, for a given vector of wages, the price indices and trade flows have to be solved simultaneously.

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.

V. Quantitative Exercise

To assess the impact that the cross-country differences in per shipment costs have on cross-country differences in the quality of exported goods, total distribution costs, and incomes, we analyze, in this section, a calibrated version of the multicountry general equilibrium model developed in Sections III and IV. We use the calibrated model to study quantitatively how much of the cross-country difference in incomes can be attributed to per shipment cost differences and how much inequality would be reduced if these costs were lowered.

A. Parameter Calibration

For a number of the model parameters, we assign values based on the existing trade literature. The value for the elasticity of substitution (σ) in the literature ranges from 2 to 10 (Anderson and van Wincoop 2004), and we pick $\sigma = 5$. Following Waugh (2010), we set value-added in production (α) at 0.3. We assume a parsimonious specification for bilateral trade costs: $\tau_{ij} - 1 = \text{dist}_{ij}^\rho$, where dist_{ij} is the distance between countries i and j and is given by the great circle distance between capital cities provided by CEPII.²⁷ Hummels (2001) uses actual freight rates to compute the distance elasticity of trade costs and we use his estimate of 0.3 for ρ . Lastly, following AKM, we assume $\delta = 0.3$.²⁸

We determine the values of three other parameters, the share of expenditure on manufactured goods (β), the size of the workforce in each country (L_i) and TFP (A_i) from standard datasets. In particular, β is set at 0.25, which is the median value for our sample of countries (World Factbook 2006). The data for the size of the workforce in each country is from the World Development Indicators (World Bank 2006), while data on TFP for 2006 comes from the Penn World Tables.²⁹

Recall that the parameters θ_1 and θ_2 capture the marginal and fixed cost of producing quality, respectively. Since our goal in this section is to understand how much of cross-country differences in export quality and in income can be explained

²⁷ Mayer and Zignago (2011). For simplicity, we ignore other conventional barriers to trade typically used in the gravity literature: tariffs, contiguity, common language, common colonial origin, etc.

²⁸ Alessandria, Kaboski, and Midrigan (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 a year. This translates into an annual depreciation rate of about 0.27. Assuming further a interest rate of 3 percent per annum, we arrive at $\delta = 0.3$.

²⁹ PWT 7.1 (2012).

by differences in per shipment costs, we maintain the assumption that θ_1 and θ_2 are identical across countries in our calibrated model.³⁰ Their values are calibrated to match two moments in the data: (i) the fraction of nonzero bilateral trade flows, and (ii) the average own import share.

Our choice of targets is motivated by two observations. First, for a fixed value of θ_2 , a lower value of θ_1 results in an increase in export quality. A higher quality reduces the ad valorem distribution cost of exporting relative to that of selling domestically—the former declines with quality, while the latter is fixed in the benchmark scenario. As a result, countries export more/purchase relatively less of their own products. For a fixed value of θ_1 , a reduction in θ_2 means that a large export market plays a less significant role in determining export quality. This means that countries with relatively high per shipment costs can produce higher quality products and so export to more destinations. This leads to a reduction in the overall fraction of zero trade flows. For our sample of countries, the fraction of nonzero trade flows in 2006 was around 0.6 (NBER-UN),³¹ while the average own import-share was around 0.5 (Nicita and Olarreaga 2007).³² We choose values for θ_1 and θ_2 to match these numbers. We summarize the set of parameter values to be used in the quantitative exercise in Table 7.

The remaining, and in many ways the most important, set of parameters to be specified are the values for K_i . In the benchmark scenario, we assume that manufacturers incur no per shipment cost when selling domestically. This is in part because the data we have are only for internationally shipped products.³³ For per shipment costs of exports, we exploit the structure of our model and estimate the values for K_i from the shipment data. Specifically, equation (7) of the model relates the values of K_i to the number of shipments and domestic sales:

$$\frac{e^{\delta/n_{ij}}}{n_{ij}} + \frac{1}{\delta} (1 - e^{\delta/n_{ij}}) = \frac{K_i}{\tau_{ij} r_i q_{ij}^{\theta_1} x_{ij}}.$$

We can rewrite this relationship in terms of the total value of imports per firm, v_{ij} , and number of shipments by noting that $v_{ij} = \tau_{ij} p_{ij}^{FOB} \times n_{ij} \times s_{ij}$. Performing this substitution, we obtain the following relationship:

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

³⁰ By making this assumption, we neutralize any direct comparative advantage in quality effects in our calibrated model. Recall that we also assume that the value of δ is independent of quality.

³¹ Feenstra et al. (2005)

³² Nicita and Olarreaga (2007).

³³ We recognize that this may be a strong assumption for developing countries, where internal shipping costs could be quite substantial. We run sensitivity checks on this assumption in Section VD.

TABLE 7—VALUE OF VARIOUS PARAMETERS

Parameter	Symbol	Value	Source
<i>Panel A. From the literature</i>			
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	World Factbook (2006)
Distance elasticity of trade	ρ	0.3	Hummels (2001)
Depreciation rate	δ	0.3	Alessandria, Kaboski, and Midrigan (2010)
<i>Panel B. Calibrated</i>			
Elasticity of marginal cost of quality	θ_1	0.7	
Elasticity of fixed cost of quality	θ_2	10	

where $\Delta = e^{\delta/n_{ij}}((\delta/n_{ij}) - 1) + 1$. For δ/n_{ij} small, the above equation can be approximated by

$$\ln(v_{ij}) - 2\ln(n_{ij}) = \ln(K_i) + \ln(\delta).$$

This equation forms the basis for our estimates of the values of the K_i s.

While our quantitative model assumes many possible export/import country pairs but only a single export product, our data contains only export flows into Chile but transactions for all HS eight-digit products and all Chilean importers. This difference between the model and data poses several challenges in estimating the values of K_i for use in our quantitative analysis. First, while the model requires, and posits, only a single value for δ , it is unlikely that δ is the same across all products in our data. If we don't allow for this possibility in our estimation procedure, then variations in countries' shipment sizes and frequencies induced by variations in inventory carrying costs will generate biased estimates of the K_i s. To control for variations in carrying costs in as flexible a way as possible, we assume that the value of δ may vary across exporting country (i), product (h), and Chilean importer (l). In particular, we assume that $\ln(\delta_{ihl}) = \ln(\delta_{HS2}) + \epsilon_{ihl}$, where δ_{HS2} is an HS two-digit product fixed effect and ϵ_{ihl} is an unobserved export country, HS eight-digit product (within HS two-digit category), Chilean importer effect. With this, we obtain the following estimating equation:

$$\ln(v_{ihl}) - 2\ln(N_{ihl}) = \ln(K_i) + \ln(\delta_{HS2}) + \epsilon_{ihl}.$$

Under the assumption that ϵ_{ihl} is orthogonal to the country and HS two-digit product fixed effects, this equation can be estimated in a consistent way by OLS. Note, however, that the country and product effects ($\ln(K_i)$ and $\ln(\delta_{HS2})$) are only identified up to a normalization. Therefore, in the estimation, we impose that the average value for δ_{HS2} is equal to our calibrated value for δ of 0.3.

