

Explaining Import Quality: the Role of the Income Distribution

Yo Chul Choi^a
David Hummels^b
Chong Xiang^b

^aThe Bank of Korea, South Korea
^bPurdue University, the United States

January 2009

Abstract: We examine a generalized version of Flam and Helpman (1987) in which consumption prices for quality differentiated goods are rising in household income. We provide propositions for aggregating this relationship across heterogeneous households to map cross-country differences in income distributions to cross-country differences in import price distributions. The theoretical predictions are examined and confirmed using disaggregated data on prices of traded goods and micro data on household income from the Luxemburg Income Study. Country pairs with more similar income distributions have more similar import price distributions, whether similarity is measured by 1st - 4th moment statistics, population and consumption shares within world income and product price quantiles, or income and price dissimilarity indices.

Keywords: quality differentiation; income distribution; import price distribution; Luxemburg Income Study.

JEL classification: F1, D3.

Acknowledgements: Yo Chul Choi, Economic Activities Analysis Team, Research Dept., the Bank of Korea, 110, 3-Ga, Namdaemun-Ro, Jung-Gu, Seoul, Korea. Email: cyc@bok.or.kr. Tel.: +82-2-759-4209. Fax: +82-2-759-4418. David Hummels (corresponding author), Economics Dept., Purdue University, 403 W. State Street, IN 47907, U.S.A. Email: hummelsd@purdue.edu. Tel.: 1-765-494-4495. Fax: 1-765-496-7434. Chong Xiang, Economics Dept., Purdue University, 403 W. State Street, IN 47907, U.S.A. Email: cxiang@purdue.edu. Tel.: 1-765-494-4499. Fax: 1-765-494-9658. The authors thank the editor and two anonymous referees for extremely helpful comments. The authors would also like to thank Purdue CIBER and the NSF for financial support and seminar participants at the Midwest International Meetings, the 2006 EIT Meetings, NBER 2006 Winter Meetings, the 2007 Econometric Society Winter Meetings, the Federal Reserve Boards of Governors and Purdue, Dartmouth, Toronto, Texas, Wisconsin, Oregon and Oregon State Universities for helpful comments.

1. Introduction

There is a large literature examining how international trade affects a nation's income distribution, but there is relatively little empirical work examining the reverse channel. This is in large part because trade models commonly rule out income effects in order to focus attention on supply considerations such as factor endowments or scale economies. To the extent that richer demand structures with non-homothetic preferences are employed they operate at the level of broad industries, for example, allowing poor countries to devote relatively large income shares to commodity foodstuffs.¹ In this paper we investigate how the distribution of income within and across countries shapes patterns of consumption and international trade in quality differentiated varieties within narrow product categories.

Our starting point is Flam and Helpman's (1987) model of quality differentiation in trade and we focus on the model's demand side implications linking consumer income to quality choice. As in Flam and Helpman (1987), goods can be quality differentiated at some cost so that higher prices reflect higher quality, and consumers use marginal income to buy higher qualities rather than higher quantities of a differentiated good. This provides an equilibrium mapping in which prices of goods consumed are rising in household income.

This prediction is consistent with household evidence on consumer durables purchases. Bils and Klenow (2001) use survey data for the US that reports household income and purchase prices and estimate positive price-income slopes (or, "Quality Engel Curves"). Our interest lies in cross-country comparisons where household consumption choices are unobservable. We show that the model can be written in terms of national income and price distributions which are, with some effort, observable. This requires aggregating heterogeneous household income and

¹ See Markusen, 1986; Hunter, 1991; Mitra and Trindade, 2005; Reimer, 2005; Dalgin, Mitra and Trindade, 2008.

consumption decisions into national income and price distributions and providing propositions linking the two.

There are many ways to empirically characterize a national price or income distribution. However, the particular statistics used to appropriately characterize these distributions and the linkage between them depends on the functional relationship between income and product prices at the household level.

Assuming identical technologies we derive a linear mapping between income and consumption prices at the household level. When aggregating, linearity preserves a tight positive correlation between the corresponding moment statistics (e.g. mean, standard deviation, skewness and kurtosis) of national price and income distributions. That is, countries whose income distributions exhibit higher means (or standard deviation, skewness, kurtosis) will have product price distributions with higher means (or standard deviation, skewness, kurtosis).

We also extend Flam and Helpman (1987) to the case of multiple differentiated goods and multiple countries with different technologies. In this case the mapping between income and consumption prices at the household level is monotonic, but not linear. As a consequence, there is no longer a clear-cut relationship between the corresponding moment conditions of national income and price distributions. (A country with higher mean income could have lower mean product prices.) However, we can still establish a linkage between national income and price distributions through probability and cumulative distribution functions. Country pairs with more similar population shares in a given world income quantile will have more similar consumption shares in the corresponding world price quantile. When we examine the probability distribution functions along the entire support, we find that country pairs with more similar income distributions have more similar product price distributions.

To examine these predictions we employ internationally comparable household income data from the Luxembourg Income Study (LIS) for 26 countries in the year 2000. The LIS provides us with household income data at percentile increments from which we construct income distributions for cross-country comparisons. We construct our price distributions using international trade data at the level of 6-digit Harmonized System (HS) products from COMTRADE for 1999-2001. Previous authors have shown that prices vary substantially across exporters and covary with exporter characteristics such as per-capita income and per worker supplies of capital and skilled labor (Schott, 2004; Hummels and Klenow, 2005). Further, countries with high export prices have larger, not smaller, shares of the markets in which they sell (Hallak, 2006). These facts point to the primacy of quality differentiation, as opposed to measurement error, as an explanation for measured price variation. For each product we observe from which exporters an importer buys, along with each exporter's quantity and value of trade, and from these construct price (= unit value) distributions for each importer and product.

Our findings are consistent with the model. The differences in importers' price distributions are correlated with differences in their income distributions, both within individual quantiles of the distributions and along the whole support. In addition, key statistics related to the 1st - 4th moments (mean, median, standard deviation, coefficient of variation, inter-decile range, skewness and kurtosis) of the income distribution are positively correlated with those of the import price distribution. In other words, countries with high incomes consume goods with high prices; countries with a greater variability in incomes over households have greater variability in prices for a particular good; and countries whose income distributions have fat or skewed tails also have price distributions with fat or skewed tails.

Our work relates most closely to a relatively new literature on the role of quality differentiation in trade. Most of the empirical work in this literature has focused on linking price variation to exporter characteristics. Some authors have provided correlations with importer characteristics, showing that within product categories, countries with high mean income per capita buy goods with higher mean prices (Hallak, 2006; Hummels and Skiba, 2004). We differ in that we provide explicit propositions showing how a correlation between income and prices at the level of heterogeneous households will aggregate to correlations between national income and price distributions. This allows us to demonstrate which are the appropriate statistics to use in linking these distributions, and to examine higher moments of income and price distributions, their individual quantiles, and differences along the entire support.

The paper proceeds as follows. Section 2 provides the theory linking a country's income and import price distributions. Section 3 discusses our empirical specifications. Section 4 explains the construction of our income and price distribution data in detail. Section 5 presents the empirical results and Section 6 concludes.

2 The Model

Flam and Helpman (1987) provide a model in which heterogeneity in household income is mapped into heterogeneity in optimal quality choice. We extend their model to a multi-country, multi-good setting, with an analysis motivated by and focused on empirical feasibility. In an international context we are unable to observe the qualities and prices of goods consumed at the household level. However, we can observe a country's income distribution, as well as the distribution of prices for imported goods. We show how to aggregate heterogeneous household income and consumption decisions to construct national income and price distributions and

provide propositions linking differences in importers' income distributions to differences in import price distributions.

We start with a model with one differentiated good and identical technologies across countries. This model's restrictive assumptions on technology and preferences imply a linear monotonic relationship between household income and the price that household will pay for a quality differentiated product. Given linearity we show that, across countries, each moment of the consumption price distribution is positively correlated with the corresponding moment of the income distribution. We then extend our analyses to multiple goods and different technologies so that the price-income relationship is monotonic but non-linear. In this case, there need not be a cross-country correlation between the corresponding moments of the price and income distributions. However, we show that there is a relationship between cross-country differences in probability distribution functions (pdf's) for prices and income. This relationship holds when examining the integral of the pdf along the entire support, or when examining sub-sets of the support, i.e. bins corresponding to "high" and "low" priced goods. To save space we post rigorous proofs of all propositions on our websites.²

2.1 Identical Technologies and One Differentiated Good

There are two goods, a homogeneous numeraire good and a vertically differentiated good. There are C countries. Each country c has population N_c , with income I distributed exogenously³ according to the pdf $g_c(\cdot)$ and cumulative distribution function (cdf) $G_c(\cdot)$ with support S_G^c .

