Income Distribution, Product Quality, and International Trade

Citation

Fajgelbaum, Pablo, Gene M. Grossman, and Elhanan Helpman. 2011. "Income Distribution, Product Quality, and International Trade." *Journal of Political Economy* 119 (4) (August): 721–765. doi:10.1086/662628.

Published Version

10.1086/662628

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:25586659>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

Income Distribution, Product Quality, and International Trade

Pablo Fajgelbaum

University of California, Los Angeles

Gene M. Grossman

Princeton University

Elhanan Helpman

Harvard University and Canadian Institute for Advanced Research

We develop a framework for studying trade in horizontally and vertically differentiated products. In our model, consumers with heterogeneous incomes and tastes purchase a homogeneous good and make a discrete choice of quality and variety of a differentiated product. The distribution of preferences generates a nested-logit demand structure such that the fraction of consumers who buy a higher-quality product rises with income. The model features a home-market effect that helps to explain why richer countries export higher-quality goods. It provides a tractable tool for studying the welfare consequences of trade and trade policy for different income groups in an economy.

We are grateful to Aykut Ahlatcıoğlu and Gabriel Kreindler for research assistance and to Juan Carlos Hallak, Jonathan Vogel, Yoram Weiss, and countless seminar participants for helpful comments and suggestions. Grossman and Helpman thank the National Science Foundation and Fajgelbaum thanks the International Economics Section at Princeton University for financial support. The paper was substantially revised while Grossman was a visiting research fellow in the Development Economics Vice Presidency at the World Bank. He thanks the World Bank for support and the Trade and Integration Team (DECTI) for its hospitality. Any opinions, findings, and conclusions or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the National Science Foundation, the World Bank Group, or any other organization.

[*Journal of Political Economy*, 2011, vol. 119, no. 4]

© 2011 by The University of Chicago. All rights reserved. 0022-3808/2011/11904-0004\$10.00

I. Introduction

International trade flows reveal systematic patterns of vertical specialization. When rich and poor countries export goods in the same product category, the richer countries sell goods with higher unit values (Schott 2004; Hummels and Klenow 2005; Hallak and Schott 2011). This suggests a positive association between per capita income and the quality of exports. Also, when a country imports goods in a product category from several sources, the higher-quality goods are imported disproportionately from the higher-income countries (Hallak 2006). Since wealthier households typically consume goods of higher quality (Bils and Klenow 2001; Broda and Romalis 2009), the pattern of vertical specialization has important implications for the distributional consequences of world trade.

In this paper, we propose a new analytic framework for studying trade in vertically differentiated products. Our approach features nonhomothetic preferences over goods of different quality, as is suggested by the observed consumption patterns. It allows trade patterns to depend on the distributions of income in trading partners, and it implies that the welfare consequences of trade vary across income groups in any country. It predicts that richer countries will be net exporters of higher-quality goods and net importers of lower-quality goods under reasonable assumptions about levels and distributions of national income. Our model implies that, in many circumstances, trade liberalization benefits the poorer households in wealthy countries and the richer households in poor countries.

We provide a demand-based explanation for the pattern of trade in goods of different quality. In this respect, our approach is reminiscent of that used by Linder (1961), who hypothesized that firms in any country produce goods suited to the predominant tastes of their local consumers and sell them worldwide to others who share these tastes.¹ Our approach complements a flourishing literature that highlights various supply-side determinants of trade in vertically differentiated goods. In Markusen (1986) or Bergstrand (1990), for example, the country with higher per capita income exports the luxury good, because that good happens to be capital intensive. Similarly, in Flam and Helpman (1987), Stokey (1991), Murphy and Shleifer (1997), and Matsuyama (2000), the

¹ Mitra and Trindade (2005) also offer a demand-based explanation for the pattern of trade in a model with nonhomothetic preferences. However, their model implies that countries will export goods that are little demanded at home in the absence of any supply-side differences between them. See also the studies by Foellmi, Hepenstrick, and Zweimüller (2010), who model trade in horizontally differentiated goods with discrete choice and nonhomothetic preferences, and Auer (2010), who models heterogeneity in consumer tastes and explores how international differences in tastes affect the pattern of trade in differentiated varieties.

pattern of trade follows from an assumption that richer countries have relative technological superiority in producing higher-quality goods.² More recently, Baldwin and Harrigan (2011) and Johnson (forthcoming) have incorporated vertically differentiated products into trade models with heterogeneous firms. They seek to explain the observation that more productive firms export higher-priced (and therefore, presumably, higher-quality) products by referencing the relatively greater incentive that such firms have to undertake quality-enhancing investments. Their approach would generate a supply-side explanation for the observed pattern of trade if richer countries are home to a disproportionate share of the high-productivity firms.³

The demand structure that we exploit has strong empirical roots. We assume that individuals consume varying quantities of a homogeneous good and a discrete choice of a product that is both horizontally and vertically differentiated. Consumers choose among different quality options for the good and from a set of distinctive products at each quality level that have idiosyncratic appeal. The assumed form of the utility function and the distribution of tastes are such that the system of aggregate demands exhibits a nested-logit structure. We draw on the theory of such demands that has been developed by McFadden (1978), Anderson, De Palma, and Thisse (1992), and Verboven (1996*b*), among others.⁴

We posit a utility function that features complementarity between the quantity of the homogeneous good and the quality of the differentiated product. This property of the assumed preferences—shared also by earlier work on vertical competition by Gabszewicz and Thisse (1979, 1980) and Shaked and Sutton (1982, 1983)—implies that the marginal value of quality is higher for households that have greater income. We add to their specification an idiosyncratic taste component that captures a consumer's personal valuation of the attributes of each of the differentiated products. With this addition, a wealthy consumer may fancy a particular low-quality variety whereas a poorer consumer favors one of the high-quality products. In the aggregate, the fraction of consumers who buy a high-quality product rises with income. This behavior gen-

² Also, in a calibration exercise using a Ricardian framework à la Eaton and Kortum (2002) with two industries and many goods, Fieler (2011) finds that the industry with a higher income elasticity of demand also has the greater spread in its productivity draws. As she shows, this gives the country with the higher technology level a comparative advantage in luxury goods.

³ Hallak and Sivadasan (2009), Manova and Zhang (2011), and Kugler and Verhoogen (forthcoming) provide empirical evidence on the relationship between firm size, firm productivity, and export unit values.

⁴ The nested-logit demand structure has been applied to international trade by Goldberg (1995) and Verboven (1996*a*) and more recently by Verhoogen (2008) and Khandelwal (2010). The latter two include a vertical dimension of product differentiation in their discussion, but their focus is very different from ours.

erates heterogeneity in income elasticities of demand across different goods. Such heterogeneity has proven useful in explaining bilateral trade flows in work by Hunter and Markusen (1988), Bergstrand (1990), Hunter (1991), and Fieler (2011). Moreover, the horizontal product differentiation validates a market structure of monopolistic competition, which simplifies the analysis greatly in comparison to the earlier literature with oligopolistic interactions.

The nonhomotheticities in demand forge a link between the shape of a country's income distribution and the pattern and intensity of its trade in vertically differentiated products. We draw out some of the implications in our analysis, much as Flam and Helpman (1987), Matsuyama (2000), and Mitra and Trindade (2005). Dalgin, Trindade, and Mitra (2008) and Choi, Hummels, and Xiang (2009) show that such links between income distribution and trade patterns are important in reality.⁵

In our model, patterns of aggregate demand translate into patterns of specialization and trade via "home-market effects." In a standard competitive model with constant returns to scale, a country exports those goods for which there is little local demand and imports goods that domestic consumers especially covet. But, as Krugman (1980) argued, when transportation is costly and production features economies of scale, a large home market lends a competitive advantage to local firms. Therefore, countries tend to export the increasing-returns goods that are in great domestic demand.⁶ In our model, the demand differences are not a matter of exogenous cross-country variations in tastes but rather derive from differences in income distribution in the face of nonhomothetic demands. We outline conditions under which a richer country, or one with a more dispersed distribution of income, has a larger home demand for high-quality goods and a smaller home demand for low-quality goods. Under such conditions, more firms enter to produce high-quality goods in the richer (or more unequal) country, whereas the opposite is true of firms producing low-quality products. Firms at a given quality charge the same ex-factory prices, so the number of producers predicts the direction of trade. Thus, our model can explain, for example, why Germany traditionally has exported high-quality cars to Korea while importing low-quality cars from there.

⁵ Choi et al. (2009) show that country pairs that share more similar income distributions also exhibit more similar distributions of import prices. Dalgin et al. (2008) find a positive correlation in a sample of developed countries between income dispersion and imports of luxury goods.

⁶ Hanson and Xiang (2004) extend Krugman's argument to a setting with many industries that differ in transport costs and the extent of product differentiation. They provide empirical support for the proposition that larger countries export more in industries with high transport costs and highly differentiated products. See Davis and Weinstein (2003) for further empirical evidence of home-market effects in the pattern of trade.

Our framework provides a tractable and parsimonious tool for studying the distributional implications of changes in transport costs or trade policy. Since different income classes in a country consume different mixes of products, the delocation of firms induced by changes in trading conditions affects the welfare of the various income classes differently. We find, for example, that trade liberalization in a rich country tends to favor the lower-income groups there, who benefit as consumers from an expansion in the range of product offerings at the low-quality level and from a transfer of income from groups that consume greater shares of the high-quality good.

In Section II, we develop our framework in the context of a closed economy. Each consumer buys one unit of some differentiated product and devotes residual income to the homogeneous good. Individuals have idiosyncratic evaluations of the various differentiated products, which also differ in quality. The distribution of taste parameters generates a nested-logit structure of aggregate demands. We combine these demands with a simple supply model that features a single factor of production, costs that vary by quality level, and free entry into the differentiated-products sector. In the monopolistically competitive equilibrium, each firm producing a differentiated product charges a markup over its unit cost that depends on the quality level of its product and a parameter describing the distribution of idiosyncratic tastes. We show in Section III for the case of two quality levels that the autarky equilibrium is unique and that it is characterized by positive numbers of producers of both low-quality and high-quality goods. We proceed to examine how changes in population size and the level and spread of the income distribution affect the numbers of producers at each quality level and the welfare of different income groups.

Section IV introduces international trade between two countries that share similar supply characteristics but differ in their levels and distributions of income. We assume that differentiated products are costly to transport internationally, with per-unit shipping costs that may vary with the quality level. When shipping costs are sufficiently high, as we assume throughout that section, each country produces and trades both low-quality and high-quality goods. We examine how country sizes and income distributions combine to determine the pattern of trade. We also investigate the distributional implications of a decline in trading costs. When such costs decline sufficiently, the production of goods of a given quality must be concentrated in a single country, as we show in Section V. For trading costs close enough to zero, each good is produced in the country that would have the larger home market in a hypothetical, integrated equilibrium. This implies, for example, that if countries are of equal size and the income distribution in one first-order stochastically dominates that in the other, then the richer country produces and

exports the higher-quality goods whereas the poorer country produces and exports the lower-quality goods.

In Section VI, we study commercial policy. Tariffs have no effect on ex-factory prices in our model. The welfare effects of a tariff derive from a composition effect and a redistribution effect. The former captures the change in the relative numbers of high- and low-quality products that results from protection. The latter reflects the transfer of tariff revenues from import purchasers to the average consumer.

Section VII extends the model to include more quality levels and more countries. By doing so, we are able to make contact with the recent empirical literature on the pattern of trade in vertically differentiated products. We assume that countries can be ranked from poorest to richest such that the income distributions in any pair of countries satisfy the monotone likelihood ratio property. We show that when two countries are of similar size and trading costs are high, the richer country has positive net exports of all the highest-quality goods and positive net imports of all the lowest-quality goods in its bilateral trading relationships with poorer countries. When trading costs are small, each quality level is generically produced in a single country and richer countries produce higher-quality goods than poorer countries. In terms of the trade pattern, we find that among countries of similar size, the richer countries export goods of higher quality, which is in keeping with the empirical findings by Schott (2004) and Hummels and Klenow (2005). When trading costs are small, a country imports higher-quality goods from richer trading partners, as Schott (2004), Khandelwal (2010), and Hallak and Schott (2011) find to be true for U.S. imports from various sources.

II. The Model

We develop a model featuring income heterogeneity and nonhomothetic preferences over goods of different quality. We describe the model in this section, characterize its autarky equilibrium in the next, and then move on to international trade in Sections IV–VII below.

Each individual consumes a homogeneous good and his optimal choice from a finite set of differentiated products. Both types of goods are produced with labor alone. The homogeneous good requires one unit of effective labor per unit of output. This good is competitively priced and serves as numeraire. The differentiated products require a fixed input of labor and a constant variable input per unit of output. Monopolistic competition prevails in this industry. We assume that the labor supply is sufficiently large relative to aggregate demand for differentiated products to ensure a positive output of the numeraire good

in any equilibrium. Then competition implies a wage rate for effective labor equal to one.

The economy is populated by a continuum of individuals who are endowed with different amounts of effective labor. This heterogeneity in endowments generates a distribution of income. We denote the income distribution by G , so that $G(y)$ is the fraction of the mass N of individuals with effective labor and wage income less than or equal to y . We assume throughout that every individual has sufficient income to purchase one unit of any variety of the differentiated product, including the most expensive, at the prevailing equilibrium prices.

Each consumer values only one unit of the differentiated product and thus faces a discrete consumption choice. Each buys the good that offers him the highest utility, considering the prices and characteristics of all available products. Varieties are distinguished by their quality level and by other attributes that affect consumers' idiosyncratic valuations. We denote by Q the finite set of available quality levels and index quality by q .

A. Preferences and Demand

Let j index the individual varieties of the differentiated product and let J_q denote the set of varieties with quality q . In this notation, the variety index identifies both the quality of the good and its other attributes, so that if j has quality q , then $j \in J_q$ and $j \notin J_{q'}$ for $q' \neq q$.

Now consider the utility u_j^h that an individual h would attain by consuming z units of the homogeneous good and variety $j \in J_q$ of the differentiated product. We assume that

$$u_j^h = zq + \varepsilon_j^h \quad \text{for } j \in J_q, \quad (1)$$

where ε_j^h is the individual's idiosyncratic evaluation of the particular attributes of variety j . Each individual has a vector of idiosyncratic evaluations, one for each of the available varieties; denote this vector by $\boldsymbol{\varepsilon}^h$. The utility function in (1) features complementarity between the quantity of the homogeneous good and the quality of the differentiated product, much as a "standard" utility function (e.g., Cobb-Douglas or constant elasticity of substitution [CES]) features complementarity between the quantities of the various goods in the individual's consumption basket.⁷ The complementarity between quantity and quality implies a greater marginal valuation of quality for those who consume more of

⁷ This feature of our formulation is quite common in the earlier literature on vertical competition in oligopoly; e.g., Gabszewicz and Thisse (1979, 1980) and Shaked and Sutton (1982, 1983) use the utility function $u = zq$ (in our notation). We add horizontal differentiation in the form of the idiosyncratic taste component, which greatly simplifies the analysis.

the homogeneous good. This property of the utility function generates the nonhomotheticity of aggregate demands in our model.⁸ Meanwhile, the additive utility component captures the other attributes of the product, which the heterogeneous consumers evaluate differently. The horizontal product differentiation validates our assumption of a monopolistically competitive market structure.

