

Zeros, Quality, and Space: Trade Theory and Trade Evidence[†]

By RICHARD BALDWIN AND JAMES HARRIGAN*

Bilateral, product-level data exhibit a number of strong patterns that can be used to evaluate international trade theories, notably the spatial incidence of “export zeros” (correlated with distance and importer size), and of export unit values (positively related to distance). We show that leading theoretical trade models fail to explain at least some of these facts, and propose a variant of the Melitz model that can account for all the facts. In our model, high quality firms are the most competitive, with heterogeneous quality increasing with firms’ heterogeneous cost. (JEL F11, F14, F40)

The gravity equation relates bilateral trade volumes to distance and country size. Countless gravity equations have been estimated, usually with “good” results, and trade theorists have proposed various theoretical explanations for gravity’s success. However, the many potential explanations for the success of the gravity equation make it a problematic tool for discriminating among trade models.¹

As a matter of arithmetic, the value of trade depends on the number of goods traded, the amount of each good that is shipped, and the prices they are sold for. Most studies of trade volumes have not distinguished among these three factors. In this paper, we show that focusing on how the number of traded goods and their prices differ as a function of trade costs and market size, turns out to be very informative about the ability of trade theory to match trade data.

We establish some new facts about United States trade using the most disaggregated data publicly available. First, most potential export flows are not present, and the incidence of these “export zeros” is strongly correlated with distance and importing country size. Second, export unit values are positively related to distance

*Baldwin: Graduate Institute, 132 rue de Lausanne CH-1202 Geneva, (e-mail: Baldwin@GraduateInstitute.ch); Harrigan: Department of Economics, University of Virginia, Charlottesville, VA 22904-4182 (e-mail: james.harrigan@virginia.edu). Both authors are Research Associates at the NBER, and Baldwin is a Research Fellow at the CEPR. This paper, the first draft of which circulated in June 2007, has benefited from audience comments at numerous academic departments and conferences in the US and Europe. We thank the referees of this journal for insightful comments and suggestions. The authors acknowledge financial help from the Swiss NCCR Trade grant (phase I).

[†] To comment on this article in the online discussion forum, or to view additional materials, visit the article page at <http://www.aeaweb.org/articles.php?doi=10.1257/mic.3.2.60>.

¹ An important exception to this principle is Robert Feenstra, James Markusen, and Andrew Rose (2001), who show and test how different trade models imply different variations on the gravity model. Our paper is similar in approach, though different in focus.

and negatively related to market size. We show that every well-known, multi-good general equilibrium trade model is inconsistent with at least some of these facts.

We conclude the paper with a variant of the heterogeneous-firms trade model pioneered by Marc J. Melitz (2003). Our model maintains the core structure of Melitz (2003), namely heterogeneity in firms' productivity with fixed market entry costs, but introduces a taste for quality so that the lowest priced goods are not necessarily the most competitive. In our model, firms' competitiveness depends upon their quality-adjusted price and, in equilibrium, higher quality goods are more costly, more profitable, and better able to penetrate distant markets. Our model's predictions are borne out by the facts established in our data analysis.

In Section II, we generate testable predictions concerning the spatial pattern of trade flows and prices. The predictions come from three multi-good general equilibrium models that are representative of a wide swath of mainstream trade theory—one based purely on comparative advantage (Jonathan Eaton and Samuel Kortum 2002), one based purely on monopolistic competition (a multi-country Elhanan Helpman and Paul R. Krugman (1985) model with trade costs), and one based on monopolistic competition with heterogeneous firms and fixed market-entry costs (a multi-country version of the Melitz (2003) model). These models predict very different spatial patterns of zeros, i.e., the impact of country size and bilateral distance on the likelihood of two nations trading a particular product. They also generate divergent predictions on how observed trade prices should vary with bilateral distance and country size.

Section III confronts these theoretical predictions with product level US data on bilateral trade flows and unit values. On the quantity side, we focus on the pattern of zeros in product-level, bilateral trade data since this data contains information that is both rich and relatively unexploited.² Another advantage of focusing on zeros (the extensive margin) rather than volumes of positive flows (the intensive margin), is that it allows us to avoid issues such as the indeterminacy of trade flows at the product level in comparative advantage theory and the lack of data on firms' cost functions. On the price side, we focus on bilateral, product-specific "free on board" (f.o.b.) unit values.

All three mainstream models fail to explain the broad outlines of the data along at least one dimension. The best performance is turned in by the heterogeneous-firms trade (HFT) model based on Melitz (2003). However, this model fails to account for the spatial pattern of trade prices, in particular the fact that average unit values clearly increase with distance while the HFT model predicts that they should decrease with distance. Section III presents the modified Melitz model. As explained, a nation's high-quality/high-price goods are the most competitive, so they more easily overcome distance-related trade costs and this accounts for why observed prices in distant markets tends to be higher.

²We are not the first to exploit zeros. Jon Haveman and David Hummels (2004) is similar in spirit to our paper, although they focus on import zeros. Feenstra and Rose (2000) and Tibor Besedeš and Thomas J. Prusa (2006a, 2006b) have looked at time-series variation in product-level zeros to test trade models.

I. Zeros and Prices in Theory

This section derives testable hypotheses concerning the spatial pattern of zeros and trade prices in models that represent a broad swath of trade theory. In the three models selected, trade is driven by comparative advantage, monopolistic competition, and monopolistic competition with heterogeneous firms.

The models we study share some assumptions and notation. There are C countries and a continuum of goods. Preferences are given by the indirect CES utility function

$$(1) \quad U = \frac{E}{P}, \quad P \equiv \left(\int_{i \in \Theta} p_i^{1-\sigma} di \right)^{\frac{1}{1-\sigma}}, \quad \sigma > 1,$$

where E is expenditure, p_i is the price of variety i , and Θ is a set of available varieties. Transport costs are assumed to be of the iceberg form, with $\tau_{od} \geq 1$ representing the amount of a good which must be shipped from “origin” nation- o for one unit to arrive in “destination” nation- d . All the models assume just one factor of production, labor, which is in fixed supply and paid a wage w .

In the theory we discuss, the definition of a product is perfectly clear. An issue of interpretation, however, arises when confronting model predictions with data. In our recap of each model’s predictions for zeros and prices across space, we specifically address the mapping of the theory predictions to the data. As a background for these mappings, we briefly consider the essential features of the data. In Sections III and IV, we work with annual data on the least aggregated trade data that is publicly available, the US ten-digit level of the Harmonized System (HS10). We refer to the HS10 categories as “codes” and recognize that the HS10 codes do not generally correspond to a single product variety. Some examples of HS10 codes exported by the United States in 2005 include: women’s or girls’ wool sweaters (HS6110110020), new motor vehicle engine between 1500–3000 cc, more than 6 cylinders (HS8703230075), bicycles with wheels greater than 63.25 cm diameter (HS8712002600), undercarriages and parts for use in civil aircraft (HS8803200010), and lasers, other than laser diodes (HS9013200000). Plainly, there is ample scope for distinct products being sold in these and most other HS10 codes. Andrew B. Bernard, J. Bradford Jensen, and Peter K. Schott (2009) report that more than a quarter million different firms exported from the US in 2000 in just over 8,500 HS10 codes, or more than 30 firms per code on average.³ As a consequence, we are careful in translating firm- and/or variety-level predictions into predictions for HS10 codes.

A. Comparative Advantage: Eaton-Kortum

Economists have been thinking about the effects of trade costs on trade in homogeneous goods since Ricardo, but we had to wait for Eaton and Kortum (2002) to get a clear, rigorous, and flexible account of how distance affects bilateral trade in a

³The exact calculation is $(265,644 \text{ exporting firms}) / (8,572 \text{ HS10 codes}) = 31.0 \text{ exporting firms per HS10 code}$. Data on number of exporting firms is from Bernard, Jensen, and Schott (2009), Table 14.4.

competitive general equilibrium trade model. Online Appendix 1 presents and solves a slightly simplified Eaton-Kortum model (EK for short) explicitly. Here we provide intuition for the EK model's predictions on the spatial pattern of zeros and prices.

Countries in the EK model compete head-to-head in every market on the basis of “cost, insurance, and freight” (c.i.f.) prices, with the low-price supplier capturing the whole market.⁴ This “winner takes all” form of competition means that the importing country buys each good from only one source. As usual in Ricardian models, the competitiveness of a country's goods in a particular market depends upon the exporting country's technology, wage, and bilateral trade costs—all relative to those of its competitors. A key novelty of the EK model is the way it describes each nation's technology. The EK model does not explicitly specify each nation's vector of unit-labor input coefficients (the a 's in Ricardian notation). Rather it views the national vectors of a 's as the result of a stochastic technology-generation process—much like the one used later by Melitz (2003). Denoting the producing nation as nation- o ($'o'$ for origin), and the unit labor coefficient for a typical good- j as $a_o(j)$, each $a_o(j)$ is an independent draw from the cumulative distribution function (cdf)⁵

$$(2) \quad F_o[a] = 1 - e^{-T_o a^\theta}, \quad a, T_o \geq 0, \quad \theta > 1, \quad o = 1, \dots, C,$$

where $T_o > 0$ is the nation-specific parameter that reflects the nation's absolute advantage, and C is the number of nations.

Equation (2) permits calculation of the probability that a particular nation has a comparative advantage in a particular market in a typical good. Since the $a_o(j)$'s for all nations are random variables, determining comparative advantage becomes a problem in applied statistics. Perfect competition implies that nation- o will offer good- j in destination nation- d at a price of $p_{od}(j) = \tau_{od} w_o a_o(j)$ where w_o is the nation- o wage. As the Appendix shows, this implies that the distribution of prices in market- d in equilibrium is

$$(3) \quad G_d[p] = 1 - \exp[-\Delta_d p^\theta], \quad \Delta_d \equiv \sum_{c=1}^C T_{cd}, \quad T_{cd} \equiv \frac{T_c}{(w_c \tau_{cd})^\theta}.$$

Given (2) and (3), the probability that origin nation- o has comparative advantage in destination nation- d for any product is $\pi_{od} = T_{od}/\Delta_d$.⁶ This probability of observing positive trade between nations o and d in a given good is the key to characterizing the spatial pattern of zeros. It reflects the relative competitiveness of nation- o 's goods in market- d . Namely, T_{od} is inversely related to o 's average unit-labor cost for goods delivered to market- d , so π_{od} is something like the ratio of o 's average unit-labor cost to that of all its competitors in market- d .⁷

⁴c.i.f. and f.o.b. stand for “cost, insurance, and freight” and “free on board,” respectively, i.e., the price with and without transport costs. Without domestic sales taxes, c.i.f. and f.o.b. correspond to the consumer and producer prices respectively.