² See links at <http://www.mgmt.purdue.edu/faculty/hummelsd/research.htm>

³ This assumption allows us to focus on the role of national and world income distributions in determining quality demand, but we abstract from the feedback channels through which trade affects income, as in Flam and Helpman's (1987) seminal work.

A consumer with income I chooses quantities of the numeraire, y , and the desired quality, $z \in [0, 1]$, of a single unit of the differentiated good in order to maximize

$$u(y, z) = ye^{\alpha z} \quad s.t. \quad y + p(z) \leq I, \quad (1)$$

where $\alpha > 0$, αz is the elasticity of utility with respect to quality, $p(z)$ is the price of the differentiated good with quality z , and the price of the numeraire is set to 1. We assume that income is sufficiently high so that every consumer consumes the differentiated good.

We initially assume that all countries produce using an identical technology. The marginal cost of producing quality is

$$MC(z) = e^{\gamma z} w. \quad (2)$$

w represents the cost component that is common to all the quality levels. $e^{\gamma z}$ represents the cost component that is unique to quality z and implies that the marginal cost increases exponentially with z . γz is the elasticity of the marginal cost with respect to quality. We assume that there are no trade costs, and that there are perfectly competitive markets at each quality level so that consumers in all countries face the same vector of prices $p(z) = MC(z)$. Solving the utility maximization problem we obtain,

$$(3) \quad z = \frac{1}{\gamma} \left[\log \frac{\alpha}{\alpha + \gamma} + \log I - \log w \right]$$

$$(4) \quad p(z) = aI, \text{ where } a = \frac{\alpha}{\alpha + \gamma}.$$

Equation (4) indicates that a consumer with income I spends a fixed fraction $a = \frac{\alpha}{\alpha + \gamma}$ of his income on (one unit of) quality z . Let $f_p(\cdot)$ denote the mapping from income into price in equation (4), which in this restrictive case is monotonic and linear.

The linear monotonicity of the household consumption decision allows us to aggregate household decisions into a relationship between the distribution of household income for a country and its distribution of consumption prices for the differentiated good. Let $h_c(\cdot)$, $H_c(\cdot)$ and S_H^c denote the pdf, cdf and support of the price distribution of the differentiated good for country c . We can use (4) to rewrite income as an inverse function of prices, or $I = p/a$. Then we have the pdf of the price distribution⁴

$$(5) \quad h_c(p(z)) = g_c\left(\frac{p(z)}{a}\right) \cdot \frac{1}{a} \quad \text{with support } S_H^c = f_p(S_G^c).$$

Equation (5) says that the price distribution is a linear mapping from the income distribution. Because the linear mapping in (5) preserves the ranking of moments we can correlate the moments of the price distribution with the moments of the income distribution. For example, suppose income is distributed log normally $L(\mu, \sigma^2)$. Then prices are also distributed log normally $L(\mu a, \sigma^2 a^2)$. The mean and variance of the price distribution are proportional to the mean and variance of the income distribution, respectively. We can also correlate higher moments of the empirical distributions of prices with those of the income distributions.

Proposition 1 Denote $M^m(P_c)$ as some moment statistic of the product price distribution for country c and $M^m(I_c)$ as the corresponding moment statistic for c 's income distribution. . Then for the statistics mean, median, standard deviation, and inter-decile range, $M^m(P_c) = aM^m(I_c)$.

For the statistics coefficient of variation, skewness and kurtosis, $M^m(P_c) = M^m(I_c)$

The direct mapping between moments of the price and income distributions is handy for

⁴ For example, if $S_G^c = [0, b]$, then $S_H^c = [0, ab]$. If $f_p(I)$ takes a more general form than aI , then equation (5) becomes $h(p(z)) = g\left(f_p^{-1}(p(z))\right) \cdot [f_p^{-1}(p(z))]'$ with support $S_H^c = f_p(S_G^c)$, provided that $f_p(I)$ is strictly increasing in income and its inverse exists and is differentiable.

empirical implementation but requires the linear monotonicity of equation (4). Suppose the relationship between income and prices were monotonic, but not necessarily linear. Monotonicity implies that for each quality z^* there is some income level $I(z^*)$ for which z^* is the optimal quality. If there is no mass in the income distribution at $I(z^*)$, then z^* is not produced or consumed in equilibrium, but for every $I(z^*)$ with positive mass, the quality z^* will be produced and consumed.⁵ Further, the share of the population in country c consuming z^* (at price p^*) is equal to the share of the population with income $I(z^*)$. This allows us to relate the cdf of the price distribution to the cdf of the income distribution.

Again, denote the monotonic mapping between prices and income as $f_p(I^*) = p^*$. The aggregate consumption share of the goods whose prices fall into the bin $[f_p(I_1^*), f_p(I_2^*)]$ equals the aggregate population share of the consumers whose income falls into the bin $[I_1^*, I_2^*]$. This does not require linearity between prices and income, only that all countries face the same monotonic price-income mapping. As a result,

Proposition 2 $G_c(I_2^*) - G_c(I_1^*) = H_c[f_p(I_2^*)] - H_c[f_p(I_1^*)]$ for all countries c and for all $I_1^*, I_2^* \in S_G^c$ and $I_1^* < I_2^*$.

For example, suppose the bin $[I_1^*, I_2^*]$ represents the top decile of worldwide income, so that $[f_p(I_1^*), f_p(I_2^*)]$ is the top decile of worldwide prices for the good. Proposition 2 says that if $G_c(I_2^*) - G_c(I_1^*) = 5\%$ of country c 's population has income in the top world decile then $H_c[f_p(I_2^*)] - H_c[f_p(I_1^*)] = 5\%$ of country c purchases of the quality-differentiated good will lie in the top decile of worldwide prices. Proposition 2 also allows us to compare the population

⁵ This is an implication of assuming no fixed costs of production and perfectly competitive markets.

and consumption shares between countries c and c' . The difference in their population shares in the top decile of the worldwide income should be equal to the difference in their consumption shares in the top decile of worldwide prices, and similarly for other income and price bins.

Therefore,

Corollary 1

$[G_c(I_2^*) - G_c(I_1^*)] - [G_{c'}(I_2^*) - G_{c'}(I_1^*)] = (H_c[f_p(I_2^*)] - H_c[f_p(I_1^*)]) - (H_{c'}[f_p(I_2^*)] - H_{c'}[f_p(I_1^*)])$ for all countries c, c' and for all $I_1^* < I_2^*$.

Proposition 2 shows how the price distribution of the differentiated good consumed in country c is mapped to its income distribution, segment by segment, and Corollary 1 shows how the cross-country differences in the distribution of income are reflected in the differences in the distribution of prices, segment by segment. We can also measure cross-country differences in income and price distributions for the whole distribution using the following dis-similarity index.⁶

Definition 1 Dis-similarity Index (DSI): The DSI for the pair of distributions with pdf's $r_1(\cdot)$

and $r_2(\cdot)$ and supports S_1 and S_2 is $DSI(r_1, r_2) \equiv \frac{1}{2} \int_{\underline{S}} |r_1(x) - r_2(x)| dx$, where $\underline{S} = S_1 \cup S_2$, $r_1(\cdot)$

is defined to be 0 for $\underline{S} - S_1$ and $r_2(\cdot)$ defined to be 0 for $\underline{S} - S_2$.

The DSI quantifies the difference between $r_1(\cdot)$ and $r_2(\cdot)$ by calculating the vertical distance between them at every point x and then aggregating these vertical distances. If $r_1(\cdot)$ and $r_2(\cdot)$ are dis-similar, they lie far away from each other, the vertical distances between them are large and so $DSI(r_1, r_2)$ is large. Because both $r_1(\cdot)$ and $r_2(\cdot)$ are pdf's, $DSI(r_1, r_2)$ exists and is bounded

⁶ The DSI is half the L_1 distance between the pdf's $r_1(\cdot)$ and $r_2(\cdot)$. Another commonly used distance metric is the $L_2 \int_{\underline{S}} [r_1(x) - r_2(x)]^2 dx$. We have chosen the L_1 metric because it enables our DSI index to fall between 0 and 1.

between 0 and 1. We can calculate a dissimilarity index for the price distributions of a country pair (PDSI), as well as a DSI for the income distributions of a country pair (IDSI). We show that

Proposition 3 $PDSI(h_c, h_{c'}) = IDSI(g_c, g_{c'})$ where c and c' represent any country pair.

The intuition for Proposition 3 is similar to Proposition 2, except that we examine differences in consumption shares at each point along the support rather than for some discrete portion of it.

