Trade in Ideal Varieties:

Theory and Evidence

David Hummels

Volodymyr Lugovskyy

March 2005

Abstract

Models with constant-elasticity of substitution (CES) preferences are commonly employed in the international trade literature because they provide a tractable way to handle product differentiation in general equilibrium. However this tractability comes at the cost of generating a set of counter-factual predictions regarding cross-country variation in export and import variety, output per variety, and prices. We examine whether a generalized version of Lancaster's 'ideal variety' model can better match facts. In this model, entry causes crowding in variety space, so that the marginal utility of new varieties falls as market size grows. Crowding is partially offset by income effects, as richer consumers will pay more for varieties closer matched to their ideal types. We show theoretically and confirm empirically that declining marginal utility of new varieties results in: a higher own-price elasticity of demand (and lower prices) in large countries and a lower own-price elasticity of demand (and higher prices) in rich countries. Model predictions about cross-country differences in the number and size of establishments are also empirically confirmed.

JEL Classification: F12, L11

Acknowledgments: We thank participants at the EIIT 2004 conference, and seminar participants at the Universities of Purdue, Vanderbilt, Texas, Indiana and Illinois for helpful comments. Hummels thanks Purdue CIBER and NSF and Lugovskyy thanks Wang CIBER at the University of Memphis for financial support.

Contact info: Hummels, Dept Economics, Purdue University, 403 W. State St, W. Lafayette IN 47907-2056. ph 765 494 4495. em: hummelsd@purdue.edu; Lugovskyy. Department of Economics, Fogelman School of Business and Economics, University of Memphis, 400 Fogelman Administrative Bldg, Memphis, 38104, Phone: (901) 678-2473. E-mail: vlgvskyy@memphis.edu.

I. Introduction

Beginning with Krugman (1979, 1980), the Dixit-Stiglitz (1977) framework of product differentiation has become a workhorse of the international trade literature. A short list of applications includes the literatures on intra-industry and north-north trade, economic geography, regional integration, gravity modeling of trade flows, and multinational firms.

The model is widely used because it is highly tractable. In its most commonly used form the model assumes constant elasticity of substitution (CES) demand, which has several important features. Varieties are not assigned to any particular "address" and product space is effectively infinite. As a consequence, differentiated varieties may exhibit a high or low degree of substitutability, but this is invariant to the number of products in the market. Further, consumers "love" variety, in the sense that increased variety improves welfare. Curiously, the marginal utility consumers derive from new variety is not declining in entry, whether new entrants represent the 20th or 200th variety on the market.

These assumptions carry pronounced normative implications. Trade economists are often bemused that the gains from trade in neoclassical models are small. However, as Romer (1994) demonstrates in a simple calibration, trade liberalization that increases the number of traded varieties can be a source of much larger welfare gains.[1] In a related empirical exercise, Feenstra (1994) shows that under a maintained hypothesis of CES demand, variety-corrected price indices for a set of US imports have fallen much lower

[1] Klenow and Rodriguez (1997) show that the welfare gains from new variety in Costa Rica are smaller than Romer suggests, but still larger than neoclassical gains.

than the uncorrected indices would seem to indicate.[2]

In sum, the CES model's core assumptions yield useful tractability, and the model's predictions about variety expansion imply tremendous welfare gains from trade in markets large and small. But are the central empirical predictions of this structure correct? Consider a few stark implications of the model that can be shown by comparing, ceteris paribus, a large and a small closed economy. The larger country will enjoy entry of new varieties at a rate proportional to its country size. Because this entry does not "crowd" variety space, the own-price (and cross-price) elasticity of demand is the same as in the small country. This implies that prices, which depend only on marginal cost and the demand elasticity, are the same in the two markets. It also implies that the quantity per variety is the same in the two markets: fixed costs of entry are the same in the two markets, and markups are the same, so the same quantity per variety clears zero-profit conditions in both places. These predictions are not incidental. It is precisely this strong symmetry in prices and quantities, and the strict proportionality between number of varieties and market size which makes these models so tractable and widely-employed.

Hummels and Klenow (2002) use cross-country data to examine how the variety and quantity per variety of imports co-vary with market size. They show that, while the number of imported varieties is greater in larger markets, variety differences are less than proportional to market size. That is, larger countries import more varieties, but also import higher quantity per variety.[3]

What could cause a less than proportional expansion of import varieties with respect to market size? Two candidates come to mind. Perhaps there are fixed costs of

[2] Broda and Weinstein (2004) extend the Feenstra calculations to a broader set of US import categories.
[3] Hummels and Klenow (2002,2005) found a similar pattern for exports: variety and quantity per variety expands with exporter size, but less than proportionally. We focus on cross-importer facts in this paper.

importing as modeled in Romer (1994), but these are rising in market size so that the cost of new varieties at the margin is higher for larger markets as in Klenow-Rodriguez Clare (1997). Alternatively, it may be that goods become more substitutable as more varieties enter the market, so that the marginal benefit of new varieties falls with market size. We emphasize the latter channel, examining the empirical implications of a generalized 'ideal variety' model.

Lancaster (1979) originally developed a model of trade in ideal varieties in which variety space is finite, and varieties have unique addresses in product space. This means that entry causes "crowding" – goods become more substitutable as more enter the market so that the own price elasticity of demand increases with market size. This has pronounced implications for the average size of firms in a closed economy, and for patterns of import variety and quantity per variety in an open economy.

Again consider a ceteris paribus comparison of a large and small country. In the large country, entry drives up the price elasticity of demand, leading to lower markups. Since firms in the large market must be able to recoup their fixed costs of entry despite lower markups, they must sell a larger quantity. This, in turn, implies that new variety expansion will be increasing in market size, but less than proportionally.

We generalize the preferences in the ideal variety framework. Lancaster assumes that the equilibrium choice of variety is independent of consumption quantities, so that consumers get no closer to their ideal regardless of expenditures. We allow the opportunity cost of the ideal variety to depend on consumers' individual consumption levels. When incomes rise, consumers increase the quantity consumed, but also place greater value on proximity to the ideal variety. The price elasticity of demand drops and prices rise. In equilibrium, the market responds by supplying more varieties, with lower output per variety. Essentially, economies of scale forsaken are compensated for by the higher markups that consumers are now willing to pay.

We examine these implications in three exercises focusing on cross-country variation in average firm size, the own-price elasticity of demand, and prices. First, we examine the model's closed economy predictions for the average size of firms. We use the UNIDO Industrial Statistics Database to measure the average value added per firm for 152 3-digit ISIC sectors in 54 countries from 1990-2000. Controlling for country and industry effects so that we exploit purely time series variation, we show that average firm size positively covaries with GDP, and negatively with GDP per worker (conditioning on GDP). This result is robust to alternative measures of firm and market size.

Next, we examine the model's open economy predictions. We already know from Hummels-Klenow (2002) that model predictions on import variety and quantity per variety match empirical facts. We show that the predictions on prices and price elasticity of demand also match trade data. Using bilateral trade data for 60 importers, 120 exporters, and 5000 products, we show that the own-price elasticity of demand is increasing in importer GDP and decreasing in importer GDP per capita. The data reveal substantial variation in these elasticities across importers. Data on prices for a smaller subset of importers (chosen because quantity units are consistently measured) match the price elasticities: prices are decreasing in importer GDP and increasing in importer GDP per capita.

This paper relates, and adds to, several literatures. First, we contribute to a relatively new but growing literature providing empirical evidence on models of product differentiation in trade. Most of these papers employ cross-exporter facts to understand Armington v. Krugman style horizontal differentiation as in Head and Ries (2001) and Acemoglu and Ventura (2002), or the importance of quality differentiation as in Schott (2003), Hallak (2004), and Hummels and Skiba (2004), or some combination of the two, as in Hummels and Klenow (2005). We emphasize cross-importer facts, and depart from the CES utility framework that dominates this literature.

Second, we contribute to a literature in which market entry affects the elasticity of demand facing a firm. Most of the theory literature has emphasized oligopoly and homogeneous goods as in Brander and Krugman (1982). The more sparse empirical literature has focused on plausibly homogeneous goods within a single country, such as the markets for gasoline, Barron, Taylor, Umbeck (2005) and concrete, Syverson (2004). In contrast, our model emphasizes free-entry monopolistic competition in a general equilibrium with many countries and differentiated goods. The model's predictions for market size and the elasticity of demand are similar to quadratic utility models as in Ottaviano and Thisse (1999). However, we allow for income effects operating through an intensity of preference for the ideal variety that can potentially counteract pure market size effects.[4] These income effects significantly improve our ability to fit the model to the data.

Finally, this paper adds to the literature on price variation across markets. The literature on pricing-to-market (see Goldberg and Knetter 1997 for an extensive review) has shown that the same goods are priced with different markups and thus have different price elasticities of demand across importing markets. We differ from, and add to, this

[4] Perloff and Salop (1985) also include preference intensity but do not link it explicitly to observable characteristics of consumers, or consider a trading equilibrium.

literature in two ways. First, we show how markups systematically vary across importers depending on market characteristics. Second, we provide a complementary explanation for the variation in markups. The pricing-to-market literature focuses on movements along the *same*, non-CES, demand curves (e.g., Feenstra 1989, Knetter 1993) so that variation in quantities caused by tariff or exchange rate shocks yields variation in the elasticity of demand. We show that variation in market characteristics (size, income per capita), yields *different* demand curves and thus different price elasticities of demand across countries.

The rest of the paper is organized as follows. Section II uses a simplified closed economy setting to motivate the generalization of Lancaster compensation function and to concentrate on the comparative statics in the model with a single differentiated product. Appendix 2 demonstrates that the key empirical predictions can also be derived in an open economy model that nests the generalized ideal variety framework into a Ricardian continuum model in the manner of Dornbusch-Fischer-Samuelson (1977). Sections III-V provide empirical examinations of model implications for average firm sizes, the own-price elasticity of demand, and prices. Section VI concludes.

II. Model

A. Demand Functions

Preferences of a consumer are defined over a homogeneous numeraire product q_0, and a differentiated product q, which is defined by a continuum of varieties indexed by $\omega \in \Omega$:

(1) $U = q_0^{1-\mu}\left[u\left(q_\omega \mid \omega \in \Omega\right)\right]^\mu \qquad 0 < \mu < 1,$

where subutility $u\left(q_\omega \mid \omega \in \Omega\right)$ is defined later in this section. The budget constraint is:

(2) $q_0 + \int_{\omega\in\Omega} q_\omega p_\omega = I,$

where p_ω are the prices of the varieties being produced and I is income in terms of the numeraire.

Varieties can be distinguished by a single attribute. We assume that all varieties can be represented by points on the circumference of a circle, with the circumference being of unit length.