A second potential source of bias in our estimates of the values of the K_i s is the possibility that the per shipment costs are exporting-importing, country-pair specific, rather than solely export-country specific as assumed. It might be, for instance, that

TABLE 8—PER SHIPMENT COSTS AND COUNTRY CHARACTERISTICS

Dependent variable: $\ln(K_i)$	(1)	(2)	(3)	(4)	(5)
$\ln(pcgdp_i)$	−0.100 (0.050)	−0.116 (0.055)	−0.131 (0.053)		−0.030 (0.063)
$\ln(gdp_i)$	−0.043 (0.031)	−0.043 (0.031)	−0.041 (0.031)		−0.058 (0.030)
$\ln(dist_i)$	0.238 (0.097)	0.143 (0.151)			
Latin America dummy		−0.185 (0.250)	−0.387 (0.125)		−0.422 (0.130)
$\ln(\text{number documents to export})$				0.015 (0.031)	0.027 (0.034)
$\ln(\text{number days to export})$				0.026 (0.007)	0.021 (0.008)
$\ln(\text{US\$ cost to export})$				−0.000 (0.000)	0.000 (0.000)
Observations	74	74	74	72	72
R^2	0.17	0.18	0.17	0.23	0.35

Note: Bootstrapped standard errors are in parentheses.

Sources: GDP and pc GDP data for 2006 are from the Penn World Tables. Distance data (distance between capital cities) is from CEPII. Data on documents, days, and cost to export are from the World Bank's Doing Business Surveys.

the per shipment cost incurred in Canada when selling to Chile is different than the one incurred when selling to the United States. If such cases are common, 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.

To check for this possibility, we investigate whether, after controlling for export country size (GDP) and income (per capita GDP), our estimated per shipment costs vary systematically with export country distance to Chile. Column 1 in Table 8 shows that our per shipment cost estimates are, in fact, correlated with export country distance to Chile. Further analysis, however, reveals that this correlation is driven entirely 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 affect 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. Inside and outside the region, the effect of distance is no longer statistically significant.

We use this fact to correct our per shipment cost estimates for "in-region" effects. This is done via the estimates reported in column 3 of the table, which has the specification with a Latin American dummy and no distance. Our estimates reveal that the impact of being in the same region as Chile is a 0.38 log-point reduction in per shipment costs. This translates into more than a 30 percent reduction in absolute per shipment costs for Latin American countries. On average, we obtain

that Latin American countries have almost 5 percent lower log per shipment costs than comparable countries—in terms of GDP and per capita GDP—from outside the region. Table A1, in Appendix A, reports the estimated in-region corrected measure of per shipment costs for all countries in our sample.³⁴ These estimates are the ones we use in the quantitative analysis below. Exporters in the median country are estimated to incur a per shipment cost of US\$4,497; the estimated per shipment cost incurred at the ninetieth percentile of the distribution is almost three times larger than at the tenth percentile of the distribution.

As a further validity check on our estimated per shipment cost values, we look at the correlation between our estimates and the World Bank Doing Business survey (Trading Across Borders 2012) measures on the costs of exporting a container from different countries.³⁵ Specifically, we examine the correlation between our estimates and measures of (i) the number of documents that need to be processed, (ii) the number of days it takes to complete all the processes, and (iii) the monetary cost incurred, for each country in our sample. These three survey metrics are highly correlated among themselves, and each metric singly is correlated with our per shipment cost estimates. When we regress our estimates on all three survey measures taken together (see column 4 in Table 8), we find that these three measures explain 23 percent of the variation in our estimated per shipment costs. Thus, while correlated, our estimates clearly contain additional trading cost effects not captured by the standard survey measures. One example of these additional trading cost effects is the in-region effect identified above. Column 5 of the table shows that the in-region dummy continues to play a major role even once we control for the survey metrics on the costs of trading. Other effects may be the time it takes for orders to be processed in different countries and/or the time it takes for processed orders to make it to the port. Whatever these effects are, column 5 suggests that they must be related to the same economic development variables captured in the survey measures: the relationship between countries' per capita GDP and our measure of per shipment cost vanishes once we include the survey measures of trade costs.

A comparison with estimates reported by other researchers reveals that our estimates of per shipment costs are also larger than ones found in the literature. The reasons for the differences are not easily determined given that our estimates are based on different data and modeling assumptions than used in other works. As an example, we measure shipment frequency and size by importer, product, and export country. AKM and Kropf and Sauré (2014), by contrast, are not able to control for the identity of the importer. By controlling for importer identity, we have a higher level of lumpiness in trade, with the median importer, HS eight-digit product, export country triplet carrying out shipments in only 3 of the 12 months of the year. This is less than half the frequency of trade we measure at the HS eight-digit product and export country level.³⁶ AKM also model demand/supply

³⁴ Per shipment costs are estimated for a sample of countries that exported at least US\$500,000 to Chile in 2006, and carry out at least 30 shipments. We exclude oil exporters.

³⁵ A container refers to a standard, 20 foot container carrying dry cargo.

³⁶ Kropf and Sauré (2014) use a selected sample of exporters and products that they argue biases down their estimates of the per shipment costs making it even harder to compare results.




FIGURE 4. RELATION BETWEEN PER SHIPMENT COST AND PER CAPITA GDP OF THE EXPORTING COUNTRY

uncertainty while we do not. As discussed earlier in the paper, we choose this approach to obtain an analytical solution that can be taken to the data. Adding uncertainty would likely reduce our estimates of the per shipment costs. One way to assess how much uncertainty might matter is to compare the ad valorem export cost (d_{ij} in our model) across studies. Section VB has a discussion of this comparison. To preview our findings, in the best comparison scenario, our estimates produce an 11 percent tariff equivalent rate, while AKM find an 8 percent rate. If some of this difference is due to the more disaggregated data we use, the effect of not modeling uncertainty on our estimates is bound to be small.

We note, finally, that the regression results in Table 8 show that, under multiple specifications, per capita GDP is negatively correlated to our estimate of per shipment cost. 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 -0.37 . 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.

B. Model Fit

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 (versus low-quality) products for rich countries. Following other studies, we use price as an observed proxy for quality and explore how well the




FIGURE 5. PRICE, QUALITY, AND EXPORTER INCOME

Notes: 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.