Finally, the results in this section extend to a case with multiple quality differentiated goods indexed by k , where utility is $u = y \exp(\sum_k \alpha_k z_k)$ and marginal production costs are given by $MC(z_k) = \exp(\gamma_k z) w$. In this case, the price-income slope is product-specific and given by

$$(6) \quad p_k(z) = a_k I, a_k = \frac{\alpha_k / \gamma_k}{1 + \sum_{l=1}^K \alpha_l / \gamma_l} \text{ for all } k.$$

While this mapping is different for each good, it is also linear and monotonic. As a result, Proposition 1 holds with good-specific constants a_k replacing the constant a , and Propositions 2 and 3 hold in precisely the same way for each good.

2.2 Multiple Countries with Different Technologies

Now we allow technologies to differ across supplying countries. For notational simplicity we return to the one differentiated-good case. Our results easily extend to the case of multiple differentiated goods. Let $j = 1 \dots C$ index supplying countries. The marginal cost of producing quality z in country j is

$$(7) \quad MC_j(z) = \exp(\gamma_j z) w_j.$$

w_j represents the cost differences (due to factor price or Ricardian technology differences) that are common to all quality levels. $\exp(\gamma_j z)$ expresses the degree to which country j has a comparative advantage in high or low qualities. We continue to assume that there are no trade

costs, so that the consumers desiring quality z buy it from the lowest cost provider. Suppose that each country j is the lowest cost supplier of some set of qualities Z_j .

Since different suppliers have a comparative advantage in different ranges of quality, there are kinks in the budget set and this creates a discontinuous relationship between prices and income. We illustrate this point using the two-country setting of Flam and Helpman (1987), where the North and the South have technologies given in equation (7). Let $\gamma_N < \gamma_S$ and $w_S < w_N$ so that the North has the comparative advantage in high qualities. The South then specializes in the low qualities $[0, z_1]$ and the North specializes in the high qualities $[z_2, 1]$ with $z_1 < z_2$. There also exists an income level I_d such that a consumer with income I_d is indifferent between buying the differentiated good from the North with quality z_2 and buying it from the South with quality z_1 .⁷

The price-income mapping is now $p(z) = \alpha I / (\alpha + \gamma_j)$, $j = N, S$. It differs from equation (4) in two aspects. First, there is no demand, in the North or the South, for the qualities between z_1 and z_2 . Second, for the qualities that are actually supplied to the market, $[0, z_1] \cup [z_2, 1]$, the price-income mapping consists of two segments: for the quality range $[0, z_1]$ it is determined by Southern technology, γ_S , and for the range $[z_2, 1]$ it is determined by Northern technology, γ_N .

We now go back to the general case of multiple supplier countries. Let S_G^j be the set of income for which some qualities in the set Z_j are the optimal quality choice. Then consumers with $I \in S_G^j$ buy the differentiated good from exporter j (S_G^j is not the support of a country's income distribution). Since every consumer buys the differentiated good from somewhere,

⁷ See Flam and Helpman (1987) or our websites for a graphical illustration.

$\cup_j S_G^j = S_G^w$, where S_G^w is the support of the world income distribution. The price-income mapping is

$$(8) \quad p(z) = a_j I \quad \text{for } I \in S_G^j, \text{ where } a_j = \frac{\alpha}{\alpha + \gamma_j}, \text{ for all } j.$$

Rather than a single line of constant slope mapping incomes into prices as in equation (4), we now have a set of lines whose slopes are determined by the technology of the lowest cost producer for the corresponding quality segment.

The price-income mapping in (8) is non-linear and so may not preserve the correlation of the price and income moments. For example, suppose the price income mapping from equation (8) has two line segments with $p(z) = 0.1I$ for $I \leq 1$ and $p(z) = 0.3I$ for $I > 1$. Country 1's population is split evenly between households earning \$1 (buying goods with price \$0.1) and households earning \$2 (buying goods with price \$0.6). Country 2's population is split evenly between households with income \$0.8 (goods price \$0.08) and households with income \$2.1 (goods price \$0.63). The mean income for country 1 is higher (\$1.5 > \$1.45) but its mean price is lower (\$0.35 < \$0.355). This simple example illustrates a general point: absent a linear relationship between income and prices at the household level there is no longer a well-defined relationship between the moments of national price and income distributions as in Proposition 1.

However, so long as the price-income relationship at the household level is monotonic Propositions 2 and 3 remain intact. Consider first the “holes” in the price distribution, those qualities that are not demanded by any income level and so not demanded by any country. These qualities carry zero weight in the price distribution and so have zero effect on consumption shares or their cross-country differences. Next consider the qualities actually supplied in the world. These segments have different price-income slopes given by (8), but along each segment

prices are still a one to one mapping from income and the same price-income mapping still applies to every country. As a result, the consumption share for the good with price $p(z^*)$ in country c equals the population share at income $I(z^*)$, and the difference in consumption shares of $p(z^*)$ for countries c and c' equals the difference in their population shares at income $I(z^*)$.

Therefore,

Proposition 4 When there are multiple supplier countries with different technologies, Propositions 2 and 3 hold.

Proof: see the Appendix.

3 Empirical Specification

The core idea of our model is that a household level prediction linking income to consumer product prices can be aggregated up to a linkage between national distributions of income and consumer product prices. However, the particular statistics used to appropriately characterize these national distributions and the linkage between them depends on the functional relationship between income and product prices at the household level. We consider two cases.

In the first case, which is examined in section 2.1, equation (4) predicts a linear monotonic relationship between income and prices. Proposition 1 indicates that each moment of the price distribution is then equal to the corresponding moment of the income distribution multiplied by a product specific constant. This constant is equal to a_k , the slope of the price-income relationship given by equation (6), or equal to 1, depending on the statistic. Let $M^m(P_{c,k})$ be a generic statistic representing a moment of the price distribution for product k in country c . $M^m(I_c)$ is the corresponding moment statistic of the income distribution in country c . We take logs of the moment statistics for country c and then calculate their differences relative to country c' to derive

the following estimating equation:

$$(9) \quad \ln M^m(P_{c,k}) - \ln M^m(P_{c',k}) = \alpha_0 + \beta_M \left(\ln M^m(I_c) - \ln M^m(I_{c'}) \right) + e_{cc',k}$$

We examine the following specific statistics: mean, median, standard deviation, coefficient of variation, inter-decile range, skewness, and kurtosis. To further address concerns that the relationship between price and income distributions is driven primarily by correlations in the first moment we also estimate (9) for the higher moment statistics using a data sub-sample of countries with similar first moments of income (those with mean per capita incomes within 25% of one another).

A second method for empirically characterizing the distributions requires only that the household relationship between income and prices is monotonic. It therefore applies to the linear case of section 2.1, as well as to the non-linear (and non-continuous) case of section 2.2. In this approach we characterize the income and price distributions for country c using bins defined by worldwide data on product prices and income.

Let $S^b(I_c)$ denote the share of country c 's population that falls within income bin b . For example, the bin $b=1$ represents incomes falling within the first decile or quartile of the world income distribution. Likewise, let $S^b(P_{c,k})$ denote the share of country c 's consumption of product k that falls within price bin b . By definition the shares sum to one over the B bins, $\sum_{b=1}^B S(P_{c,k})^b = 1$ and $\sum_{b=1}^B S(I_c)^b = 1$. Corollary 1 of Proposition 2 states that, for each product, the difference in bin b consumption shares for countries c and c' is equal to the difference in bin b income shares for countries c and c' . We thus estimate

$$(10) \quad S^b(P_{c,k}) - S^b(P_{c',k}) = \alpha_0 + \beta_B \left[S^b(I_c) - S^b(I_{c'}) \right] + e_{c,c',k}$$

We estimate (10) separately for each bin b and we pool over all products k in the estimation.

We consider two cases, B=4 (quartiles) and B=10 (deciles). We experiment with defining the comparison country c' as the whole world or as an individual country. In the latter case we take pairwise combinations of countries c and c' in our data. We also estimate (10) using aggregated data. Let $S^b(P_c)$ denote the share of country c 's aggregate consumption (across all products) that falls into price bin b (where the price cutoffs for each bin are still specific to each product k). The resulting estimating equation is

$$(11) \quad S^b(P_c) - S^b(P_{c'}) = \alpha_0 + \beta_B [S^b(I_c) - S^b(I_{c'})] + e_{c,c'}$$

Regression (11) is particularly useful in cases where we see only a few price data points for a c - k pair so that the product level distribution is lumpy.