Each point of the circumference represents a different variety. Each consumer has his most preferred type, which we call his 'ideal' variety, and which we denote as $\tilde{\omega}$. It is ideal in the sense that given a choice between equal amounts of his ideal variety $\tilde{\omega}$ and any other variety ω consumer will always choose $\tilde{\omega}$. Moreover, utility is decreasing in distance from $\tilde{\omega}$: the further is the product from the ideal variety the less preferable it is for the consumer. These assumptions are usually incorporated in the formal model with a help of Lancaster's compensation function $h\left(v_{\omega,\tilde{\omega}}\right)$, defined for $0 \le v_{\omega,\tilde{\omega}} \le 1$. Lancaster's compensation function is defined such that the consumer is indifferent between q units of his ideal variety $\tilde{\omega}$ and $h\left(v_{\omega,\tilde{\omega}}\right)q$ units of some other variety ω, where $v_{\omega,\tilde{\omega}}$ is the shortest arc distance between $\tilde{\omega}$ and ω. It is assumed that:

(3) $h(0)=1$, $h'(0)=0$, and $h'\left(v_{\omega,\tilde{\omega}}\right) > 0$, $h''\left(v_{\omega,\tilde{\omega}}\right) > 0$ for $v_{\omega,\tilde{\omega}} > 0$.

The subutility of variety ω for consumer whose ideal variety is $\tilde{\omega}$ is usually assumed to have the following separable form (e.g., Lancaster 1979, 1984, Helpman and

Krugman 1985):

$$u\left(q_\omega, \omega, \tilde{\omega}\right) = \frac{q_\omega}{h\left(v_{\omega,\tilde{\omega}}\right)}$$

The generalized subutility function, which includes all varieties $\omega \in \Omega$, can then be formulated as

$$(4) \qquad u\left(q_\omega \mid \omega \in \Omega\right) = \max_{\omega \in \Omega}\left[\frac{q_\omega}{h\left(v_{\omega,\tilde{\omega}}\right)}\right]$$

Given the weak separability of the utility function (1), we can use a two-stage budgeting procedure. From the second stage we find that the consumer spends μI on the differentiated product. In the first stage, we can maximize the subutility subject to the budget constraint, μI, and given the prices of differentiated varieties, p_ω. The solution to this problem is:

$$(5) \qquad \begin{aligned} &q_{\omega'} = \frac{\mu I}{p_{\omega'}}, \qquad q_\omega = 0 \quad \text{for } \omega \neq \omega', \\ &\text{where} \quad \omega' = \arg\min\left[p_\omega h\left(v_{\omega,\tilde{\omega}}\right) \mid \omega \in \Omega\right]. \end{aligned}$$

In (5), the utility maximizing variety is independent of expenditures: given prices, a consumer will move no closer to his ideal variety if he buys one unit or a thousand. For example, imagine that the consumer's ideal variety is apple juice, the price of which is five times higher than the price of water: $p_{AJ} = 5 p_W$. Equation (5) suggests that the consumer will buy $\frac{\mu I}{p_W}$ units of water if $5 > h\left(v_{W,AJ}\right)$. This answer holds whether income allows him to buy five cups of water or five gallons of water.

Consider a more general formulation in which the strength of preference for the ideal variety depends on quantities consumed. Formally, we define a generalized compensation function, $h\left(q_{\omega}^{\gamma}, v_{\omega,\tilde{\omega}}\right)$, having the following properties:

(6) $h_2\left(q_{\omega}^{\gamma}, v_{\omega,\tilde{\omega}}\right) > 0$, $h_{22}\left(q_{\omega}^{\gamma}, v_{\omega,\tilde{\omega}}\right) > 0$ for $v_{\omega,\tilde{\omega}} > 0$, $h\left(q_{\omega}^{\gamma}, 0\right) = 1$, $h_2\left(q_{\omega}^{\gamma}, 0\right) = 0$

(7) $h\left(0, v_{\omega,\tilde{\omega}}\right) = 1$, $h_{12}\left(q_{\omega}^{\gamma}, v_{\omega,\tilde{\omega}}\right) > 0$ for $q_{\omega}, \gamma, v_{\omega,\tilde{\omega}} > 0$

(8) $h\left(q_{\omega}^{0}, v_{\omega,\tilde{\omega}}\right) = h\left(v_{\omega,\tilde{\omega}}\right)$

where the parameter $\gamma \geq 0$ defines the degree to which the consumer is finicky, or willing to forego consumption to get closer to the ideal.

The standard properties associated with the distance from the ideal variety are represented by (6). By (7) we assume that the consumer is not finicky at all at a zero consumption level, but when his consumption of a differentiated good increases he becomes increasingly finicky. Finally, (8) nests Lancaster's compensation function: if $\gamma = 0$, the compensation function does not depend on consumption volumes. An additional condition needs to be introduced to address the fact that in the generalized compensation function, the quantity of the chosen variety appears both in the nominator and in the denominator of the subutility function (4). Consequently, while the quantity consumed increases, the cost of being distanced from the ideal variety might increase so fast that it outweighs utility gains from the higher consumption level of this variety. This would contradict the standard assumption of the non-decreasing (in quantity) utility function. It is easy to show that the necessary and sufficient condition for utility to be increasing in the quantity consumed is:

(9) $h\left(q_{\omega}^{\gamma}, v_{\omega,\tilde{\omega}}\right) - \gamma p_{\omega} h_1\left(q_{\omega}^{\gamma}, v_{\omega,\tilde{\omega}}\right) > 0 \qquad \forall \omega \in \Omega$

The difference between the Lancaster's and generalized compensation functions is illustrated by Figure 1.

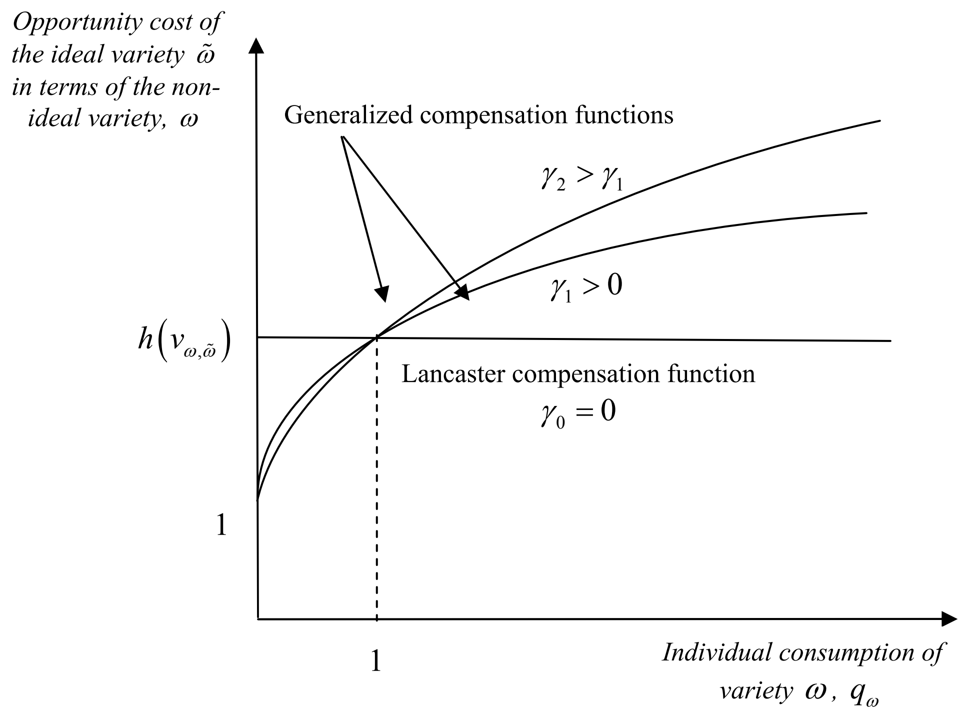

Figure 1. Lancaster and generalized compensation functions.

In order to be able to derive the closed form solution of the model, we chose a specific functional form of the generalized compensation function:

$$h\left(q_\omega^\gamma, v_{\omega,\tilde{\omega}}\right) = 1 + q_\omega^\gamma v_{\omega,\tilde{\omega}}^\beta \qquad \beta > 1,\ 0 \le \gamma \le 1, \tag{10}$$

It is easy to verify that the restrictions imposed on the parameters β and γ in (10) are necessary and sufficient for properties (6) – (9) to hold. The corresponding subutility function is then

$$u\left(q_\omega \mid \omega \in \Omega\right) = \max_{\omega \in \Omega}\left(\frac{q_\omega}{1 + q_\omega^\gamma v_{\omega,\tilde{\omega}}^\beta}\right). \tag{11}$$

Now we can apply the two-stage budgeting procedure in order to maximize (1) subject to (2) using a subutility function defined by (11). Given that the upper-case utility function is Cobb-Douglas, the consumption of the homogeneous good and expenditure on the chosen differentiated varieties are:

(12) $q_0 = (1-\mu) I$,

(13) $\int_{\omega \in \Omega} q_\omega p_\omega = \mu I$.

Consumption of a differentiated variety ω' is found by maximizing the subutility (11) subject to budget constraint (13):

(14)
$$q_{\omega'} = \frac{\mu I}{p_{\omega'}}, \qquad q_\omega = 0 \quad \text{for } \omega \neq \omega'$$
$$\text{where} \quad \omega' = \arg\min \left[p_\omega \left(1 + q_\omega^\gamma v_{\omega,\tilde{\omega}}^\beta \right) | \, \omega \in \Omega \right].$$

According to this solution, when the consumer's income increases, he becomes less price sensitive, and he values more highly the proximity to the ideal variety.

B. Market Equilibrium

We assume that the varieties of the differentiated product are produced by monopolistically competitive firms. The firms play a non-cooperative game, choosing a variety and its price under the assumption of perfect information. Each variety is produced by one firm, and firms are free to enter and exit. All varieties are produced using the same technology which is characterized by constant marginal cost and *flow* fixed cost, which is incurred so long as the variety is produced. Finally, preferences for ideal variety are uniformly distributed over the unit length circumference of the circle and the population density on the circumference is equal to L .

Under these assumptions, it is possible to show that *all* existing equilibria are zero-profit Nash equilibria. Moreover, there will exist symmetric Nash equilibria characterized by identical prices and output levels for the individual firms. In these equilibria, the specification of the varieties produced will be evenly spaced along the spectrum.[5] In the following analysis, We will focus exclusively on such symmetric equilibria in which all varieties are equally priced and equally distributed on the circumference of the circle.

Each individual is endowed with z efficient units of labor, which he supplies inelastically in the perfectly competitive labor market. The homogeneous good is produced with constant returns to scale with labor requirement equal to one. Consequently, the wage is equal to one in terms of the numeraire and an individual's income is equal to his labor endowment:

(15) $I = z$.

The production of each variety is subject to a fixed labor requirement α and marginal labor requirement c . Given that the wage equals one, α and c are also interpreted as fixed and marginal costs.