⁵EK work with $z = 1/a$, so their cdf is $\exp(-T/z^\theta)$.

⁶Technology draws are independent across goods, so this is valid for all goods. See the online Appendix for details.

⁷Nation- o 's unit-labor cost, averaged over all goods, for goods delivered to d is $t_{od} w_o / T^{1/\theta}$; this equals $1/(T_{od})^\theta$.

The EK model does not yield closed-form solutions for equilibrium wages, so a closed form solution for π_{od} is unavailable. We can, however, link the T_{od} s to observable variables by employing the market clearing conditions for all nations (see Appendix for details). In particular, wages must adjust to the point where every nation can sell all its output and this implies

$$(4) \quad \pi_{od} = Y_o \left(\frac{P_d}{\tau_{od}} \right)^\theta \left(\frac{1}{Y_d P_d^\theta + \sum_{c \neq d} Y_c (P_c / \tau_{oc})^\theta} \right),$$

where the Y 's are nations' total output (GDP) and the P 's are nations' price indices from (1) for the continuum of goods $\Theta = [0, 1]$. While this is not a closed form solution, the export probability is expressed in terms of endogenous variables for which we have data or proxies.

Equation (4) gives the sensible prediction that the probability of o successfully competing in market- d is decreasing in bilateral transport costs. The incidence of export zeros (that is, products exported to at least one but not all potential markets) should be increasing in distance if distance is correlated with trade costs, and import zeros (products imported from at least one potential source but not all) should predominate since each importer buys each good from just one supplier. Finally, Equation (4) also shows that *all* of the trade costs facing nation- d affect country o 's competitiveness in d , through the effect of bilateral trade costs on d 's price level P_d . The expression shows the probability of trade is higher when the destination market d has a high price level, and d 's price level depends on all of its bilateral trade costs through the term Δ_d in equation (3). The intuition is that if d faces high average trade costs, it will have high local prices and thus be relatively easy to penetrate for any given bilateral trade cost. In what follows, we will refer to this relative location effect as a country's remoteness.

The role of market size in determining the probability of exporting can also be studied with (4). The bigger is market- d , as measured by its GDP Y_d , the *smaller* is the probability that o successfully sells in d . There are two elements explaining this counterintuitive result. Large countries must sell a lot so they need, on average, low unit-labor-costs (as measured by w_o^θ / T_o). This means that large nations are often their own low-cost supplier. The second is that there are no fixed market-entry costs, so an exporter will supply every market where it has the lowest c.i.f. price—regardless of how tiny the market. Expression (4) also predicts that nation- d imports more goods from larger exporters, with size measured by Y_o .

The EK model makes simple predictions for the spatial distribution of import prices. The distribution of prices inside nation- d is given by (3) and each exporting nation has a constant probability of being the supplier of any given good. Consequently, the c.i.f. price of nation- o 's exports to nation- d is just a random sample from (3), which means (3) also describes the distribution of import prices for every exporting nation. The average c.i.f. price of goods imported from every partner should be identical and related to nation- d price index by (see Appendix for details)

$$(5) \quad p_{od}^e = P_d \left(\frac{\Gamma[(1 - \sigma + \theta)/\theta]}{\Gamma[(1 + \theta)/\theta]} \right),$$

where $\Gamma[\cdot]$ is the Gamma function. Since trade costs are fully passed on under perfect competition, the average bilateral export f.o.b. price, p_{od}^e/τ_{od} , should be increasing in the destination nation's price index (remoteness) and declining in bilateral distance.

The EK model is a multi-country Ricardian model with trade costs. In all Ricardian models, the locus of competition is destination nations; exporters must meet the competitive demands in each nation if they are to export successfully. Given this basic structure, the prediction of equal average import prices from all source nations is quite robust. Putting it differently, highly competitive nations export a wider range of goods than less competitive nations but the average import price of their goods does not vary with exporter's competitiveness, size or distance from the importing market. Staying in the Ricardian-Walrasian framework limits the range of modifications and extensions, so most extensions and modifications of the EK model lead to quite similar spatial predictions for zeros and prices.

One important extension of EK is Bernard et al. (2003). This model introduces imperfect competition into the EK framework with the low-cost firm in each market engaging in limit pricing. Limit pricing ties the market price to the marginal cost of the second-best firm, rather than the first-best as in EK. However with randomly generated technology, the outcome for the spatial pattern of zeros and prices is qualitatively unaltered.

In terms of HS10 codes, the export-zero prediction is that the probability that no firm exports in a given code from o to d is increasing in the distance between o and d and in the size of d . The export price prediction is that, within a given code, higher-cost products are less likely to be shipped longer distances and to less remote markets. As a consequence of this composition effect, the average export price within a code will be decreasing in distance, increasing in remoteness, and unrelated to size.

B. Monopolistic Competition

Monopolistic competition (MC) models constitute another major strand in trade theory. The particular MC model that we focus on has C countries, iceberg trade costs, and a single factor of production L . Goods are produced under conditions of increasing returns and Dixit and Stiglitz' (1977) monopolistic competition. Firms are homogenous in that they all face the same unit-labor requirement, a . As the model is very standard, we move quickly (see Appendix for details).

Dixit and Stiglitz' (1977) competition implies that "mill pricing" is optimal, so firms charge the same f.o.b. export price regardless of destination (iceberg trade costs are fully passed on to consumers). Using this and the demand function implied by (1), the value of bilateral exports for each good is

$$(6) \quad v_{od}[a] = \phi_{od} \left(\frac{w_o a}{1 - 1/\sigma} \right)^{1-\sigma} B_d; \quad \phi_{od} \equiv \tau^{1-\sigma} \in [0, 1], \quad B_d \equiv \frac{w_d L_d}{P_d^{1-\sigma}},$$

where $v_{od}[a]$ is the value of bilateral exports for a nation- o firm with marginal cost a , ϕ_{od} reflects the "freeness" of bilateral trade (ϕ ranges from zero when τ is prohibitive to unity under costless trade, i.e., $\tau = 1$), B_d is the per-firm demand-shifter in market- d , and P_d is as in (1), with Θ being the set of goods sold in d .

Spatial patterns are exceedingly simple in the MC model. Consumers buy some of every good with a finite price, so there should be no zeros in the trade matrix—if a good is exported to one country it is exported to all. Neither the size of the destination market nor the distance between the origin and destination markets has any bearing on the probability of observing a zero. This outcome is, of course, due to the model's dramatic simplifications meant to explain pervasive intra-industry trade (the widespread existence of zeros was not known when the model was developed in the 1980s).

The predictions for the spatial pattern of trade prices are equally sharp. As mill-pricing is optimal, trade costs are passed fully on to consumers, and since all products are sold in all markets, the f.o.b. export prices to all destinations should be identical (and thus unrelated to distance, market size, or remoteness).

These stark predictions depend upon standard simplifying assumptions. Some of these have been relaxed in the theory literature, so we consider the predictions that are robust across well-known variants of the MC model.

The core elements of the MC model are imperfect competition, increasing returns, and homogenous firms. Since imperfect competition can take many forms, many variants of the standard MC model are possible. The general formula for optimal pricing under monopolistic competition sets consumer price to $a\tau_{od}/(1 - \varepsilon_d^{-1})$, where ε_d is the perceived elasticity of demand in market- d . Different frameworks link ε_d to different parameters. Under Dixit-Stiglitz monopolistic competition, ε_d always equals σ —a constancy that is the heart of the simple spatial pattern of zeros and prices in the standard MC model.

Under the Gianmarco Ottaviano, Takatoshi Tabuchi, and Jacques-Francois Thisse (2002) monopolistic competition ε_d framework, firms face linear demand so ε_d rises as firms move up the demand curve. This creates a link between bilateral pricing and distance via bilateral markups. In other words, producers absorb some of the trade costs, so the f.o.b. export prices should be lower for more distant markets. Linear demand also features a choke-price for consumption and this opens the door to zeros in the spatial trade pattern, even without beachhead costs (fixed market-entry costs). The predicted pattern of zeros with homogeneous firms is very stark. Nation- o 's export matrix has either no zero with respect to nation- d , or all zeros, with the likelihood of observing such zeros rising with the distance to the destination market. Moreover, as the equilibrium number of varieties consumed is higher in bigger markets, and a rise in this number makes the market more competitive (i.e., it lowers the demand curve intercept), it should be more likely to observe zeros for larger destination-markets. Likewise, as remoteness makes a destination market less competitive, *ceteris paribus*, there should be a negative conditional correlation between zeros and remoteness.

Finally, if demand curves are sufficiently convex, higher bilateral trade costs raise the markup and this implies that f.o.b. prices can rise with distance. Because this degree of convexity implies a counterfactual more-than-full pass-through of cost shocks (e.g., more than 100 percent exchange rate pass-through) such demand structures are not typically viewed as part of the standard MC model.

In terms of HS10 codes, the export zero prediction with linear demand is that the probability that no firm exports in code i from o to d is decreasing in the distance

between o and d . Unlike in the EK model, the export price prediction with linear demand has nothing to do with a composition effect. Rather, the prediction is driven by the reduction in markups with distance.