We take the ϵ terms to be distributed independently across the population of consumers according to a GEV distribution, which we denote by $G_\epsilon(\epsilon)$. That is,

$$G_\epsilon(\epsilon) = \exp \left[- \sum_{q \in Q} \left(\sum_{j \in J_q} e^{-\epsilon_j / \theta_q} \right)^{\theta_q} \right],$$

with $\theta_q \in (0, 1)$ for all $q \in Q$. This distribution of taste parameters is common in the discrete-choice literature, following Ben-Akiva (1973) and McFadden (1978), because it generates a convenient and empirically estimable system of demands (see also Verboven 1996*b*; Train 2003, chap. 4).

Now consider the optimization problem facing an individual with income y^h and vector of taste parameters ϵ^h . Of course, this individual simply chooses the quality and variety that yield the highest utility among all available options, that is, the q and the $j \in J_q$ that maximize $(y^h - p_j)q + \epsilon_j^h$, where p_j is the price of variety j . Here $y^h - p_j$ represents the amount of (residual) income that the individual devotes to spending on the numeraire good after buying one unit of his most preferred variety of the differentiated product. The calculations in McFadden (1978) and elsewhere imply that, with ϵ distributed according to a GEV, the fraction of individuals with income y who choose variety j with quality q is given by

⁸ In our formulation, the household's indirect utility function is given by $v_j^h = q(I^h - p_j) + \epsilon_j^h$, which is related to but different from the specification commonly used to model consumer demand in the recent industrial organization (IO) literature. There, researchers typically take $v_j^h = \eta^h \xi_j + \mu^h(I^h - p_j) + \epsilon_j^h$, where ξ_j is some attribute of product j (such as quality); η^h and μ^h are parameters that represent the household's valuation of the attribute ξ and the marginal utility of income, respectively; and ϵ_j^h is an idiosyncratic taste component like the one used here. In the literature that follows McFadden (1978), authors have assumed that $\eta^h = \eta$ and $\mu^h = \mu$ for all h and that ϵ_j^h has a generalized extreme value (GEV) distribution. In the literature that follows Berry, Levinsohn, and Pakes (1995), authors have assumed that the taste parameters η^h and μ^h are heterogeneous across households and that ϵ_j^h has a type I extreme value distribution. In either case, the implied (direct) utility function does not allow for interactions between utility from the numeraire good and that from specific attributes of the differentiated product such as "quality." As Nevo (2011) notes, the separability of $I^h - p_j$ and ξ_j is implied by quasi-linear preferences, which may be an appropriate assumption in the partial-equilibrium applications that interest IO economists, but less so for general equilibrium concerns at issue here. In particular, a complementarity between quality of the differentiated product and consumption of the homogeneous good provides a parsimonious explanation for consumption patterns across the income distribution and for the resulting patterns of international trade.

$$\rho_j(y) = \rho_{j|q} \cdot \rho_q(y) \quad \text{for } j \in J_q, \quad (2)$$

where

$$\rho_{j|q} = \frac{e^{-p_j q / \theta_q}}{\sum_{i \in J_q} e^{-q p_i / \theta_q}} \quad (3)$$

is the fraction of consumers who buy variety j among those who purchase a differentiated product with quality q and

$$\rho_q(y) = \frac{[\sum_{j \in J_q} e^{(y - p_j) q / \theta_q}]^{\theta_q}}{\sum_{\omega \in Q} [\sum_{j \in J_\omega} e^{(y - p_j) \omega / \theta_\omega}]^{\theta_\omega}} \quad (4)$$

is the fraction of consumers with income y who opt for a product of this quality. The fraction of individuals who buy a good with quality q varies by income level and with the vector of all product prices, whereas the fraction who buy a particular variety $j \in J_q$ conditional on the choice of quality q depends only on the prices of the goods in this quality segment.

Readers familiar with the empirical literature on discrete-choice modeling will recognize the implied demand system as a *nested logit*, with choice over quality levels (the “nest”) and over horizontally differentiated varieties with a given quality. In that literature, θ_q is known as the dissimilarity parameter; it measures the degree of heterogeneity in preferences over the varieties in the set J_q .⁹ The greater θ_q is, the smaller the correlation between ε_j and $\varepsilon_{j'}$ for j and j' in J_q (see McFadden 1978), and therefore the greater the perceived differences among the various varieties with quality q . It is typically the case that higher-quality products embody richer sets of product characteristics, which expands the scope for horizontal differentiation. If so, the varieties of a lower-quality good will be closer substitutes for one another than the varieties of a higher-quality good. We shall assume that this is the case in what follows; that is, we adopt the following assumption.

ASSUMPTION 1. θ_q is increasing in q .

Variation in the spending pattern across income groups arises solely from variation in the fraction of individuals who purchase the products at different levels of quality q , as reflected by the functions $\rho_q(y)$. It follows that the market share of a good j with quality q varies across income groups according to

$$\frac{1}{\rho_j(y)} \frac{d\rho_j(y)}{dy} = \frac{1}{\rho_q(y)} \frac{d\rho_q(y)}{dy} = q - q_a(y) \quad \text{for } j \in J_q, \quad (5)$$

⁹ Readers familiar with the trade literature will also recognize a similarity between the distribution of preference shocks here and the distribution of productivity shocks in the Ricardian model of Eaton and Kortum (2002). In their work, the productivity shocks are assumed to have a type II extreme value distribution in which θ parameterizes the dissimilarity of productivity levels across goods.

where

$$q_a(y) \equiv \sum_{q \in Q} q \rho_q(y)$$

is the average quality consumed by all individuals with income y . Equation (5) implies that the fraction of individuals who purchase variety j of quality q rises with income if and only if $q > q_a(y)$, that is, if and only if q is above the average quality consumed by individuals in this income group. In particular, in the case of two quality levels, $Q = \{H, L\}$, we have $H > q_a(y) > L$ for all y , so that the fraction of individuals who purchase high-quality products rises with income at all income levels. This is the key property of these nonhomothetic preferences that will guide our analysis of the trade flows.

B. Pricing and Profits

Firms can enter freely into the differentiated-products sector by choosing any quality level $q \in Q$ and by employing a fixed input of f_q units of effective labor to develop a particular variety. A producer uses c_q units of labor per unit of output to produce a good of quality q . Firms set prices to maximize profits taking aggregate price indexes as given. Entry at each quality level proceeds until the next entrant would fail to cover its fixed cost. Let n_q be the number of firms that produce goods of quality q . The demand structure requires the number of varieties in each quality class to be a finite integer, but we will take liberty in treating n_q as if it were a continuous variable to facilitate the exposition.

A firm that produces a variety j of the differentiated product with quality q earns profits of $\pi_j = d_j(p_j - c_q) - f_q$, where $d_j = N\mathbb{E}[\rho_j(y)]$ is the aggregate demand for variety $j \in J_q$ and $\mathbb{E}$ is the expectations operator with respect to the distribution of income, that is, $\mathbb{E}[B(y)] \equiv \int B(y)dG(y)$. Note that demand can be expressed as a function of prices using (2), (3), and (4). If the number of active producers of each quality level is large, then terms in the various sums in (3) and (4) vary only slightly with a firm's own price. We assume that the firm ignores this dependence, as is common in models of monopolistic competition. Then a firm producing any variety j with quality q maximizes profits by setting the price

$$p_q = c_q + \frac{\theta_q}{q} \quad \text{for } q \in Q \quad (6)$$

Evidently, the (absolute) markup over marginal cost differs for goods of different qualities. The markup reflects two properties of the class of goods. First, the higher q is, the greater the marginal utility from consumption of the homogeneous good, because of the complementarity between z and q reflected in (1). A higher marginal utility from

consumption of the homogeneous good makes consumers more sensitive to price differences when choosing among the different brands in J_q . Second, the greater θ_q is, the greater the perceived differences among the various brands with quality q , as we have noted before. This greater degree of product differentiation tends to make demands less sensitive to price changes. These two forces work in opposition as they affect price setting; the markup on high-quality goods will be greater than that on low-quality goods if and only if θ_q/q is increasing in q .

With common prices, the firms that produce different varieties of a good in a given quality segment achieve similar volumes of sales. Let d_q be the total quantity demanded of a typical variety with quality q when all goods are priced according to (6). Then

$$d_q = \frac{N}{n_q} \mathbb{E} \left[\frac{n_q^{\theta_q} \phi_q(y)}{\sum_{\omega \in Q} n_{\omega}^{\theta_{\omega}} \phi_{\omega}(y)} \right] \quad \text{for } q \in Q, \quad (7)$$

where

$$\phi_q(y) \equiv e^{(y - c_q)q - \theta_q}$$

captures the effect of income on demand. The markups of θ_q/q on sales of d_q yield a common profit π_q to all producers of varieties with quality q , where

$$\pi_q \equiv \frac{\theta_q}{q} \frac{N}{n_q} \mathbb{E} \left[\frac{n_q^{\theta_q} \phi_q(y)}{\sum_{\omega \in Q} n_{\omega}^{\theta_{\omega}} \phi_{\omega}(y)} \right] - f_q \quad \text{for } q \in Q. \quad (8)$$

These functions determine the profitability of entry at each quality level. In equilibrium, $n_q > 0$ implies $\pi_q = 0$ whereas $\pi_q \leq 0$ when $n_q = 0$. In the next section, we will use these free-entry conditions to characterize an equilibrium in a closed economy. Once the number of firms producing at a given quality level is known, sales of all varieties of the differentiated product are also determined. Together, the firms selling goods with quality q capture aggregate sales of $n_q d_q$, so that aggregate output of all differentiated products is $\sum_{q \in Q} n_q d_q = N$. This equality reflects the fact that each of the mass N of consumers buys one unit of some product.

The differentiated-products industry employs a total of units of effective labor of $\sum_{q \in Q} n_q (d_q c_q + f_q)$. The difference between aggregate labor supply—which equals N times the mean value of y —and labor use in the differentiated-products industry gives the labor used in producing homogeneous goods. The market for homogeneous products clears by Walras' law. Therefore, once we solve for the number of firms of each type in the differentiated-products industry, the remaining variables determined in the general equilibrium are readily found.

III. Autarky Equilibrium

To characterize an equilibrium in a closed economy, we define x_q as the quantity that a firm producing a brand with quality q must sell in order to break even when it prices according to (6), that is,

$$x_q = \frac{f_q q}{\theta_q} \quad \text{for } q \in Q. \quad (9)$$

Notice that the break-even volume depends only on the magnitude of the fixed cost and the size of the profit margin, as in Krugman (1980). So, (9) will pin down the output per variety for any quality of good that is available in equilibrium.

In an autarky equilibrium, if some positive number of firms produce goods with quality q , the demand per brand must reach the break-even level. Otherwise, no firm producing this quality can profitably enter. In other words, if $n_q > 0$, $d_q = x_q$, whereas $d_q < x_q$ implies $n_q = 0$. In any case, the aggregate output of all differentiated products matches the population size N , or

$$\sum_{q \in Q} n_q x_q = N. \quad (10)$$

We will refer to this equation as the *aggregate demand* condition. It implies, of course, that n_q must be positive for some $q \in Q$.

But notice from (7) that as n_q approaches zero with $n_{q'} > 0$ for some $q' \neq q$, the demand for a typical brand with quality q grows infinitely large. This means that a producer of a brand with quality q will certainly be able to achieve the break-even scale when the number of its competitors offering a similar quality is sufficiently small. In equilibrium, some positive number of firms will be active in every segment of the market.¹⁰

Now that we know that n_q must be positive for all $q \in Q$, market clearing for each brand requires $x_q = d_q$ or

$$x_q = N \mathbb{E} \left[\frac{n_q^{\theta_q - 1} \phi_q(y)}{\sum_{\omega \in Q} n_\omega^{\theta_\omega} \phi_\omega(y)} \right] \quad \text{for } q \in Q. \quad (11)$$

With x_q from (9), this system of equations allows us to solve for the number of varieties at each quality level in an autarkic equilibrium.¹¹

We turn now to the special case with two quality levels, H and L , where

¹⁰ In making this statement, we have ignored the integer constraint. The equilibrium "solution" for some n_q might be a fraction, in which case it might not be profitable for the first "whole" firm to enter in a quality segment. Moreover, we have assumed that many firms compete in order to justify our assumption that firms take price indexes as given. We will not divert attention to these details but instead restrict ourselves to parameters for which our focus on an equilibrium with $n_q > 0$ for all $q \in Q$ is well justified.

¹¹ Note that the weighted sum of x_q from (11) implies (10), which means that only the equations in (11) need be used to solve for the equilibrium numbers of varieties.

$H > L$; we will return to the more general case with an arbitrary number of qualities in Section VII below. With only two quality levels, (11) represents a pair of equations that together determine n_H and n_L . In the Appendix we show that these equations have a unique solution, which is characterized by positive values for n_H and n_L . This establishes the following proposition.

PROPOSITION 1. If $Q = \{H, L\}$, there exists a unique autarky equilibrium. In the autarky equilibrium, $n_H > 0$ and $n_L > 0$.

In the remainder of this section, we describe how the autarky equilibrium reflects the size of the economy and its income distribution.¹² We also show how the model can be used to examine the welfare implications of changes in the economic environment for different income groups. These properties of the model will aid us in understanding the direction and distributional implications of trade in the sections that follow.

The size of the economy is captured by the parameter N . As N increases, the demand in each quality segment grows, given the initial numbers of firms; see (7). Were it the case that the two quality segments offered similarly differentiated products ($\theta_H = \theta_L$), the demand expansion would induce equiproportionate entry by both types of firms. Inasmuch as high-quality products are more dissimilar than low-quality products by assumption 1, there must be proportionately more entry of firms that produce the former goods relative to the latter, that is, $\hat{n}_H > \hat{N} > \hat{n}_L$.¹³ As in other contexts, growth in market size causes an expansion in variety of the more horizontally differentiated products.¹⁴ However, it is readily shown that even the number of low-quality varieties must rise in response to market growth ($\hat{n}_L > 0$).

Now consider an upward shift in the income distribution in the sense of first-order stochastic dominance; that is, at every income level y , the fraction of the population with income less than or equal to y declines. Added income makes consumers more likely to buy a high-quality product across the entire income distribution. Thus, at the initial numbers of firms, demand for high-quality varieties grows and that for low-quality varieties shrinks. This shift in demand induces entry of firms that produce high-quality goods and exit of producers of low-quality products, that is, $\hat{n}_H > 0 > \hat{n}_L$.

Finally, consider an increase in income inequality, as represented by a mean-preserving spread of the distribution $G(\cdot)$. The effect on relative demand is in general ambiguous, as those at the top end of the distri-

¹² The algebra of the comparative statics that we describe here is derived more formally in the Appendix.

¹³ We use a circumflex to denote a proportional increase, i.e., $\hat{Z} = dZ/Z$.