Now we would like to find the aggregate demand function and the price elasticity of demand for the produced variety ω . The solution to this problem is described by Lancaster (1984) and Helpman and Krugman (1985), and for completeness is included in Appendix 1. In the symmetric equilibrium, in which the prices of all varieties are the same and all varieties are equally distanced from each other, the demand for any produced variety $\omega \in \Omega$ is:

(16) $$Q = \frac{d\mu z L}{p},$$

where d is the shortest arc distance between any two available varieties, and p is the price of each available variety. The corresponding price elasticity of demand is:

[5] The proof of existence and the detailed characterization of equilibria is provided by Lancaster (1979). An extension of Lancaster's proof for the form of the utility function in (11) is available upon request from the author.

(17) $$\varepsilon = 1 + \frac{1}{2\beta}\left(\frac{p}{\mu z}\right)^{\gamma}\left(\frac{2}{d}\right)^{\beta} + \frac{1-\gamma}{2\beta} > 1.$$

Knowing the cost structure and the price elasticity of demand, we can find the profit-maximizing price and zero-profit quantity for each produced variety:

(18) $$p = \frac{c\varepsilon}{\varepsilon - 1} \qquad\qquad Q = \frac{\alpha}{c}\left(\varepsilon - 1\right).$$

Knowing the expenditure on each product from (13), and the size of firms from (18), we can find the equilibrium number of monopolistic firms:

(19) $$n = \frac{\mu z L}{\alpha\varepsilon}.$$

The circumference length is equal to one, so the distance between the closest varieties is:

(20) $$d = \frac{1}{n} = \frac{\alpha\varepsilon}{\mu z L}.$$

Now we can rewrite (17) using (18)–(20):

(21) $$\varepsilon = 1 + \frac{1}{2\beta}\left[\frac{c\varepsilon}{\mu z\left(\varepsilon - 1\right)}\right]^{\gamma}\left(\frac{2\mu z L}{\alpha\varepsilon}\right)^{\beta} + \frac{1-\gamma}{2\beta}.$$

The equilibrium value of the price elasticity of demand is unique, since the LHS of (21) is increasing in ε, while the RHS is decreasing in ε. From (18) and (19) we can show that the equilibrium price, quantity per variety, and number of varieties are also unique.

C. Comparative Statics

We turn to a discussion of how the equilibrium is affected by the following parameters of the model: population density L, individual labor endowment z, and the share of expenditure on differentiated good μ. The equilibrium variables of interest include the price elasticity of demand, price, output per variety, and the number of varieties. The comparative statics will be conducted for generalized and Lancaster ideal variety models and used to generate empirical predictions for Section III of this paper.

We start with the population density and its role in the generalized model. By implicit derivation of (21) we can get

$$\frac{\partial \varepsilon}{\partial L}\frac{L}{\varepsilon} = \left\{ 2\varepsilon \left[\frac{\mu z(\varepsilon - 1)}{c\varepsilon} \right]^{\gamma} \left(\frac{2\mu z L}{\alpha \varepsilon} \right)^{-\beta} + \frac{\gamma}{\beta(\varepsilon - 1)} + 1 \right\}^{-1} . \tag{22}$$

Since the resulting expression is strictly positive and strictly less than one, we can conclude that the price elasticity of demand is increasing in population density, and that it increases less than proportionally.

Let us begin the explanation of this result by defining the *market width* of variety ω, as a portion of the total spectrum of consumers buying this variety rather than some other variety.[6] The extreme values of market width in this model are one and zero, which approximate pure monopoly and perfect competition. An increase in L increases purchasing power on each interval of the spectrum, and thus each firms needs a smaller interval to get the same total revenue. As a result, in the new zero-profit equilibrium, the market width for each produced variety shrinks. Consequently, the distance between the neighboring varieties decreases, thus making consumers more sensitive to the variation in price.

Note that an increase in the price elasticity of demand decreases the equilibrium price per variety [see (18)]. This induces the secondary effect: each consumer is now choosing between higher volumes of the differentiated varieties, making him less price-

[6] This definition is identical in its meaning to Lancaster's definition (Lancaster 1979, p. 142), and it was only adopted for the notation and terminology of this paper.

sensitive. Still, the main effect is stronger than the secondary effect, and expression (22) is strictly positive. In the Lancaster model, the secondary effect does not exist since the Lancaster compensation function does not depend on prices. The corresponding derivative and percentage change can be found by setting $\gamma = 0$ in (22):

$$\text{(23)} \qquad \left.\frac{L}{\varepsilon}\frac{\partial \varepsilon}{\partial L}\right|_{\gamma=0} = \left[2\varepsilon\left(\frac{2\mu z L}{\alpha\varepsilon}\right)^{-\beta} + 1\right]^{-1}.$$

As in the generalized model, the price elasticity of demand is increasing in L and is increasing less than proportionally. However, to compare the magnitudes of the effect in the Lancaster and the generalized models, we need to know volume of individual consumption. As illustrated by Figure 1, the generalized compensation function is normalized such that its value is equal to the value of the Lancaster compensation function when the individual consumption of the differentiated good is one. The generalized compensation function is also increasing in the consumption volumes. Thus, an increase in price elasticity is smaller in the Lancaster model for the low levels of individual consumption and larger for the high levels.

Note that an increase in the population density increases not only the number of varieties [as can be inferred from (19)], but also the output per variety [see (18)]. The logic is straightforward. Entry drives up the price elasticity, lowering the markups firms can charge. Since prices are lower (and entry costs are constant) each firm must sell a higher quantity to break even. Consequently, the number of varieties increases less than proportionally to the labor force.

Next let us find the corresponding changes in equilibrium due to an increase in the individual labor endowment z. In this model an increase in z can be interpreted both as an increase in productivity and in income per capita. By implicit derivation of (21), we can find the effect of z on the price elasticity of demand:

(24)
$$\frac{\partial \varepsilon}{\partial z}\frac{z}{\varepsilon}=\left(1-\frac{\gamma}{\beta}\right)\left\{1+\frac{\gamma}{\beta\left(\varepsilon-1\right)}+2\varepsilon\left[\frac{\mu z\left(\varepsilon-1\right)}{c\varepsilon}\right]^{\gamma}\left(\frac{2\mu zL}{\alpha\varepsilon}\right)^{-\beta}\right\}^{-1},$$
$$0<\frac{\partial \varepsilon}{\partial z}\frac{z}{\varepsilon}<1.$$

An increase in z affects the equilibrium through aggregate and individual channels. The first channel can be thought of as an increase in the aggregate efficient labor endowment while keeping the individual productivity constant. The effect of this channel on the price elasticity is contained by the inverse portion of (24) and it increases the elasticity. By comparing this expression with (22), we can see that, in percentage terms, this channel predicts exactly the same changes in all variables of interest as an increase in population density does.

The second channel is associated with an increase in individual consumption levels while keeping the aggregate consumption level constant. This channel is contained by the expression $\frac{-\gamma}{\beta}\{...\}^{-1}$ of (24), and it decreases the elasticity. It can be imagined as a comparison of two countries with the same GDP where country 1 has smaller population and more productive workers – which in this model means also richer consumers – than country 2. The elasticity of demand in country 1 is lower, which means higher prices and lower output per variety. However, the number of available varieties in country 1 will be higher.

This result is interesting because it indicates that, *ceteris paribus*, an identical variety produced in both poorer and richer countries will be priced higher in the richer country. The reason is that with an increase in income, a consumer is not only increasing the volume of consumption, but he also values more highly the proximity to the ideal variety. Thus, richer consumers are willing to pay a higher price for the larger degree of diversification. The market responds by supplying more varieties, even though the economies of scale are utilized to a lesser degree for produced varieties.

The second channel does not exist in the Lancaster model. Thus, in percentage terms, variations in productivity and in labor density have an identical effect on the price elasticity of demand[7]:

$$(25)\quad \left.\frac{z}{\varepsilon}\frac{\partial \varepsilon}{\partial z}\right|_{\gamma=0} = \left.\frac{L}{\varepsilon}\frac{\partial \varepsilon}{\partial L}\right|_{\gamma=0} = \left\{2\varepsilon\left(\frac{2\mu z L}{\alpha\varepsilon}\right)^{-\beta} + 1\right\}^{-1}.$$

Finally, let us discuss the importance of the share of expenditures on the differentiated product. In percentage terms, changes in μ affect the price elasticity of demand identically to changes in z. This is true for both generalized and Lancaster ideal variety models:

$$(26)\quad \frac{\partial \varepsilon}{\partial \mu}\frac{\mu}{\varepsilon} = \frac{\partial \varepsilon}{\partial z}\frac{z}{\varepsilon} \qquad \left.\frac{\mu}{\varepsilon}\frac{\partial \varepsilon}{\partial \mu}\right|_{\gamma=0} = \left.\frac{z}{\varepsilon}\frac{\partial \varepsilon}{\partial z}\right|_{\gamma=0}.$$

Consequently, all other variables of interest will be affected by changes in μ in the same way as they are affected by changes in z. The only real difference will be noticed in the consumption of the numeraire good: while an increase in z leads to a higher consumption level of the numeraire good, an increase in μ – to a lower one.

[7] As well as on the prices quantities and number of varieties.

D. Open economy

To this point, we have focused on closed economy comparative statics. In the closed economy, our predictions for the number of varieties, quantity per variety, price elasticity, and prices refer to domestic output, which is also domestic consumption. Unfortunately, it is not possible to obtain domestic data on these variables in sufficient detail for many countries to test the model. Trade data are better in this regard, but to use them we must re-interpret the model in an open economy context. This extension is derived in detail in Appendix 2, and described briefly here.

A key weakness of the ideal variety is that analytics rely on the symmetric equilibrium in which varieties are equally spaced and have the same prices. A trivial way to extend the model to the open economy is to maintain this symmetry for foreign and domestic firms. That is, foreign varieties pay some fixed cost to enter the market that is identical to that of domestic firms, and that the delivered marginal cost of foreign varieties (inclusive of production, tariffs or other trade costs) matches domestic production costs. But even this solution is problematic: what if tariffs change? And, does foreign entry lead to domestic exit?

Instead, we rely on a solution that has been employed previously in the literature in a two-country setting by Dornbusch, Fischer, and Samuelson (1977), and in a multi-country setting by Eaton and Kortum (2002). In these models there is a continuum of homogeneous goods arranged according to strength of comparative advantage. At any point on the continuum we can compare the price of domestic goods to the delivered price of foreign goods and discern whether the products will be imported, exported, or neither.

We show in the appendix that it is possible to nest the ideal variety framework into the continuum; each good is no longer a single point on the continuum, it is now a circle. The key features of the equilibrium are that, as in DFS and EK, each country will buy goods on a particular circle from only one supplier. This ensures that all varieties in a particular product space are symmetric. In addition, it is possible to do comparative statics on how tariffs and tariff liberalization affect the equilibrium. We show that all our model predictions from the closed economy case extend to the open economy case, except that the predictions now refer to variation in import behavior across markets, not domestic production and consumption behavior.

III. Empirics – The Number and Average Size of Firms

In this section we examine predictions from the closed economy model regarding the number and size of firms.[8] Simply, the model predicts that the number of firms expands with market size, but less than proportionally so that the average size of firms is rising in market size. Conditioning on market size, growth in income per capita leads consumers to prefer a closer match to their ideal types which increases the number of firms, and lowers their average size.