C. The Heterogeneous-Firms Trade Model

One of the beauties of the Melitz (2003) heterogeneous-firms trade (HFT) model is that it provides a clean and convincing story about why some products are not exported at all. But since Melitz (2003) works with symmetric countries, it cannot address the spatial pattern of export zeros or export prices. To generate testable hypotheses concerning zeros and prices, Melitz (2003) must be modified to allow asymmetric countries and arbitrary bilateral trade costs. Such modifications, however, pose difficult general equilibrium problems linked to the simultaneous determination of wages, export thresholds and the mass of firms in each nation. The standard modification that allows country asymmetries—introduced by Helpman, Melitz, and Stephen R. Yeaple (2004)—gets around the general equilibrium problems by including a costlessly-traded, constant-returns numeraire sector that pins down wage rates globally. This theoretically convenient artifice, however, implies that the per-firm demand levels—the B_d —are independent of market size (see Appendix). This in turn leads to the unintuitive prediction that, controlling for distance, a given export firm sells the same to every market regardless of market size.⁸ To avoid these pitfalls, we take an approach that is closer to the original Melitz (2003) set up.

The Melitz model embraces all of the features of the baseline monopolistic competition model and adds in two new elements—beachhead costs denoted as F , and heterogeneous, firm-level, marginal costs. This firm-level marginal-cost heterogeneity is introduced—as in the EK model—via a stochastic technology-generation process. When a firm pays its set-up cost, it simultaneously draws a unit labor coefficient from the Pareto distribution⁹

$$(7) \quad G[a] = \left(\frac{a}{a_0}\right)^\kappa; \quad \kappa > 1, \quad 0 \leq a \leq a_0,$$

where κ is the shape parameter and a_0 is the maximum a ; as a regularity condition, $\kappa > \sigma - 1$. After seeing its a , the firm decides how many markets to enter.

The assumed Dixit-Stiglitz competition means that: (i) a firm's optimal price is proportional to its marginal cost, (ii) its operating profit is proportional to its revenue, and (iii) its revenue is inversely proportional to its relative price. Thus, a firm that draws a relatively high marginal cost earns little if it produces. If this amount is insufficient to cover the beachhead cost F in any market, the firm never produces. Firms that draw lower marginal costs may find it profitable to enter some markets (especially their local market where the absence of trade costs provides them with a relative

⁸The economic logic rests on entry. The B 's in larger markets have higher numerators (higher wL), but they also have more firms in equilibrium so the denominator is also higher. Indeed, free entry exactly offsets country-size differences, equalizing the equilibrium B 's across all nations.

⁹The HFT model is easily solved with an exponential cdf as in the EK model, but the Pareto is more traditional and can be justified by reference to data on the US firm size distribution (Robert L. Axtell 2001).

cost advantage). More generally, for each export market there is a threshold marginal cost for every origin nation; firms with marginal cost draws above the threshold do not sell their production in the market concerned as doing so would yield operating profit insufficient to cover the beachhead cost. Using (6), which shows the revenue a nation- o firm would earn if it sold in market- d , and the constant operating-profit share (i.e., $1/\sigma$), we can find the cut-off conditions for each trade pair. Defining the bilateral threshold for firms based in nation- o and selling in nation- d as a_{od} , defined by $v_{od}[a_{od}]/\sigma = F$, where v_{od} is as in (6).¹⁰ Rearranging, the pairwise cutoff is

$$(8) \quad w_o a_{od} = \frac{1}{\tau_{od}} \left(\frac{B_d}{f} \right)^{\frac{1}{\sigma-1}},$$

where f is the collection of parameters, $\sigma F(1 - 1/\sigma)^{1-\sigma}$. The HFT model also involves a free entry condition (see Appendix), but this does not come into play in characterizing the spatial pattern of zeros and prices.

As we show in the Appendix, it is not possible to solve this model analytically, but numerical simulation is straightforward. We show that larger countries have larger per-firm demands, and as a result any given firm sells more to larger countries than to smaller ones (controlling for trade costs). The mechanism is that larger countries have endogenously higher wages, which leads to less entry than there would be if wages were equalized. A way to see the contrast between the two solutions is to note that incipient entry raises the demand for labor in the larger country, but that this has different effects depending on the behavior of wages. When wages are fixed by the numeraire sector, all the adjustment takes place through entry. Without the numeraire sector, part of the adjustment comes through higher wages in the larger country which in turn dampens entry, leaving per-firm demand higher in larger countries (see Appendix for details).

The spatial pattern of zeros comes from the cut-off conditions. The probability that a firm producing variety- j with marginal costs $w_o a_o(j)$ will export to nation- d is the probability that its marginal cost is less than the threshold defined in (8), namely

$$(9) \quad \Pr \left\{ a(j) < \frac{B_d^{1/(\sigma-1)}}{\tau_{od} w_o^\sigma f^{1/(\sigma-1)}} \right\} = \frac{B_d^\beta}{\tau_{od}^\kappa w_o^{\sigma\kappa} f^\beta a_0^\kappa},$$

where β equals $\kappa/(\sigma - 1)$ and we use (7) to evaluate the likelihood.

In our empirics, we only have data on products that are actually exported to at least one market so it is useful to derive the expression for the conditional probability, i.e., the probability that a firm exports to market d given that it exports to at least one market. This conditional probability of exports from o to d by a typical firm is

$$(10) \quad \frac{\tau_{od}^{-\kappa} B_d}{\min_{i \neq 0} \tau_{oi}^{-\kappa} B_i}.$$

¹⁰Here we have chosen units (without loss of generality) such that $a_0 = 1$. We also assume the beachhead costs are the same in all markets (no difference between home and foreign markets) for notational simplicity.

Since we work with data for a single exporting nation, the denominator here will be a constant for all products. Plainly, the probability of a positive trade flow decreases with τ (proxied by bilateral distance in our empirics), and increases in the firm-specific demand shifter, B . As B is increasing in a destination market's size and price index P , it is more likely to observe positive flows with destination markets that are large and remote, when other factors are controlled for.¹¹

The spatial pattern of prices is also simple to derive. We consider both the export (f.o.b.) price for a particular good exported to several markets, and the average export (f.o.b.) price for all varieties exported by a particular nation. Because mill pricing is optimal for every firm, a given firm's f.o.b. export price is predicted to be identical for all destinations. In particular, the export price for any given good should be unrelated to bilateral distance and unrelated to the size or remoteness of the origin and destination nations. However, the range of goods exported does depend on bilateral distance, size, and remoteness, so the average bilateral f.o.b. export price will also depend on these. The cut-off conditions (8) imply

$$(11) \quad \bar{p}_{od} = \int_0^{a_{od}} \left(\frac{\tau_{od} w_o a}{1 - 1/\sigma} \right)^{1-\sigma} dG(a_{od} \mid a \leq a_{od}),$$

where \bar{p}_{od} is the weighted average of f.o.b. prices of all varieties exported from nation- o to nation- d . Using (7) to evaluate the integral and eliminating the cutoff a_{od} using (8), this becomes

$$(12) \quad \bar{p}_{od} = \frac{\delta_o}{\tau_{od}^\kappa} \left(\frac{B_d}{f} \right)^{\frac{1+\kappa-\sigma}{\sigma-1}},$$

where δ_o is a function of parameters and country- o variables. Equation (12) implies that the ratio of average export price to any two markets, say c and d , from a single origin nation- o depends only on relative distance from nation- o and the relative firm-level demand shifters (the B 's),

$$(13) \quad \frac{\bar{p}_{od}}{\bar{p}_{oc}} = \left(\frac{\tau_{oc}}{\tau_{od}} \right)^\kappa \left(\frac{B_d}{B_c} \right)^{\frac{\kappa+\sigma-1}{\sigma-1}}.$$

The logic of (13) is that since the cut-off marginal cost, a_{od} , falls with bilateral distance, and increases with market size and remoteness, the average export price of nation- o varieties to nation- d should be decreasing in distance, and increasing in the size and remoteness of the export market. The intuition is that the cheapest goods are the most competitive in the Melitz (2003) model, so they travel the farthest.

¹¹ As c.i.f. prices in remote markets tend to be higher, the price index tends to be higher, so B tends to be higher.

The HFT model, like the monopolistic competition model, has imperfect competition and increasing returns as core elements. As noted above, there are many different forms of imperfect competition and scale economies. The other core elements of the HFT model are beachhead costs or choke-prices (these explain why not all varieties are sold in all markets) and heterogeneous marginal costs (these explain why some nation-*o* firms can sell in a market but others cannot). This suggests three dimensions along which HFT models can vary: market structure, source of scale economies, and source of heterogeneity.

For example, Eaton, Kortum, and Francis Kramarz (2008) present a model that incorporates beachhead and iceberg costs in a setting that nests the Ricardian framework of Eaton and Kortum (2002) and Bernard et al. (2003) with the monopolistic competition approach of Melitz (2003). The model, however, does not assume set-up costs and does not deal with entry since firms are assumed to be endowed with a technology draw. Another difference is that they work with a technology-generating function from the exponential family. Svetlana Demidova (2008), and Rod Falvey, David Greenaway, and Zhihong Yu (2006) allow for technological asymmetry among nations but embrace Dixit-Stiglitz competition with iceberg, beachhead and set-up costs. Yeaple (2005) assumes the source of the heterogeneous marginal costs stems from workers who are endowed with heterogeneous productivity and works with Dixit-Stiglitz monopolistic competition with iceberg and beachhead costs. As these models all involve Dixit-Stiglitz monopolistic competition, iceberg and beachhead costs, their spatial predictions for zeros and price are qualitatively in line with the model laid out above.

Melitz and Ottaviano (2008) work with the Ottaviano, Tabuchi, and Thisse (2002) monopolistic competition framework with linear demands. The implied choke-price substitutes for beachhead costs in shutting off the trade of high-cost varieties. The Melitz and Ottaviano (2008) prediction for the spatial pattern of zeros matches that of the asymmetric HFT model with respect to bilateral distance, but not with respect to market size. As our Appendix illustrates, the cut-off marginal cost in Melitz and Ottaviano (2008) is tied to the *y*-axis intercept of the linear residual demand curve facing a typical firm. More intense competition lowers this intercept (this is how pure profits are eliminated in the model) and thus the cut-off, a_{od} , falls with the degree of competition. Since large nations and nations that are not very remote have more intense competition from local and imported varieties, Melitz and Ottaviano (2008) predicts that large countries should have lower cut-offs, controlling for the nation's remoteness. In other words, Melitz and Ottaviano (2008) predicts a positive relationship between the size of the partner nation and the number of zeros in an exporter's matrix of bilateral, product-level exports.