¹⁴ See Hanson and Xiang (2004) and Epifani and Gancia (2006) for similar results in a different context.

bution collectively buy more of the high-quality goods whereas those at the bottom end do just the opposite. However, if the initial equilibrium is such that a majority of every income class purchases low-quality products, the relative demand for high-quality goods is a concave function of y for given n_H and n_L . Then, a mean-preserving spread in the distribution of y causes the relative demand for high-quality goods to expand, inducing entry of producers of these varieties and exit of producers of low-quality products. A spread in income distribution in a poor economy (one in which $\rho_L(y) > \rho_H(y)$ for all y) induces a shift in the composition of firms toward producers of high-quality products.¹⁵

We can readily examine the implications of these shifts for the welfare of different income groups. As McFadden (1978) has shown, the expected welfare among those with income y increases with

$$v(y) \equiv n_H^{\theta_H} \phi_H(y) + n_L^{\theta_L} \phi_L(y). \quad (12)$$

As market conditions change,

$$\hat{v}(y) = \rho_H(y) \theta_H \hat{n}_H + \rho_L(y) \theta_L \hat{n}_L.$$

In words, the change in average welfare at income y weights the changes in the number of products in each quality class by the probability that a consumer with income y purchases a good of that class times the degree of horizontal differentiation (dissimilarity) within the class.

From the aggregate demand condition (10), $\rho_H \hat{n}_H + \rho_L \hat{n}_L = \hat{N}$, where $\rho_q = n_q d_q / N = n_q x_q / N$ is the fraction of the overall population that purchases a good of quality q . Using this equation, we can write

$$\hat{v}(y) = \left[\theta_L \frac{\rho_L(y)}{\rho_L} + \theta_H \frac{\rho_H(y)}{\rho_H} \right] \hat{N} + \rho_H \rho_L \left[\theta_H \frac{\rho_H(y)}{\rho_H} - \theta_L \frac{\rho_L(y)}{\rho_L} \right] (\hat{n}_H - \hat{n}_L). \quad (13)$$

The first term in the expression for $\hat{v}(y)$ is a pure *scale effect*. With the relative number of high-quality and low-quality products held constant, an expansion of scale benefits consumers at all income levels, because it increases the number of varieties and therefore increases the likelihood that an individual will find one to his liking. The second term is a pure *composition effect*. For a given scale, an increase in the relative number of high-quality products benefits those who are more likely than average to consume such a product and harms those who are more likely than average to consume a low-quality product. An increase in the variety of high-quality products relative to the variety of low-quality products is more likely to benefit a given income group the more dissimilar the brands of high-quality products and the more similar the brands of low-quality products.

Now let us examine the distributional implications of the market changes we described above. An increase in population size generates

¹⁵ See the Appendix for the details.

a scale effect that benefits all income groups and a composition effect that especially benefits the wealthy (since growth in market size generates an increase in the relative number of high-quality varieties when $\theta_H > \theta_L$). The richest consumers in the economy, who have income $y_{\max}$, are more likely to purchase the high-quality good than the average consumer and are less likely to purchase the low-quality good, which implies that $\rho_H(y_{\max})/\rho_H > 1 > \rho_L(y_{\max})/\rho_L$. By (13), these wealthy individuals must gain on average from population growth. The poorest consumers—who are more likely than average to purchase the low-quality good and less likely than average to consume the high-quality good—will also benefit from the expansion in variety in both quality segments, but their gain will be more modest.

An upward shift in the income distribution (or a spread of the distribution in a poor economy) generates a shift in the composition of differentiated products toward high-quality goods, without changing the output-weighted number of products. With $\theta_H > \theta_L$, the associated composition effect must benefit the members of the highest-income group (on average). As for the poorest consumers, they too may benefit if high-quality goods are substantially more dissimilar than low-quality goods but will lose if θ_H and θ_L are quite close in size. Although the low-income individuals are more likely to consume a low-quality product, the contraction of variety in this market segment will not hurt them so much if these goods are relatively similar to one another; meanwhile, the expansion in the variety of high-quality products can be quite advantageous even to these consumers (on average) if the idiosyncratic tastes for the various high-quality products are little correlated. If there are income groups that lose from a change in the composition of products that favors high-quality goods, it will be all groups with income less than or equal to some critical value.

IV. Trade with Diversified Production

In this section, we introduce international trade. We assume for the time being that there are two countries that differ in size and in their distributions of effective labor. We do not allow for any supply-side determinants of the trade pattern—as would arise from comparative cost advantages—in order to focus more sharply on those that derive from differences in income in the face of nonhomothetic preferences. We designate the countries as $\mathcal{R}$ and $\mathcal{P}$ to suggest “rich” and “poor,” although we do not insist on any particular relationship between their sizes or their income distributions except in some special cases. In Section VII, we will extend the analysis to many countries in order to make contact with the empirical evidence cited in the introduction.

We assume throughout that both (or all) countries have sufficient

supplies of effective labor relative to the equilibrium labor demands by their producers of differentiated products so that some labor in each country is used to produce the homogeneous, numeraire good. This ensures that the wage of a unit of effective labor is equal to one in both (all) countries.

We assume that the differentiated products are costly to trade.¹⁶ In particular, it takes τ_q units of effective labor to ship one unit of a variety with quality q from one country to another.¹⁷ As τ_q grows large, national outputs converge on those of the autarky equilibria. In such a setting, as we now know, both countries produce goods in both quality segments. We will find that such incomplete specialization characterizes the trade equilibrium whenever trading costs are sufficiently high. These are the circumstances that we consider now, whereas in the next section we will study equilibria in which each quality level is produced in only one country, as happens almost surely when shipping costs are small.¹⁸

Shipping requirements raise the cost of serving foreign consumers relative to domestic consumers. For a good with quality q , the marginal cost of a delivered export unit is $c_q + \tau_q$, whereas local consumers can be supplied at a cost of c_q . The arguments from Section II.B now imply that a firm producing a variety with quality q maximizes profits by charging foreign consumers the price $c_q + \tau_q + \theta_q/q$, whereas domestic consumers are charged the lower price $c_q + \theta_q/q$ (see [6]). In other words, profit margins are θ_q/q for all sales, as firms fully pass on their shipping costs to their foreign customers.¹⁹

Demands for domestic goods of quality q in country k reflect the

¹⁶ Davis (1998) has shown that a home-market effect may not exist if differentiated products and homogeneous goods bear similar trading costs. But Amiti (1998) and Hanson and Xiang (2004) demonstrate that the home-market effect requires only that transport costs differ across sectors.

¹⁷ Unlike models of monopolistic competition that rely on the CES demand system, our model easily allows for transport costs that are incurred on a per-unit basis. Although we consider this to be more realistic than the popular assumption of "iceberg" transport costs, there is no meaningful difference between the two in our model. In either case, firms price their exports at a fixed markup per unit over the delivered cost.

¹⁸ It is common in trade models featuring a home-market effect that production is geographically dispersed when transport costs are large, but each good is produced in only one location when transport costs are small; see, e.g., Krugman (1991*a*, 1991*b*) and Rossi-Hansberg (2005). The implications for the trade pattern are somewhat different in these two regimes—especially when the number of quality levels and countries is large—which explains why we consider both possibilities.

¹⁹ If transportation costs instead took the iceberg form—such that the delivery of one unit of a good of quality q to an export market required that $\gamma_q > 1$ units be shipped—then the profit-maximizing price for export sales would instead be $c_q\gamma_q + \theta_q/q$. The analysis would proceed exactly as in what follows. The ability of our model to accommodate both per-unit and iceberg shipping costs reflects the fact that, under the nested-logit demand system, the optimal gap between price and unit cost is constant and independent of the delivered cost. In contrast, models with CES demands imply an absolute markup that is proportional to the delivered cost, in which case per-unit and proportional shipping costs have different effects on export prices.

prices of these goods, the prices of competing import goods, and the numbers of local and imported varieties at each quality level. Let d_q^k represent the aggregate demand by domestic consumers for a typical good of quality q produced in country k when all goods are priced optimally. Then (4) implies

$$d_q^k = \frac{N^k}{\tilde{n}_q^k} \mathbb{E}^k \left[\frac{(\tilde{n}_q^k)^{\theta_q} \phi_q(y)}{\sum_{\omega \in Q} (\tilde{n}_\omega^k)^{\theta_\omega} \phi_\omega(y)} \right], \quad q = H, L \text{ and } k = \mathcal{R}, \mathcal{P}, \quad (14)$$

where

$$\tilde{n}_q^k = n_q^k + \lambda_q n_q^l, \quad l \neq k, \quad \lambda_q \equiv e^{-\tau_q q / \theta_q},$$

n_q^k is the number of varieties of quality q produced in country k , N^k is the population in country k , and $\mathbb{E}^k$ is the expectation with respect to the income distribution there. Notice the similarity between (14) and (7). Now, domestic brands share the market with both domestic and foreign rivals, but inasmuch as imports of a given quality bear a higher price because of shipping costs, the foreign varieties are less effective competitors. For local firms, domestic demand is the same as it would be in autarky with $\tilde{n}_q^k$ local competitors producing quality q . The foreign firms are discounted in this measurement of “effective competitors” by an amount $\lambda_q \in (0, 1)$ that reflects the trading cost for goods of quality q as well as the quality and dissimilarity of these products. Also, (4) implies that per capita demand for an imported variety of quality q in country k equals $\lambda_q d_q^k / N^k$; that is, it is a fraction λ_q of the per capita demand faced by a local firm.

A firm producing a variety with quality q in either country earns profits per sale of θ_q/q , considering the fixed absolute markup it charges over delivered cost. In order to break even, such a firm, no matter where it is located, must make sales totaling $x_q = f_q q / \theta_q$ units, as per (9). In an equilibrium with producers of both qualities in both countries, we must have

$$x_q = d_q^k + \lambda_q d_q^l, \quad k, l = \mathcal{R}, \mathcal{P}, \quad l \neq k, \quad q = H, L.$$

The right-hand side of this equation represents total sales by a firm located in country k , comprising domestic sales and exports sales, where the latter are a fraction λ_q of what a local producer in country l makes of domestic sales. For these equations to hold for both $k = \mathcal{P}$ and $k = \mathcal{R}$, it must be that $d_q^{\mathcal{R}} = d_q^{\mathcal{P}}$ for $q = H, L$; that is, firms in each country must achieve the same volume of domestic sales. Since the size and the income distributions in the two countries may differ, the equality must be achieved by adjustment in the numbers of effective firms in each market. In particular, the equality between the required volume of total sales and the total demand faced by producers in market k implies $x_q = d_q^k(1 + \lambda_q)$ or

$$N^k \mathbb{E}^k \left[\frac{(\tilde{n}_q^k)^{\theta_q - 1} \phi_q(y)}{\sum_{\omega \in Q} (\tilde{n}_\omega^k)^{\theta_\omega} \phi_\omega(y)} \right] = \frac{1}{1 + \lambda_q} \frac{f_q q}{\theta_q}, \quad q = H, L, k = \mathcal{R}, \mathcal{P}. \quad (15)$$

The equations in (15) provide four independent relationships, two for country $\mathcal{R}$ that jointly determine $\tilde{n}_L^{\mathcal{R}}$ and $\tilde{n}_H^{\mathcal{R}}$ and two for country $\mathcal{P}$ that jointly determine $\tilde{n}_L^{\mathcal{P}}$ and $\tilde{n}_H^{\mathcal{P}}$. These equations have exactly the same form as those in (11) that describe the autarky equilibrium, except that x_q for the closed economy is replaced by $x_q/(1 + \lambda_q)$ for the open economy. The proof of proposition 1 guarantees that (15) has a unique solution with a positive number of effective firms of each type in each country.

It is not enough, of course, that the number of effective firms in each country be positive for the solutions to (15) to represent a legitimate trade equilibrium. We require in addition that the actual number of varieties in each country be positive. Given the values of $\tilde{n}_L^{\mathcal{R}}$, $\tilde{n}_H^{\mathcal{R}}$, $\tilde{n}_L^{\mathcal{P}}$, and $\tilde{n}_H^{\mathcal{P}}$ that result from the solution of (15), we can solve for n_L^k and n_H^k using $\tilde{n}_q^k = n_q^k + \lambda_q n_q^l$. This gives

$$n_q^k = \frac{\tilde{n}_q^k - \lambda_q \tilde{n}_q^l}{1 - (\lambda_q)^2}, \quad k, l = \mathcal{R}, \mathcal{P}, l \neq k, q = H, L. \quad (16)$$

A positive solution for n_q^k for all q and k requires

$$\frac{1}{\lambda_q} > \frac{\tilde{n}_q^k}{\tilde{n}_q^l} > \lambda_q, \quad k, l = \mathcal{R}, \mathcal{P}, l \neq k, q = H, L,$$

which is always satisfied when λ_q is close to zero but rarely satisfied when λ_q is close to one.²⁰ This justifies our claim that a trade equilibrium with incomplete specialization always exists when transport costs are sufficiently high but fails to exist (generically) when transport costs are low.

The trade equilibrium with incomplete specialization features intraindustry trade at each quality level. Some consumers in $\mathcal{R}$ opt for preferred varieties of the high-quality good produced in $\mathcal{P}$, despite their higher price that includes a charge for shipping. Similarly, some consumers in $\mathcal{R}$ choose to import a favorite foreign variety of the low-quality good. Consumers in $\mathcal{P}$ will likewise import high-quality and low-quality goods produced in $\mathcal{R}$.

In fact, we know that the export sales by a typical producer of quality q are the same in both locations. Therefore, country $\mathcal{R}$ exports more of goods of quality q to $\mathcal{P}$ than it imports of that quality if and only if country $\mathcal{R}$ has more firms producing goods of quality q than country $\mathcal{P}$ does. But, from (16), $n_q^{\mathcal{R}} > n_q^{\mathcal{P}}$ if and only if $\tilde{n}_q^{\mathcal{R}} > \tilde{n}_q^{\mathcal{P}}$. Therefore, we can identify the equilibrium trade balance in each quality segment by

²⁰ When λ_q is close to one, the pair of inequalities can be satisfied only if $\tilde{n}_q^{\mathcal{R}} \approx \tilde{n}_q^{\mathcal{P}}$, which happens only under exceptional circumstances. For example, $N^{\mathcal{R}} = N^{\mathcal{P}}$ and $G_s^{\mathcal{R}} = G_s^{\mathcal{P}}$ imply $\tilde{n}_q^{\mathcal{R}} = \tilde{n}_q^{\mathcal{P}}$.

comparing the effective number of sellers of that quality in the two countries.²¹

The pair of equations that determine $\tilde{n}_q^{\mathcal{R}}$ and $\tilde{n}_q^{\mathcal{P}}$ are identical to those that determine the autarky numbers of producers of quality q in $\mathcal{R}$ and $\mathcal{P}$, except that x_q in the latter is replaced by $x_q/(1 + \lambda_q)$ in the former. Therefore, we can use the comparative statics of the autarky equilibrium to identify the sectoral imbalances of the trade equilibrium with incomplete specialization. For example, suppose that the countries have the same distributions of income ($G^{\mathcal{R}} = G^{\mathcal{P}}$) but country $\mathcal{R}$ is larger than country $\mathcal{P}$ (i.e., $N^{\mathcal{R}} > N^{\mathcal{P}}$). We have seen that the larger country has in autarky a greater relative abundance of firms that produce high-quality goods, because market growth generates biased entry in favor of the more horizontally differentiated products. It follows that the larger country must have absolutely more producers of high-quality goods in autarky, whereas it may support fewer (or more) producers of low-quality products. These comparisons carry over to the numbers of effective sellers in a trade equilibrium with incomplete specialization. That is, $\tilde{n}_H^{\mathcal{R}} > \tilde{n}_H^{\mathcal{P}}$, whereas the comparison of effective numbers of producers of low-quality goods can run in either direction. In such circumstances, the larger country is a net exporter of high-quality goods but may be a net exporter or a net importer of low-quality products.