TABLE 9—PRICE ELASTICITIES

Panel A. Data		Panel B. Model		
All HS8 codes	Top 30 HS8 codes	Benchmark	$K_i = 0$	$K_i = K_{US}$
0.26 (0.05)	0.28 (0.04)	0.06 (0.006)	-0.08 (0.002)	-0.01 (0.004)

Notes: 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 are in parentheses.

model predicts the correlation between exports price and per capita income.³⁷ Figure 5, panel A plots prices p_{ij} , averaged across import destination 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 (see Figure 5, panel B).

To get some sense of how significant our mechanism is in explaining the relationship between price and per capita GDP in the data, we compare estimates of the income elasticity of export prices using our model data and the actual data. The result is reported in Table 9. Recall from Section II that, within Chilean importer-product pairs, a 10 percent increase in per capita income of the source country is related to a 2.6 percent average increase in the prices of the varieties purchased from this country (2.8 percent average increase in the prices of the Top 30 HS-eight products). Table 9 shows the analogous elasticity for our model,

³⁷ In a recent paper, Feenstra and Romalis (2014) shows that a large fraction of the variation in unit values across exporting countries is due to differences in quality.

derived from a regression of $\log p_{ij}$ on $\log \text{GDP}$, $\log \text{per capita GDP}$, and $\log \text{bilateral distance}$.³⁸

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 (the calibrated model), the price elasticity with respect to exporters' per capita income is equal to 0.06 (statistically significant at the 5 percent level). The model, therefore, generates a degree of sensitivity of product price to exporter income that is about a quarter of the one in 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. To get some additional sense of the importance of per shipment costs on this elasticity, we ask what happens if there are no per shipment costs ($K_i = 0$ for all i). In this case, the model generates a price elasticity of -0.08 , meaning that the mechanism in our model—per shipment costs in the exporting country coupled with inventory management costs in importing countries—generates a 0.14 increase in the elasticity of export prices with respect to per capita income.

A second important insight from our model is that there is a link among per shipment costs, comparative advantage in quality and income levels. With cross-country heterogeneity only in per shipment costs, size, TFP, and distance from trading partners, our calibrated model produces (real) GDP and (real) per capita GDP values that have a correlation coefficient with the actual data of 0.9 and 0.8, respectively. The variance of the log-income across countries in the model is 0.88; the country at the ninetieth percentile of the income distribution is 6.4 times richer than the country at the tenth percentile. In the data, for the same set of countries, the variance of log-income across countries is 0.68, while the ratio of the ninetieth to the tenth percentile of the country income distribution is equal to 9.4. These results are reported in Table 10.³⁹ In the absence of per shipment costs ($K_i = 0$ for all i), inequality, as measured by the ratio of the ninetieth to tenth percentile, falls to 2.8. This means that cross-country differences in per shipment costs account for 38 percent of the difference between the incomes of the country at the ninetieth and at the tenth percentile of the income distribution in the data and more than half the inequality produced by the model.

Waugh (2010) shows that, to explain trade flow and price data, trade costs must (i) take an ad valorem form, (ii) be asymmetric with the asymmetry originating in the exporting country, and (iii) be decreasing in the exporting country's level of development. As discussed in Section III, per shipment costs, in the presence of an inventory management problem, translate into ad valorem trade cost equivalents that share these three characteristics. In essence, our model provides a potential micro-foundation for Waugh's approach. The question is, of course, how much of the variation in Waugh's estimated asymmetric trade costs can the model generate.

³⁸ 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.

³⁹ There are a few African countries in the sample that have very low per capita incomes in the benchmark specification and, thus, the ratio of the ninetieth to the tenth percentile of the country income distribution is a more meaningful metric for the level of inequality across countries produced by the model.

TABLE 10—INCOME INEQUALITY

	Data	Benchmark	$K_i = 0$	$K_i = K_{US}$
$\text{var}[\log(\text{Income})]$	0.68	0.88	0.15	0.30
y_{90}/y_{10}	9.4	6.4	2.8	3.8

Notes: y_{90}/y_{10} denotes the ratio of the ninetieth to the tenth percentile real wage. The inequality measures have been computed for the 74 countries in the sample.

A challenge in calculating the average ad valorem export cost equivalent for any given exporting country in our model is that these costs are equilibrium outcomes that depend on the number and size of export transactions *to each foreign country*. If firms in countries with high per shipment costs choose not to export to hard-to-reach destinations, for example, this “selection effect” will bias the estimates of the implied ad valorem trade cost produced by our mechanism. We deal with this potential bias by assigning one export transaction to country pairs that, in the model simulation, had no trade (zero transactions). We, then, compute the *un-weighted* average of the implied ad valorem export-specific trade cost. This gives a lower bound to the implied ad valorem trade costs produced by the model.

The median ad valorem export cost implied by our model, d_{ij} , is 28 percent. In a similar exercise AKM find a tariff equivalent rate of 20 percent. However, given that AKM calibrate their model to US exports data and we assume no shipping lags, the best comparison would be our estimate of 11 percent tariff-equivalent for the United States to AKM’s no-shipping-lag tariff-equivalent rate of 8 percent.

More importantly, we find a strong, negative correlation between the exporting country’s per capita income and ad valorem, export-specific distribution cost implied by our model (see Figure 6). Moreover, Figure 7 shows that the ad valorem trade costs implied by our model correlate strongly with Waugh’s estimates (simple correlation equal to 0.53); and cover the same range of relative values as the ones estimated by Waugh. Lastly, Waugh’s counterfactual exercises that eliminate the asymmetry in trade costs between countries produce reductions in the “ninetieth to tenth ratio” metric of between 23 percent and 33 percent.⁴⁰ As a rough comparison, when we eliminate asymmetries by setting $K_i = K_{US}$, the “ninetieth to tenth ratio” metric of inequality across countries is reduced by 27 percent. Combined, we take this as suggestive that our mechanism can generate a large fraction of the effects captured by Waugh’s asymmetric trade costs.

C. Counterfactuals

As discussed in the introduction, and in contrast to other models of quality differentiation and trade, in our model, policies that reduce a country’s per shipment costs can have a direct impact on the quality of the country’s export products and, via this, on income. So, for instance, trade facilitation policies, such as the WTO Trade Facilitation Agreement, that “. . . contains provisions for expediting the

⁴⁰ See Waugh (2010, 2117) for a description of these exercises.




FIGURE 6. AD VALOREM EXPORT COST AND COUNTRY PER CAPITA INCOME




FIGURE 7. AD VALOREM EXPORT COST: MODEL VERSUS WAUGH (2010)

Note: Ad valorem export costs are normalized to $\tau_{US} = 1$.

movement, release and clearance of goods, . . . ,”⁴¹ and “Single Window” initiatives aimed at reducing the paper work and bureaucracy involved in international trade transactions, can impact not just trade costs but country income levels and income inequality. Using our calibrated model, we are able to quantify this impact and

⁴¹ https://www.wto.org/english/tratop_e/tradfa_e/tradfa_e.htm.