Regressions (10)-(11) examine consumption and income shares bin by bin. They allow us to ask, for example, whether the consumption share of goods in the first world price decile is high for countries with a large population share in the first world income decile. We can also examine the difference in income and price distributions along the whole support using the income and price dis-similarity indices (IDSI and PDSI) as Proposition 3 suggests. We first consider a discrete measure of IDSI and PDSI for a country pair c, c' by aggregating over the bins in regression (10).

$$(12) \quad \sum_{b=1}^B |S^b(P_{c,k}) - S^b(P_{c',k})| = \alpha_0 + \beta_P \sum_{b=1}^B |S^b(I_c) - S^b(I_{c'})| + e_{c,c'}$$

In regression (12) we can have c' as another country or the whole world. In the latter case, the consumption and income shares of c' , $S^b(P_{c',k})$ and $S^b(I_{c'})$, equal $1/B$, where B is the number of bins. We next consider the continuous measure of IDSI and PDSI.

$$(13) \quad PDSI(h_c, h_{c'}) = \alpha_0 + \beta_P IDSI(g_c, g_{c'}) + e_{cc'}$$

As compared with regressions (10)-(11), regressions (12) and (13) characterize the difference in

distributions over the entire support. However, (12) and (13) may fail to distinguish the direction of the difference. Suppose the income distributions for countries 1, 2, and 3 are $N(\mu-d, \sigma^2)$, $N(\mu, \sigma^2)$, and $N(\mu+d, \sigma^2)$. Then $\text{IDSI}(1,2) = \text{IDSI}(2,3)$. In contrast, regressions (9)-(11) match moments of price and income distributions or match consumption and population shares by bin and so they clearly distinguish the direction as well as the magnitude of the difference in distribution.

4 Data

4.1 Income Data

To construct cross-country comparisons of income distributions, we employ wave 5 (year 2000) of the Luxembourg Income Study (LIS) data.⁸ When matched to our price data we have 26 countries, listed in Table 1. The LIS data are a compilation of national income survey data files, made comparable by rearranging or reclassifying the income measures from national household budget surveys. For each country LIS provides disposable household income (monetary income after direct taxes and transfer payments) at percentile increments, and it also allows us to make adjustments to account for differences in family size.⁹ Since the household income data are in local currency units we convert them to U.S. dollars using current exchange rates.

For several exercises we compare a country's income distribution to the world. Henceforth the “world” consists of the 26 LIS countries. To construct the world income distribution we begin with the country distributions, then weight them by the corresponding populations and

⁸ In previous drafts of this paper we employed LIS data for waves 1-5 for the period 1979-2000. In this draft we employ much better price data, but this constrains us to using a single wave of income data.

⁹ See the Data Appendix for details, including a discussion of the advantages of these data relative to other cross-country data on income distribution.

aggregate.¹⁰ We can then identify income cutoff points corresponding to world income quartiles or deciles and calculate the share of each country's population within each quartile or decile.

Table 1 provides income moment statistics for each of our LIS countries along with the population share for each country within each worldwide income quartile. Figure 1 illustrates the income dispersion for the LIS countries. We normalize each country's income relative to the US median income (\$23,916 = 100) and plot the range of income starting at the 10th percentile (P10) and ending at the 90th percentile (P90). We arrange the countries in ascending order of their decile ratios (P90/P10).

Our implementation of equation (13) requires us to construct and then compare income pdf's across countries. To construct a continuous income distribution from the discrete household income data we perform a non-parametric kernel estimation using the "kdensity" command in STATA (Deaton 1997). We use STATA's default kernel, the Epanechnikov, and STATA's default bandwidth,¹¹ and evaluate the densities of the distributions of all the countries at the same income levels of \$100, \$200 ... \$150,000. We then calculate the differences in income distributions, both pair wise and relative to the world. IDSI achieves its maximal value of 1 if two countries have completely disjoint distributions and achieves its minimal value of 0 if two countries have identical distributions.

4.2 Data on Import Prices and Their Distributions

Our theory refers to the distribution of consumption prices for all goods consumed in the economy. Unfortunately, it is not possible to get cross-country data on all sales prices for

¹⁰ E.g. In 2000 the US and Canadian populations were 275.4 million and 30.7 million. When constructing the world income distribution, each percentile of the US distribution is given a population weight of 2.754 million and each percentile of the Canadian population is given a population weight of 0.307 million.

¹¹ The choice of kernel tends to be relatively unimportant in practice (e.g., DiNardo and Tobias 2001) and STATA's default bandwidth is based on Silverman (1986)'s optimal bandwidth.

narrowly defined consumer goods; i.e. we are unable to observe domestic purchases of domestically produced goods.¹² Instead, we use import price data to approximate the price distributions. Given the technology assumed in equation (7) each exporting country j will specialize in a range of qualities, with a corresponding range of export prices. By knowing the prices and import shares for each importer c , exporter j , product k transaction, we can calculate the desired moments of the price distributions for equation (9) as well as the consumption shares for each price bin b in equations (10), (11) and (12), and the price kernel densities for equation (13).

We employ bilateral trade value and quantities at the six digit level of the Harmonized System taken from the COMTRADE data. We have 26 importers' purchases from all exporters worldwide in 4759 HS 6 products. We use the value and quantity data to construct unit values (henceforth, prices) as value/quantity. Unit values can be noisy and so we employ the following screening procedures.

1. We employ pooled data for years 1999-2001.¹³
2. We employ only those observations for which we observe the quantity units, and for which these quantity units are measured in kilograms.¹⁴ This eliminates 15.8 % of our observations.
3. We discard the observations for which quantity = 1, as well as those for which price is either less than 10% of the worldwide median price for that commodity or more than 10 times the median price. This eliminates another 5.7 % of our observations.

After these data screens we have 4.99 million price and quantity observations from which we

¹² In the results section we remark on how the absence of domestic sales data will affect our estimates.

¹³ If exporter j sells the same product k to importer c in each year we use this as 3 separate observations. Results are similar when averaging over multiple observations for a single exporter, or when using a single year of data.

¹⁴ Note that a kilogram may not be an especially meaningful quantity unit for cross-product comparison. However, since all our statistics involve within-product calculations we can think of these as subsuming a product specific conversion of kilograms to a meaningful quantity unit.

calculate various price distribution statistics for each importer c - HS6 product k combination. For the median c - k we have 61 transactions involving 27 unique exporters. Roughly 10 percent of our c - k observations involve 10 or fewer transactions, which might be especially problematic for calculating higher order statistics. Dropping these observations has little effect on results so we employ the entire sample.

5. Results and Discussions

5.1 Results

The top panel of Table 2 reports the results of estimating equation (9) using the logs of each of the following seven moment statistics: mean, median, standard deviation, coefficient of variation, inter-decile range (the 90th percentile minus the 10th percentile), skewness and kurtosis. We pool over all HS 6 products and report standard errors clustered on the country pair. The coefficient β_M is positive and highly significant for all six statistics, though less than the unitary elasticity implied by the theory. In other words, countries with high incomes consume goods with high prices; countries with a greater variability in incomes over households have greater variability in prices for a particular good; and countries whose income distributions have fat or skewed tails also have price distributions with fat or skewed tails.

One might be concerned that the primary relationship is between the first moments and that this is also driving the correlation in the higher moments.¹⁵ If this were the case, we should get weaker results for the coefficient of variation, which equals the standard deviation normalized by the mean. However, our result for the coefficient of variation ($\beta_M = 0.38$) is at least as strong as those for the mean ($\beta_M = 0.138$) and standard deviation ($\beta_M = 0.131$). In addition, we also find

¹⁵ The sample correlation between income means and other moments is: st. dev (.86), range (.92), skewness (.12) and kurtosis (.16).

positive correlations for skewness ($\beta_M = 0.146$) and kurtosis ($\beta_M = 0.108$), whose sample correlation with income means is statistically insignificant. These strongly suggest that we have uncovered much more than just first-moment correlations in the data.

To further address this concern we restrict our sample to the 95 country pairs with similar first moments (income means within 25% of one another) and report estimates in the bottom panel of Table 2. In this restricted sample, the higher moments of the price and income distributions remain significantly positively correlated with one another even though the first moment correlations (mean and median) are no longer significant.

In Table 3 we report the following robustness checks for regression (9): (a) eliminating two countries (France and Australia) with imputed income data;¹⁶ (b) eliminating homogeneous commodities (primarily agriculture and mining) in order to focus on goods for which our theory of vertical differentiation may be more plausible; and (c) eliminating textile and apparel products to avoid the distortionary effects of the MFA from contaminating our sample. In all cases the coefficients are quite similar to the full sample estimates from Table 2.