Now suppose that the two countries are identical in size ($N^{\mathcal{R}} = N^{\mathcal{P}}$) but the income distribution in the richer $\mathcal{R}$ first-order stochastically dominates that in poorer $\mathcal{P}$. Then, in autarky, the rich country has more firms producing high-quality goods and the poor country has more firms producing low-quality goods. These comparisons carry over to the effective numbers of firms in the trade equilibrium with incomplete specialization, so that $\tilde{n}_H^{\mathcal{R}} > \tilde{n}_H^{\mathcal{P}}$ and $\tilde{n}_L^{\mathcal{P}} > \tilde{n}_L^{\mathcal{R}}$. It follows that the rich country $\mathcal{R}$ is a net exporter of high-quality goods and a net importer of low-quality goods.

Finally, suppose that a majority of consumers at every income level in both countries purchase low-quality goods. Let the countries be of similar size and with similar mean income, but suppose that the income distribution in $\mathcal{R}$ is more spread than that in $\mathcal{P}$. As we have seen before, $\mathcal{R}$ has more producers of high-quality goods and fewer producers of low-quality goods than $\mathcal{P}$ does in autarky. With costly trade, $\mathcal{R}$ becomes

²¹ Net exports from $\mathcal{R}$ to $\mathcal{P}$ of goods of quality q are given by

$$\begin{aligned} \lambda_q d_q^{\mathcal{P}} n_q^{\mathcal{R}} - \lambda_q d_q^{\mathcal{R}} n_q^{\mathcal{P}} &= \frac{\lambda_q}{1 + \lambda_q} x_q (n_q^{\mathcal{R}} - n_q^{\mathcal{P}}) \\ &= \frac{\lambda_q}{1 - (\lambda_q)^2} \frac{f_q q}{\theta_q} (\tilde{n}_q^{\mathcal{R}} - \tilde{n}_q^{\mathcal{P}}). \end{aligned}$$

a net exporter of high-quality goods and a net importer of low-quality goods.

We summarize our findings about the pattern of trade in the following proposition.

PROPOSITION 2. If trade costs are sufficiently high, there exists a unique trade equilibrium in which each country produces both high- and low-quality differentiated products. In this equilibrium, (i) if $N^{\mathcal{R}} > N^{\mathcal{P}}$ and $G^{\mathcal{R}}(y) = G^{\mathcal{P}}(y)$ for all y , then $\mathcal{R}$ exports on net the high-quality goods but may export or import on net the low-quality goods; (ii) if $N^{\mathcal{R}} = N^{\mathcal{P}}$ and $G^{\mathcal{R}}(y) < G^{\mathcal{P}}(y)$ for all y , then $\mathcal{R}$ exports on net the high-quality goods and imports on net the low-quality goods; (iii) if $N^{\mathcal{R}} = N^{\mathcal{P}}$, $\rho_L(y) > \rho_H(y)$ for all income groups in $\mathcal{R}$ and $\mathcal{P}$, and $G^{\mathcal{R}}(\cdot)$ is a mean-preserving spread of $G^{\mathcal{P}}(\cdot)$, then $\mathcal{R}$ exports on net the high-quality goods and imports on net the low-quality goods.

Proposition 2 can be understood in terms of the “home-market effect” described by Krugman (1980). Take, for example, the case in which the countries are of similar size but the income distribution in $\mathcal{R}$ first-order stochastically dominates that in $\mathcal{P}$. The greater income in $\mathcal{R}$ compared to $\mathcal{P}$ provides this country with a larger home market for high-quality goods. If the same numbers of producers of high-quality goods were to enter in both countries, those in $\mathcal{R}$ would earn greater profits than those in $\mathcal{P}$, thanks to their ability to serve more consumers with sales that do not bear shipping costs. In order that producers of high-quality goods in both countries break even, there must be greater entry of such producers in the rich country, so that their finer division of the market offsets their local-market advantage. The same is true in the market for low-quality goods, where producers in $\mathcal{P}$ enjoy an advantage due to their closer proximity to the larger market. Access to a large home market affords a competitive advantage that induces entry and ultimately dictates the pattern of trade.

We turn now to the effects of a reduction in trade costs, focusing particularly on the distributional consequences. For concreteness, consider first a decline in the cost of transporting high-quality goods.²² A fall in τ_H induces an increase in λ_H . It is clear from (15) that such an increase in λ_H generates the same outcomes as a reduction in the fixed cost of entry for producers of high-quality products, f_H . As τ_H falls, profitability rises for firms that produce high-quality varieties. The number of effective producers of such varieties rises in each country. This expansion in $\tilde{n}_H^{\mathcal{R}}$ and $\tilde{n}_H^{\mathcal{P}}$ reduces demand for low-quality goods in each country, and so there is effective exit from this market segment. In the new trade equilibrium, there is a more effective variety of high-quality goods in each country and a less effective variety of low-quality products.

²² The details of the algebra are provided in the Appendix.

The effects of a reduction in the cost of transporting low-quality goods are analogous.²³

What are the welfare implications of these induced changes in the effective numbers of varieties? In a world with costly trade, the average welfare of those with income y in country k increases with

$$v^k(y) = (\tilde{n}_H^k)^{\theta_H} \phi_H(y) + (\tilde{n}_L^k)^{\theta_L} \phi_L(y) \quad \text{for } k = \mathcal{R}, \mathcal{P}.$$

Welfare of individuals in country k depends on the effective numbers of varieties available there, with foreign brands carrying less weight than domestic brands because of their higher prices. Differentiating the expression for $v^k(y)$ and rearranging terms, we can derive an expression for the change in average welfare of an income group analogous to (13), namely

$$\begin{aligned} \hat{v}^k(y) = & \left[\theta_L \frac{\rho_L^k(y)}{\rho_L^k} + \theta_H \frac{\rho_H^k(y)}{\rho_H^k} \right] [\rho_H^k(1 + \widehat{\lambda}_H) + \rho_L^k(1 + \widehat{\lambda}_L)] \\ & + \rho_H^k \rho_L^k \left[\theta_H \frac{\rho_H^k(y)}{\rho_H^k} - \theta_L \frac{\rho_L^k(y)}{\rho_L^k} \right] (\hat{n}_H^k - \hat{n}_L^k) \quad \text{for } k = \mathcal{R}, \mathcal{P}, \end{aligned} \quad (17)$$

where $\rho_q^k(y)$ is the fraction of consumers in country k with income y who buy a good with quality q and ρ_q^k is the fraction of all consumers in country k who buy a good with quality q . The term in the first line of (17) is a pure *cost-savings effect*, analogous to the scale effect in (13). The term in the second line of (17) is a pure *composition effect*, analogous to the similarly named term in (13). The cost-savings effect benefits consumers at all levels of income; it reflects the fact that, for given relative numbers of effective brands of each quality, a fall in the cost of trade facilitates entry of new producers, which expands the range of available varieties and so the probability that a consumer will find one especially to his liking. The composition effect affects different income classes differently. An expansion in the effective variety of high-quality goods relative to the effective variety of low-quality goods benefits those who are more likely to consume a high-quality product but harms those who are more likely to consume a low-quality product; and, of course, the likelihood of consuming a high-quality good rises with income.

Let us return to the effects of a reduction in trade costs. Consider first a decline in the cost of transporting high-quality goods. As we have seen, such a decline in τ_H expands the effective number of high-quality

²³ It is also evident from (15) that an equiproportionate rise in $1 + \lambda_H$ and $1 + \lambda_L$ has the same impact on the effective number of high- and low-quality products in country k as a similar percentage increase in that country's population would. From our analysis of the autarky equilibrium, we know that the effective number of high-quality products expands more than in proportion to the increase in $1 + \lambda_H$ and $1 + \lambda_L$, whereas the effective number of low-quality products rises less than proportionately.

varieties in each country while contracting the effective number of low-quality varieties. The cost-savings effect benefits all consumers. Since $\theta_H > \theta_L$, the composition effect must benefit the average member of the highest-income group in each economy, but it may harm the average member of the lowest-income group. It follows that a fall in τ_H augments the average welfare of the wealthiest consumers in each country but may bring harm to income groups below some critical level.²⁴

Our analysis also sheds lights on the distribution of the gains from trade. The autarky equilibrium for either country is the solution to (15) with $\lambda_H = \lambda_L = 0$. The effects of trade can be found by integrating the increases in λ_H and λ_L from zero to their actual levels. This generates a cost-savings effect that benefits all consumers. It also generates a composition effect that may benefit some income groups at the expense of others. If the effective number of brands at both quality levels rises as a result of trade, then all consumers must gain. If the effective number of brands of some quality level declines, then income groups that buy this good with a probability that exceeds the economywide average may lose. Although trade may not benefit every income group, it always benefits some such groups.²⁵

We summarize our discussion of the distributional consequences of a reduction in trade costs in the following proposition.

PROPOSITION 3. In a trade equilibrium with incomplete specialization, a decline in the trade cost τ_q raises the effective number of brands of quality q and reduces the effective number of brands of quality q' , $q' \neq q$, in both countries. Any reduction in trade costs must benefit the average member of some income group. If, as a result of a reduction in trade costs, the effective number of high-quality (low-quality) varieties falls in some country, then the highest-income (lowest-income) groups in that country may lose.

In our working paper (Fajgelbaum, Grossman, and Helpman 2009), we present two numerical examples to illustrate the alternative possible welfare outcomes. For one set of parameter values, the average member of every income group in country $\mathcal{R}$ gains from a fall in the cost of shipping high-quality goods. An alternative set of parameter values il-

²⁴ There must be some income groups that gain from a reduction in τ_H . To see this, suppose that the opposite were true. Then, the left-hand side of (15) increases for $q = L$ inasmuch as the numerator increases at every y (because $\tilde{n}_L^k$ falls) and the denominator falls at every y (because average welfare has been assumed to fall). But the right-hand side of (15) is unchanged, which contradicts the requirement for equilibrium in the market for low-quality goods.

Other reductions in trade costs can be analyzed similarly. For example, declines in τ_H and τ_L that increase $1 + \lambda_H$ and $1 + \lambda_L$ by the same proportions must benefit all income groups because such a fall in transport costs results in larger numbers of both low-quality and high-quality products.

²⁵ The proof of this statement follows along lines similar to that used in n. 24.

FIG. 1.—Patterns of specialization

illustrates the possibility of distributional conflict from falling trade costs. In this case, the high-quality products are less dissimilar than in the first example (smaller θ_H) and the low-quality products are more dissimilar (larger θ_L), so that the composition effect is more damaging to the poor. Here, the average member of the median-income group in country $\mathcal{R}$ gains from a reduction in the cost of trading high-quality goods, but the lowest-income group in $\mathcal{R}$ loses.

Figure 1 illustrates the patterns of specialization for different values of λ_H and λ_L for a particular set of parameter values.²⁶ In this case, the countries are similar in size but consumers in $\mathcal{R}$ are richer than their counterparts in $\mathcal{P}$. When both trading costs are reasonably large, so that λ_H and λ_L are small, both countries are incompletely specialized, much as they are in autarky. A sufficient reduction in the cost of trading the

²⁶ In this figure $N^{\mathcal{R}} = N^{\mathcal{P}} = 1,000$, $H = 1.05$, $L = 0.9$, $\theta_H = 0.7$, $\theta_L = 0.5$, $f_H = 5$, $f_L = 1.5$, $c_H = 0.3$, and $c_L = 0.05$. The distributions of income are such that $y - 1$ has a gamma distribution in each country, with a coefficient of variation equal to one in each case. We take the scale parameter in $\mathcal{R}$ to be 6 and that in $\mathcal{P}$ to be 2, so that mean incomes are 7 and 3, respectively.

high-quality goods, with λ_L held fixed at a reasonably low level, generates an equilibrium in which the poor country $\mathcal{P}$ produces only low-quality goods whereas the rich country $\mathcal{R}$ produces both high- and low-quality goods. Similarly, a sufficient reduction in the cost of trading low-quality goods, with λ_H held at a reasonably low level, results in an equilibrium in which $\mathcal{R}$ produces only high-quality goods and $\mathcal{P}$ produces goods in both quality classes. If the cost of transporting both goods is sufficiently small, each class of goods is produced in a single location. We study this latter type of equilibrium in greater depth in the next section.

V. Trade with Specialization

In a trade equilibrium, high transport costs allow firms in each country to enter profitably in both quality segments of the market for differentiated products. Even if there are relatively many foreign producers of a given quality level, local firms can enter to sell to local customers thanks to the protection afforded by the high shipping costs. As we have seen, when λ_H and λ_L are sufficiently close to zero, the trade equilibrium is characterized by incomplete specialization in both countries.

As transport costs fall, it becomes more difficult for firms in a smaller market to overcome the disadvantage of their lesser local demand. Eventually, as λ_q rises toward one, the number of producers of quality q in some country must fall to zero, as is implied by equation (16). For still smaller transport costs, all of the varieties with quality q are produced in a single country. In this section, we study trade equilibria with specialization of this sort. We are particularly interested in the limiting equilibrium, as transport costs approach zero. We will see that this equilibrium is unique and has a readily understood pattern of trade. Before we begin this analysis, however, it will prove useful to have a brief discussion of the integrated equilibrium, when transport costs for both quality levels are literally zero.

A. The Integrated Equilibrium

Suppose that $\tau_H = \tau_L = 0$ so that $\lambda_H = \lambda_L = 1$. With no supply-side sources of comparative advantage, there is nothing in our model to pin down the location of production. Factor-price equalization and zero transport costs mean that the different goods can be produced in various combinations in the two countries, without consequence for any aggregate variables or anyone's welfare level. Although we cannot say anything about the pattern of trade, we can nonetheless characterize the integrated equilibrium in terms of the total numbers of brands of each quality that are produced and the average welfare of the different income groups.

In the absence of trade costs, the effective number of varieties with quality q is the same in both countries, that is, $\tilde{n}_q^R = \tilde{n}_q^P = n_q^R + n_q^P$ for $q = H, L$. We can solve for these aggregate numbers of varieties using the autarky equilibrium conditions for an economy with population $N^R + N^P$ and an income distribution that is the composite of the separate distributions in the two countries. This gives $\bar{n}_H$ and $\bar{n}_L$, the aggregate numbers of high- and low-quality products, respectively, that are produced in the integrated global economy. Armed with these variables, we can calculate aggregate demand in country k for a typical variety with quality q , which we denote by $\bar{d}_q^k$. That is,

$$\bar{d}_q^k = \frac{N^k}{\bar{n}_q} \mathbb{E}^k \left[\frac{(\bar{n}_q)^{\theta_q} \phi_q(y)}{\sum_{\omega \in Q} (\bar{n}_\omega)^{\theta_\omega} \phi_\omega(y)} \right]. \quad (18)$$

The impact of trade with zero transport costs on the welfare of an income group y in country k reflects a scale effect and a composition effect, as before. The scale effect—which arises because the integrated economy has a larger population than either separate economy—works to the benefit of all income groups in both countries. The composition effect benefits the high-income groups in country k if the relative number of high-quality varieties in the integrated equilibrium exceeds the relative number of high-quality varieties in the country's autarky equilibrium. Otherwise, the composition effect benefits the low-income groups in country k . The effect of an opening of trade on the relative numbers of varieties of the different quality levels reflects both the biased nature of growth due to $\theta_H > \theta_L$ and the demand effects of a change in income distribution from one with the properties of the local economy to one with the properties of the global economy.