To examine these predictions we employ data from the UNIDO Industrial Statistics Database. We have data on the number of establishments, total employment, value added, and gross output by 152 ISIC 3-digit industries for 54 countries from 1990-2000. There are undoubtedly important differences across countries and industries in industrial and government regulatory structure, as well as subtle differences in data

[8] We know from Hummels-Klenow (2002,2005) that open economy model predictions for traded variety and quantity per variety match the data since the model was built to match these facts.

definitions (e.g. what constitutes an "establishment"). As a consequence, cross-sectional differences in the number and size of firms are likely to be extremely noisy. Instead, we exploit the panel structure of the data to examine how changes in market size and income per capita affect changes in the number and size of establishments.

Our estimating equations for number of establishments is

$$(27) \quad \ln N_{it}^{k} = \alpha_i^k + \beta_1 \ln \frac{Y_{it}}{L_{it}} + \beta_2 \ln SIZE_{it}^{k} + e_{it}^{k}$$

where the dependent variable N_i^k is the number of establishments in country i, industry k, at time t, $SIZE_{it}^k$ is the size of the industry (measured variously as total employment, gross output, or value added), $\frac{Y_{it}}{L_{it}}$ is real GDP per capita taken from the Penn World Tables, and α_i^k is a country-industry fixed effect.

We also examine average establishment size,

$$(28) \quad \ln\left(\frac{SIZE_{it}^{k}}{N_{it}^{k}}\right) = \alpha_i^k + \beta_1 \ln \frac{Y_{it}}{L_{it}} + \beta_2 \ln SIZE_{it}^{k} + e_{it}^{k}$$

where $SIZE_{it}^k$ is again measured variously as employment, gross output, or value added. We also employ total market size for all sectors (ln GDP) in place of sector specific measures on the right hand side.

Results are reported in Table 1. In the first 3 columns we see that the prediction of the generalized ideal variety model for number of establishments is borne out. Conditioning on a country and industry, growth in the total size of an industry leads to a less than proportional expansion in the number of establishments. This holds for all three measures of industry size. The number of establishments is also rising in income per capita.

The next six columns examine the average size of firms. Regardless of the measure employed, we see that average size of establishments is increasing in industry size and decreasing in income per capita. The same holds true when we employ GDP instead of sector specific measures, though the regression fits are much lower in this case.

IV. Empirics – Own Price Elasticity of Demand

In this section we examine the generalized ideal variety model's predictions for how the own-price elasticity of demand varies across markets. Unfortunately, this model does not yield a convenient structural form for estimating the own-price elasticity. Our approach is to take as the null hypothesis that import demand is derived from a CES utility function with a common price-elasticity of demand across all markets. We then examine whether we can reject this null in favor of a model in which the elasticity varies systematically across markets. In particular, equations (22) and (24) predict that the own-price elasticity of demand is higher in large markets, and lower in rich markets (conditional on market size).

A. Methodology

The subutility function for product k (k = 6 digit HS good), for importer i, facing j = 1…J exporting sources for k is given by $u_i^k = \left(\sum_{j=1}^{J} \lambda_j^k (q_{ij}^k)^{\theta^k}\right)^{1/\theta^k}$ where $\theta^k = (\sigma_k - 1)/\sigma_k$, and λ_j^k is a demand shifter, which could represent quality differences, or (unobserved) differences in the number of distinct varieties available from each exporter. As is well known, we can write the import demands as

$$(29) \quad q_{ij}^k = \frac{E_i^k}{\Pi_i^k}\left(\frac{p_{ij}^k}{\lambda_j^k}\right)^{-\sigma^k}$$

Where E_i^k denotes expenditures, Π_i^k is the CES price index. Under the CES null, the elasticity is constant across all markets, so we can write the delivered price in market i as a function of the factory gate price at j, multiplied by ad-valorem trade costs, $p_{ij}^k = p_j^k t_{ij}^k$. When estimating this at the for k = HS 6 digit level of aggregation, everything in (29) is unobservable except the nominal value of bilateral trade and trade costs. To isolate these terms, we multiply both sides of (29) by exporter prices, and sum over all importers $c \neq i$ to get j's exports to rest of the world, r.

$$(pq)_{rj}^k = \sum\nolimits_{c\neq i}(pq)_{cj}^k = \left(\lambda_j^k\right)^{\sigma^k}\left(p_j^k\right)^{1-\sigma^k}\sum\nolimits_{c\neq i}\frac{E_c^k}{\Pi_c^k}\left(t_{cj}^k\right)^{-\sigma^k}$$

Express i's imports from j as a share of rest of world imports from j,

$$(30) \quad \ln s_{ij}^k = \ln\frac{(pq)_{ij}^k}{(pq)_{rj}^k} = \ln\frac{E_i^k}{\Pi_i^k} - \sigma^k \ln t_{ij}^k - \sum\nolimits_{c\neq i}\frac{E_c^k}{\Pi_c^k}\left(t_{cj}^k\right)^{-\sigma^k}$$

Writing this in share terms eliminates unobserved price and quality (variety) shifters specific to j.[9] We assume trade costs take the form $\ln t_{ij}^k = \ln(1+\tau_i^k) + \delta_k \ln d_{ij}$, where τ_i^k is an MFN tariff facing all exporters in importer i, product k, d_{ij} is the distance between countries, and δ^k is the elasticity of trade costs with respect to distance.

[9] Alternatively, we could also write (30) by expressing i's imports relative to any particular importer, or set of importers, rather than the world.

To simplify this expression, we employ importer i – product k fixed effects α_i^k (implemented by mean differencing) which eliminates the importer expenditure share, the CES price index, and MFN tariff rates. This leaves variation in bilateral distance to trace out the variation in trade costs. The final term we assume to be orthogonal to the included trade barriers and include it in the error.[10] We now have

$$(31) \quad \ln s_{ij}^k = \alpha_i^k - \sigma^k \delta^k \ln d_{ij} + e_{ij}^k$$

In the CES model, we can interpret the coefficient on distance as $\beta^k = -\sigma^k \delta^k$, which is invariant to the importer. We will test whether the constant elasticity is rejected by the data in favor of a form consistent with the generalized ideal variety model, by interacting distance with importer GDP and GDP per worker.

$$(32) \quad \ln s_{ij}^k = \alpha_i^k - \beta_1^k \ln d_{ij} + \beta_2^k \ln d_{ij} \ln Y_i + \beta_3^k \ln d_{ij} \ln \frac{Y_i}{L_i} + e_{ij}^k$$

Before proceeding to the results, a few notes regarding interpretation are in order. Ideally, we would estimate (32) separately for each exporter and commodity in order to examine how the own-price elasticity of demand varies across markets for the same product. However, in order to identify the importer-commodity fixed effects it is necessary to pool over multiple exporters. This pooling is not the same thing as identifying a cross-price elasticity, i.e. how imports of Japanese televisions change when

[10] Since we cannot measure the price indices or the elasticity of substitution it is difficult to include this last term explicitly. We cannot verify that trade costs between i and j are orthogonal to the real expenditure weighted sum of trade costs between j and all other countries. However, simple proxies for this term such as a sum over nominal GDP weighted distances are very weakly correlated with distances and tariffs between i and j. We later show that our results are robust to an alternative specification in which this omitted term does not appear.

the price of Korean televisions rise. Instead, it is equivalent to restricting the own-price elasticity to be the same across all exporter and products over which we pool, i.e. imports of Japanese TVs respond to a change in the price of Japanese TVs in the same way that imports of Korean TVs respond to a change in the price of Korean TVs. In the estimates that follow, we employ two pooling strategies. For simplicity, we first pool over all exporters and 6 digit products. Then, we pool over all exporters and 6 digit products within a particular 2 digit aggregate. In both cases, the importer fixed effects are still calculated with respect to the 6 digit product.

Second, the use of bilaterally varying trade costs identifies price variation under the CES null, but not in the variable elasticity case. With variable elasticity preferences a rise in trade costs will be partially offset by a fall in the factory gate price so that only a part of the trade cost is passed through to the final price. That is, the true destination price includes a pricing-to-market adjustment, which is an omitted variable in our specification that is negatively correlated with trade costs. This omission will create a bias in the price elasticity toward zero. For a similar reason, if the interaction terms are significant, PTM will cause a bias in these estimates toward zero. This is problematic if we want to precisely identify own-price elasticity of demand. It is less concerning if our primary interest lies in testing the CES null since we will be biased toward not finding a significant interaction between tariffs and importer characteristics.

B. Results

Table 2 reports the results of estimating equation (32) by pooling over all exporters and products. We can immediately reject the hypothesis that the response of imports to price changes (via trade costs) is the same in all markets, as both interaction terms are significant, with signs matching the theory. To gauge the impact of the interaction terms, the elasticity of the import share to distance (evaluated at the means of lnY and lnY/L) is -1.243. Starting at the mean and doubling GDP increases the distance elasticity to -1.3, while doubling GDP per capita reduces it -1.14.

Of course, not all products are likely to fit the model equally well, and the pooling restrictions necessary to take Table 2 estimates seriously the are unlikely to be met. Accordingly, we estimate equation (32) separately for each 2 digit HS product. Full details for each 2 digit regression are in Appendix Table 1, and we summarize the distribution of the interaction terms in Figure 2. These are histograms on interaction term point estimates for 96 HS2 regressions, reported both on a simple count basis, and weighting the products by their value share in our sample of world trade. Nearly all the mass for the GDP interaction terms lies between 0 and -0.2. At the 10% level, 55 of the 96 industries, representing 84 percent of trade by value have negative signs and are significant (5 percent by value are positively signed and significant). Nearly all the mass for the GDP per capita interaction term lies between 0 and 0.2. At the 10% level, 54 of the industries, representing 76 percent of trade by value have positive signs and are significant (4 percent are negatively signed and significant). In 67 percent of trade (by value) both interaction terms are right signed and significant. It is clear from these figures that, while the effect differs significantly across industries, the basic message of the interaction from the pooled regression comes through.

The model performs well on sign and significance, but does it imply significant differences in the price elasticity of demand across markets? A problem with interpreting these interaction terms is that we have a product of the price elasticity and an elasticity of trade costs with respect to distance.

$$(33)\quad \hat{\beta}_1^k \ln d + \hat{\beta}_2^k \ln d \ln Y + \hat{\beta}_3^k \ln d \ln\left(Y/L\right) = \delta^k (\sigma^k + \sigma_y^k \ln Y + \sigma_{y/l}^k \ln\left(Y/L\right))$$

To isolate the price elasticity, we can express the combined distance and interaction terms as a ratio for countries of different size and income. For countries 1 and 2, we have

$$(34)\quad \frac{\delta^k}{\delta^k} \frac{(\sigma^k + \sigma_y^k \ln Y_1 + \sigma_{y/l}^k \ln\left(Y/L\right)_1)}{(\sigma^k + \sigma_y^k \ln Y_2 + \sigma_{y/l}^k \ln\left(Y/L\right)_2)}$$

Note that the elasticity of trade costs with respect to distance falls out, leaving only the elasticity of substitution and any interaction effects with importer Y and Y/L. We provide several calculations designed to show the range of price elasticities over countries in the sample.