Predictions for the spatial pattern of prices are also decided by the intercept, which is itself affected by the local price index, P_d . Thus, Melitz and Ottaviano (2008) predicts that export prices should decline with the market's distance and with market size. As remoteness dampens competition in nation-*d*, the export price from nation-*o* should rise with *d*'s remoteness.

In terms of HS10 codes, the export-zero prediction refers to the probability that no variety is exported in a particular code from *o* to *d*. The export-price prediction is driven by a composition effect within HS10 codes, as was the export-price prediction in the EK model.

TABLE 1—SUMMARY OF MODEL PREDICTIONS

	Pr(export zero)			f.o.b. export price		
	Distance	Importer		Distance	Importer	
		Size	Remoteness		Size	Remoteness
Eaton-Kortum	+	+	−	−	0	+
Monopolistic competition, CES	0	0	0	0	0	0
Monopolistic competition, linear demand	+	0	−	−	0	+
Heterogeneous firms, CES	+	−	−	−	+	+
Heterogeneous firms, linear demand	+	+	−	−	−	+
Heterogeneous firms, CES, quality competition	+	−	−	+	−	−

Notes: The first five rows of the table summarize the theoretical comparative static predictions discussed in Section I, with the last row giving the predictions of the model that we develop in Section III. The six models under discussion are listed in the first column. Each entry reports the effect of an increase in distance or importer size on the probability of an export zero or f.o.b. export price. An export zero is defined to occur when a country exports a good to one country but not all.

II. Zeros and Prices in Trade Data

The models described in the previous section all make predictions about detailed trade data in a many country world. These predictions are collected for easy reference in the first five rows of Table 1.

To evaluate these predictions, we use detailed data on imports and exports, specifically the trade data collected by the US Customs Service and made available on CD-ROM. For both US imports and US exports, the census reports data for all trading partners classified by the 10-digit Harmonized System (HS). For each country-HS10 record, census reports value, quantity, and shipping mode. In addition, the import data include shipping costs and tariff charges. Our data analysis also includes information on distance and various macro variables, which come from the Penn-World Tables.

The census data are censored from below, which means that very small trade flows are not reported. For imports, the cut-off is \$250, so the smallest value of trade reported is \$251. For exports, the cut-off is 10 times higher, at \$2,500. This relatively high censoring level for exports is a potential problem, since there might be many economically meaningful export relationships which are inappropriately coded as nonexistent. One hint that this problem is not too prevalent comes from the import data, where only 0.8 percent of the nonzero trade flows are between \$250 and \$2,500.

A. How Many Zeros?

We define a zero as a trade flow which could have occurred but did not. For exports, a zero occurs when the US exports an HS10 code to at least one country but not all. The interpretation of an export zero is simple, since they are defined only for goods actually produced. For imports, a zero is an HS10 code which is imported from at least one country but not all. The interpretation of an import zero is not as simple; they may be defined in cases when the country in question does not even produce the good (e.g., the US imports no bananas from Canada).

TABLE 2—INCIDENCE OF ZEROS IN US TRADE, 2005

	Imports	Exports
All countries		
Trading partners	228	230
HS10 products	16,843	8,880
Partners \times products	3,840,204	2,042,400
Percent zeros	92.6	82.2
100 largest countries		
HS10 products	16,843	8,880
Partners \times products	1,684,300	888,000
Percent zeros	85.5	70.0

The incidence of zeros in US trade in 2005 is reported in Table 2. The US imports nearly 17,000 different HS10 categories from 228 countries, for a total of over 3.8 million potential trade flows. Over 90 percent of these potential trade flows are zeros. The median number of supplier countries was 12, with a quarter of goods being supplied by at least 23 countries. Only 5 percent of codes have a unique supplier. In principle this pattern of imports is consistent with a homogeneous goods model, since if we define a good narrowly enough it will have just one supplier. However, the large number of suppliers for the majority of narrowly defined codes seems instead to be suggestive of product differentiation.¹² This well-known phenomenon is part of what motivated the development of monopolistic competition trade models.

Zeros are almost as common in the export data as in the import data. Table 2 shows that in 2005 the US exported 8,880 HS10 codes to 230 different destinations, for a total of more than 2 million potential trade flows. Of these, 82 percent are zeros. Unlike the import zeros these have an unambiguous interpretation, since a zero is defined only if a good is exported to at least one country, which necessarily means the good is produced in the US. The median number of export markets is 35, with a quarter of codes exported to at least 59 markets. Only 1 percent of codes are sent to a unique partner.

Many of the 230 places that the US trades with are tiny to the point of insignificance (Andorra, Falkland Islands, Nauru, Pitcairn, Vatican City, and the like). Restricting attention to the 100 large countries for which we have at least some macroeconomic data reduces the incidence of zeros somewhat (86 percent for imports, 70 percent for exports), but does not fundamentally change the message that zeros predominate.

From the standpoint of theory, the predominance of HS10 code zeros almost certainly understates the number of zeros at the level of firms and/or products. This is because, as noted above, each HS10 code contains products from an average of 31 different firms, and each firm-product in a given code might be exported to a proper subset of the countries with positive purchases in that code.

The predominance of export zeros suggests that the baseline monopolistic competition model with CES preferences discussed in Section II misses important aspects of the world trade pattern (the model was not designed with zeros in mind as these

¹²The largest number of trading partners for any product is 125, for product 6204.52.2020, "Women's trousers and breeches, of cotton, not knitted or crocheted." This is not a homogeneous goods category.

facts were unknown when it was developed). This suggests that including either beachhead costs or choke-off prices is essential to mapping even the first order features of real world trade patterns.

B. Export Zeros Across Space

The gravity equation offers a flexible and ubiquitous statistical explanation for the aggregate volume of trade between countries. The basic logic of the gravity equation is simplicity itself: bilateral trade volumes depend positively on country size and negatively on distance. As emphasized by James E. Anderson and Eric van Wincoop (2003), countries' relative distance from their trading partners also affects bilateral trade volumes, an effect visible in the theory models discussed in the previous section and that we are calling "remoteness." Here we adopt the gravity approach to explain not the *volume* of trade but rather its *incidence*. This descriptive statistical exercise is intended to help us understand the pattern of export zeros summarized in Table 2.

We emphasize that the purpose of our statistical exercises in this and the following section is to document robust reduced form relationships in the data. Any structural model of trade should be consistent with these correlations. We contrast our approach with the well-known recent paper by Helpman, Melitz, and Yona Rubinstein (2008); that paper structurally estimates the distance elasticity of trade in the context of a particular theoretical model. They also focus only on country-level zeros, that is, pairs of countries that have no trade in *any* product. As Table 2 and Figure 1 illustrate, product level zeros are a first-order feature of the data for US trade, while the US has positive trade in *some* products with virtually every country in the world.

We focus on US export zeros because of their unambiguous interpretation as products which the US produces and ships to at least one, but not all, countries. Extending the gravity logic suggests that exports should be more likely the larger the production of the good, the larger and more remote the export market, and the shorter the distance the good would have to travel. We have no information on production volumes by good, so we focus on within-product variation across export partners. As indicated in Table 2, the dimensions of the data in 2005 are 8,880 HS10 codes and 100 countries.

Figure 1 offers a first look at the data. The vertical axis shows market size (measured as log real GDP of the importing country) and the horizontal axis shows distance from the United States (measured as log kilometers between Chicago and the capital city of the importer). Each point is represented by a circle, where the size of the circle is proportional to the number of HS10 codes exported from the US to the importer. Canada and Mexico are identified, with other country names left out in the interest of legibility.

A number of patterns are visible in Figure 1. First, controlling for distance, size is associated with more codes being exported (to see this, pick a point on the log distance axis such as 8 or 9 and observe that the circles generally get bigger from bottom to top). Second, the small, distant exporters clustered in the lower right-hand corner of the plot have very small circles, meaning that few codes are exported to these markets. Overall, Figure 1 suggests a gravity relationship: the number of codes exported increases with market size and decreases with distance.

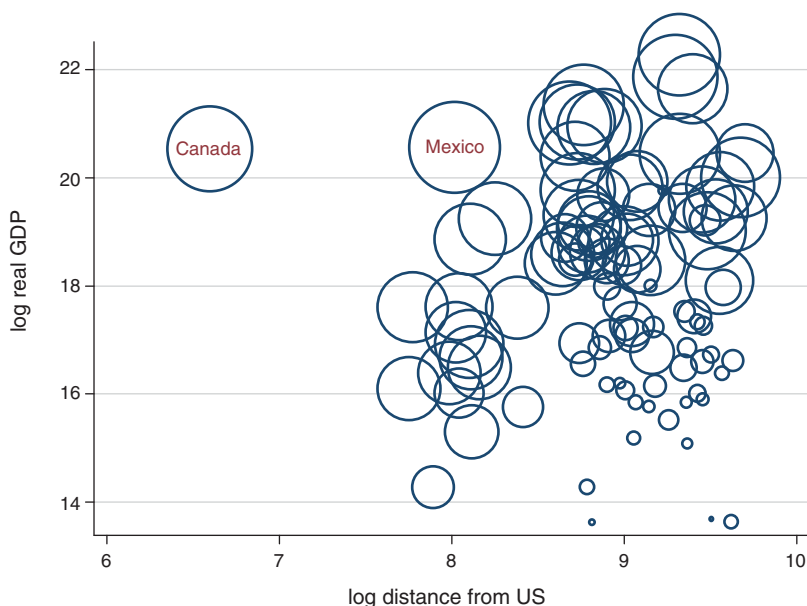


FIGURE 1. POSITIVE TRADE FLOWS, DISTANCE, AND MARKET SIZE

Notes: Each circle in the plot represents a US export market, and the size of the circle is proportional to the number of HS10 products exported by the US to that destination in 2005. Distance is measured as kilometers between Chicago and the capital city of each country.