B. Trade Equilibrium with Small (but Positive) Trade Costs

Now we are ready to characterize the trade equilibrium when transport costs are positive but small. If a firm producing quality q in country k is to break even, it must attain total worldwide sales of $x_q = f_q q / \theta_q$. Each firm's sales comprise its home sales— d_q^k for a firm in country k —and its export sales, which are a fraction λ_q of the domestic sales of a foreign firm. For firms producing quality q to achieve the break-even volume of sales in both countries given the required relationship between the home sales of one and the export sales of the other requires that domestic sales be common to the two countries, that is, $d_q^R = d_q^P$, as we have noted before.

But note that the aggregate demand in country k for a typical variety with quality q approaches $\bar{d}_q^k$ as transport costs go to zero. The aggregate demands of the integrated equilibrium are given by (18) and are uniquely determined by parameters of the world economy. Only excep-

tionally will it happen that $\bar{d}_q^R = \bar{d}_q^P$ for $q = H$ or $q = L$. In other words, only exceptionally will it happen that firms in both countries producing a given quality can break even when transport costs are sufficiently small. Otherwise, goods of a particular quality are produced in a single country, whereas a potential entrant at that quality level in the other country finds insufficient demand (at its optimal price) to cover its fixed costs.²⁷

Which country produces each class of goods when trade costs are small? To answer this question, we look at national demands for products of a given quality in the integrated equilibrium. Suppose, for example, that $\bar{d}_q^R > \bar{d}_q^P$ for products of quality q ; that is, the typical producer of a good with quality q makes greater sales in country $\mathcal{R}$ than in country $\mathcal{P}$. With positive trade costs and optimal pricing, each firm's exports are a fraction of sales by a local producer in the destination market. It follows that when transport costs are sufficiently small, profits per firm for a producer of a brand with quality q in country $\mathcal{R}$ must exceed those for a producer of that quality in country $\mathcal{P}$.²⁸ More generally, all production of goods with quality q takes place in the country with the larger domestic market for goods of that quality in the integrated equilibrium. We summarize in proposition 4.

PROPOSITION 4. Suppose $\bar{d}_q^k > \bar{d}_q^l$ for $q \in \{H, L\}$, $k, l = \mathcal{R}, \mathcal{P}$, and $l \neq k$. Then, for τ_H and τ_L sufficiently close to zero, all goods of quality q are produced in country k .

Let us apply proposition 4 to some special cases that we have considered previously. Suppose, for example, that $G^R = G^P$ and $N^R > N^P$; that is, the countries share the same income distribution but differ in size. Then, by (18), $\bar{d}_H^R > \bar{d}_H^P$ and $\bar{d}_L^R > \bar{d}_L^P$, so the larger country produces and exports all varieties of both the high-quality and low-quality differentiated products. Now suppose instead that $N^R = N^P$ whereas $G^R(\cdot)$ first-order stochastically dominates $G^P(\cdot)$. Then $\bar{d}_H^R > \bar{d}_H^P$ and $\bar{d}_L^R < \bar{d}_L^P$, so the richer country produces all the high-quality goods and the poorer coun-

²⁷ In the literature on the new economic geography, it is common to have diversification for high transport costs but specialization for low transport costs. In these models, if the locations have no inherent productivity or cost advantages, the equilibrium with zero transport costs is indeterminate; see, e.g., Rossi-Hansberg (2005). More generally, Krugman (1991*a*, 1991*b*) was the first to point out an inverted-U-shaped relationship between specialization and transport costs when the various locations have inherent advantages. See Aiginger and Rossi-Hansberg (2006) for a fuller discussion of this issue.

²⁸ That is, for λ_q close to one,

$$\pi_q^R \approx \frac{\theta_q}{q} (\bar{d}_q^R + \lambda_q \bar{d}_q^P) - f_q$$

and

$$\pi_q^P \approx \frac{\theta_q}{q} (\lambda_q \bar{d}_q^R + \bar{d}_q^P) - f_q,$$

so $\bar{d}_q^R > \bar{d}_q^P$ implies $\pi_q^R > \pi_q^P$, where π_q^k is the net profit of a typical producer of quality q in country k .

try produces all the low-quality goods. We record these results in corollary 1.

COROLLARY 1. Suppose that transport costs are small. (i) If $G^{\mathcal{R}}(y) = G^{\mathcal{P}}(y)$ for all y and $N^{\mathcal{R}} > N^{\mathcal{P}}$, then $n_H^{\mathcal{P}} = n_L^{\mathcal{P}} = 0$ and only $\mathcal{R}$ produces and exports goods of quality H and L . (ii) If $N^{\mathcal{R}} = N^{\mathcal{P}}$ and $G^{\mathcal{R}}(y) < G^{\mathcal{P}}(y)$ for all y , then $n_H^{\mathcal{P}} = n_L^{\mathcal{R}} = 0$, only $\mathcal{R}$ produces and exports goods of quality H , and only $\mathcal{P}$ produces and exports goods of quality L .

We can also readily examine the effects of a fall in trading costs in an equilibrium with specialization by quality level. Suppose, for example, that only $\mathcal{R}$ produces high-quality goods whereas only $\mathcal{P}$ produces low-quality goods, as when the countries are of similar size and the income distribution in $\mathcal{R}$ first-order stochastically dominates that in $\mathcal{P}$. Since every consumer buys one unit of the differentiated product of some quality level or another,

$$n_L x_L + n_H x_H = N^{\mathcal{R}} + N^{\mathcal{P}}, \quad (19)$$

where n_q is the equilibrium number of varieties of quality q , all produced in $\mathcal{R}$ for $q = H$ and all produced in $\mathcal{P}$ for $q = L$, and x_q is the break-even quantity of sales per firm for producers of goods with quality q as before. This linear relationship implies that a change in trade costs that induces entry at one quality level in the country where that quality is produced also forces exit of producers of the other quality level in the other country.

Next we can use (14) to calculate per capita domestic sales for each type of firm, recognizing that $\tilde{n}_H^{\mathcal{R}} = n_H$, $\tilde{n}_L^{\mathcal{R}} = \lambda_L n_L$, $\tilde{n}_H^{\mathcal{P}} = \lambda_H n_H$, and $\tilde{n}_L^{\mathcal{P}} = n_L$. Also, export sales per firm are $(N^{\mathcal{P}} - n_L d_L^{\mathcal{P}})/n_H$ for a typical producer of a high-quality good in $\mathcal{R}$ and $(N^{\mathcal{R}} - n_H d_H^{\mathcal{R}})/n_L$ for a typical producer of a low-quality good in $\mathcal{P}$. The fact that total sales by each type of firm must attain the break-even level gives us two more equations for n_H and n_L , one of which is redundant given (19).²⁹

Suppose now that the cost of transporting high-quality goods falls. This shifts demand toward the (imported) high-quality products in $\mathcal{P}$, without affecting demands in $\mathcal{R}$ (at the initial numbers of producers). The shift in the composition of demand induces additional entry by producers of high-quality goods in $\mathcal{R}$, whereas some producers of low-quality goods in $\mathcal{P}$ are forced to exit the market. Then, in both countries, the effective number of high-quality varieties rises whereas the effective number of low-quality varieties falls. From (17), we see that the wealthiest consumers in both countries must gain whereas the poorest consumers in both countries may gain or lose. The consequences of a reduction in the cost of shipping low-quality goods are analogous: the variety of

²⁹ Details of these equations and their comparative statics can be found in the Appendix.

low-quality goods expands and that of high-quality goods contracts in each country, to the benefit of poor consumers and the possible detriment of those who are well off.

VI. Commercial Policy

We turn to the impact of commercial policy. With fixed markups over delivered costs, tariffs do not alter the terms of trade. All welfare effects of tariffs emanate from the induced entry and exit of firms producing at different quality levels and from the shifts in the distribution of income that result from the disposition of tariff revenue.

The novel effects of tariffs in our model are well illustrated in a simple setting in which transportation costs are positive but close to zero and each quality class is produced in only one country. We assume that $\mathcal{R}$ produces all varieties of high-quality products whereas $\mathcal{P}$ produces all varieties of low-quality products and examine a specific tariff of t per unit that is introduced in country $\mathcal{R}$.³⁰ In this setting with nonhomothetic demands, the manner of redistribution of government revenues influences the effects of tariffs on aggregate demand. For concreteness, we assume that tariffs are redistributed to consumers on an equal per capita basis. Each consumer receives additional income of $r = tn_L d_L^{\mathcal{R}}/N^{\mathcal{R}}$, where $d_L^{\mathcal{R}}$ is the aggregate demand for a typical (imported) low-quality product in country $\mathcal{R}$.

A tariff raises the relative price of low-quality goods in country $\mathcal{R}$. Considering both the price hike and the redistributed proceeds, (4) implies that the aggregate demand for a typical low-quality product there becomes

$$d_L^{\mathcal{R}} = \frac{N^{\mathcal{R}}}{n_L} \mathbb{E}^{\mathcal{R}} \left[\frac{(n_L)^{\theta_L} \phi_L(y) e^{(r-t)L}}{(n_L)^{\theta_L} \phi_L(y) e^{(r-t)L} + (n_H)^{\theta_H} \phi_H(y) e^{rH}} \right].$$

The per capita demand for a typical high-quality product in $\mathcal{R}$ is

$$d_H^{\mathcal{R}} = \frac{N^{\mathcal{R}}}{n_H} \mathbb{E}^{\mathcal{R}} \left[\frac{(n_H)^{\theta_H} \phi_H(y) e^{rH}}{(n_L)^{\theta_L} \phi_L(y) e^{(r-t)L} + (n_H)^{\theta_H} \phi_H(y) e^{rH}} \right].$$

Note that $r = tn_L d_L^{\mathcal{R}}/N^{\mathcal{R}} = t\rho_L^{\mathcal{R}} > 0$ whereas $r - t = -t(1 - \rho_L^{\mathcal{R}}) < 0$, so the tariff-cum-redistribution shifts demand in $\mathcal{R}$ from low-quality goods to high-quality goods, for given numbers of each type of product. Consequently, the tariff induces entry of firms producing high-quality goods in $\mathcal{R}$ and exit by firms producing low-quality goods in $\mathcal{P}$.

How does the tariff affect the welfare of different income groups in country $\mathcal{R}$? The average welfare of individuals with income y in country

³⁰ In this setting with specialization, a tariff on all imports in country $\mathcal{R}$ is indistinguishable from a tariff on imports of low-quality goods.

$\mathcal{R}$ can be written analogously to (12) as

$$v^{\mathcal{R}}(y) = n_H^{\theta_H} \phi_H(y) e^{\tau_H} + n_L^{\theta_L} \phi_L(y) e^{(\tau - \theta)L}.$$

Differentiating the expression for $v^{\mathcal{R}}(y)$ at $t = 0$ and using (19) and the market-clearing condition, $n_q x_q = \rho_q^{\mathcal{R}} N^{\mathcal{R}} + \rho_q^{\mathcal{P}} N^{\mathcal{P}}$, we can express the impact of a small tariff on the average welfare of those with income y in country $\mathcal{R}$ as

$$\begin{aligned} \hat{v}^{\mathcal{R}}(y)|_{t=0} &= \bar{\rho}_H \bar{\rho}_L \left[\theta_H \frac{\rho_H^{\mathcal{R}}(y)}{\bar{\rho}_H} - \theta_L \frac{\rho_L^{\mathcal{R}}(y)}{\bar{\rho}_L} \right] (\hat{n}_H - \hat{n}_L) \\ &\quad + \rho_H^{\mathcal{R}} \rho_L^{\mathcal{R}} \left[\frac{\rho_H^{\mathcal{R}}(y)}{\rho_H^{\mathcal{R}}} H - \frac{\rho_L^{\mathcal{R}}(y)}{\rho_L^{\mathcal{R}}} L \right] dt, \end{aligned} \quad (20)$$

where $\bar{\rho}_q = n_q x_q / (N^{\mathcal{R}} + N^{\mathcal{P}})$ is the fraction of consumers worldwide who buy a differentiated product with quality q . The right-hand side of (20) combines two terms, a *composition effect* and a *redistribution effect*. The composition effect should be familiar. It reflects the rise in the number of high-quality varieties and the fall in the number of low-quality varieties induced by the tariff. The richest income group in $\mathcal{R}$ buys a greater fraction of high-quality goods and a smaller fraction of low-quality goods than the average consumer in the world economy, that is, $\rho_H^{\mathcal{R}}(y_{\max}) > \bar{\rho}_H$ and $\rho_L^{\mathcal{R}}(y_{\max}) < \bar{\rho}_L$. Therefore, since $\theta_H > \theta_L$, the composition effect certainly benefits the richest income group in country $\mathcal{R}$. However, the poorest income group in that country may well lose from the change in the composition of differentiated products.

The redistribution effect reflects the transfers of income implied by the lump-sum redistribution of tariff revenues. The tariff transfers income from those who choose to purchase an imported, low-quality product to those who choose to purchase a domestic, high-quality product. The rich are more likely to buy a high-quality product than the poor, so they are most likely to benefit from these transfers. Indeed, for the richest income group, $\rho_H^{\mathcal{R}}(y_{\max}) > \rho_H^{\mathcal{R}}$ and $\rho_L^{\mathcal{R}}(y_{\max}) < \rho_L^{\mathcal{R}}$ and the members of this group must gain on average from the redistribution effect as well. But notice that the redistribution effect might also benefit the poor since $H > L$. The reason is that the tariff transfers income from those in any income class who happen to prefer one of the low-quality varieties to those in that same class who happen to prefer one of the high-quality varieties. The latter group has a higher marginal utility of income due to the complementarity in preferences between quality of the differentiated good and quantity of the numeraire good. If the quality difference between differentiated products is large, it may be that $[\rho_H^{\mathcal{R}}(y_{\min})/\rho_H^{\mathcal{R}}]H > [\rho_L^{\mathcal{R}}(y_{\min})/\rho_L^{\mathcal{R}}]L$, in which case the redistribution effects of a small tariff serve to benefit even the poorest income group in $\mathcal{R}$. Indeed, the combined composition and redistribution effect can be

positive for those with income $y_{\min}$, in which case a tariff would raise the average welfare of every income group in $\mathcal{R}$, despite the absence of any terms-of-trade improvement.³¹

VII. Trade with Many Quality Levels and Many Countries

In this section, we extend our analysis of the trade pattern to the general case with many quality levels and many countries. In so doing, we are able to make contact with the empirical evidence on vertical specialization cited in the introduction, much of which bears on a country's relative exports to different markets or its relative imports from different suppliers. To address such issues, we need to examine a setting with more than two countries.

We now assume that there is an arbitrary set Q of quality levels indexed by q and K countries indexed by k . We suppose that all countries are of equal size and normalize their populations so that $N^k = 1$ for $k = 1, 2, \dots, K$. By neglecting variation in country size, we can focus on differences in the level and distribution of income without invoking a taxonomy of cases. Our key assumption in this section is that countries can be ordered unambiguously from poorest to richest. In particular, we will assume that the income distributions in any pair of countries k and k' satisfy the monotone likelihood ratio property (MLRP), as in the following assumption.