FIGURE 8. CHANGE IN SHIPMENT FREQUENCY AND SIZE

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

examine how reductions in per shipment costs affect distribution (shipping plus inventory) costs, the qualities and prices of countries' exports, income and income inequality, and welfare. This quantification exercise is, to our knowledge, the first general equilibrium treatment of the welfare implications of these types of policies.

We consider, first, the impacts of a policy that changes (in almost all cases, lowers) the per shipment costs for all countries in our sample to those of the United States.⁴² The impact on shipment size and frequency is shown in Figure 8. As expected, countries increase shipment frequency in response to the now lower per shipment costs they face, with the median exporter increasing shipment frequency by 85 percent. While some countries decrease shipment sizes, the increase in shipment frequency tends to be the dominant effect, leading to increases in total trade. The countries experiencing the largest increases in shipment frequency are the ones that were initially poor, as illustrated in Figure 9. On average, these countries experience the largest reductions in their per shipment costs.

Table 11 reports the effects of changing per shipment costs to US levels (as well as $K_i = 0$) on aggregate and individual countries' real income. At the aggregate level, the increase is small—about 1 percent. For the country with the median real per capita income in the sample, however, real per capita income increases by 13.7 percent. This difference reflects the extent of heterogeneity in real income impacts across countries, with poor countries benefiting strongly, while the richest countries actually lose due to the added competition in high-quality varieties. Indeed, among the poorest 25 percent of countries in the sample, the median

⁴² Only one of the 74 countries in our sample has per shipment costs lower than the United States—Japan.




FIGURE 9. CHANGE IN SHIPMENT FREQUENCY AND INITIAL PER CAPITA GDP

Note: In this exercise, we reduce per shipment costs from the benchmark values to K_{US} . Values for per capita GDP are in logs and expressed as deviations from the mean.

TABLE 11—INCOME EFFECTS OF COUNTERFACTUAL EXERCISES

	$K_i = 0$	$K_i = K_{US}$
<i>Percentage change in real income:</i>		
Aggregate	1.0	1.6
Median country	25.8	13.7
Median rich country	0.1	−1.9
Median poor country	93.8	45

Note: The sample has 68 countries, so median values are calculated as the average of the two countries at the median points.

country experiences a 45 percent increase in real income as a result of the equalization of per shipment costs to US levels. Among the richest 25 percent of countries, the median country experiences a 1.9 percent 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.

The same policy also has a large effect on the export price elasticity with respect to exporting country per capita GDP (Table 9) and, as discussed before, reduces our preferred metric of inequality across countries by 27 percent (Table 10).

Finally, we can decompose the welfare effects of a change in per shipment costs into the change due to changes in total shipment costs, changes in total inventory costs, and price changes (consumption changes due to general equilibrium changes in wages and quality adjusted prices). The effect on total shipment costs can be further decomposed into the part due directly to the change in per shipment cost,




FIGURE 10. WELFARE EFFECT OF $K_i = K_{US}$

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

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 total 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 12 reports this decomposition for the country with the median change in real income under a policy that changes per shipment costs in all countries to US levels. This country experiences an increase in real income of 12.1 percent.⁴³ The decomposition shows that this change has a small direct effect on per shipment cost amounting to a savings of 0.36 percent of real income. This direct cost savings is more than offset by an indirect cost effect resulting from the fact that the lower per shipment costs induce exporters to send more shipments, resulting in an indirect cost effect of -0.38 percent of real income. Counteracting this is a cost savings due to the fact that the additional shipments result in smaller shipments and, hence, reduced inventory costs. This reduced inventory cost effect amounts to an offsetting 0.23 percent increase in real income. Finally, because the lower distribution costs result in a larger total volume of exports and in exports of higher quality, total inventory costs increase in an amount equal to -0.14 percent of real income. Overall, for the median country, the impact on distribution costs is small, with the per shipment cost reduction leading to a total distribution cost savings of 0.12 percent of real income.

⁴³ This is actually a country at the sixtieth percentile of the income distribution in the benchmark scenario. The country at the median of the real per capita income distribution has similar numbers for this decomposition.

TABLE 12—DECOMPOSITION OF WELFARE EFFECTS

	Total	Distribution costs				Price adjustment effect
		Per shipment cost		Inventory cost		
		$\frac{\partial(K \times n)}{\partial K} \Big _{\bar{n}}$	$\frac{\partial(K \times n)}{\partial n} \Big _{K_{US}}$	$\frac{\partial(IC)}{\partial N} \Big _{\bar{x}}$	$\frac{\partial(IC)}{\partial x} \Big _{n_{US}}$	
<i>Panel A. Variable quality</i>						
Median change (percent)	12.1	0.36	−0.38	0.23	−0.14	12.08
<i>Panel B. Fixed quality</i>						
Median change (percent)	0.50	0.34	−0.13	0.21	−0.00	0.08

Notes: In this exercise, we reduce per shipment costs from the benchmark values to K_{US} . Under the fixed-quality scenario, quality is kept unchanged at the benchmark levels. $\partial(K \times n)/\partial K|_{\bar{n}}$ is the change in the per shipment component of costs, nK , due to a change in K , evaluated at the benchmark level of n ; $\partial(K \times n)/\partial n|_{K_{US}}$ is the change in the per shipment component of costs due to the optimal adjustment in n , evaluated at the new per shipment costs; $\partial(IC)/\partial n|_{\bar{x}}$ is the change in total inventory cost due to the adjustment in n , holding total sales constant at the benchmark level; $\partial(IC)/\partial x|_{n_{US}}$ is the change in total inventory costs due to the optimal adjustment in total sales, holding N constant at its new level. IC is given by $\tau_{ij} r_i q_{ij}^{\theta_i} (n_{ij} s_{ij} - x_{ij})$.

The major impact of the reduction in per shipment costs is in the wage and quality-adjusted price changes due to the quality upgrading that reduced distribution costs permit. The second row of Table 12 confirms this finding. It shows the effects of the same policy—changes in per shipment costs to US levels—when quality is not allowed to vary. The total effect of the policy change is one order of magnitude smaller than when firms are allowed to upgrade quality.

The reason that the reduction in per shipment costs to US levels has so little impact, quality fixed, is that the reduction translates into very little change in implied variable trade costs (in contrast to Waugh). In our model, the (endogenous) variable trade cost depends on the number and size of transactions firms choose to make. Firms in poor countries export little—they have low productivity and lower quality products. For these firms, inventory costs do not represent a significant part of their overall cost of distribution (recall the discussion in Section IVA). As a result, the reduction in per shipment cost, quality fixed, has little effect on their number of shipments and export volumes, and so on implied variable trade costs and incomes. As quality upgrading occurs, however, poor countries can exploit the benefits of lower per shipment costs, increasing the number of shipments and exports, dramatically reducing effective variable trade cost and increasing real income. To illustrate, take the country experiencing the median change in income depicted in Table 12. A change in per shipment cost to US levels represents a 35 percent reduction for this country. This results in a 23 percent reduction in this country's variable trade cost, that leads to a 8.4 percent increase in exports and to the 12.1 percent increase in real income reported above. Quality fixed, however, the same change in per shipment cost translates into a reduction of this country's variable trade cost by a mere 1 percent, and virtually no change in real income.