First, for each HS 2 product we take the regression point estimates, and combine them with importer data on Y and Y/L in order to calculate the combined interaction effects for each country in (33). We then rank them from most to least elastic, and express the elasticity ratio in (34) using the 90th percentile / 10th percentile country. This gives, for each HS 2 product, a measure of the range of elasticity over importers in the sample. In Figure 3a we plot a distribution of this statistic over all HS 2 products. A value of two means that the price elasticity of demand is twice as high as in the 10th percentile country. Most of the distribution lies between 1.2 and 2.5.

We also want to separately isolate the impact of Y variation and Y/L variation on the elasticity. First we hold lnY fixed at the sample mean and generate variation in (33) from importer variation in lnY/L, rank countries as before, then express the ratio in (34) for the 90th percentile / 10th percentile country in each HS 2 product. Figure 3b plots a distribution over HS 2 products, showing the variation in elasticity coming only from Y/L variation. Figure 3c repeats this exercise, except that we hold Y/L fixed at the sample mean and generate variation in (33) from importer variation in Y. While the magnitude of the interaction effects on Y and Y/L were similar, Figures 5 and 6 show that Y variation generates more dispersion in the elasticity. This is because there is much greater sample variation in Y than in Y/L.

C. Robustness: tariffs as an alternative trade cost measure

Ideally we would have data on bilaterally varying ad-valorem trade costs, rather than using a distance proxy for them. Our TRAINS also includes tariff rates for each trade observation, but has insufficient variation across exporters for a given importer-product for us to use this variable in estimating (32). There is some variation in the tariff *schedule* across export sources for a given importer and six digit HS, but most of the high tariff observations correspond to zero trade values. Of the pairs where trade is observed, in about 90 percent of the cases $TAR_{ij}^{k} = \text{median}(TAR_{i}^{k})$. This means that tariffs in the data can reasonably be treated as identical across exporters as we model above, in which case employing importer-product (ik) fixed effects eliminates tariffs as a useful source of variation for trade costs

We experimented with an alternative approach in which we used the tariff variation for trade costs. Starting from (29), we multiply both sides by exporter prices to get (observable) nominal values for bilateral trade. We include exporter j – commodity k fixed effects, α_j^k, to eliminate all exporter-specific effects, including the last term in (30) which we had omitted as unmeasurable in our primary specification. Finally, we proxy for real expenditures in equation (30) using importer GDP and GDP per capita so that the estimating equation becomes

$$(35) \quad \begin{aligned} \ln pq_{ij}^k &= \alpha_j^k + \gamma_1 \ln Y_i + \gamma_1 \ln \frac{Y_i}{L_i} + \beta_4^k \ln d_{ij} \\ &+ \beta_1^k \ln \tau_{ij} + \beta_2^k \ln \tau_{ij} \ln Y_i + \beta_3^k \ln \tau_{ij} \ln \frac{Y_i}{L_i} + e_{ij}^k \end{aligned}$$

The advantage of this specification is that one can read the coefficients on tariffs and its interactions directly in terms of a price elasticity σ^k, rather than as the product of two elasticities $\sigma^k \delta^k$. The disadvantage is that we omit the importer price index, which could plausibly affect our estimated interaction terms.[11]

We estimate (35) by pooling over all exporters and HS6 products within an HS2 aggregate . Figure 4 reports the distributions on the interaction terms. The same basic message from Figures 2 and 3 goes through. For roughly two-thirds of trade by value, the signs match the predictions of the generalized ideal variety model. Larger countries exhibit a higher price elasticity of demand. Conditioning on size, richer countries exhibit a lower price elasticity of demand.

Unlike the specification employing distance as a measure of trade costs, here we can directly interpret the interaction coefficients in terms of their impact on the absolute price elasticity. A coefficient of -0.5 on the GDP x tariff term implies that doubling GDP yields demand that is 0.5 percentage points more elastic. Given the range of GDP variation in the world, this represents substantial variation in the price elasticity.

[11] To explain, large countries are likely to have larger domestic industries. If we believe the CES null, this translates into a lower price index due to the value of greater variety (see Feenstra 1994).

V. Empirics -- Cross-Importer Variation in Prices

In this section we empirically examine the theoretical predictions regarding cross-importer prices: prices should be lower in large markets and higher in rich markets. Unfortunately, we cannot employ the full TRAINS dataset as quantity units are inconsistently measured across importers. Instead, we use a small subset of the data, taken from Hummels and Skiba (2004), for which we have consistent quantity units and other necessary variables.

The data cover the bilateral trade of six importers (Argentina, Brazil, Chile, Paraguay, Uruguay, and the United States) with all exporters worldwide, measured at 6-digit HS level. Quantity is measured using the shipment weight, so prices are the f.o.b. value of trade divided by weight. Since we pool over commodities, commodity fixed effects can be thought of as a conversion of a common price measure (value per pound) into commodity specific units.

The regression specification examines variation in prices across import market characteristics (Y, Y/L), after including fixed effects α_j^k to absorb any differences in quality or price that are specific to the exporter-commodity. Hummels-Skiba (2004) show that these prices will also depend on tariffs, τ_{ij}^k, and per unit transportation costs f_{ij}^k due to Alchian-Allen effects. Tariffs are taken from TRAINS in Section IV. These six importers also report freight paid on each shipment, enabling us to calculate a per unit transportation cost. This gives us

$$\ln p_{ij}^k = \alpha_j^k + b_1 \ln Y_i + b_2 \ln \frac{Y_i}{L_i} + b_3 \ln f_{ij}^k + b_4 \ln(1+\tau_{ij}^k) + e_{ij}^k \quad (36)$$

Finally, the freight bill may be also rising in prices due to higher insurance and handling costs. To deal with the endogeneity problem, we follow Hummels-Skiba (2004) in instrumenting the freight bill with distance shipped and the shipment weight.

The IV estimates of equation (36) are reported in table 2. The first row reports results using the full sample of all importers and exporters, the second omits the US. The elasticity of prices with respect to market size is (-0.014). (INTERPRETATION: compare to above?)

The elasticity of price with respect to importer Y/L is 0.36, consistent with our prediction that firms increase markups in high income countries. Of course, some of this effect may be due to non-homothetic demands for quality. Recall however that our estimates condition on exporter-commodity, which sweeps out much of the cross-exporter quality variation found, for example, in Schott (2003), Hallak (2004), and Hummels-Klenow (2005).

VI. Conclusion

We derive a generalized version of the Lancaster (1979) ideal variety model in order to match known empirical facts regarding the co-variation between importer size and imported variety and quantity per variety. In this model, entry leads to a "crowding" of variety space, so that larger markets exhibit a higher own-price elasticity of demand for differentiated goods, lower prices, and a larger average firm size. Working against this crowding is an income effect: as consumers grow rich and quantities consumed rise, their strength of preference for their ideal variety also rises. This gives firms greater pricing power over consumers. Conditioning on market size, richer markets see a lower own-price elasticity, higher prices, and fewer firms.

We provide new evidence supporting the model's predictions regarding average firm size, prices and the own-price elasticity of demand. Conditioning on country and industry and exploiting time series variation, average firm size is rising for countries with rising GDP, and falling for countries with rising GDP per worker. Conditioning on an exporter and product and exploiting cross-importer variation, the own-price elasticity of demand is higher in large markets, and lower in rich markets. In an exercise exploiting data variation similarly but using a smaller sample of importers, we find that prices are lower in large markets, and higher in rich markets.

We see three implications of these findings. First, the theoretical and empirical literature on product differentiation in trade has relied almost exclusively on constant-elasticity-of-substitution utility functions. While these models are highly tractable, they yield counter-factual implications on central empirical questions.

Second, as has been pointed out by Romer (1994) and Feenstra (1994) and the literature they have inspired, CES utility models imply substantial welfare gains from trade in new varieties. Evaluating the welfare implication of new varieties in the generalized ideal variety model is beyond the scope of the current paper. However, our results suggest two important qualifications for existing welfare studies. First, variety space does appear to fill up with entry, suggesting that the welfare gains from new variety may be substantially lower in large countries than in small. Second, the news is not all bad in the sense that income effects partially trump the crowding effect for some goods. Rich consumers want, and are willing to pay for, varieties very closely matched to their ideal preferences. GDP growth that occurs primarily through growth in output per worker will still lead to substantial variety gains for some goods, albeit at the cost of lowered economies of scale and higher prices.

Finally, we know that prices are systematically higher in rich than in poor countries, a fact that has typically been ascribed to cross-country differences in the prices of non-traded goods as in Balassa and Samuelson. Our results show that price elasticities and prices of traded goods also systematically covary with incomes. Whether these traded goods price differences are a significant contributor to national price levels as a whole, and constitute a challenge to the centrality of non-traded goods in explaining them, we leave for subsequent work.

References

Acemoglu, D and Ventura, J "The World Income Distribution," *Quarterly Journal of Economics* 117 (May 2002), 659-694.

Barron, J., Taylor, B., and Umbeck, J. *International Journal of Industrial Organization* 2004 forthcoming.

Brander, J and Krugman, P. "A 'Reciprocal Dumping' Model of International Trade", *Journal of International Economics*, 15 (1982), pp. 313-321

Broda, C. and Weinstein, D. "Globalization and the Gains from Variety" NBER WP 10314, 2004.

Dixit, A. and Stiglitz, J. E. "Monopolistic Competition and Optimum Product Diversity." *American Economic Review*, June 1977, *67*(3), 297-308.

Dornbusch R., Fischer S., and Samuelson P. A. "Comparative Advantage, Trade, and Payments in a Ricardian Model with a Continuum of Goods." *American Economic Review*, December 1977, *67* (5), 823-39.

Eaton J. and & Kortum S. "Technology, Geography, and Trade." *Econometrica*, September 2002, *70*(5), 1741-79.

Feenstra, R. C. "Symmetric Pass-Trough of Tariffs and Exchange Rates Under Imperfect Competition: An Empirical Test." *Journal of International Economics*, 1989, *16*, 227-42.

Feenstra, R.C. "New Product Varieties and the Measurement of International Prices," *American Economic Review* 84, March 1994, 157-177.

Goldberg, P. and Knetter, M. "Goods Prices and Exchange Rates: What Have We Learned?" *Journal of Economic Literature,* September 1997, *35*(3), 1243-72.

Head, K. and Ries, J. "Increasing Returns versus National Product Differentiation as an Explanation for the Pattern of U.S.-Canada Trade." *American Economic Review*, September 2001, *91* (4), 858-76.

Hallak, JC "Product Quality, Linder, and the Direction of Trade" NBER WP 10877, 2004.

Helpman, E. and Krugman, P. R. *Market Structure and Foreign Trade*, Cambridge MA: MIT Press, 1985.

Hummels, D. and Klenow, P. J. "The Variety and Quality of a Nation's Trade." *NBER WP 8712* 2002.