In Tables 4-5 we provide estimates of equations (10) and (11), regressing cross-country differences in consumption shares for price bin b on cross-country differences in population shares for income bin b . We experiment with using individual importers as the comparison country c' as well as using the world as the comparison country, and using bin shares calculated both at the product level (10) and after aggregating over products to arrive at a single bin share for each importer (11).

In Table 4 we define bins as quartiles of the world income and product price distributions. We estimate separate regressions for each bin and (for the product level bins) report clustered

¹⁶ See the Data Appendix for the details of the imputation.

standard errors. The coefficient β_B is positive and significant in all but two cases. To better illustrate the results in Table 4 we use the aggregated data to plot the differences in consumption shares against the difference in income shares by quartile in Figure 2. These two variables show a clear positive correlation that is stronger for the 1st and 4th quartiles than for the 2nd and 3rd quartiles. Consistent with our theory, having a large population share in income quartile b results in having a large consumption share in the corresponding price quartile b , with the effects being most pronounced at the high and low end.

Table 5 is organized similarly to Table 4, except that we organize the data in income and price deciles. All of the coefficients are positive, though some coefficients in the middle deciles are not significant.

Tables 4 and 5 also show that β_B is much larger at either end of the income and price distributions than in the middle of these distributions. With four bins the estimates of β_B for the 1st and 4th quartiles range from 0.2 to 0.4. In other words, if the population share in the 1st or 4th world income quartile increases by 10 percentage points, the consumption share in the same world price quartile will increase by 2 to 4 percentage points. In contrast, the β_B estimates for the 2nd and 3rd quartiles range from 0.058 to 0.127. With 10 bins β_B reaches 0.696 for the 10th decile, roughly one standard error away from the unitary value predicted by theory. However, β_B ranges from 0.01 to 0.130 for the 3rd - 7th deciles.

The last columns of Tables 4 and 5 report the results of estimating regression (12), the discrete version of the PDSI regression. When we use individual importers as the comparison country c' , the coefficient β_p is positive and significant in all cases and it ranges from 0.088 (product data with ten bins) to 0.144 (aggregate data with ten bins). When we use the world as

the comparison country β_p is positive and significant with product data, ranging from 0.227 (with ten bins) to 0.267 (with four bins), but it is negative and insignificant with aggregate data. For regression (13), the continuous version of the PDSI regression, we obtain $\beta_p = 0.071$ (significant) when we compare with individual importers and $\beta_p = 0.059$ (insignificant) when we compare with the world. Consistent with our theory, the country pairs with similar income distributions have similar price distributions, though this relationship appears weaker when estimated using the continuous distributions rather than the discrete bin approach.

Finally, in Table 6 we conduct the same robustness exercises as in Table 3 for the bin and PDSI regressions: (a) eliminating France and Australia; (b) dropping homogeneous goods; and (c) dropping textile and apparel products. We report the results with specifications (11) and (12) (the results for the other cases are similar). In all cases the coefficients are quite similar to the full sample estimates from Table 4.

5.2 Discussions

Our results are generally supportive of the theory. The 1st-4th moments of the product price distribution for a country are correlated with the corresponding moment of the income distribution, and consumption shares by price bin are correlated with income shares in the corresponding income bin. However, in each case the coefficients are lower than the unitary value predicted by the theory. Why is this?

For the moment regression (9) the theoretical prediction of $\beta_M = 1$ holds only under the restrictive setting of identical production technologies in which the price-income mapping at the household level is linear. As we discussed in section 2.2 $\beta_M = 1$ may not hold under the more general setting of different production technologies. In fact, it is easy to construct examples

where the coefficients could be zero or even negative. However, for regressions (10)-(13) the theoretical predictions of $\beta_B = 1$ and $\beta_p = 1$ remain valid under the more general setting. We offer two conjectures, one theoretical and one empirical, for why our estimates of β_B and β_p are less than 1.

We start by examining the bin regressions in Tables 4-5 and the scatterplots in Figure 2. We notice two features. One, within a bin consumption shares rise more slowly than income shares. Two, across bins consumption shares are more dispersed than income shares. That is, countries with population shares concentrated in the upper quantiles still consume goods with prices in the lower quantiles.

Our first conjecture relates to the assumption in our theory that each household consumes a single unit of the quality differentiated good. Suppose instead that households buy a quality portfolio. For example, a household might buy an expensive television for the media room and less expensive televisions for bedrooms or the basement. Similarly, a consumer might buy high quality apparel for dressy occasions along with low quality apparel for casual or recreational occasions. This would cause consumption shares to be spread more widely over price bins.

Our second conjecture relates to the absence of domestically produced goods from our price data. Suppose that countries specialize in producing goods that will be heavily demanded by domestic consumers. This could be due to factor-based specialization as in Flam-Helpman (1987), in which a country well-endowed with skilled labor will have a comparative advantage in high quality goods and consume high quality goods as well (due to high income). Or, in the presence of trade costs, firms may target the portion of the quality spectrum demanded most intensively by local consumers. Then, a high population share in income bin b would yield a high share of domestically produced goods falling in price bin b . Since domestic sales are

missing from our data we would systematically understate the rise in consumption shares in price bin b corresponding to a rise in population shares for income bin b .

It is also instructive to compare results for individual bins with those for the entire support. The theory works especially well at the upper and lower tails of the distribution, but poorly through the middle of the distribution. When integrating over the entire support the stronger tail results are watered down by the weak middle.

Why are the effects so much stronger at the tails of the distribution? To further explore this feature we calculate, for each product, the price dispersion within each bin as the percentage change in prices from the minimum to the maximum price.¹⁷ Using values for the median product, we then plot the price dispersion within each bin in the top panel of Figure 3 (a value of one is a 100% price increase from min to max). For both quartiles and deciles we see a U-shaped pattern of large price changes in the tails (over 300% within the 1st and 4th quartiles and over 140% for the 1st and 10th deciles) and much smaller price changes in the middle bins. These results suggest that price variations are more prominent at the upper and lower tails of the import price distributions. This may explain the stronger correlations we find in the tails. By organizing the data into equal sized consumption bins we are compressing the price variation available to explain in those middle bins. This could be especially problematic for our estimates if the prices are noisy.

Another way to explore price variation by bin is to calculate the share of each bin in the total range in prices. (Since bins are defined by consumption shares, they need not contribute equally to the total range in prices.) Suppose that prices for a commodity vary from $p=1$ to $p=101$, and the first quartile of prices ranges from $p=1$ to $p=51$. In that case the first quartile is responsible

¹⁷ For the top bin we use the 99th percentile price for the max and for the bottom bin we use the 1st percentile price. This avoids a problem of large outliers driving the variation.

for half of the total range. We calculate the contribution of each bin to total price range for each product, then plot the median products in the bottom panels of Figure 3. In both cases, bins 1 to B-1 contribute relatively little to the total price range while the last bin contributes the lion's share. Essentially, prices rise slowly over the bins until, in the last bin, prices turn up sharply. This feature of the data is reminiscent of the theoretical results of Flam and Helpman (1987) as illustrated in Figure 3. Prices rise slowly at low income levels and then turn up sharply once reaching the North's specialization region.

6. Conclusion

In this paper, we investigate how the distribution of income shapes patterns of consumption and international trade in quality differentiated varieties within narrow product categories. We extend Flam and Helpman (1987) to the case of multiple differentiated goods and multiple countries with different technologies. We show that cross-country differences in the moments and cumulative distribution functions of income lead to corresponding differences in the distribution of product prices. By deriving results in terms of national income and price distributions we are able to evaluate a model that predicts heterogeneity in household consumption decisions without needing household consumption data.

To test these predictions we employ microdata on income from household surveys for 26 countries to construct income distributions within and across countries. Our findings are consistent with the predictions of our model. The pairs of importers whose income distributions look more similar have more similar import price distributions, whether similarity is measured by 1st - 4th moment statistics, population and consumption shares within world income and product price quartiles and deciles, or income and price dis-similarity indices.

Our findings show that a country's income distribution shapes its import demand in important ways. They lend support to Murphy and Shleifer's (1997) insight that developing countries may have limited access to developed countries' markets because the goods they produce lack the high qualities that high-income consumers demand. Further, this view of trade patterns lies in stark contrast to the dominant models of horizontal product differentiation in the trade literature, which provide no role for heterogeneous consumers or income differences in explaining trade patterns. These models invoke fixed costs of trade to explain why, in the data, countries import only small subsets of available varieties. They further imply that restricting the variety set leads to first order welfare losses. In the quality differentiation model, a household desires a single quality differentiated variety, while an economy as a whole desires subsets of the world's varieties dictated by its income distribution. Consuming a narrow range of imported varieties may simply reflect a narrow range of income, with no particular welfare loss. We leave to future work the question of how calculations of the variety gains from trade should be qualified by the insights of quality differentiation models.