D. Sensitivity Analysis

In deriving the results of the previous section, we assume that poor countries faced with a reduction in per shipment costs have the same ability to upgrade quality

as do rich countries: rich and poor countries have identical values for θ_1 and θ_2 . We also assume that rich and poor countries have identical inventory management technologies—the value of δ is the same across countries—and that there are no per shipment costs for internal shipments. In this section, we examine the extent to which our results change if we relax some of these assumptions. In general, we find that the results continue to hold in large measure but, of course, the magnitudes of the impacts change.

Different θ_1 .—A potentially critical assumption for the counterfactual analysis is that (rich) countries have no natural comparative advantage in producing high quality goods. Were poor countries assumed to have higher costs of upgrading quality than rich countries, then one can imagine that reductions in the per shipment costs of exporting would not promote as much quality upgrading, and so income growth, in poor countries nor as much competition in high-quality goods to rich countries.⁴⁴ To check the robustness of our quantitative findings with respect to this assumption, we allow θ_1 , the marginal cost of producing quality, to vary between rich and poor countries.

The extent of rich countries' natural comparative advantage in quality is unclear. For this analysis, we assume that $\theta_1^{poor} = a\theta_1^{rich}$, where θ_1^{rich} is the same as in the benchmark specification, and $a > 1$ is such that the model matches the ninetyth-to-tenth inequality ratio in the data. Recall that, in our sample, the country in the ninetyth percentile of the income distribution is 9.4 times richer than the country at the tenth percentile of the same distribution. The benchmark specification generates a ratio equal to 6.4. Calibrated in this way, we find that $a = 1.1$ can produce the inequality in the data, when added to the benchmark specification of our model.

Column 3 in Table 13 reports results of this sensitivity exercise. Poor countries' higher cost to produce quality widens inequality across countries, primarily by making poor countries even poorer, although rich countries also benefit slightly from less competition in high-quality varieties. A policy that reduces all countries' per shipment costs to US levels leaves cross-country income inequality slightly larger than in the benchmark counterfactual (4.0 versus 3.8). However, the reduction of per shipment costs to US levels actually has a larger effect on reducing inequality than in the benchmark case (58 percent versus 27 percent). The reason is that, when poor countries have a 10 percent higher marginal cost of producing quality, they produce and trade little, mostly because of the low price of their exports. In this scenario, the marginal effect of reducing per shipment costs on trade volumes and, thus, on the efficiency of these countries' distribution systems is quite large.

⁴⁴ Note that the existing literature on quality and trade argues that rich countries' comparative advantage in quality can be the result of: (i) quality being intensive in physical and human capital (Markusen 1986, Bergstrand 1990); (ii) quality being a luxury "good" coupled with economies of scale in the production of quality (Linder 1961; Fajgelbaum, Grossman, and Helpman 2011); (iii) high quality being complementary with high productivity (Flam and Helpman 1987, Matsuyama 2000, Baldwin and Harrigan 2011, Johnson 2012, and Feenstra and Romalis 2014).

TABLE 13—SENSITIVITY ANALYSIS

	Data	Benchmark	θ_1^{poor} $= 1.1 \times \theta_1^{rich}$	$\delta_{rich} = 0.3;$ $\delta_{poor} = 0.9$	$K_{dom.}$ $= \frac{1}{4}K_{int.}$
y_{90}/y_{10}	9.4	6.4	9.5	7.6	19.4
Export price elasticity of income	0.26	0.06	0.06	0.05	0.10
Median ad valorem export cost (%)		28%	28%	26%	28%
<i>Changing K to K_{US}:</i>					
y_{90}/y_{10}		3.8	4.0	4.8	3.7
<i>Percentage change in real income</i>					
Aggregate		1.0	1.0	0.5	0.8
Median country		13.7	15.9	7.8	10.0
Median rich country		0.1	−2.0	2.8	2.1
Median poor country		93.8	95.3	35.5	225.1

Different δ .—Another potentially important source of heterogeneity between rich and poor countries is in inventory management. The benchmark counterfactuals assume that all countries have the same inventory management technology (value of δ is assumed the same across importing countries). Because δ captures country-specific components of the cost of physically storing goods, it is reasonable to imagine that these costs would be higher in developing countries. (Guasch and Kogan 2001), for instance, argue that poor countries have domestic depreciation rates that are up to three times larger than those in rich countries. If this is the case, then reductions in per shipment costs may benefit poor countries even more by making smaller and more frequent shipments to these countries more economically feasible. To check for this, we redo the counterfactuals assuming that the value of δ is three times larger for countries in the bottom twenty-fifth percentile of the income distribution than it is for all other countries.

Column 4 in Table 13 reports results of this sensitivity exercise. The higher depreciation rate in poor countries reduces these countries' income and increases inequality across countries. In this scenario, poor countries import very little, due to the high level of wastage in inventory and, thus, export very little. The reduction in per shipment cost benefits poor countries but, as in the previous case, income inequality continues to be higher than in the benchmark case. Again, as in the previous case, the reduction in per shipment costs creates a larger percentage reduction in income inequality (38 percent versus 27 percent in the benchmark) for the reasons outlined.

Different domestic K.—Finally, we assumed in our benchmark specification that there are no domestic per shipment costs, implying that producers face no inventory management problem when selling domestically. Although it is generally agreed that domestic trade costs are smaller than international ones, assuming away all domestic inventory management problems is clearly a strong assumption. AKM uses data on differential inventory-sales ratios between domestic and imported inputs to conclude

that international per shipment costs are about four times as large as domestic per shipment costs. If this is the case, then poor countries have inefficient domestic distribution systems as well, and this makes policies to reduce distribution costs even more important. To provide a rough quantification of the implications of inefficiencies in countries' domestic distribution systems, we redo our analysis assuming that the domestic per shipment cost in country i is one-fourth of its international per shipment cost, K_i . Now, the counterfactual exercise that reduces international per shipment cost to US levels also reduces the countries' domestic distribution costs.

Column 5 in Table 13 reports results of this analysis. With domestic and international per shipment costs, the model generates twice as much cross-country income inequality as in the data. The reason, primarily, is that poor countries are significantly damaged by their inefficient domestic distribution systems. A reduction in per shipment costs, both international and domestic, to US levels creates an even larger benefit for poor countries than in the benchmark case.