We next turn to statistical analysis of the probability that a code is exported from the United States to a particular market. Our statistical model is motivated by the theory summarized in Table 1. All of the models suggest that both export prices and the probability of trade depend on trade costs between the US and the importer, importer market size, and importer remoteness. Here is how we measure these theoretical concepts:

Trade Costs.—The most obvious proxy for trade costs is distance, which we measure as kilometers from Chicago to the capital city of the importer.

While trade costs are likely to be weakly monotonic in distance, there is no reason to expect them to have any particular functional form, so we specify the distance proxy in two ways. The first is simply log distance. The second breaks distance down into bins, derived from looking for natural breaks in distance among US trading partners. These bins are described in Appendix Table 1. We also include several country-specific indicators of trade costs that have been often used in the gravity literature: dummies equal to one if the importer is landlocked, an island, shares a border with the US, and if the primary language is English.¹³

Market Size.—Our measure of market size is real GDP, from the Penn World Tables. We also include real GDP per capita as a demand-related control.

¹³Our source for the distance and indicator variables is <http://www.cepii.fr/anglaisgraph/bdd/distances.htm>.

Remoteness.—The proper specification of remoteness is model-specific, which creates a potentially important measurement issue for our purposes since we are interested in comparing the performance of various models in explaining a common set of facts. However, most theoretically consistent measures of remoteness have a common structure as they all work via the average price of goods sold in a destination market. This average price in turn depends upon the number of varieties produced locally in the destination market, and the number of imported varieties and the bilateral trade costs they face. As the number of varieties coming from each exporting nation is—roughly speaking—related to the origin-nation’s size, a reasonable proxy for remoteness involves market-size weighted sums of an inverse power function of trade costs. Following this logic, we adopt the following measure of remoteness in our empirical work,

$$R_d = \left[\sum_{o=1}^c Y_o d_{oc}^{-\eta} \right]^{-1}.$$

Harrigan (2003) shows that this remoteness index is an approximation to the model-specific measures of Anderson and van Wincoop (2003) and Eaton and Kortum (2002). Empirical implementation of R_d involves some potentially important choices about how to measure within-country distance d_{dd} and what value to use for the exponent η .¹⁴ Fortunately, our empirical results are entirely insensitive to any reasonable choice of how to construct R_d , and in what follows we include within country distance as reported in the CEPII data (cf. footnote 13), and set η equal to 1. The reason for this robustness is simply that the cross section variation in R_d is overwhelmingly dominated by differences in the GDP-weighted raw distances (consider New Zealand versus Belgium), so that different choices about including own distance and what value to choose for η lead to very highly correlated measures.

Our statistical approach models the incidence of positive US exports as a function of a linear index of distance, market size, and remoteness,

$$(14) \quad \Pr(y_{ic} = 1) = F(\beta \mathbf{x}_c + \varepsilon_{ic}),$$

where $y_{ic} = 1$ if positive exports of HS10 code i from the US to country c , \mathbf{x}_c is a vector of country c characteristics, F is a probability distribution function, β is a vector of parameters to be estimated, and ε_{ic} is an error term.

There are three sets of econometric issues involved in estimating (14): the functional form of F , the specification of the index function $\beta \mathbf{x}_c$, and computation of the covariance matrix of $\hat{\beta}$ (the estimate of β).

Functional Form.—The key specification issue is how to control for HS10-specific effects. This is quite important a priori, since different HS10 codes are likely to have

¹⁴In a CES monopolistic competition model, η is the elasticity of substitution minus 1, while in the Eaton-Kortum model η equals the dispersion parameter θ of the Fréchet distribution.

systematically different probabilities of being exported. The probit and linear probability models have different advantages and disadvantages, so we use both. The advantage of the probit is that it produces fitted probabilities that are between zero and one, but the disadvantage is that we cannot include HS10 fixed effects.¹⁵ With the linear probability model we can estimate HS10 fixed effects, but the assumption that the error term has unlimited range is incorrect. We also estimate a random effects probit model, which treats HS10 effects as random and orthogonal to the country-level regressors. Finally, to allow for different parameter values for different subsamples, we estimate the model on two samples: all HS codes, and only codes that map to SITC 6 (manufactured goods), 7 (machinery and transport equipment), and 8 (miscellaneous manufactures).

Index Function.—We specify three variants of the index function. The first simply enters distance, real GDP, real GDP per capita, and remoteness log-linearly. The second adds an interaction between log distance and log real GDP. The third replaces the log linear measure of distance with the distance bins discussed above (the excluded distance category includes Mexico and Canada, so we cannot estimate a separate border effect in this specification).

Covariance Matrix of $\hat{\beta}$.—Our data has a balanced panel structure, with each HS10 code i potentially exported to each importer c . There is sure to be population correlation in the errors ε_{ic} both across i for a given c , and across c for a given i . Recent econometric theory results due to A. Colin Cameron, Jonah B. Gelbach, and Douglas L. Miller (forthcoming) show how to handle such a situation using multi-way clustering of the errors by HS code i and importer c , and we implement this procedure for the probit and linear probability specifications. Put simply, the Cameron, Gelbach, and Miller (forthcoming) procedure allows us to assume arbitrary and nonconstant correlation within both c and i clusters, and we assume independence only across clusters, as is standard. Robust covariance matrix estimation is not applicable with the random effects probit estimator (see Wooldridge 2002, 485–86), so we report standard errors conditional on the (untenable) assumption that the errors are independent within, as well as across, i and c .

To summarize, we estimate 18 different versions of (14), namely three specifications \times three estimators \times two samples. The results are reported in Table 3 and Appendix Table 3. The first column of Table 3 reports our baseline results, which only includes manufacturing trade. The coefficients are the marginal effects on the probability of a code being exported to a particular country, conditional on being exported to at least one country, evaluated at the sample mean. The continuous variables are scaled to have sample standard deviation one, so the coefficients answer the question: how much does the probability of exporting increase with a one standard deviation increase in the explanatory variable? Distance has a powerful negative effect on the probability of exporting, with a one standard deviation increase in distance reducing the probability by about 0.13. Market size

¹⁵This is not a computational issue but a statistical one: the probit estimator with fixed effects is not a consistent estimator of β (see, for example, Jeffrey M. Wooldridge 2002, 483–84).

TABLE 3—STATISTICAL DETERMINANTS OF NONZERO US EXPORTS, 2005

Specification	Linear distance			Distance step function		
	Probit ctry., HS none	OLS ctry., HS fixed	RE probit none random	Probit ctry., HS none	OLS ctry., HS fixed	RE probit none random
Log distance	-0.1247*** (0.0156)	-0.0943*** (0.0144)	-0.1466*** (0.0018)			
1 < km ≤ 4,000				0.1056 (0.132)	0.0132 (0.0730)	0.1574*** (0.00712)
4,000 < km ≤ 7,800				-0.1507 (0.104)	-0.195*** (0.0732)	-0.1522*** (0.00473)
7,800 < km ≤ 14,000				-0.2169* (0.110)	-0.253*** (0.0744)	-0.2344*** (0.00548)
14,000 < km				-0.2050** (0.0662)	-0.288*** (0.0778)	-0.1890*** (0.00357)
Log real GDP	0.2032*** (0.0147)	0.1696*** (0.0146)	0.2400*** (0.00283)	0.1983*** (0.01425)	0.1679*** (0.07304)	0.2346*** (0.00281)
Log real GDP per worker	0.1208*** (0.0155)	0.0892*** (0.0158)	0.1424*** (0.00184)	0.1244*** (0.0178)	0.0895*** (0.0171)	0.1461*** (0.00194)
Log remoteness	0.0727*** (0.0094)	0.0633*** (0.01193)	0.0872*** (0.0012)	0.0528*** (0.0135)	0.0477*** (0.0136)	0.0636*** (0.00108)
NAFTA	-0.1653 (0.129)	-0.1020 (0.162)	-0.1653*** (0.00316)			
Landlocked	-0.04502 (0.0337)	-0.0207 (0.0296)	-0.0482*** (0.00222)	-0.0444 (0.0358)	-0.0208 (0.0310)	-0.0470*** (0.00228)
Island	-0.01311 (0.0310)	0.0299 (0.0327)	-0.0138*** (0.00217)	-0.01819 (0.0288)	0.0172 (0.0304)	-0.0200*** (0.00213)
English language	0.0963*** (0.0322)	0.0551** (0.0277)	0.1168*** (0.00214)	0.0768*** (0.0288)	0.0467* (0.0262)	0.0938*** (0.00202)

Notes: Dependent variable is indicator for positive exports in particular HS code to particular importing country. Manufacturing observations only (5,834 HS10 codes × 100 importing countries). Reported probit and random effects probit results are marginal effects, evaluated at sample mean for continuous variables. Continuous variables are scaled to have mean zero and standard deviation one, so coefficient gives effect on probability of a one standard deviation increase in the continuous RHS variable. For indicator variables, coefficient is change in probability from turning on indicator. Robust standard errors in parentheses, asterisks denote statistical significance. Except for random effects probit, standard errors are clustered on both country and HS10.

***Significant at the 1 percent level.

**Significant at the 5 percent level.

*Significant at the 10 percent level.

as measured by real GDP has a similarly powerful positive effect, with a one standard deviation increase in market size raising the probability of exporting by about 0.20, with the effect of real GDP per capita a bit smaller at 0.13. Remoteness also has a large positive effect, with a one standard deviation increase raising the export probability by 0.073. English-language importers have a 0.1 greater probability of importing a particular product. The other indicators (common border, landlocked, and island) do not have statistically significant effects on export probability. These inferences are very similar across the linear and random effects probit specifications, although as noted above the very small standard errors in the latter specification should be discounted. The effects are also quite similar, though a bit smaller, when the sample includes nonmanufacturing codes only (see Appendix Table 3).