Finally, there is a rich theoretical literature on quality differentiation in trade in which authors combine vertical differentiation with non-homothetic preferences and income distributions to shed light on many questions that are difficult for horizontal differentiation models to answer. They show that one country's income re-distribution policy may affect another country's income distribution (Flam and Helpman, 1987; Matsuyama, 2000), that absolute poverty and per capita growth can be sustained simultaneously in a fully integrated world economy (Funk, 1998), that an export boom may push a country into industrialization in the presence of a large middle class (Murphy, Shleifer and Vishny, 1989), and that an improvement in the productivity of one industry may trigger the take-off of a series of industries

one after another (Matsuyama, 2002). While we do not directly address these implications, our paper is a first step in taking the common elements of these models—the interactions of vertical differentiation with non-homothetic preferences and income distribution—to the data.

References

- Atkinson, A. B., Brandolini, A., 2001. Promise and Pitfalls in the Use of ‘Secondary’ Data-Sets: Income Inequality in OECD Countries as a Case Study. *Journal of Economic Literature*, 39, 771-799.
- Atkinson, A.B., Rainwater, L., Smeeding, T.M.,1995. Income Distribution in OECD countries: Evidence from the Luxembourg Income Study, in *Social Policy Studies* No. 18. OECD, Paris.
- Bils, M., Klenow, P. J., 2001. Quantifying Quality Growth. *American Economic Review*, 91,1006-2001.
- Buhmann, B., Rainwater, L., Schmaus, G., Smeeding, T.M.,1988. Equivalence Scales, Well-Being, Inequality, and Poverty: Sensitivity Estimates across Ten Countries Using the Luxembourg Income Study (LIS) Database. *Review of Income and Wealth*, 34, 115-42.
- Dalgin, M., Mitra, D., Trindade, V., 2008. Inequality, Non-homothetic Preferences, and Trade: A Gravity Approach. *Southern Economic Journal* 74, 747-774.
- Deaton, A., 2003. Health, Inequality, and Economic Development. *Journal of Economic Literature*, 41, 113-158.
- Deaton, A., 1997, *The Analysis of Household Surveys: A Microeconometric Approach to Development Policy*. The Johns Hopkins University Press, Baltimore.
- Deininger, K. and Squire, L., 1996. A New Data Set Measuring Income Inequality. *World Bank Economic Review*, 10, 565-591.
- DiNardo, J. ,Tobias, J.L., 2001. Nonparametric Density and Regression Estimation. *Journal of Economic Perspectives*, 15, 11-28.
- Dollar, D., Kraay, A., 2002. Growth is Good for the Poor. *Journal of Economic Growth*, 7, 195-225.
- Flam, H., Helpman, E., 1987. Vertical Product Differentiation and North-South Trade. *American Economic Review*, 77, 810-822.
- Funk, P., 1998. Satiation and Underdevelopment. *Journal of Development Economics* 57, 319-

341.

Hallak, J. C., 2006. Product Quality and the Direction of Trade. *Journal of International Economics* 68, 238-265.

Heston, A., Summers R., Aten, B., 2002. Penn World Table Version 6.1. Center for International Comparisons, University of Pennsylvania (CICUP).

Hummels, D., Klenow, P. J., 2005. The Variety and Quality of a Nation's Trade. *American Economic Review* 95, 704-723.

Hummels, D., Skiba, XXXX

Hunter, L., 1991. The Contribution of Nonhomothetic Preferences to Trade, *Journal of International Economics* 30, 345-358.

Markusen, J., 1986. Explaining the Volume of Trade: An Eclectic Approach, *American Economic Review* 76, 1002-1011.

Matsuyama, K., 2000. A Ricardian Model with a Continuum of Goods under Nonhomothetic Preferences: Demand Complementarities, Income Distribution, and North-South Trade. *Journal of Political Economy* 108, 1093-1120.

Matsuyama, K., 2002. The Rise of Mass Consumption Societies. *Journal of Political Economy*, 110, 1035-1070.

Mitra, D., Trindade, V., 2005. Inequality and Trade. *Canadian Journal of Economics* 38, 1253-1271.

Murphy, K. M., Shleifer, A., 1997. Quality and Trade. *Journal of Development Economics* 53, 1-15.

Murphy, K. M., Shleifer, A. Vishny, R., 1989. Income Distribution, Market Size, and Industrialization. *Quarterly Journal of Economics* 104, 537-564.

Reimer, J., 2005. Non-Homothetic Preferences and International Trade, working paper, Oregon State University.

Sala-i-Martin, X., 2006. The World Distribution of Income: Falling Poverty and ... Convergence, Period. *Quarterly Journal of Economics* 121, 51-397.

Schott, P. K., 2004. Across-Product versus Within-Product Specialization in International Trade. *Quarterly Journal of Economics* 119, 647-678.

Silverman, B.W., 1986. Density Estimation for Statistics and Data Analysis. Chapman and Hall, London.

Data Appendix

We employ wave 5 (2000) of the Luxembourg Income Study (LIS) to generate income distribution data for 26 countries. We use the LIS data, rather than another widely used dataset on cross-country income distribution (Deninger and Squire 1996 and its extensions by the World Bank, henceforth DWSB) for two reasons. One, the LIS is more consistent and better suited for cross-country comparisons of income distributions (Atkinson and Brandolini 2001, Deaton 2003). It has a more complete measure of disposable household income (detailed below), and allows us to make adjustments to account for differences in family size. Two, the LIS allows us to calculate household income at single *percentile* increments while the DWSB provides *quintile* level income shares. This is critical for our calculations involving higher income moments.

From the LIS we extract disposable household income (DPI), a commonly used measure in the analysis of income inequality. DPI includes monetary income after direct taxes and transfer payments, and omits indirect taxes, benefits from public spending (e.g. health care, education, most housing subsidies) and wealth, except to the extent that it is represented by cash interest, rent, and dividends. The data are in local currency units and we convert them to US dollars using the current exchange rate data from Penn World Tables 6.1.

The DPI data are at the level of households rather than consumers. Since household sizes vary, and consumption needs vary by age, we adjust DPI using an adult equivalence scale (AES). Total household income is divided by the number of equivalent adults in order to get a measure of household “equivalent” income. Buhmann et al. (1988) propose a succinct parametric approximation to equivalence scales that summarizes the wide range of scales in use:

$$\text{Adjusted Income} = \text{DPI} / \text{Household Size}^E.$$

The equivalence elasticity $E \in [0,1]$ represents economies of scale in household size. We employ the LIS Equivalence Scale ($E = 0.5$), a commonly used scale among researchers who study income inequality using the LIS data (e.g. Atkinson et al. 1995). An alternative popular approach explicitly employs data on the numbers of adults and children in the household. This approach is only feasible for a limited subset of our data.

The wave 5 (2000) LIS data is available for Austria, Belgium, Canada, Estonia, Finland, Germany, Greece, Hungary, Ireland, Israel, Italy, Luxembourg, Mexico, Netherlands, Norway, Poland, Russia, Slovenia, Spain, Sweden, Switzerland, Taiwan, UK and US. The wave 5 LIS data is missing for Australia and France. Since the literature shows that quantile income levels within a country tend to follow smooth trends over time (e.g. Dollar and Kraay 2002, Sala-i-Martin 2006), we estimate linear income trends by percentile using waves 1-4 data for Australia and France. The average R^2 for these income trends is 0.958 for Australia and 0.965 for France. In our tables we experiment with dropping Australia and France from the regressions with little change in results.

Finally, it should be noted that higher moments of the income statistics, especially skewness and kurtosis, are significantly impacted by how we handle the first and last percentile of income for each country. Our LIS extract contains data for min, mean, and max income within each income percentile. For percentiles 1-99 in all countries, the min, mean and max income values are very similar. For the first percentile, min income is commonly negative. For the last percentile the max income might be two orders of magnitude larger than the mean (e.g. for Finland the mean is 567790.8 Euros and the max is 16.2 million Euros) For this reason we use the mean income within each percentile for calculating distributions. Alternative choices (e.g.

using the max income) have large effects on the magnitudes of the skewness and kurtosis values in Table 1 and their coefficients in Tables 2 and 3, but do not change the statistical significance of the coefficients in Tables 2 and 3.