ASSUMPTION 2. If $k' > k$, then

$$\frac{g^{k'}(y')}{g^{k'}(y)} > \frac{g^k(y')}{g^k(y)}$$

for all $y' > y \geq y_{\min}$.

In assumption 2, $g^k(y) \equiv dG^k(y)/dy$ is the probability density function for income in country k . In words, assumption 2 says that in a comparison of any two countries and any two income levels, the richer country (with the higher country index) has relatively more consumers at the higher level of income. This is a somewhat stronger assumption on the ordering of income distributions than that of first-order stochastic dominance—which we used in the preceding sections—but we shall need it for the statements that follow.

The trading environment is a straightforward extension of the closed-economy model described in Sections II and III. All countries share the same technologies and the same distributions of consumer preferences over varieties. All exports of goods of quality q incur a per-unit shipping cost of τ_q , regardless of source or destination. These shipping costs give

³¹ A small tariff raises the average welfare of all income groups, e.g., when the parameters take the values that were used to generate fig. 1 in Fajgelbaum et al. (2009).

rise to a parameter $\lambda_q = e^{-q\tau_q/\theta_q}$ such that the effective variety of goods of quality q in country k is given by

$$\tilde{n}_q^k = n_q^k + \lambda_q \sum_{l, l \neq k} n_q^l, \quad q \in Q, k, l = 1, \dots, K. \quad (21)$$

In country k , the aggregate demand for a local variety with quality q is

$$d_q^k = \frac{1}{\tilde{n}_q^k} \mathbb{E}^k \left[\frac{(\tilde{n}_q^k)^{\theta_q} \phi_q(y)}{\sum_{\omega \in Q} (\tilde{n}_\omega^k)^{\theta_\omega} \phi_\omega(y)} \right], \quad q \in Q, k = 1, \dots, K, \quad (22)$$

whereas purchases of a typical imported variety with this quality are $\lambda_q d_q^k$. We denote the right-hand side of (22) by $\Gamma_q^k(\tilde{\mathbf{n}}^k)$, where $\tilde{\mathbf{n}}^k$ is the vector of effective varieties available in country k with typical element $\tilde{n}_q^k$. An important property of these aggregate demands is captured in the following lemma.³²

LEMMA 1. Under assumption 2, the function $\Gamma_q^k(\tilde{\mathbf{n}})$ is log supermodular in (q, k) , so that

$$\frac{\Gamma_{q'}^{k'}(\tilde{\mathbf{n}})}{\Gamma_q^k(\tilde{\mathbf{n}})} > \frac{\Gamma_{q'}^k(\tilde{\mathbf{n}})}{\Gamma_q^k(\tilde{\mathbf{n}})} \quad \text{for all } k' > k \text{ and } q' > q.$$

The lemma says that if the same effective variety is available in two countries, the relative demand for the higher-quality goods will be greater in the richer country.³³ This property of the demand system puts a great deal of structure on the home-market effects, which in turn guide the trade patterns in a world economy satisfying the MLRP in incomes.

As before, we consider in turn the two extreme possibilities concerning the size of transport costs. First, we assume that trading costs are sufficiently large that all quality levels are produced in every country. Then we assume that trading costs are sufficiently small that every quality level is produced in only one country. In each of these cases, we are able to derive strong predictions about the bilateral trade flows that mirror patterns found in the data.

A. Large Transport Costs

We know, as before, that all qualities are produced in every country in the autarky equilibrium. Such a pattern of diversified production also obtains in a trade equilibrium when transport costs are sufficiently high. We examine now an equilibrium in which every country produces some varieties in all quality segments.

³² The proof, which relies on lemma 2 in Athey (2002) that is originally due to Ahlswede and Daykin (1978), is given in the Appendix.

³³ Note that this statement does not rely on the countries being of equal size. That is, we could instead define $\Gamma_q^k(\tilde{\mathbf{n}})$ as the right-hand side of (22) multiplied by N^k , and the inequality cited in the lemma would still apply.

If production of quality q takes place in country k , then total sales by firms in this country must match the break-even quantity, $x_q = f_q q / \theta_q$. This implies, as in the case with two countries, that all firms producing quality q make the same local sales and the same export sales (in total and in each individual market), no matter where in the world they are located.³⁴ In the bilateral trade between country k and country k' , the latter will be a net exporter of varieties of quality q if and only if there are more varieties with this quality produced in k' than in k , that is, if and only if $n_q^{k'} > n_q^k$. But it is straightforward to show, as we did for the case of two countries, that $n_q^{k'} > n_q^k$ if and only if $\tilde{n}_q^{k'} > \tilde{n}_q^k$.³⁵ To find the pattern of bilateral trade between any pair of countries for any quality level, we need only find which country has access to a larger effective number of varieties in the trade equilibrium.

In the Appendix we establish the following lemma concerning the relative effective variety of two different quality levels available in any pair of countries.

LEMMA 2. Suppose that some varieties of q and q' are produced in countries k and k' . Then, under assumptions 1 and 2,

$$\frac{\tilde{n}_{q'}^{k'}}{\tilde{n}_q^{k'}} > \frac{\tilde{n}_{q'}^k}{\tilde{n}_q^k} \quad \text{for } k' > k \text{ and } q' > q.$$

The lemma says that a relatively greater effective number of varieties of a higher-quality good are consumed in a richer country than in a poorer country. Intuitively, if the effective variety were the same in the two countries, then by lemma 1, the richer country would have the relatively larger domestic market for the higher-quality product. This would mean that producers of the higher-quality good could make relatively greater profits in the richer country than in the poorer country thanks to their savings on transport costs. But this implies that there would be incentives for entry of suppliers of the higher-quality good in the richer country or incentives for entry of suppliers of the lower-quality good in the poorer country. Of course, in equilibrium, all firms must make zero profits in all market segments and locations. Under assump-

³⁴ The break-even conditions require

$$d_q^k + \lambda_q \sum_{l, l \neq k} d_q^l = x_q \quad \text{for } q \in Q \text{ and } k, l = 1, \dots, K.$$

This implies

$$d_q^k = \frac{x_q}{1 + \lambda_q(K-1)} \quad \text{for all } q \in Q \text{ and all } k = 1, \dots, K.$$

A producer of some variety of quality q in country k exports $\lambda_q d_q^k = \lambda_q d_q^k$ units of this good to every foreign market.

³⁵ From the definition of effective variety in k and k' given in (21), subtraction yields

$$n_q^{k'} - n_q^k = \frac{1}{1 - \lambda_q} (\tilde{n}_q^{k'} - \tilde{n}_q^k) \quad \text{for all } k \text{ and } k' \text{ and for all } q \in Q.$$

tion 1, which ensures that higher-quality goods also are more dissimilar, the relative incentives for entry at $\tilde{n}^k = \tilde{n}^{k'}$ are mirrored in the relative effective numbers of varieties in the trade equilibrium.

Now we can describe the bilateral trade pattern between k and k' . Suppose that country k' has positive net exports of goods with quality q in its trade with country k . This requires, as we have seen, that there be more effective varieties of q consumed in k' than in k , that is, $\tilde{n}_q^{k'} > \tilde{n}_q^k$. But then lemma 2 implies that the richer country also consumes a more effective variety of quality q' for $q' > q$ than the poorer country, which in turn implies that the richer country also has positive net exports of quality q' in its trade with the poorer country. More formally, we have the following proposition.

PROPOSITION 5. Suppose that all countries are of equal size and all quality levels are produced in all countries. Then, under assumptions 1 and 2, for every pair of countries k and k' , $k' > k$, there exists a quality level $q^*(k, k')$ such that country k' exports on net to country k all goods with quality greater than or equal to $q^*(k, k')$ and imports on net all goods with quality less than $q^*(k, k')$.

Proposition 5 generalizes part ii of proposition 2 to the case of many countries and many quality levels. It provides sufficient conditions under which a richer country has positive net exports of all the highest-quality goods and positive net imports of all the lowest-quality goods in its bilateral trade with a poorer country.

Proposition 5 implies that the average quality of a country's exports is increasing in k , in keeping with the findings by Schott (2004) and Hummels and Klenow (2005). In fact, the proposition offers a stronger prediction: it suggests that in all bilateral trades, the basket of goods that a richer country exports to a poorer country should be of higher average quality than the basket of goods that flow in the opposite direction.

B. Small Transport Costs

We now consider the opposite extreme case in which transport costs are quite small. As we know, if two countries both produce varieties in the same quality segment, then firms in each country must make the same domestic sales.³⁶ But, in the limit as the vector of transport costs becomes vanishingly small, the domestic sales of each firm approach those of the integrated equilibrium. Only exceptionally will the aggregate demands for a given quality level coincide in two countries in the integrated equilibrium. It follows that, generically, varieties of a given quality level are produced in only one country in a trade equilibrium

³⁶ The argument is the same as in n. 34

with small transport costs. Moreover, even if we allow for exceptional income distributions, any two countries can produce at most one quality level in common.³⁷

Which quality levels are produced in which countries? The answer is provided in the following proposition, the proof of which appears in the Appendix.

PROPOSITION 6. Suppose that all countries are of equal size and each quality level is produced in only one country. Then, under assumption 2, if country k produces quality q and country $k' > k$ produces quality q' , then $q' > q$.

The proposition does not rule out the possibility that some countries produce more than one quality level or that others produce none at all. In particular, if there are more countries than quality levels, there will be some countries that specialize in producing the homogeneous good and that produce none of the differentiated products. But the proposition does indicate that the higher the quality level of a product, the richer the country that produces it. The result follows from the home-market effect, which, as in the case with two countries, becomes overwhelmingly strong as transport costs go to zero. Production of a good is concentrated in the country that has the largest home market, and richer countries have larger markets than poorer countries for higher-quality goods.

Proposition 6 has strong implications for the trade flows. First, the pattern of specialization implies immediately that every country imports higher-quality goods from richer countries. This is consistent with the findings by Schott (2004), Khandelwal (2010), and Hallak and Schott (2011) of a positive correlation between the quality of goods imported by the United States and the per capita GDP of the exporting country. Second, richer countries import relatively more from countries that produce goods of higher quality. This result, which is in keeping with the empirical evidence provided by Hallak (2006), follows from the fact that, when transport costs are small, the relative demand for a higher-quality good must be greater in a richer country, as indicated in lemma 1. Third, suppose that a country exports goods in two quality segments, q and q' , with $q' > q$. This country will export relatively more of the higher-quality goods to the richer market. This too follows from lemma 1, which implies that, when transport costs are small, the relative demand for the higher-quality good must be greater in the richer country. More formally, we have established proposition 7.

³⁷ Suppose, to the contrary, that countries k and k' both produce varieties of quality q and q' . Then $d_q^k = d_q^{k'}$ and $d_{q'}^k = d_{q'}^{k'}$. In the limit, as all transport costs vanish ($\lambda_q \rightarrow 1$ for all q), this implies $\Gamma_q^k(\bar{\mathbf{n}}) = \Gamma_q^{k'}(\bar{\mathbf{n}})$ and $\Gamma_{q'}^k(\bar{\mathbf{n}}) = \Gamma_{q'}^{k'}(\bar{\mathbf{n}})$, where $\bar{\mathbf{n}}$ is the vector of total varieties in the integrated equilibrium. But this pair of equalities contradicts lemma 1, so the two countries cannot in fact share in production of these two quality levels.

PROPOSITION 7. Suppose that $k' > k$ and that transport costs are vanishingly small. Under assumption 2, if countries k and k' import goods of quality q and q' with $q' > q$, then k' imports relatively more of quality q' .

Notice that this proposition does not require that countries k and k' be of equal size, for the reason stated in footnote 33; that is, the ordering of relative aggregate demands for alternative quality levels in two markets is independent of the sizes of those markets. Proposition 7 yields the following immediate corollary concerning a country's exports to multiple markets.

COROLLARY 2. Suppose that transport costs are vanishingly small. If country k produces varieties in two quality segments, it exports relatively more of the higher-quality goods to a relatively richer market.

Finally, consider a country that exports varieties of a given quality to two different markets. If we control for the sizes of the importing countries, the exporter's sales are greater in the country whose income ranking is more similar to its own.³⁸ This prediction is in the spirit of the Linder hypothesis and the empirical evidence provided by Hallak (2010).

VIII. Concluding Remarks

We have developed a tractable model of trade in vertically and horizontally differentiated products. The model features discrete quality choices by consumers who differ in income levels and nonhomothetic aggregate demands for goods of different qualities. The nonhomotheticity in demand reflects a complementarity in individual preferences between the quality of the differentiated product and the quantity of a homogeneous good. Consumers have idiosyncratic components in their evaluations of the available varieties of the differentiated product. The distribution of taste parameters in the population generates a nested-logit system of product demands.

We have embedded such consumers in a simple, supply-side environment. Goods are produced from labor alone, with constant returns to scale in the homogeneous-good industry and fixed and constant-variable costs for the varieties of the differentiated products. The number of

³⁸ Suppose that $k'' > k' > k$ and consider the exports of k of some quality q to k' and k'' when transport costs are small. By lemma 1,

$$\frac{d_q^{k'}}{d_{q'}^{k'}} > \frac{d_q^{k''}}{d_{q'}^{k''}},$$

where $q' > q$ is the quality of some varieties produced and exported only by country k' . Since q' is produced in k' but not in k'' , $d_q^{k'} > d_{q'}^{k''}$. Therefore, $d_q^{k'} > d_q^{k''}$; i.e., the market for goods of quality q must be larger in k' than in k'' . The same argument applies when $k > k' > k''$ once we note that $q > q'$ in this case.

varieties at each quality level is determined by free entry in a monopolistically competitive, general equilibrium. Transport costs impede trade between countries that differ in size and in their income distributions but are otherwise similar. In this setting, a large home market for goods of a given quality confers a competitive advantage to firms located there, which renders them as net exporters in the trade equilibrium.

Our model yields predictions about the pattern of trade that are consistent with the empirical evidence. For example, we find that, among countries of similar size, the richer countries export goods of higher average quality. This is in keeping with the empirical findings by Schott (2004) and Hummels and Klenow (2005). When trading costs are small, a country imports higher-quality goods from richer trading partners, as Schott (2004), Khandelwal (2010), and Hallak and Schott (2011) have found to be true in the U.S. bilateral trade data. When a country exports varieties of a given quality to two different markets of similar size, it exports a greater volume to the country whose income ranking is more similar to its own.

Our framework lends itself readily to welfare analysis. We can decompose the welfare impact on a particular income group of, for example, reductions in trading costs into a cost-savings effect and a composition effect. The former tends to benefit all consumers, whereas the latter—reflecting the induced change in the relative numbers of low- and high-quality products—often benefits consumers at one end of the income distribution at the expense of those at the other. We find that, in the absence of supply-side determinants of comparative advantage, trade between countries at different levels of income tends to benefit on average the poorer consumers in the richer country and the richer consumers in the poorer country. These income groups gain from integration not only because imports provide a wider variety of choices but also because trade shifts the composition of available products toward those goods that they are most likely to buy.

Our nested-logit demand system is familiar from the empirical literature. For this reason, we believe that it would be possible to estimate key parameters of the model from data on household income and spending. Armed with such estimates and observable data on expenditure patterns, one could utilize our expressions for the average welfare changes in different income groups to calculate the gains from trade across the income distribution.³⁹

³⁹ See Porto (2006) for an ambitious attempt to estimate the distributional impacts of trade policy in a model that includes heterogeneous consumer gains as well as wage and income effects. Porto does not allow for nonhomothetic demands, but rather he takes the variation in budget shares for different households to be an exogenous reflection of heterogeneous tastes.