The last three columns of Table 3 measure distance using a step function rather than as log kilometers.¹⁶ In the probit specification, distance less than 4,000 kilometers (which includes the Caribbean basin) has no effect on export probability, an effect which becomes -0.16 to -0.21 once distance is greater than 4,000 kilometers with no statistically significant differences in the effects of the three distance categories above 4,000 kilometers. The estimated effects are similar in the linear probability and random effects probit specifications, and when the sample is restricted to manufacturing goods. These results suggest a simple interpretation: when goods must cross an ocean the probability of exporting falls by about 0.2, but other variation in distance is of secondary importance. In fact, a simple specification where distance is replaced by a single indicator for crossing an ocean, which is equivalent to restricting the coefficients on the three indicators for distance greater than 4,000 kilometers to be the same, fits quite well (we don't report these results to save space). The estimated effects of the other explanatory variables are quite similar to their estimates from the linear distance specifications.

Appendix Table 3 reports results from the specification that includes an interaction between log distance and log GDP per capita. In the linear probability model all three parameter estimates are statistically insignificant, and in the probit specification the coefficient on real GDP is significant and negative at just the 10 percent level. Our interpretation of these imprecise results is that the data cannot reliably estimate an interaction effect.

In summary, the 18 different estimated versions of (14) tell a clear and consistent story: distance has a big negative effect on the probability that a good is exported, while larger, richer, and more remote markets are much more likely to import. While in many respects not surprising, it is worth noting that the results of Table 3 are *not* consistent with most of the models summarized in Table 1. Only the HFT model with CES preferences is consistent with the positive market size and remoteness effects and negative distance effects identified so strongly in the data.

In a recent paper, Roc Armenter and Miklós Koren (2008) show that correlations of zeros with distance and importer size such as those documented in Table 3 can be generated by a purely random assignment of trade shipments to product categories, conditional on aggregate bilateral trade volumes. Armenter and Koren (2008) argue that their results make it difficult to use empirical results about zeros to discriminate between trade models. As shown in Section I above and summarized in Table 1, the different models we study *do not* all predict that HS10-level zeros are correlated with distance and importer size, so the Armenter and Koren (2008) results are not relevant to how we interpret the empirical results of Table 3.

C. Export Unit Values Across Space

We now turn to a descriptive analysis of export unit values. The statistical model is very similar to the previous section, with within-HS10 code variation in unit

¹⁶The excluded distance category is zero kilometers (Canada and Mexico), which is why we cannot separately estimate a border effect in this specification.

values regressed on characteristics of the importing country including distance, market size, and remoteness:

$$(15) \quad p_{ic} = \alpha_i + \beta \mathbf{x}_c + \varepsilon_{ic},$$

where p_{ic} is the log of the f.o.b. average unit value of code i shipped to country c and the vector \mathbf{x}_c is the same as in the previous section. The code fixed effect α_i controls both for the average unit value of products (industrial diamonds versus peanuts) and differences in units (kilos versus bushels) across codes. Because this analysis uses only nonzero export observations, the sample size is much smaller than in the previous section, and the panel is highly unbalanced because the incidence of zeros varies widely across products. Estimation is by OLS with HS10 fixed effects, and standard errors are clustered by HS10 code and importing country. As in the previous section, we consider three different specifications estimated on different subsamples.

The residual in (15) can be given several interpretations. Probably the most natural is simple measurement error stemming from the fact that unit values are not the same as true prices. A related interpretation is that the residual arises due to aggregation within a product-importer category of many different goods that may be horizontally as well as vertically differentiated.

The most common definition of units in the US export data unit is a simple count, with the second most common being weight in kilograms (some records report two unit definitions, in which case the second unit is almost always kilos). Other units include bushels, barrels, square meters, grams, and the like. While the code fixed effects sweep out differences in units across products, there may still be a difficulty in comparing the effects of distance, market size, and remoteness on unit values not in common units. To address this concern, we run the regression above on the subset of data for which kilograms are the unit, so that unit value is simply the value/weight ratio. We also report results where the sample is restricted to manufactured goods, as in the previous section.

Product level unit values are notoriously noisy indicators of prices, particularly for very small trade flows. To make sure that our results are not overly influenced by noisy and economically unimportant observations we trim the estimation sample by discarding all export flows of less than \$10,000, which eliminates the smallest 7 percent of observations.

Table 4 and Appendix Table 4 report the results of our export unit value regressions. A striking message is that *distance has a very large positive effect on unit values*. Looking at the first column in Table 4, which is our baseline specification, the distance elasticity of unit values is 0.27. This elasticity is a bit smaller when non-manufactured goods are included in the sample and for goods measured in kilos, though the differences across samples are not statistically significant (see Appendix Table 4).

The market size elasticity is small and negative at -0.05 , an effect which disappears when the sample is restricted to goods measured in kilos. Turning to the specification with a distance-market size interaction, reported in Appendix Table 4, the implied distance elasticity in the full sample is smaller for richer countries, but remains positive for 97 out of 100 of the countries in the sample. The implied market size elasticity

TABLE 4—STATISTICAL DETERMINANTS OF US EXPORT UNIT VALUES, 2005

Linear distance	Distance step function	
Log distance	0.272*** (0.0667)	
1 < km ≤ 4,000		0.0436 (0.103)
4,000 < km ≤ 7,800		0.427*** (0.0894)
7,800 < km ≤ 14,000		0.516*** (0.0905)
14,000 < km		0.644*** (0.0920)
Log real GDP	-0.0523*** (0.0177)	-0.0649*** (0.0146)
Log real GDP per worker	-0.0560** (0.0274)	-0.0766*** (0.0268)
Log remoteness	-2.706*** (0.328)	-2.413*** (0.396)
NAFTA	-0.0141 (0.231)	
Landlocked	0.0709 (0.0646)	0.0365 (0.0666)
Island	0.0319 (0.0582)	0.0700 (0.0515)
English language	-0.0380 (0.0625)	-0.0732* (0.0390)
R ² within Sample	0.030 123,547 observations on 4,886 HS10 codes	0.035

Notes: Dependent variable is log unit value of exports by HS10 product and export destination, relative to HS10 code mean (that is, HS10 code fixed effects are included in all regressions). Manufacturing observations only, 123,547 observations on 4,886 HS10 codes. Includes only trade flows over \$10,000. Independent variables are characteristics of export destinations. Estimator is OLS with HS10 product fixed effects, and errors are clustered by HS10 and country. Robust standard errors clustered by HS10 code and importing country in parentheses, asterisks denote statistical significance.

***Significant at the 1 percent level.

**Significant at the 5 percent level.

*Significant at the 10 percent level.

becomes more negative with distance, and is negative for 85 of the 100 countries in the sample (similar results apply to the manufacturing and kilos samples).

When distance is measured using a step function rather than as log kilometers, the same pattern that was observed in the zeros results reappears: for importers within 4,000 kilometers of Chicago, distance does not have a statistically significant effect (except when the sample is restricted to goods measured in kilos, see Appendix Table 4). When reaching an export market requires crossing an ocean, the effect of distance is very large and mildly increasing in distance, with an effect of about 40 to 60 log points. These are extremely large effects, since a 50 log point effect corresponds to distance raising export unit values by 65 percent.

The effect of remoteness is large and consistently negative across specifications and samples, with the range of elasticities between -1.5 and -2.7 (though it should

be noted that the sample standard deviation of log remoteness is only 0.05, so not much of the variation in the data is explained by remoteness). The effect of aggregate productivity is fragile: there is a small positive effect when the sample is restricted to products measured in kilos, and a small negative effect for manufactured goods, with the effect for the sample as a whole being indistinguishable from zero. Finally, the other indicator variables (border, landlocked, island, and English language) have inconsistent and generally statistically insignificant effects.

The Table 4 results are consistent with the findings of Hummels and Alexandre Skiba (2004). In particular, equation (12) in Hummels and Skiba (2004) is quite similar to the specification estimated in Table 4. The biggest difference is that we include market size and remoteness as explanatory variables for log unit values while Hummels and Skiba (2004) do not. Hummels and Skiba (2004) interpret the price-distance correlation as evidence of a composition effect in demand due to lower proportionate transport costs on higher quality goods, an effect whose strength increases in distance. The Hummels-Skiba explanation is based on their evidence that transport costs are per unit rather than *ad valorem*, which is to say that the usual iceberg transport cost assumption, which is embodied in all of the theory models that we consider, is wrong. In the following section we retain the iceberg transport cost assumption for comparability with the literature and to allow analytical solution of the general equilibrium model. The quality-sorting mechanism that we propose to explain the price-distance relationship operates through a different mechanism than the relative demand effect proposed by Hummels and Skiba (2004), but the two explanations work in the same direction and we regard them as complementary.

The strong positive relationship between export unit values and distance seen in Table 4 is inconsistent with *all* of the models presented in Section II. The baseline monopolistic competition model predict a zero relationship, while the other models predict a negative relationship between export unit values and distance, the exact opposite of what shows up so strongly in the data. Only the Melitz-Ottaviano model is consistent with the negative market size effect on prices.

We emphasize that the effect of distance on unit values found here cannot be given a clear structural interpretation. Most obviously, since we know from the results of Table 3 that distance has a strong influence on the probability of zeros, the effect estimated here conflates a selection effect (which markets are exported to) and a direct effect (conditional on positive exports, what is the effect of distance on unit value). However, recent work by Harrigan and Victor Shlychkov (2010) strongly suggest that selection is what explains the price-distance effect in our data. Using confidential firm-level data that underlies the product level data used here, Harrigan and Shlychkov (2010) show that firms do not systematically charge different export prices as a function of distance. Instead, they find that high-price firms are more likely to sell in more distant markets.

III. Trade with Heterogeneous Quality

The empirical evidence presented above has a clear message: the Melitz model does a good job of explaining export zeros, but cannot explain spatial variation in prices. In this section we introduce a slight variant of the HFT model that can

account for the spatial facts on prices. The basic change is that firms will compete on the basis of quality as well as price, so we refer to it as the quality heterogeneous-firms model, or QHFT for short.