Figure 1 Income Dispersion Across Countries

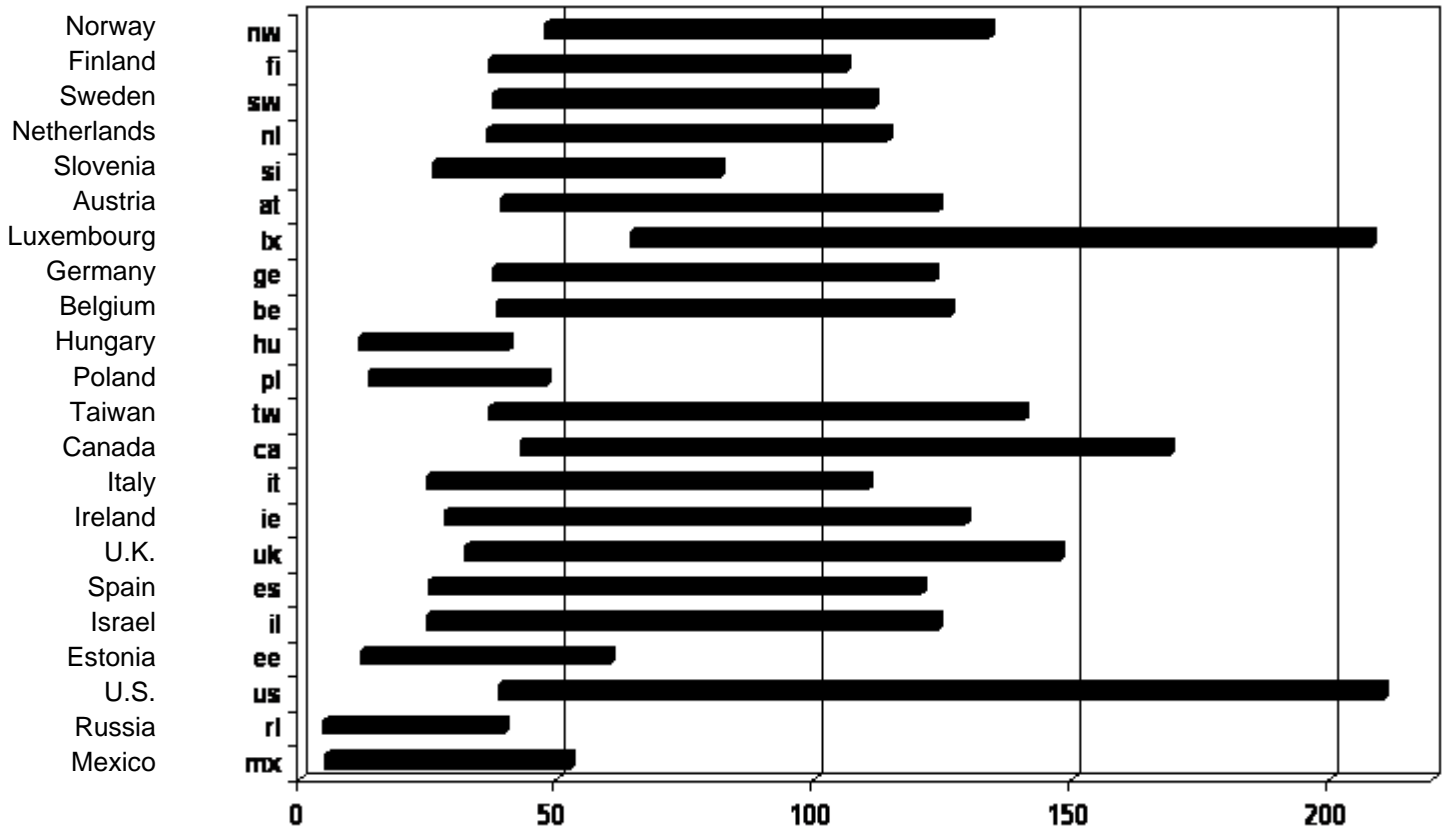


Figure 2: Consumption Price and Income Shares
Pairwise Differences

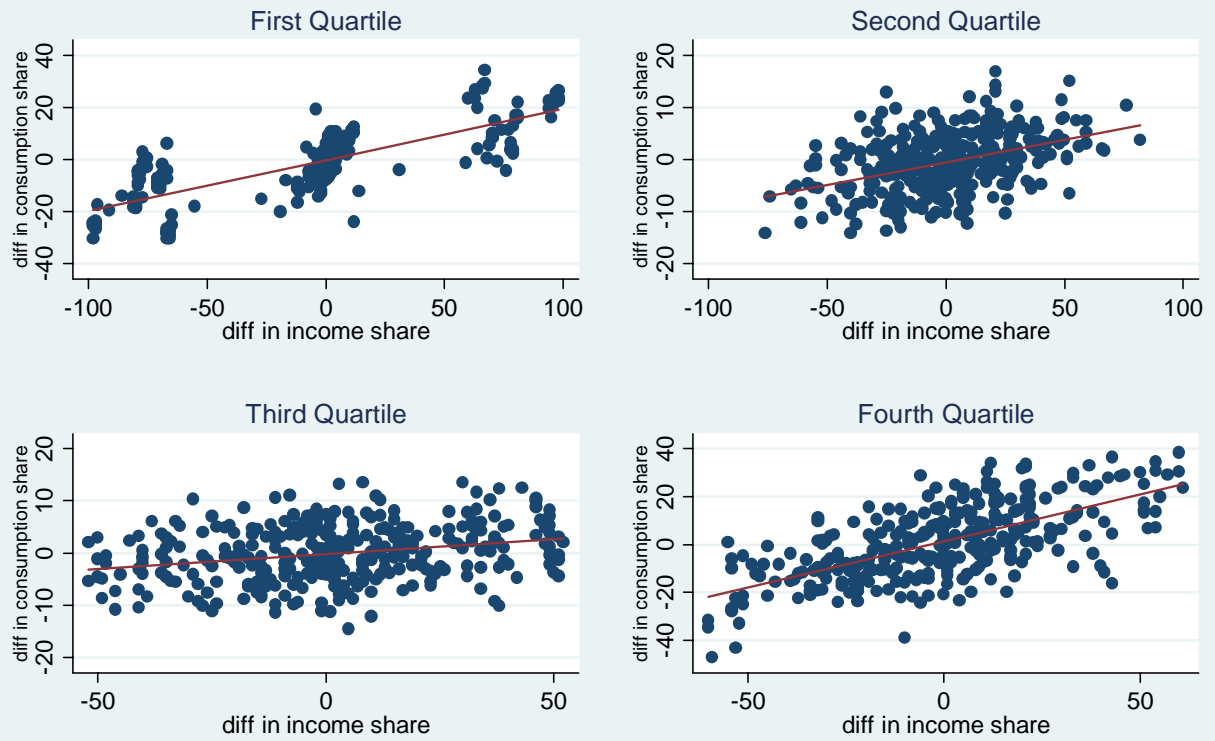


Figure 3 Price Dispersion by Bin

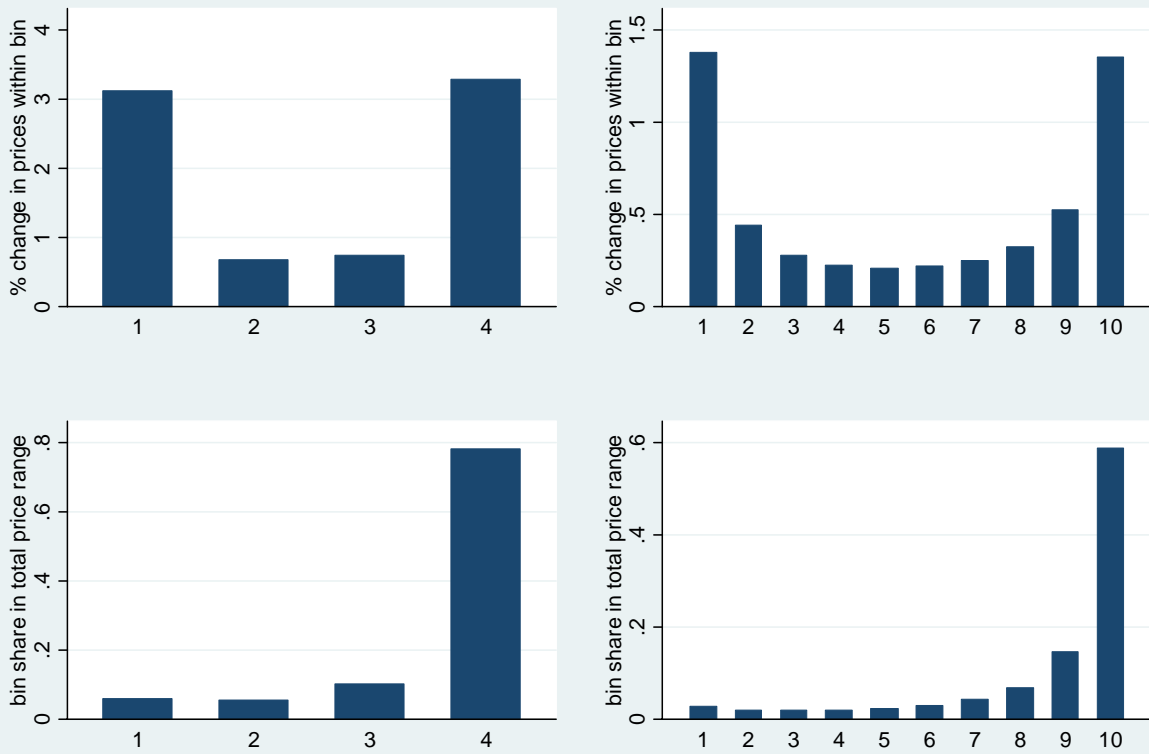


Table 1: Income Statistics

	Distribution of Household Income						Population Share by World Income Quartiles				
	Mean	Median	St. Dev	Inter-decile		Skewness	Kurtosis	1	2	3	4
				Range							
Australia	15710	12875	16318	19642		6.65	57.76	3	44	36	17
Austria	17799	16100	9517	18837		2.33	12.76	1	27	49	23
Belgium	17854	15532	14456	18645		6.09	50.90	1	31	48	20
Canada	19483	17069	12231	23809		2.62	15.23	1	27	41	31
Switzerland	29156	25111	19690	31858		4.21	30.21	1	6	31	62
Germany	17724	15861	9331	19411		1.98	10.16	1	27	49	23
Spain	12159	10167	8750	16704		2.79	15.75	4	59	28	9
Estonia	2523	1969	2001	3681		3.00	16.54	79	20	1	0
Finland	16460	14845	9750	15841		4.23	30.43	1	33	51	15
France	21720	16967	28261	22889		8.14	75.68	0	25	45	30
UK	21516	17381	17366	29044		3.89	25.99	1	29	34	36
Greece	8915	7557	6161	12614		2.24	11.19	10	70	15	4
Hungary	2736	2306	1753	3192		2.93	16.20	77	22	1	0
Ireland	17391	15204	11915	22226		2.96	17.44	1	36	40	23
Israel	15472	12678	11830	21810		3.24	20.12	2	46	33	19
Italy	12848	10769	9524	16507		3.29	20.07	4	55	31	10
Luxemborg	26958	23735	13845	29844		1.88	8.71	0	7	37	56
Mexico	3810	2308	5307	6754		4.89	33.98	66	30	3	1
Netherlands	17944	16604	8907	19196		1.53	8.34	2	23	50	25
Norway	25167	22685	17240	23059		5.70	46.59	0	9	38	53
Poland	3133	2691	2095	3598		3.60	23.14	67	32	1	0
Russia	851	596	956	1446		4.57	32.03	98	2	0	0
Slovenia	7942	7328	3816	8318		1.51	7.38	6	84	9	1
Sweden	18338	16644	9937	18461		2.94	18.71	1	23	52	24
Taiwan	13842	11842	8357	16999		2.27	11.38	2	51	35	12
USA	29055	23916	23889	40796		3.29	19.22	2	16	27	55

Table 2: Moments of the Price and Income Distributions

	Full Sample					
	Mean	Median	SD	Range	Skewness	Kurtosis
beta	0.138	0.160	0.131	0.190	0.146	0.108
t-stat	(15.68)	(16.29)	(7.42)	(7.80)	(5.30)	(4.77)
R2	0.076	0.091	0.024	0.037	0.008	0.011
N	1145518	1145518	1115180	1115180	979885	1115180

	Restricted Sample: Similar First Moments					
	Mean	Median	SD	Range	Skewness	Kurtosis
beta			0.255	0.317*	0.146	0.098
t-stat			(2.99)	(1.71)	(2.58)	(2.21)
R2			0.014	0.006	0.01	0.01
N			278253	278253	245264	278253

Notes:

Estimation of equation XX

standard errors adjusted for clustering at country or country pair level

t-stats in parentheses. bold: significant at 5% level; * significant at 10% level

Table 3: Robustness Checks for Price and Income Moments

Restricted Sample: no imputed income data						
	Mean	Median	SD	Range	Skewness	Kurtosis
beta	0.144	0.166	0.139	0.183	0.090	0.051*
t-stat	(15.60)	(16.33)	(6.84)	(7.09)	(2.73)	(1.92)
R2	0.085	0.100	0.025	0.036	0.002	0.002
N	971932	998779	971932	971932	852046	971932

Restricted Sample: No Homogeneous Commodities						
	Mean	Median	SD	Range	Skewness	Kurtosis