Finally, our framework is simple enough to allow for extensions and variations. For example, it is straightforward to introduce direct foreign investment as an alternative means for firms to serve foreign markets. Then the model could shed light on the spread of Chinese and Indian multinational corporations to other developing countries (see Boston Consulting Group 2006). We are pursuing such an extension in our ongoing research.

Appendix

A. Proof of Proposition 1

The pair of equations in (11) implies

$$\frac{\mathbb{E}[\rho_L(y)]}{n_L x_L} = \frac{1}{N} \quad (\text{A1})$$

and

$$\frac{\mathbb{E}[\rho_H(y)]}{n_H x_H} = \frac{1}{N}, \quad (\text{A2})$$

where

$$\rho_L(y) = \frac{n_L^{\theta_L} \phi_L(y)}{n_L^{\theta_L} \phi_L(y) + n_H^{\theta_H} \phi_H(y)}$$

and

$$\rho_H(y) = 1 - \rho_L(y).$$

These equations are depicted as *LL* and *HH*, respectively, in figure A1. They represent combinations of the number of low-quality varieties and the number of high-quality varieties consistent with zero profits and market clearing in the two market segments. Both curves are downward sloping inasmuch as an increase in the number of competitors of any type necessitates a reduction in the number of competitors of the opposite type in order to preserve profitability. Along *LL*, $n_L \rightarrow N/x_L$ as $n_H \rightarrow 0$ and $n_L \rightarrow 0$ as $n_H \rightarrow \infty$, whereas along *HH*, $n_H \rightarrow N/x_H$ as $n_L \rightarrow 0$ and $n_H \rightarrow 0$ as $n_L \rightarrow \infty$. Therefore, the two curves must intersect in the positive quadrant; that is, an equilibrium exists with $n_H > 0$ and $n_L > 0$.

Note that

$$\begin{aligned} \frac{\partial \rho_L(y)}{\partial n_L} &= \frac{\theta_L}{n_L} \rho_L(y) \rho_H(y), \\ \frac{\partial \rho_L(y)}{\partial n_H} &= -\frac{\theta_H}{n_H} \rho_L(y) \rho_H(y). \end{aligned}$$

Using these expressions and totally differentiating (A1) and (A2) with respect to n_L and n_H , we find

FIG. A1.—Autarky equilibrium

$$\left\{ \theta_L \frac{\mathbb{E}[\rho_L(y)\rho_H(y)]}{\mathbb{E}[\rho_L(y)]} - 1 \right\} \hat{n}_L - \theta_H \frac{\mathbb{E}[\rho_L(y)\rho_H(y)]}{\mathbb{E}[\rho_L(y)]} \hat{n}_H = 0,$$

$$-\theta_L \frac{\mathbb{E}[\rho_L(y)\rho_H(y)]}{\mathbb{E}[\rho_H(y)]} \hat{n}_L + \left\{ \theta_H \frac{\mathbb{E}[\rho_L(y)\rho_H(y)]}{\mathbb{E}[\rho_H(y)]} - 1 \right\} \hat{n}_H = 0.$$

All elements in the Jacobian matrix are negative.⁴⁰ Denoting the Jacobian determinant by Δ , we have

$$\begin{aligned} \Delta &= 1 - \theta_L \frac{\mathbb{E}[\rho_L(y)\rho_H(y)]}{\mathbb{E}[\rho_L(y)]} - \theta_H \frac{\mathbb{E}[\rho_L(y)\rho_H(y)]}{\mathbb{E}[\rho_H(y)]} \\ &= 1 - \frac{\mathbb{E}[\rho_L(y)\rho_H(y)]}{\mathbb{E}[\rho_H(y)]\mathbb{E}[\rho_L(y)]} \{ \theta_L \mathbb{E}[\rho_H(y)] + \theta_H \mathbb{E}[\rho_L(y)] \}. \end{aligned} \quad (\text{A3})$$

We see that $\Delta > 0$ because $\rho_L(y)$ and $\rho_H(y)$ are negatively correlated and a weighted average of θ_L and θ_H is smaller than one. This implies that LL must be steeper

⁴⁰ Note that

$$\theta_L \frac{\mathbb{E}[\rho_L(y)\rho_H(y)]}{\mathbb{E}[\rho_L(y)]} < \theta_L \mathbb{E}[\rho_H(y)] < 1,$$

because ρ_L and ρ_H are negatively correlated and $\rho_H(y)$ and θ_L are both fractions. Similarly,

$$\theta_H \frac{\mathbb{E}[\rho_L(y)\rho_H(y)]}{\mathbb{E}[\rho_H(y)]} < 1.$$

than HH at any intersection point, which in turn implies that the equilibrium is unique.

B. Comparative Statics of the Autarky Equilibrium

1. Population Size

Now differentiate (A1) and (A2) totally with respect to n_L , n_H , and N . The solution is

$$\begin{aligned}\Delta \frac{\hat{n}_L}{\hat{N}} &= 1 - \theta_H \frac{\mathbb{E}[\rho_L(y)\rho_H(y)]}{\mathbb{E}[\rho_H(y)]} - \theta_H \frac{\mathbb{E}[\rho_L(y)\rho_H(y)]}{\mathbb{E}[\rho_L(y)]} \\ &= 1 - \theta_H \frac{\mathbb{E}[\rho_L(y)\rho_H(y)]}{\mathbb{E}[\rho_H(y)]\mathbb{E}[\rho_L(y)]} > 0\end{aligned}$$

and

$$\begin{aligned}\Delta \frac{\hat{n}_H}{\hat{N}} &= 1 - \theta_L \frac{\mathbb{E}[\rho_L(y)\rho_H(y)]}{\mathbb{E}[\rho_L(y)]} - \theta_L \frac{\mathbb{E}[\rho_L(y)\rho_H(y)]}{\mathbb{E}[\rho_H(y)]} \\ &= 1 - \theta_L \frac{\mathbb{E}[\rho_L(y)\rho_H(y)]}{\mathbb{E}[\rho_H(y)]\mathbb{E}[\rho_L(y)]} > 0.\end{aligned}$$

Note that

$$\frac{\hat{n}_H - \hat{n}_L}{\hat{N}} = \frac{1}{\Delta} (\theta_H - \theta_L) \frac{\mathbb{E}[\rho_L(y)\rho_H(y)]}{\mathbb{E}[\rho_H(y)]\mathbb{E}[\rho_L(y)]} > 0.$$

In terms of figure A1, an increase in N shifts both curves to the right, but the HH curve shifts by more along a ray through the initial intersection point.

2. First-Order Stochastic Dominance

Since $\rho_H(y)$ is monotonically increasing in y and $\rho_L(y)$ is monotonically decreasing, a first-order shift in the income distribution raises $\mathbb{E}[\rho_H(y)]$ and reduces $\mathbb{E}[\rho_L(y)]$. By (A1), the LL curve in figure A1 shifts to the left, whereas by (A2), the HH curve shifts to the right. As a result, $\hat{n}_H > 0$ and $\hat{n}_L < 0$.

3. Mean-Preserving Spread

A mean-preserving spread in $G(\cdot)$ raises $\mathbb{E}[\rho_H(y)]$ if $\rho_H(y)$ is convex in y at all y . From equation (5), we have

$$\rho'_H(y) = \rho_H(y)[H - q_a(y)]$$

and therefore

$$\rho''_H(y) = -\rho_H(y)q'_a(y) + \rho_H(y)[H - q_a(y)]^2.$$

After substituting the definition $q_a(y) \equiv \rho_H(y)H + \rho_L(y)L$, noting that $q'_a(y) = \rho'_H(y)(H - L)$, and rearranging terms, we have

$$\rho''_H(y) = \rho_H(y)[H - q_a(y)](H - L)[\rho_L(y) - \rho_H(y)].$$

Therefore, $\rho''_H(y) > 0$ if and only if $\rho_L(y) > \rho_H(y)$ for all y . This is sufficient for a

mean-preserving spread in income to induce an increase in $\mathbb{E}[\rho_H(y)]$ and a decline in $\mathbb{E}[\rho_L(y)]$. In such circumstances, the LL curve shifts to the left and the HH curve shifts to the right, so that $\hat{n}_H > 0$ and $\hat{n}_L < 0$.

C. Comparative Statics of Trade Equilibrium with Diversified Production

Now (15) implies

$$\frac{\mathbb{E}^k[\rho_L^k(y)]}{\tilde{n}_L^k x_L / (1 + \lambda_L)} = \frac{1}{N^k} \quad (\text{A4})$$

and

$$\frac{\mathbb{E}^k[\rho_H^k(y)]}{\tilde{n}_H^k x_H / (1 + \lambda_H)} = \frac{1}{N^k} \quad (\text{A5})$$

for $k = \mathcal{R}, \mathcal{P}$, where

$$\rho_L^k(y) = \frac{(\tilde{n}_L^k)^{\theta_L} \phi_L(y)}{(\tilde{n}_L^k)^{\theta_L} \phi_L(y) + (\tilde{n}_H^k)^{\theta_H} \phi_H(y)}$$

and

$$\rho_H^k(y) = 1 - \rho_L^k(y).$$

The equilibrium values of $\tilde{n}_L^k$ and $\tilde{n}_H^k$ are found in a figure analogous to figure A1, in which $x_q/(1 + \lambda_q)$ now plays the role that x_q played previously and N^k plays the role of N .

It is clear that an increase in λ_H , for example, shifts the new HH curve to the right, leading to a rise in $\tilde{n}_H^k$ and a fall in $\tilde{n}_L^k$ for $k = \mathcal{R}, \mathcal{P}$. An equiproportionate rise in $1 + \lambda_H$ and $1 + \lambda_L$ shifts both curves to the right, but the HH curve shifts by more along a ray through the initial intersection point. The result is an increase in $\tilde{n}_H^k$ and $\tilde{n}_L^k$, with $\tilde{n}_H^k$ rising proportionately more, for $k = \mathcal{R}, \mathcal{P}$.

D. Equilibrium Numbers of Brands with Small Trade Costs

When high-quality varieties are produced only in $\mathcal{R}$ and low-quality varieties are produced only in $\mathcal{P}$, the zero-profit conditions imply

$$\frac{N^{\mathcal{P}} \mathbb{E}^{\mathcal{P}}[\rho_L^{\mathcal{P}}(y)] + \lambda_L N^{\mathcal{R}} \mathbb{E}^{\mathcal{R}}[\rho_L^{\mathcal{R}}(y)]}{n_L x_L} = 1$$

and

$$\frac{\lambda_H N^{\mathcal{P}} \mathbb{E}^{\mathcal{P}}[\rho_H^{\mathcal{P}}(y)] + N^{\mathcal{R}} \mathbb{E}^{\mathcal{R}}[\rho_H^{\mathcal{R}}(y)]}{n_H x_H} = 1,$$

where

$$\begin{aligned}\rho_L^{\mathcal{P}}(y) &= \frac{n_L^{\theta_L} \phi_L(y)}{(\lambda_H n_H)^{\theta_H} \phi_H(y) + n_L^{\theta_L} \phi_L(y)}, \\ \rho_L^{\mathcal{R}}(y) &= \frac{(\lambda_L n_L)^{\theta_L} \phi_L(y)}{n_H^{\theta_H} \phi_H(y) + (\lambda_L n_L)^{\theta_L} \phi_L(y)}, \\ \rho_H^{\mathcal{P}}(y) &= 1 - \rho_L^{\mathcal{P}}(y), \\ \rho_H^{\mathcal{R}}(y) &= 1 - \rho_L^{\mathcal{R}}(y).\end{aligned}$$

We can use these equations to draw LL and HH curves analogous to those in figure A1, with similar properties. Again, the curves must intersect and the LL curve must be steeper at any point of intersection, which implies that the equilibrium is unique.

A decline in trading costs for high-quality products ($d\lambda_H > 0$) causes the new LL curve to shift to the left and the new HH curve to shift to the right, leading to a rise in the number of high-quality varieties produced in $\mathcal{R}$ ($\hat{n}_H > 0$) and a fall in the number of low-quality varieties produced in $\mathcal{P}$ ($\hat{n}_L < 0$).

E. Many Countries and Many Quality Levels

1. Proof of Lemma 1

Note that $\phi_q(y) = e^{(y - c_q)q - \theta_q}$ is strictly log supermodular in (y, q) ; that is, $\phi_q(y')/\phi_q(y'') > \phi_{q'}(y')/\phi_{q'}(y'')$ for all $q, q' \in Q$, $q' > q$ and $y' > y'' \geq y_{\min}$. Therefore, assumption 2 implies that $\phi_q(y)g^k(y)$ is strictly log supermodular in (y, q, k) and so is $\phi_q(y)g^k(y)/(\tilde{n}_q)^{1-\theta_q} \sum_{\omega \in Q} (\tilde{n}_\omega)^{\theta_\omega} \phi_\omega(y)$. It follows from lemma 2 in Athey (2002), due originally to Ahlswede and Daykin (1978), that $\Gamma_q^k(\tilde{\mathbf{n}})$ is strictly log supermodular in (q, k) inasmuch as integration preserves log supermodularity.