Hummels, D. and Klenow, P. J. “The Variety and Quality of a Nation’s Exports.” *American Economic Review*, forthcoming, 2005.

Hummels, D. and Skiba, A. “Shipping the Good Apples Out? An Empirical Confirmation of the Alchian-Allen Conjecture.” *Journal of Political Economy*, 2004.

Klenow, P. and Rodriguez-Clare, A. “Quantifying Variety Gains from Trade Liberalization” mimeo. University of Chicago, 1997.

Knetter, M. M. “International Comparison of Pricing-to Market Behavior.” *American Economic Review*, June 1993, *83*(3), 473-89.

Krugman, P. R. “Scale Economies, Product Differentiation, and the Pattern of Trade.” *Journal of International Economics*, November 1979, *9*, 469-80.

______. ”Increasing Returns, Monopolistic Competition, and International Trade.” *American Economic Review*, December 1980, *70*, 950-9.

Lancaster, K. *Variety, Equity, and Efficiency*. New York: Columbia University Press, 1979.

______. “Intra-Industry Trade under Perfect Monopolistic Competition.” *Journal of International Economics*, 1980, *11*, 858-64.

______. “Protection and Product Differentiation.” In H. Kierzkowski (ed.), *Monopolistic Competition and International Trade.* Oxford: Clarendon Press, 1984.

Perloff, J.M. and Salop, S.C. “Equilibrium with Product Differentiation” *Review of Economic Studies,* Vol. 52, No. 1 (Jan., 1985), pp. 107-120

Romer, P. M. “New Goods, Old Theory, and the Welfare Costs of Trade Restrictions.” *Journal of Development Economics*, February 1994, *43*(1), pp. 5-38.

Schott, Peter K., “Across-Product versus Within-Product Specialization in International Trade,” *Quarterly Journal of Economics* 2003.

Syverson, C. “Market Structure and Productivity: A Concrete Example”, *Journal of Political Economy*, Vol 112, Num 6, Dec 2004, pp.1181-1222.

Venables, A. J. “Integration and the Export Behaviour of Firms: Trade Costs, Trade Volumes and Welfare.” *Weltwirtschaftliches Archiv,* 1994, *130* (1), 118-32.

Table 1: The Number and Size of Firms

	Number of Establishments			Average Establishment Size: Employment		Gross Output		Value Added	
GDP per capita	1.279 (0.037)	0.774 (0.040)	0.864 (0.040)	-1.631 (0.050)	-1.264 (0.036)	-0.134 (0.049)	-0.703 (0.039)	-0.256 (0.052)	-0.804 (0.040)
GDP				0.246 (0.019)		0.040 (0.015)		0.110 (0.016)	
Industry Total Employment	0.669 (0.009)				0.300 (0.009)				
Industry Total Gross Output		0.449 (0.008)					0.531 (0.008)		
Industry Total Value Added			0.367 (0.008)						0.605 (0.008)
Constant	-0.044 (0.003)	-0.037 (0.003)	-0.037 (0.003)	0.038 (0.003)	0.039 (0.002)	0.023 (0.003)	0.038 (0.003)	0.025 (0.003)	0.038 (0.003)
Observations	21983	21014	21043	21983	21983	21014	21014	21043	21043
R2	0.229	0.163	0.131	0.053	0.090	0.000	0.166	0.002	0.213

Notes:
1. Panel regression of equations XX and YY includes country-ISIC3 industry fixed effects
2. All variables are in logs, standard errors are in parentheses.
3. Regression R2 are net of fixed effects

Table 2: Trade Responses to Trade Costs and Importer Characteristics

Dep var: $\ln\left(s_{ij}^{k}\right)$	$\ln dist_{ij}$	$lang_{ij}$	$\ln Y_i \times \ln dist_{jk}$	$\ln\left(Y/L\right)_i \times \ln dist_{ij}$	R^2	Obs.
	-.668 (.027)	.533 (.005)	-.057 (.001)	.101 (.003)	.17	1,183,696

Notes:
1. Estimates of equation (), pooled over all exporters and HS codes.
2. Standard errors are in parentheses, all coefficients significant at 1% level.
3. Regression R2 are net of importer-HS6 fixed effects.

Table 3: Price Variation and Importer Characteristics

Dep var: ln(price)	Variables (in logs)				R^2	Obs
	Per-unit freight cost	Tariff rate	GDP (importer)	GDP per capita (importer)		
Full Sample	.851 (.002)	-1.161 (.031)	-.014 (.001)	.360 (.004)	.53	275,398
Sample w/o US	.837 (.003)	-.961 (0.38)	-.015 (.002)	.340 (.005)	.49	179,381

Notes:
1. IV Estimates of equation (), including exporter-hs6 fixed effects.
2. Shipment distance and weight are used to instrument per unit freight cost.
3. Standard errors in parentheses, all coefficients significant at 1% level.

Figure 2. Importer Characteristics Interacted with Distance

Point Estimates (weighted by count and by value)

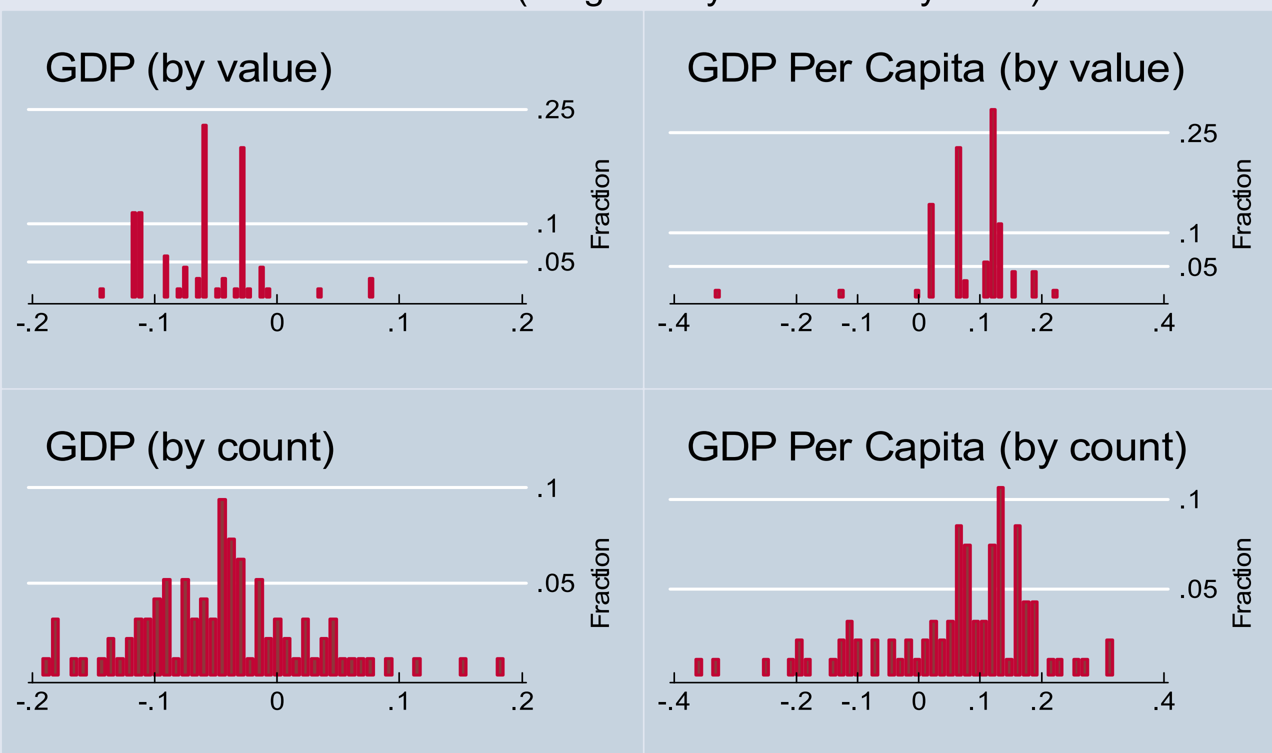

Figure 3. Implied Range of Elasticity Over Importers

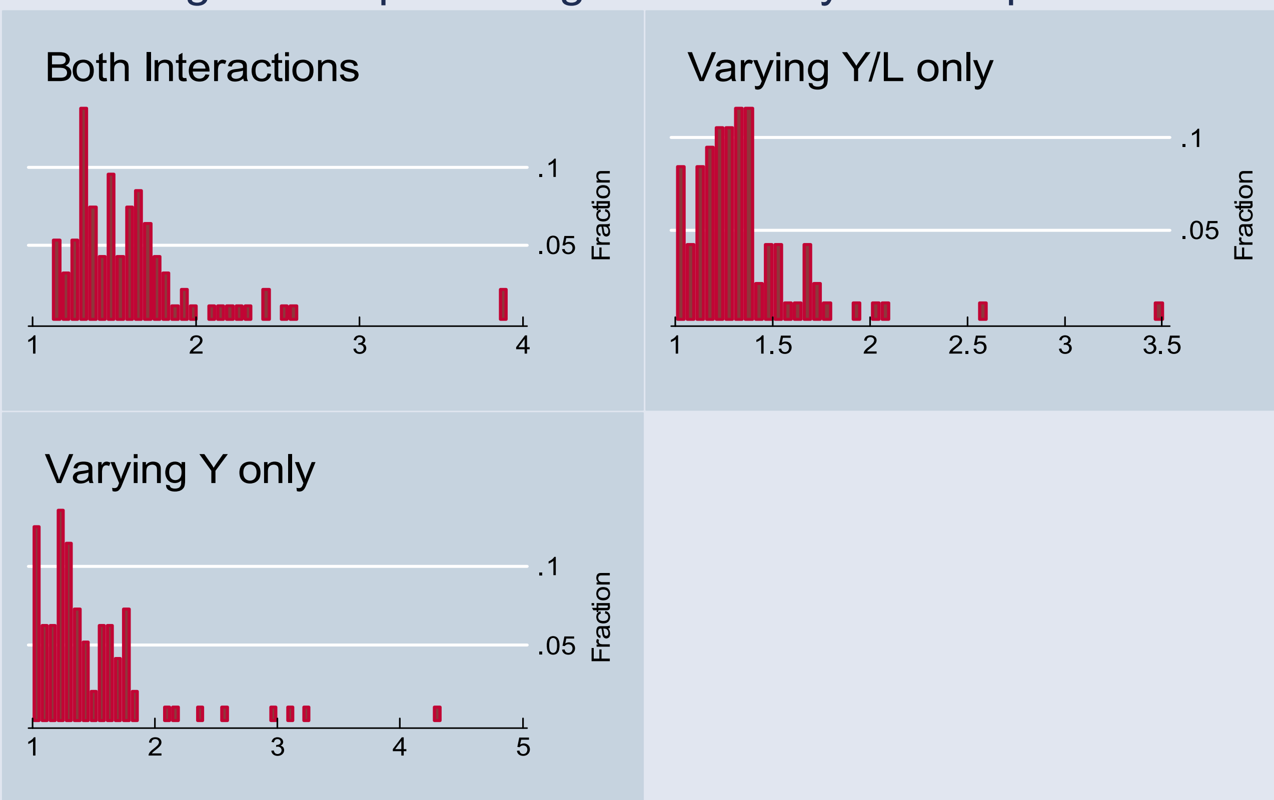

Figure 4. Importer Characteristics Interacted with Tariff

Point Estimates (weighted by value)