To get a sense of how much of the above mentioned changes is being driven by the difference between international and domestic per shipment costs, we perform a counterfactual where only the international per shipment costs are reduced to the domestic levels. In this case, cross-country income inequality falls from 19.4 to 7.2, while aggregate income rises by 0.7 percent. Thus, a bulk of the increase in income and decrease in inequality is brought about by eliminating the *additional* barriers that impede the movement of shipments across international borders. The results are probably not surprising because we assume that domestic per shipment costs are one-fourth of the external costs. Consequently, for the high per shipment cost countries, a lowering of the external costs to the corresponding domestic levels represents a large reduction in absolute terms.

VI. Conclusion

Trade frictions are an important reason for why standards of living vary so widely across countries, and they matter because they affect the countries' ability to specialize in their most productive sectors (Waugh 2010). In this paper, we argue that they also matter because they affect the countries' ability to upgrade the quality of the products they produce.

A key finding of our analysis is that countries with efficient distribution systems (low per shipment cost) export goods with high-inventory costs. In our model, this has the implication that countries with low per shipment costs export high-quality varieties: high-quality varieties of a product have high unit values and so high inventory costs. An alternative implication of this finding, for a multi-product world, is that countries with low per shipment costs export products with high rates of depreciation while in inventory. We did not pursue this implication in the paper, given our single product, multicountry setting. Nevertheless, our estimation procedure produces estimates for the values of depreciation rates while in inventory for HS two-digit product categories. Some preliminary analysis suggests that this alternative implication is consistent with the trade flow data. We will be pursuing more work on this issue in future research.

APPENDIX A

TABLE A1—PER SHIPMENT COSTS (2006 US DOLLARS) AT
EXPORTING COUNTRY: K_i

Country	US Dollars	Country	US Dollars	Country	US Dollars
JPN	2,484	TWN	4,493	LVA	5,789
USA	2,664	DNK	4,494	PAN	5,861
MAR	2,689	TUR	4,541	ZAF	6,037
SWE	2,753	CAN	4,629	MEX	6,126
HUN	2,827	CRI	4,660	IND	6,192
DEU	2,843	PER	4,675	HND	6,252
TUN	3,021	MYS	4,692	COL	6,263
CHE	3,027	FIN	4,745	PRT	6,484
BGR	3,106	ISL	4,835	SVK	6,539
KOR	3,201	ECU	4,841	SLV*	6,625
CHN	3,293	PHL	5,003	LKA	6,669
SGP	3,365	NLD	5,016	VNM*	7,163
FRA	3,406	ARG	5,075	PRY	7,412
GBR	3,440	NZL	5,101	GTM	7,981
EGY	3,553	AUS	5,121	URY	8,112
MLT	3,670	POL	5,190	SVN	8,304
BOL	3,784	BGD*	5,200	VEN	9,079
LTU	3,819	DOM	5,247	HKG	9,362
ITA	3,970	IDN	5,251	RUS	9,790
ESP	4,145	GRC	5,253	EST	12,298
BEL	4,171	IRL	5,259	PAK*	13,248
AUT	4,412	THA	5,381	UKR	14,922
CZE	4,419	ISR	5,458	IRN	14,962
BLR*	4,425	HRV	5,628	KHM*	19,948
ROU	4,455	BRA	5,701		

* Indicates countries not included in calibrations exercise due to lack of required data on TFP.

APPENDIX B

PROOF OF PROPOSITION 1:

Note that we can write the implicit solution to the inventory management problem as

$$e^{\delta/n} \left(\frac{1}{\delta} - \frac{1}{n} \right) = \frac{1}{\delta} - \chi.$$

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 2:

Note that d can be written as $d = f(s) \times A$, where A is independent of both X and K , 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 . But s is decreasing in X and increasing in K . Hence, (ii) and (iii) follow. ■

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 neighborhood 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 4:

Note that q can be written as $q = g(s) \times B$, where B is exogenously given and

$$g(s) = \left(\frac{e^s - 1}{s} \right)^{\frac{1-\sigma}{\xi_2}},$$

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

$$\frac{\partial g(s)}{\partial s} = \frac{1-\sigma}{\xi_2} \left(\frac{e^s - 1}{s} \right)^{\frac{1-\sigma}{\xi_2} - 1} \frac{1}{s^2} [e^s(s-1) + 1].$$

From $e^s > s + 1$ and $\sigma > 1$, we have

$$\frac{\partial g(s)}{\partial s} < 0.$$

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 . ■

Derivation of Equilibrium Price.—The exporter's problem is

$$\max_p (p - ce^{\delta/n}) p^{-\sigma},$$

where c is the marginal cost of production. The first-order condition with respect to p yields

$$p = \frac{\sigma}{\sigma - 1} ce^{\delta/n} \left(1 + \frac{p}{\sigma} \frac{\delta}{n^2} \frac{\partial n}{\partial p} \right).$$

Notice that optimal n and p are related by

$$\frac{1}{n} e^{\delta/n} + \frac{1}{\delta} (1 - e^{\delta/n}) = \frac{K}{cp^{-\sigma}}.$$

From the above equation, we obtain the following:

$$\frac{p}{\sigma} \frac{\delta}{n^2} \frac{\partial n}{\partial p} = \frac{n}{\delta} - 1 - \frac{n}{\delta} \frac{1}{e^{\delta/n}}.$$

Replacing this in the expression for equilibrium p , we have

$$p = \frac{\sigma}{\sigma - 1} \frac{e^{\delta/n} - 1}{\delta/n} c.$$

Derivation of Equilibrium Quality.—The first-order condition of the exporter with respect to quality can be rewritten as

$$(p - e^{\delta/n} \tau r q^{\theta_1}) \frac{\partial x}{\partial q} - \theta_1 e^{\delta/n} \tau r q^{\theta_1-1} x + \frac{\delta}{n^2} e^{\delta/n} \tau r q^{\theta_1} x \frac{\partial n}{\partial q} = r^{1-\sigma} \theta_2 q^{\theta_2-1}.$$

Notice that optimal n and q are related by

$$\frac{1}{n} e^{\delta/n} + \frac{1}{\delta} (1 - e^{\delta/n}) = \frac{K}{\tau r q^{\theta_1} x}.$$

From the above equation, we obtain the following:

$$\frac{\delta}{n^2} e^{\delta/n} \tau r q^{\theta_1} x \frac{\partial n}{\partial q} = \frac{nK}{x} \left(\frac{\theta_1}{q} + \frac{\sigma}{q} \right).$$

Replacing this in the expression for equilibrium q and using $\partial X/\partial q$, we have

$$(p - e^{\delta/n} \tau r q^{\theta_1} + \frac{nK}{x}) \frac{\sigma}{q} - (e^{\delta/n} \tau r q^{\theta_1} - \frac{nK}{x}) \frac{\theta_2}{q} = r^{1-\sigma} \theta_2 q^{\theta_2-1}.$$

Now, from the first-order condition for n , we have

$$\begin{aligned} e^{\delta/n} \tau r q^{\theta_1} - \frac{nK}{x} &= \frac{e^{\delta/n} - 1}{\delta/n} \tau r q^{\theta_1} \\ &= \frac{\sigma - 1}{\sigma} p. \end{aligned}$$

Replacing this in the equation above and using $x = q^\sigma p^{-\sigma}$, we can solve for q .