2. Proof of Lemma 2

Suppose that a set of quality levels $\mathbb{Q} \subseteq Q$ are produced in countries $k \in \mathbb{K} \subseteq \{1, 2, \dots, K\}$; then the demand for a typical local variety of quality $q \in \mathbb{Q}$ is the same in all these countries, that is,

$$d_q^k = \frac{x_q}{1 + \lambda_q(K-1)} \quad \text{for all } q \in \mathbb{Q}; k \in \mathbb{K}.$$

This implies

$$\left(\frac{\tilde{n}_q^{k'}}{\tilde{n}_q^k} \right)^{1-\theta_q} = \frac{\int_{y_{\min}}^{\infty} \Lambda_q^{k'}(\tilde{\mathbf{n}}^{k'}, y) dy}{\int_{y_{\min}}^{\infty} \Lambda_q^k(\tilde{\mathbf{n}}^k, y) dy} \quad \text{for } q \in \mathbb{Q}; k, k' \in \mathbb{K}, \quad (\text{A6})$$

where

$$\Lambda_q^k(\tilde{\mathbf{n}}, y) = \frac{\phi_q(y)g^k(y)}{\sum_{\omega \in Q} (\tilde{n}_\omega)^{\theta_\omega} \phi_\omega(y)}.$$

Note that

$$\frac{\Lambda_q^{k'}(\tilde{\mathbf{n}}^{k'}, y)}{\Lambda_q^k(\tilde{\mathbf{n}}^k, y)} = \frac{\phi_{q'}(y)}{\phi_q(y)} = \frac{\Lambda_{q'}(\tilde{\mathbf{n}}^k, y)}{\Lambda_q^k(\tilde{\mathbf{n}}^k, y)}, \quad (\text{A7})$$

$$\frac{\Lambda_q^{k'}(\tilde{\mathbf{n}}^k, y')}{\Lambda_q^k(\tilde{\mathbf{n}}^k, y'')} = \frac{\phi_q(y')g^{k'}(y')}{\phi_q(y'')g^{k'}(y'')} > \frac{\phi_q(y')g^k(y')}{\phi_q(y'')g^k(y'')} = \frac{\Lambda_q^k(\tilde{\mathbf{n}}^k, y')}{\Lambda_q^k(\tilde{\mathbf{n}}^k, y'')} \quad (\text{A8})$$

for $y' > y''$ and $k' > k$, where the inequality results from log supermodularity of $g^k(y)$. Next note that (A6) together with (A7) imply that for $q, q' \in \mathbb{Q}$ and $k, k' \in \mathbb{K}$

$$\begin{aligned} \frac{(\tilde{n}_q^{k'}/\tilde{n}_q^k)^{1-\theta_q}}{(\tilde{n}_q^{k'}/\tilde{n}_q^k)^{1-\theta_q}} &= \frac{\int_{y_{\min}}^{\infty} \Lambda_q^{k'}(\tilde{\mathbf{n}}^k, y) [\phi_q(y)/\phi_q(y)] dy / \int_{y_{\min}}^{\infty} \Lambda_q^{k'}(\tilde{\mathbf{n}}^k, y) dy}{\int_{y_{\min}}^{\infty} \Lambda_q^k(\tilde{\mathbf{n}}^k, y) [\phi_q(y)/\phi_q(y)] dy / \int_{y_{\min}}^{\infty} \Lambda_q^k(\tilde{\mathbf{n}}^k, y) dy} \\ &= \frac{\int_{y_{\min}}^{\infty} \Psi_q^{k'}(\tilde{\mathbf{n}}^k, y) [\phi_q(y)/\phi_q(y)] dy}{\int_{y_{\min}}^{\infty} \Psi_q^k(\tilde{\mathbf{n}}^k, y) [\phi_q(y)/\phi_q(y)] dy}, \end{aligned}$$

where

$$\Psi_q^k(\tilde{\mathbf{n}}, y) = \frac{\Lambda_q^k(\tilde{\mathbf{n}}, y)}{\int_{y_{\min}}^{\infty} \Lambda_q^k(\tilde{\mathbf{n}}, z) dz}$$

is a density, so that $\int_{y_{\min}}^{\infty} \Psi_q^k(\tilde{\mathbf{n}}^k, y) [\phi_q(y)/\phi_q(y)] dy$ represents a weighted average of $\phi_q(y)/\phi_q(y)$. However, (A8) implies that

$$\frac{\Psi_q^{k'}(\tilde{\mathbf{n}}^k, y')}{\Psi_q^{k'}(\tilde{\mathbf{n}}^k, y'')} > \frac{\Psi_q^k(\tilde{\mathbf{n}}^k, y')}{\Psi_q^k(\tilde{\mathbf{n}}^k, y'')},$$

so that in country k' (where $k' > k$) the weights are relatively larger for larger income levels. Moreover, $\phi_q(y)/\phi_q(y)$ is increasing in y , and therefore

$$\int_{y_{\min}}^{\infty} \Psi_q^{k'}(\tilde{\mathbf{n}}^k, y) [\phi_q(y)/\phi_q(y)] dy > \int_{y_{\min}}^{\infty} \Psi_q^k(\tilde{\mathbf{n}}^k, y) [\phi_q(y)/\phi_q(y)] dy,$$

which implies that

$$\left(\frac{\tilde{n}_q^{k'}}{\tilde{n}_q^k} \right)^{1-\theta_q} > \left(\frac{\tilde{n}_q^{k'}}{\tilde{n}_q^k} \right)^{1-\theta_q}.$$

It follows from assumption 1 that

$$\frac{\tilde{n}_q^{k'}}{\tilde{n}_q^k} > \frac{\tilde{n}_q^k}{\tilde{n}_q^k} \quad \text{for } q, q' \in \mathbb{Q}; k, k' \in \mathbb{K}; q' > q \text{ and } k' > k.$$

3. Proof of Proposition 6

Suppose, to the contrary, that $q' < q$. It follows from the fact that producers of quality q in k and of quality q' in k' must break even and from the fact that there must be no-profit opportunities for producers of quality q in k' and of quality q' in k that

$$\begin{aligned}
d_q^k + \lambda_q d_q^{k'} + \lambda_q \sum_{l=1, l \neq k, l \neq k'}^K d_q^l &= x_q, \\
\lambda_q d_q^k + d_q^{k'} + \lambda_q \sum_{l=1, l \neq k, l \neq k'}^K d_q^l &\leq x_q, \\
d_{q'}^{k'} + \lambda_q d_{q'}^k + \lambda_{q'} \sum_{l=1, l \neq k, l \neq k'}^K d_{q'}^l &= x_{q'}, \\
\lambda_q d_{q'}^{k'} + d_{q'}^k + \lambda_{q'} \sum_{l=1, l \neq k, l \neq k'}^K d_{q'}^l &\leq x_{q'}.
\end{aligned}$$

These no-profit conditions imply $d_q^k \geq d_q^{k'}$, $d_q^{k'} \geq d_{q'}^{k'}$, and therefore

$$\frac{d_q^{k'}}{d_q^k} \geq \frac{d_{q'}^k}{d_{q'}^{k'}},$$

or by (22),

$$\frac{\Gamma_q^{k'}(\tilde{\mathbf{n}}^{k'})}{\Gamma_q^{k'}(\tilde{\mathbf{n}}^{k'})} \geq \frac{\Gamma_q^k(\tilde{\mathbf{n}}^k)}{\Gamma_q^k(\tilde{\mathbf{n}}^k)}.$$

For $\lambda \downarrow (1, 1, \dots, 1)$ we have $\tilde{\mathbf{n}}^k \rightarrow \bar{\mathbf{n}}$ and $\tilde{\mathbf{n}}^{k'} \rightarrow \bar{\mathbf{n}}$, which contradicts lemma 1 when $q' < q$. Therefore, the richer country must produce the goods of higher quality.

References

- Ahlsvede, Rudolf, and David E. Daykin. 1978. "An Inequality for Weights of Two Families of Sets, Their Unions and Intersections." *Probability Theory and Related Fields* 43 (3): 183–85.
- Aiginger, Karl, and Esteban Rossi-Hansberg. 2006. "Specialization versus Concentration: A Note on Theory and Evidence." *Empirica* 33 (4): 255–66.
- Amiti, Mary. 1998. "Inter-industry Trade in Manufactures: Does Country Size Matter?" *J. Internat. Econ.* 44 (2): 231–55.
- Anderson, Simon P., Andre De Palma, and Jacques-Francois Thisse. 1992. *Discrete Choice Theory of Product Differentiation*. Cambridge, MA: MIT Press.
- Athey, Susan. 2002. "Monotone Comparative Statics under Uncertainty." *Q.J.E.* 117 (1): 187–223.
- Auer, Raphael A. 2010. "Consumer Heterogeneity and the Impact of Trade Liberalization: How Representative Is the Representative Agent Framework?" Working Paper no. 2010-13, Swiss National Bank, Zurich.
- Baldwin, Richard E., and James Harrigan. 2011. "Zeros, Quality and Space: Trade Theory and Trade Evidence." *American Econ. J.: Microeconomics* 3 (1): 60–88.
- Ben-Akiva, Moshe E. 1973. "Structure of Passenger Travel Demand Models." PhD diss., Massachusetts Inst. Tech.
- Bergstrand, Jeffrey H. 1990. "The Heckscher-Ohlin-Samuelson Model, the Linder Hypothesis and the Determinants of Bilateral Intra-industry Trade." *Econ. J.* 100 (403): 1216–29.
- Berry, Steven, James Levinsohn, and Ariel Pakes. 1995. "Automobile Prices in Market Equilibrium." *Econometrica* 63 (4): 841–90.
- Bils, Mark, and Peter J. Klenow. 2001. "Quantifying Quality Growth." *A.E.R.* 91 (4): 1006–30.
- Boston Consulting Group. 2006. *The New Global Challengers: How 100 Top Com-*

- panies from the Rapidly Developing Economies Are Changing the World*. Boston: Boston Consulting Group.
- Broda, Christian, and John Romalis. 2009. "The Welfare Implications of Rising Price Dispersion." Manuscript, Univ. Chicago, http://faculty.chicagobooth.edu/john.romalis?Research/Draft_v7.pdf.
- Choi, Yo Chul, David Hummels, and Chong Xiang. 2009. "Explaining Import Quality: The Role of the Income Distribution." *J. Internat. Econ.* 77 (2): 265–75.
- Dalgin, Muhammed, Vitor Trindade, and Devashish Mitra. 2008. "Inequality, Nonhomothetic Preferences, and Trade: A Gravity Approach." *Southern Econ. J.* 74 (3): 747–74.
- Davis, Donald R. 1998. "The Home Market, Trade, and Industrial Structure." *A.E.R.* 88 (5): 1264–76.
- Davis, Donald R., and David E. Weinstein. 2003. "Market Access, Economic Geography and Comparative Advantage: An Empirical Test." *J. Internat. Econ.* 59 (1): 1–23.
- Eaton, Jonathan, and Samuel Kortum. 2002. "Technology, Geography and Trade." *Econometrica* 70 (5): 1741–79.
- Epifani, Paolo, and Gino Gancia. 2006. "Increasing Returns, Imperfect Competition, and Factor Prices." *Rev. Econ. and Statis.* 88 (4): 583–98.
- Fajgelbaum, Pablo D., Gene M. Grossman, and Elhanan Helpman. 2009. "Income Distribution, Product Quality, and International Trade." Working Paper no. 15329, NBER, Cambridge, MA.
- Fieler, Ana Cecilia. 2011. "Non-homotheticity and Bilateral Trade: Evidence and a Quantitative Explanation." *Econometrica* 79 (4): 1069–1101.
- Flam, Harry, and Elhanan Helpman. 1987. "Vertical Product Differentiation and North-South Trade." *A.E.R.* 77 (5): 810–22.
- Foellmi, Reto, Christian Hepenstrick, and Josef Zweimüller. 2010. "Non-homothetic Preferences, Parallel Imports and the Extensive Margin of International Trade." Working Paper no. 1424-0459, Inst. Empirical Res. Econ., Zurich.
- Gabszewicz, J. Jaskold, and Jacques-Francois Thisse. 1979. "Price Competition, Quality, and Income Disparities." *J. Econ. Theory* 20 (3): 340–59.
- . 1980. "Entry (and Exit) in a Differentiated Industry." *J. Econ. Theory* 22 (2): 327–38.
- Goldberg, Pinelopi K. 1995. "Product Differentiation and Oligopoly in International Markets: The Case of the U.S. Automobile Industry." *Econometrica* 63 (4): 891–951.
- Hallak, Juan Carlos. 2006. "Product Quality and the Direction of Trade." *J. Internat. Econ.* 68 (1): 238–65.
- . 2010. "A Product-Quality View of the Linder Hypothesis." *Rev. Econ. and Statis.* 92 (3): 453–66.
- Hallak, Juan Carlos, and Peter K. Schott. 2011. "Estimating Cross-Country Differences in Product Quality." *Q.J.E.* 126 (1): 417–74.
- Hallak, Juan Carlos, and Jagadeesh Sivadasan. 2009. "Firm's Exporting Behavior under Quality Constraints." Working Paper no. 88, Internat. Policy Center, Ann Arbor, MI.
- Hanson, Gordon C., and Chong Xiang. 2004. "The Home-Market Effect and Bilateral Trade Patterns." *A.E.R.* 94 (4): 1108–29.
- Hummels, David, and Peter J. Klenow. 2005. "The Variety and Quality of a Nation's Exports." *A.E.R.* 95 (3): 704–23.
- Hunter, Linda. 1991. "The Contribution of Nonhomothetic Preferences to Trade." *J. Internat. Econ.* 30 (3–4): 345–58.

- Hunter, Linda, and James R. Markusen. 1988. "Per-Capita Income as a Determinant of Trade." In *Empirical Methods for International Trade*, edited by Robert C. Feenstra, 89–109. Cambridge, MA: MIT Press.
- Johnson, Robert C. Forthcoming. "Trade and Prices with Heterogeneous Firms." *J. Internat. Econ.*
- Khandelwal, Amit. 2010. "The Long and Short (of) Quality Ladders." *Rev. Econ. Studies* 77 (4): 1450–76.
- Krugman, Paul R. 1980. "Scale Economies, Product Differentiation, and the Pattern of Trade." *A.E.R.* 70 (5): 950–59.
- . 1991a. *Geography and Trade*. Cambridge, MA: MIT Press.
- . 1991b. "Increasing Returns and Economic Geography." *J.P.E.* 99 (3): 483–99.
- Kugler, Maurice, and Eric Verhoogen. Forthcoming. "Prices, Plant Size, and Product Quality." *Rev. Econ. Studies*.
- Linder, Staffan Burenstam. 1961. *An Essay on Trade and Transformation*. Stockholm: Almqvist & Wiksell.
- Manova, Kalina, and Zhiwei Zhang. 2011. "Export Prices across Firms and Destinations." *Q.J.E.*, forthcoming.
- Markusen, James R. 1986. "Explaining the Volume of Trade: An Eclectic Approach." *A.E.R.* 76 (5): 1002–11.
- Matsuyama, Kiminori. 2000. "A Ricardian Model with a Continuum of Goods under Nonhomothetic Preferences: Demand Complementarities, Income Distribution, and North-South Trade." *J.P.E.* 108 (6): 1093–1120.
- McFadden, Daniel. 1978. "Modelling the Choice of Residential Location." In *Spatial Interaction Theory and Planning Models*, edited by A. Karlqvist, L. Lundqvist, F. Snickars, and J. Weibull, 75–96. Amsterdam: North-Holland.
- Mitra, Devashish, and Vitor Trindade. 2005. "Inequality and Trade." *Canadian J. Econ.* 38 (4): 1253–71.
- Murphy, Kevin M., and Andrei Shleifer. 1997. "Quality and Trade." *J. Development Econ.* 53 (1): 1–15.
- Nevo, Aviv. 2011. "Empirical Models of Consumer Behavior." *Ann. Rev. Econ.* 3: 51–75.
- Porto, Guido G. 2006. "Using Survey Data to Assess the Distributional Effects of Trade Policy." *J. Internat. Econ.* 70 (1): 140–60.
- Rossi-Hansberg, Esteban. 2005. "A Spatial Theory of Trade." *A.E.R.* 95 (5): 1461–91.
- Schott, Peter K. 2004. "Across-Product versus Within-Product Specialization in International Trade." *Q.J.E.* 119 (2): 647–78.
- Shaked, Avner, and John Sutton. 1982. "Relaxing Price Competition through Product Differentiation." *Rev. Econ. Studies* 49 (1): 3–14.
- . 1983. "Natural Oligopolies." *Econometrica* 51 (5): 1469–83.
- Stokey, Nancy L. 1991. "The Volume and Composition of Trade between Rich and Poor Countries." *Rev. Econ. Studies* 58 (1): 63–80.
- Train, Kenneth E. 2003. *Discrete Choice Methods with Simulation*. New York: Cambridge Univ. Press.
- Verboven, Frank. 1996a. "International Price Discrimination in the European Car Market." *RAND J. Econ.* 27 (2): 240–68.
- . 1996b. "The Nested Logit Model and Representative Consumer Theory." *Econ. Letters* 50 (1): 57–63.
- Verhoogen, Eric. 2008. "Trade, Quality Upgrading, and Wage Inequality in the Mexican Manufacturing Sector." *Q.J.E.* 123 (2): 489–530.