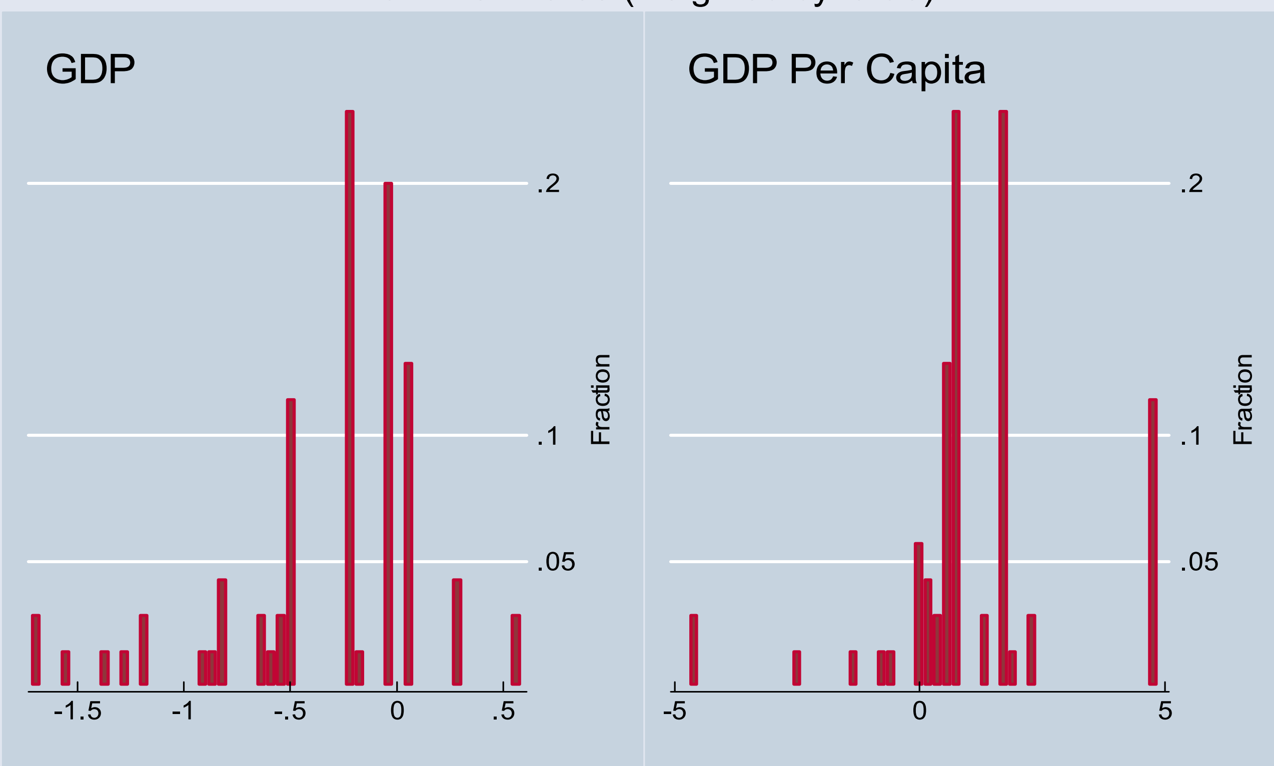

Appendix 1. Derivation of the aggregate demand and price elasticity of demand for the differentiated varieties in the symmetric equilibrium.

The derivation of the market demand and of the price elasticity of demand in this appendix can be considered as a modification of the corresponding derivations provided by Helpman and Krugman (1985, Chapter 6).

First we would like to find the aggregate demand function for variety $\hat{\omega}$ given that its closest competitor to the left is variety $\underset{\frown}{\omega}$, and its closest variety to the right is $\breve{\omega}$. The corresponding prices are denoted as $p_{\underset{\frown}{\omega}}$, $p_{\hat{\omega}}$ and $p_{\breve{\omega}}$. Next let us choose the varieties $\underline{\omega}, \overline{\omega} \in d^{*}\left(\underset{\frown}{\omega}, \breve{\omega}\right)$ such that

(A1) $$\begin{aligned} p_{\underset{\frown}{\omega}}\left(1+q^{\gamma}_{\underset{\frown}{\omega}} v^{\beta}_{\underset{\frown}{\omega},\underline{\omega}}\right) &= p_{\hat{\omega}}\left(1+q^{\gamma}_{\hat{\omega}} v^{\beta}_{\hat{\omega},\underline{\omega}}\right) \\ p_{\breve{\omega}}\left(1+q^{\gamma}_{\breve{\omega}} v^{\beta}_{\breve{\omega},\overline{\omega}}\right) &= p_{\hat{\omega}}\left(1+q^{\gamma}_{\hat{\omega}} v^{\beta}_{\hat{\omega},\overline{\omega}}\right), \end{aligned}$$

where $d^{*}\left(\underset{\frown}{\omega}, \breve{\omega}\right)$ is the shortest arc between $\underset{\frown}{\omega}$ and $\breve{\omega}$. From (5) we know that all prices are symmetric. Consequently, the market clientele for variety $\hat{\omega}$ is a compact set of consumers whose ideal varieties range from $\underline{\omega}$ to $\overline{\omega}$. Note that from the first stage of the two-stage budgeting procedure we know the individual consumption levels for each produced variety $\widehat{\omega}$:

(A2) $$q_{\hat{\omega}} = \frac{\mu I}{p_{\hat{\omega}}}$$

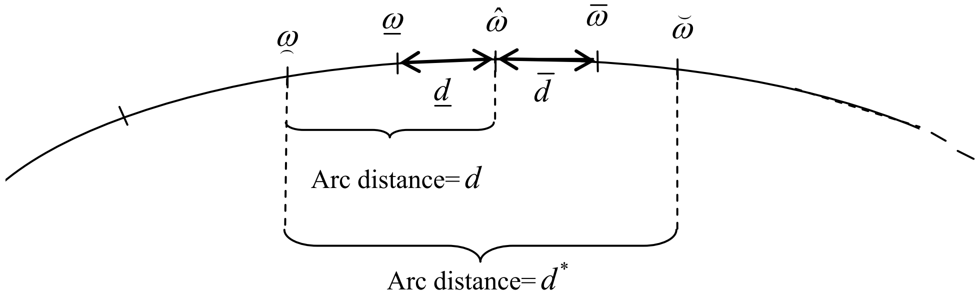

Figure A1.

In what follows, all varieties are identified by the shortest arc distance from variety $\underset{\frown}{\omega}$: variety $\hat{\omega}$ is represented by d, variety $\underline{\omega}$ is represented by $\left(d-\underline{d}\right)$ where $\underline{d}=v_{\underline{\omega},\hat{\omega}}$, and variety $\overline{\omega}$ is represented by $d+\overline{d}=v_{\underline{\omega},\overline{\omega}}$ where $\overline{d}=v_{\hat{\omega},\overline{\omega}}$. Figure A1 illustrates these identifications graphically. Now we can update our notation and substitute (A2) into (A1):

$$\text{(A3)}\quad \begin{aligned} p_{\underline{\omega}}\left[1+\mu^{\gamma}I^{\gamma}p_{\underline{\omega}}^{-\gamma}\left(d-\underline{d}\right)^{\beta}\right] &= p_{\hat{\omega}}\left(1+\mu^{\gamma}I^{\gamma}p_{\hat{\omega}}^{-\gamma}\underline{d}^{\beta}\right) \\ p_{\breve{\omega}}\left[1+\mu^{\gamma}I^{\gamma}p_{\breve{\omega}}^{-\gamma}\left(d^{*}-d-\overline{d}\right)^{\beta}\right] &= p_{\hat{\omega}}\left(1+\mu^{\gamma}I^{\gamma}p_{\hat{\omega}}^{-\gamma}\overline{d}^{\beta}\right) \end{aligned},$$

where $p_{\underline{\omega}}$, $p_{\breve{\omega}}$, and $p_{\hat{\omega}}$ denote the prices of the corresponding varieties.

From (A3) we can express the boundaries of the firm's clientele as a function of the distance between its closest competitors' varieties, their pricing $\left(p_{\underline{\omega}} \text{ and } p_{\breve{\omega}}\right)$, the firm's own pricing $\left(p_{\hat{\omega}}\right)$ and variety choice (as measured by d), and individual income spent on the differentiated good:

$$\text{(A4)}\qquad \begin{aligned} \underline{d} &= \underline{v}\left[p_{\hat{\omega}}, d, p_{\underset{\smile}{\omega}}, p_{\breve{\omega}}, d^{*}, \mu I\right] \\ \overline{d} &= \overline{v}\left[p_{\hat{\omega}}, d, p_{\underset{\smile}{\omega}}, p_{\breve{\omega}}, d^{*}, \mu I\right] \end{aligned}$$

Thus we can write the demand function faced by a firm producing variety $\hat{\omega}$ as:

$$\text{(A5)}\qquad Q_{\hat{\omega}} = \frac{\left[\underline{v}(.) + \overline{v}(.)\right]\mu IL}{p_{\hat{\omega}}}$$

where μIL is the aggregate expenditure on the differentiated varieties.

Next let us derive the price elasticity of demand function defined by (A5). To do it, we will first apply the implicit derivation to (A4) in order to find the response of the market width towards an increase in price:

$$\text{(A6)}\qquad \begin{aligned} \frac{\partial \underline{v}}{\partial p_{\hat{\omega}}} &= -\frac{1-(1-\gamma)\, p_{\hat{\omega}}^{-\gamma}\, \underline{d}^{\beta}}{p_{\underset{\smile}{\omega}} \beta p_{\underset{\smile}{\omega}}^{-\gamma}\left(d-\underline{d}\right)^{\beta-1} + p_{\hat{\omega}} \beta p_{\hat{\omega}}^{-\gamma}\left(\underline{d}\right)^{\beta-1}} < 0 \\ \frac{\partial \overline{v}}{\partial p_{\hat{\omega}}} &= -\frac{1-(1-\gamma)\, p_{\hat{\omega}}^{-\gamma}\, \underline{d}^{\beta}}{p_{\breve{\omega}} \beta p_{\breve{\omega}}^{-\gamma}\left(d^{*}-d-\overline{d}\right)^{\beta-1} + p_{\hat{\omega}} \beta p_{\hat{\omega}}^{-\gamma}\left(\overline{d}\right)^{\beta-1}} < 0 \end{aligned}$$

where the nominators of both fractions are strictly positive according to (9). Recall that we are focusing on the symmetric equilibria, and thus all prices are symmetric and $\left(d^{*}-d-\overline{d}\right) = \left(d-\underline{d}\right) = \underline{d} = \overline{d} = \frac{d}{2}$. Combining this fact with (A6), we can derive the price elasticity of demand from (A5):

$$\text{(A7)}\qquad \varepsilon = 1 + \frac{1}{2\beta}\left(\frac{p}{\mu I}\right)^{\gamma}\left(\frac{2}{d}\right)^{\beta} + \frac{1-\gamma}{2\beta} > 1$$

Appendix 2. Open Economy

The model outlined in this section can be considered as a discrete version of Dornbusch-Fischer-Samuelson (1977) model. In particular, consumer's preferences are defined over a finite number of products, each of which is defined over a continuum of varieties. Each country is assumed to have a distribution of technologies across products, so that only the varieties of the same product are produced with the same technology. Moreover, these distributions are assumed to be asymmetric across countries, so that there are incentives for inter-product international trade. In equilibrium, all varieties of the same product will be priced symmetrically within each country, though prices of the same varieties might differ across countries. This modeling strategy avoids the substantial technical complications arising with the asymmetry of prices in the standard Lancaster model of international trade (see Lancaster 1984).