beta	0.149	0.181	0.123	0.175	0.153	0.114
t-stat	(17.01)	(17.70)	(7.60)	(7.97)	(5.79)	(5.20)
R2	0.105	0.133	0.027	0.04	0.01	0.013
N	846365	846365	831108	831108	754644	831108

Restricted Sample: No Textile and Apparel						
	Mean	Median	SD	Range	Skewness	Kurtosis
beta	0.124	0.144	0.121	0.174	0.133	0.100
t-stat	(14.82)	(15.49)	(7.03)	(7.24)	(5.03)	(4.59)
R2	0.062	0.074	0.020	0.031	0.007	0.010
N	951302	951302	926431	926431	813996	926431

Notes:

Estimation of equation XX

standard errors adjusted for clustering at country or country pair level

t-stats in parentheses. bold: significant at 5% level; * significant at 10% level

Table 4	Quartiles				Total
	1	2	3	4	
Comparison: rest of world					
Aggregated Bins	0.202 (6.04)	0.091 (2.38)	0.058 (1.36)	0.381 (4.23)	-0.225 (-1.61)
Product Level Bins	0.275 (3.42)	0.070* (1.91)	0.126 (2.77)	0.402 (3.98)	0.267 (3.10)
Comparison: individual importer					
Aggregated Bins	0.202 (22.15)	0.092 (8.74)	0.058 (5.00)	0.380 (15.51)	0.144 (8.76)
Product Level Bins	0.273 (14.27)	0.071 (6.74)	0.127 (10.26)	0.397 (14.22)	0.114 (9.42)

Notes:

Estimation of equation XX

standard errors adjusted for clustering at country or country pair level

t-stats in parentheses. bold: significant at 5% level; * significant at 10% level

Table 5

	Deciles										Total
	1	2	3	4	5	6	7	8	9	10	Total
Comparison: rest of world											
Aggregated Bins	0.250 (5.42)	0.139 (4.47)	0.061 (2.75)	0.065* (1.87)	0.040 (0.94)	0.031 (0.53)	0.080* (1.81)	0.121 (2.54)	0.235 (3.81)	0.575 (3.11)	-0.198 (-1.67)
Product Level Bins	0.498 (5.67)	0.110 (2.43)	0.042 (2.91)	0.039 (1.48)	.086* (1.97)	0.114 (2.76)	0.130 (2.92)	0.140 (2.52)	0.196 (3.33)	0.696 (2.38)	0.227 (2.94)
Comparison: individual importer											
Aggregated Bins	0.251 (19.99)	0.139 (16.46)	0.061 (10.12)	0.065 (6.81)	0.040 (3.46)	0.01* (1.96)	0.080 (6.61)	0.121 (9.34)	0.235 (14.03)	0.575 (11.40)	0.119 (8.16)
Product Level Bins	0.501 (18.31)	0.107 (9.86)	0.041 (8.39)	0.039 (5.08)	0.089 (7.47)	0.116 (10.10)	0.130 (10.76)	0.137 (9.69)	0.192 (12.43)	0.687 (9.53)	0.088 (9.14)

Notes:

Estimation of equation XX

standard errors adjusted for clustering at country or country pair level

t-stats in parentheses. bold: significant at 5% level; * significant at 10% level

Tables 6: Robustness

	Quartiles				Total
	1	2	3	4	Total
Aggregated Bins : individual importer comparison					
All Data	0.202 (22.15)	0.092 (8.74)	0.058 (5.00)	0.380 (15.51)	0.144 (8.76)
No Imputed Income Data	0.205 (21.43)	0.092 (8.25)	0.051 (4.12)	0.388 (14.82)	0.153 (8.43)
No Homogeneous Commodities	0.184 (18.05)	0.096 (9.02)	0.047 (4.14)	0.361 (14.72)	0.130 (7.66)
No Textile and Apparel	0.203 (22.45)	0.098 (9.56)	0.063 (5.39)	0.397 (16.29)	0.147 (9.18)

Notes:

Estimation of equation XX; comparison group

standard errors adjusted for clustering at country or country pair level

t-stats in parentheses. bold: significant at 5% level; * significant at 10% level