A. Quality and Trade: Antecedents and Recent Work

The idea that exporting firms compete on quality as well as price has a long history in international trade economics. In examining the linkages between quality and national trade patterns, we follow a number of important recent papers, including primarily empirical papers by Schott (2004), Hummels and Peter J. Klenow (2005), and Juan Carlos Hallak (2006). Schott's (2004) importance lies in the finding of a strong positive relationship between exporter GDP per capita and the average unit value of HS10 codes sold to the United States (in contrast, our results reported in Table 4 look at the opposite flow and find no robust relationship between US export unit values and the GDP per capita of the importer). Hummels and Klenow (2005) sketch the empirical implications of a number of homogenous-firms trade theories that take quality as national attributes. As such, the Melitz-like selection effect of distance-related trade costs is not in operation; bilateral export patterns are marked by an all-varieties-or-none pattern. A key empirical finding of Hummels and Klenow (2005) is an empirical association between exporter GDP per capita and export unit value. While Schott (2004) and Hummels and Klenow (2005) find a relationship between exporter GDP per capita and quality, Hallak (2006) looks at the demand side, and finds that the demand for quality is related to importer GDP per capita.

In an important recent paper, Hallak and Schott (2008) take a more nuanced approach to estimating exporter product quality, both across countries and over time, by developing a method to disentangle price and quality variation in unit value data. Among their findings is a confirmation that the level of quality is correlated with the level of development, but the relationship is somewhat weaker in growth rates. Amit Khandelwal (2010) also separates price and quality, and finds that this decomposition enables provocative insights into the nature of US manufacturer's exposure to competition from imports. Khandelwal (2010) also confirms the association between exporter level of development and export quality.

There are a number of different approaches to modeling the demand for quality. The most common is to model preferences for what might be called box-size-quality: the utility of consuming two boxes of variety- j with quality 1 is identical to consuming one box of variety- j with quality 2, just as if it were a bigger box of the same good.¹⁷ In the standard CES monopolistic competition setting this means consumers make their decisions on quality-adjusted price rather than the observed price (i.e., they care about the observed price adjusted for box-size). An older theoretical literature, exemplified by Kevin M. Murphy and Andrei Shleifer (1997) and Gene M. Grossman and Helpman (1991) worked with vertical quality models where several firms compete on price and quality to win a single market. A separate

¹⁷Hummels and Klenow (2005), Hallak (2006), Hallak and Schott (2008) and Maurice Kugler and Eric A. Verhoogen (2008), among many others, use variations on this approach.

approach is found in the Industrial Organization “buy only one” models of quality demand used in the trade literature by Verhoogen (2008), Khandelwal (2010), and John Sutton (2007).

Models of the supply of quality by firms differ in their details, but invariably deliver a mapping between an exogenous parameter (Sutton (2008) calls it “capability,” Khandelwal (2010) calls it “ability,” Verhoogen (2008) calls it “productivity,” etc) and the possibly endogenous supply of quality, a characteristic valued by consumers.¹⁸ Market equilibrium in these models is usually some variant of Nash equilibrium (including monopolistic competition).

The value added of our QHFT model developed below lies in three main dimensions. First, we work out a general equilibrium trade model with trade costs and many nations where quality is a basis of comparative advantage and quality is linked in a very simple manner to fundamentals in each nation. As part of this, we show that adding in quality as part of a firm’s competitiveness is best thought of as a conceptual amendment to the standard HFT models rather than a separate model. Second, the model allows for firm heterogeneity in productivity and quality in line with empirical evidence. Third, we develop specific, testable hypotheses and, referring back to Section I above, show that the QHFT model is more consistent with the data than the models discussed in Section I above.

Since the first draft of our paper was widely disseminated in March 2007, a number of “quality heterogeneous-firms trade” models have been developed or are in development. Several of these use the box-size-quality preferences over quality and monopolistic competition (Robert C. Johnson 2008, Kugler and Verhoogen 2008), while others work in the buy-only-one-unit preferences or the linear preferences taken from the tradition of partial equilibrium models of industrial organization theory (Sutton 2007). The only paper that works out the full general equilibrium with trade costs is Alexis Antoniadis (2008, unpublished), who incorporates quality competition into the Melitz-Ottaviano model. Others assume away trade costs, which are essential to our empirically testable predictions (Kugler and Verhoogen 2008, Sutton 2007).

B. *Quality, Heterogeneous Firms, and Trade: The QHFT Model*

Most of the assumptions and notation of our QHFT model are in keeping with the HFT model introduced above. There are two main changes. On the demand side consumers now care about quality, and on the supply side firms produce varieties of different quality. More precisely, consumers regard some varieties as superior to others. This superiority could be regarded purely as a matter of taste, but we will interpret superiority as a matter of “quality.” The utility function is

$$(16) \quad U = \left(\int_{\epsilon \in \Theta} (c_i q_i)^{1-1/\sigma} di \right)^{1/(1-1/\sigma)} ; \quad \sigma > 1,$$

¹⁸An exception is Mathias Helble and Toshihiro Okubo (2006), whose model assumes that quality has no effect on demand, but that high-quality products face lower beachhead costs.

where c and q are the units consumed and quality of a typical variety and Θ is the set of consumed varieties. The optimal expenditure on a typical variety is

$$(17) \quad p_j c_j = \left(\frac{p_j}{q_j}\right)^{1-\sigma} B; \quad B \equiv \frac{E}{P^{1-\sigma}}, \quad P^{1-\sigma} \equiv \int_{\epsilon \in \Theta} \left(\frac{p_i}{q_i}\right)^{1-\sigma} di,$$

where p_j/q_j has the interpretation of a quality-adjusted price of good- j , P is the CES index of quality-adjusted prices, and Θ the set of consumed varieties. The standard CES preferences are a special case of (16) with $q_j = 1$, for all j .

As in the standard HFT model, manufacturing firms draw their a from a random distribution after paying a fixed innovation cost of F_I units of labor (see Appendix for details). In the QHFT model, however, high costs are not all bad news, for higher quality is assumed to come with higher marginal cost. In particular

$$(18) \quad q_j = a_j^{1+\theta}, \quad \theta > -1,$$

where $1 + \theta$ is the “quality elasticity,” namely the extent to which higher marginal costs are related to higher quality (setting $\theta = -1$ reduces this to the standard HFT model). Observe that, although this link is assumed for simplicity’s sake, allowing firms to optimally choose quality would not enrich the model. In any model where heterogeneity only stemmed from firms’ productivity, optimization would yield a link between a firm’s optimized q and its a that would resemble (18). Indeed, a similar positive (and often log-linear) relationship between quality and marginal cost is common to many of the papers discussed above, including Khandelwal (2010), Verhoogen (2008), Johnson (2008), and Antoniadis (2008). Some of these papers derive this relationship as the solution of the firm’s optimal quality choice problem, but the reduced forms plainly depend only upon a firm’s randomly assigned productivity. The point is that all firms face identical parameters (factor costs, etc.) except for their firm-specific productivity draw. Obviously then the optimal, firm-specific quality choice will vary only with firm-specific productivity.

The important contribution of Kugler and Verhoogen (2008) provides direct empirical support for the relationship between quality and cost. Using a remarkable dataset from Colombia which records information on the cost of firms’ inputs, Kugler and Verhoogen (2008) show that higher cost inputs are systematically associated with higher quality outputs, as we assume in (18). Also see the plant-level evidence for this mechanism in Eileen L. Brooks (2006), and Verhoogen (2008).

At the time it chooses prices, the typical firm takes its quality and marginal cost as given, so the standard Dixit-Stiglitz results apply. Mill-pricing is optimal for all firms in all markets, and operating profit is a constant fraction, $1/\sigma$, of firm revenue. Using mill-pricing and (18) in (17), the optimized operating profit for a typical nation- o firm selling in nation- d is

$$(19) \quad \frac{v_{od}[a]}{\sigma} = \left(\frac{\tau_{od} a^{-\theta} w_o}{1 - 1/\sigma}\right)^{1-\sigma} \frac{B_d}{\sigma}.$$

The only difference between this and the corresponding expression for profits without quality differences is the θ in the exponent on a . Plainly, the properties of this model crucially depend on how elastic quality is with respect to marginal cost. For $\theta \in [-1, 0)$, quality increases slowly with cost and the optimal quality-adjusted consumer price increases with a firm's a . In this case, a firm's revenue and operating profit fall with its marginal cost. For $\theta > 0$, by contrast, quality increases quickly enough so that the quality-adjusted price falls as a rises. This means that higher a 's are associated with higher operating profit. Henceforth we focus on the $\theta > 0$ case because, as the empirics above suggested, it is the case that is most consistent with the data.

Given (19), $v_{od}[a_{od}]/\sigma = F$ defines the cut-off for nation- o sales to market- d ; this can be rewritten as

$$(20) \quad w_o a_{od}^\theta = \left(\frac{\phi_{od} B_d}{f} \right)^{\frac{1}{1-\sigma}}.$$

Assuming $\theta > 0$, this tells us that only firms with sufficiently high-price/high-quality goods find it worthwhile to sell to distant markets. This is the *opposite* of Melitz (2003) and all other HFT models. In standard HFT models, competition depends only on price, so it is the lowest priced goods that make it to the most distant markets. In the QHFT model, competition depends on quality-adjusted prices and with $\theta > 0$, the most competitive varieties are high-price/high-quality. This means that distance selects for high-priced varieties rather than low-priced varieties as in the HFT model.

Since the QHFT model can be reduced to the HFT model with quality adjusted prices, the general equilibrium solution for the QHFT model is isomorphic to that of the HFT model (see Appendix for details). More elaborate and nuanced models of quality choice (e.g., Johnson 2008 and Sutton 2007) are well suited to addressing partial equilibrium industry-competition issues, but they are not amenable to the general equilibrium analysis necessary to yield testable predictions for the spatial pattern of zeros and prices. For this reason, we do not pursue such extensions here.