A. Setup

Imagine that the world consists of two countries, Home and Foreign, which are indexed by superscripts H and F, respectively. The utility function is the same for both countries, but it is slightly different from (1): now we assume that the preferences of a consumer are defined over a homogeneous product q_0 and over M differentiated products q_k, $k = 1, 2, ..., M$, where each of the differentiated products is defined by a continuum of varieties indexed by $\omega \in \Omega_k$:

(A8) $$U = q_0^{1-\mu} \prod_{k=1}^{M} \left[u_k \left(q_{i\omega} \mid \omega \in \Omega_k \right) \right]^{\frac{\mu}{M}} \qquad 0 < \mu < 1,$$

where subutility $u_k \left(q_{k\omega} \mid \omega \in \Omega_k \right)$ is defined as

(A9) $$u_k\left(q_{k\omega} \mid \omega \in \Omega_k\right) = \max_{\omega \in \Omega_k}\left(\frac{q_{k\omega}}{1 + q_{k\omega}^{\gamma} v_{\omega,\tilde{\omega}}^{\beta}}\right),$$

and the budget constraint is

(A10) $$q_0 + \sum_{k=1}^{M} \int_{\omega \in \Omega_k} q_{k\omega} p_{k\omega} = I .$$

The differentiated varieties are still assumed to be produced by monopolistically competitive firms, and the technologies are characterized by the fixed and marginal labor requirements. However, we now introduce some additional assumptions about the structure of the costs. We now interpret α not as the fixed cost of production, but as the cost of adjustment to each market. Consequently, the fixed cost is incurred for each market the firm chooses to enter. The fixed cost of market adjustment is assumed to be symmetric across the countries, products, and varieties:

(A11) $$\alpha_{k\omega}^{H} = \alpha_{k\omega}^{F} = \alpha \qquad \forall \omega \in \Omega_k ,\ k = 1, 2, ..., M .$$

In contrast, the marginal costs are assumed to differ across products and countries, while remaining the same for all varieties of a given product and country. In particular, we assume that the Home's marginal cost is linearly increasing in the index of the product:

(A12) $$c_k^{H} = ck \qquad c > 0 ,\ k = 1, 2, ..., M ,$$

while, for Foreign, the order of marginal costs across products is the reverse:

(A13) $$c_k^{F} = c\left(M - k + 1\right) \qquad c > 0 ,\ k = 1, 2, ..., M .$$

Given the distribution of costs, each country has a Ricardian comparative advantage in the production of a certain subset of products, which stimulates inter-product trade. The degree of comparative advantage differs across products: it is decreasing in k for Home and increasing in k for Foreign.

The transportation costs for the differentiated goods are of the 'iceberg' form, and they are identical for all varieties of all the products and for both directions of trade:

(A14) $t_{k\omega}^{HF} = t_{k\omega}^{FH} = t \geq 1$ for any $\omega \in \Omega_k$, $k = 1, 2, ..., M$.

The homogeneous good can be traded at no cost. It is included in the model to guarantee balanced trade and equality of wages across countries: the numeraire sector is assumed to be large enough for each country to produce homogenous product under free trade.

B. Market Equilibrium

Now, in addition to the domestic varieties, consumers can potentially access the imported varieties. Hence, our first step is to find out which varieties will be traded internationally. Let us start the solution of this problem by considering a single product $l \in \{1, 2, ..., M\}$. We are interested in establishing whether the consumption of product l in Home consists of only domestic varieties, of only imported varieties, or both. It is possible to show that the answer depends on the comparison of Home's marginal cost to Foreign's marginal cost adjusted for the trade costs.

First, consider the case when $c_l^H > tc_l^F$. Assume that there exists a domestic firm producing variety $\omega \in \Omega_l$ and earning nonnegative profit by selling the amount $Q_{l\omega}$ at price $p_{l\omega}$ in the Home's market. Then there exists $p < p_{l\omega}$ such that a Foreign's producer can sell the same amount $Q_{l\omega}$ at Home's market at price p and earn a strictly positive profit. Moreover, this entry will always occur due to the free entry condition. By undercutting the Home producer's price, Foreign's producer will crowd out Home's producer from the market. Thus, by contradiction, none of Home's producers is able to earn nonnegative profit at Home's market of product l, and all varieties of this product in the Home's market will be imported from Foreign. In a similar fashion, it is possible to show that, if $c_l^H < tc_l^F$, no varieties of product l will be imported from Foreign by Home. Finally, if $c_l^H = tc_l^F$, the varieties of product l consumed at Home, can be produced both in Home and in Foreign. Similar analysis for Foreign can be conducted to determine the production location for each good l consumed in Foreign.

Figure A2 provides the graphical representation of the trade equilibrium on the product level.

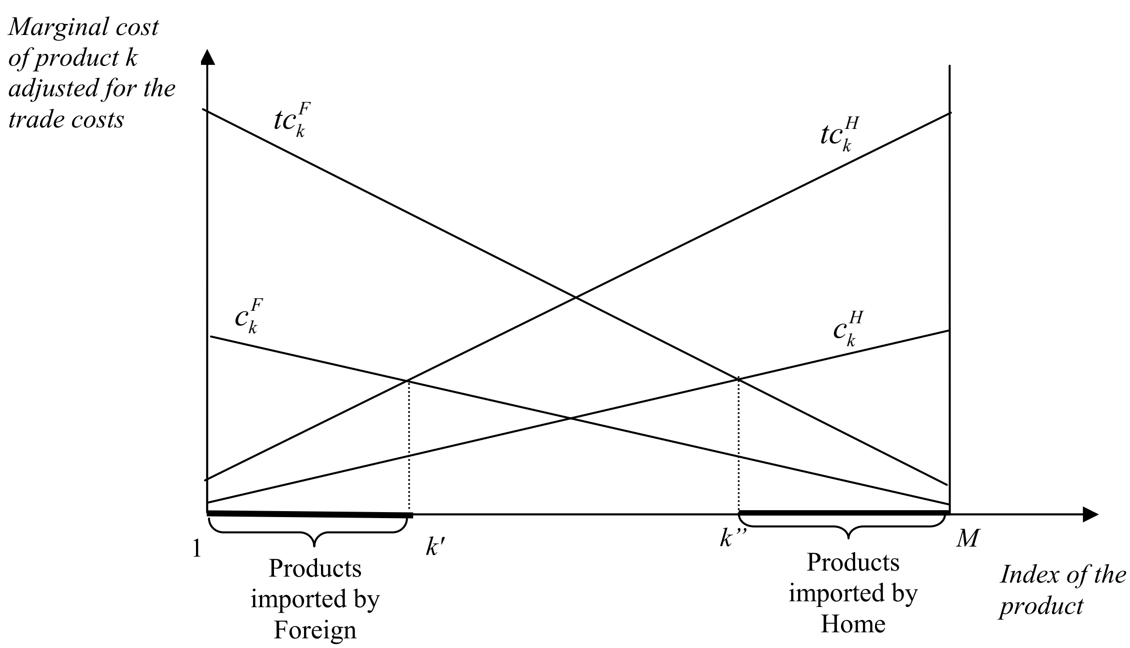

Figure A2. Trade Equilibrium on the Product Level

In Figure 2, Home imports products indexed from $k"$ to M, while Foreign imports products indexed from 1 to k'. Ignoring the divisibility problem we can find the number of imported products for each country by finding k' and k'':

(A15) $k' = \frac{M+1}{t+1}$ $\qquad$ $k'' = \frac{t(M+1)}{1+t}$.

Due to the reverse order of marginal costs and the same transportation cost, the number of products imported by Home, N^H, is equal to the number of products, imported by Foreign, N^F:

(A16) $N^H = N^F = \frac{M+1}{1+t}$.

Now let us concentrate on deriving the equilibrium values for Home. The equilibrium values for Foreign can then be easily found by the corresponding adjustment of notation.

We have established that, while choosing among varieties of product $k = 1, 2, ..., M$, consumers are choosing from the varieties produced with the same technology. In particular, if we consider product k, such that $k \leq k''$, the partial equilibrium for this product will be the same as in the case of the closed economy, and it can be characterized by (16) – (21) if we adjust the notation properly. Specifically, the notation of product-specific variables is adjusted by index k, and notation of the country-specific variables is adjusted by superscripts H. On the contrary, if the product's index is greater or equal to k'', it will be imported from Foreign. The solution for the partial equilibrium of the imported products will have the same functional form as the solution for the partial product of domestically produced and consumed products. The only difference between these two solutions will be in determining the marginal cost of product k. For the imported products the marginal cost equals the corresponding marginal cost of Foreign's producers adjusted for trade costs and barriers:

(A17) $c_k = tc\left(M - k + 1\right)$.

Consequently, in all formulas (16)–(21), we should substitute ck with $c\left(M - k + 1\right)$.

Note that the set of traded varieties differs in this model compared to the Krugman (1980) model, in which all produced varieties are traded. This difference is obtained due to assuming the fixed cost of adjusting to markets rather than the fixed cost of production. It can be shown that an introduction of the fixed cost of adjusting to markets reduces the set of traded varieties also in the Dixit-Stiglitz framework (e.g., Romer 1994, Venables 1994). To normalize the comparison of the models in section "Empirical Implications", we standardize all of them by assuming the fixed cost of adjusting to the market rather than the fixed cost of production.

C. Comparative Statics

Note that the number of imported products depends neither on the aggregate, nor on the individual endowment of labor. It is also not affected by the share of consumption of all differentiated products, μ. Thus, changes in these parameters will not affect the trade equilibrium at the product level.

As we established, the partial equilibria on each country-product-specific market will be identical to the equilibrium derived for the closed economy if properly adjusted for notation. Consequently, on the variety level (within each product category), the comparative statics with respect to L, z, and μ will have the same signs, magnitudes and interpretation as the corresponding comparative statics for the closed economy.

Table 1: Estimating the own-price elasticity of demand