Uniqueness of Equilibrium Quality.—For any collection of aggregate price indices, P_j fixed for all j , equilibrium quality satisfies

$$q = \chi \left[\left(\frac{e^{\delta/n} - 1}{\delta/n} \right)^{1-\sigma} x \right]^{\frac{1}{\xi_2}},$$

where χ is independent of q . Conditional on x , we then have

$$q \approx \left(\frac{e^{\delta/n} - 1}{\delta/n} \right)^{\frac{1-\sigma}{\xi_2}},$$

where n solves the optimal inventory management problem for a given q and x . It is easy to show that $\partial n/\partial q > 0$. Now, the function $(e^z - 1)/z$ is increasing in z .

Therefore, as q increases, the right-hand side of the above equation falls. Furthermore, as $q \rightarrow 0$, the right-hand side converges to infinity, while as $q \rightarrow \infty$, the right-hand side converges to one. Hence, conditional on x , there exists a unique q .

Because $x = q^\sigma p^{-\sigma}$, we have the following:

$$x \approx q^{\sigma(1-\theta_1)} \left(\frac{e^{\delta/n} - 1}{\delta/n} \right)^{-\sigma}.$$

Replacing the value of q , and re-arranging terms, we have

$$x = \left(\frac{e^{\delta/n} - 1}{\delta/n} \right)^{-\frac{\sigma\theta_2 - \sigma(1-\theta_1)}{\theta_2 - \sigma(1-\theta_1)}}.$$

The above equation can be rewritten as $x = 1/f(x)^\psi$, where $\psi > 1$ if θ_2 is large enough. Furthermore, $f(x)$ is convex in x . Hence, $f(x)^\psi$ is convex in x and $1/f(x)^\psi$ is concave in x . It is straightforward to show that the right-hand side of the above equation has a slope greater than one as x goes to zero. Therefore, x has a unique solution.

APPENDIX C

In this section, we solve the inventory management problem of a firm in a setting where demand is uncertain. We are particularly interested in the relation between the number and size of orders (or shipments) and per shipment costs. To obtain sharp characterization of the equilibrium number and size of shipments, we develop an extremely stylized model.

Firms live for two periods. Demand in the two periods, q_1 and q_2 , is uncertain and is drawn independently from a distribution with pdf $\phi(q)$. 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. One can think of h as capturing how fast goods depreciate, with lower h corresponding to higher depreciation rate. Each unit normally sells for p , costs c ($c < p$), and the per shipment cost is K . Observe that if $h > c$, the firm would simply place an order at the beginning of period one that equals the maximum possible total demand in periods one and two, and then take no other action. In this scenario, there is effectively no penalty for carrying an overstock. In order to get a non-trivial solution for the inventory management problem, we therefore assume that h is strictly less than c . Because of the finite horizon, we solve the firm's problem recursively, starting with period 2.

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 beginning of 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 $y_2 - q_2$. Thus $E[r_2]$, the expected revenue in period 2, is given by

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

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

$$E[r_2] = py_2 - \frac{1}{2}py_2^2 + \frac{1}{2}hy_2^2.$$

The corresponding expected marginal revenue in period 2 is $p - (p - h)y_2$. The firm orders a shipment only if $y_2 > x$. The marginal cost incurred by the firm in period 2 is then given by

$$MC_2 = \begin{cases} c & \text{if } y_2 > x \\ 0 & \text{if } y_2 \leq x \end{cases}$$

Given the above marginal revenue and marginal cost functions, the optimum level of inventories in period 2 is given by

$$y_2^* = \begin{cases} \frac{p}{p-h} & \text{if } x \geq \frac{p}{p-h} \\ x & \text{if } \frac{p-c}{p-h} \leq x < \frac{p}{p-h} \\ \frac{p-c}{p-h} & \text{if } x < \frac{p-c}{p-h} \end{cases}$$

When $x < (p - c)/(p - h)$, the firm has to incur a cost of $c(y_2^* - x) + K$. The corresponding expected profit levels are

$$E[\pi_2] = \begin{cases} \frac{p^2}{2(p-h)} & \text{if } x \geq \frac{p}{p-h} \\ px - \frac{1}{2}(p-h)x^2 & \text{if } \frac{p-c}{p-h} \leq x < \frac{p}{p-h} \\ \frac{(p-c)^2}{2(p-h)} + cx - K & \text{if } x < \frac{p-c}{p-h} \end{cases}$$

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. The firm's period 1 expected profit is then given by

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

The firm's expected total profit for the two periods is

$$\begin{aligned}
 E[\pi_1 + \pi_2] &= E[\pi_1] \\
 &+ E\left[\pi_2 \mid x \geq \frac{p}{p-h}\right] \cdot \Pr\left(x \geq \frac{p}{p-h}\right) \\
 &+ E\left[\pi_2 \mid \frac{p-c}{p-h} \leq x < \frac{p}{p-h}\right] \cdot \Pr\left(\frac{p-c}{p-h} \leq x < \frac{p}{p-h}\right) \\
 &+ E\left[\pi_2 \mid x < \frac{p-c}{p-h}\right] \cdot \Pr\left(x < \frac{p-c}{p-h}\right).
 \end{aligned}$$

Using the expressions for profits in the two periods and using the result that $x = y_1 - q_1$, we have

$$\begin{aligned}
 E[\pi_1 + \pi_2] &= py_1 - \frac{1}{2}py_1^2 - cy_1 - K \\
 &+ \frac{p^2}{2(p-h)}\left(y_1 - \frac{p}{p-h}\right) \\
 &+ \left[p\left(y_1 - \frac{1}{2}\right) - \frac{1}{2}(p-h)\left(y_1^2 - y_1 + \frac{1}{3}\right)\right]\frac{c}{p-h} \\
 &+ \left[\frac{(p-c)^2}{2(p-h)} + c\left(y_1 - \frac{1}{2}\right) - K\right]\left(1 - y_1 + \frac{p-c}{p-h}\right).
 \end{aligned}$$

Maximizing the above expression with respect to y_1 yields

$$y_1 = \frac{K + \Xi}{p + 3c},$$

where Ξ is independent of K . 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 two with probability $\Pr(y_1 - q_1 > (p-c)/(p-h)) = 1 - y_1 + (p-c)/(p-h)$. Because y_1 is increasing in K , this probability is decreasing in K . Hence, the number of period two shipments decreases when K rises.