C. Quality HFT's Spatial Pattern of Zeros and Prices

The spatial pattern of zeros in the QHFT model conforms to those of the HFT model, as a comparison of the cut-off conditions of the two models, (8) and (20), reveals. The key, new implication has to do with the relationship between prices and distance. Since a high observed price indicates high competitiveness (i.e., a low quality-adjusted price), the marginal cost thresholds are increasing in distance, rather than decreasing as in the HFT model. Given that mill pricing is optimal, this means that the selection produces average c.i.f. and f.o.b. prices that are increasing with bilateral distance. Of course, the logic of the model turns on the fact that the average f.o.b. *quality-adjusted* prices are decreasing in distance, but since the data report only average unit values this is not a testable implication.

Like all of the other models reviewed in this paper, the QHFT model assumes *ad valorem* trade costs. As discussed in the previous section, the partial equilibrium quality-sorting explanation of the price-distance relationship proposed by Hummels

and Skiba (2004) is complementary to the general equilibrium quality-selection mechanism in the QHFT model. Incorporating per-unit trade costs into our CES-based general equilibrium model of quality competition would make an analytical solution impossible and would also break the close connection to Melitz (2003), so we do not pursue this extension here.

Equation (20) also implies that the relationship between average f.o.b. prices and market size and remoteness is decreasing, the opposite of the relationship given by the baseline HFT model. The reason for the different prediction is that as export market becomes more attractive—either due to size or remoteness—selection leads to lower quality firms to enter; their entry lowers the average f.o.b. price. Harrigan and Shlychkov (2010) provide direct evidence for this channel using confidential US firm-level data.

These predictions are noted in the last line of Table 2. Once again, the QHFT model is the only one that we considered that matches the findings of the data analysis in Section III.

IV. Concluding Remarks

This paper has shown that existing models of bilateral trade all fail to explain key features of the product-level data that we document using highly disaggregated US trade data:

- Most products are exported to only a few destinations.
- The incidence of these “export zeros” is positively related to distance and negatively related to market size and remoteness.
- The average unit value of exports is positively related to distance.

We finished the paper by proposing a slight modification of the Melitz model which fits all of the facts just summarized. The modification assumes that consumers care about quality and firms’ heterogeneity in productivity shows up in heterogeneous quality in a specific way. The most competitive firms—i.e., those with the lowest quality-adjusted price—charge the highest market prices. This accounts for the fact that the most expensive products seem to be sold to the most distant markets.

REFERENCES

- Anderson, James E., and Eric van Wincoop.** 2003. “Gravity with Gravitas: A Solution to the Border Puzzle.” *American Economic Review*, 93(1): 170–92.
- Antoniades, Alexis.** 2008. “Heterogeneous Firms, Quality and Trade.” <http://www.eiit.org/WorkingPapers/Papers/FirmLevelTrade/FREIT004.pdf>.
- Armenter, Roc, and Miklós Koren.** 2008. “A Balls-and-Bins Model of Trade.” http://resources.cefig.eu/archive/2008/06/papers/balls_and_bins.pdf.
- Axtell, Robert L.** 2001. “Zipf Distribution of U.S. Firm Sizes.” *Science*, 293(5536): 1818–20.
- Baldwin, Richard.** 2005. “Heterogeneous Firms and Trade: Testable and Untestable Properties of the Melitz Model.” National Bureau of Economic Research Working Paper 11471.
- Bernard, Andrew B., Jonathan Eaton, J. Bradford Jensen, and Samuel Kortum.** 2003. “Plants and Productivity in International Trade.” *American Economic Review*, 93(4): 1268–90.
- Bernard, Andrew B., J. Bradford Jensen, and Peter K. Schott.** 2009. “Importers, Exporters, and Multinationals: A Portrait of Firms in the U.S. that Trade Goods.” In *Producer Dynamics: New Evidence*

- from *Micro Data*, ed. Timothy Dunne, J. Bradford Jensen, and Mark J. Roberts, 513–56. Chicago: University of Chicago Press.
- Besedeš, Tibor, and Thomas J. Prusa.** 2006a. “Ins, Outs, and the Duration of Trade.” *Canadian Journal of Economics*, 39(1): 266–95.
- Besedeš, Tibor, and Thomas J. Prusa.** 2006b. “Product Differentiation and Duration of US Import Trade.” *Journal of International Economics*, 70(2): 339–58.
- Brooks, Eileen L.** 2006. “Why Don’t Firms Export More? Product Quality and Colombian Plants.” *Journal of Development Economics*, 80(1): 160–78.
- Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller.** Forthcoming. “Robust Inference with Multi-way Clustering.” *Journal of Business and Economic Statistics*.
- Choi, Yo Chul, David Hummels, and Chong Xiang.** 2006. “Explaining Import Variety and Quality: The Role of the Income Distribution.” National Bureau of Economic Research Working Paper 12531.
- Demidova, Svetlana.** 2008. “Productivity Improvements and Falling Trade Costs: Boon or Bane?” *International Economic Review*, 49(4): 1437–62.
- Eaton, Jonathan, and Samuel Kortum.** 2002. “Technology, Geography, and Trade.” *Econometrica*, 70(5): 1741–79.
- Eaton, Jonathan, Samuel Kortum, and Francis Kramarz.** 2008. “An Anatomy of International Trade: Evidence from French Firms.” National Bureau of Economic Research Working Paper 14610.
- Falvey, Rod, David Greenaway, and Zhihong Yu.** 2006. “Extending the Melitz Model to Asymmetric Countries.” University of Nottingham Globalisation and Economic Policy Research Paper 06/07.
- Feenstra, Robert C., James R. Markusen, and Andrew K. Rose.** 2001. “Using the Gravity Equation to Differentiate among Alternative Theories of Trade.” *Canadian Journal of Economics*, 34(2): 430–47.
- Feenstra, Robert C., and Andrew K. Rose.** 2000. “Putting Things in Order: Trade Dynamics and Product Cycles.” *Review of Economics and Statistics*, 82(3): 369–82.
- Grossman, Gene M., and Elhanan Helpman.** 1991. *Innovation and Growth in the Global Economy*. Cambridge, MA: MIT Press.
- Hallak, Juan Carlos.** 2006. “Product Quality and the Direction of Trade.” *Journal of International Economics*, 68(1): 238–65.
- Hallak, Juan Carlos, and Peter K. Schott.** 2008. “Estimating Cross-Country Differences in Product Quality.” National Bureau of Economic Research Working Paper 13807.
- Hansen, Christian B.** 2007. “Asymptotic Properties of a Robust Variance Matrix Estimator for Panel Data When T Is Large.” *Journal of Econometrics*, 141(2): 597–620.
- Harrigan, James.** 2003. “Specialization and the Volume of Trade: Do the Data Obey the Laws?” In *Handbook of International Trade*, ed. E. Kwan Choi and James Harrigan, 85–118. Malden, MA: Blackwell Publishing.
- Harrigan, James, and Victor Shlychkov.** 2010. “Export Pricing by U.S. Manufacturing Firms.” Unpublished.
- Haveman, Jon, and David Hummels.** 2004. “Alternative Hypotheses and the Volume of Trade: The Gravity Equation and the Extent of Specialization.” *Canadian Journal of Economics*, 37(1): 199–218.
- Helble, Mathias, and Toshihiro Okubo.** 2006. “Heterogeneous Quality and Trade Costs.” <http://www.etsg.org/ETSG2006/papers/helble.pdf>.
- Helpman, Elhanan, and Paul R. Krugman.** 1985. *Market Structure and Foreign Trade: Increasing Returns, Imperfect Competition, and the International Economy*. Cambridge, MA: MIT Press.
- Helpman, Elhanan, Marc Melitz, and Yona Rubinstein.** 2008. “Estimating Trade Flows: Trading Partners and Trading Volumes.” *Quarterly Journal of Economics*, 123(2): 441–87.
- Helpman, Elhanan, Marc J. Melitz, and Stephen R. Yeaple.** 2004. “Export versus FDI with Heterogeneous Firms.” *American Economic Review*, 94(1): 300–316.
- Hummels, David.** 1999. “Have International Transportation Costs Declined?” http://www.aerohabitat.eu/uploads/media/11-01-2006_-_D_Hummels_Transportation_cost_declines.pdf.
- Hummels, David, and Peter J. Klenow.** 2005. “The Variety and Quality of a Nation’s Exports.” *American Economic Review*, 95(3): 704–23.
- 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.** 2008. “Trade and Prices with Heterogeneous Firms.” http://www4.gsb.columbia.edu/null/download?&exclusive=filemgr.download&file_id=34166.
- Khandelwal, Amit.** 2010. “The Long and Short (of) Quality Ladders.” Dartmouth College Working Paper.
- Kugler, Maurice, and Eric A. Verhoogen.** Forthcoming. “Prices, Plant Size, and Product Quality.” *Review of Economic Studies*.
- Melitz, Marc J.** 2003. “The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity.” *Econometrica*, 71(6): 1695–1725.

- Melitz, Marc J., and Giancarlo I. P. Ottaviano.** 2008. "Market Size, Trade, and Productivity." *Review of Economic Studies*, 75(1): 295–316.
- Murphy, Kevin M., and Andrei Shleifer.** 1997. "Quality and Trade." *Journal of Development Economics*, 53(1): 1–15.
- Ottaviano, Gianmarco, Takatoshi Tabuchi, and Jacques-Francois Thisse.** 2002. "Agglomeration and Trade Revisited." *International Economic Review*, 43(2): 409–35.
- Schott, Peter K.** 2004. "Across-Product versus Within-Product Specialization in International Trade." *Quarterly Journal of Economics*, 119(2): 647–78.
- Sutton, John.** 2007. "Quality, Trade and the Moving Window: The Globalisation Process." *Economic Journal*, 117(524): F469–98.
- Verhoogen, Eric A.** 2008. "Trade, Quality Upgrading, and Wage Inequality in the Mexican Manufacturing Sector." *Quarterly Journal of Economics*, 123(2): 489–530.
- Wooldridge, Jeffrey M.** 2002. *Econometric Analysis of Cross Section and Panel Data*. Cambridge, MA: MIT Press.
- Yeaple, Stephen Ross.** 2005. "A Simple Model of Firm Heterogeneity, International Trade, and Wages." *Journal of International Economics*, 65(1): 1–20.