REFERENCES

- Alchian, Armen A., and William R. Allen. 1964. *University Economics*. Belmont, CA: Wadsworth.
- Alessandria, George, Joseph Kaboski, and Virgiliu Midrigan. 2010. "Inventories, Lumpy Trade, and Large Devaluations." *American Economic Review* 100 (5): 2304–39.

- Alvarez, Fernando, and Francesco Lippi. 2009. "Financial innovation and the transactions demand for cash." *Econometrica* 77 (2): 363–402.
- Alvarez, Fernando, and Robert E. Lucas, Jr. 2007. "General equilibrium analysis of the Eaton-Kortum model of international trade." *Journal of Monetary Economics* 54 (6): 1726–68.
- Anderson, James E., and Eric van Wincoop. 2004. "Trade Costs." *Journal of Economic Literature* 42 (3): 691–751.
- Antràs, Pol, Teresa C. Fort, and Felix Tintelnot. 2017. "The Margins of Global Sourcing: Theory and Evidence from U.S. Firms." *American Economic Review* 107 (9): 2514–64.
- Arrow, Kenneth J., Theodore Harris, and Jacob Marschak. 1951. "Optimal Inventory Policy." *Econometrica* 19 (3): 250–72.
- Baldwin, Richard, and James Harrigan. 2011. "Zeros, Quality and Space: Trade Theory and Trade Evidence." *American Economic Journal: Microeconomics* 3 (2): 60–88.
- Baumol, William J. 1952. "The Transactions Demand for Cash: An Inventory Theoretic Approach." *Quarterly Journal of Economics* 66 (4): 545–56.
- Békés, Gabor, Lionel Fontagné, Balazs Muraközy, and Vincent Vicard. 2017. "Shipment frequency of exporters and demand uncertainty." *Review of World Economics* 153 (4): 779–807.
- Bergstrand, Jeffrey H. 1990. "The Heckscher-Ohlin-Samuelson Model, The Linder Hypothesis and the Determinants of Bilateral Intra-Industry Trade." *Economic Journal* 100 (403): 1216–29.
- Blum, Bernardo S., Sebastian Claro, Kunal Dasgupta, and Ignatius J. Horstmann. 2019. "Inventory Management, Product Quality, and Cross-Country Income Differences: Dataset." *American Economic Journal: Macroeconomics*. <https://doi.org/10.1257/mac.20170080>.
- Chaney, Thomas. 2008. "Distorted Gravity: The Intensive and Extensive Margins of International Trade." *American Economic Review* 98 (4): 1707–21.
- Evans, Carolyn L., and James Harrigan. 2005. "Distance, Time, and Specialization: Lean Retailing in General Equilibrium." *American Economic Review* 95 (1): 292–313.
- Fajgelbaum, Pablo, Gene M. Grossman, and Elhanan Helpman. 2011. "Income Distribution, Product Quality, and International Trade." *Journal of Political Economy* 119 (4): 721–65.
- Feenstra, Robert C., Robert E. Lipsey, Haiyan Deng, Alyson C. Ma, and Hengyong Mo. 2005. "World Trade Flows: 1962–2000." National Bureau of Economic Research (NBER) Working Paper 11040.
- Feenstra, Robert C., and John Romalis. 2014. "International Prices and Endogenous Quality." *Quarterly Journal of Economics* 129 (2): 477–527.
- Flam, Harry, and Elhanan Helpman. 1987. "Vertical Product Differentiation and North-South Trade." *American Economic Review* 77 (5): 810–22.
- Guasch, Jose Luis, and Joseph Kogan. 2001. "Inventories in Developing Countries: Levels and Determinants: A Red Flag for Competitiveness and Growth." World Bank Policy Research Working Paper 2552.
- Hornok, Cecília, and Miklós Koren. 2015. "Per-Shipment Costs and the Lumpiness of International Trade." *Review of Economics and Statistics* 97 (2): 525–30.
- Hummels, David. 2001. "Toward a Geography of Trade Costs." <http://www.krannert.purdue.edu/faculty/hummelsd/research/toward/tgtc.pdf>.
- Hummels, David, and Alexandre Skiba. 2004. "Shipping the Good Apples Out? An Empirical Confirmation of the Alchian-Allen Conjecture." *Journal of Political Economy* 112 (6): 1384–1402.
- Johnson, Robert C. 2012. "Trade and prices with heterogeneous firms." *Journal of International Economics* 86 (1): 43–56.
- Kropf, Andreas, and Philip Sauré. 2014. "Fixed costs per shipment." *Journal of International Economics* 92 (1): 166–84.
- Krugman, Paul. 1980. "Scale Economies, Product Differentiation, and the Pattern of Trade." *American Economic Review* 70 (5): 950–59.
- Linder, Staffan Burenstam. 1961. *An Essay on Trade and Transformation*. Uppsala: Almqvist and Wiksell.
- Lugovskyy, Volodymyr, and Alexandre Skiba. 2016. "Positive and negative effects of distance on export prices." *Journal of Economic Behavior and Organization* 127: 155–81.
- Markusen, James R. 1986. "Explaining the Volume of Trade: An Eclectic Approach." *American Economic Review* 76 (5): 1002–11.
- Martin, Julien. 2012. "Markups, quality, and transport costs." *European Economic Review* 56 (4): 777–91.
- Matsuyama, Kiminori. 2000. "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.

- Mayer, Thierry, and Soledad Zignago.** 2011. "Notes on CEPII's Distance Measures: The GeoDist Database." Centre d'Études Prospectives et d'Informations (CEPII) Working Paper 2011–25.
- Nicita, Alessandro, and Marcelo Olarreaga.** 2007. "Trade, Production, and Protection Database, 1976–2004." *World Bank Economic Review* 21 (1): 165–71.
- PWT 7.1.** 2012. Center for International Comparisons of Production, Income and Prices, University of Pennsylvania. <https://www.rug.nl/ggdc/productivity/pwt/pwt-releases/pwt-7.1>.
- Santos-Silva, J. M. C., and Silvana Tenreyro.** 2006. "The Log of Gravity." *Review of Economics and Statistics* 88 (4): 641–58.
- Tobin, James.** 1956. "The Interest-Elasticity of Transactions Demand For Cash." *Review of Economics and Statistics* 38 (3): 241–47.
- Trading across Borders.** 2012. World Bank. <http://www.doingbusiness.org/data/exploretopics/trading-across-borders>.
- Waugh, Michael E.** 2010. "International Trade and Income Differences." *American Economic Review* 100 (5): 2093–2124.
- World Bank.** 2006. *2006 World Development Indicators*. World Bank Development Data Center. Washington, DC, April.
- World Factbook.** 2006. US Central Intelligence Agency. <https://www.cia.gov/library/publications/the-world-factbook/>.